July 2025

Reference Manual

Blackmagicdesign

Fusion 20



Leading the Creative Video Revolution



Welcome

Welcome to Fusion for Mac, Linux and Windows!

Fusion is the world's most advanced compositing software for visual effects artists, broadcast and motion graphic designers and 3D animators. With over 30 years of development, Fusion has been used on over 1000 major Hollywood blockbuster feature films! Fusion features an easy and powerful node based interface so you can construct complex effects simply by connecting various types of processing together. That's super easy and extremely fast! You get a massive range of features and effects included, so you can create exciting broadcast graphics, television commercials, dramatic title sequences and even major feature film visual effects!

Fusion Studio customers can also use DaVinci Resolve Studio to get a complete set of editing, advanced color correction and professional Fairlight audio post production tools. Clips in the DaVinci Resolve timeline can be shared with Fusion so you can collaborate on your complex compositions within your VFX team and then render the result back directly into the DaVinci Resolve timeline. We hope you enjoy reading this manual and we can't wait to see the work you produce with Fusion.

The Fusion Engineering Team



Grant Petty
CEO Blackmagic Design

Contents

Fusion Fundamentals

1 Introduction to Compositing in Fusion 7
2 Exploring the Fusion Interface 13
3 Getting Clips into Fusion 61
4 Rendering Using Saver Nodes 88
5 Working in the Node Editor 113
6 Node Groups, Macros, and Fusion Templates 163
7 Using Viewers 186
8 Editing Parameters in the Inspector 229
9 Animating in Fusion's Keyframes Editor 255
10 Animating in Fusion's Spline Editor 273
11 Animating with Motion Paths 305
12 Using Modifiers, Expressions, and Custom Controls 323
13 Bins 335
14 Fusion Connect 355
15 Preferences 369

2D Compositing

16 Controlling Image Processing and Resolution 419
17 Managing Color for Visual Effects 430
18 Understanding Image Channels 443
19 Compositing Layers in Fusion 479
20 Rotoscoping with Masks 504
21 Paint 528
22 Using the Tracker Node 556
23 Planar Tracking 592
24 Using Open FX, Resolve FX, and Fuse Plugins 599

3D Compositing

25 3D Compositing Basics 605

26 3D Camera Tracking 660

27 Particle Systems 678

Advanced Compositing Techniques

28 Optical Flow and Stereoscopic Nodes 688

Fusion Page Effects

29 3D Nodes 701

30 3D Light Nodes 819

31 3D Material Nodes 838

32 3D Texture Nodes 867

33 Blur Nodes 892

34 Color Nodes 917

35 Composite Nodes 992

36 Deep Image Nodes 1012

37 Deep Pixel Nodes 1031

38 Effect Nodes 1047

39 Film Nodes 1084

40 Filter Nodes 1103

41 Flow Nodes 1120

42 Flow Organizational Nodes 1123

43 Fuses 1128

44 Generator Nodes 1130

45 I/O Nodes 1170

46 Layer Nodes 1197

47 LUT Nodes 1211

48 Mask Nodes 1220

49 Matte Nodes 1261

50 Metadata Nodes 1332

51 Miscellaneous Nodes 1340

52 Optical Flow 1379

53 Paint Node 1399
54 Particle Nodes 1409
55 Position Nodes 1469
56 Resolve Connect 1488
57 Shape Nodes 1494
58 Stereo Nodes 1534
59 Tracking Nodes 1566
60 Transform Nodes 1622
61 USD Nodes 1650
62 VR Nodes 1703
63 Warp Nodes 1714
64 Modifiers 1744

Menu Description

Fusion 1787



Fusion Fundamentals

FUSION

CONTENTS

1	Introduction to Compositing in Fusion	7
2	Exploring the Fusion Interface	13
3	Getting Clips into Fusion	61
4	Rendering Using Saver Nodes	88
5	Working in the Node Editor	113
6	Node Groups, Macros, and Fusion Templates	163
7	Using Viewers	186
8	Editing Parameters in the Inspector	229
9	Animating in Fusion's Keyframes Editor	255
10	Animating in Fusion's Spline Editor	273
11	Animating with Motion Paths	305
12	Using Modifiers, Expressions, and Custom Controls	323
13	Bins	335
14	Fusion Connect	355
15	Preferences	369

Chapter 1

Introduction to Compositing in Fusion

This introduction is designed explicitly to help users who are new to Fusion get started learning this exceptionally powerful environment for creating and editing visual effects and motion graphics right from within DaVinci Resolve or using the stand-alone Fusion Studio application.

This documentation covers both the Fusion Page inside DaVinci Resolve and the stand-alone Fusion Studio application.

FUSION

Contents

What Is Fusion? 8
The Fusion Page within DaVinci Resolve 8
The Fusion Studio Stand-Alone Application 10
What Kinds of Effects Does Fusion Offer? 10
How Hard Will This Be to Learn? 12

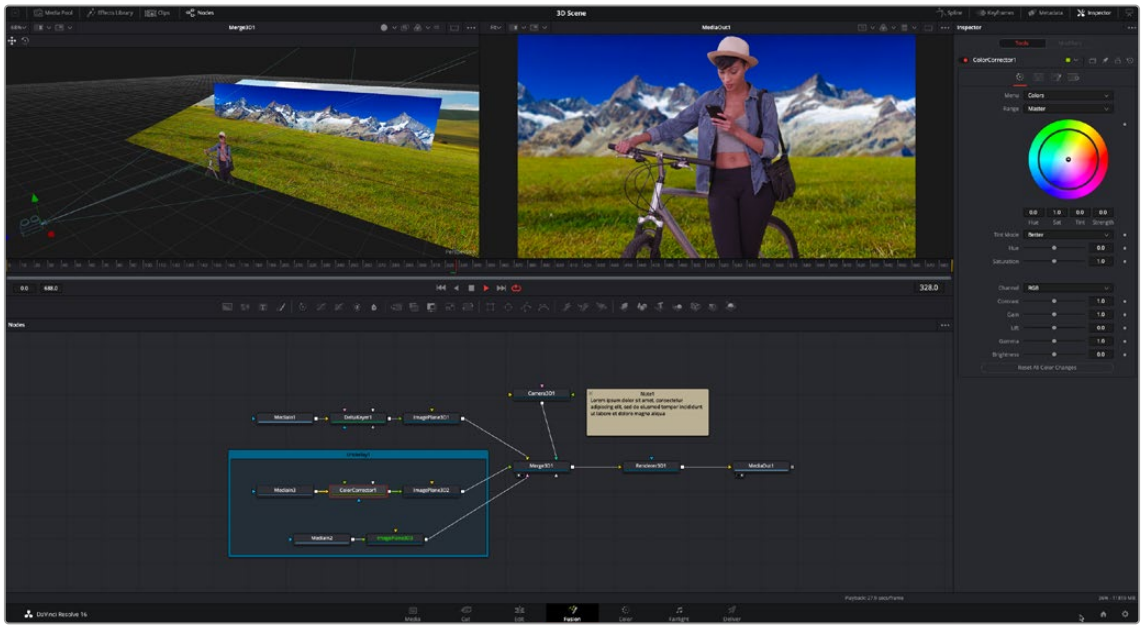
What Is Fusion?

Blackmagic Design Fusion is powerful 2D and 3D visual effects compositing software with over thirty years of evolution serving the motion picture and broadcast industry, creating effects seen in countless films and television series. It is available as a stand-alone application as well as a page within DaVinci Resolve.

In its purest form, Fusion is a collection of image-processing engines called nodes. These nodes represent effects like blurs and color correctors, as well as images, 3D models, and spline masks. Similar to effects you may be familiar with, each node includes a set of parameters that can be adjusted and animated over time. Stringing different nodes together in a graphical user interface called a node tree allows you to create sophisticated visual effects. The nodes, node trees, and all settings you create are saved in a document called a Composition, or “comp” for short.

The Fusion Page within DaVinci Resolve

Merged right into DaVinci Resolve, the Fusion page makes it possible to jump immediately from editing right into compositing, with no need to export media, relink files, or launch another application to get your work done. Now everything you need lives right inside DaVinci Resolve.

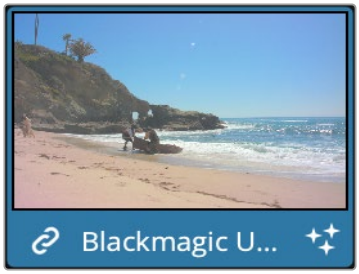


The Fusion page in DaVinci Resolve, showing viewers, the Node Editor, and the Inspector

How Do I Use the Fusion Page?

The relationship between the Edit page and the Fusion page is similar to the relationship between the Edit page and the Color page. Each clip can have a grade applied to it in the Color page, and similarly every clip can have a composition applied to it in the Fusion page.

If you use the Fusion page to create any kind of effect or composite, a badge appears on that clip in the Timeline to show that clip has a composition applied to it.



Clips with Fusion page compositions have a Fusion badge to the right of the name.

To create an effect in the Fusion page of DaVinci Resolve, you need only park the playhead over a clip in the Edit or Cut page and then click the Fusion page button. Your clip is immediately available as a MediaIn node in the Fusion page, ready for you to add a variety of stylistic effects. You can paint out an unwanted blemish or feature, build a quick composite to add graphics or text, or accomplish any other visual effect you can imagine, built from the Fusion page's toolkit of effects.

Alternatively, in DaVinci Resolve, you have the option of editing together all the clips you want to use, superimposing and lining up every piece of media you'll need with the correct timing, before selecting them and creating a Fusion clip. A Fusion clip functions as a single item in the Edit or Cut page timeline, but once in the Fusion page, each piece of media you've assembled is revealed in a fully built Fusion composition, ready for you to start adding nodes to customize for whatever effect you need to create.

Whichever way you want to work, all this happens on the very same timeline as editing, grading, and audio post, for a seamless back and forth as you edit, refine, and finish your projects.

How Do Fusion Effects Differ from Edit Page Effects?

When using DaVinci Resolve, you can create numerous effects in the Edit page. Transitions, fades, superimpositions, over-the-shoulder picture-in-picture effects, time remapping, and lower third titles are some of the effects that can be quickly and more efficiently created in the Edit or Cut page timeline. However, the Fusion page's node-based interface lets you go deep into the details of a composition to create sophisticated 2D and 3D effects with precise control and endless customization. Effects that include more than two or three layers can be much more manageable in Fusion. Green or blue screen composites, sky replacements, and object removal are all effects better suited for Fusion's more advanced toolset.

How Do Fusion Effects Differ from Color Page Effects?

The Color page in DaVinci Resolve can also handle some visual effects work. Effects that blur the line between color grading and finishing can be very fast and intuitive in the Color Page, especially for people already familiar with the Color page toolset. Beauty work and small image repairs can make efficient use of the Color page's straightforward Tracking tool, Face Refinement, and Patch Replacer effects. However, when it comes to more challenging blue/green screen compositing, the tools built around Fusion's powerful Delta keyer are more capable of handling these shots. Integrating 3D objects into live-action scenes, split-screen effects, motion graphics, and precise keyframing are all better suited to the Fusion page.

The Fusion Studio Stand-Alone Application

Creating visual effects with the stand-alone Fusion Studio software begins with opening Fusion, creating a new composition, importing some clips via Loader nodes, and building out your composite with effects. Just like the Fusion Page in DaVinci Resolve, you add effects using different nodes from the Effects Library, and you combine multiple layers of imagery using Merge nodes. Once you've created the desired result, add a Saver node to the end of the tree of nodes you've created to render your final result.

Rendering Out Your Final Result

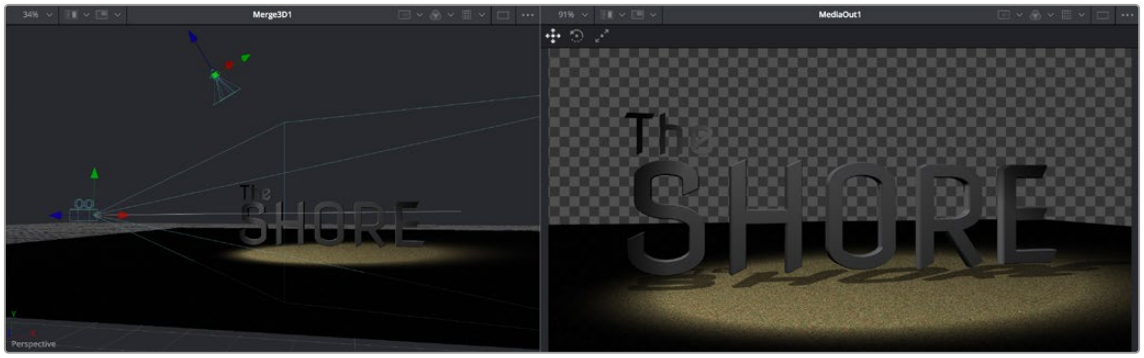
Unlike the Fusion Page in DaVinci Resolve, which renders directly back into the Edit or Cut page timeline, the final step in Fusion Studio is to render the finished effect to disk as a movie file or image sequence. The last node in every node tree is a Saver node. Saver nodes configure the output file format and render the file to disk. You can use as many Saver nodes in a composite as you need. For instance, you might use multiple Saver nodes to render out intermediate areas of a composite or to output a composite in multiple formats.

What Kinds of Effects Does Fusion Offer?

In addition to the kinds of robust compositing, paint, rotoscoping, and keying effects you'd expect from a fully-featured 2D compositing environment, Fusion offers much more.

3D Compositing

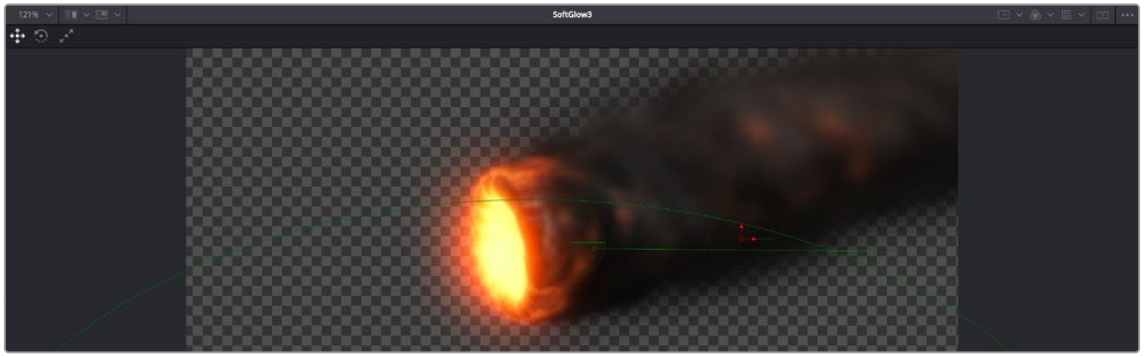
Fusion has powerful 3D nodes that include extruded 3D text, simple geometry, and the ability to import 3D models. Once you've assembled a 3D scene, you can add cameras, lighting, and material shaders, and then render the result with depth-of-field effects and auxiliary channels to integrate with more conventional layers of 2D compositing, for a sophisticated blending of 3D and 2D operations in the very same node tree.



A 3D scene with textured 3D text, created entirely within Fusion

Particles

Fusion also has an extensive set of nodes for creating particle systems that have been used in major motion pictures, with particle generators capable of spawning other generators, 3D particle generation, complex simulation behaviors that interact with 3D objects, and endless options for experimentation and customization. You can create particle system simulations for VFX or more abstract particle effects for motion graphics.



A 3D particle system, also created entirely within Fusion

Text

The Text tools in Fusion are exceptional, giving you layout and animation options in both 2D and 3D. Furthermore, within DaVinci Resolve, these Text tools have been incorporated into the Edit and Cut pages as Fusion Titles. These title templates are compositions saved from Fusion as macros with published controls that are visible in the Edit or Cut page Inspector for easy customization, even if you're working with people who don't know Fusion.



A multi-layered text composite integrating video clips and Fusion-generated elements

And Lots More

The list goes on. With Stereo and VR adjustment nodes, Planar Tracking, Deep Pixel nodes for re-compositing rendered 3D scenes using Auxiliary Channel data, powerful Masking and Rotoscoping nodes, and Warping effects, Fusion is an impressively featured environment for building worlds, fixing problems, and flying multi-layered motion graphics animations through your programs.

How Hard Will This Be to Learn?

That depends on what you want to do, but honestly it's not so bad with this PDF at your side, helping guide the way. It's worth repeating that this Fusion documentation was developed specifically to help users who've never before worked with Fusion learn the core concepts needed to perform the basics, in preparation for learning the rest of the application on your own.

Fusion is a deep, production-driven product that's had decades of development, so its feature set is deep and comprehensive. You won't learn it in an hour, but much of what you'll find won't be so very different from other compositing applications you may have used. And if you've familiarized yourself with the node-based grading workflow of the DaVinci Resolve Color page, you've already got a leg up on understanding the central operational concept of compositing in Fusion.

Chapter 2

Exploring the Fusion Interface

This chapter provides an orientation on the Fusion user interface, providing a quick tour of what tools are available, where to find things, and how the different panels fit together to help you build and refine compositions in this powerful node-based environment.

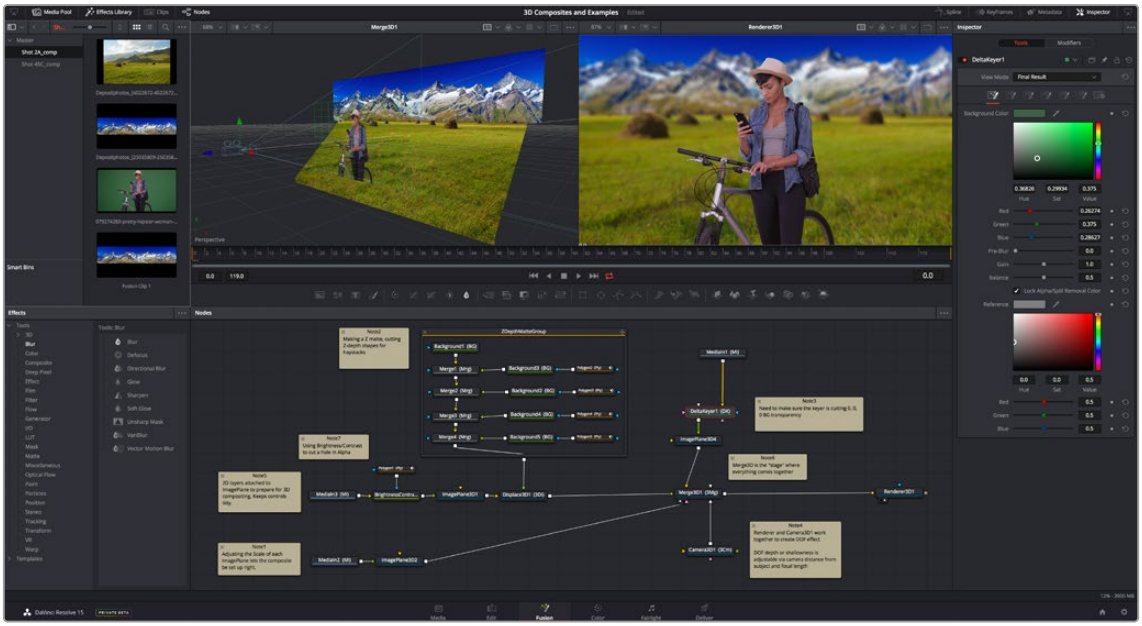
Contents

The Fusion User Interface 15
The Work Area 16
Interface Toolbar 16
Choosing Which Panel Has Focus 17
Viewers 18
Zooming and Panning into Viewers 19
Loading Nodes Into Viewers 19
Clearing Viewers 20
Viewer Controls 20
Time Ruler and Transport Controls 22
Time Ruler Controls in the Fusion Page 22
Time Ruler Controls in Fusion Studio 23
The Playhead 24
Zoom and Scroll Bar 24
Transport Controls in the Fusion Page 24
Audio Monitoring 26
Transport Controls in Fusion Studio 28
Changing the Time Display Format 32
Keyframe Display in the Time Ruler 32
The Fusion RAM Cache for Playback 32
Toolbar 34
Customizing the Toolbar 34

Node Editor 36
Adding Nodes to Your Composition 36
Removing Nodes from Your Composition 37
Identifying Node Inputs and Node Outputs 37
Node Editing Essentials 37
Navigating the Node Editor 39
Vertical Node Editor Layouts 40
Keeping Organized 40
Status Bar 41
Effects Library 42
The Inspector 43
The Tools and Modifiers Panels 43
Parameter Header Controls 43
Parameter Tabs 44
Keyframes Editor 44
Keyframes Editor Control Summary 45
Adjusting Clip Timings 45
Adjusting Effect Timings 45
Adjusting Keyframe Timings 46
Spline Editor 47
Spline Editor Control Summary 48
Choosing Which Parameters to Show 48
Essential Spline Editing 48
Essential Spline Editing Tools and Modes 49
Thumbnail Timeline in the Fusion Page 51
The Media Pool in the Fusion Page 52
Importing Media Into the Media Pool on the Fusion Page 52
Bins in Fusion Studio 53
Bins Interface 54
The Bin Studio Player 55
The Console 56
Customizing Fusion 57
The Fusion Settings Window 57
Showing and Hiding Panels 58
Resizing Panels 58
Fusion Studio Floating Frame 58
Fusion Keyboard Remapping 59
Fusion Studio Keyboard Remapping 59
Undo and Redo 60

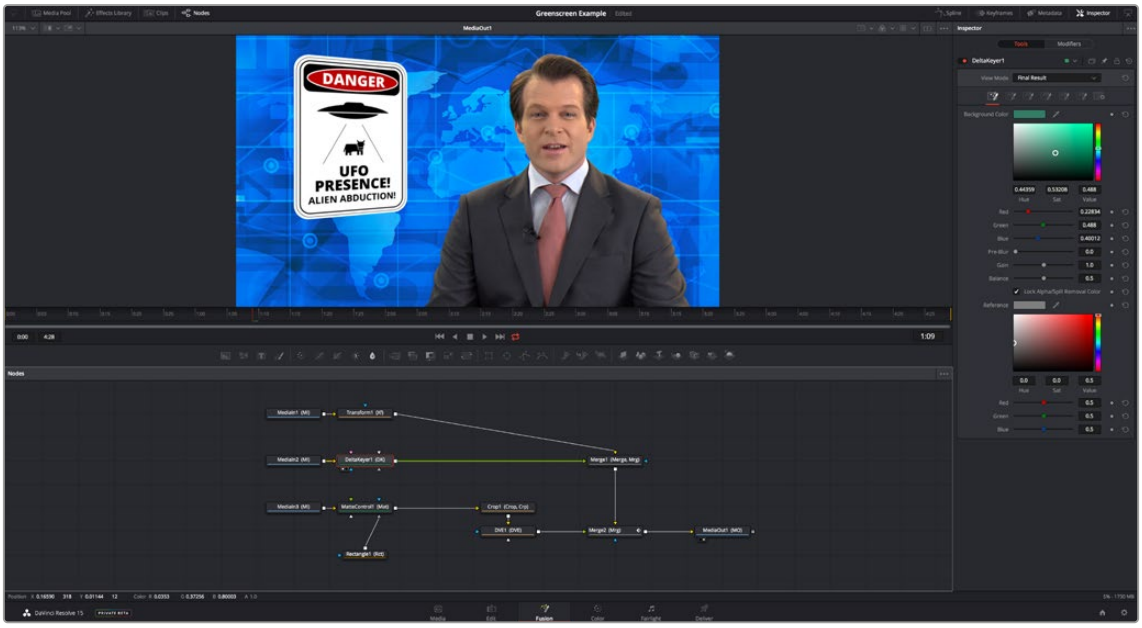
The Fusion User Interface

If you open up everything at once, Fusion is divided into four principal regions designed to help you make fast work of node-based compositing. The viewer(s) are at the top, the work area is at the bottom, the Inspector is at the right, and the Effects Library is the area found at the left. In DaVinci Resolve's Fusion page, the Effects Library shares space with the Media Pool. All these panels work together to let you add effects, paint to correct issues, create motion graphics or title sequences, or build sophisticated 3D and multi-layered composites.



The Fusion user interface shown completely

However, Fusion doesn't have to be that complicated, and in truth, you can work very nicely with only the viewer, Node Editor, and Inspector open for a simplified experience.



A simplified set of Fusion controls for everyday working

The Work Area

You probably won't see the term "the work area" used much, in favor of the specific panels within the work area that you'll be using. Still, the area referred to as the work area is the region at the bottom half of the Fusion user interface, within which you can expose the three main panels used to construct compositions and edit animations in Fusion. These are the Node Editor, the Spline Editor, and the Keyframes Editor. By default, the Node Editor is the first thing you'll see, and the main area you'll be working within, but it can sit side-by-side with the Spline Editor and Keyframes Editor as necessary. You can make more horizontal room on your display for these three panels by putting the Effects Library and Inspector into a half-height mode, if necessary.



The work area showing the Node Editor, the Spline Editor, and the Keyframes Editor

Interface Toolbar

At the very top of Fusion is a toolbar with buttons that let you show and hide different parts of the user interface (UI). Buttons with labels identify which parts of the UI can be shown or hidden. In DaVinci Resolve's Fusion page, if you right-click anywhere within this toolbar, you have the option of displaying this bar with or without text labels.



The UI toolbar of the Fusion page



The UI toolbar of Fusion Studio

These buttons are as follows, from left to right:

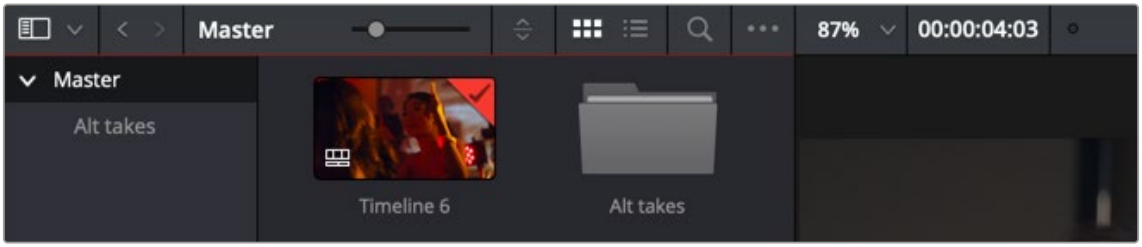
- **Media Pool/Effects Library Full Height:** Lets you set the area used by the Media Pool (DaVinci Resolve only) and/or Effects Library to take up the full height of your display, giving you more area for browsing at the expense of a narrower Node Editor and viewer area. At half-height, the Media Pool/Templates/Effects Library are restricted to the top half of the UI along with the viewers (you can only show one at a time), and the Node Editor takes up the full width of your display.
- **Media Pool:** (DaVinci Resolve only): Shows and hides the Media Pool, from which you can drag additional clips into the Node Editor to use them in your Fusion page composition.
- **Effects Library:** Opens or hides the repository of all node tools available to use in Fusion. From here, you can click nodes to add them after the currently selected node in the Node Editor, or you can drag and drop nodes to any part of the node tree you like.
- **Clips:** (DaVinci Resolve only): Opens and closes the Thumbnail timeline, which lets you navigate your program, create and manage multiple versions of compositions, and reset the current composition.
- **Nodes:** Opens and closes the Node Editor, where you build and edit your compositions.

- **Console (Fusion Studio only):** The Console is a window in which you can see the error, log, script, and input messages that may explain something Fusion is trying to do in greater detail. The Console is also where you can read FusionScript outputs, or input FusionScripts directly.
- **Spline:** Opens and closes the Spline Editor, where you can edit the curves that interpolate keyframe animations to customize and perfect their timing. Each keyframed parameter appears hierarchically within the effect in which it appears in a list to the left.
- **Keyframes:** Opens and closes the Keyframes Editor, which shows each clip and effects node in your Fusion composition as a layer. You can use the Keyframes Editor to edit and adjust the timing of keyframes that have been added to various effects in your composition. You can also use the Keyframes Editor to slide the relative timing of clips that have been added to Fusion, as well as to trim their In and Out points. A spreadsheet can be shown and hidden within which you can numerically edit keyframe values for selected effects.
- **Metadata (DaVinci Resolve only):** Hides or shows the Metadata Editor, which lets you read and edit the available clip and project metadata associated with any piece of media within a composite.
- **Inspector:** Shows or hides the Inspector, which shows you all the editable parameters and controls that correspond to selected nodes in the Node Editor. You can show the parameters for multiple nodes at once, and even pin the parameters of nodes you need to continue editing so that they're displayed even if those nodes aren't selected.
- **Inspector Height:** Lets you open the Inspector to be half height (the height of the viewer area) or full height (the height of your entire display). Half height allows more room for the Node Editor, Spline Editor, and/or Keyframes Editor, but full height lets you simultaneously edit more node parameters or have enough room to display the parameters of multiple nodes at once.

Choosing Which Panel Has Focus

Whenever you click somewhere on the Fusion interface using the pointer or using a keyboard shortcut to "select" a particular panel, you give that panel of the user interface "focus." A panel with focus captures specific keyboard shortcuts to do something within that panel, as opposed to doing something elsewhere in the interface.

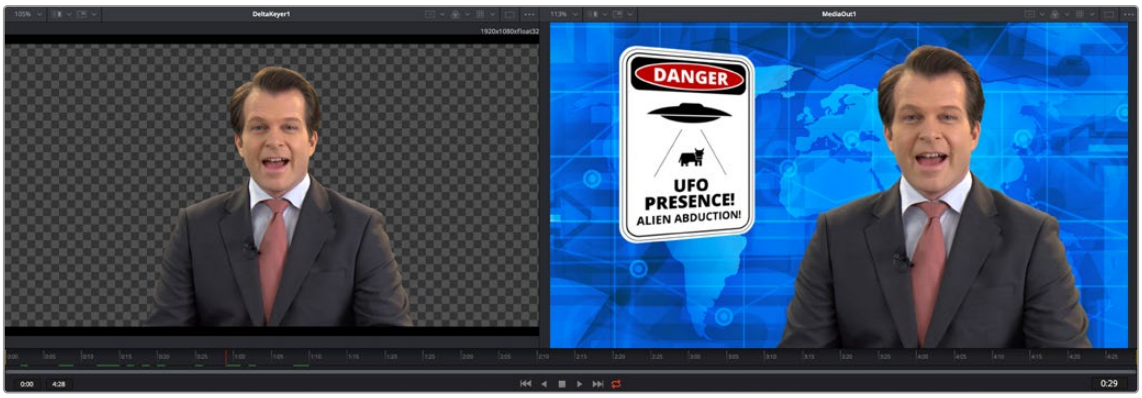
To make it easier to keep track of which panel has focus, a highlight appears at the top edge of whichever panel has focus. In DaVinci Resolve, you must turn on "Show focus indicators in the User Interface" in the UI Settings panel of the User Preferences to see the highlight.



The focus indicator shown at the top edge of the Media Pool, shown next to a viewer that doesn't have focus

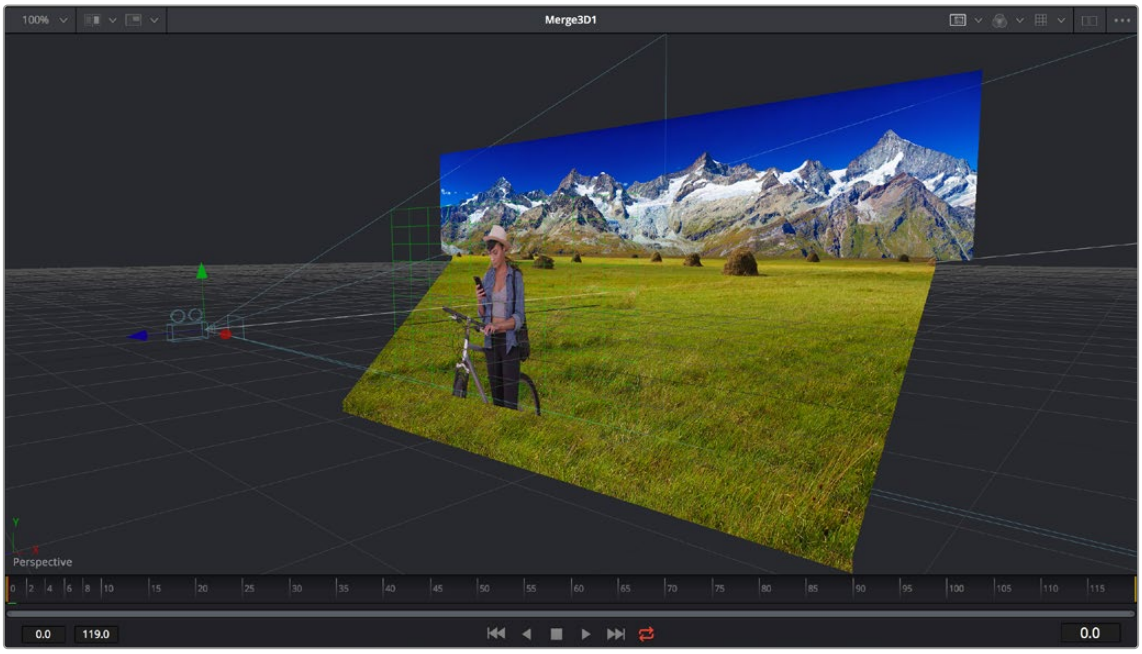
Viewers

The viewer area displays either one or two viewers at the top of the Fusion page, and this is determined via the Viewer button at the far right of the Viewer title bar. Each viewer can show a single node's output from anywhere in the node tree. You assign which node is displayed in which viewer. This makes it easy to load separate nodes into each viewer for comparison. For example, you can load a Keyer node into the left viewer and the final composite into the right viewer, so you can see the image you're adjusting and the final result at the same time.



Dual viewers let you edit an upstream node in one while seeing its effect on the overall composition in the other.

Ordinarily, each viewer shows 2D nodes from your composition as a single image. However, when you're viewing a 3D node, you have the option to set that viewer to one of several 3D views. A perspective view gives you a repositionable stage on which to arrange the elements of the world you're creating. Alternatively, a quad view lets you see your composition from four angles, making it easier to arrange and edit objects and layers within the XYZ axes of the 3D space in which you're working.



Loading a 3D node into a viewer switches on a Perspective view

TIP: In Perspective view, you can hold down the middle and right mouse buttons, then drag in the viewer to pivot the view around the center of the world. All other methods of navigating viewers work the same.

The viewers have a variety of capabilities you can use to compare and evaluate images. This section provides a short overview of viewer capabilities to get you started.

Zooming and Panning into Viewers

There are standardized methods of zooming into and panning around viewers when you need a closer look at the situation. These methods also work with the Node Editor, Spline Editor, and Keyframes Editor.

Methods of panning viewers:

- Middle-click and drag to pan around the viewer.
- Hold down Shift and Command and drag the viewer to pan.
- Drag with two fingers on a track pad to pan.

Methods of scaling viewers:

- Click a viewer, and press the equals key (=) to zoom in, and the minus key (-) to zoom out.
- Press the middle and left mouse buttons simultaneously and drag left or right to resize the viewer.
- Hold down the Command key and use your pointer's scroll control to zoom in and out of the viewer.
- Hold down the middle mouse button, and then click the left mouse button to zoom in, or click the right button to zoom out. The scaling uses a fixed amount, centered on the position of the cursor.
- Click a viewer and press Command-1 to resize the image in the viewer to 100 percent.
- Click a viewer and press Command-2 to resize the image in the viewer to 200 percent.
- Click a viewer and press F to reset the image in the viewer to fit the viewer.
- Click the Scale viewer menu and choose Fit or a percentage.
- Right-click on a viewer and choose an option from the Scale submenu of the contextual menu. This includes a Custom Scale command that lets you type your own scale percentage.
- Hold down the Command key and drag with two fingers on a track pad to zoom in and out of the viewer.

Methods of spinning 3D viewers:

- In 3D Perspective view, hold down the middle and right mouse button, then drag to spin the stage around.
- In 3D Perspective view, hold down the Shift key and drag with two fingers on a track pad to spin the stage around.

Loading Nodes Into Viewers

When you first open the Fusion page in DaVinci Resolve, the output of the current empty composition (the MediaOut1 node) is usually showing in viewer 2. If you're in Dual-viewer mode, viewer 1 remains empty until you assign a node to one of them.

When using Fusion Studio, nothing is loaded into either of the viewers until you assign a node to one of them.

To load specific nodes into specific viewers:

- Hover the pointer over a node, and click one of two buttons that appear at the bottom left of the node.
- Click once to select a node, and press 1 (for the left viewer) or 2 (for the right viewer).
- Right-click a node and choose View On > None/Left View/Right View in the contextual menu.
- Right-click the control header of a node in the Inspector, and choose View On > None/Left View/Right View from the contextual menu.
- Drag a node and drop it over the viewer you'd like to load it into (this is great for tablet users).

When a node is being viewed, a View Indicator button appears at the bottom left. This is the same control that appears when you hover the pointer over a node. Not only does this control let you know which nodes are loaded into which viewer, but they also expose little round buttons for switching between viewers.



Viewer assignment buttons at the bottom of nodes indicate when they're being viewed.

Clearing Viewers

To clear an image from a viewer, click in the viewer to make it active; a thin red highlight is displayed at the top of the active viewer. With the viewer active, press the ` (accent) key. This key is usually found to the left of the 1 key on U.S. keyboards. The fastest way to remove all the images from all the viewers is to make sure none of the viewers is the active panel, and then press the Tilde key.

You can also select the node that is currently showing in the viewer, and press the viewer number again (1 or 2 respectively) to clear the viewer.

Viewer Controls

A series of buttons and pop-up menus in the viewer's title bar provides several quick ways of customizing the viewer display.



Controls in the viewer title bar

- **Zoom menu:** Lets you zoom in on the image in the viewer to get a closer look, or zoom out to get more room around the edges of the frame for rotoscoping or positioning different layers. Choose Fit to automatically fit the overall image to the available dimensions of the viewer.
- **Split Wipe button and A/B Buffer menu:** You can actually load two nodes into a single viewer using that viewer's A/B buffers by choosing a buffer from the menu and loading a node into the viewer. Turning on the Split Wipe button (press Forward Slash) shows a split wipe between the two buffers, which can be dragged left or right via the handle of the onscreen control, or rotated by dragging anywhere on the dividing line on the onscreen control. Alternatively, you can switch between each full-screen buffer to compare them (or to dismiss a split-screen) by pressing Comma (A buffer) and Period (B buffer).

- **SubView type:** (These aren't available in 3D viewers.) Clicking the icon itself enables or disables the current "SubView" option you've selected, while using the menu lets you choose which SubView is enabled. This menu serves one of two purposes. When displaying ordinary 2D nodes, it lets you open up SubViews, which are viewer "accessories" within a little pane that can be used to evaluate images in different ways. These include an Image Navigator (for navigating when zoomed far into an image), Magnifier, 2D viewer (a mini-view of the image), 3D Histogram scope, Color Inspector, Histogram scope, Image Info tooltip, Metadata tooltip, Vectorscope, or Waveform scope. The Swap option (Shift-V) lets you switch what's displayed in the viewer with what's being displayed in the Accessory pane. When displaying 3D nodes, this button lets you have access to an additional mini 3D viewer.
- **Node name:** The name of the currently viewed node is displayed at the center of the viewer's title bar.
- **ROI controls:** Clicking the icon itself enables or disables RoI (Region of Interest) limiting in the viewer, while using the menu lets you choose the region of the RoI. RoI lets you define the region of the viewer in which pixels actually need to be updated. When a node renders, it intersects the current RoI with the current Domain of Definition (DoD) to determine what pixels should be affected. When enabled, you can position a rectangle to restrict rendering to a small region of the image, which can significantly speed up performance when you're working on very high resolution or complex compositions. Auto (the default) sets the region to whatever is visible at the current zoom/pan level in the viewer. Choosing Set lets you draw a custom region within the frame by dragging a rectangle that defaults to the size of the viewer, which is resizable by dragging the corners or sides of the onscreen control. Choosing Lock prevents changes from being made to the current RoI. Choosing Reset resets the RoI to the whole viewer.
- **Color controls:** Lets you choose which color and/or image channels to display in the viewer. Clicking the icon itself toggles between Color (RGB) and Alpha, the two most common things you want to see (pressing C or A also toggles between Color and Alpha). Opening the menu displays every possible channel that can be displayed for the currently viewed node, commonly including RGB, Red, Green, Blue, and Alpha (available from the keyboard by pressing R, G, B, or A). For certain media and nodes, additional auxiliary channels are available to be viewed, including Z-depth, Object ID, Material ID, XYZ Normals, and so on.
- **Viewer LUT:** Clicking the icon itself toggles LUT (LookUp Table) display on or off, while the menu lets you choose which of the many available color space conversions to apply to the viewer. The top options let you choose Fusion controls that can be customized via the Edit item at the top of this menu. The rest of this menu shows all LUTs installed in the LUT directory to use for viewing.

  By default, when using DaVinci Resolve, the viewers in the Fusion page show you the image prior to any grading done in the Color page, since the Fusion page comes before the Color page in the DaVinci Resolve image processing pipeline. When you're working on clips that have been converted to linear color space for compositing, it is desirable to composite and make adjustments to the image relative to a normalized version of the image that appears close to what the final will be. Enabling the LUT display lets you do this as a preview, without permanently applying color adjustments to the image.

- **Option menu:** This menu contains various settings that pertain to the viewers in Fusion.

  Snap to Pixel: When drawing or adjusting a polyline mask or spline, the control points will snap to pixel locations.

  Show Controls: Toggles whatever onscreen controls are visible for the currently selected node.

  Region: Provides all the settings for the Region of Interest in the viewer.

  Smooth Resize: This option uses a smoother bilinear interpolated resizing method when zooming into an image in the viewer; otherwise, scaling uses the nearest neighbor method and shows noticeable aliasing artifacts. However, this is more useful when you zoom in at a pixel level since there is no interpolation.

Show Square Pixels: Overrides the auto aspect correction when using formats with non-square pixels.

Checker Underlay: Toggles a checkerboard underlay that makes it easy to see areas of transparency.

Normalized Color Range: Allows for the visualization of brightness values outside of the normal viewing range, particularly when working with floating-point images or auxiliary channels.

Gain/Gamma: Exposes a simple pair of Gain and Gamma sliders that let you adjust the viewer's brightness.

360 View: Used to properly display spherical imagery in a variety of formats, selectable from this submenu.

Stereo: Used to properly display stereoscopic imagery in a variety of formats, selectable from this submenu.

Time Ruler and Transport Controls

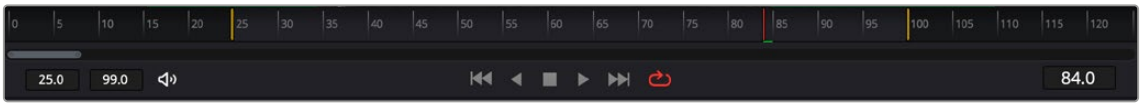
The Time Ruler, located beneath the viewer area, is based on the total duration of the composition. What it represents depends on which version of Fusion you're using:

- For DaVinci Resolve users, the duration displayed in the Time Ruler range depends on what's currently selected in the Edit or Cut page timeline.
- In Fusion Studio, the Time Ruler depends on the Global Start and End values set in the Fusion Studio Preferences > Defaults.

The transport controls under the Time Ruler include playback controls, audio monitoring, as well as number fields for the composition duration and playback range. Additional controls enable motion blur and proxy settings.

Time Ruler Controls in the Fusion Page

If you've selected a single clip in the Edit or Cut page Timeline, then the global range shown in the Time Ruler is based on the total source duration for that clip. You cannot move the playhead outside the global range. The yellow lines, called the render range, identify the current In and Out points for the clip and are the only frames visible in the Fusion page. All frames outside of this range constitute the unused head and tail handles of that source clip.



The Time Ruler displaying ranges for a clip in the Timeline via yellow marks (the playhead is red)

If you've created a Fusion clip or a compound clip, then the "working range" reflects the entire duration of that clip.



The Time Ruler displaying ranges for a Fusion clip in the Timeline

Render Range

The render range determines the range of frames that are visible in the Fusion page and that are used for interactive playback, disk caches, and previews. Frames outside the default render range are not visible in the Fusion page and are not rendered or played.

You can modify the duration of the render range for preview and playback only. Making the range shorter or longer does not trim the clip in the Edit or Cut page Timelines.

You can change the render range in the Time Ruler by doing one of the following:

- Hold down the Command key and drag a new range within the Time Ruler.
- Drag either the start or end yellow line to modify the start or end of the range.
- Right-click within the Time Ruler and choose Set Render Range from the contextual menu.
- Enter new ranges in the Range In and Out fields to the left of the transport controls.
- Drag a node from the Node Editor to the Time Ruler to set the range to the duration of that node.

You can return the render range to the In and Out points of the timeline clip by doing one of the following.

- Right-click within the Time Ruler and choose Auto Render Range.
- Click back in the Edit or Cut page, and then return to the Fusion page.

Time Ruler Controls in Fusion Studio

The Time Ruler, located beneath the viewer area, displays two different frame ranges: one for the entire composition, called the global range, and the other called the render range, which determines what to render and what to cache in memory for previews. The global start and end range takes up the entire Time Ruler and sets the total duration of a composition. You cannot move the playhead outside the global range.



The Time Ruler displaying ranges for a clip in the Timeline via yellow marks (the playhead is red)

Global Start and End Range

The global start and end range is simply the total duration of the current composition.

You can change the global range by doing one of the following:

- To change the global range for all new compositions, choose Fusion Studio > Preferences on macOS or File > Preferences on Windows or Linux. In the Global and Default Settings panel, enter a new range in the Global range fields.
- To change the Global range for the current composition, enter a new range in the Global Start and End fields to the left of the transport controls.
- Dragging a node from the Node Editor to the Time Ruler automatically sets the Global and Render Range to the extent of the node.

Render Range

The render range determines the range of frames used for interactive playback, disk caches, and previews. Frames outside the render range are not rendered or played, although you can still drag the playhead to these frames to see the unused frames.

To preview or render a specific range of a composition, you can modify the render range in a variety of ways.

You can set the render range in the Time Ruler by doing one of the following:

- Hold down the Command key and drag a new range within the Time Ruler.
- Right-click within the Time Ruler and choose Set Render Range from the contextual menu to set the Render Range based on the selected Node's duration.
- Enter new ranges in the Range In and Out fields to the left of the transport controls.
- Drag a node from the Node Editor to the Time Ruler to set the range to the duration of that node.

The Playhead

A red playhead within the Time Ruler indicates the currently viewed frame. Clicking anywhere within the Time Ruler jumps the playhead to that frame, and dragging within the Time Ruler drags the playhead within the available duration of that clip or composition.

Zoom and Scroll Bar

A two-handled gray scroll bar lets you zoom into the range shown by the Time Ruler, which is useful if you're using a very large Global range such that the Render range is a tiny sliver in the Time Ruler. Dragging the left or right handles of this bar zooms relative to the opposite handle, enlarging the width of each displayed frame. Once you've zoomed in, you can drag the scroll bar left or right to scroll through the composition.

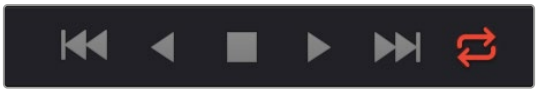
> **TIP:** Holding the middle mouse button and dragging in the Time Ruler lets you scroll the visible range.

Transport Controls in the Fusion Page

The Transport Controls in DaVinci Resolve's Fusion page include buttons that control playback as well as time fields on the left side for setting the render range and the current time on the right side. Additional controls are available in the right-click menu.

Controlling Playback

There are six transport controls underneath the Time Ruler in the Fusion page. These buttons include Composition First Frame, Play Reverse, Stop, Play Forward, Composition Last Frame, and Loop.



The Fusion page controls for playback

Navigation Shortcuts

Many standard transport control keyboard shortcuts you may be familiar with work in Fusion, but some are specific to Fusion's particular needs.

To move the playhead in the Time Ruler using the keyboard, do one of the following:

Spacebar: Toggles forward playback on and off.

JKL: Basic JKL playback is supported, including J to play backward, K to stop, and L to play forward.

Back Arrow: Moves 1 frame backward.

Forward Arrow: Moves 1 frame forward.

Shift-Back Arrow: Moves to the clip's Global Start frame.

Shift-Forward Arrow: Moves to the clip's Global End frame.

Command-Back Arrow: Jumps to the Render Range In point.

Command-Forward Arrow: Jumps to the Render Range Out point.

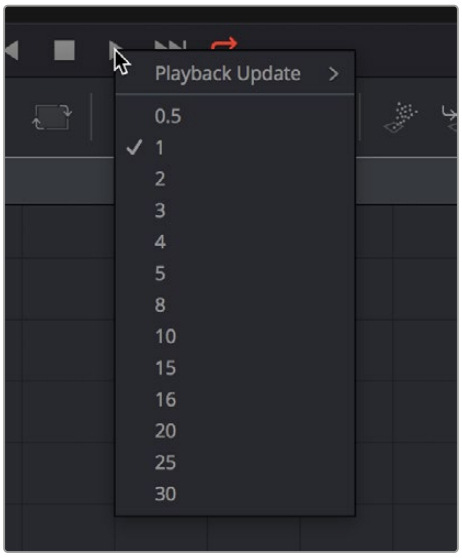
Real-Time Playback Not Guaranteed

Because many of the effects you can create in the Fusion page are processor-intensive, there is no guarantee of real-time playback at your project's full frame rate unless you've cached your composition first (discussed later).

Frame Increment Options

Right-clicking either the Play Reverse or Play Forward buttons opens a contextual menu. This menu contains options to set a frame increment value, which lets you use a keyboard shortcut to move the playhead in sub-frame or multi-frame increments.

Moving the playhead in multi-frame increments can be useful when rotoscoping. Moving the playhead in sub-frame increments can be useful when rotoscoping or inspecting interlaced frames one field at a time (0.5 of a frame).



Right-click the Play Forward or Play Backward buttons to choose a frame increment in which to move the playhead.

Looping Options

The Loop button can be toggled to enable or disable looping during playback. You can right-click this button to choose the looping method that's used:

Playback Loop: The playhead plays to the end of the Time Ruler and starts from the beginning again.

Ping-pong Loop: When the playhead reaches the end of the Time Ruler, playback reverses until the playhead reaches the beginning of the Time Ruler, and then continues to ping-pong back and forth.

Render Range Fields

The two time fields on the left side of the transport controls are used to modify the Render Range. You can enter time values in frames, or click and drag inside the fields to modify the In and Out of the render range for previews and caching.



The Render Start and Render End time fields

Audio Monitoring

Playing a composition in DaVinci Resolve's Fusion page will play the audio from the Edit or Cut page Timeline. You can choose to hear the audio or mute it using the Audio toolbar button to the left of the transport controls. The audio waveforms are displayed in the Keyframes Editor to assist in the timing of your animations.

TIP: If the Mute button is enabled on any Timeline tracks, audio from those tracks will not be heard in Fusion.

For Fusion Studio, audio can be loaded using the Loader node's Audio tab. The audio functionality is included in Fusion Studio for scratch track (aligning effects to audio and clip timing) purposes. Final renders should almost always be performed without audio. Audio can be heard if it is brought in through a Loader node.

To hear the audio from a specific Loader node:

Right-click over the Speaker icon and choose the file name that contains the audio you want to hear.

Audio Toolbar Button

The Audio button in the toolbar is a toggle that can be used to enable or mute audio playback associated with the clip. Additionally, right-clicking this button displays a contextual menu that can be used to select a MediaIn node in the Fusion page or an external WAV file in Fusion Studio.

The Current Time Field

The Current Time field at the right of the transport controls displays the frame number for the playhead position, which corresponds to the frame seen in the viewer. Clicking and dragging in this field scrubs the playhead position back and forth. However, you can also enter time values into this field to move the playhead by specific amounts.

When setting ranges and entering frame numbers to move to a specific frame, numbers can be entered in sub-frame increments. You can set a range to be –145.6 to 451.75 or set the playhead to 115.22. This can be very helpful when animating parameters because you can set keyframes where they actually need to occur, rather than on a frame boundary, so you get more natural animation. Having sub-frame time lets you use time remapping nodes or just scale keyframes in the Spline view and maintain precision.

The Fusion Page Viewer Quality and Proxy Options

Right-clicking anywhere in the transport control area other than over the Play Forward/Play Reverse buttons lets you turn on and off Fusion quality controls. You can either enable high-quality playback at the expense of more significant processing times or enter various proxy modes that temporarily lower the display quality of your composition to speed processing as you work.

Rendering for final output is always done at the highest quality, regardless of these settings.

High Quality

As you build a composition, often the quality of the displayed image is less important than the speed at which you can work. The High Quality setting gives you the option to either display images with faster interactivity or at final render quality. When you turn off High Quality, complex and time-consuming operations such as area sampling, anti-aliasing, and interpolation are skipped to render the image to the viewer more quickly. Enabling High Quality forces a full-quality render to the viewer that's identical to what is output during final delivery.

Motion Blur

The Motion Blur button is a global setting. Turning off Motion Blur temporarily disables motion blur throughout the composition, regardless of any individual nodes for which it's enabled. This can significantly speed up renders to the viewer. Individual nodes must first have motion blur enabled before this button has any effect.

Proxy

The Proxy setting is a draft mode used to speed processing while you're building your composite. Turning on Proxy reduces the resolution of the images that are rendered to the viewer, speeding render times by causing only one out of every x pixels to be processed, rather than processing every pixel. The value of x is decided by adjusting a slider in the Proxy section in the Fusion > Fusion Settings > General panel.

Auto Proxy

The Auto Proxy setting is a draft mode used to speed processing while you're building your composite. Turning on Auto Proxy reduces the resolution of the image while you click and drag to adjust a parameter. Once you release that control, the image snaps back to its original resolution. This lets you adjust processor-intensive operations more smoothly, without the wait for every frame to render at full quality causing jerkiness. You can set the auto proxy ratio by adjusting a slider in the Proxy section of the Fusion > Fusion Settings > General panel.

Selective Updates

When working in Fusion, only the tools needed to display the images in the viewer are updated. The Selective Update options select the mode used during previews and final renders.

The options are available in the Proxy section of the Fusion > Fusion Settings > General panel.

The three options are:

- **Update All (All):** Forces all the nodes in the current node tree to render. This is primarily used when you want to update all the thumbnails displayed in the Node Editor.
- **Selective (Some):** Causes only nodes that directly contribute to the current image to be rendered. So named because only selective nodes are rendered. This is the default setting.
- **No Update (None):** Prevents rendering altogether, which can be handy for making many changes to a slow-to-render composition.

Transport Controls in Fusion Studio

The transport controls in Fusion Studio include buttons to control playback, time fields on the left side for setting the global range and render range, and a Render button to initiate rendering of the composite. There are also controls on the right side for proxy and motion blur. The time field on the far right is used for the current time.



Fusion Studio transport controls

Controlling Playback

There are eight transport controls underneath the Time Ruler in Fusion Studio. These buttons include Composition First Frame, Step Backward, Play Reverse, Stop, Play Forward, Step Forward, Composition Last Frame, and Loop.



Fusion Studio transport controls

Navigation Shortcuts

Many standard transport control keyboard shortcuts you may be familiar with work in Fusion, but there are some keyboard shortcuts specific to Fusion's particular needs.

To move the playhead in the Time Ruler using the keyboard, do one of the following:

- **Spacebar:** Toggles forward playback on and off.
- **JKL:** Basic JKL playback is supported, including J to play backward, K to stop, and L to play forward.
- **Back Arrow:** Moves 1 frame backward.
- **Forward Arrow:** Moves 1 frame forward.

- **Shift-Back Arrow:** Moves to the clip's Global End frame.
- **Shift-Forward Arrow:** Moves to the clip's Global Start frame.
- **Command-Back Arrow:** Jumps to the Render Range In point.
- **Command-Forward Arrow:** Jumps to the Render Range Out point.

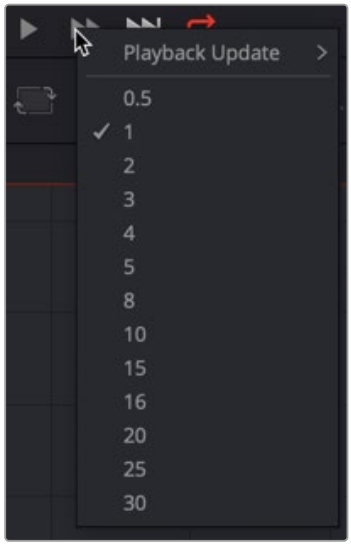
Real-Time Playback Not Guaranteed

Because many of the effects you can create in the Fusion page are processor-intensive, there is no guarantee of real-time playback at your project's full frame rate, unless you've cached your composition first. For more information, see the "Fusion RAM Cache for Playback" section later in this chapter.

Frame Increment Options

Right-clicking the Step Backward, Play Reverse, Play Forward, or Step Forward buttons opens a drop-down menu with options to set a frame increment value. Selecting a frame number from the menu lets you move the playhead in sub-frame or multi-frame increments whenever you use a keyboard shortcut or press the Step Forward/Backward buttons.

Moving the playhead in multi-frame increments can be useful when rotoscoping. Moving the playhead in sub-frame increments can be useful when rotoscoping or inspecting interlaced frames one field at a time (0.5 of a frame).



Right-click the Step Forward or Step Backward buttons to choose a frame increment in which to move the playhead.

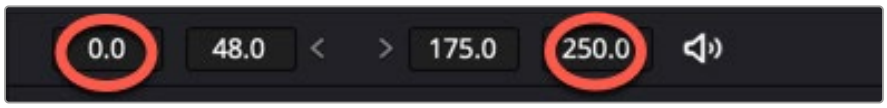
Looping Options

The Loop button can be toggled to enable or disable looping during playback. You can right-click this button to choose the looping method that's used:

- **Playback Loop:** The playhead plays to the end of the Time Ruler and starts from the beginning again.
- **Ping-pong Loop:** When the playhead reaches the end of the Time Ruler, playback reverses until the playhead reaches the beginning of the Time Ruler, and then continues to ping-pong back and forth.

Range Fields

The four time fields on the left side of the transport controls are used to quickly modify the global range and render range in Fusion Studio.



Global Time Fields to the left of the transport controls

Audio

The Audio button is a toggle that mutes or enables any audio associated with the clip. Additionally, right-clicking on this button displays a drop-down menu that can be used to select a WAV file, which can be played along with the composition, and to assign an offset to the audio playback.

Render

Clicking the Render button in the transport controls displays the composition's Render Settings dialog. This dialog is used to configure the render options and initiate rendering of any Saver nodes in the composition. Shift-clicking on the button skips the dialog, using default render values (full resolution, high quality, motion blur enabled).

The Current Time

The Current Time field at the right of the transport controls shows the frame at the position of the playhead, which corresponds to the frame seen in the viewer. However, you can also enter time values into this field to move the playhead by specific amounts.

When setting ranges and entering frame numbers to move to a specific frame, numbers can be entered in sub-frame increments. You can set a range to be –145.6 to 451.75 or set the playhead to 115.22. This can be very helpful when animating parameters because you can set keyframes where they actually need to occur, rather than on a frame boundary, so you get more natural animation. Having sub-frame time lets you use time remapping nodes or just scale keyframes in the Spline view and maintain precision.

NOTE: Many fields in Fusion can evaluate mathematical expressions that you type into them. For example, typing 2 + 4 into most fields results in the value 6.0 being entered. Because Feet + Frames uses the + symbol as a separator symbol rather than a mathematical symbol, the Current Time field will not correctly evaluate mathematical expressions that use the + symbol, even when the display format is set to Frames mode.

Fusion Studio Viewer Quality and Proxy Options

Five buttons along the right side of the transport controls let you either enable high-quality playback at the expense of greater processing times, or enter various proxy modes that temporarily lower the display quality in order to speed processing as you work.

Rendering for final output is always done at the highest quality, regardless of these settings.



Five buttons control the viewer quality, motion blur, proxy options, and image-processing update settings.

HiQ

As you build a composition, often the quality of the displayed image is less important than the speed at which you can work. The High Quality setting gives you the option to either display images with faster interactivity or at final render quality. When you turn off High Quality, complex and time-consuming operations such as area sampling, anti-aliasing, and interpolation are skipped to render the image to the viewer more quickly. Enabling High Quality forces a full-quality render to the viewer that's identical to what will be output during final delivery.

MB

The Motion Blur button is a global setting. Turning off Motion Blur temporarily disables motion blur throughout the composition, regardless of any individual nodes for which it's enabled. This can significantly speed up renders to the viewer. Individual nodes must first have motion blur enabled before this button has any effect.

Prx

A draft mode to speed processing while you're building your composite. Turning on Proxy reduces the resolution of the images that are rendered to the viewer, speeding render times by causing only one out of every x pixels to be processed, rather than processing every pixel. The value of x is decided by adjusting a slider in the General panel of the Fusion Preferences, found under the Fusion menu on macOS or the File menu on Windows and Linux.

Aprx

A draft mode to speed processing while you're building your composite. Turning on Auto Proxy reduces the resolution of the image while you click and drag to adjust a parameter. Once you release that control, the image snaps back to its original resolution. This lets you adjust processor-intensive operations more smoothly, without the wait for every frame to render at full quality causing jerkiness. You can set the auto proxy ratio by adjusting a slider in the General panel of the Fusion Preferences, found under the Fusion menu on macOS or the File menu on Windows and Linux.

Selective Updates

The last of the five buttons on the right of the transport controls is a three-way toggle that determines when nodes update images in the viewer. By default, when working in Fusion, any node needed to display the image in the viewer is updated. The Selective Update button can change this behavior during previews and final renders.

The three options are:

- **Update All (All):** Forces all the nodes in the current node tree to render. This is primarily used when you want to update all the thumbnails displayed in the Node Editor.
- **Selective (Some):** Causes only nodes that directly contribute to the current image to be rendered. So named because only selective nodes are rendered. This is the default setting.
- **No Update (None):** Prevents rendering altogether, which can be handy for making a lot of changes to a slow-to-render composition.

The options are also available in the Fusion Preferences General panel.

Changing the Time Display Format

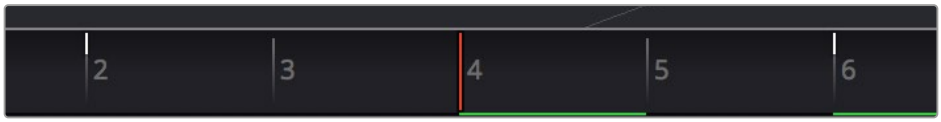
By default, all time fields and markers in Fusion count in frames, but you can also set the time display to SMPTE timecode or Feet + Frames.

To change the time display format:

1. Choose Fusion > Fusion Settings in DaVinci Resolve, choose Fusion Studio > Preferences in Fusion Studio on macOS, or choose File > Preferences in Fusion Studio on Windows or Linux.
2. When the Fusion settings dialog opens, select the Defaults panel and choose a Timecode option.
3. Select the Frame Format panel. If you're using timecode, choose a frame rate and turn on the "has fields" checkbox if your project is interlaced. If you're using feet and frames, set the Film Size value to match the number of frames found in a foot of film in the format used in your project.
4. Click Save.

Keyframe Display in the Time Ruler

When you select a node with keyframed parameters, those keyframes appear in the Time Ruler as little white tic marks, letting you navigate among and edit keyframes without having to open the Keyframes Editor or Spline Editor to see them.



The Time Ruler displaying keyframe marks

To move the playhead in the Time Ruler among keyframes:

- Press Option-Left Bracket ([) to jump to the next keyframe to the left.
- Press Option-Right Bracket (]) to jump to the next keyframe to the right.

The Fusion RAM Cache for Playback

When assembling a node tree, all image processing operations are rendered live to display the final result in the viewers. However, as each frame is rendered, and especially as you initiate playback forward or backward, these images are automatically stored to a RAM cache as they're processed so you can replay those frames in real time. The actual frame rate achieved during playback is displayed in the Status bar at the bottom of the Fusion window during playback. Of course, when you play beyond the cached area of the Time Ruler, uncached frames need to be rendered before being added to the cache.

Priority is given to caching nodes that are currently being displayed, based on which nodes are loaded to which viewers. However, other nodes may also be cached, depending on available memory and on how processor-intensive those nodes happen to be, among other factors.

Memory Limits of the RAM Cache

There is a single setting in DaVinci Resolve for limiting the RAM used for caching. This setting is located in the DaVinci Resolve Preferences Memory and GPU panel.

- **Limit Fusion Memory Cache To:** This slider sets the maximum amount of RAM that Fusion can access for caching. It is a subset of the RAM allocated to DaVinci Resolve. You can assign a maximum of 75% to Fusion from DaVinci Resolve's total RAM allocation. When not using the Fusion page, the RAM is released for other pages in DaVinci Resolve.

There are two settings in Fusion Studio for limiting the RAM used for caching. These settings are located in the Preferences Memory panel.

Limit Caching To: This slider sets the maximum amount of RAM used for caching. The 60% default setting on a 32-GB system limits the cache to 19.2 GB. The maximum amount you can assign to Fusion Studio is limited to 80% of the total system memory. This leaves a minimum amount of memory for other applications and the operating system.

Leave at least # MBytes: This number field further limits caching in cases where the system's available free RAM drops below the entered value. For instance, setting this to 200 MB attempts to keep 200 MB of RAM free for the OS or other applications. Setting the number field to 0 allows Fusion Studio to use the full amount of RAM specified by the Limit Caching To setting, ignoring other apps.

When the size of the cache reaches the Fusion Caching/Memory Limits setting found in the Memory panel of the Preferences, then lower-priority cache frames are automatically discarded to make room for new caching. You can keep track of the RAM cache usage via a percentage indicator on the far right of the Status bar at the bottom of the Fusion window.

Displaying Cached Frames

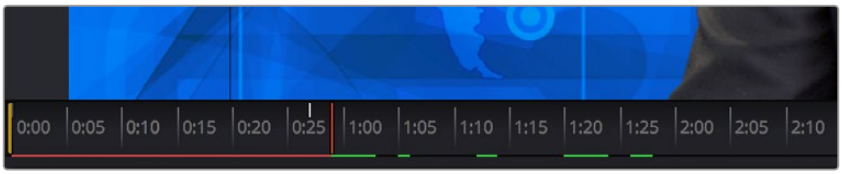
All cached frames for the currently viewed node are indicated by a green line at the bottom of the Time Ruler. Any green section of the Time Ruler should play back in real time.



The green lines indicate frames that have been cached for playback.

Temporarily Preserving the Cache When Changing Quality or Proxy Settings

If you toggle the composition's quality settings or proxy options, the cache is not immediately discarded. The green line instead turns red to let you know the cache is being preserved and can be used again when you go back to the original level of quality or disable proxy mode. However, if you play through those frames at the new quality or proxy settings, this preserved cache is overwritten with a new cache at the current quality or proxy setting.



A red line indicates that cached frames from a different quality or proxy setting are being preserved.

There's one exception to this, however. When you cache frames at the High Quality setting, and you then turn off High Quality, the green frames won't turn red. Instead, the High Quality cached frames are used even though the HiQ setting has been disabled.

Toolbar

The toolbar, located underneath the Time Ruler, contains buttons that let you quickly add commonly used nodes to the Node Editor. Clicking any of these buttons adds that node after the currently selected node in the node tree, or adds an unconnected instance of that node if no nodes are selected. The Toolbar can be customized and saved for specific tasks.



The toolbar has buttons for adding commonly used nodes to the Node Editor.

The default toolbar is divided into sections that group commonly used nodes together. As you hover the pointer over any button, a tooltip shows you that node's name.

- **Loader/Saver nodes (Fusion Studio Only):** The Loader node is the primary node used to select and load clips from the hard drive. The Saver node is used to write or render your composition to disk.
- **Generator/Title/Paint nodes:** The Background and FastNoise generators are commonly used to create all kinds of effects, and the Title generator is obviously a ubiquitous tool, as is Paint.
- **Color/Blur nodes:** ColorCorrector, ColorCurves, HueCurves, and BrightnessContrast are the four most commonly used color adjustment nodes, while the Blur node is ubiquitous.
- **Compositing/Transform nodes:** The Merge node is the primary node used to composite one image against another. ChannelBooleans and MatteControl are both essential for reassigning channels from one node to another. Resize alters the resolution of the image, permanently altering the available resolution, while Transform applies pan/tilt/rotate/zoom effects in a resolution-independent fashion that traces back to the original resolution available to the source image.
- **Mask nodes:** Rectangle, Ellipse, Polygon, and BSpline mask nodes let you create shapes to use for rotoscoping, creating garbage masks, or other uses.
- **Particle system nodes:** Three particle nodes let you create complete particle systems when you click them from left to right. pEmitter emits particles in 3D space, while pMerge lets you merge multiple emitters and particle effects to create more complex systems. pRender renders a 2D result that can be composited against other 2D images.
- **3D nodes:** Seven 3D nodes let you build sophisticated 3D scenes. These nodes auto attach to one another to create a quick 3D template when you click from left to right. ImagePlane3D lets you connect 2D stills and movies for compositing into 3D scenes. Shape3D lets you create geometric primitives of different kinds. Text3D lets you build 3D text objects. Merge3D lets you composite multiple 3D image planes, primitive shapes, and 3D text together to create complex scenes, while Camera3D lets you frame the scene in whatever ways you like. SpotLight lets you light the scenes in different ways and Renderer3D renders the final scene and outputs 2D images and auxiliary channels that can be used to composite 3D output against other 2D layers.

When you're first learning to use Fusion, these nodes are really all you need to build most common composites. Once you've become a more advanced user, you'll still find that these are truly the most common operations you'll use.

Customizing the Toolbar

You can add and remove tools from the Fusion toolbar and then save the custom toolbar as a preset. New tools can be added by dragging them from the Effects Library or the Node Editor, and dividers can be added to group tool sets together.

To create a new toolbar, do the following:

1. Right click in an empty area of the toolbar and choose Customize > Create Toolbar from the contextual menu.
2. Enter a name for the toolbar in the dialog box, and click OK.

To rearrange the tools in the toolbar, do the following:

1. Create a new custom toolbar or select an existing custom toolbar.
2. Drag a node in the toolbar to a new location.

To add a tool to the toolbar:

1. Create a new custom toolbar or select an existing custom toolbar.
2. Do one of the following:
   - Drag a node from the Effects Library to the location on the toolbar where you want it added.
   - Drag a node from the Node Editor to the location on the toolbar where you want it added.

To add a divider to a toolbar, do the following:

1. Create a new custom toolbar or select an existing custom toolbar.
2. Right-click over any tool and choose Customize > Add Divider. A divider is added to the right of the tool.

To remove a tool from a toolbar, do the following:

1. Create a new custom toolbar or select an existing custom toolbar.
2. Right-click over any tool and choose Remove [name of tool].

To remove a group of tools between two dividers, do the following:

1. Create a new custom toolbar or select an existing custom toolbar.
2. Right-click over any tool in a group and choose Remove Group.

To prevent a custom toolbar from being modified:

Right-click over the toolbar and choose Lock from the contextual menu.

To switch between toolbars:

Right-click over the toolbar and choose the custom toolbar name or choose Default to return to Fusion's default toolbar.

To rename a custom toolbar, do the following:

1. Right-click over the toolbar and choose the name of the custom toolbar you want to rename.
2. Right-click over the toolbar again and choose Customize > Rename [toolbar name].
3. Enter a new name for the toolbar.

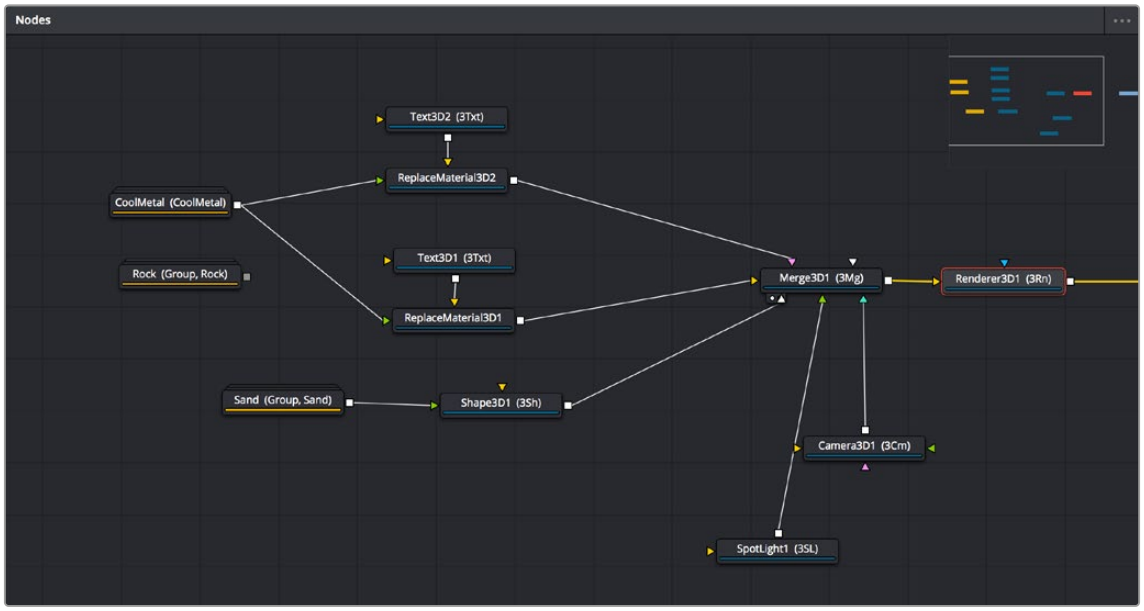
To delete a custom toolbar, do the following:

1. Right-click over the toolbar and choose the name of the custom toolbar you want to delete.
2. Right-click over the toolbar again and choose Customize > Remove [toolbar name].

TIP: Adding and deleting tools from a custom toolbar is not undoable. If you are creating a complex toolset, make a new custom toolbar based on your current toolbar in between major changes and work off that. That way if you make an error, you can revert back to the last known good toolbar. Once you have the final toolbar the way you want it, you can go back and remove all the interim custom toolbars you made.

Node Editor

The Node Editor is the heart of Fusion because it's where you build the tree of nodes that makes up each composition. Each node you add to the node tree adds a specific operation that creates one effect, whether it's blurring the image, adjusting color, painting strokes, drawing and adding a mask, extracting a key, creating text, or compositing two images into one.



The Node Editor displaying a node tree creating a composition

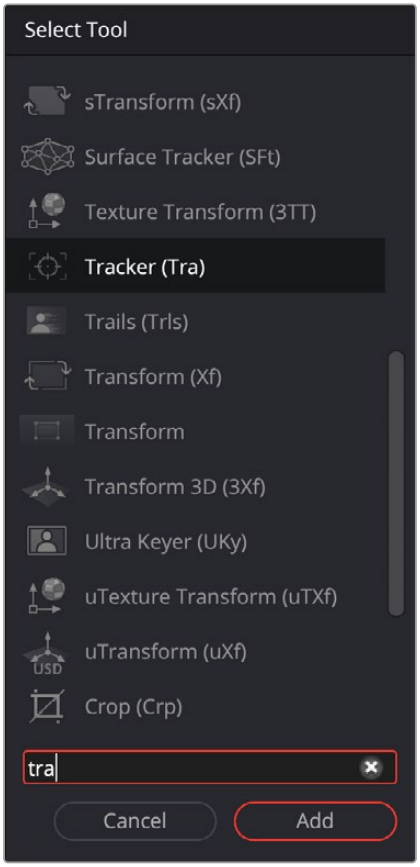
You can think of each node as a layer in an effects stack, except that you have the freedom to route image data in any direction to branch and merge different segments of your composite in completely nonlinear ways. This makes it easy to build complex effects, but it also makes it easy to see what's happening, since the node tree doubles as a flowchart that clearly shows you everything that's happening, once you learn to read it.

Adding Nodes to Your Composition

Depending on your mood, there are a few ways you can add nodes from the Effects Library to your composition. For most of these methods, if there's a single selected node in the Node Editor, new nodes are automatically added after it, but if there are no selected nodes or multiple selected nodes, then new nodes are added as disconnected from anything else.

Methods of adding nodes include:

- Click a button in the toolbar.
- Open the Effects Library, find the node you want in the relevant category, and click once on a node you'd like to add.
- Right-click on a node and choose Insert Tool from the drop-down menu to add it after the node you've right-clicked on. Or, you can right-click on the background of the Node Editor to use that submenu to add a disconnected node.
- Press Shift-Spacebar to open a Select Tool dialog, type characters corresponding to the name of the node you're looking for, and press the Return key (or click OK) when it's found. Once you learn this method, it'll probably become one of your most frequently used ways of adding nodes.



The Select Tool dialog lets you find any node quickly if you know its name.

Removing Nodes from Your Composition

Removing nodes is as simple as selecting one or more nodes, and then pressing the Delete or Backspace keys.

Identifying Node Inputs and Node Outputs

Each node displays small colored connections around the edges. One or more arrows represent inputs, and the square connection represents the tool's processed output, of which there is always only one. If you hover the pointer over any of a node's inputs or output, the name of that input or output immediately appears in the Status bar. If you wait for a few more moments, a floating tooltip displays the same name right over the node.

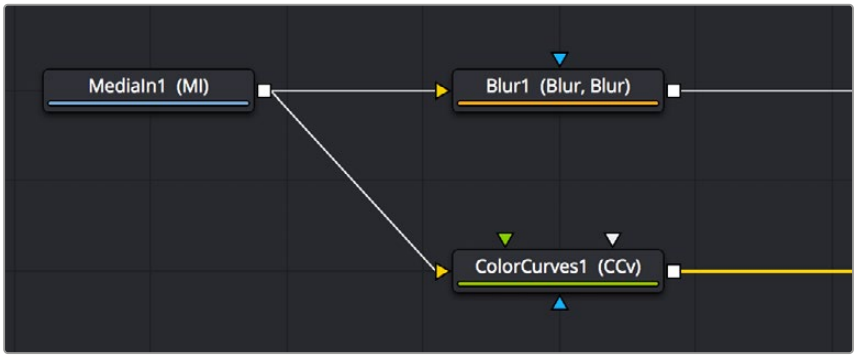
Node Editing Essentials

Each node has inputs and outputs that are "wired together" using connections. The inputs are represented by arrows that indicate the flow of image data from one node to the next, as each node applies its effect and feeds the result (via the square output) to the next node in the tree. In this way, you can quickly build complex results from a series of relatively simple operations.



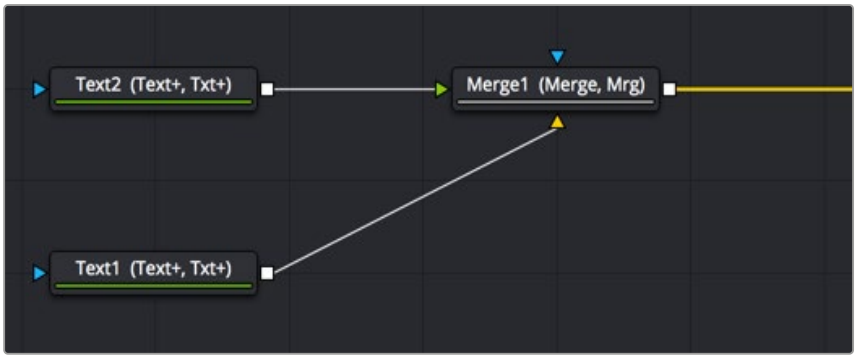
Three nodes connected together

You can connect a single node's output to the inputs of multiple nodes (called "branching").



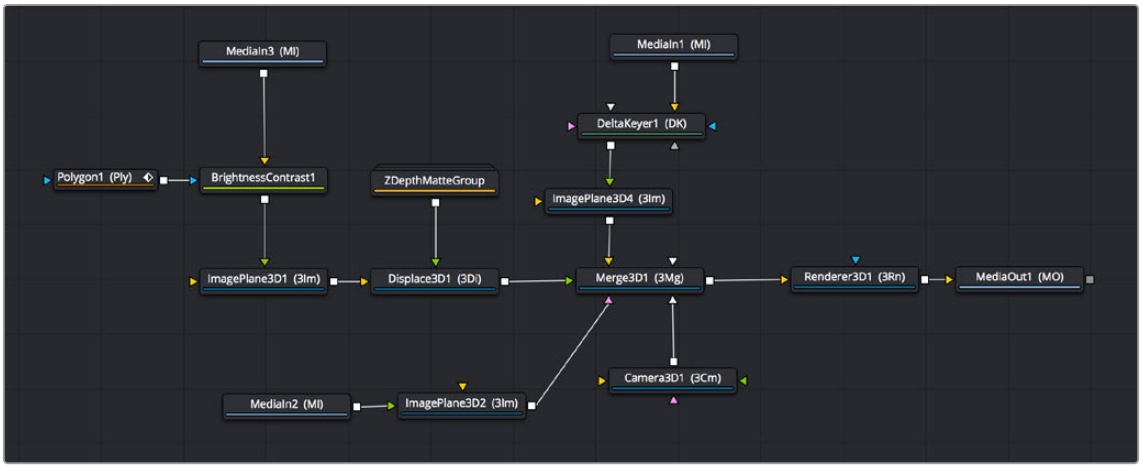
One node branching to two to split the image to two operations

You can then composite images together by connecting the output from multiple nodes to certain nodes such as the Merge node that combines multiple inputs into a single output.



Two nodes being merged together into one to create a composite

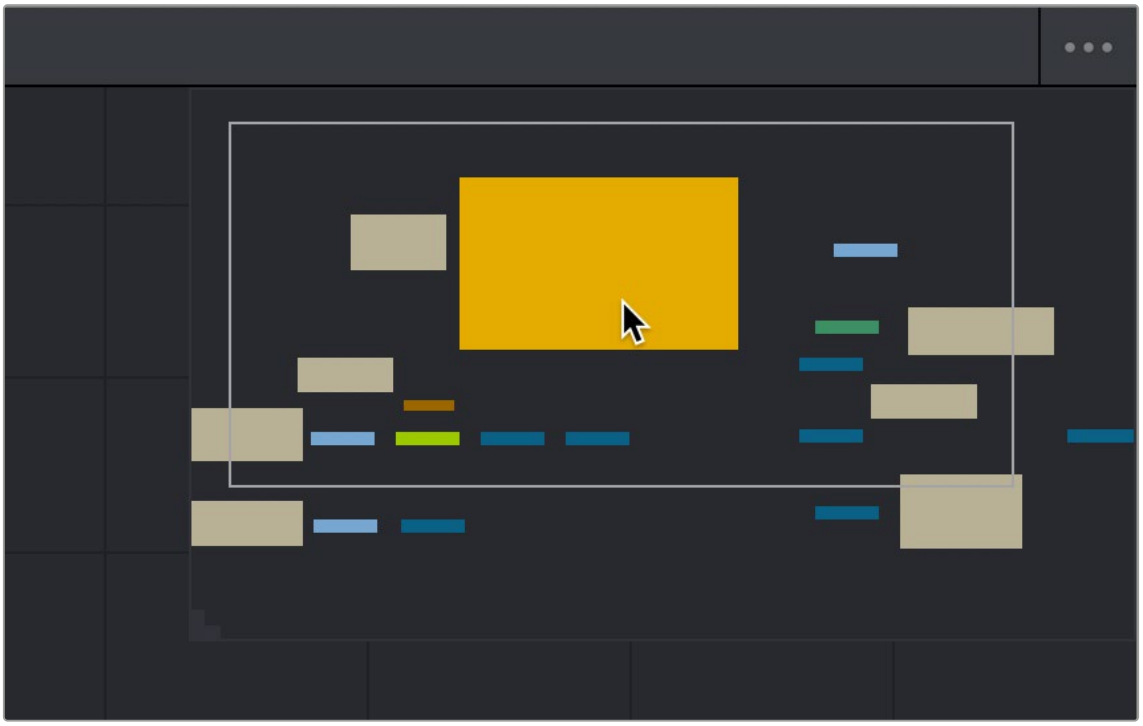
By default, new nodes are added from left to right in the Node Editor, but they can also flow from top to bottom, right to left, bottom to top, or in all directions simultaneously. Connections automatically reorient themselves along all four sides of each node to maintain the cleanest possible presentation as you rearrange other connected nodes.



Nodes can be oriented in any direction; the input arrows let you follow the flow of image data.

Navigating the Node Editor

As your Node tree gets larger, parts inevitably go offscreen. When a portion of the node tree is offscreen, a resizable Navigator pane appears in the upper-right corner. The Navigator is a miniature representation of the entire node tree that you can drag within to pan to different parts of your composition quickly. You can resize the navigator using a handle in the lower-left corner, and you can choose to show or hide the Navigator by pressing the V key, or by right-clicking the Node Editor to access the Options submenu of the contextual menu.



The Navigator pane for accessing offscreen parameters or tools

There are other standard methods of panning and zooming around the Node Editor.

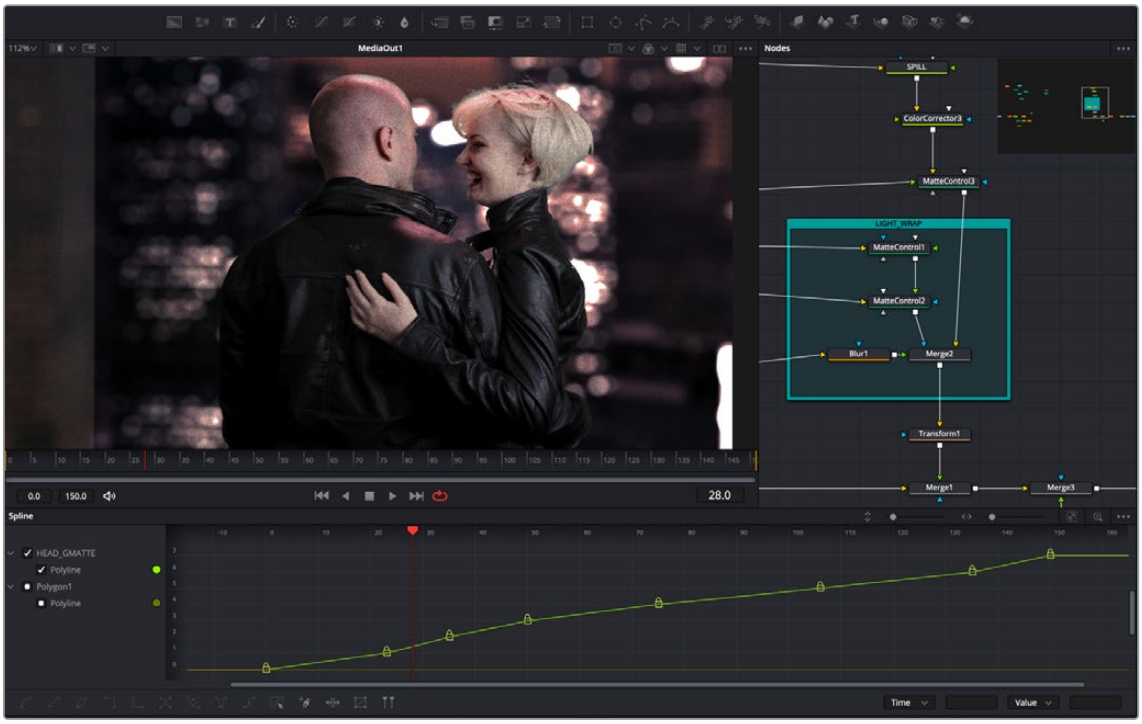
Methods of navigating the Node Editor:

- Middle-click and drag to pan around the Node Editor.
- Hold down Shift and Command and drag the Node Editor to pan.
- Press the Middle and Left mouse buttons simultaneously and drag to resize the Node Editor.
- Hold down the Command key, and use your mouse's scroll control to resize the Node Editor.
- Right-click the Node Editor and choose an option from the Scale submenu of the contextual menu.
- Press Command-1 to reset the Node Editor to its default size.
- Drag two fingers on a track pad to pan.
- Hold the Command key down and drag two fingers on a track pad to resize the Node Editor.

Vertical Node Editor Layouts

Alternative Node View Layout presets located in the Fusion page allow for positioning the Node Editor vertically, either alongside the Inspector, or along the left side of the screen. This can be very helpful when animating in the Spline Editor or Keyframes Editor.

When in the Fusion page, you can choose the layouts from Workspace > Layout Presets. Choosing a vertical layout allows the node tree to flow from top to bottom, leaving much more room along the lower half of the screen for the Spline Editor or Keyframes Editor.



The Mid Flow Vertical layout preset used with the Vertical Flow direction setting.

When using the vertical layouts, enabling the Flow > Build Direction > Vertical option in the Fusion settings will cause all new Node trees to build vertically, leaving maximum room for Fusion's animation tools.

You can then save alternative layouts based on these two vertical presets using the Workspace > Layout Presets submenu.

When you want to return to the default horizontal Node Editor layout, just choose Workspace > Layout Presets > Fusion Presets > Default.

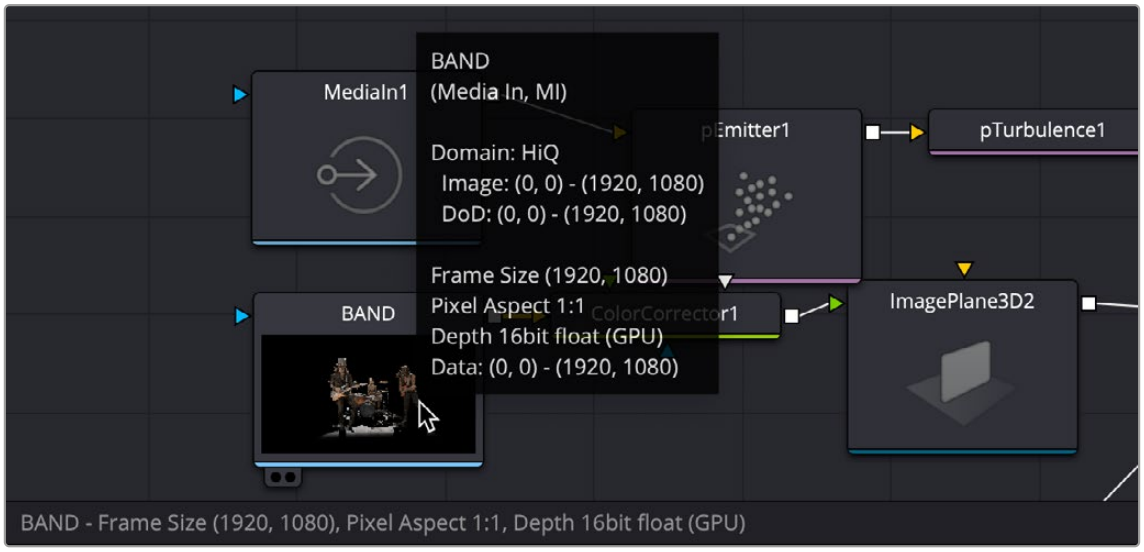
These Layout options are not available in Fusion Studio, however, you can use the Floating Frame to position the Node Editor wherever you like.

Keeping Organized

As you work, it's important to keep the node trees that you create tidy to facilitate a clear understanding of what's happening. Fortunately, the Fusion Node Editor provides a variety of methods and options to help you with this, found within the Options and Arrange Tools submenus of the Node Editor contextual menu.

Status Bar

The Status bar in the lower-left corner of the Fusion window shows you a variety of up-to-date information about things you're selecting and what's happening in Fusion. For example, hovering the pointer over a node displays information about that node in the Status bar. Additionally, the currently achieved frame rate appears whenever you initiate playback, and the percentage of the RAM cache that's used appears at all times. Other information, updates, and warnings appear in this area as you work.



The Status bar under the Node Editor showing you information about a node under the pointer, as well as a popup box with even more relevant info.

Occasionally the Status bar will display a badge to let you know there's a message in the console you might be interested in. The message could be a log, script message, or error.



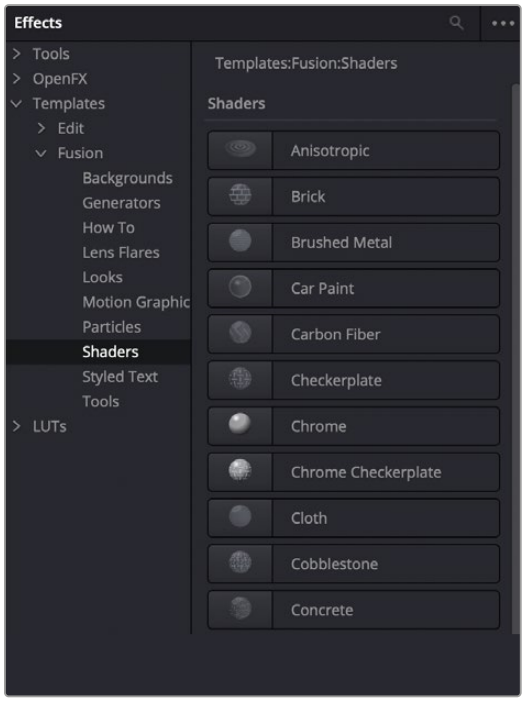
A notification that there's a message in the Console

Effects Library

The Effects Library in Fusion shows all the nodes and effects available in Fusion, including third-party OFX plugins, if installed. If you are using DaVinci Resolve, Resolve FX also appear in the OFX category. While the toolbar shows many of the most common nodes you'll use in any composite, the Effects Library contains every single tool available in Fusion, organized by category, with each node ready to be quickly added to the Node Editor. Suffice it to say that there are many, many more nodes available in the Effects Library than on the toolbar, spanning a wide range of uses.



The Effects Library with Tools open



The Templates section of the Effects Library

The hierarchical category browser of the Effects Library is divided into several sections depending on whether you are using Fusion Studio or the Fusion page within DaVinci Resolve. The Tools section is the most often used since it contains every node that represents an elemental image-processing operation in Fusion. Hovering the pointer over a specific tool will reveal a tool-tip explaining its functionality at the bottom right of the DaVinci Resolve interface. The OpenFX section contains third-party plugins, and if you are using the Fusion page, it also contains ResolveFX, which are included with DaVinci Resolve. A third section, only visible when using the Fusion page in DaVinci Resolve, is the Templates section. The Template section contains a variety of additional content including templates for Lens Flares, Backgrounds, Generators, Particle Systems, Shaders (for texturing 3D objects), and other resources for use in your composites.

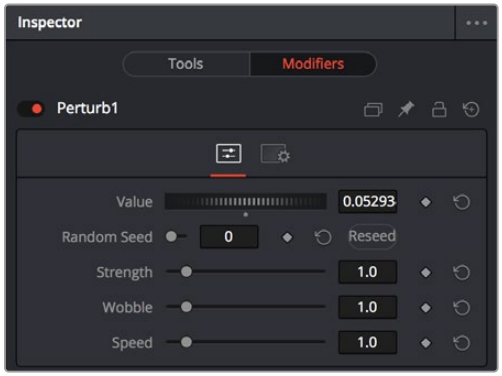
> The Effects Library's list can be made full height or half height using a button at the far left of the UI toolbar.

The Inspector

The Inspector is a panel on the right side of the Fusion window that you use to display and manipulate the parameters of one or more selected nodes. When a node is selected in the Node Editor, its parameters and settings appear in the Inspector.



The Inspector shows parameters from one or more selected nodes.



The Modifier panel showing a Perturb modifier

The Tools and Modifiers Panels

The Fusion Inspector is divided into two panels. The Tools panel is the main panel that shows you the parameters of selected nodes. The Modifiers panel shows optional extensions to the standard parameters of a tool. In the following image, a Perturb modifier has been added to a parameter to add random animation to that parameter, and the controls found on the Modifier panel let you customize what kind of randomness is being added.

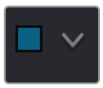
> Other nodes display node-specific items here. For example, Paint nodes show each brush stroke as an individual set of controls in the Modifiers panel, available for further editing or animating.

Parameter Header Controls

A cluster of controls appears at the top of every node's controls in the Inspector.



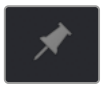
Common Inspector controls



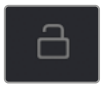
Set Color: A pop-up menu that lets you assign one of 16 colors to a node, overriding a node's own color.



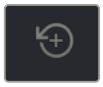
Versions: Clicking Versions reveals another toolbar with six buttons. Each button can hold an individual set of adjustments for that node that you can use to store multiple versions of an effect.



Pin: The Inspector is also capable of simultaneously displaying all parameters for multiple nodes you've selected in the Node Editor. Furthermore, a Pin button in the title bar of each node's parameters lets you "pin" that node's parameters into the Inspector so that they remain there even when that node is deselected, which is valuable for key nodes that you need to adjust even while inspecting other nodes of your composition.



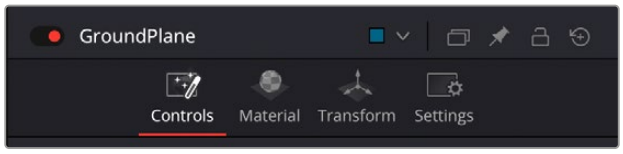
Lock: Locks that node so that no changes can be made to it.



Reset: Resets all parameters within that node.

Parameter Tabs

Many nodes expose multiple tabs' worth of controls in the Inspector, seen as icons at the top of the parameter section for each node. Click any tab to expose that set of controls.



Nodes with several tabs' worth of parameters

Keyframes Editor

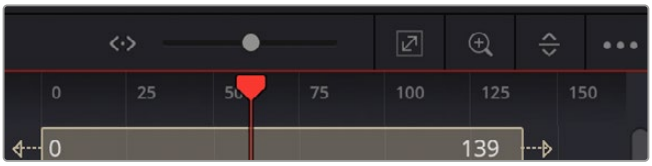
The Keyframes Editor displays each node in the current composition as a stack of layers within a miniature timeline. The order of the layers is largely irrelevant as the order and flow of connections in the node tree dictate the order of image-processing operations. You use the Keyframes Editor to trim, extend, or slide Loader, MediaIn, and effects nodes, or to adjust the timing of keyframes, which appear superimposed over each effect node unless you open them up into their editable track.



The Keyframes Editor is used to adjust the timing of clips, effects, and keyframes.

Keyframes Editor Control Summary

At the top, a series of zoom and framing controls let you adjust the work area containing the layers.



These controls let you manipulate the Keyframes Editor view.

A Horizontal zoom control lets you scale the size of the editor.

- A Zoom to Fit button fits the width of all layers to the current width of the Keyframes Editor.
- A Zoom to Rect tool lets you draw a rectangle to define an area of the Keyframes Editor to zoom into.
- A Sort pop-up menu lets you sort or filter the tracks in various ways.
- An Option menu provides access to many other ways of filtering tracks and controlling visible options.

A timeline ruler provides a time reference, as well as a place in which you can scrub the playhead.

At the left, a track header contains the name of each layer, as well as controls governing that layer.

- A lock button lets you prevent a particular layer from being changed.
- Nodes that have been keyframed have a disclosure control, which when opened displays a keyframe track for each animated parameter.

In the middle, the actual editing area displays all layers and keyframe tracks available in the current composition.

At the bottom-left, Time Stretch and Spreadsheet mode controls provide additional ways to manipulate keyframes.

At the bottom right, the Time/TOffset/TScale drop-down menu and value fields let you numerically alter the position of selected keyframes either absolutely, relatively, or based on their distance from the playhead.

Adjusting Clip Timings

Each Loader or MediaIn node that represents a clip used in a composition is represented as a layer in this miniature timeline. You can edit a layer's In or Out points by positioning the pointer over the beginning or end of a segment and using the resize cursor to drag that point to a new location. You can slide a layer by dragging it to the left or right, to better line up with the timing of other elements in your composition.

The Keyframes Editor also lets you adjust the timing of elements that you've added from directly within Fusion.

Adjusting Effect Timings

Each effect node also appears as a layer, just like clips. You can resize the In and Out points of an effect layer, and slide the entire layer forward or backward in time, just like a Loader or MediaIn layers. If you trim an effect layer to be shorter than the duration of the composition, the effect cuts in at whichever frame the layer begins, and cuts out after the last frame of that layer, just like a clip on a timeline.

Adjusting Keyframe Timings

When you've animated an effect by adding keyframes to a parameter in the Inspector, the Keyframes Editor is used to edit the timing of keyframes easily. By default, all keyframes applied to parameters within a particular node's layer appear superimposed in one flat track over the top of that layer.

To edit keyframes, you can click the disclosure control to the left of any animated layer's name in the track header, which opens up keyframe tracks for every keyframed parameter within that layer.



Keyframe tracks exposed

Keyframe Editing Essentials

Here's a short list of keyframe editing methods that will get you started.

Methods of adjusting keyframes:

- You can click on a single keyframe to select it.
- You can drag a bounding box over a series of keyframes to select them all.
- You can drag keyframes left and right to reposition them in time.
- You can right-click one or more selected keyframes and use commands from the drop-down menu to change keyframe interpolation, copy/paste keyframes, or even create new keyframes.
- You can Command-drag one or more selected keyframes to drag a duplicate of them to another position in the keyframe track.

To change the position of a keyframe using the toolbar, do one of the following:

- Select a keyframe, and then enter a new frame number in the Time Edit box.
- Choose T Offset from the Time Editor pop-up, select one or more keyframes, and enter a frame offset.
- Choose T Scale from the Time Editor pop-up, select one or more keyframes, and enter a multiplier added to the current playhead frame position. For instance, if the playhead is on frame 10 and the keyframe is on frame 30, entering the TScale value of 2 will position the keyframe on frame 50. The distance between the playhead and original keyframe is 20, so (20 x 2) = 40, which is then added to the playhead position.

Time Stretching Keyframes

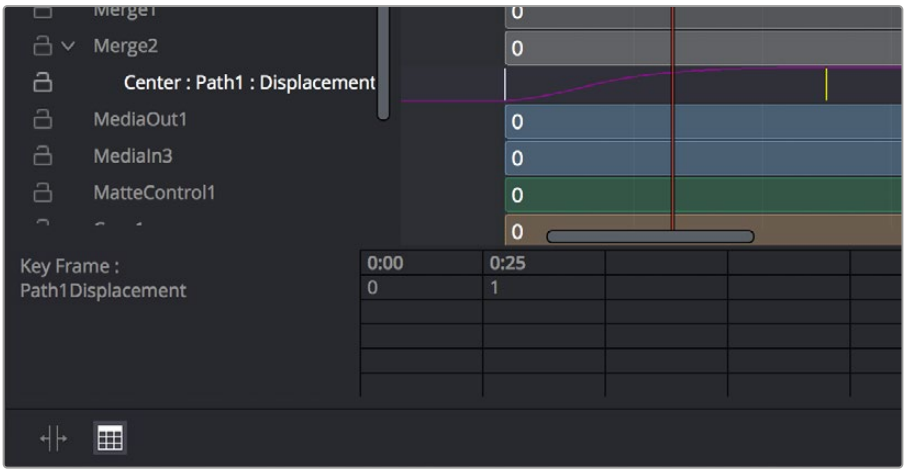
If you select a range of keyframes in a keyframe track, you can turn on the Time Stretch tool to show a box used to squeeze and stretch the entire range of keyframes relative to one another. The Time Stretcher changes the overall timing of a sequence of keyframes without losing the relative timing from one keyframe to the next. Alternatively, you can turn on Time Stretch and draw a bounding box around the keyframes you want to adjust to create a time-stretching boundary that way. Click the Time Stretch tool again to turn it off.



Time stretching keyframes

The Keyframe Spreadsheet

If you turn on the Spreadsheet and then click on the name of a layer in the keyframe track, the numeric time position and value (or values if it's a multi-dimensional parameter) of each keyframe appear as entries in the cells of the Spreadsheet. Each column represents one keyframe, while each row represents a single aspect of each keyframe.



Editing keyframes in the Spreadsheet

For example, if you're animating a motion path, then the "Key Frame" row shows the frame each keyframe is positioned at, and the "Path1Displacement" row shows the position along the path at each keyframe. If you change the Key Frame value of any keyframe, you'll move that keyframe to a new frame of the Timeline.

Spline Editor

The Spline Editor provides a more detailed environment for editing the timing, value, and interpolation of keyframes. Using control points at each keyframe connected by splines (also called curves), you can adjust how animated values change over time. The Spline Editor has four main areas: the Zoom and Framing controls at the top, the Parameter list at the left, the Graph Editor in the middle, and the toolbar at the bottom.



The Spline Editor is divided into the Zoom controls at top, Parameter list at left, Graph Editor, and toolbar.

Spline Editor Control Summary

At the top, a series of Zoom and Framing controls let you adjust the work area containing the curves.

- Vertical and horizontal zoom controls let you scale the size of the editor.
- A Zoom to Fit button fits the width of all curves to the current width of the Spline Editor.
- A Zoom to Rect tool lets you draw a rectangle to define an area of the Spline Editor to zoom into.

A Timeline ruler provides a time reference, as well as a place in which you can scrub the playhead.

The Parameter list at the left is where you decide which splines are visible in the Graph view. By default, the Parameter list shows every parameter of every node in a hierarchical list. Checkboxes beside each name are used to show or hide the curves for different keyframed parameters. Color controls let you customize each spline's tint to make splines easier to see in a crowded situation.

The Graph view that takes up most of this panel shows the animation spline along two axes. The horizontal axis represents time, and the vertical axis represents the spline's value. Selected control points show their values in the edit fields at the bottom of the graph.

Lastly, the toolbar at the bottom of the Spline Editor provides controls to set control point interpolation, spline looping, or choose Spline editing tools for different purposes.

Choosing Which Parameters to Show

Before you start editing splines to customize or create animation, you need to choose which parameter's splines you want to work on.

To show every parameter in every node:

- Click the Splines Editor Option menu and choose Expose All Controls. Toggle this control off again to go back to viewing what you were looking at before.

To show splines for the currently selected node:

- Click the Splines Editor Option menu and choose Show Only Selected Tool.

Essential Spline Editing

The Spline Editor is a deep and sophisticated environment for keyframe and spline editing and retiming, but the following overview will get you started using this tool for creating and refining animation.

To select one or more control points:

- Click any control point to select it.
- Command-click multiple control points to select them.
- Drag a bounding box around multiple control points to select them as a group.

To edit control points and splines:

- Click anywhere on a spline to add a control point.
- Drag one or more selected control points to reshape the spline.
- Shift-drag a control point to constrain its motion vertically or horizontally.

To edit Bézier curves:

- Select any control point to make its Bézier handles visible, and drag the Bézier handles.
- Command-drag a Bézier handle to break the angle between the left and right handles.

To delete control points:

- Select one or more control points and press the Delete or Backspace key.

Essential Spline Editing Tools and Modes

The Spline Editor toolbar at the bottom contains a mix of control point interpolation buttons, Spline loop modes, and Spline editing tools.

Control Point Interpolation

The first six buttons let you adjust the interpolation of one or more selected control points.



Control point interpolation controls



Smooth: Creates automatically adjusted Bézier curves to create smoothly interpolating animation.



Linear: Creates linear interpolation between control points.



Invert: Inverts the vertical position of non-animated LUT splines. This does not operate on animation splines.



Step In: For each keyframe, creates sudden changes in value at the next keyframe to the right. Similar to a hold keyframe in After Effects® or a static keyframe in the DaVinci Resolve Color page.



Step Out: Creates sudden changes in value at every keyframe for which there's a change in value at the next keyframe to the right. Similar to a hold keyframe in After Effects or a static keyframe in the DaVinci Resolve Color page.



Reverse: Reverses the horizontal position of selected keyframes in time, so the keyframes are backward.

Spline Loop Modes

The next three buttons let you set up spline looping after the last control point on a parameter's spline, enabling a limited pattern of keyframes to animate over a far longer duration. Only the control points you've selected are looped.



Spline Loop modes



Set Loop: Repeats the same pattern of keyframes over and over.



Set Ping Pong: Repeats a reversed set of the selected keyframes and then a duplicate set of the selected keyframes to create a more seamless pattern of animation.



Set Relative: Repeats the same pattern of selected keyframes but with the values of each repeated pattern of keyframes being incremented or decremented by the trend of all keyframes in the selection. This results in a loop of keyframes where the value either steadily increases or decreases with each subsequent loop.

Spline Editing Tools

The next five buttons provide specialized Spline editing tools.



Spline editing controls



Select All: Selects every keyframe currently available in the Spline Editor.



Click Append: Click once to select this tool and click again to de-select it. This tool lets you add or adjust keyframes and spline segments (sections of splines between two keyframes), depending on the keyframe mode you're in. With Smooth or Linear keyframes, clicking anywhere above or below a spline segment adds a new keyframe to the segment at the location where you clicked. With Step In or Step Out keyframes, clicking anywhere above or below a line segment moves that segment to where you've clicked.



Time Stretch: If you select a range of keyframes, you can turn on the Time Stretch tool to show a box you can use to squeeze and stretch the entire range of keyframes relative to one another, to change the overall timing of a sequence of keyframes without losing the relative timing from one keyframe to the next. Alternatively, you can turn on Time Stretch and draw a bounding box around the keyframes you want to adjust to create a time-stretching boundary that way. Click Time Stretch a second time to turn it off.



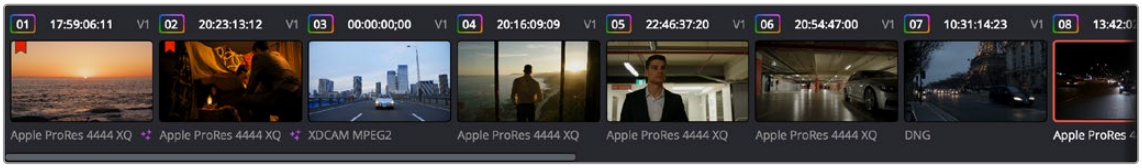
Shape Box: Turn on the Shape Box to draw a bounding box around a group of control points you want to adjust in order to horizontally squish and stretch (using the top/bottom/left/right handles), corner pin (using the corner handles), move (dragging on the box boundary), or corner stretch (Command-drag the corner handles).



Show Key Markers: Turning on this control shows keyframes in the top ruler that correspond to the frame at which each visible control point appears. The colors of these keyframes correspond to the color of the control points they're indicating.

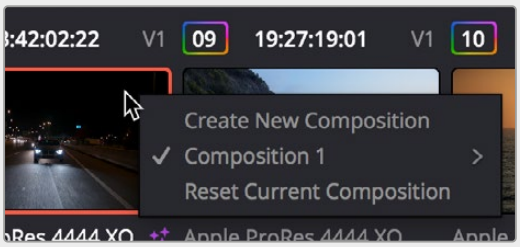
Thumbnail Timeline in the Fusion Page

In the Fusion page of DaVinci Resolve, the Thumbnail timeline (hidden by default) can be opened by clicking the Clips button in the UI toolbar and appears underneath the Node Editor when it's open. The Thumbnail timeline shows every clip in the current Timeline, giving you a way to navigate from one clip to another. Each thumbnail has a pop-up menu for creating and switching among multiple versions of compositions, and resetting the current composition, when necessary.



The Thumbnail timeline lets you navigate the Timeline and manage versions of compositions.

Right-clicking on any thumbnail exposes a contextual menu.



The contextual menu for the Thumbnail timeline

To open another clip:

- Click any thumbnail to jump to that clip's composition. The current clip is outlined in orange.

To navigate the Thumbnail timeline:

- Middle-click and drag the Timeline left or right to scroll through the clips in the Timeline.
- Click and drag the slider at the bottom of the Thumbnail timeline to scroll through the clips in the Timeline.

To create and manage versions of compositions:

- **To create a new version of a composition:** Right-click the current thumbnail, and choose Create New Composition from the contextual menu.
- **To load a different composition:** Right-click the current thumbnail, and choose "NameOfVersion" > Load from the contextual menu.
- **To delete a composition:** Right-click the current thumbnail, and choose "NameOfVersion" > Delete from the contextual menu.
- **To rename a composition:** Right-click the current thumbnail, and choose "NameOfVersion" > Rename from the contextual menu.

To reset the current composition:

- Right-click the current thumbnail, and choose Reset Current Composition from the contextual menu.

To change how thumbnails are identified:

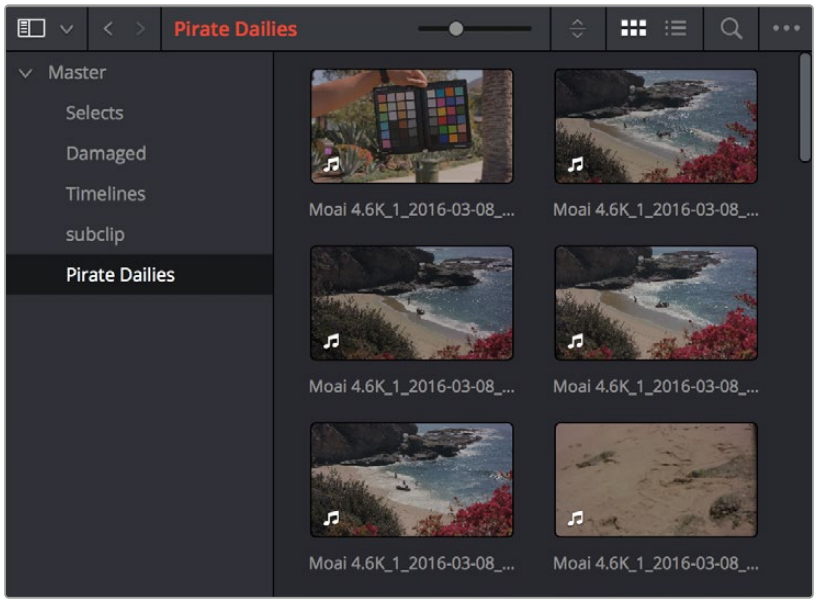
- Double-click the area underneath any thumbnail to toggle among clip format, clip name, and the composition version name.

The Media Pool in the Fusion Page

In DaVinci Resolve's Fusion page, the Media Pool continues to serve its purpose as the repository of all media you've imported into your project. This makes it easy to add additional clips to your compositions simply by dragging the clip you want from the Media Pool into the Node Editor. The media you add appears as a new MediaIn node in your composition, ready to be integrated into your node tree however you need it.

TIP: If you drag one or more clips from the Media Pool onto a connection line between two nodes in the Node Editor, the clips are automatically connected to that line via enough Merge nodes to connect them all.

For more information on using the myriad features of the Media Pool, see Chapter 18, "Adding and Organizing Media with the Media Pool," in the *DaVinci Resolve 20 Reference Manual*.



The Media Pool in Thumbnail mode showing video clips

Importing Media Into the Media Pool on the Fusion Page

If you find yourself in the Fusion page and you need to quickly import a few clips for immediate use, you can do so in a couple of different ways.

To add media by dragging one or more clips from the Finder to the Fusion page Media Pool (macOS only):

1. Select one or more clips in the Finder.
2. Drag those clips into the Media Pool of DaVinci Resolve, or to a bin in the Bin list. Those clips are added to the Media Pool of your project.

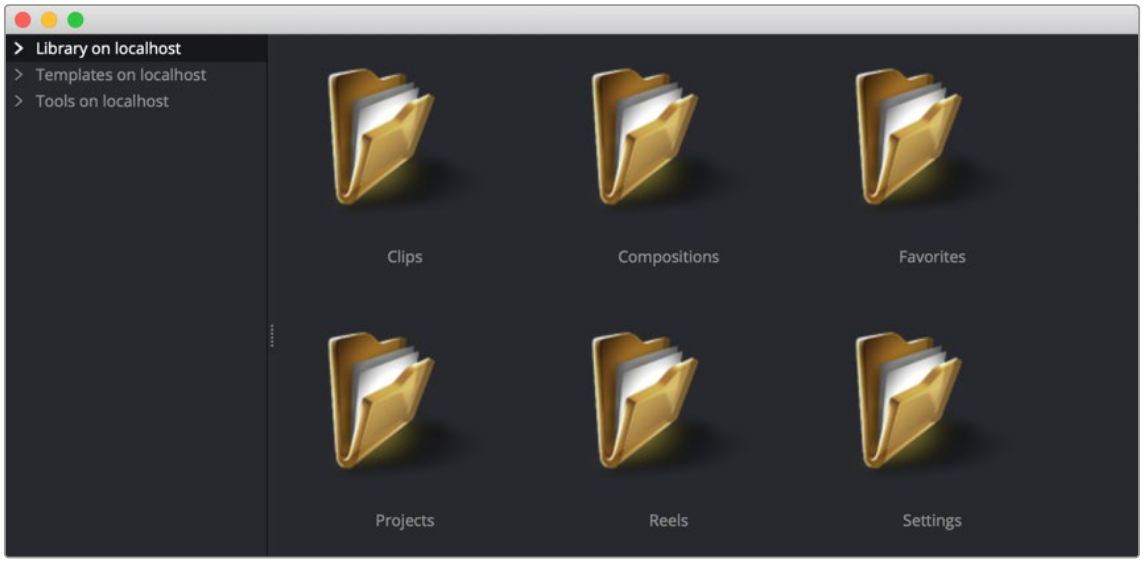
To use the Import Media command in the Fusion page Media Pool:

1. With the Fusion page open, right-click anywhere in the Media Pool, and choose Import Media.
2. Use the Import dialog to select one or more clips to import, and click Open. Those clips are added to the Media Pool of your project.

For more information on importing media using the myriad features of the Media page, *see Chapter 15, "Preferences,"*

Bins in Fusion Studio

Bins in Fusion Studio are similar to the Media Pool in DaVinci Resolve. Bins are organizational panels that provide an easy way of accessing commonly used tools, settings, macros, compositions, and media content. They can keep all your custom content and resources close at hand, so you can use them without searching through your hard drives. Bins can also be shared over a network to improve a collaborative workflow with other Fusion Studio artists.



The Bins window

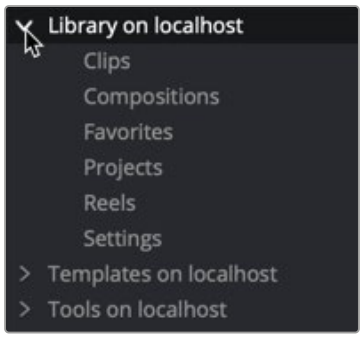
To open the Bins window:

Choose File > Bins from the menu bar.

Similar to the Media Pool in DaVinci Resolve, when adding an item to the Fusion bins, a link is created between the item on disk and the bins. Fusion does not copy the file into its own cache or hard drive space. The file remains in its original format and in its original location.

Bins Interface

The Bins window is actually a separate application used to save content you may want to reuse at a later time. The Bins window is separated into two panels. The sidebar on the left is a bin list where items are placed into categories, while the panel on the right displays the selected bin's content.



The Bin list sidebar

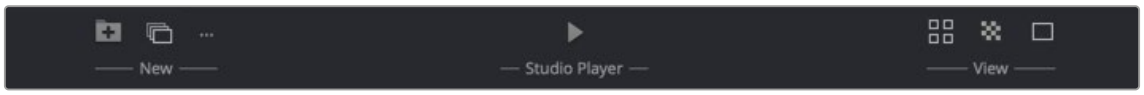
The Bin list organizes content into bins or folders using a hierarchical List view. These folders can be organized to suit your workflow, but standard folders are provided for Clips, Compositions, Favorites, Settings, and Tools. Parent folders contain subfolders that hold the content. For instance, the Tools bin is a parent folder to all the categories of tools. To access subfolders, click the disclosure arrow to the left of the parent folder's name.

When you select a bin from the bin list, the contents of the folder are displayed in the Contents panel as thumbnail icons.



The Bins icon view

A toolbar along the bottom of the bin provides access to organization, playback, and editing controls.



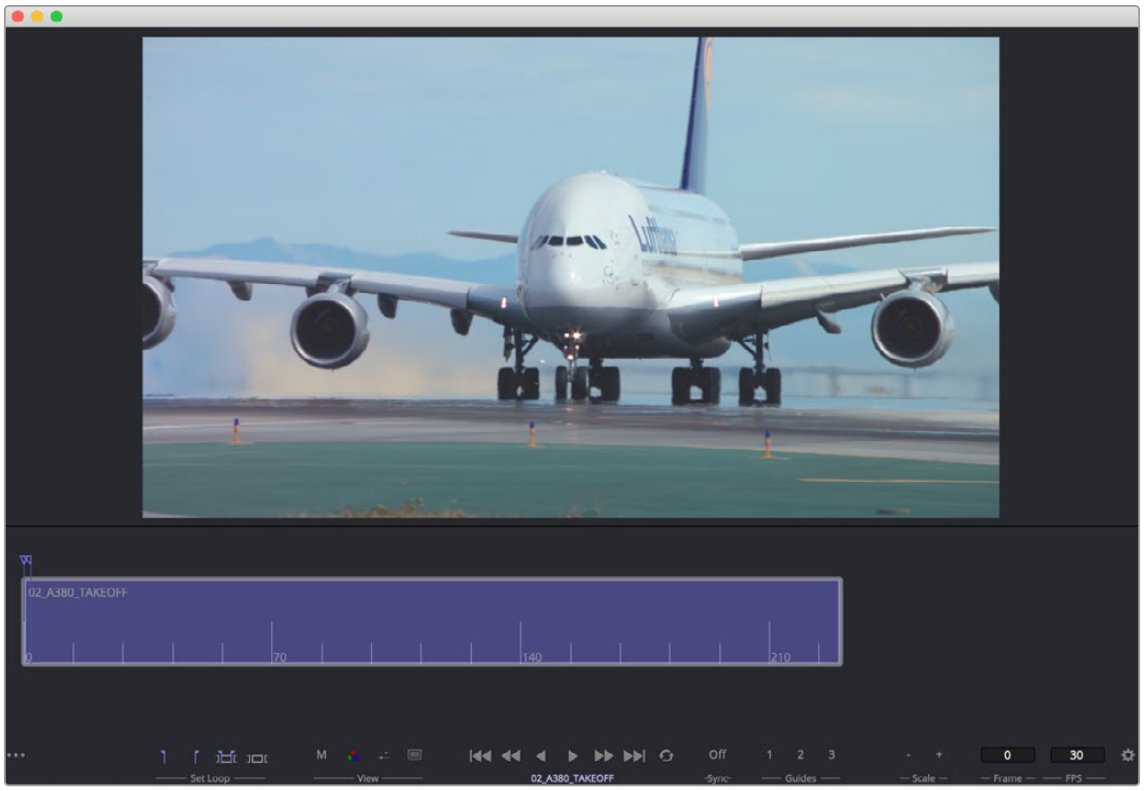
The Bins toolbar

- **New Folder:** Creates a new folder in the current window.
- **New Reel:** Creates an empty reel that can contain multiple clips edited together into a timeline.
- **New Clip:** Opens a dialog to link a new media file into a bin.
- **Studio Player:** Opens a playback viewer for a selected clip.
- **Icon/List view:** This button toggles between showing contents of a bin in thumbnail view and list view.
- **Checkerboard:** Shows a checkerboard pattern in a clip thumbnail to signify transparency.
- **Thumbnail size:** Provides a few preset sizes for thumbnail icons.

Right-clicking in any area of the bin window displays a pop-up menu to access most of a bin's features. Right-clicking on an item in a bin shows the same menu with additional options for renaming, playing, or deleting the item.

The Bin Studio Player

Selecting a clip in a bin and clicking the Studio Player button or double-clicking the clip opens the Studio Player. The Studio Player can be used to view clips, view metadata, and add notes.



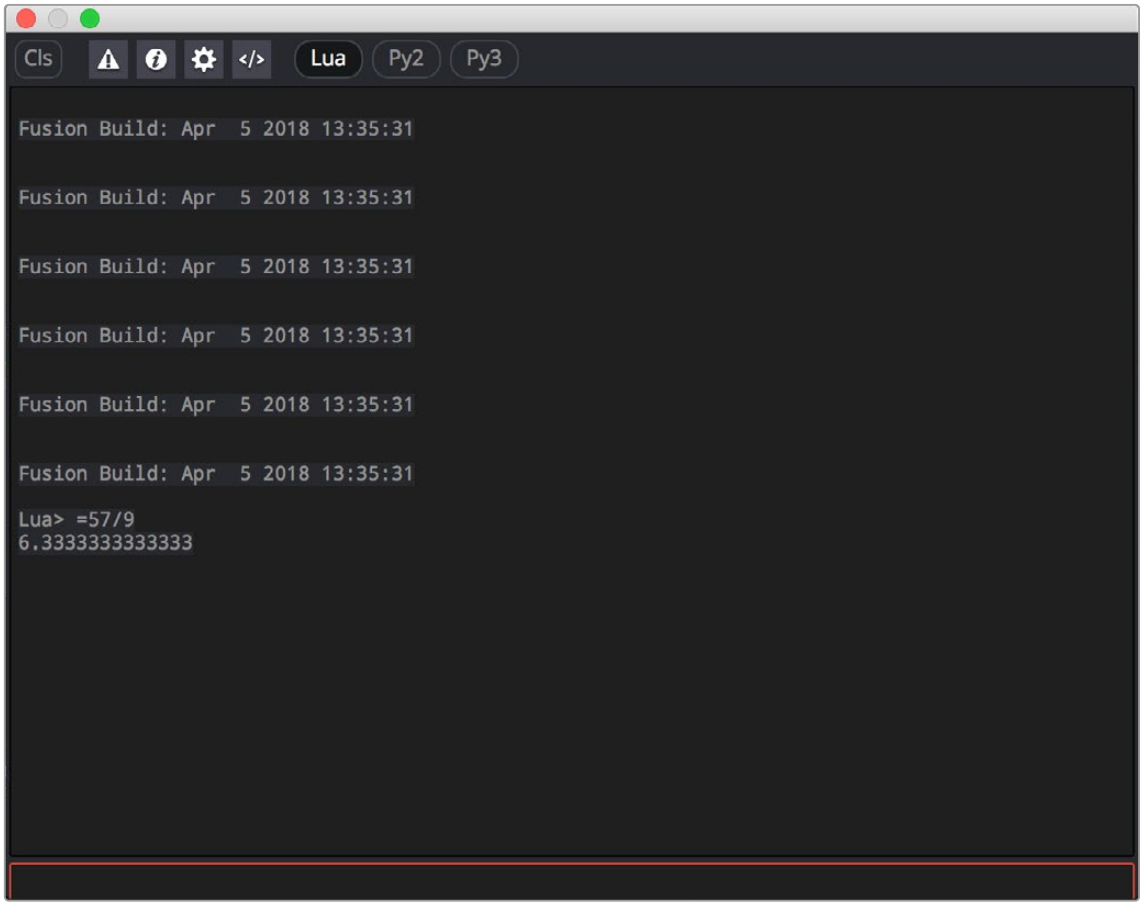
The Bin Studio Player

For more information on Bins, *see Chapter 13, "Bins"* in the Fusion Reference Manual."

The Console

The Console is a window in which you can see the error, log, script, and input messages that may explain something Fusion is trying to do in greater detail. The Console is also where you can read FusionScript outputs, or input FusionScripts directly. In DaVinci Resolve, the Console is available by choosing Workspace > Console or choosing View > Console in Fusion Studio. There is also a Console button in the Fusion Studio User Interface toolbar.

Occasionally the Status bar displays a badge to let you know there's a message in the Console you might be interested in.



The Console window

A toolbar at the top of the Console contains controls governing what the Console shows. At the top left, the Clear Screen button clears the contents of the Console. The next four buttons toggle the visibility of error messages, log messages, script messages, and input echoing. Showing only a particular kind of message can help you find what you're looking for when you're under the gun at 3:00 in the morning. The next three buttons let you choose the input script language. Lua 5.1 is the default and is installed with Fusion. Python 2.x and Python 3.x require that you install the appropriate Python environment on your computer. Because scripts in the Console are executed immediately, you can switch between input languages at any time.

At the bottom of the Console is an Entry field. You can type scripting commands here for execution in the current comp context. Scripts are entered one line at a time, and are executed immediately. For more information on scripting, see the Fusion Scripting Manual.

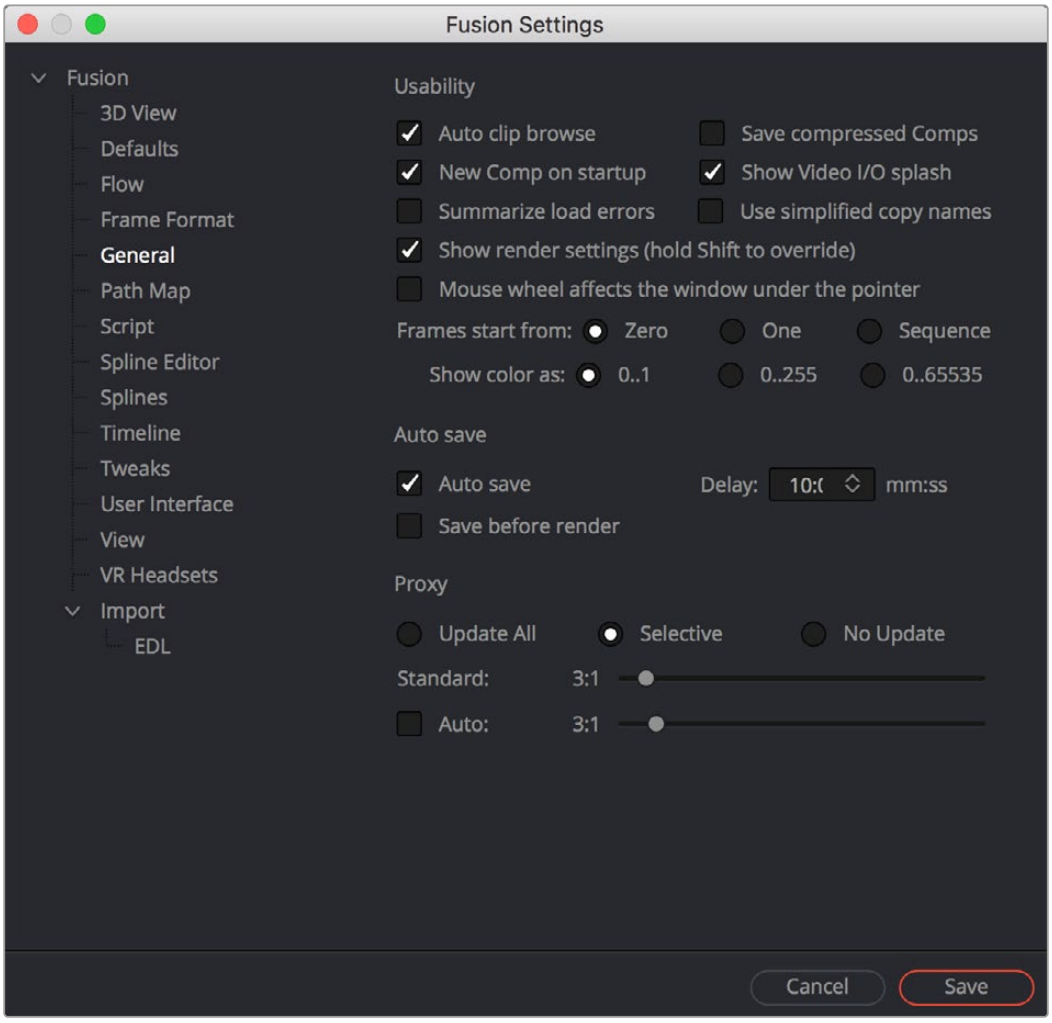
Customizing Fusion

This section explains how you can customize Fusion to accommodate whatever workflow you're pursuing.

The Fusion Settings Window

Fusion has its own settings window, accessible by choosing Fusion > Fusion Settings in DaVinci Resolve, or in Fusion Studio by choosing Fusion > Preferences on macOS or File > Preferences on Windows or Linux. This window has a variety of options for customizing the Fusion experience.

For more information, *see Chapter 15, "Preferences,"* in the Fusion Reference Manual.



The Fusion Settings window set to the General panel

Saving Fusion Layouts

It is possible to customize the layout and configuration of panels to suit the size of the desktop and monitor, or to match personal preferences.

In DaVinci Resolve, configure and resize the panels you want displayed and then:

Choose Workspace > Layout Presets > Save Layout Presets.

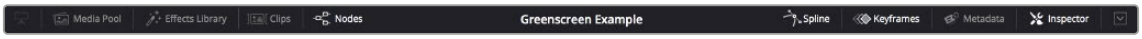
In Fusion Studio, configure and resize the panels you want displayed and then:

Click the Grab Document Layout button in the Preferences > Layout panel to save the layout for all new Compositions.

Click the Grab Program Layout button to remember the size and position of any floating views, and enable the Create Floating Views checkbox to automatically create the floating windows when Fusion restarts.

Showing and Hiding Panels

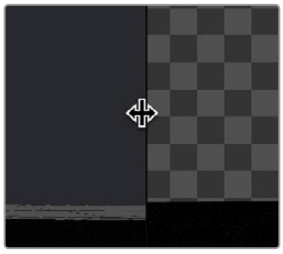
The UI toolbar at the top of the screen lets you open panels you need and hide those you don't. It's the simplest way to create a layout for your particular needs at the moment.



The UI toolbar of the Fusion page

Resizing Panels

You can change the overall size of each panel using preset configurations, or you can adjust them manually. The viewers and Work panel are inverse of each other. The more space used to display the Work panel, the less space available for the viewers. To resize a panel, manually drag anywhere along the raised border surrounding the edges of the panel.



Dragging the edge between two viewers to resize it

Fusion Studio Floating Frame

Fusion Studio includes a Floating Frame window that can be used to house any panel.

To place a panel in the Floating Frame, do the following:

1. In Fusion Studio, choose Window > New Floating Frame.
2. Right-click in the Floating Frame and choose the panel from Add View submenu.

When using multiple monitors, you can choose to have floating panels spread across your displays for greater flexibility.

Fusion Keyboard Remapping

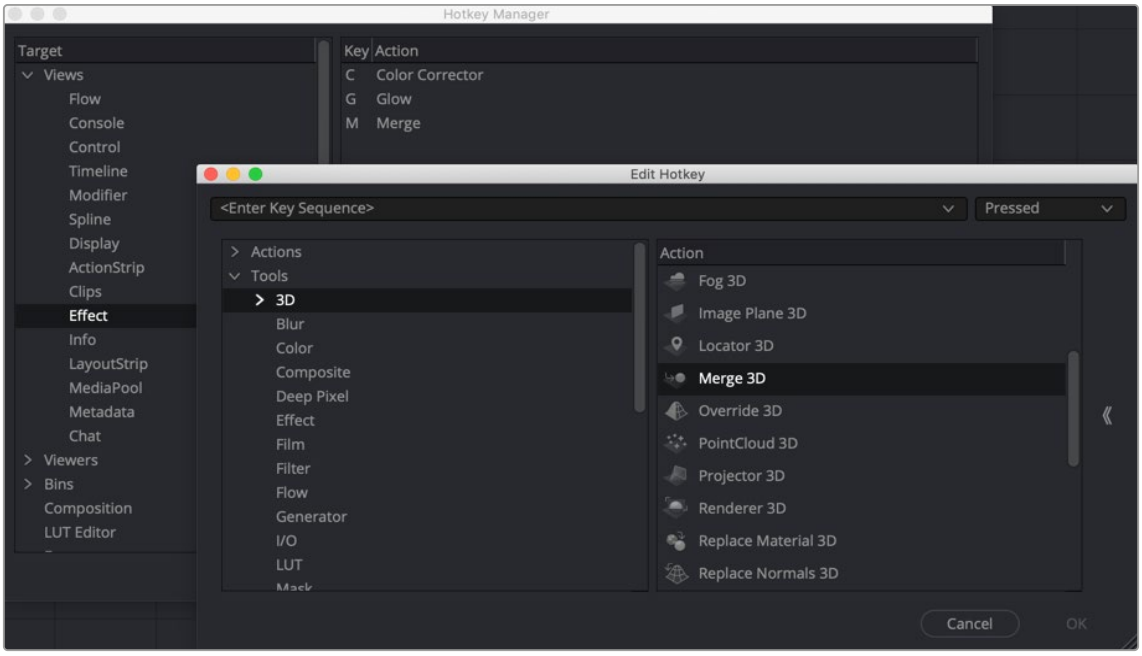
When using the Fusion page, functions and tools can be mapped to hot keys on your keyboard by choosing DaVinci Resolve > Keyboard Customization. In the Commands column of the Keyboard Customization window, select Panels > Fusion Page.

> For more information on using the Keyboard Customization window, see Chapter 4, "System and User Preferences," in the *DaVinci Resolve 20 Reference Manual*.

Fusion Studio Keyboard Remapping

When using Fusion Studio, functions and tools can be mapped to hot keys on your keyboard by choosing Views > Customize Hotkeys.

The Fusion Hotkey Manager dialog is divided into two sections. The left is where you select the functional area where you want to assign a keyboard shortcut. The right side displays the keyboard shortcut if one exists. You can use the New button at the bottom of the dialog to add a new keyboard shortcut.



The Fusion Studio Hotkey Manager window

For instance, if you want to add a shortcut for a specific node:

1. Open the Keyboard Hotkey Manager.
2. Select Views > Effect from the Target area of the Hotkey Manager.
3. Below the Key/Action area, click the New button to create a new keyboard shortcut for the Node.
4. In the Edit Hotkey window, click the Tools disclosure arrow, and then select Blur to display all the Blur-related nodes.

5 Select Glow in the Action panel.
6 At the top of the Edit Hotkey window, type G as the shortcut for the Glow node, and then click OK.
7 Glow and the G hotkey will now appear in the Key/Action area on the right.
8 Click OK to close the Hotkey Manager.
9 Click in the Node Editor and press G to add a Glow node.

Undo and Redo

Undo and Redo commands let you back out of steps you've taken or commands you've executed and reapply them if you change your mind. Fusion is capable of undoing the entire history of things you've done since creating or opening a particular project. When you close a project, its entire undo history is purged. The next time you begin work on a project, its undo history starts anew.

There is no practical limit to the number of steps that are undoable (although there may be a limit to what you can remember).

To simply undo or redo changes you've made one at a time:

Choose Edit > Undo (Command-Z) to undo the previous change.

Choose Edit > Redo (Shift-Command-Z) to redo to the next change.

Chapter 3

Getting Clips into Fusion

This chapter details the various ways you can move clips into Fusion as you build your compositions.

Contents

Preparing Compositions in the Fusion Page 62
Working on Single Clips in the Fusion Page 62
Turning One or More Clips into Fusion Clips 63
Fusion Referenced Compositions 65
Adding Fusion Composition Generators 68
Creating a Fusion Composition Clip in a Bin 69
Resetting a Fusion Clip 69
Using Fusion Transitions 69
Adding Clips from the Media Pool 70
Finding Clips in the Media Pool from the Node Editor 71
Adding Clips from the File System 71
Using MediaIn Nodes 71
MediaIn Node Inputs 71
Inspector Properties of MediaIn Nodes 72
Using Loader and Saver Nodes in the Fusion Page 75
Preparing Compositions in Fusion Studio 77
Setting Up a Composition 79
Reading Clips into Fusion Studio 81
Aligning Clips in a Fusion Studio Composition 82
Loader Node Inputs 83
Using Proxies for Better Performance 83
Presetting Proxy Quality 84
File Format Options 85
Loading Audio WAV Files in Fusion Studio 87

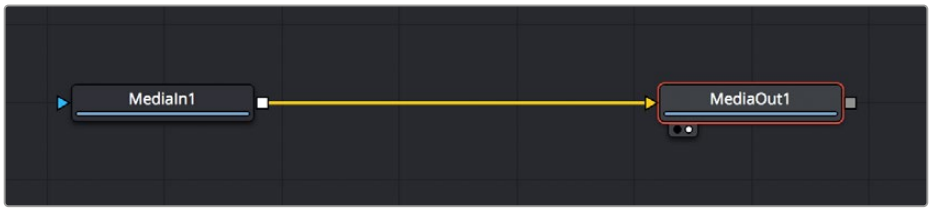
Preparing Compositions in the Fusion Page

Ordinarily, clips come into the Fusion page from the Edit or Cut page Timeline as either a single clip, or as multiple layers contained within a Fusion clip. You can also add clips to a composition directly from DaVinci Resolve's Media Pool. How clips find their way into a Fusion composition can determine how they function within that composition and what resolution that composition outputs to the rest of DaVinci Resolve.

Working on Single Clips in the Fusion Page

Positioning the playhead over a clip in the Edit page or Cut page Timeline and clicking the Fusion page button causes that clip to appear in the Fusion page as a single MediaIn node connected to a MediaOut node. Only the topmost visible clip is taken into Fusion. Clips that aren't visible because they're on lower tracks with clips above them are ignored, unless you disable the clips or tracks that appear above. These very simple default compositions are referred to unofficially in this manual as "single-clip compositions."

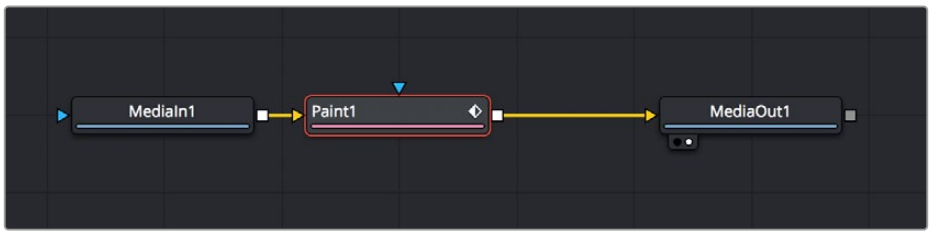
The MediaIn node represents the image that's fed to the Fusion page for further work, and the MediaOut node represents the final output that's fed onward to the Color page for grading.



The default node tree that appears when you first open the Fusion page while the playhead is parked on a clip.

This initial node structure makes it easy to quickly use the Fusion page to create relatively simple effects using the procedural flexibility of node-based compositing.

For example, if you have a clip that's an establishing shot, with no camera motion, that needs some fast paint to cover up a bit of garbage in the background, you can open the Fusion page, add a Paint node, and use the Clone mode of the Stroke tool to paint it out quickly.



A simple paint effect applied to a shot with no camera motion.

TIP: The resolution of a single clip brought into Fusion via the Edit or Cut page Timeline is the resolution of the source clip, not the Timeline resolution.

Once you've finished, simply go back to the Edit or Cut page and continue editing, because the entire Fusion composition is encapsulated within that clip, similarly to how grades in the Color page are also encapsulated within a clip. However you slip, slide, ripple, roll, or resize that clip, the Fusion effects you've created and the Color page grades you've made follow that clip's journey through your edited Timeline.

> **TIP:** While you'll likely want to do all the compositing for a greenscreen style effect in the Fusion page, it's also possible to add a keyer, such as the excellent DeltaKeyer node, between the MediaIn and MediaOut nodes, all by itself. When you pull a key this way, the alpha channel is added to the MediaOut node, so your clip on the Edit page has transparency, letting you add a background clip on a lower track of your Edit page Timeline.

How Nodes Are Named

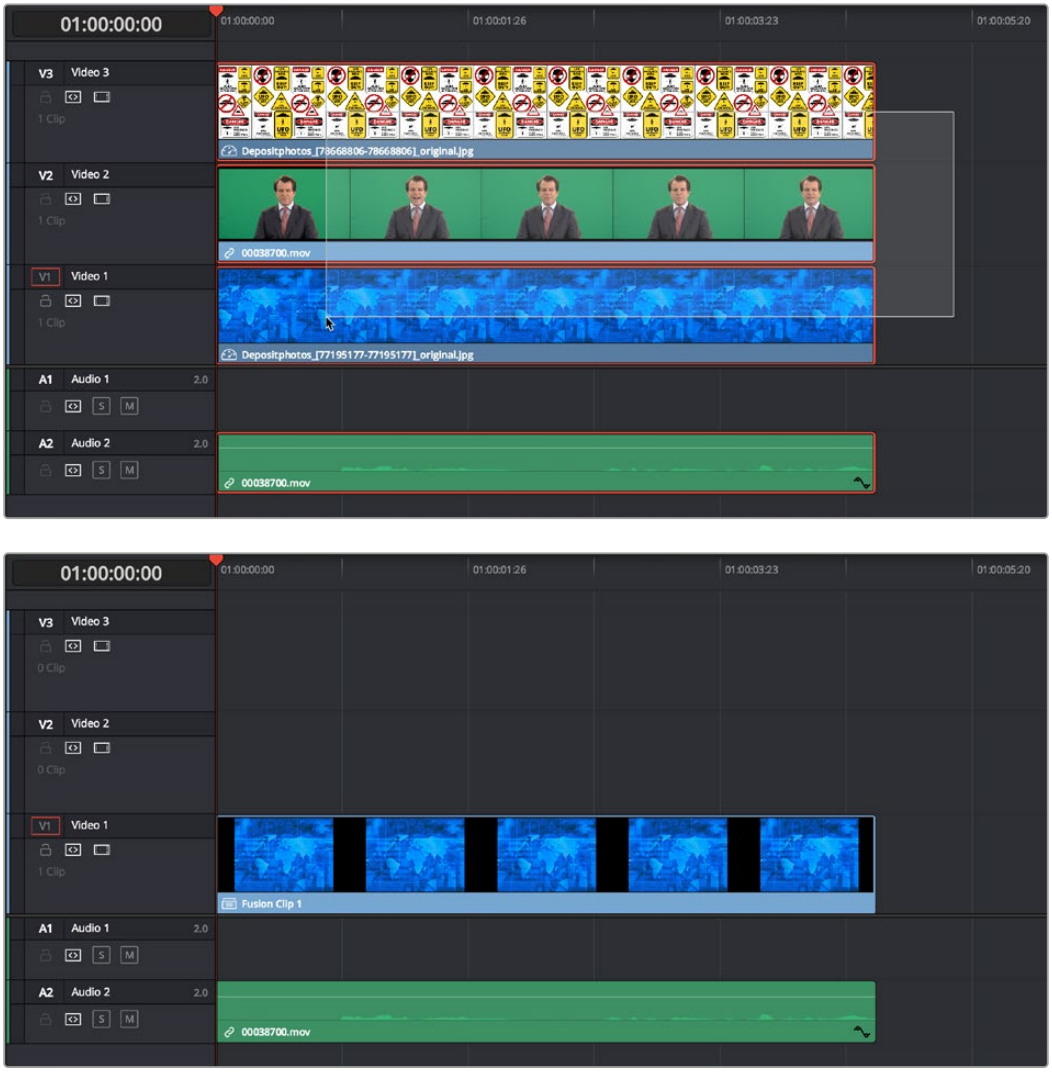
While the documentation refers to nodes by their regular name, such as "MediaIn," the actual names of nodes in the Fusion Node Editor have a number appended to them, to indicate which node is which when you have multiple instances of a particular type of node.

Turning One or More Clips into Fusion Clips

For situations where you know you're creating a more ambitious composited effect that requires multiple layers edited together with very specific timings, you can create a "Fusion clip" right from the Timeline. For example, if you have a foreground greenscreen clip, a background clip, and an additional graphic clip, you can stack them all on the Timeline as superimposed clips. You can then use the Edit page slip and slide features to align their timings so they work together as necessary. You can also edit multiple consecutive clips together that you want to use in a composition as a series of clips. Once that's done, you can select every clip in the stack to create a Fusion clip, so you can easily use all these superimposed layers within a Fusion composite.

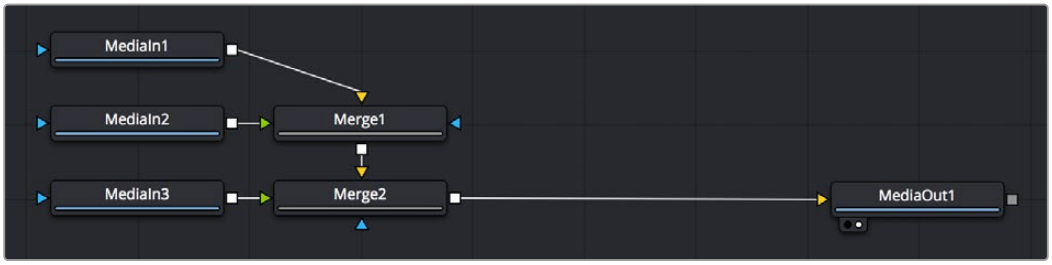
To create a Fusion clip:

1. Edit all the clips you want to use in the Edit page Timeline.
2. Select all clips you want to be in the same composite at once.
3. Right-click one of the selected clips and choose New Fusion Clip from the contextual menu.
4. A new clip, named "Fusion Clip X" (where X is an automatically incrementing number) appears in the currently selected bin of the Media Pool and in the Timeline to replace the previously selected clips.
5. With the playhead parked over that clip, open the Fusion page to see the new arrangement of those clips in the Fusion page Node Editor.



A stack of clips to use in a composite (Top), and turning that stack into a Fusion clip in the Edit page (Bottom).

The nice thing about creating a Fusion clip is that every superimposed clip in a stack is automatically connected into a cascading series of Merge nodes that create the desired arrangement of clips. Note that whatever clips were in the bottom of the stack in the Edit page appear at the top of the Node Editor in the Fusion page, but the arrangement of background and foreground input connections is appropriate to recreate the same compositional order.



The initial node tree of the three clips we turned into a Fusion clip.

TIP: Fusion clips change the working resolution of the individual clips to match the Timeline resolution. For instance, if two 4K clips are stacked one on top of the other in an HD Timeline, creating a Fusion clip resizes the clips to HD. The full resolution of the individual 4K clips is not available in Fusion. To maintain the full resolution of course clips, bring only one clip into the Fusion composition from the Edit or Cut page Timeline, and then bring other clips into the Fusion composition using the Media Pool.

Fusion Referenced Compositions

There is a newer powerful way of applying Fusion composites to your clips called a "Referenced Composition." This method allows you to reuse your work across multiple clips or even timelines.

You create a referenced composition on the Edit page by selecting a single clip, or a stack of clips, and then select "Create Referenced Composition" from the contextual menu. This generates a Fusion Referenced Composition in the Media Pool. You can enter the Fusion page to start compositing or double-click the Referenced Composition in the Media Pool to open it.

Referenced comps can be linked to other clips in any timeline by selecting the referenced comp, then right-clicking the target clip and selecting "Link to Referenced Composition." Any changes you make are applied across all connected clips because they are all "referred" or "linked" to the same composition.

Referenced Compositions are an effective and secure way to keep your work backed up and organized, as they live in the Media Pool and will persist even if the original referenced clip is deleted from the timeline.

Creating and Managing Referenced Compositions

To create a Referenced Composition:

1. Select one or more timeline clips that are "stacked" on tracks in an Edit window timeline (clips have to overlap in time at least partially).
2. Right-click and select "Create Referenced Fusion Composition."
3. The new reference composition, associated with the top selected clip, now appears in the Media Pool.

You can open the referenced composition in the Fusion page by double-clicking on it in the Media Pool. Media nodes match the relative locations of the source track clip arrangement in the Edit timeline.

The contents of the video track are treated as a source, allowing creative manipulation of clips before they're passed on to Fusion for compositing.

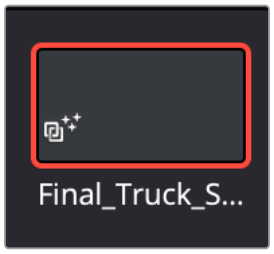
You can right-click on a referenced composition in the Media Pool and choose the Usage option to view and navigate to its target linked clips.

You can duplicate your referenced composition for iteration purposes.

In the example below we create a simple reference composition out of three clips, replacing a logo on a truck.

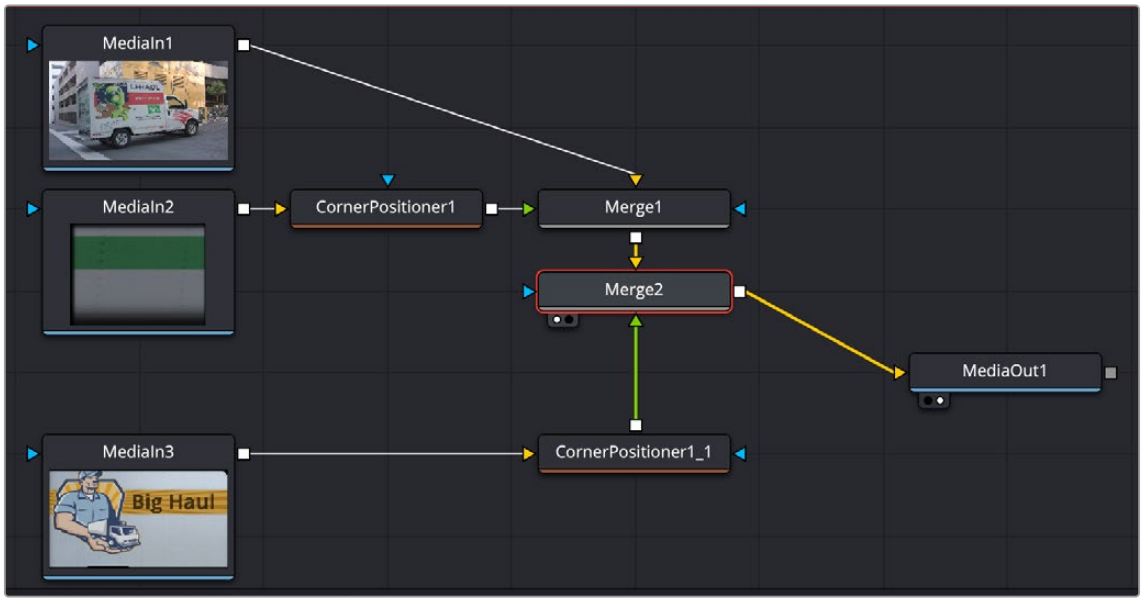


Creating a reference composition by selecting multiple overlapping clips, and then Create Referenced Composition from the contextual menu.



The Reference Composition thumbnail in the Media Pool

Double-clicking on the Reference Node brings us into Fusion, where we add some Corner Positioner nodes to line up the green stripe and logos to fit.

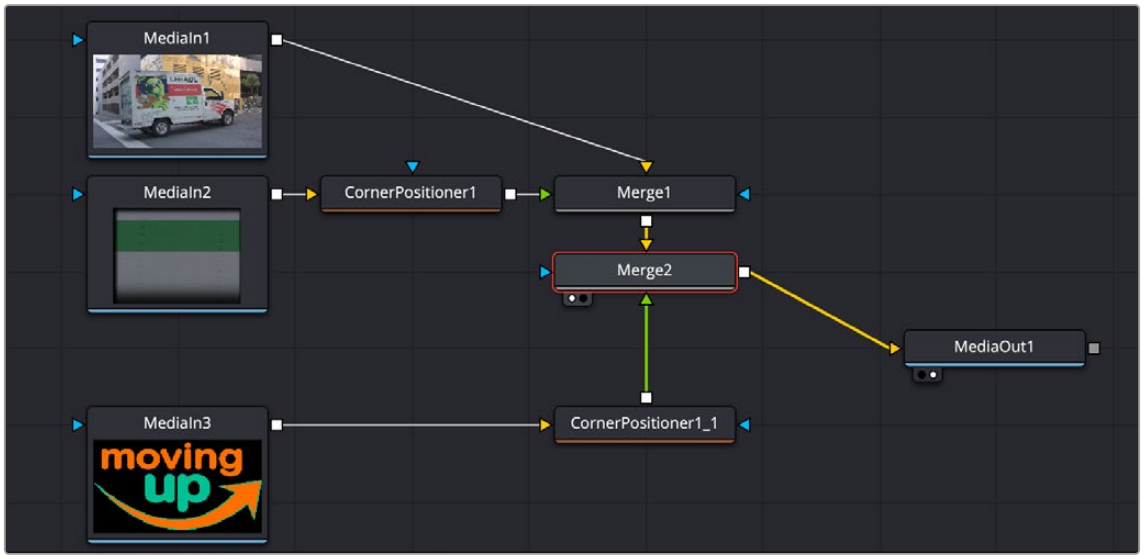




Now that we have the referenced composition set, we can easily change the logo, just by performing a replace edit with the new logo to V3 in the timeline. The new logo is automatically applied in the Fusion comp with the same settings.



The sBSpline controls





Adding Additional Referenced Clips

Once you've established a referenced composition, you can add clip(s) that can use the referenced composition to affect the other clips.

To add clip references to a referenced composition:

1. Select the referenced composition that you want to refer to another clip.
2. Select a target clip in the timeline, right click and choose "Link to Referenced Composition."
3. Add other clips, one a time if you wish, including clips from other timelines.

Referenced Composition Behaviors and Rules

There are some additional behaviors and rules to be aware of when using referenced compositions.

- Referenced compositions maintain the original source resolution of your clips.
- When selecting multiple clips, ensure that the clips have some overlap with the top selected clip. Without this overlap, the context menu action to create a referenced composition across layers is not available.
- When creating a project composition from multiple overlapping clips, a single composition associated with the top selected clip is created.
- The referenced composition features additional Media In nodes (one per track). Media In 1 refers to the top selected clip. Subsequent Media In nodes reference timeline tracks (layers) below the top clip.
- Removing a timeline clip will not remove the associated composition.
- When you edit the contents in a referenced composition's track, the corresponding Media In refreshes to automatically to show the new clips, edits, effects, and transitions.

From the timeline clip's context menu, you can:

- **Reset Reference Composition:** To reset the composition. This affects all referenced clips.
- **Unlink Reference Composition:** To remove the composition reference from the selected clip. This retains the composition in the Media Pool, but the clip effect is cleared. Useful if you want to link to another referenced composition.
- **Find Reference Composition in Media Pool:** To find the reference composition the clip is a part of in the Media Pool.

NOTE: Clips of differing frame rates or clips with speed changes are not supported in referenced compositions.

Adding Fusion Composition Generators

The Generator category of the Edit page Effects Library has a Fusion Composition generator. It's useful for creating an empty placeholder in the Timeline that you later want to work on in the Fusion page to create a more fully-featured Fusion composition.

To create a blank Fusion clip in the Edit page:

1. Open the Effects Library, and select the Effects category.
2. Edit a Fusion Composition clip into the Timeline in whichever way is most convenient.
    a) You can drag a Fusion Composition clip into the Timeline, which will result in a clip that's the length of the "Standard generator duration" preference, which is 5 seconds by default.
    b) You can set In and Out points in the Timeline, and drag the Fusion Composition clip onto any of the editing overlays of the Timeline viewer to perform that sort of edit to insert, overwrite, "place on top," or ripple overwrite it into a specific place in the Timeline, for a specific duration.
3. A new clip named "Fusion Composition" appears in the Timeline. It initially displays only black in the Timeline viewer, since it's a blank composition with no contents.
4. With the playhead parked over that clip, open the Fusion page. Since this composition is blank, there's only a single MediaOut node in the Node Editor. At this point, you can add whatever media, generators, and other effects you require.

Creating a Fusion Composition Clip in a Bin

You can create an empty Fusion Composition clip in any bin in the Media Pool without creating a Timeline. This method can be useful for creating motion graphics or titles when there is no Timeline or when you plan on using the clip in multiple Timelines.

To create a blank Fusion Composition clip in a bin:

1. Select the bin in the Media Pool where you want to save the Fusion Composition.
2. Right click in an empty area of the bin and choose New Fusion Composition.
3. In the New Fusion Composition clip dialog, enter a Name for the clip, a duration, and a frame rate, and then click Create.
4. The clip will appear in the bin. To open it in Fusion, do one of the following:
   - Double-click the Fusion Composition
   - Right-click over the Fusion Composition clip and choose Open in Fusion Page

Resetting a Fusion Clip

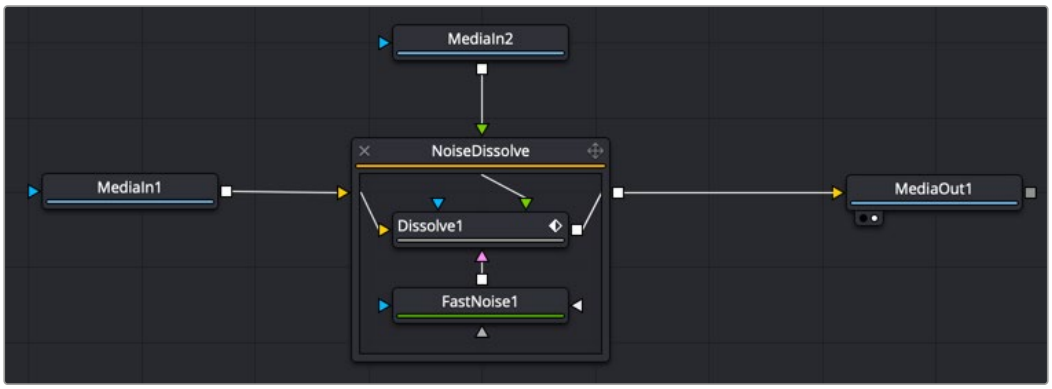
To reset a Fusion clip back to its initial default state, right-click on the Fusion clip in the Edit page Timeline, and select Reset Fusion Composition. You can select multiple Fusion clips and reset them at once.

Using Fusion Transitions

Specific Fusion transitions are available in the Edit page Effects Library. You can use these transitions to bring two clips and the transition between them into the Fusion Page. These transitions can then be modified and saved back to the Edit page Timelines or saved as a new, reusable Fusion Transition that appears in the Edit Page Effects llbrary.

To apply and open a Fusion Transition:

1. Open the Effects Library, and select the Video Transitions category.
2. Scroll to the bottom of the transition list and drag one of the Fusion Transitions onto a cut in the Timeline.
3. Right-click over the Fusion Transition in the Timeline and choose Open in Fusion Page.
4. The Fusion page opens with two MediaIn nodes representing the two sides of the transitions. The MediaIn nodes connect to a cross dissolve or a group of nodes used to create the transition. At this point, you can modify the transition using masks or other nodes and return to the Edit page to see the results.



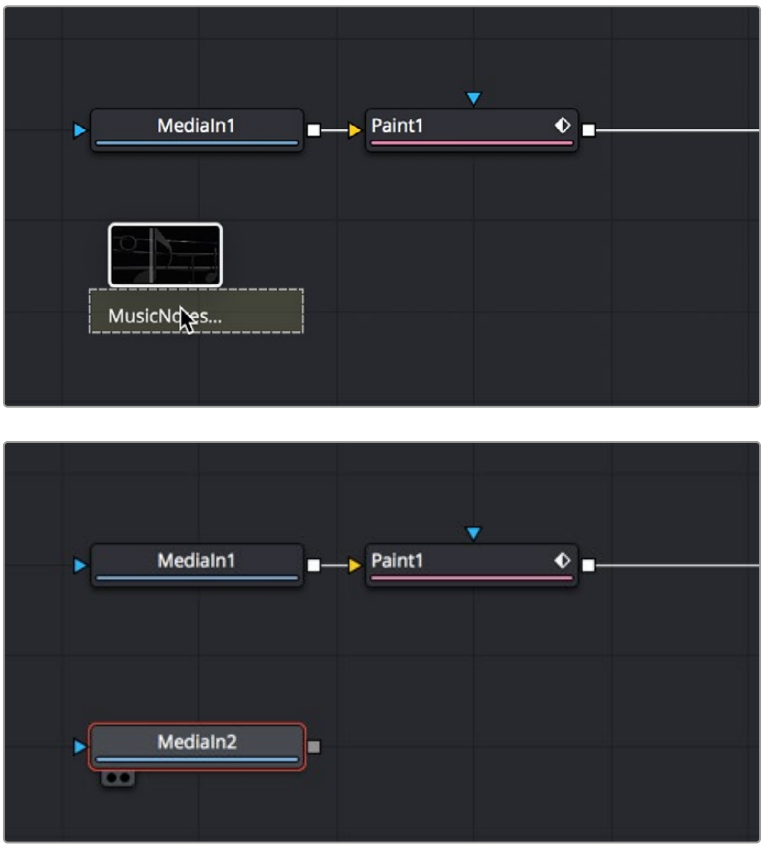
A Noise Dissolve Transition opened in the Fusion page.

To learn about creating custom Fusion Transitions that appear in the Effects Library, see *Chapter 6, "Node Groups, Macros, and Fusion Templates,"* in the Fusion Reference Manual.

Adding Clips from the Media Pool

You can open the Media Pool on the Fusion page and drag clips directly to the Node Editor to add them to your node tree.

When you add a clip by dragging it into an empty area of the Node Editor, it becomes another MediaIn node, disconnected, that's ready for you to merge into your current composite in any one of a variety of ways.



Dragging a clip from the Media Pool (Top), and dropping it onto your composition (Bottom).

TIP: Dragging a clip from the Media Pool on top of a connection line between two other nodes in the Node Editor adds that clip as the foreground clip to a Merge node.

When you add additional clips from the Media Pool, those clips becomes a part of the composition, similar to how Ext Matte nodes you add to the Color page Node Editor become part of that clip's grade.

Audio with Media Pool Clips

Audio from a clip brought in through the Media Pool is muted by default. Hearing the audio from a Media Pool clip is a two step process.

To hear audio from a clip brought in through the Media Pool, do the following:

1. Select the clip in the Node Editor.
2. In the Inspector, click the Audio tab and select the clip name from the Audio Track drop-down menu.
3. Right-click the speaker icon in the toolbar, then choose the MediaIn for the Media Pool clip to solo its audio.

You can now use the speaker icon contextual menu to switch back and forth between all the MediaIn nodes.

Finding Clips in the Media Pool from the Node Editor

Complex Fusion Compositions can have multiple MediaIn nodes to keep track of. It is possible to find the exact clip in the Media Pool assigned to a MediaIn node by selecting it, and then choosing Clip > Find Clip in Media Pool (Option-F).

Adding Clips from the File System

You also have the option of dragging clips from the file system directly into the Node Editor. When you do this, they'll be added to the currently selected bin of the Media Pool automatically. So, if you have a library of stock animated background textures and you've just found the one you want to use using your file system's search tools, you can simply drag it straight into the Node Editor to use it right away.

Using MediaIn Nodes

The MediaIn nodes in the Fusion Page are the foundation of every composition you create. This section provides more detail about the controls available for adjusting MediaIn and Loader nodes.

MediaIn Node Inputs

MediaIn nodes have one Effects mask input and one output. In the case of the Effects mask input, connecting a mask node such as a Polygon or B-Spline node creates an alpha channel in the MediaIn node.

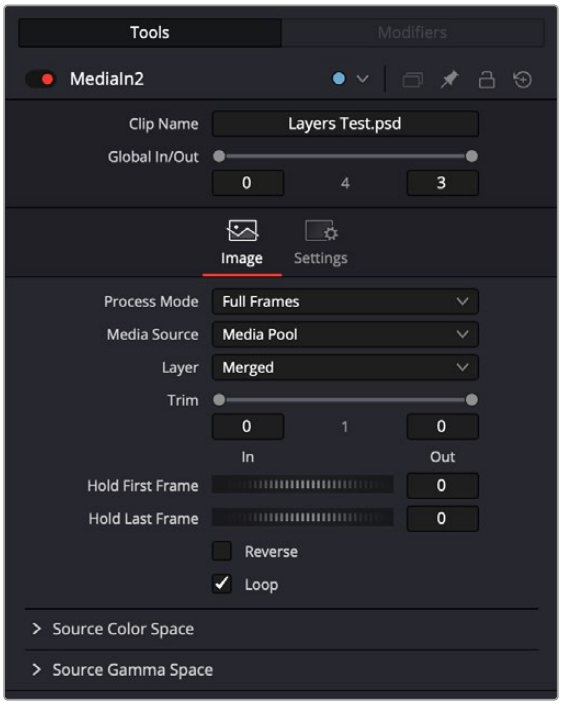
TIP: If you connect a mask node without any shapes drawn, that mask outputs full transparency, with the result that the image output by the MediaIn node is uselessly blank. If you want to rotoscope over a MediaIn node, first create a disconnected mask node, and with the mask node selected (exposing its controls in the Inspector) and the MediaIn node loaded into the viewer, draw your mask. Once the shape you're drawing has been closed, you can connect the mask node to the MediaIn node's input, and you're good to go.

Inspector Properties of MediaIn Nodes

Which Inspector options are available for MediaIn nodes of your composition depends on how you imported the media.

MediaIn Node Parameters for Clips in a Timeline

When you create a composition using clips from the Edit page Timeline, the MediaIn nodes for those clips display fewer parameters than those imported directly from the Media Pool, because the timing of clips that have already been edited into a Timeline is already set.



The Inspector Image tab parameters for a clip from the Timeline

Image Tab

- **Clip Name:** Displays the name of that clip.
- **Process Mode:** Lets you choose whether the clip represented by that node will be processed as Full Frames, or via one of the specified interlaced methods.
- **Media Source:** This control allows you to choose where media is sourced from; the three available options include: Timeline, Background, and Media Pool.

  Timeline: Links the clip from the Edit Timeline (default).

  Background: This pulls frames from the lower video tracks in the edit. When toggling the Media Source to Background, the composited result of all the lower tracks is visible.
- **Media Pool:** Links the clip directly from the Media Pool, bypassing the Edit Timeline.
- **Layer:** In the case of Multilayer or Photoshop PSD files, this selects the layer in the file to use. For a Fusion Clip created using multiple video layers on the timeline, this will refer to the index number of each of the layers.

Source Color Space

The Color Space Type menu sets the color space of the media to help achieve a linear workflow. Unlike the Gamut tool, this doesn't perform any actual color space conversion but rather adds the source space data into the metadata, if that metadata doesn't already exist. The metadata can then be used downstream by a Gamut tool with the From Image option, or in a Saver if explicit output spaces are defined there.

- **Auto:** Passes along any metadata that might be in the incoming image.
- **Space:** Allows the user to set the color space from a variety of options.

Source Gamma Space

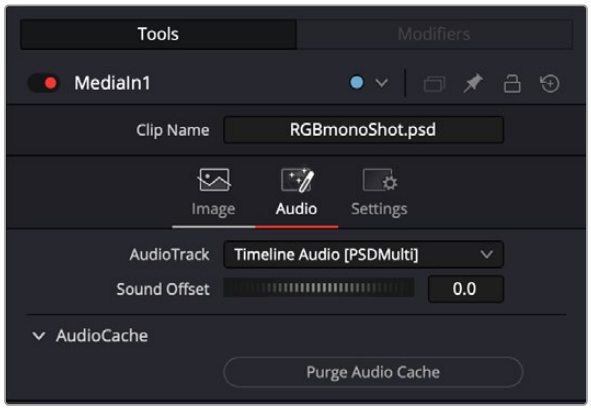
The Curve Type menu automatically determines or allows you to choose the Gamma setting for the image and allows the option to remove the Gamma curve to help achieve a linear workflow.

- **Auto:** Passes along any metadata that might be in the incoming image.
- **Space:** Allows you to choose a specific setting from a Gamma Space drop-down menu, while a visual graph lets you see a representation of the gamma setting you've selected.
- **Log:** Similar to the Log-Lin node, this option reveals specific log-encoded gamma profiles so that you can select the one that matches your content. A visual graph shows a representation of the log setting you've selected. When Cineon is selected from the Log Type menu, additional Lock RGB, Level, Soft Clip, Film Stock Gamma, Conversion Gamma, and Conversion table options are presented to finesse the gamma output.
- **Remove Curve:** Depending on the selected gamma space or on the gamma space found in Auto mode, the associated gamma curve is removed from the material, effectively converting it to output in a linear color space.
- **Pre-Divide/Post-Mujltiply:** Lets you convert "straight" alpha channels into pre-multiplied alpha channels, when necessary.

TIP: All content in the DaVinci Resolve Fusion page is processed using 32-bit floating-point bit depth, regardless of the content's actual bit depth.

Audio Tab

The Inspector for the MediaIn node contains an Audio tab, where you can choose to solo the audio from the clip or hear all the audio tracks in the Timeline.



The MediaIn Audio tab

The Audio tab in the MediaIn node is used to select the track for playback, slip the audio timing, and reset the audio cache

If the audio is out of sync when playing back in Fusion, the Audio tab's Sound Offset wheel allows you to slip the audio in subframe frame increments. The slipped audio is only modified in the Fusion page. All other pages retain the original audio placement.

Purging the Audio Cache

The audio and its settings are cached for faster performance. If you change which audio tracks you want to play back in Fusion, or you use the Sound Offset wheel to slip the audio tracks, you need to purge the audio cache. Also, if you return to the Edit, Cut, or Fairlight page and modify the audio levels, you need to purge the audio cache.

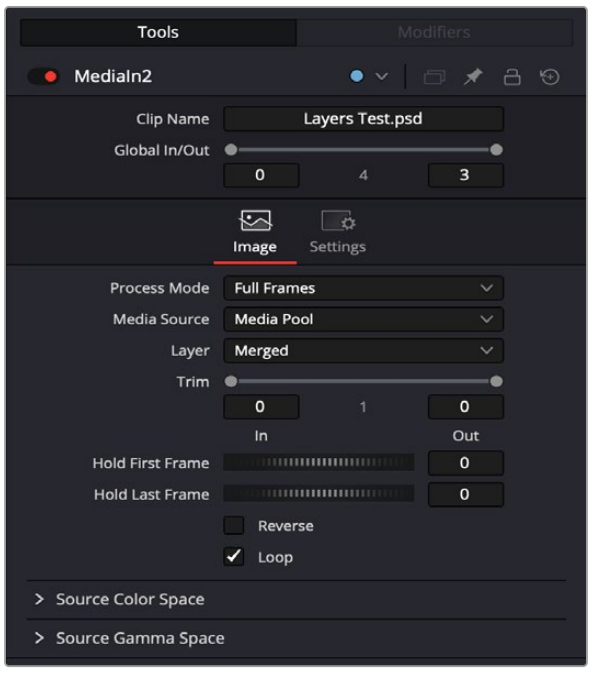
To purge the audio cache after any change to the audio playback:

Click the Purge Audio Cache button in the Inspector.

The audio will be updated when you next play back the composition.

Aligning Clips from the Media Pool

When you add a clip from the Media Pool or your file system directly into a composition, the resulting MediaIn node has more options than clips from the Edit page Timeline. These additional options make it easier to align the Media Pool clips with other clips from the Edit or Cut page Timeline. You can also trim clips, hold the first or last frame for a longer duration than the original media, and reverse or loop the clip to get more range for your composition.



The Inspector parameters for a clip imported from the Media Pool.

Below is a list of controls that are added beyond the controls that appear when a clip is added from the Edit or Cut page Timeline.

- **Global In/Out:** Use this control to specify the position of this node within the composition. For instance, when the clip is added to the comp from the Media Pool, it is added at frame 0. However, the MediaIn node from the Edit page Timeline may not start until a much later frame, based on where it is edited into the Timeline. Use Global In to specify the frame on which that the clip starts so that it aligned with media from the Edit page Timeline. It is easiest to view and change the alignment of different clips in the comp while viewing the Keyframes Editor.

  To slide the clip in time or align it to other clips without changing its length, place the mouse pointer in the middle of the range control and drag it to the new location, or enter the value manually in the Global In value control.

  If the Global In and Out values are decreased to the point where the range between the In and Out values is smaller than the n number of available frames in the clip, Fusion automatically trims the clip by adjusting the Clip Time range control. If the Global In/Out values are increased to the point where the range between the In and Out values is larger than the number of available frames in the clip, Fusion automatically lengthens the clip by adjusting the Hold First/Last Frame controls. Extended frames are visually represented in the range control by changing the color of the held frames to purple in the control.
- **Trim:** The Trim range control is used to trim frames from the start or end of a clip. Adjust the Trim In to remove frames from the start and set Trim Out to specify the last frame of the clip. The values used here are offsets. A value of 5 in Trim In would use the 5th frame in the sequence as the start, ignoring the first four frames. A Trim Out value of 95 would stop loading frames after the 95th.
- **Hold First Frame/Hold Last Frame:** The Hold First Frame and Hold Last Frame controls will hold the first or last frame of the clip for the specified amount of frames. Held frames are included in a loop if the footage is looped.
- **Reverse:** Select this checkbox to reverse the footage so that the last frame is played first and the first frame is played last.
- **Loop:** Select this checkbox to loop the footage until the end of the project. Any lengthening of the clip using Hold First/Last Frame or shortening using Trim In/Out is included in the looped clip.

Using Loader and Saver Nodes in the Fusion Page

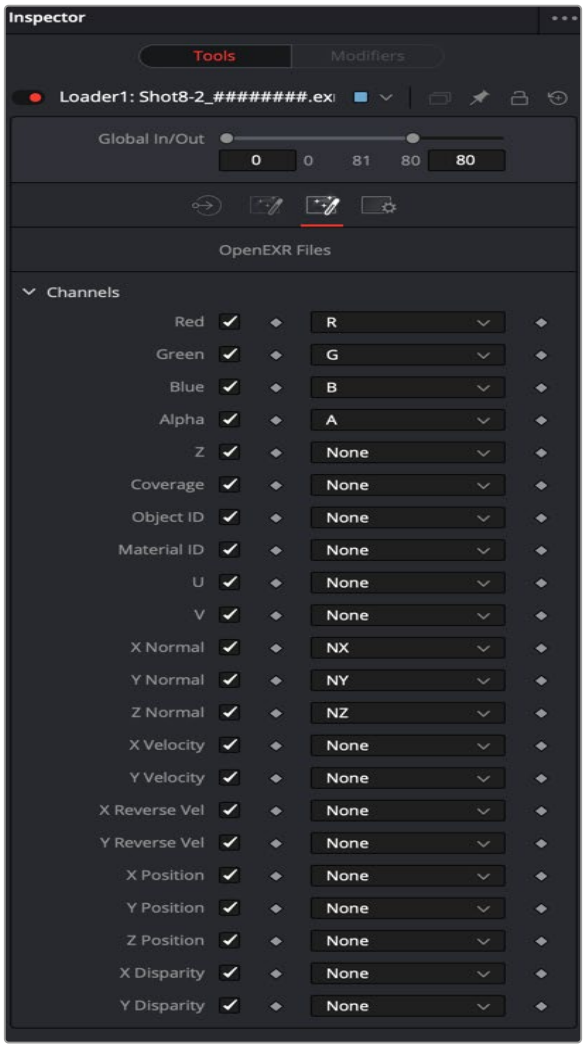
The Loader and Saver nodes in Fusion Page are specifically used for workflows that center around multi-channel EXR files. OpenEXR media can contain high-quality floating-point image data, multiple matte channels, as well as auxiliary channels rendered from 3D software. The Loader node lets you add OpenEXR files to a composition directly from the file system, retaining all the channels that are embedded within that file. Saver nodes enable you to render either all or part of a composition as EXR files directly to disk, bypassing the DaVinci Resolve Deliver page.

To import a multi-channel EXR image sequence into the Fusion page:

1. Open the Fusion page.
2. Open the Effects Library, and select the Tools > I/O category, and click the Loader node.
3. In the OS navigation window that opens, select the EXR image sequence you want to import, and click Open.
4. A Loader node linked to the EXR file will appear in the Fusion page, although a clip will not appear in the Media Pool.

Loader Node Parameters

The image tab of Loader nodes shares parameters with MediaIn nodes, as described previously in this chapter. However, when using a Loader node for EXR files, the Format tab selectively enables and disables the use of specific auxiliary channels contained in the file.



The Format tab in a Loader node Inspector displays Aux channels in EXR files.

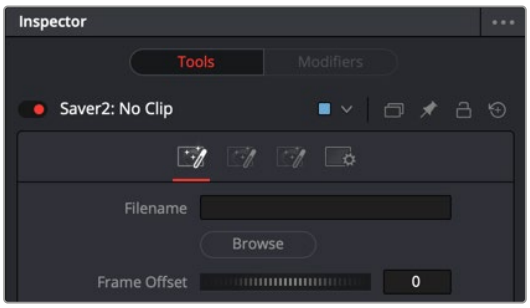
Outputting Images Using Saver Nodes

Saver nodes render OpenEXR image sequences to disk directly from the Fusion page. Saver nodes can be added to any branch of a node tree, allowing you to export one or more subsets of nodes in a composition. You can add as many Saver nodes as you like to whichever branches of your composition's node tree you need to output, to export multiple parts of a composition. For example, if a particular composition is made up of one branch of nodes working together to create a complex background, and another branch of nodes creating a foreground character with transparency, you could export the background and foreground branches as separate OpenEXR files using two Saver nodes.

To do this, simply create a Saver node after any set of nodes you want to output, and then open the Inspector and click Browse to choose a name and location for the rendered result.

When naming the file, you must add the .exr file extension. Fusion sets the output format accordingly. A four-digit frame number is automatically added before the filename extension. However, you can specify the frame padding by adding several zeroes to indicate the number of digits. For example, 000 signifies 001.

Once the Saver node is set up, output one or more Saver nodes, and choose Fusion > Render All Savers.



The Inspector parameters for a Saver node.

Manual Disk Caching Using Loader and Saver Nodes

The Loader and Saver nodes in the Fusion page are also useful for optimizing extremely complex and processor-intensive compositions. For example, you can render out specific branches of a node tree that no longer require frequent adjustment to OpenEXR via a Saver node, then reimport the result to take the place of the original branch of nodes in order to improve the performance of your composition. Used this way, Loader and Saver nodes provide a bulletproof manual workflow for caching using media files that will never be automatically purged unless you specifically delete them. As long as you retain the original branch of nodes, you can always readjust and re-render these manually cached parts of a composition, if necessary.

Preparing Compositions in Fusion Studio

The next few sections in this chapter cover preparing a project and adding clips into a composition when using Fusion Studio. The term composition, or comp, is used to refer to the Fusion project file. By default, opening the Fusion Studio application creates a new empty composition when it's launched. A composition can contain single frames, image sequences, or movie files at various resolutions and bit depths. Knowing which files you can load in, how to set up a composition to handle them, and finally, reading those files in are the first steps in beginning to composite.

Opening, Closing, and Saving Compositions

As soon as you open Fusion Studio, a new empty composition is created. If necessary, you can also create or open multiple compositions at once. Each additional composition is opened as a tab to the main Fusion Studio window.



Three compositions opened as tabs in Fusion Studio.

To create a second new composition:

- Choose File > New.

To open an existing composition, do one of the following:

- Choose File > Open.
- Choose File > Open Recent, and choose from the list of recently opened comps.
- Drag a composition file from an OS file browser into the tabbed composition area at the top of the Fusion Studio window.
- Double-click on a composition file in the OS file browser.

The following methods can be used to close the current composition:

- Choose File > Close from the menu at the top of the Fusion window.
- Click the Close X icon on the right of the composition's tab.

If the composition has unsaved changes, a dialog box appears allowing you to save before closing.

> **TIP:** Compositions that have unsaved changes will display an asterisk (*) next to the composition's name in the Fusion Studio title bar and in the composition's tab.

To save the current composition, you can do the following:

- Choose File > Save and enter a name if the comp has yet to be named.
- Choose File > Save As to save under a new name.
- Choose File > Save Version to save the current composition with an added three-digit version number at the end of the name. Each time you save a version, the number automatically increments and the comp file is saved in the same location as the first version.

Auto Save

Auto save automatically saves the composition to a temporary file at preset intervals. Auto saves help to protect you from loss of work due to power loss, software issues, or accidental closure.

To enable auto save for new compositions, choose Fusion Studio > Preferences, and then locate Global > General > Auto Save in the Preferences dialog.

An auto-save file does not overwrite the current composition in the file system. A file with the same name is created in the same folder as the composition but with the extension .autosave instead of .comp. Unsaved compositions will place the autosave file in the default folder specified by the Comp: path in the Paths panel of the Global Preferences.

If an auto-save file is present when Fusion Studio loads a composition, a dialog will appear asking to load the auto-saved or original version of the composition.

The Composition File Format

Composition files are saved as readable plain text files. Using plain text files to describe a composition makes it easier to integrate Fusion into structured, visual effects pipeline and asset management solutions.

Composition files can be opened and edited using any standard text editing program. However, it is never a good idea to open the file using a word processor, such as Microsoft Word or Apple Pages, as these will generally save additional formatting information which will make the composition unreadable to Fusion.

Importing and Exporting a Composition from DaVinci Resolve

Although the compositions created in DaVinci Resolve's Fusion page are saved in the DaVinci Resolve project library as .drp project files, you can import and export Fusion composition files when in the Fusion page. This makes it very easy to share Fusion compositions between the different applications.

To export a Fusion composition from DaVinci Resolve:

1. From within DaVinci Resolve, switch to the Fusion page with the composition you want to export.
2. Choose File > Export Fusion Composition.
3. A Save dialog appears in which you can enter a name and location from the exported Fusion composition.

A .comp extension is added to the end of the filename. Only the node tree created in the Fusion page is exported. Clips not added to the Node Editor will not appear in the Fusion Studio bins. ResolveFX added to the comp will also not translate from the Fusion page to Fusion Studio.

MediaIn nodes from DaVinci Resolve are automatically converted to Loader nodes, and if the file path remains identical, the media is automatically relinked.

MediaOut nodes are converted to Saver nodes.

The return trip can also be performed, saving a composition file from Fusion Studio and importing it into the Fusion page within DaVinci Resolve.

To import a composition from Fusion Studio into the Fusion page within DaVinci Resolve:

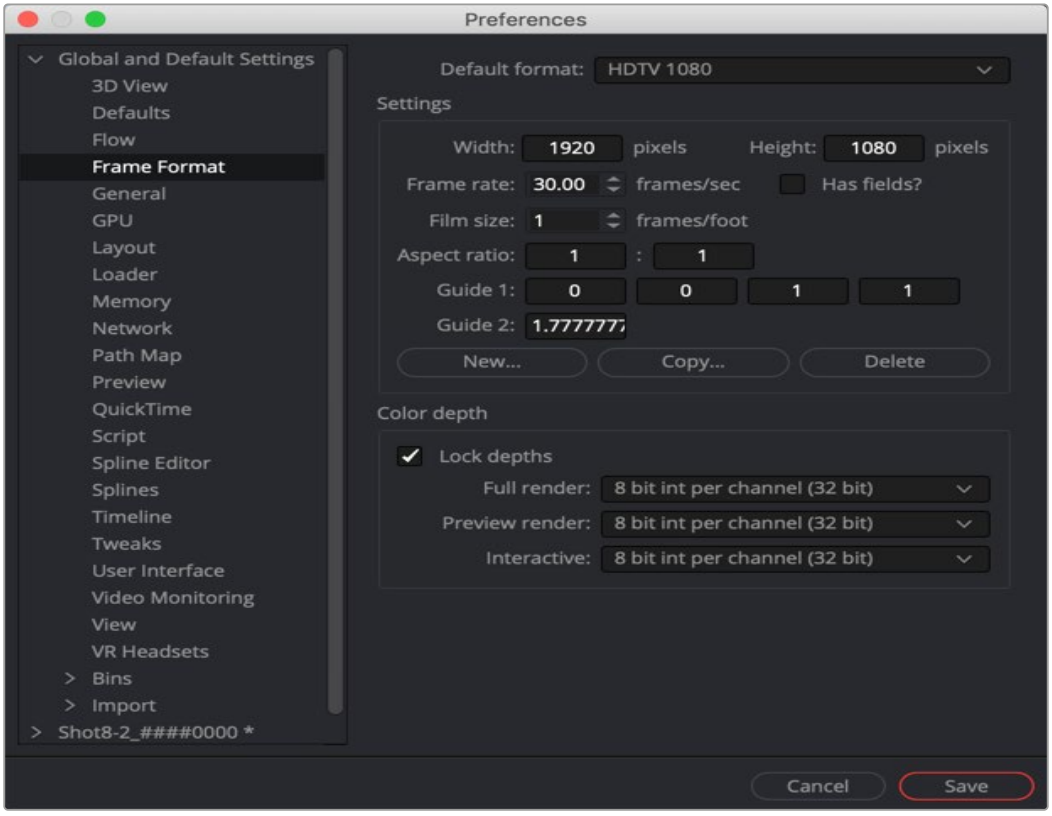
1. From within Fusion Studio, open the composition you want to move into the Fusion page.
2. From within DaVinci Resolve, switch to the Fusion page with an empty composition. The composition you import will completely replace the existing composition in the Fusion page Node Editor.
3. Choose File > Import Fusion Composition.
4. In the Open dialog, navigate to the Fusion comp and click Open.
5. The new comp is loaded into the Node Editor, replacing the previously existing composition.

TIP: To keep an existing comp in the Fusion page and merge a new comp from Fusion Studio, open Fusion Studio, select all the nodes in the Node Editor, and press Command-C to copy the selected nodes. Then, open DaVinci Resolve and switch the Fusion page with the composition you want, click in an empty location in the Node Editor, and press Command-V to paste the Fusion Studio nodes. Proceed to connect the pasted node tree into the existing one using a Merge or Merge 3D node.

Setting Up a Composition

Source media can come in a variety of formats, including HD, UHD, and 4K or larger. Often you will have different formats within a single comp. Each format has different properties, from resolution to color depth and gamma curve. Fusion can mix and match material of different formats together in a single composite, but it is important to note how Fusion Studio configures and combines materials of different formats when loading and merging them together.

When you open Fusion Studio, an empty composition is created. The first thing you do when starting on a new composition is to set the preferences to match the intended final output format. The preferences are organized into separate groups: one for global preferences, and one for the preferences of the currently opened composition.



Frame Format preferences for images generated in Fusion.

Although the final output resolution is determined in the Node Editor, the Frame Format preferences are used to determine the default resolution used for new Creator tools (i.e., text, background, fractals, etc.), aspect ratio, as well as the frame rate used for playback.

If the same frame format is used day after day, the global Frame Format preferences should match the most commonly used footage. For example, on a project where the majority of the source content will be 1080p high definition, it makes sense to set up the global preferences to match the frame format of the HD source content you typically use.

To set up the default Frame Format for new compositions, do the following:

1. Choose Fusion Studio > Preferences.
2. Click the Global and Default Settings disclosure triangle in the sidebar to open the Globals group.
3. Select the Frame Format category to display its options.

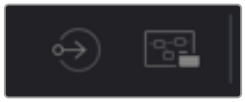
When you set options in the Global Frame Format category, they determine the default frame format for any new composition you create. They do not affect existing compositions or the composition currently open. If you want to make changes to existing compositions, you must open the comp. You can then select the Frame Format controls listed under the comp's name in the sidebar.

For more information on preferences, *see Chapter 15, "Preferences,"* in the Fusion Reference Manual.

Reading Clips into Fusion Studio

Once the Frame Format preferences are set, you usually begin to create a composite by reading in source media. When Fusion reads in media, it doesn't convert or move the original files; it simply reads the files in place, from whichever storage volume they're on. You are always dealing with the original source files in their original location.

Source media is read into a comp using a Loader tool. Although there are other tools within Fusion Studio that you can use to generate images like gradients, fractals, or text, each still image, image sequence, or movie file must be added to your comp using a Loader tool.



Loader and Saver tools are used to add media to Fusion Studio and render it out.

To add media to your comp, do one of the following:

- Click Effects to open the Effects Library, and then select Tools > I/O > Loader.
- Click the Loader icon in the toolbar.
- Right-click over the Node Editor, and then choose Add Tool > I/O > Loader.
- Drag a file from an OS file browser window into the Node Editor.

If multiple files are dragged into the Node Editor, a separate Loader is added for each file. However, if you drag a single frame from an image sequence, the entire series of the image sequence is read into the comp using one Loader, as long as the numbers are sequential.

To add only one frame of an image sequence to your comp:

- Hold Shift while you drag a single frame from an image sequence into the Node Editor.

  This comes in handy when you want to read in photographs from a digital camera that are numbered sequentially.

A Loader represents any clip, image file, or graphic that you bring into Fusion. However, other types of media can also be brought into Fusion Studio. Photoshop PSD files, SVG splines, and 3D models in the Alembic, FBX, and OBJ format can be imported using the File > Import menu.

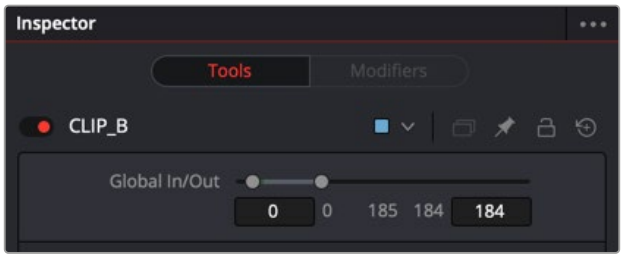
TIP: Using File > Import > Footage creates a new composition along with a Loader node for the footage. The selected media is automatically used for the name of the composition.

For more information about the Loader node, *see Chapter 45, "I/O Nodes,"* in the Fusion Reference Manual.

Aligning Clips in a Fusion Studio Composition

When you add a clip into a composition, the resulting Loader node is added at frame 0 of the composition. However, the vital portion of the clip you are interested in may not start until a few frames or even seconds later. To ensure you can align the timing of each piece of media, each Loader includes timing and trimming options in the Inspector. You can also hold the first or last frame for a longer duration than the original media, and reverse or loop the clip to get more range for your composition.

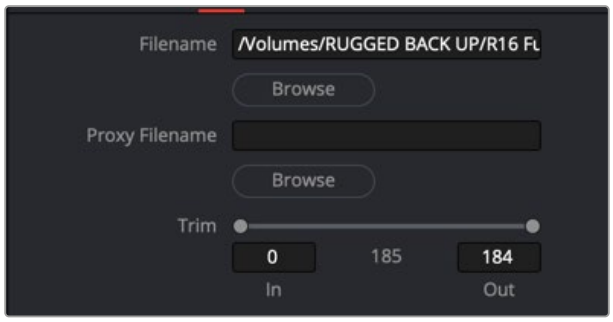
At the top of the Inspector are the Global In and Global Out settings. This range slider determines when in your composition the clip begins and ends. It is the equivalent of sliding a clip along a track in a Timeline. The Hold First Frame and Hold Last Frame dials at the bottom of the Inspector allow you to freeze frames in case the clip is shorter than the composition's global time.



Use the Global In/Out range slider to slide a clip in time to have it appear at the correct time in a comp.

Below the filename in the Inspector is a Trim In and Out range slider. This range slider determines the start frame and end frame of the clip. Dragging the Trim In will remove frames from the start of the clip, and dragging the Trim Out will remove frames from the end of the clip.

Although you may remove frames from the start of a clip, the Global In always determines where in time the clip begins in your comp. For instance, if the Loader has a Global In starting on frame 0, and you trim the clip to start on frame 10, then frame 10 of the source clip will appear at the comp's starting point on frame 0.



Use the Trim In/Out to remove unnecessary frames from the start or end of a clip.

Instead of using the Inspector to adjust timing, it is visually more obvious if you use the Keyframes Editor. For more information on the Keyframes Editor and adjusting a clip's time, *see Chapter 9, "Animating in Fusion's Keyframes Editor,"* in the Fusion Reference Manual.

Loader Node Inputs

Loader nodes have one Effects mask input and one output. In the case of the Effects mask input, connecting a mask node such as a Polygon or B-Spline node automatically creates an alpha channel in the Loader node.

TIP: If you connect a Mask node without any shapes drawn, that mask outputs full transparency, so the result is that the image output by the MediaIn node is blank. If you want to rotoscope over a MediaIn node, first create a disconnected Mask node, and with the Mask node selected and the Media In node loaded into the viewer, draw your mask. Once the shape you're drawing has been closed, connect the Mask node to the MediaIn node's input, and you're good to go.

Using Proxies for Better Performance

For increased performance, you can do one of two things:

- Generate smaller media files and write them to disk using Optimized Media in DaVinci Resolve
- Render out proxy files using Saver nodes in Fusion Studio

Both applications also allow you to generate proxies on-the-fly without rendering new files to disk using the Proxy and Auto Proxy options in the transport controls area.

To enable the Proxy and Auto Proxy options, you can do one of two things, depending on the version of Fusion you are using:

- In the Fusion page, right-click the empty area behind the transport controls to enable the Proxy option.



Proxy and Auto Proxy options in the transport controls right-click menu.

- In Fusion Studio, click the Proxy (Prx) button in the transport area to enable the usage of proxies.

The Proxy option reduces the resolution of the images as you view and work with them. Instead of displaying every pixel, the Proxy option processes one out of every x pixels interactively. In Fusion Studio, the value of x is determined by right-clicking the Prx button and selecting a proxy ratio from the drop-down menu. For instance, choosing 5 from the menu sets the ratio at 5:1. In the Fusion page, the proxy ratio is set by choosing Fusion > Fusion Settings and setting the Proxy slider in the General panel.



The Proxy menu sets the ratio for skipping pixels when processing.

The Auto Proxy button enables Fusion to interactively degrade the image only while adjustments are made. The image returns to normal resolution when the control is released. Similar to the Prx button in Fusion Studio, you can set the Auto Proxy ratio by right-clicking the APrx button and choosing a ratio from the menu.

When a Loader node is selected, the Inspector includes a Proxy Filename field where you can specify a clip that will be loaded when the Proxy mode is enabled. This allows smaller versions of the image to be loaded to speed up file I/O from disk and processing. This is particularly useful when working with high resolution files like EXR that might be stored on a remote server. Lower resolution versions of the elements can be stored locally, reducing network bandwidth, interactive render times, and memory usage.

The proxy clip that you create must have the same number of frames as the original clip, and if using image sequences, the sequence numbers for the clip must start and end on the same frame numbers. If the proxies are the same format as the original files, the proxies will use the same format options in the Inspector as the originals.

Presetting Proxy Quality

When using Fusion Studio, rather than right-clicking over the Proxy button to set the proxy quality, you can preset the standard and Auto Proxy quality in the Fusion Preferences window. The General pane in the Preferences window includes sliders for both standard Proxy files and Auto Proxy files. These sliders designate the default ratio used to create proxies when the Proxy and Auto Proxy modes are turned on. These settings do not affect the final render quality.

TIP: Even though the proxies are being processed smaller than their original size, the viewers scale the images so they refer to original resolutions.

File Format Options

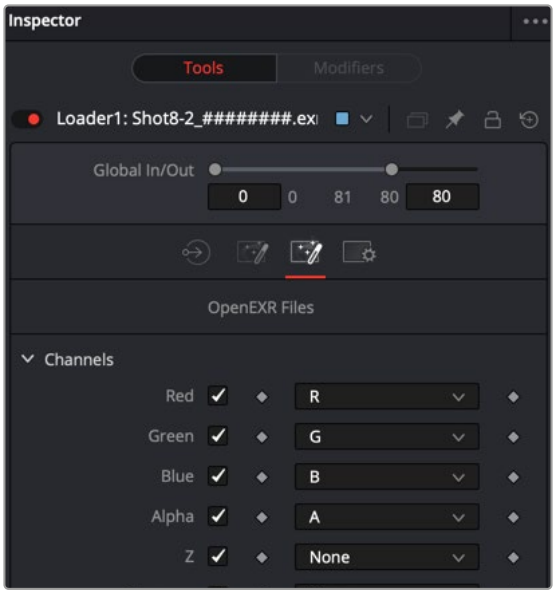
The Fusion interface in DaVinci Resolve and Fusion Studio display specific options for various file formats in slightly different ways. Where Fusion Studio displays most file-specific options in the Loader's Format tab, the Fusion page in most cases displays these options in the main Image tab of the MediaIn node. The only exception being the OpenEXR format. Its extensive options are displayed in a separate tab even in the MediaIn node of the Fusion page. Not all file formats have options. Only the DPX, OpenEXR, PSD, and QuickTime formats provide additional options when loaded.

DPX

The Format tab in Fusion Studio's Loader node for DPX files is used to convert image data from logarithmic to linear. These settings are often left in bypass mode, and the Log to Linear conversion is handled using a Cineon Log node.

OpenEXR

The OpenEXR format provides a compact and flexible high dynamic range (float) format. The format supports a variety of extra non-RGBA channels and metadata. These channels can be viewed and enabled in the Format tab of the Inspector.



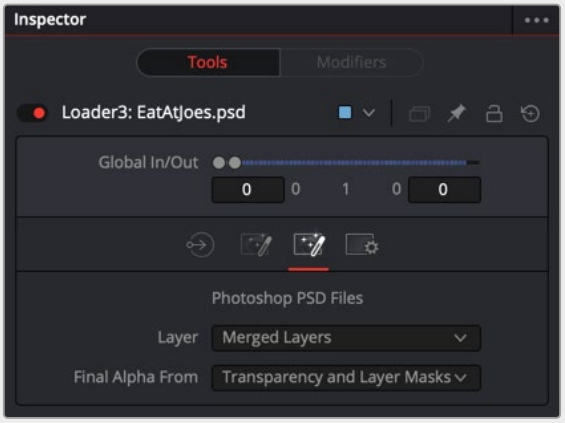
The Format tab in a Loader node Inspector displays Aux channels in EXR files.

Photoshop PSD

There are two methods for importing Photoshop PSD files. You can either import the PSD file and have it represented as a single node in the Node Editor or import the PSD and have each layer represented as a node in the Node Editor. If you do not need independent control over each layer or blend modes are not used when creating the PSD file, then importing the PSD file as a single node will make for a more manageable experience. If you do need control over each layer or Blend modes used in the PSD file are critical, then you should import the file so each layer becomes a node in the Node Editor. Each method is explained below for the Fusion page and Fusion Studio.

Using the Media Pool in DaVinci Resolve's Fusion page: Any PSD file added to the Media Pool in DaVinci Resolve can be accessed from the Fusion page. After dragging the PSD file from the Media Pool into the Node Editor, the image appears as a MediaIn node. From there, you can select which layer to use from the PSD file from the Layer drop-down menu in the Inspector.

Using a Loader node in Fusion Studio: This lets you read in Photoshop PSD files with the ability to select the layer in the PSD file that is used in the comp. Fusion can load any one of the individual layers stored in the PSD file, or the completed image with all layers. Transformation and adjustment layers are not supported.

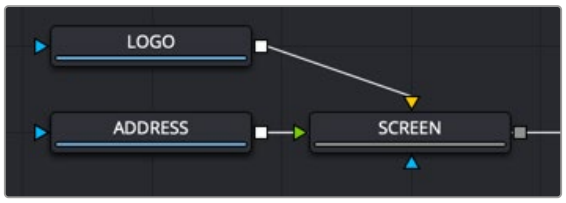


The Format tab in the Inspector displays specific controls for a Photoshop PSD file.

To load all layers individually from a PSD file, with appropriate blend modes, do one of the following:

- In DaVinci Resolve, switch to the Fusion page and choose Fusion > Import > PSD.
- In Fusion Studio, choose File > Import > PSD.

Using either of the methods above creates a node tree where each PSD layer is represented by a node and one or more Merge nodes are used to combine the layers. The Merge nodes are set to the Apply mode used in the PSD file and automatically named based on the Apply mode setting.



The two layers of a Photoshop PSD file are imported and connected to a Merge node set to a Screen Apply mode.

QuickTime

QuickTime files can potentially contain multiple tracks. You can use the Format tab in the Inspector to select one of the tracks.

Loading Audio WAV Files in Fusion Studio

You can load WAV format audio-only files into Fusion Studio. The entire WAV file is loaded into RAM in order to quickly display the waveform in the Keyframe Editor. That being the case, it is best to use the shortest possible audio file you need for the Comp, so as not to use up more memory than necessary.

> **TIP:** AIFF files can be loaded on macOS.

You can either load the audio file independent of any nodes, or load an audio file into the Saver node. The benefit of using a Saver node to load the audio is that you can view the audio waveforms in the Keyframes Editor.

To load a WAV audio file, do the following:

1. Right-click over the speaker icon and select Choose from the contextual menu.
2. In the file browser window, select the audio WAV file track to be used.

To load a WAV audio file using a Saver node, do the following:

1. Add a Saver node to the Node Editor.
2. In the Inspector, click the Audio tab and click the Browse button.
3. In the file browser window, select the audio WAV file track to be used.

To view the Audio Waveform in Fusion Studio, do the following:

1. Open the Keyframes Editor.
2. Expand the Saver track to view the audio waveform.

When you want to find the precise location of an audio beat, transient, or cue, you can slowly drag over the audio waveform to hear the audio.

Chapter 4

Rendering Using Saver Nodes

This chapter covers how to render compositions using Saver nodes in Fusion Studio and the Fusion page in DaVinci Resolve. It also covers how to render using multiple computers over a network when using Fusion Studio.

Contents

Rendering Overview 89
Rendering in the Fusion Page 89
Rendering in Fusion Studio 89
Rendering with the Saver Node 90
Setting Filenames for Export 91
Using the Render Settings Dialog 92
Render Settings Dialog Options 93
Rendering Previews 94
Setting Up Network Rendering in Fusion Studio 95
Licensing for Network Rendering 95
Configuring the Render Master and Render nodes 97
Setting Up the Render Manager 98
Submitting Comps to Network Render 101
Using the Render Settings Dialog for Network Rendering 101
Using the Render Manager Window for Network Rendering 101
Working with Node Groups 102
Viewing the Render Log 103
Using Third-Party Render Managers with Fusion Studio 103
Preparing Compositions for Network Rendering 105
Using Relative Paths 106
Using Mapped Drives 107
Installing All Fonts on Render Nodes 107
Installing Third-Party Plugins on Render Nodes 107

Other Uses of Network Rendering 108
Flipbook Previews 108
Disk Cache 108
When Renders Fail 108
Automatic Rejoining of the Queue 108
Relaunching Render Nodes with Fusion Server 109
Frame Timeouts 109
Heartbeats 110
Managing Memory Use 110
Override Composition Settings 110
Render Several Frames at Once 110
Simultaneous Branching 110
Limitations of Render Nodes 111
Time Stretching 111
Linear Tools 111
Saving to Multi-Frame Formats 111
Troubleshooting 111
Checking the Render Log 111
Check the Composition 112

Rendering Overview

When you have finished creating a composition in Fusion, you need to render the files out to disk for playback and integration into a larger timeline. Fusion Studio and the Fusion page in DaVinci Resolve use very different workflows for rendering. To finish a composite in the Fusion Page, you use a MediaOut node to cache the results into the Edit or Cut page Timeline. The DaVinci Resolve Deliver page handles the final render of the entire Timeline. To get completed composites out of Fusion Studio, you configure and render them starting with a Saver node in the Node Editor. Fusion Studio is also capable of distributing a variety of rendering tasks to other machines on a network.

Rendering in the Fusion Page

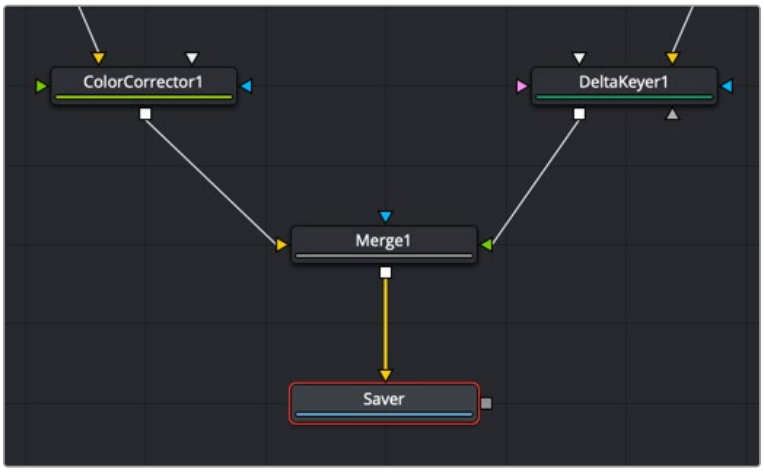
In the Fusion page, a MediaOut node is required for getting your composite from the Fusion page back into the Edit or Cut page Timeline. Whatever the MediaOut node displays when you see it in the viewer is what gets rendered back into the Edit or Cut page. This process is semi-automatic in DaVinci Resolve, where the Smart Render Cache setting begins caching the MediaOut node almost immediately when you return to the Edit or Cut page Timeline. The cache file format and any resolution scaling to fit the composition into the Timeline Resolution is handled in the DaVinci Resolve Project Settings.

Rendering in Fusion Studio

In Fusion Studio, all rendering goes through Saver nodes. Similar to MediaOut nodes in the Fusion page, Saver nodes are most often appended to the end of a node tree to render the final composite. The Saver node determines the name, format, and location of the rendered files.

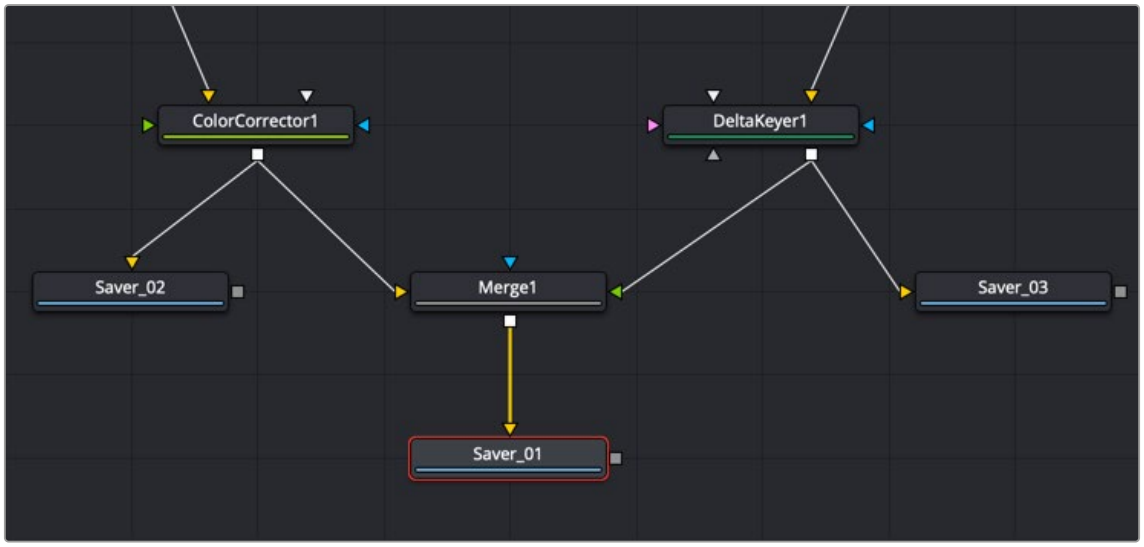
Rendering with the Saver Node

To begin rendering in Fusion Studio, you must add at least one Saver node to the node tree. Most of the time, you will place at least one Saver node at the very end of your tree to render the final image.



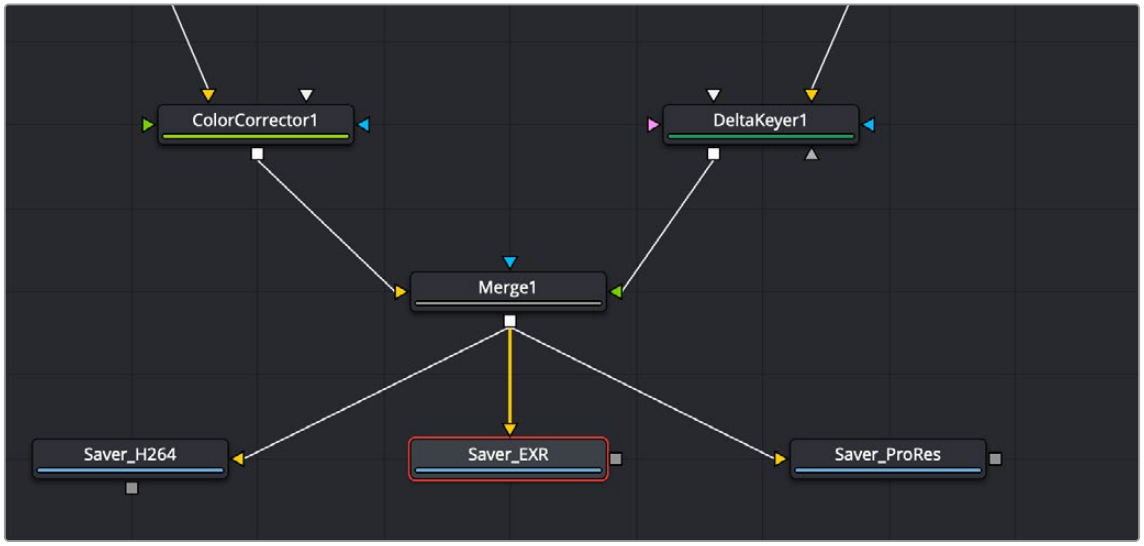
A single Saver node is added to the end of a node tree to render the final composite.

You can attach multiple Saver nodes anywhere along the node tree to render out different parts of a composite. In the example below, three Saver nodes are added at different points in the node tree. The top two render out each half of the composite while the bottom renders the results of the entire composite.



Multiple Saver nodes can be added to different parts of a node tree.

You can also use multiple Saver nodes stemming from the same node in order to create several output formats. The below example uses the three Savers to export different formats of the same shot.



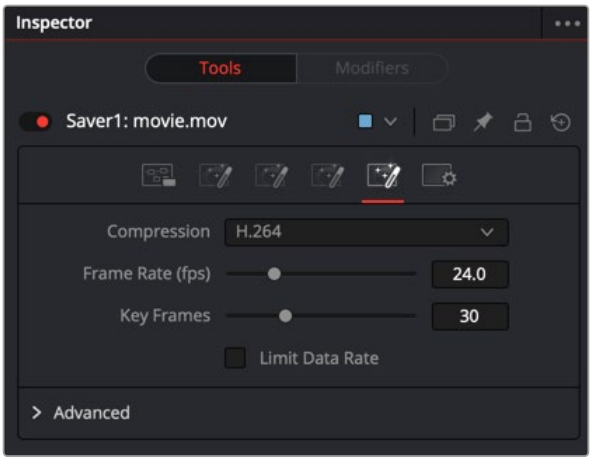
Multiple Saver nodes can be added to create different formats for output.

Adding a Saver node to a node tree automatically opens a Save dialog where you name the file and navigate to where the exported file is saved. You can then use the Inspector to configure the output format.

> For more information on the Saver node, *see Chapter 45, "I/O Nodes,"* in the Fusion Reference Manual.

Setting Filenames for Export

If you use a file extension when naming the file, Fusion will set the output format accordingly. For example, naming your file image_name.exr will set the Inspector to output an EXR file or naming a file image_name.mov will set the Inspector for an H264 QuickTime movie. If you decide to change or modify the setting of the file type, the Saver's format tab in the Inspector contains the specific parameters for the selected format.



The Saver's Format tab with controls for QuickTime

If you decide to output an image sequence, a four-digit frame number is automatically added before the filename extension. For example, naming your file image_name.exr results in files named image_name0000.exr, image_name0001.exr, and so on. You can specify the frame padding by adding several zeroes to indicate the number of digits. For example, entering a filename as image_name_000.exr results in a sequence of images with the names Image_name_000.exr, Image_name_001.exr, Image_name_002.exr, and so on.

NOTE: The starting frame number always uses the Time Ruler start frame number.

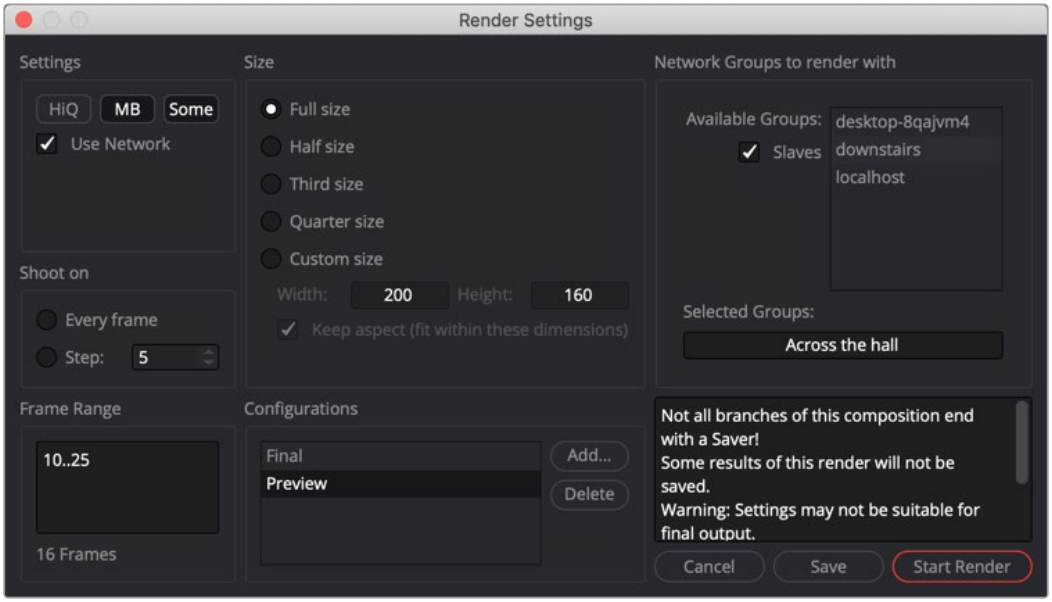
Using the Render Settings Dialog

In Fusion Studio, you initiate rendering by clicking the Render button and opening the Render Settings dialog. This dialog configures the quality, frame range, and network usage of the rendering.

To Render a comp in Fusion Studio:

1. Connect a Saver node at the end of your composition.
2. Enter a name and a location for the saved file(s) in the Save window.
3. Set the format using the Format tab in the Inspector, if necessary.
4. Click the Render button in the transport controls area or choose File > Render All Savers.

The Render Settings dialog opens providing options for the rendered output.



The Render Settings dialog options

Ensure that the frame range and other parameters are correct and click Start Render.

Saver nodes in the DaVinci Resolve Fusion page

The Fusion page also includes a Saver node, although it is used for a different purpose than the Saver node in Fusion Studio. Rendering from the Fusion page is handled primarily through the MediaOut node. There is no Render Settings dialog since the rendering location and format is predetermined by DaVinci Resolve's cache settings. Saver nodes in the Fusion page are a special case, and are used to render OpenEXR files only. Once you add a Saver node, you enter the filename with the extension. exr. Click the Browse button to select a location for the EXR sequence, and then choose Fusion > Render All Savers.

Using the Saver node is useful for optimizing extremely complex and processor-intensive compositions. For example, you can render out specific branches of a node tree that no longer requires frequent adjustment to OpenEXR via a Saver node, and then reimport the result to take the place of the original branch of nodes in order to improve the performance of your composition.

Alternatively, you can render out multi-channel mattes or EXR images containing Arbitrary Output Variables (AOVs) to bring into other applications.

Render Settings Dialog Options

Many of the options in the Fusion Studio Render Settings dialog are used when you need to create a quick preview or a test render. The options in this dialog allow you to increase performance by disabling some of the image-processing operations that are time consuming but deliver higher quality results. Often the first settings you set are those in the Configurations section. This section determines whether you want to produce a final high quality render or a faster preview render. Selecting Final prevents you from modifying the options that will limit the quality.

Settings

When the Configuration section is set to Preview, the Settings section of the Render dialog includes three options that determine the overall quality and appearance of your final output. These buttons also have a significant impact on render times. When the Configurations setting is set to Final, these options cannot be disabled

HiQ: When enabled, this setting renders in full image quality. If you need to see what the final output of a node would look like, then you would enable the HiQ setting. If you are producing a rough preview to test animation, you can save yourself time by disabling this setting.

MB: The MB in this setting stands for Motion Blur. When enabled, this setting renders with motion blur applied if any node is set to produce motion blur. If you are generating a rough preview and you aren't concerned with the motion blur for animated elements, then you can save yourself time by disabling this setting.

Some: When Some is enabled, only the nodes specifically needed to produce the image of the node you're previewing are rendered.

Size

When the Configurations section is set to Preview, you can use the Size options to render out frame sizes lower than full resolution. This is helpful when using the Render dialog to create proxies or just creating a smaller file size.

Network

The Network setting controls the distribution of rendering to multiple computers. For more information, see the network rendering section in this chapter.

Shoot On

Again, this option is only available when Configurations is set to Preview. The Shoot On setting allows you to skip frames when rendering. You can choose to render every second, third, or fourth to save render time and get faster feedback. You can use the Step parameter to determine the interval at which frames are rendered.

Frame Range

Regardless of whether the Configurations is set to Final or Preview, this option defaults to the current Render In/Out Range set in the Time Ruler to determine the start and end frames for rendering. You can modify the range to render more or fewer frames.

Configurations

When set to Final, the Render Settings are set to deliver the highest quality results, and you cannot modify most of the options in this dialog. When set to Preview, you can set the options to gain faster rendering performance. Once you've created a useful preview configuration, you can save it for later use by clicking the Add button, giving it a name, and clicking OK.

Rendering Previews

You can render Flipbook previews into a viewer. These Flipbook previews exist entirely within RAM. They are created by right-clicking over a node in the Node Editor and choosing Create > Play/Preview on > Left viewer/Right viewer from the drop-down menu. The Render Settings dialog appears where you can configure the preview and initiate the rendering. You can also Option-drag a node directly from the Node Editor into a viewer. The Render Settings dialog will be displayed, and the preview will appear on the viewer you target.

For more information on rendering RAM previews, *see Chapter 7, "Using Viewers."* in the Fusion Reference Manual.

TIP: Option-Shift-dragging a node into a viewer will skip the Render dialog and previously used settings.

Setting Up Network Rendering in Fusion Studio

Fusion Studio is capable of distributing a variety of rendering tasks to an unlimited number of computers on a network, allowing multiple computers to assist with creating network-rendered previews, disk caches, and final renders.

Using the Render Settings dialog or the built-in Render Manager, you can submit compositions to be rendered by other copies of Fusion Studio, as well as to one or more Fusion Render nodes. Rendering can also be controlled through the command line for integration with third-party render managers like Deadline, Rush, and Smedge.

Render nodes are computers that do not have the full Fusion application installed but do have Fusion Render node software installed. The Render node software is not installed by default when you install Fusion Studio, but it can be installed at any time using the Fusion Render node Installer. The installer is located in the Blackmagic Fusion Studio installer.dmg on macOS and the Blackmagic Fusion Studio.zip on Linux and Windows. Fusion Studio is licensed for an unlimited number of Render nodes, so you can install the Render node software on as many macOS, Windows, and Linux computers that you want involved in network rendering.

To install a Render node:

1. Download and unzip the Blackmagic Fusion Studio archive. Then locate the Render Node Installer.
2. Copy the Install Fusion Render Node [version] to each computer on the network that you want to perform rendering operations.
3. Install the Render node.

By default, the Render node application will be added to the Start Menu on Windows under Blackmagic Design. On macOS, it is added to the menu bar, and on Linux it appears in the app launcher. Each time you log in to the computer, the Render node application will run automatically.

To disable the Render node application from starting up automatically, choose Quit from the Render node icon in the macOS menu. On Linux, right-click over the icon and choose Kill Process, and on Windows, delete the shortcut from the Windows Startup directory.

Licensing for Network Rendering

Most versions of Fusion Studio are licensed by connecting a single-seat hardware key (dongle) to the same computer where Fusion is installed. Each dongle includes an unlimited number of cross-platform Render node licenses, which you can install on as many macOS, Windows, and Linux computers as you need. For Fusion Studio to access the Render nodes, the computer with the Fusion Studio dongle needs to be on the same local network subnet as the Render nodes. The network licensing does not require individual license files; instead, the Render nodes automatically search for the dongle on the subnet, making it easy to set up. Single-seat dongles do not “float” over a network; they must be connected to the same computer where Fusion Studio operates.

Multi-License Dongles

Using a multi-license dongle, you can license 10 copies of Fusion Studio by connecting the dongle to any computer on the same subnet. Since these licenses “float” over a network, Fusion Studio does not have to be running on the same computer where the dongle is connected. As long as Fusion Studio is on the same subnet, it can automatically find the license server and check out an available license.

Multi-seat dongles can be combined together to tailor the number of Fusion seats in a larger facility. For example, three dongles each licensed for 10 Fusion Studios would serve up 30 licenses. This also allows for redundancy. For instance, in the example above, three computers can act as license servers. If the first server fails for some reason, Fusion Studio will automatically try the next server. Alternatively, multiple dongles can also be plugged into a single computer.

Setting Up a License Server

Setting up the license for network rendering begins by connecting either a single-seat or multi-seat dongle to a computer that will act as the host for the license server. The Render node installer installs Fusion Server, which acts as the license server, although the Render node does not have to run on that same computer. The Fusion Server is set up to launch at startup and run unobtrusively in the background as a service/daemon, ready to serve licenses and Fusion bins. The Fusion Server is used to serve up multiple licenses so it must be running whenever you want to operate Fusion Studio or use the Render nodes for network rendering. Its default is to start up automatically and remain running as long as a license is being used by another computer or a Render node is operating. If nothing is using the Fusion Server, it will quit after 30 seconds.

You will need your network administrator to set firewall rules allowing the Fusion Server, FusionScript, and the Fusion Render node applications to communicate and confirm licensing with the computer that has the Fusion Studio dongle.

If for some reason you remove a dongle or the network drops out, the licenses of any connected Fusion Studio application will also drop. Upon losing its license, Fusion Studio will start searching for another license, locally or on a different machine. If no license is found, Fusion pauses rendering and displays a dialog with options to retry the search or autosave the comp and quit. Render nodes only check for a license on the network once during startup, so they are not affected by removing the dongle or network issues.

Setting Up a License Server with Environment Variables

Environment variables provide a way to specify flexible or “variable” configuration options. When network rendering with Fusion, environment variables can be useful for temporarily setting a location or choosing a preference file. Using the FUSION_LICENSE_SERVER environment variable, you can set different locations for the File Server.

Instead of looking in a single location for the Fusion Server, you can set up multiple license servers separated by semicolons. For instance, fu:SetPrefs("Global.EnvironmentVars.FUSION_LICENSE_SERVER", "192.168.1.12; 192.168.10.55;*")

You can also use the environment variable to scan for license servers within a subnet—for example, “bobs-mac.local;10.0.0.23;*;license.mystudio.com”. Including an asterisk (*) indicates a broadcast search of the local subnet.

Like most environment variables, you can put the license server in the Global Preferences via the Prefs text file. EnvironmentVars: fu:SetPrefs("Global.EnvironmentVars. FUSION_LICENSE_SERVER", "10.0.0.23;*") fu:SavePrefs(), *see Chapter 15, “Preferences,”* in the Fusion Reference Manual for more information on using environment variables.

NOTE: The use of straight quotes (" ") in the environment variables above are intentional and should not be replaced with typographer’s, or curly, quotes (“ ”).

Configuring the Render Master and Render nodes

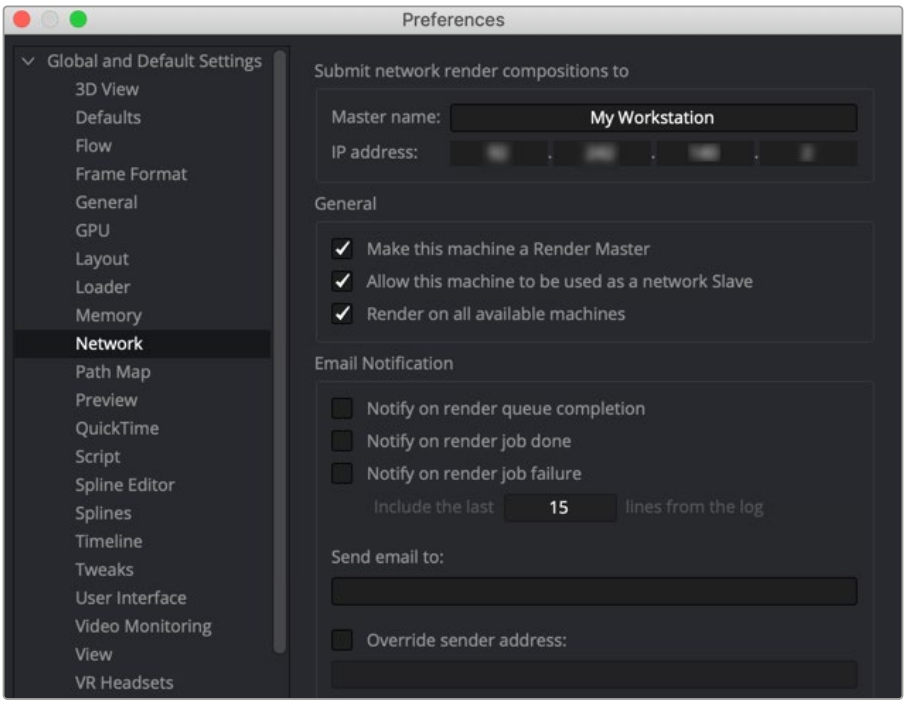
There are two roles played by computers involved in network rendering.

- The Render Master manages the list of compositions to be rendered (the queue) and allocates frames to Render nodes for rendering. Metaphorically speaking, the Render Master is the traffic cop of this process.
- The Render nodes are the main computers used for the rendering process. All computers involved in network rendering must be on the same network subnet, and they all must have access to the various files (including Fonts and third-party plugins) used to create the composite. The path to the files must be the same for each computer involved in rendering.

Preparing the Render Master

The Render Master manages the list of compositions to be rendered (the queue) and allocates frames to Render nodes for rendering. The Render Master is also used to maintain the Render node list and to push updates to the various Render nodes when needed. At least one computer in the render farm must be configured to act as the Render Master.

Any copy of Fusion can act as a Render Master by setting up the Fusion Network Preferences.



The Network panel sets up computers as Render Masters.

Acting as a Render Master has no significant impact on render performance. The system resources consumed are insignificant. However, there are specific steps you must take to make one of your computers a Render Master.

To set up the Render Master:

1. Install a copy of Fusion Studio on the computer you want to be the Render Master.
2. In Fusion Studio, choose Fusion Studio > Preferences on macOS or File > Preferences on Windows and Linux
3. In the Preferences dialog, select the Global > Network Preferences panel.
4. Enter the name of the Render Master in the Name field and enter the IP address.

5 Enable the Make This Machine a Render Master checkbox.

6 If you want to use this computer as part of the render farm, enable the Allow This Machine to Be Used as a Render Node checkbox as well.

To have the Render Node act as the Render Manager:

- Select a node in the Render Manager and choose Set Default Master from the contextual menu.

Once a computer is enabled to act as the master, use the Render Manager to add the Render nodes it will manage. The Render Manager dialog is described in detail later in this chapter.

Preparing Render Nodes

Before you can begin rendering on the network, the Render nodes must be set up to accept instructions from the Render Master.

In Fusion Studio, you can enable the computer to be used as a Render node in two ways:

- Choose File > Allow Network Renders.
- Enable the Allow This Machine to Be Used as a Network Render Node in the Global > Network Preferences.

On a macOS Render Node computer:

- Click the Render Node icon in the menu bar and choose Allow Network Renders.



The Render Node menu accessed in the macOS menu bar

On a Linux Render Node computer:

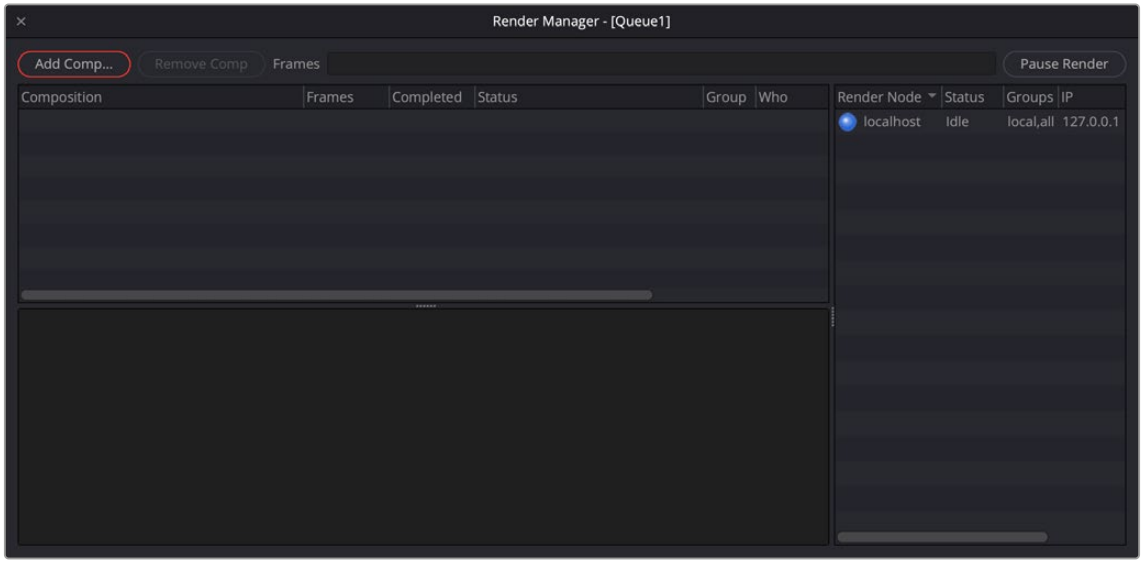
- Right-click the Fusion Render Node icon in the App Launcher and choose Allow Network Renders.

On a Windows Render Node computer:

- Right-click the Fusion Render Node icon in the taskbar Notification area and choose Allow Network Renders.

Setting Up the Render Manager

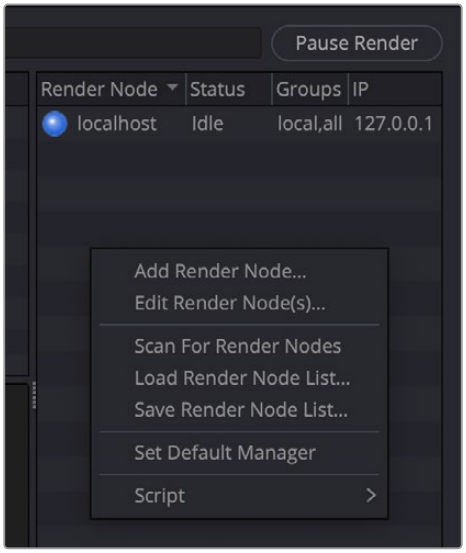
The Render Manager window is used to monitor the progress of rendering. It can be used to reorder, add, or remove compositions from a queue, and to manage the list of Render nodes used for rendering. To open the Render Manager window in Fusion Studio, choose File > Render Manager.



The Render Manager is used to reorder, add, and remove compositions from a render queue.

The Render Master is always listed as the first computer in the Render Node list along the right side. This allows the Render Manager to render local queues without using the network. For the Render Master to control additional Render nodes, the nodes must be added to the Render Node list.

Right-clicking in the Render Node list allows you to add Render nodes by entering the Render node's name or IP address. You can also choose Scan to have the Render Manager look for Render nodes on the local network.



The Render Manager is used to add Render nodes.

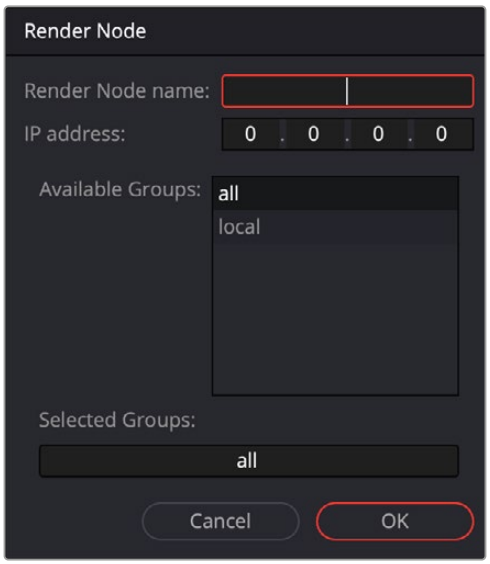
Scanning for Render Nodes

With the Render Manager open, you can scan for Render nodes by right-clicking in the Render Manager's Render Node list and choosing Scan for Render Nodes from the drop-down menu.

Scanning looks through all IP addresses on the network subnet to determine whether any other computers in the local network are actively responding on the port Fusion uses for network rendering. A copy of the Fusion Render node must be running on the remote computer in order for it to be detected by the scan.

Manually Adding Render Nodes

To manually add a Render node to the Render Node list, right-click in the Render Manager's Render Node list and choose Add Render Node from the drop-down menu.



The Add Render Node dialog allows you to manually enter a Render Node name or IP address to locate it on the subnet.

In the Add Render Node dialog that opens, enter the name or the IP address of the remote Render node. The Render Manager will attempt to resolve names into IP addresses and IP addresses into names automatically. You can use this method to add Render nodes to the list when they are not currently available on the network.

Removing Render Nodes

To remove a computer from the Render Node list, right-click over the Render Node in the Render Node list and choose Remove Render Node(s) from the pop-up menu. You can use Command on macOS or Ctrl on Windows and Linux to select multiple Render nodes for removal.

Loading and Saving Render Node Lists

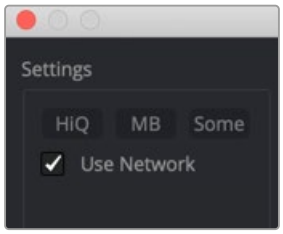
The list of Render nodes is automatically saved in the Documents > Blackmagic Design > Fusion > Queue folder when you quit the Render Manager. You can save and reload alternative lists of Render nodes by choosing Render Node > Save Render Node List and Load Render Node list from the menu.

Submitting Comps to Network Render

To submit a comp to render on the network, you can use the Render Manager, the Render Settings dialog, or a third-party render farm application. The Render Settings dialog is quicker, while the Render Manager and third-party render farm applications can provide more feedback and control over the process.

Using the Render Settings Dialog for Network Rendering

When starting a preview or a final render, selecting the Use network checkbox from the Render Settings dialog and submitting the render adds a composition to the end of the current queue in the Render Manager. The Render Master used is based on the Fusion preferences from the workstation submitting the comp.



The Use Network checkbox enables network rendering from the Render Settings dialog.

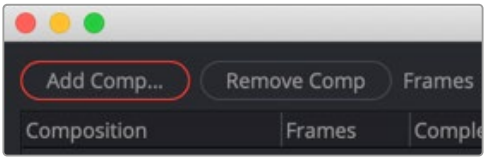
NOTE: Distributed network rendering works for image sequences like EXR, TIFF, and DPX. You cannot use network rendering for Quicktime, H264, ProRes, or MXF files.

Using the Render Manager Window for Network Rendering

The Render Manager uses a render queue that lets you batch render comps. Compositions are rendered in the order in which they are listed in the Render Manager, with the top entry rendered first, followed by the next item down, and so on. Multiple comps in a queue may render simultaneously depending on the group of Render nodes they are using and the priority assigned to each comp.

To add a comp to the queue in the Render Manager:

- Click the Add Comp button and navigate to the comp on your hard drive.
- Right-click in the queue list and select Add Comp from the drop-down menu, and then navigate to the comp on your hard drive.
- Drag a comp file from an OS window into the Render Manager's queue list.



The Add Comp button in the Render Manager adds a comp to the queue for batch rendering over the network.

Removing a Composition from the Queue

To remove a composition from the queue, select the composition in the queue list and press the Backspace/Delete key or right-click over the comp in the queue list and choose Remove Composition from the drop-down menu.

Saving and Reloading Queue Lists

It can be useful to save a queue list to reuse at a later time. The current queue list is saved in the Documents > Blackmagic Design > Fusion > Queue folder. To save the current queue with a new name, choose File > Save Queue As in the Render Manager menu bar. To reload a saved Queue, choose File > Load Queue and navigate to the saved location.

Reordering the Queue Lists

In the middle of a job, priorities for finishing a composition may change. Shifting deadlines may require that a composition further down in the queue be rendered sooner. You can move comps to a new position in the queue by dragging them in the queue list. If a composition with a status set to Done is moved lower in the queue, it does not re-render. To re-render a comp, right-click on the comp in the queue and choose Clear Completed frames from the drop-down menu.

Working with Node Groups

Render nodes can be configured into Groups, which are then used when submitting compositions. For example, imagine you have five Render nodes. All the Render nodes are members of the group named All. Two of the Render nodes include more memory and faster processors, so you create a new group called Hi-Performance.

New Render nodes are automatically added to All, but you can assign them to other groups as well.

To assign a Render node to a Group:

1. Open the Render Manager and select the Render nodes to assign to a group.
2. Choose Render Node > Assign Group or right-click over the Render nodes and choose Assign Group from the drop-down menu.
3. In the Choose Group dialog, enter a name for the group.
4. To assign Render nodes to multiple groups, separate the name of each group using a comma (e.g., All, Local, or Hi-Performance). The order of the groups determines the priority. See “Using Multiple Groups” below.

When a render is submitted to the network, it is automatically sent to the All group. However, you can choose to submit it to other groups in the list.

To submit a comp to a group from the Render Manager:

1. Open the Render Manager.
2. Submit the comp.
3. Click the Pause Render button.
4. Right-click over the comp in the queue list and select Assign Group.
5. In the Assign Group dialog, select an existing group to render the comp and click OK.
6. Click the Resume Render button.

To submit a comp to a group from the Render Settings dialog:

1. Click the Render button in the transport controls area.
2. In the Render Settings dialog, enable the Use Network checkbox.
3. Select an existing group from the Available groups list.
4. Click the Start Render button.

Continuing with the group example above, five Render nodes are contained in the All group, and two of those Render nodes are also in the Hi-Performance group. If you submit a render to the Hi-Performance group, only two of the computers on the network are used for rendering. If a composition is then submitted to the All group, the remaining three machines will start rendering the new composition. Once the two Render nodes in the Hi_Performance group complete the first job, they join the render in progress on the All group.

Groups are optional and do not have to be used. However, groups can make managing large networks of Render nodes easier and more efficient.

Using Multiple Node Groups

A single Render node can be a member of multiple groups. A single composition can also be submitted to multiple groups. Submitting a composition to multiple groups results in it rendering on all Render nodes in the selected groups.

When a Render node is a member of multiple groups, the order of the groups is important because the order defines the priority for that Render node.

For example, if groups are assigned to a Render node as All, Hi-Performance, then renders submitted to the All group take priority. Any renders in progress that were submitted to the Hi-Performance group will be overridden. If the order is changed to Hi-Performance, All, then the priority is reversed.

Viewing the Render Log

The Render Log is displayed in the lower half of the Render Manager window, although it can also be displayed in the console window. The text in the log displays the Render Manager activities, including which frame is assigned to which Render node, which Render nodes have loaded the compositions in the queue, and statistics for each render after completion.

To view the Render Log in the console:

- Open the Render Manager and choose Misc > Show Render Log.
- There are two modes for the Render Log: a Verbose mode and a Brief mode. Verbose mode logs all events from the Render Manager, while Brief mode logs only which frames are assigned to each Render node and when they are completed.

To disable Verbose mode:

- Choose Misc > Verbose Logging from the Render Manager's menu bar.

Using Third-Party Render Managers with Fusion Studio

You can make use of third-party render manager software to control network rendering. This allows for efficient sharing of your computer resources between the many applications that may make use of them. Examples of such managers are Smedge from Uberware LLC, Rush from Seriss, and Deadline from GetRender. Generally, these render managers use a command line renderer. By default, Fusion's Render nodes operate as a service to the Fusion internal Render Manager. However, you can also run the Render nodes via the command line for third-party render managers.

Keep in mind that using a third-party render manager will prevent the use of some of Fusion's network rendering features, such as the ability to create network rendered Flipbook Previews and disk caches.

Command Line Rendering

For studios using third-party render farm managers like Deadline, Smedge, or Rush, the Fusion Render node can be called via command line passing arguments and file paths. In this Windows example, a Render node is called to load a composition called exampleV001, and render 10 frames:

```
//pathtoRN/FusionRenderNode.exe //pathtoProject/exampleV001.comp -render -start
101 -end 110 -quit
```

This would start up, render frames from 101 to 110, and then quit.

The following table lists additional command line features.

Command	Description
"Fusion Server -i"	Installs the license server as a service or daemon, launching it on startup before user login.
"Fusion Server -S" (capital S)	Causes the Fusion Server to run persistently in the background, until force-quit.
<filename.comp>	Full path and comp name, like /storage/project/episode/shot/filename.comp.
-render	Tells the Render Node to render.
-frames <frameset>	Passes a series of frame ranges to be rendered—e.g., 101..110,120,121,130..150.
-start <frame>	Sets the start frame of the render.
-end <frame>	Sets the last frame of the render.
-step <step>	Normally set to 1, step skip frames for rendering. For instance, 2 would render every second frame
-quit	Causes the Render Node to quit after the render is complete.
-join <host>	Prompts the node to connect to a manager at <hostname,IP>, and (re)join any ongoing renders.
-listen	The node remains running and waits for incoming requests from a manager.
-log <filename>	Causes the Render Node to output information about the render to a log file. This appends to the end of an existing log file.
-cleanlog	Clears existing text from a log file.
-verbose	Outputs more detailed information into the log file.

Command	Description
-quiet	Suppresses pop-ups and interface buttons from displaying and needing interaction.
-version	Returns the Render Node version number.
-pri high|above|normal|below|idle	Sets the node's process priority to high, above, normal, below, or idle.
-args <arg1> [, <arg2> ...]	Allows storing custom values that can be fetched by calling the script function GetArgs(), which will return a table of { <arg1>, <arg2>, ... }

TIP: An X11 virtual frame buffer is required to make a headless Linux command line interface work.

Preparing Compositions for Network Rendering

The way you construct a composition in Fusion Studio can help or hinder network rendering. The media you read in, where plugins are installed, and the mix of operating systems on your networked computers all play a part in how smoothly your network rendering goes. Your setup must include several essential parts before network rendering will work:

- License dongle, Render Master, and Render nodes must be on the same local network (subnet).
- Fusion Server must be running as a background service on the same computer where the dongle is installed.
- All source media from the comp should be placed on a network volume.
- The network volume must be mounted on each Render node.
- Loaders must point to the media on the mounted volumes.
- Savers must write to a drive that is mounted on each Render node.
- The Fusion comp must be saved to a volume that is mounted on each Render node.
- All Render Nodes and Render Masters need read and write access to any volumes specified as a source media location or render destination.
- Make sure all fonts used in the comp for Text+ and 3D text nodes are installed on all the Render nodes
- Make sure all Render nodes have third-party OFX plugins installed if any are used in the comp.

Below are more details about some of these items.

Using Relative Paths

The file paths used to load a composition and its media, and to save the composition's rendered results, are critical to the operation of network rendering. Each computer used for network rendering must be able to access the media location for each Loader in the comp. Savers must be set up to save to folders that all Render nodes can see and to which all Render nodes have write permission. Even the composition must be saved in a folder accessible to all Render nodes, and it should be added to the queue list using a path that's visible to all Render nodes.

For example, if you open a composition located at c:\compositions\test1.comp in Fusion Studio and add the composition to the network rendering queue, the Render Manager sends a message to each Render node to load the composition and render it. The problem is that each computer is likely to have its own c:\drive that does not contain the comp you created. In most cases, the Render nodes will be unable to load the composition, causing it to fail.

Path Maps located in Fusion Preferences are virtual paths used to replace segments of file paths. They can change the absolute paths used by Loader and Saver nodes to relative paths. There are a number of Path Maps already in Fusion, but you can also create your own. The most common path to use is the Comp:\ path.

Comp:\ is a shortcut for the folder where the actual composition is saved. So, using Comp:\ in a Loader makes the path to the media file relative, based on the saved location of the comp file. As long as all your source media is stored in the same folder or subfolder as your comp file, Fusion locates the media regardless of the actual hard drive name.

Here's an example of a file structure that enables you to use relative file references. The composition is stored in the following file path:

`Volumes\Project\Shot0810\Fusion\Shot0810.comp`

And your source media is stored here:

`Volumes\Project\Shot0810\Fusion\Greenscreen\0810Green_0000.exr`

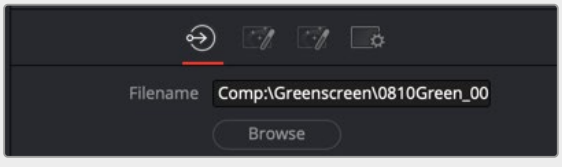
This overall directory structure can be seen in the following screenshot:



File paths can use relative paths based on the location of the saved comp file.

In this situation, using the Comp:\ path means your media location starts from your comp file's location. The relative path set in the Loader node would then be:

`Comp:\Greenscreen\0810Green_0000.exr`



Replacing the Loader's path to start with Comp:\ creates a relative path from the comp file's location.

If your source media's actual file path uses a subfolder in the same folder as the comp file's folder:

```
Volumes\Project\Shot0810\Footage\Greenscreen\0810Green_0000.exr
```

The relative path set in the Loader node would then be:

```
Comp:\..\Footage\Greenscreen\0810Green_0000.exr
```

The two dots .. instruct the path to go up one folder.

TIP: Some Path Maps are not set up on a Fusion Render Node automatically. For instance, you must manually add an entry for macros if you are using macros in your comp.

Using Mapped Drives

Having the comp and media set to relative paths solves only part of the problem. Each computer involved in rendering needs to have access to the drive where the comp and source media are located. Mapping a drive provides permanent access to a folder on another computer or network storage device.

Windows assigns a new drive letter to the folder, and it can be accessed just like any other drive connected to your computer. Mapped drives assign a letter of the alphabet to a shared network resource. Your shared drives must be the same drive letters on all Render nodes. For example, if your media is on drive Z, then the network drive must appear as the letter Z on each of the Render nodes.

On macOS, you can map a network drive using Connect to Server from the Go menu. Entering the smb:// path to the drive will mount it on the computer. Using Accounts > LogIn Items, you can have the network drive auto-mount after a reboot as well.

Installing All Fonts on Render Nodes

All fonts used by Text tools in the composition must be available to all nodes participating in the render. The render will otherwise fail on the Render nodes that do not have the font installed.

Installing Third-Party Plugins on Render Nodes

All third-party plugins and tools used by a composition must be installed in the plugins directory of each Render node. A Render node attempting to render a composition that uses a plugin that's not installed will fail to render. Licensed plugins are required on each Render node.

Other Uses of Network Rendering

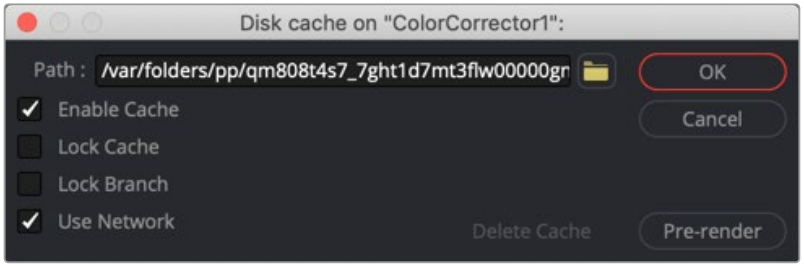
Although you will probably set up network rendering for the purpose of accelerating the output of your final renders, Fusion is capable of using the network for other purposes as well. You can use the Render nodes on your network to accelerate the creation of Flipbook Previews and disk caching as well.

Flipbook Previews

Fusion Studio is able to use Render nodes to accelerate the production of Flipbook Previews, allowing for lightning fast previews. Frames for previews that are not network rendered are rendered directly into memory. Select the Use Network checkbox and the Render nodes to render the preview frames to the folder set in the Preferences Global > Path > Preview Renders. This folder should be accessible for all Render nodes that will participate in the network render. The default value is Temp\, which is a virtual path pointing to the system's default temp folder. This will need to be changed before network rendered previews can function. Once the preview render is completed, the frames that are produced by each Render node are spooled into memory on the local workstation. As each frame is copied into memory, it is deleted from disk.

Disk Cache

Right-clicking a node in the Node Editor and choosing Cache to Disk opens a dialog used to create the disk cache. If you enabled the Use Network checkbox and click the Pre-Render button to submit the disk cache, the network Render nodes are used to accelerate the creation of the disk cache.



Render nodes can be used for disk caching as well as final renders.

When Renders Fail

It is a fact of life that render queues occasionally fail. The composition has an error, the power goes out, or a computer is accidentally disconnected from the network are some causes for failure. If no one is available to monitor the render, the risk that an entire queue may sit inactive for several hours may become a serious problem.

Fusion Studio includes a variety of measures to protect the queue and ensure that the render continues even under some of the worst conditions.

Automatic Rejoining of the Queue

If a Render node becomes unavailable to the Render Master for any reason, frames assigned to that Render node are reassigned among the remaining Render nodes in the list.

When the Render node becomes available for rendering again, it will signal the Render Master that it is ready to render again, and new frames will be assigned to that Render node.

This is why it is important to set the Render Master in the network preferences of the Render nodes. If the master is not set, the Render node will not know what master to contact when it becomes available.

In the Fusion Render Node Preferences, select the Tweaks panel. Using the Last Render Node Restart Timeout field, you can enter the number of seconds Fusion waits after the last Render node goes offline before aborting that queue and waiting for direct intervention.

Relaunching Render Nodes with Fusion Server

Fusion Server is a small utility installed with Fusion Studio and the Render node. The application is Fusion Server is a small utility installed with Fusion Studio and the Render node. The application is silently launched by each Fusion Render node when started.

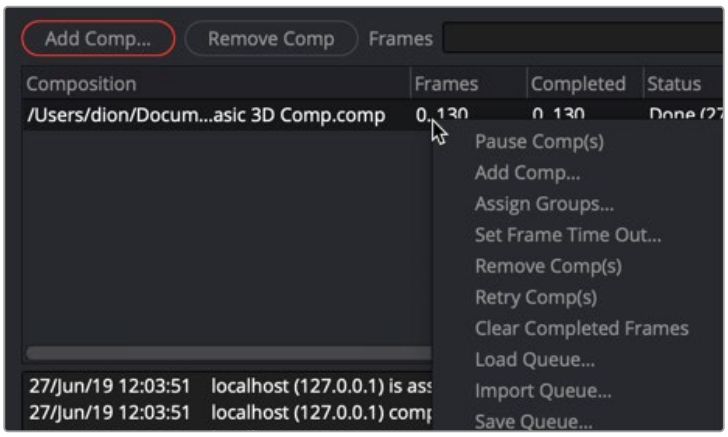
Fusion Server monitors the Render node to ensure that the Render node is still running during a render. It consumes almost no CPU cycles and very little RAM. If the monitored Render node disappears from the system's process list without issuing a proper shutdown signal, as can happen after a crash, the Fusion Server relaunches the Render node, allowing it to rejoin the render.

Fusion Server will only detect situations where the Render node has exited abnormally. If the Render node is still in the process list but has become unresponsive for some reason, the Fusion Server cannot detect the problem. Hung processes like this are detected and handled by frame timeouts, as described below.

Frame Timeouts

Frame timeouts are a fail-safe method of canceling a Render node's render if a frame takes longer than the specified time (with a default of 60 minutes, or one hour). The frame timeout ensures that an overnight render will continue if a composition hangs or begins swapping excessively and fails to complete its assigned frame.

The timeout is set per composition in the queue. To change the timeout value for a composition from the default of 60 minutes, right-click on the composition in the Render Manager's queue list and select Set Frame Timeout from the contextual menu.



Right-click over a comp in the Render Manager to set a timeout value.

To change the frame timeout value, choose Set Frame Time Out from the Render Manager's Misc menu and enter the number of seconds you want for the Time Out.

Heartbeats

The Render Master regularly sends out heartbeat signals to each node, awaiting the node's reply. A heartbeat is basically a message from the manager to the Render node asking if the node is still responsive and healthy. If the Render node fails to respond to several consecutive heartbeats, Fusion will assume the Render node is no longer available. The frames assigned to that Render node will be reassigned to other Render nodes in the list.

The number of heartbeats in a row that must be missed before a Render node is removed from the list by the manager, as well as the interval of time between heartbeats, can be configured in the Network Preferences panel of the master. The default settings for these options are fine for 90% of cases.

If the compositions that are rendered tend to use more memory than is physically installed, this will cause swapping of memory to disk. It may be preferable to increase these two settings somewhat to compensate for the sluggish response time until more RAM can be added to the Render node.

Managing Memory Use

Often, the network environment is made up of computers with a variety of CPU and memory configurations. The memory settings used on the workstation that created a composition may not be appropriate for all the Render nodes in the network. The Render node software offers the ability to override the memory settings stored in the composition and use custom settings more suited to the system configuration of a specific Render node.

To access preferences for a node, right-click on the icon in the Windows Notification area or from the macOS menu bar and choose Preferences. In the Preferences dialog, select the Memory panel.

Override Composition Settings

Enable this option to use the Render node's local settings to render any incoming compositions. Disable it to use the default settings that are saved into the composition.

Render Several Frames at Once

Fusion has the ability to render multiple frames at once for increased render throughput. This slider controls how many frames are rendered simultaneously. The value displayed multiplies the memory usage (a setting of 3 requires three times as much memory as a setting of 1).

Normal values are 2 or 3, although machines with a lot of memory may benefit from higher values, whereas machines with less memory may require the value to be 1.

Simultaneous Branching

Enable this option to render every layer in parallel. This can offer substantial gains in throughput but may also use considerably more memory, especially if many layers are used in the composition. Machines with limited memory may need to have Simultaneous Branching disabled when rendering compositions with many layers.

Limitations of Render Nodes

There are a few important limitations to remember while setting up compositions and rendering over a network.

Time Stretching

Compositions using the Time Stretcher and Time Speed tools may encounter difficulties when rendered over the network. Speeding up or slowing down compositions and clips requires fetching multiple frames before and after the current frame that is being rendered, resulting in increased I/O to the file server. This may worsen bottlenecks over the network and lead to inefficient rendering. If the composition uses the Time Stretcher or Time Speed tools, make certain that the network is up to the load or pre-render that part of the composition before network rendering.

Linear Tools

Certain tools cannot be network rendered properly. Particle systems from third-party vendors, such as Genarts's Smoke and Rain, and the Fusion Trails node cannot render properly over the network. These tools generally store the previously rendered result and use it as part of the next frame's render, so every frame is dependent on the one rendered before it. This data is local to the tool, so these tools do not render correctly over a network.

Saving to Multi-Frame Formats

Multiple machines cannot render a single QuickTime file. Always render to separate sequential file formats like EXR, DPX, JPEG, and so on. Once the render is complete, a single workstation can load the image sequence in order and save to the desired compiled format.

NOTE: The above does not apply to network rendered previews, which are previews created over the network that employ spooling to allow multi-frame formats to render successfully. Only final renders are affected by this limitation.

Troubleshooting

There are some common pitfalls when rendering across a network. Virtually all problems with network rendering have to do with path names or plugins. Return to the "Preparing Compositions for Network Rendering" section in this chapter to review some of the essential setup requirements. Verify that all Render nodes can load the compositions and the media, and that all Render nodes have installed the plugins used in the composition.

If some difficulties persist, contact Blackmagic Design's technical support using the support section on the Blackmagic Design website. Save a copy of the render.log file to send to technical support.

Checking the Render Log

The log file shown in the Render Manager dialog displays messages that can assist with diagnosing why a render or node has failed. The render log shows a step-by-step account of what happened (or didn't happen) during a render. If a Render node cannot be found, fails to load a composition or render a frame, or simply stops responding, it will be recorded here.

Check the Composition

The Render Manager's Status field in the render log indicates if a composition fails to render. Some possible causes of this are as follows:

- **No Render Nodes Could Be Found:** On the Preferences Network tab, make sure that there is at least one Render node available, running and enabled. If all Render nodes are listed as Offline when they are not, check the network.
- **The Composition Could Not Be Loaded:** Some Render nodes may not be able to load a composition while others can. This could be because the Render node could not find the composition (check that the path name of the composition is valid for that Render node) or because the composition uses plugins that the Render node does not recognize.
- **The Render Nodes Stop Responding:** If a network link fails, or a Render node goes down for some reason, the Render node will be removed from the active list and its frames will be reassigned. If no more Render nodes are available, the composition will fail after a short delay (configurable in network preferences). If this happens, check the render log for clues as to which Render nodes failed and why.
- **The Render Nodes Failed to Render a Frame:** Sometimes a Render node simply cannot render a particular frame. This could be because the Render node could not find all the source frames it needed, or the disk it was saving to become full or because of any other reason for which Fusion might normally be unable to render a frame. In this case, the Render Manager will attempt to reassign that failed frame to a different Render node. If no Render node can render the frame, the render will fail. Try manually rendering that frame on a single machine and observe what happens.
- **Check the Render Nodes:** Fusion's Render Manager incorporates a number of methods to ensure the reliability of network renders. Periodically, the Render Manager will send signals called Heartbeats, generated at regular intervals, to detect network or machine failures. In this event, a failed Render node's outstanding frames are reassigned to other Render nodes where possible. In rare cases, a Render node may fail in a way that the heartbeat continues even though the Render node is no longer processing. If a Render node failed (although the Render Master may not have detected it) and you do not want to wait for the Frame Timeout, simply restart the Fusion workstation or Fusion Render Node that has hung. This triggers the heartbeat check, reassigns the frames on which that Render node was working, and the render should continue. Heartbeats may fail if the system that is performing the render is making extremely heavy use of the Swap file or is spending an extraordinary amount of time waiting for images to become available over a badly lagged network. The solution is to provide the Render node with more RAM, adjust memory settings for that node, or upgrade the network bandwidth.
- **Check the Network:** At the Render Master, bring up the Network tab of the Preferences dialog box and click Scan. If a Render node is not listed as running, the Render Master will not be able to contact it for network rendering. Alternatively, bring up a command prompt and ping the Render nodes manually. If the remote systems do not respond when they are up and running, the network is not functioning and should be examined further.

Chapter 5

Working in the Node Editor

This chapter discusses how to work in the Node Editor, including multiple ways to add, connect, rearrange, and remove nodes to create any effect you can think of.

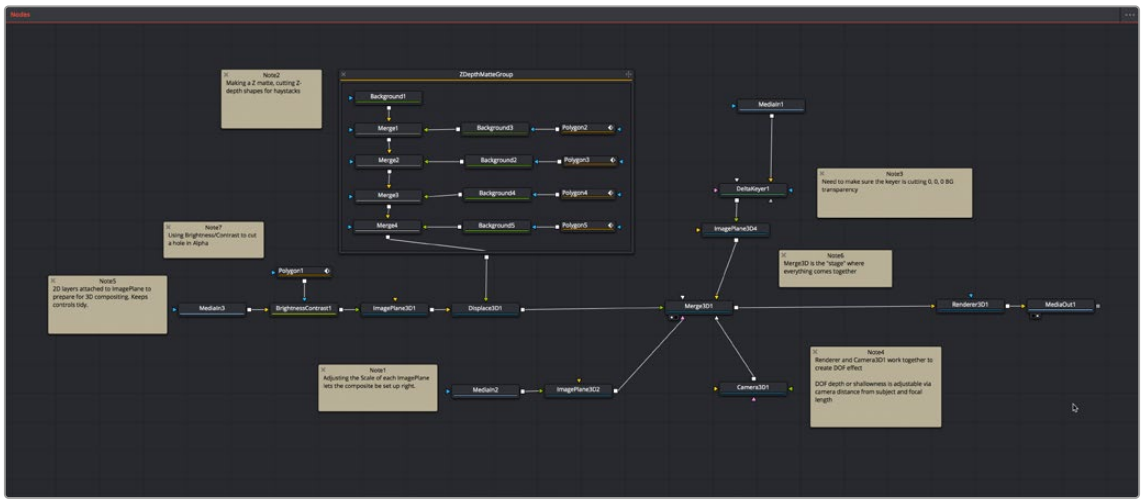
Contents

Learning to Use the Node Editor 115
Navigating within the Node Editor 116
Automatic Node Editor Navigation 116
Using the Node Navigator 116
Node View Bookmarks 117
Adding Nodes to a Composition 119
Adding, Inserting, and Replacing Nodes Using the Toolbar 119
Adding Nodes Quickly Using the Select Tool Window 120
Adding Nodes from the Effects Library 121
Adding, Inserting, and Replacing Nodes Using the Contextual Menu 124
Deleting Nodes 125
Disconnected Nodes 125
Selecting and Deselecting Nodes 125
Loading Nodes into Viewers 126
Viewed Nodes When You First Open Fusion 127
Node View Indicators 128
Drag and Drop Nodes into a Viewer 128
Using the Contextual Menu 128
Clearing Viewers 129
Create/Play Preview 129
Connecting and Disconnecting Nodes 129
Node Basics 129
How to Connect Nodes 130
Dropping Connections on Top of Nodes 130
Identifying Node Inputs 132
Node Order Matters 132

Upstream and Downstream Nodes 135
Disconnecting and Reconnecting Nodes 135
Tracing Connections Through the Node Tree 136
Branching 137
Connecting Merge Nodes 138
Automatically Creating a Merge Node When Adding Nodes 139
Automatically Creating a Merge Node by Connecting Two Outputs 140
Connection Options and Routers 141
Using Routers to Reshape and Branch Connections 142
Swapping Node Inputs 143
Extracting and Inserting Nodes 144
Cut, Copy, and Paste Nodes 145
Cut, Copy, and Paste in the Node Editor 145
Pasting Node Settings 146
Copying and Pasting Nodes to and from Any Text Editor 146
Using Multilayer Nodes (Studio Version Only) 147
Viewer 147
Nodes and Inspector 148
Input Layer Controls 149
Instancing Nodes 149
Using Instanced Nodes 150
De-Instancing and Re-Instancing Specific Parameters 151
Keeping Node Trees Organized 151
Moving Nodes 151
Renaming Nodes 152
Changing Node Colors 153
Using Sticky Notes 153
Using Underlay Boxes 154
Node Thumbnails 155
Choosing Which Nodes Show Thumbnails 156
Switching Thumbnails between Images and Icons 157
Finding Nodes 157
Performing Simple Searches 158
Using Regular Expressions 158
Custom Node Settings 159
Managing Saved Settings 159
Resetting Defaults 160
Saving and Loading Alternate Node Settings 160
Adding Saved Settings from the File System 160
Node Modes Including Disable and Lock 161
Node Editor Options 161
Node Tooltips and the Status Bar 162

Learning to Use the Node Editor

The Node Editor (formerly called the Flow or Flow Editor) is the heart of Fusion's compositing interface. It uses a flowchart structure called a node tree that lets you build a composition out of interconnected nodes, as opposed to using layers in a layer list. Each clip you add to the composition, and each image-processing operation you apply to those clips, is added as a node, all of which are joined together with connections that propagate image data from one node to the next. Each individual node performs a relatively simple operation, but collectively they combine to let you create wonderfully complex results.

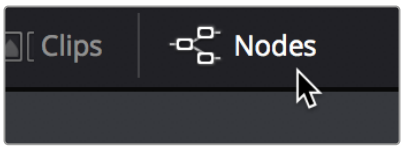


The Node Editor.

This chapter discusses how to work in the Node Editor in greater detail, showing you how to add, connect, rearrange, and remove nodes to create any effect you can think of.

To display the Node Editor:

- Click the Nodes button on the UI toolbar.



The Nodes button in the UI toolbar.

Navigating within the Node Editor

The Node Editor is the place where everything relating to nodes and the construction of your composites happens. The more you learn about how to navigate within the Node Editor, the faster you'll be able to work. There are a variety of standard methods of panning and zooming around the Node Editor, many of which are shared with other panels in Fusion.

Methods of panning the Node Editor:

- Middle-click and drag to pan around the Node Editor.
- Hold down Shift and Command, and then click and drag within the Node Editor to pan.
- Drag with two fingers on a track pad to pan in the Node Editor

Methods of zooming the Nod4e Editor:

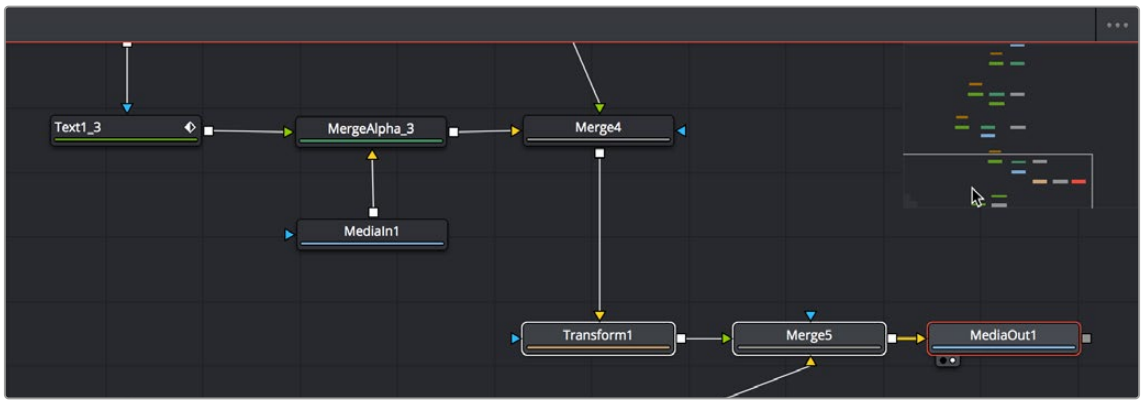
- Press the Middle and Left buttons simultaneously and drag to resize the Node Editor.
- Hold down the Command key and use your pointer's scroll control to resize the Node Editor.
- Right-click the Node Editor and choose an option from the Scale submenu of the contextual menu.
- Press Command-1 to reset the Node Editor to its default size.
- Hold down the Command key and drag with two fingers on a track pad to resize the Node Editor.

Automatic Node Editor Navigation

If a node that is not visible in the Node Editor becomes selected, either by using the Find command or by selecting a node's header in the Inspector, the Node Editor will automatically pan to display the node in the visible area.

Using the Node Navigator

Another useful way to pan around the Node Editor is to use the Node Navigator. The Node Navigator is a small rectangular overview in the upper-right corner of the Node Editor. It gives a bird's eye view of the entire composition, with an inner outline that indicates the portion of the composition that is visible in the panel. You can use the Node Navigator when you are zoomed in on a node tree and want to pan around a composition.



The Node Navigator.

To display or hide the Node Navigator, do one of the following:

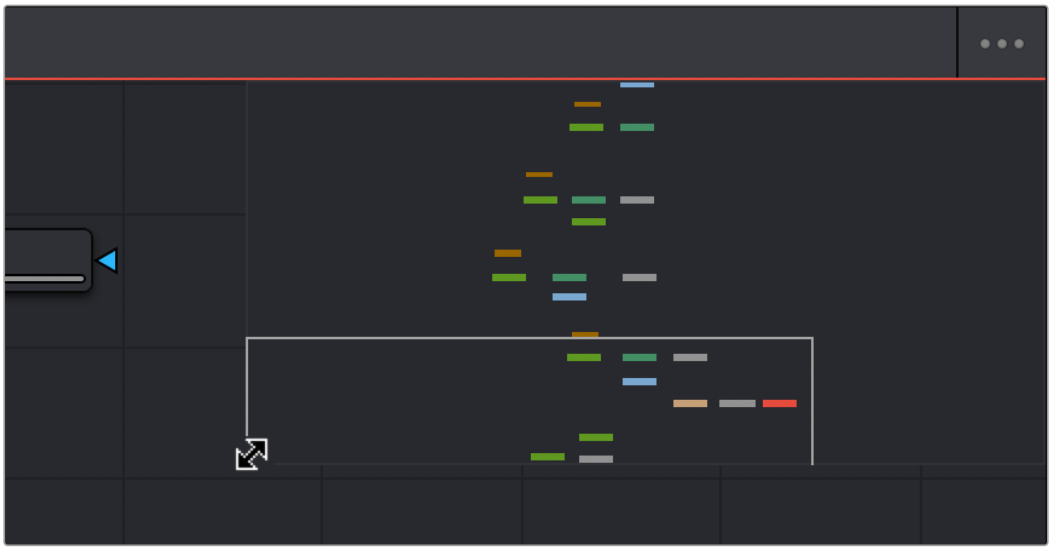
- Right-click in an empty area of the Node Editor, and then choose Options > Show Navigator.
- Press the V key.

To have the Node Navigator resume displaying automatically when needed after you've closed it:

- Right-click in an empty area of the Node Editor, and then choose Options > Auto Navigator.

To change the size of Node Navigator, do the following:

- Drag the lower-left corner of the Navigator to resize it.



Drag the corner to resize the Navigator.

To return to the default Node Navigator size, do the following:

- Right-click anywhere within the Node Navigator and choose Reset Size.

To pan the Node Editor using the Node Navigator, do the following:

- Drag within the Node Navigator to move around different parts of your node tree.
- Within the Navigator, drag with two fingers on a track pad to move around different parts of your node tree.

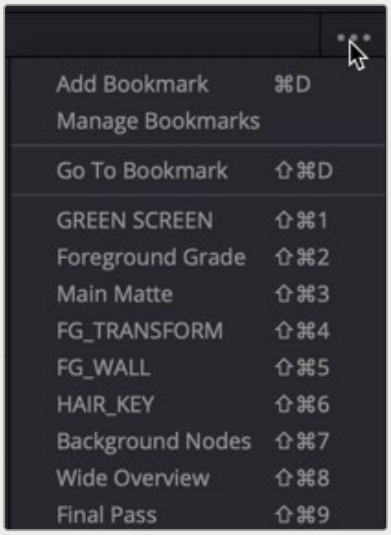
Node View Bookmarks

Bookmarks are another way of navigating the Node Editor. Bookmarks save the position and scale of the Node Editor, so you can quickly and precisely jump from viewing one group of nodes to viewing another.

To add a bookmark, do the following:

1. Pan and scale in the Node Editor to view a group of nodes you are interested in.
2. From the Options menu in the upper right of the Node Editor, choose Add Bookmark, or press Cmd-D.
3. In the Manage Bookmarks dialog that opens, enter a name for the bookmark and click the Add button.

The first nine saved bookmarks are given keyboard shortcuts and listed in the Options menu. They are also listed in the Go To Bookmarks dialog along with any saved bookmarks beyond the initial nine.



The Node Editor Options menu with nine bookmarks added

TIP: You can return the Node Editor to the default scale by right-clicking in the Node Editor and choosing Scale > Default Scale or pressing Cmd-1.

If your Node Tree changes and you want to update Bookmark names or delete bookmarks, those tasks can be done in the Manage Bookmarks dialog.

To rename or delete a bookmark, do the following:

1. From the Options menu in the upper right of the Node Editor, choose Manage Bookmarks.
2. In the Manage Bookmarks dialog that opens, right-click over the bookmark and choose Rename or Remove.
3. Click OK to close the Manage Bookmarks dialog.

Using Bookmarks

You can jump to a Bookmark view by selecting a bookmark listed in the Options menu or choosing Go To Bookmarks to open the Go To Bookmarks dialog. The Go To Bookmarks dialog has all the bookmarks listed in the order they were created in the current composition. Double-clicking on any entry in the dialog will update the Node Editor to that view and close the Go To Bookmarks dialog.

If you have a long list of bookmarks, you can use the search field at the bottom of the dialog to enter the name of the bookmark you want to find.

Changing the Sort Order and Assigning Keyboard Shortcuts

Bookmarks appear in the Options menu and in the Go To Bookmarks dialog in the order they were created. The top nine bookmarks listed are assigned keyboard short cuts. If you want to change the

sort order of the list you can do that in the Manage Bookmarks dialog. This can be useful if you want some bookmarks to have specific keyboard shortcuts, or you want a bookmark you made early in the process to not have a keyboard shortcut assigned to it.

To change the order of the listed bookmarks, do the following:

1. From the Options menu in the upper right of the Node Editor, choose Manage Bookmarks.
2. In the Manage Bookmarks dialog drag a bookmark up or down in the list.
3. An insert line appears where the bookmark will be inserted. Release the mouse button when the insert line is where you want the bookmark to be listed.

The keyboard short cuts will update to reflect the new order.

> **TIP:** You can hold down the Shift key to select multiple bookmarks and move them simultaneously up or down in the Manage Bookmark list.

Using Underlays as Bookmarks

Underlays added to the Node Editor are automatically added as bookmarks. They are listed in the Options menu below the list of your custom bookmarks. You can omit Underlays from showing in the list of bookmarks by opening the Go To Bookmarks dialog and disabling the checkbox to Show Underlays. When the checkbox is disabled, Underlays will not show in the Go To Bookmarks dialog, and they will not appear in the Options menu.

Adding Nodes to a Composition

You can add nodes to the Node Editor in a variety of ways, depending on the type of node you're adding, and how much guidance you need to find what you're looking for. Additionally, the way you add nodes to a composition may also be dictated by how you need to attach that node to the current node tree.

Make Sure You're Adding Compatible Nodes

It's a good rule of thumb to make sure that whenever you're adding or inserting new nodes to the node tree, that you're adding nodes that are compatible with the nodes you're trying to attach to. For example, you'll have no problem inserting a Blur, Color, Filter, Paint, or Position node after almost any 2D operation. However, if you try to add a Merge3D node after a Glow node, it won't automatically connect, because those two nodes cannot be connected directly.

Adding, Inserting, and Replacing Nodes Using the Toolbar

The Fusion toolbar, located above the Node Editor, displays a selection of frequently-used nodes, displayed as buttons with distinct icons. These buttons make it fast to add Merge, Background, Paint, Mask, Transform, and many other commonly used nodes with the click of a button, or the drag of your pointer.



The Fusion page toolbar.

TIP: If you don't know which node a particular icon corresponds to, just hover your pointer over any toolbar button and a tooltip will display the full name of that tool.

Methods of adding nodes by clicking toolbar buttons:

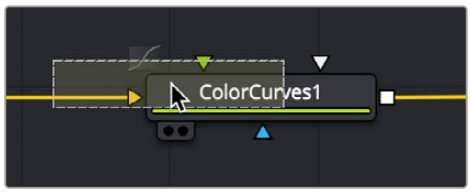
- **To add a node after a selected node:** Select a node in the Node Editor and then click a toolbar button.
- **To add a disconnected node to the Node Editor:** Deselect all nodes in the Node Editor and then click a toolbar button.

Methods of adding nodes by dragging toolbar buttons:

- **To insert a new node into the node tree:** Drag a toolbar button into the Node Editor and onto the connection line between any two compatible nodes. When the connection highlights as the node is over it, drop the node and it'll be inserted.
- **To create a disconnected node:** Drag a toolbar button into an empty part of the Node Editor. Dragging a toolbar button into the Inspector also creates a disconnected node.
- **To insert a new node after a node loaded into a viewer:** Drag a toolbar button onto a viewer to insert a new node after whichever node is viewed, regardless of whether any nodes are selected.

To replace a node in the Node Editor with a node from the toolbar:

1 Drag a button from the toolbar so that it's directly over the node in the Node Editor that you want replaced. When the node underneath is highlighted, drop the node.



Dragging a node from the toolbar to replace an existing tool.

2 Click OK in the dialog to confirm the replacement.

TIP: When you replace one node with another, any settings that are identical between the two nodes are copied into the new node. For example, replacing a Transform node with a Merge will copy the existing center and angle values from the Transform to the Merge.

Adding Nodes Quickly Using the Select Tool Window

The next fastest way of adding or inserting nodes to the Node Editor is using the Select Tool window, which lets you search for any node available to Fusion by typing a few characters. Once you learn this method, it'll probably become one of your most frequently-used ways of adding nodes.

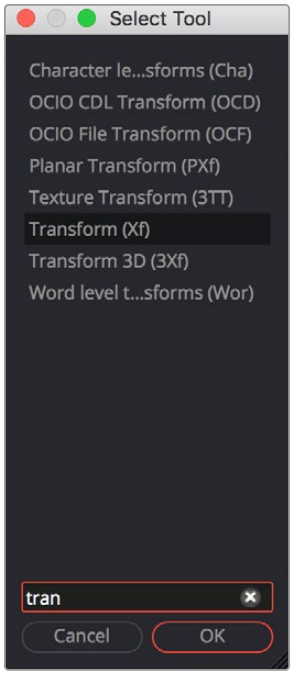
To use the Select Tool window to add nodes:

1 Do one of the following to determine if you want to insert a node or create a disconnected node:
   - If you want to insert a node, select a node that's compatible with the one you'll be creating, and the new node will be inserted after it.
   - If you want to create a disconnected node, deselect all nodes.

2 Press Shift-Spacebar to open the Select Tool dialog.

3 When the window appears, type characters corresponding to the name of the node you're looking for. A list automatically appears with likely candidates, and you can use the Up and Down arrow keys to select the correct node (if it's not selected already). If you don't know the name of the exact tool you're looking for, you can also type in common keywords and categories (i.e., keyer, materials, etc.) to get a selection of tools based on your input.

4 When you've selected the correct node, press the Return key (or click OK), and that node will be either inserted or added.



The Select Tool dialog lets you find any node quickly if you know its name.

TIP: Whenever you use the Select Tool window, the text you entered is remembered the next time you open it, so if you want to add another node of the same kind—for example, if you want to add two Blur nodes in a row—you can just press Shift-Spacebar and then press Return to add the second Blur node.

Adding Nodes from the Effects Library

While the toolbar shows many of the most common nodes you'll be using in any composition, the Effects Library contains every single tool available in Fusion, organized by category, with each node ready to be quickly added to the Node Editor. If you need more guidance to find the node you're looking for, or if you just want to browse around and see what's available, the Effects Library is the perfect place to start.

To open the Effects Library:

- Click the Effects Library button in the UI toolbar at the top of the Fusion window.

The Effects Library appears at the upper-left corner of the Fusion window, and consists of two panels. A category list at the left shows all categories of nodes and presets that are available, and a list at the right shows the full contents of each selected category.



The Tools bin of the Effects Library exposing 3D nodes.

By default, the category list shows the primary sets of effects: Tools, Open FX, Templates, and LUTs; with disclosure controls to the left that hierarchically show all subcategories within each category.

The categories are:

- **Tools:** Tools consist of all the effects nodes that you use to build compositions, organized by categories such as 3D, Blur, Filter, Mask, Particles, and so on.
- **Open FX:** All Resolve FX and any installed third party Open FX plugins will appear here.
- **Templates:** When using the Fusion page in DaVinci Resolve, templates consist of presets, macros, and utilities that have been created to get you started quickly. For example, Backgrounds consists of a variety of customizable generators that have been created using a combination of Fusion tools. Lens flares presents a wide variety of multi-element lens flares that you can add to any composition. Particles has a selection of pre-made particle systems that you can customize for your own use. Shaders has a variety of materials that you can use as texture maps for 3D text and geometry that you create in Fusion. And there are many, many other categories' worth of useful presets and macros that you can learn from and use in your own projects.
- **LUTs:** An assortment of pre-installed Look Up Tables for gamma, gamut, and colorspace conversions can be found here. Inserting a LUT into the Node Editor creates a new FileLUT (FLUT) node with the selected LUT file preloaded.

For more information on adding your own LUTs to this list, see Chapter 150, “Using LUTs,” in the *DaVinci Resolve 20 Reference Manual*.

Adding, Inserting, and Replacing Tools Using the Effects Library

Adding nodes to the Node Editor from the Tools category of the Effects Library is very similar to adding nodes from the toolbar.

Methods of adding nodes by clicking in the Effects Library:

- **To add a node after a selected node:** Select a node in the Node Editor and then click a node in the browser of the Effects Library.
- **To add a disconnected node to the Node Editor:** Deselect all nodes in the Node Editor and then click a node in the browser of the Effects Library.

Methods of adding nodes by dragging from the Effects Library:

- **To insert a new node into the node tree:** Drag a node from the browser of the Effects Library into the Node Editor and onto the connection line between any two compatible nodes. When the connection highlights as the node is over it, drop the node and it’ll be inserted.
- **To create a disconnected node:** Drag a node from the browser of the Effects Library into an empty part of the Node Editor. Dragging a toolbar button into the Inspector also creates a disconnected node.
- **To insert a new node after a node loaded into a viewer:** Drag a node from the browser of the Effects Library onto a viewer to insert a new node after whichever node is viewed, regardless of whether any nodes are selected.

To replace a node in the Node Editor with a node from the Effects Library:

1. Drag a node from the browser of the Effects Library so it’s directly over the node in the Node Editor that you want replaced. When that node is highlighted, drop it.
2. Click OK in the dialog to confirm the replacement.

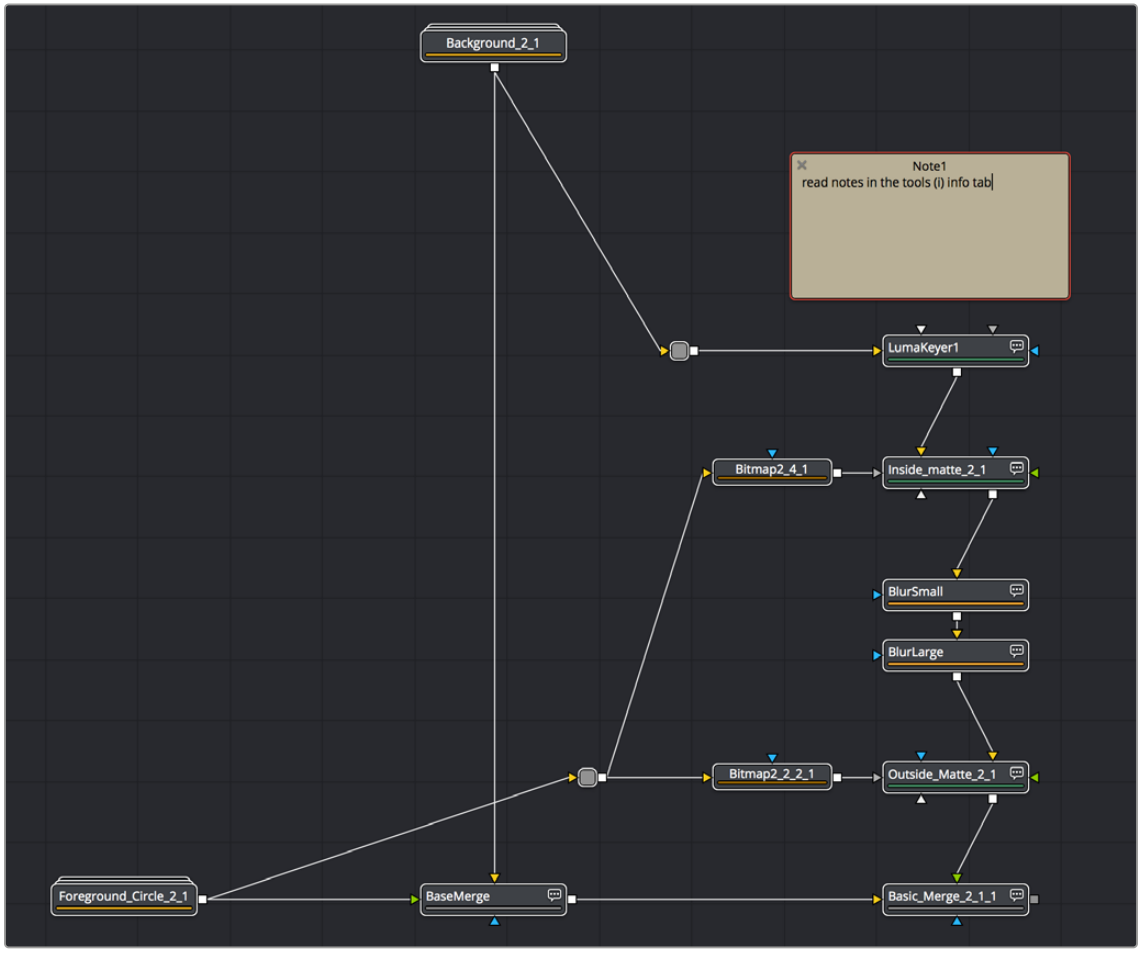
Adding, Inserting, and Replacing Templates Using the Effects Library

Adding items from DaVinci Resolve’s Fusion page Templates category is often a bit different. Sometimes, as when adding a Lens Flare, a single node can be added or inserted into the Node Editor. When this is the case, adding nodes works the same as when adding from the Tools category.



Adding a Lens Flare effect.

Other times, such as when adding an item from the “How to” category, dragging a single item from the Node Editor results in a whole node tree being added to the Node Editor. Fortunately, all nodes of the incoming node tree are automatically selected when you do this, so it’s easy to drag the entire node tree to another location in the Node Editor where there’s more room. When this happens, the nodes of the incoming effect are exposed so you can reconnect and reconfigure it as necessary to integrate the effect with the rest of your composition.



Adding a LightWrap effect from the “How to” bin of the Templates category of the Effects Library.

Adding, Inserting, and Replacing Nodes Using the Contextual Menu

Another way of adding, inserting, and replacing nodes is to use the Node Editor’s contextual menu, which has dedicated submenus that let you create any kind of node available in Fusion. This can be a convenient when the pointer is already in the Node Editor selecting, moving, or connecting nodes.

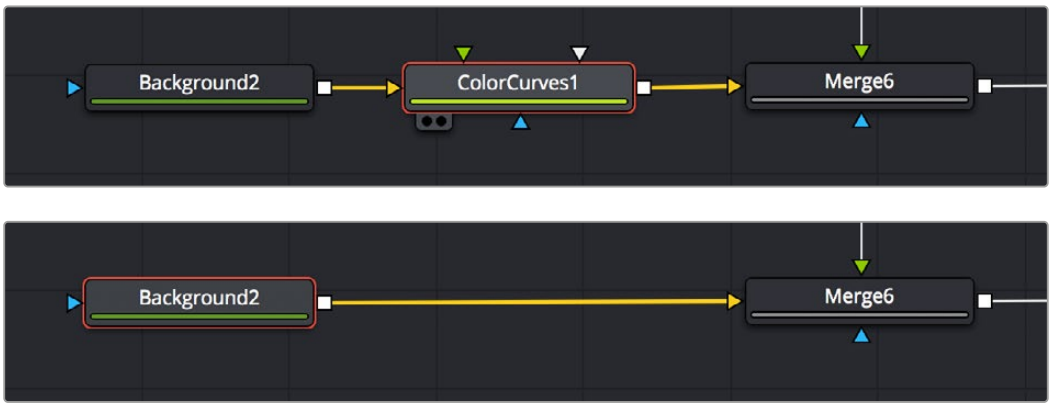
Methods of adding nodes using the contextual menu:

- **To add a node:** Right-click in an empty area of the Node Editor, and choose a node from the Add Tool submenu.
- **To insert a node:** Right-click a node in the Node Editor, and choose a node from the Insert Tool submenu.
- **To replace a node:** Right-click a node in the Node Editor, and choose a node from the Replace Tool submenu.

TIP: When you replace one node with another, any settings that are identical between the two nodes are copied into the new node. For example, replacing a Transform node with a Merge will copy the existing center and angle values from the Transform to the Merge.

Deleting Nodes

To delete one or more selected nodes, press Delete (macOS) or Backspace (Windows), or right-click one or more selected nodes and choose Delete from the contextual menu. The node is removed from the Node Editor, and whichever nodes are connected to its primary input and output are now connected together. Nodes connected to other inputs (such as mask inputs) become disconnected.



Before deleting a node from a node tree (top), and after upstream and downstream nodes have automatically reconnected (bottom).

Disconnected Nodes

It's perfectly fine to have disconnected nodes, or even entire disconnected branches of a node tree, in the Node Editor alongside the rest of a composition. All disconnected nodes are simply ignored while being saved for possible future use. This can be useful when you're saving nodes that you've customized but later decided you don't need. It's also useful for saving branches of trees that you've since exported to be self-contained media that's re-imported to take the place of the original effect, but you want to save the original nodes just in case you need to redo your work.

Selecting and Deselecting Nodes

In order to work with nodes in the Node Editor in any way, or modify node parameters in the Inspector, you first need to learn to select the node or nodes you want to work with.

Selecting Nodes

Selecting nodes is one of the most fundamental things you can do to move nodes or target them for different operations. There are a variety of methods you can use.

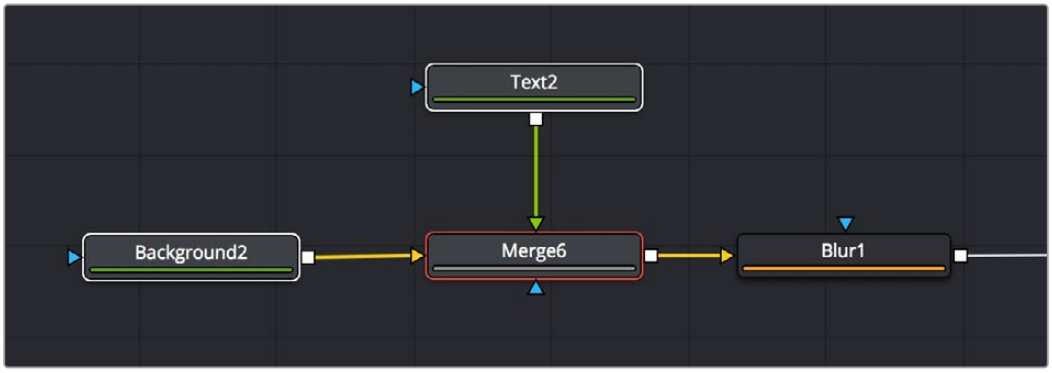
Methods of selecting nodes:

- **To select a single node:** Click any node in the Node Editor.
- **To select multiple nodes one at a time:** Command-click each node you want to select.
- **To select a whole region of nodes:** Drag a bounding box around all nodes you want to select.
- **To select all upstream or downstream nodes:** Right-click a node and choose Select > Upstream Nodes/Downstream Nodes from the contextual menu.
- **To select all nodes in the Node Editor:** Press Command-A.
- **To select a node from the Keyframe Editor:** Click any layer in the Keyframe Editor to select the corresponding node in the Node Editor.

The Active Node

When you select a single node using any of the methods described above, the selected node is known as the active node, and is highlighted orange to indicate that its parameters are currently editable in the Inspector (if the Inspector is open). This also indicates that node will be targeted for specific operations (such as inserting new nodes).

While multiple nodes can be selected, only one node will be the active node. To indicate the difference, the active node remains highlighted with orange, while all other selected nodes are highlighted with white. Unselected nodes have simple black outlines.



The active node is highlighted orange, while other selected nodes are highlighted white.

To set the active node when there are multiple selected nodes:

- Option-click one of the selected nodes in the Node Editor to make that one the active node.
- Open the Inspector (if necessary), and click a node's header bar to make it the active node.

Deselecting Nodes

Deselecting nodes, when necessary, works pretty much as you would expect.

Methods of deselecting nodes:

- Click once in the background of the Node Editor to deselect all nodes.
- Press Command-Shift-A to deselect all nodes.
- Command-click to deselect multiple nodes one at a time.
- Command-drag a bounding box to deselect a group of selected nodes at one time.

Loading Nodes into Viewers

Once you've started building a composition, the next thing you need to learn is how to view specific nodes that you want to work on. This is important because the combination of which node is being viewed and which node is currently selected (these aren't always the same node) often determines which onscreen controls are available and how they appear.

In the following example, you're set up to rotoscope an image using a Polygon node that's attached to the garbage mask input of a MatteControl node which is inserting the mask as an alpha channel.



A node tree for doing a simple rotoscoping job.

As seen in the screenshot above, you'll want to load the upstream MediaIn or Loader node into a viewer while the Polygon node is selected for editing in order to see the full image you're rotoscoping while keeping the Polygon node's spline visible.

Viewed Nodes When You First Open Fusion

When you first open the Fusion page in DaVinci Resolve, the output of the current empty composition (the MediaOut1 node) is usually showing in viewer 2. If you're in Dual-viewer mode, viewer 1 remains empty until you assign a node to one of them.

When you first open Fusion Studio with an empty comp, both viewers remain empty even after reading in media using a Loader node. The viewers only display content when you assign a node to one of them.

There are several different ways to display a node in a viewer. Which ones you use depends on how you like to work.

Node View Indicators

The View indicators are displayed under each node, and serve two purposes. First, they're a clickable interface for displaying a node's output in one of the viewers. Second, they're an indication of which nodes in the Node Editor are being viewed. By default, there are two round indicators, representing the two viewers. The left and right indicators correspond to the left and right viewers, regardless of whether both viewers are visible or just one.



A viewer indicator enabled for the right viewer and disabled for the left viewer.

To load a node into a viewer using the Node View indicators:

- Clicking an indicator turns it white to show that node is currently loaded in the corresponding viewer. Clicking it again turns the indicator black and removes it from the viewer. Nodes only display View indicators if they're currently being viewed. If you want to view indicators, hovering the pointer over the node makes the indicators visible and available for clicking.
- You can also use keyboard shortcuts to toggle each View indicator. The default two viewers are assigned numeric keyboard shortcuts 1 and 2. Pressing the corresponding number once displays the selected node in the appropriate display view, while pressing it again clears that display.

For complex compositions, you may need to open additional viewers. For example, one viewer may be used to display the end result of the final comp, while another viewer displays the source, a third viewer displays a mask, and a fourth viewer might be a broadcast monitor connected via a Blackmagic DeckLink card or other display hardware. When you have more than two viewers, additional View indicators are added and each one is assigned a consecutive number between 3 and 9.

The more viewers you add, the more you may need help remembering which viewer is represented by which View indicator. Positioning the pointer over the View indicator in question will display a tooltip with the name of the viewer it represents.

Drag and Drop Nodes into a Viewer

If the View indicators are too small of a target for you to click on reliably and you are not keyboard oriented, another way to load a node into a viewer is to drag and drop it onto the viewer you want to load it into. This offers a quick explicit way to assign a node to a viewer, especially for pen and tablet users. Please note that as you drag, the node will appear to move at first, but it'll snap back into its original location once the pointer leaves the Node Editor.

Using the Contextual Menu

You can also right-click a node, and then choose View On > Left or Right to display the node on the appropriate viewer.

Clearing Viewers

Whenever you load a node into a viewer, you prompt that node, all upstream nodes, and other related nodes to be rendered. If you load nodes into both viewers, this is doubly true. If you want to prevent your computer from processing views that aren't currently necessary, you can clear each viewer.

Methods of clearing viewers:

- Press 1 or 2 to empty the left or right viewers if they're filled.
- Press ` (the Accent key) to empty both viewers.

Create/Play Preview

You can right-click a node, and choose an option from the Create/Preview Play On submenu of the contextual menu to render and play a preview of any node's output on one of the available viewers. The Render Settings dialog is displayed, and after accepting the settings, the tool will be rendered and the resulting frames stored in RAM for fast playback on that view.

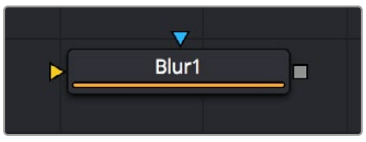
> **TIP:** Hold the Shift key when selecting the viewer from the menu to bypass the Render dialog and to start creating the preview immediately using the default settings or the last settings used to create a preview.

Connecting and Disconnecting Nodes

Once you've started to add nodes to your composition, you need to connect them to perform their intended operations.

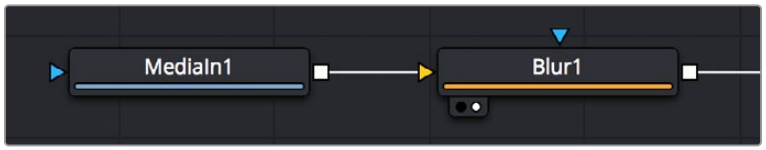
Node Basics

Each node displays small colored knots around the edges. One or more arrows represent inputs, and the square represent the tool's processed output, of which there is always only one. Outputs are white if they're connected properly, gray if they're disconnected, or red to let you know that something's wrong and the node cannot process properly.



A Blur node with a Foreground Input, Mask Input, and Output.

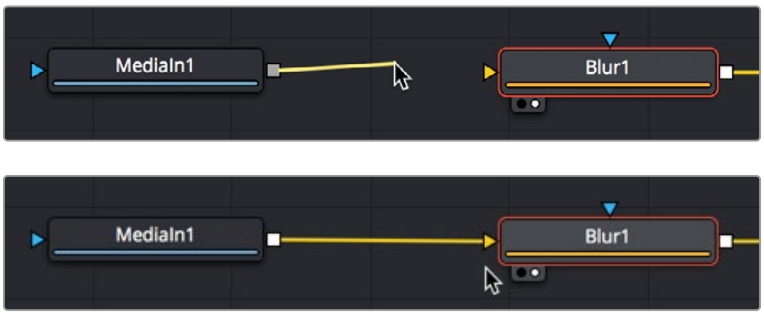
Each node takes as its input the output of the node before it. By connecting a MediaIn node's output to a Blur node, you move image data from the MediaIn node to the Blur node, which does something to process the image before the Blur node's output is in turn passed to the next node in the tree.



Two nodes connected together.

How to Connect Nodes

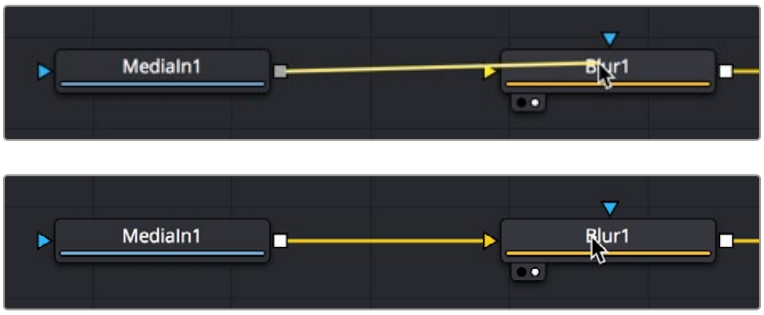
To manually connect one node to another, click on one node's output and drag a connection line out to drop on another node's input. The order in which you drag node connections is not important; you can just as easily drag a connection from one node's input to another node's output and get the same results.



Before (top), and after (bottom) dragging a connection line and dropping it to connect two nodes.

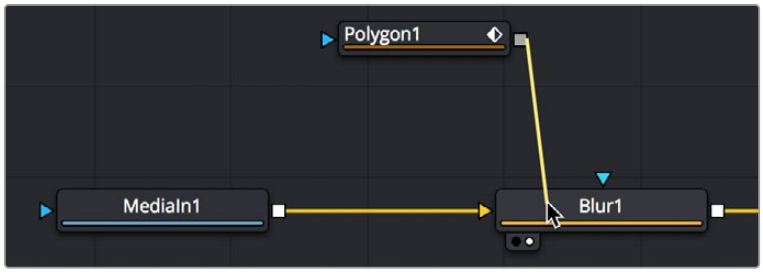
Dropping Connections on Top of Nodes

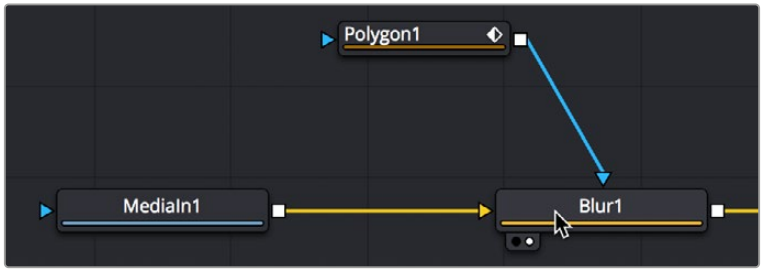
To make your life a bit easier, you can also drag a connection line and drop it directly on top of the tile of a node to automatically connect to the default input of that node, which is usually labeled "background" or "input." In the following example, a connection is dragged from the output of a MediaIn node and dropped onto the tile of a Blur1 node, and the background input is connected first.



Before (top), and after (bottom) dragging a connection line and dropping it on top of a node.

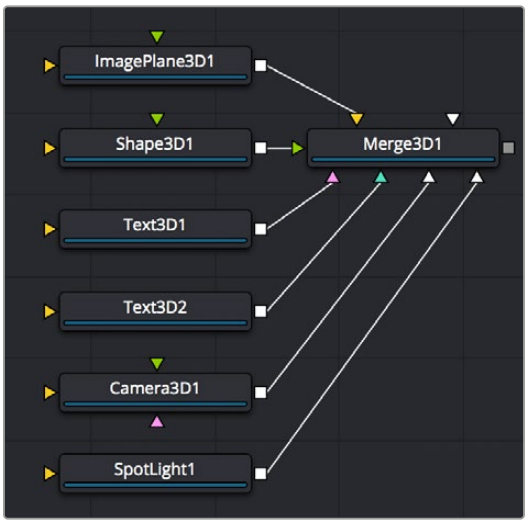
If you drop a connection on top of a node that already has the background input connected, then the second most important connection will be attached, which for multi-input nodes is the foreground input, and for other single-use nodes may be the Effects Mask input.





Before (previous page), and after (above) dragging a connection line and dropping it on top of a node that has the background input already connected.

Some multi-input nodes are capable of adding inputs to accommodate many connections, such as the Merge3D node. These nodes simply add another input whenever you drop a connection onto them.



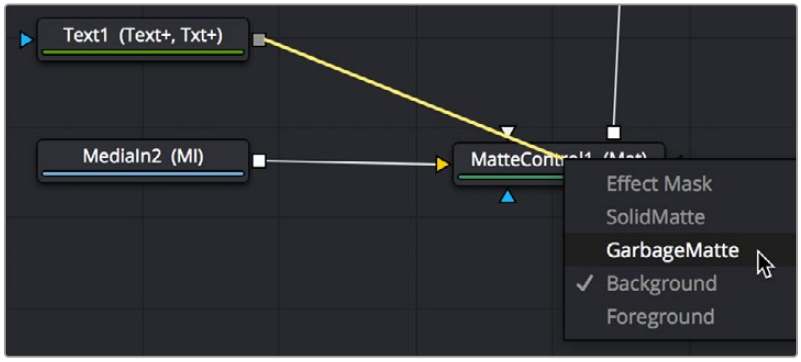
After dragging a connection line and dropping it on top of a Merge3D node.

Attaching Connections to Specific Inputs

If you want to make sure you don't attach a connection to the default input of a node, then you need to drop it right on top of the specific node input you want to attach it to. If you can see the input's label in the tooltip bar, then you know you're correctly positioned to make a good connection.

However, there's an alternate method of connecting nodes together in instances where there are several inputs to choose from and you want to make sure you're choosing the correct one. Hold down the Option key while dragging a connection from one node's output and dropping it onto the body of another node. This opens a pop-up menu from which you can choose the specific input you want to connect to, by name. Please note that this menu only appears after you've dropped the connection on the node and released your pointing device's button.





Option-dragging a node connection to drop onto another node exposes a node input menu.

Automatically and Manually Attaching Mask Nodes

Mask nodes, such as the Polygon, B-Spline, Ellipse, or Rectangle, have a different automatic behavior when you connect them to other nodes. If you drag a connection from a Mask node onto the body of another node, it will automatically connect itself to the default mask input, which is usually the effect mask input. The assumption is that you're using the mask to limit the node's effect somehow. However, this isn't always the case, so you'll need to be careful of this behavior to make sure you're attaching your mask to the input that will actually create the effect you need.



Before (left) and after (right) dragging a connection from a Mask node and dropping it on top of a MatteControl node.

Identifying Node Inputs

While you are still figuring out all the nodes and their inputs, hovering the pointer over any knot will display a node tip with the knot's name.

TIP: Rather than remembering the different knot types, press the right mouse button, hold Option, and drag from the output of a node to the center of another tool. When you release the mouse, a tooltip will appear allowing you to select the knot you want to connect to.

Node Order Matters

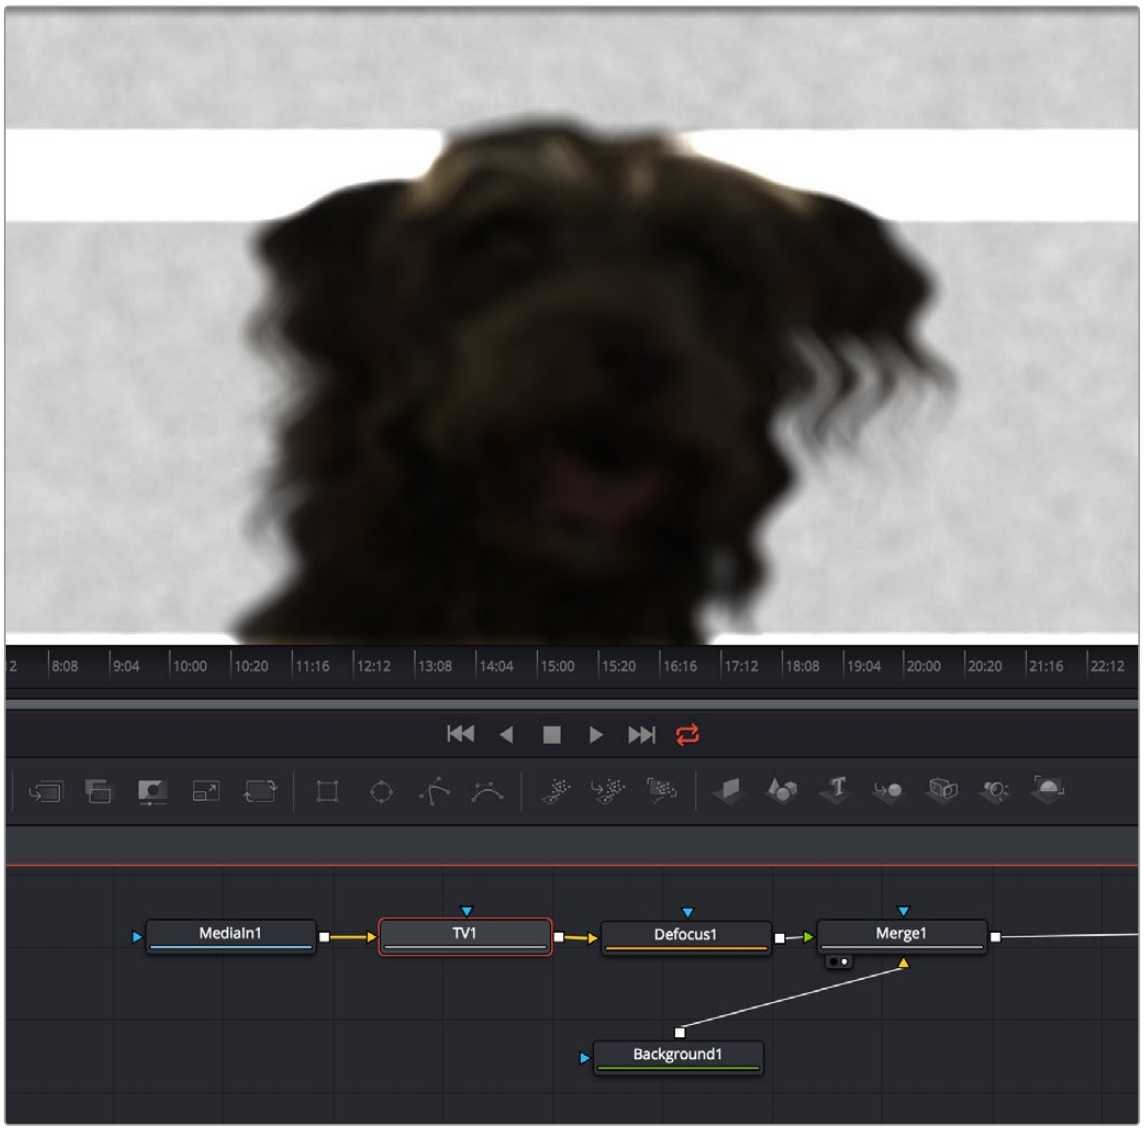
The order in which nodes are attached defines the order in which each image-processing operation is applied to the image.

In the following example, a MediaIn node adds a clip to the composition, while a Defocus node blurs the image, and then a TV node adds scanlines and vertical distortion. Those effect nodes are then connected to the MediaOut node in the Fusion page in DaVinci Resolve or a Saver node in Fusion Studio.



Adding a Defocus effect first, then the TV node second.

As you can see above, connecting the Defocus node first, followed by the TV node, means that while the initial image is softened, the TV effect is sharp. However, if you reverse the order of these two nodes, then the TV effect distorts the image, but the Defocus node now blurs the overall result, so that the TV effect is just as soft as the image it's applied to. The explicit order of operations you apply makes a big difference.



Adding a TV effect first, and a Defocus second.

As you can see, the node tree that comprises each composition is a schematic of operations with tremendous flexibility. Additionally, the node tree structure facilitates compositing by giving you the ability to direct each node's output into separate branches, which can be independently processed and later recombined in many different ways, to create increasingly complex composites while eliminating the need to precompose, nest, or otherwise compound layers together, which would impair the legibility of your composition.

In the following example, several graphics layers are individually transformed and combined with a series of Merge nodes. The result of the last Merge node is then transformed, allowing you to move the entire collection of previous layers around at once. Because each of these operations is clearly represented via the node tree, it's easy to see everything that's happening, and why.

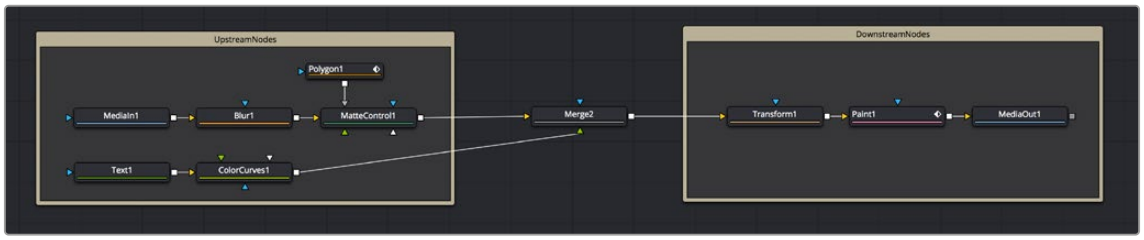


The output of five Text nodes being combined using Merge nodes is modified by a single Transform node.

Upstream and Downstream Nodes

Since nodes can be positioned anywhere in the Node Editor, and added in any direction, nodes are referred to as being upstream and downstream of one another. Once you select a node, all other nodes that directly or indirectly connect to its input are considered to be upstream. Any other nodes that are directly or indirectly connected to the output are said to be downstream.

This is an important distinction to make because, unlike layer-based systems, the visual positioning of nodes in your node tree has no bearing on the order of operations in that composition. The only thing that matters is whether nodes are upstream or downstream of each other.



Tools upstream (left) and downstream (right) of the Merge node.

TIP: To help you stay organized, there are Select > Upstream/Downstream commands in the Node Editor contextual menu for selecting all upstream or downstream nodes to move them, group them, or perform other organizational tasks.

Disconnecting and Reconnecting Nodes

Node trees are a continuous work in progress, requiring constant revision and rearrangement as you discover new details that need to be finessed, or things that you can do better once the overall composition has taken shape. To facilitate quick changes, each connection between two nodes is divided into two halves: the output half (connected to the upstream node's output) and the input half (connected to the downstream node's input). This can only be seen when you hover the pointer over a connection. The half your pointer is over is highlighted in blue.



The two halves of a connection line that are revealed when you hover your pointer over it.

By clicking and/or dragging these two halves, it's possible to quickly disconnect, reconnect, and overwrite node connections, which is essential to rearranging your node tree quickly and efficiently.

To disconnect two nodes, do one of the following:

- Click once on the input half of the connection between two nodes.
- Click on the input arrow to which a connection is attached, then drag to pull the connection away from the tool and drop it anywhere in an empty area of the Node Editor.

To overwrite a current connection:

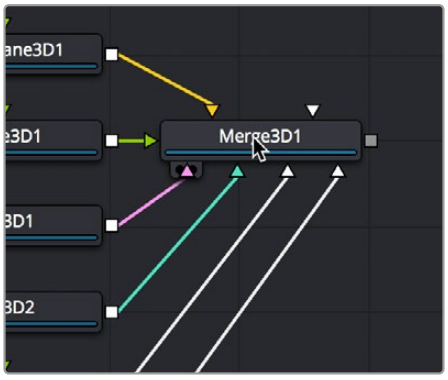
- Drag the output or input half of a connection, and drop it directly onto another node's input or output. This simultaneously disconnects the previous connection and connects the one you're dragging.

To reconnect a connection from one node to another:

- Drag the output or input half of a connection to disconnect it from one node, and drop the connection directly on another node's input or output.

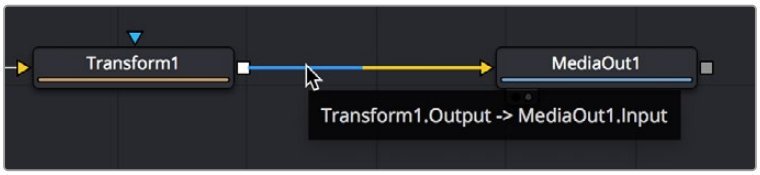
Tracing Connections Through the Node Tree

Positioning the pointer over a node causes the connections attached to that node to become highlighted, which makes it easier to see which nodes it's attached to. Additionally, highlighted connections display the color of the input they're connected to, which makes it easy to see if they're connected to a foreground, a background, or a particular kind of mask.



Hovering the pointer over a node highlights the color of all connections, telling you what kinds of inputs are connected.

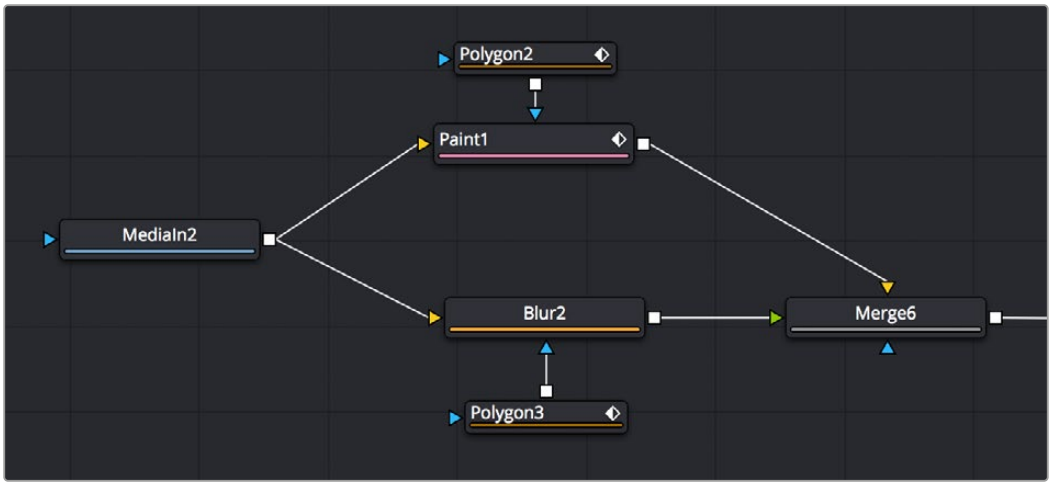
Additionally, positioning the pointer over a connection causes a tooltip to appear that displays the output and input that connection is attached to.



Hovering the pointer over a node highlights the connection between it and other nodes.

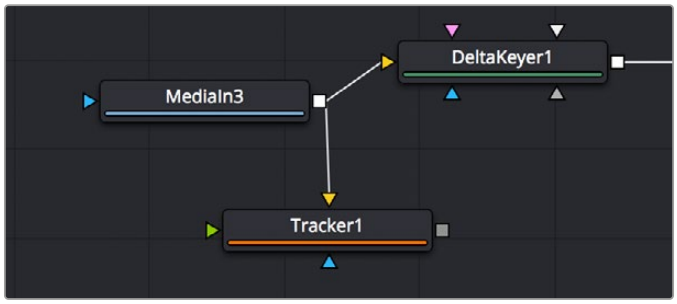
Branching

A node's input can only have one connection attached to it. However, a tool's output can be connected to inputs on as many nodes as you require. Splitting a node's output to inputs on multiple nodes is called branching. There are innumerable reasons why you might want to branch a node's output. A simple example is to process an image in several different ways before recombining these results later on in the node tree.



A MediaIn node branched to two node operations and then recombined using a Merge node.

Alternatively, it lets you use one image in several different ways—for example, feeding the RGB to one branch for keying and compositing, while feeding the A channel to the Effects Mask input of another node to limit its effect, or feeding RGB to a tracker to extract motion information.



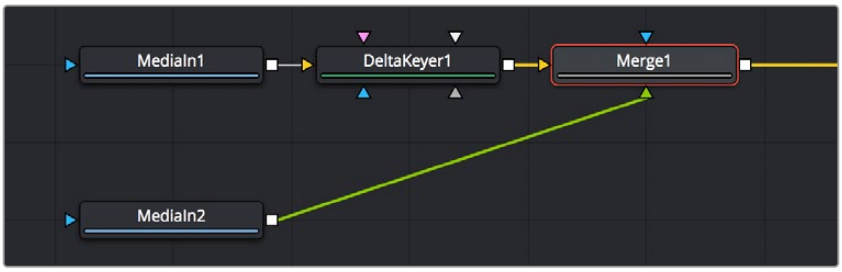
A MediaIn node branched to two different kinds of inputs, used separately.

Connecting Merge Nodes

The Merge node is the primary tool available for compositing images together. Each Merge node is capable of combining two inputs to create a third, using standard compositing methods and composite modes.

For more extensive information about the Merge node, *see Chapter 45, "I/O Nodes,"* in the Fusion Reference Manual.

For this chapter, all you need to know is that if you attach a background image to the Background input (such as a landscape), and a foreground image with an alpha channel to the Foreground input (such as a graphic with an alpha channel), the Merge node will combine them into a single image for further compositing.



Two MediaIn nodes and a DeltaKeyer node attached to a Merge node, creating a composite.

Each Merge node has three inputs:

- **Background (orange):** The default input. Whichever image is connected to this input defines the output resolution of the Merge node.
- **Foreground (green):** The secondary input, meant for whichever image you want to be "on top."
- **Effect Mask (blue):** An optional input you can use to attach a mask or matte with which to limit the effect of the Merge node.

It's important to make sure you're attaching the correct nodes to the correct inputs to ensure you're getting the result you want, and it's important to keep these inputs in mind when you connect to a Merge node. Of course, you can always drag a connection to a specific input to make sure you're connecting things the way you need. However, if you're in a hurry and you simply drag connections right on top of a Merge node:

- The first connection will be made to the background input.
- The second connection will be made to the foreground input.
- The third connection will be made to the effect mask input.

TIP: When you add a Merge node after a selected node by clicking the Merge button on the toolbar, by clicking on the Merge icon in the Effects Library, or by right-clicking a node in the node tree and choosing Insert Tool > Composite > Merge from the contextual menu, the new Merge node is always added with the background connected to the upstream node coming before it.

Automatically Creating a Merge Node When Adding Nodes

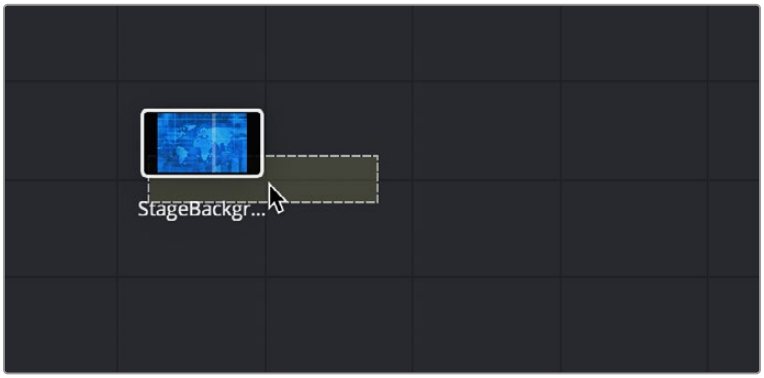
There's a nice shortcut for connecting Merge nodes if you want to connect the incoming clip immediately to your node tree as the top layer of a composite, and that's to drag a clip from an Operating System window or a Generator from the Effects Library right on top of any connection line.

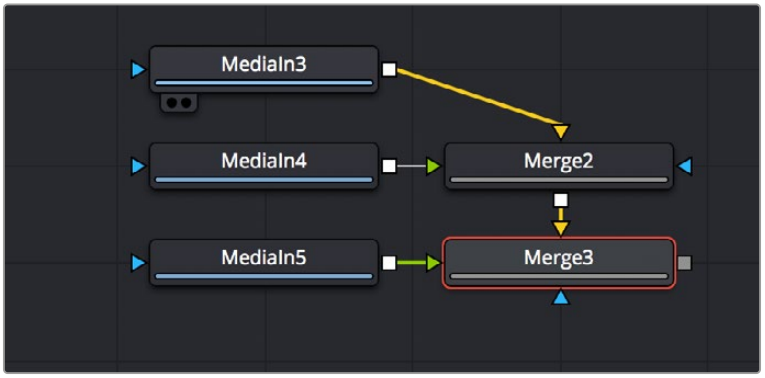
When you drop the resulting node, this automatically creates a Merge node, the background input of which is connected to the next node to the left of the connection you dropped the clip onto, and the foreground input of which is connected to the new node that represents the clip or Generator you've just added.



Dragging a node from the Media Pool onto a connection (top), and dropping it to create a Merge node composite (bottom).

Additionally, If you drag two or more nodes from an OS window into the Node Editor at the same time, Merge nodes will be automatically created to connect them all, making this a fast way to initially build a composite.





By dragging three nodes from an OS window to the Node Editor (top), Merge nodes are automatically created to connect them all (bottom).

Automatically Creating a Merge Node by Connecting Two Outputs

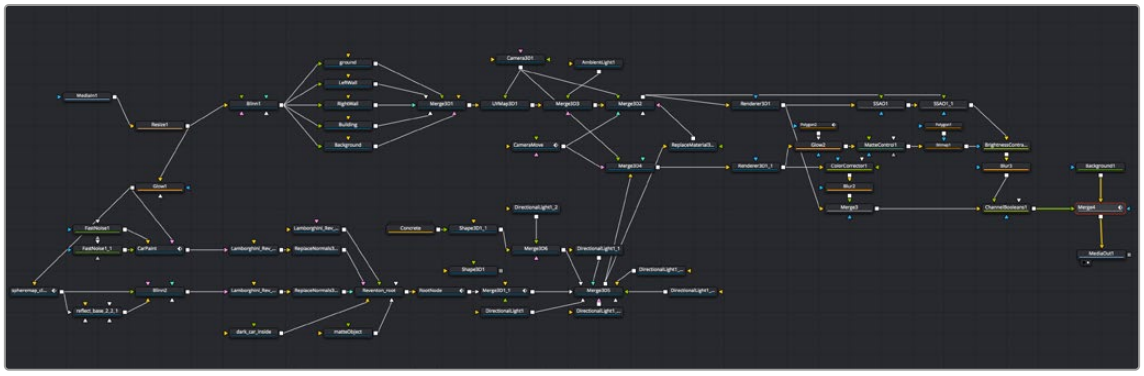
Here's an endlessly useful shortcut for when you have a disconnected node that you want to composite over another node. Drag a connection from the output of the node you want to be the foreground layer, and drop it on top of the output of the node you want to be the background layer, and a Merge node will be automatically created to build that composite.



Dragging a connection from a disconnected node to another node's output (top), and dropping it to create a Merge node composite (bottom).

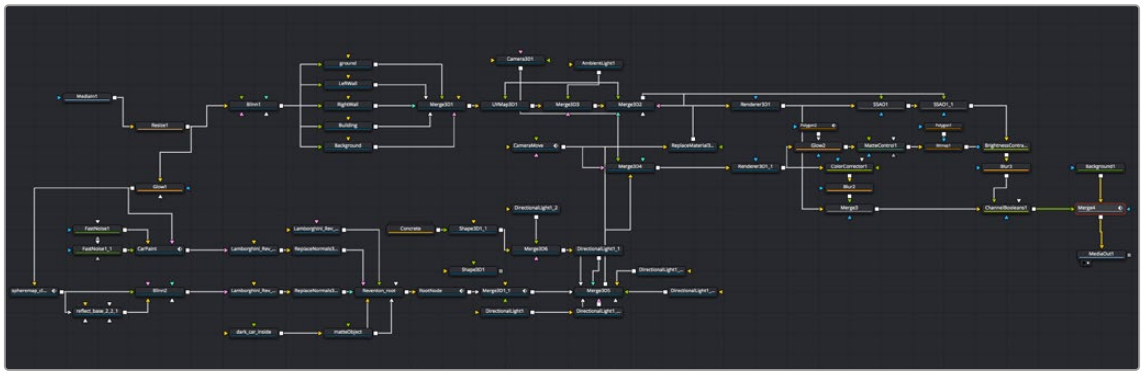
Connection Options and Routers

By default, the Node Editor uses linear connections that are drawn straight between any two connected nodes. While efficient, this sometimes causes connection lines to overlap nodes, which some people feel interferes with the view of the Node Editor.



Linear connections between nodes.

If you like, you can change how connections are drawn by enabling *orthogonal* connections, which automatically draws lines with right angles to avoid having connections overlap nodes.



Optional orthogonal connections between nodes.

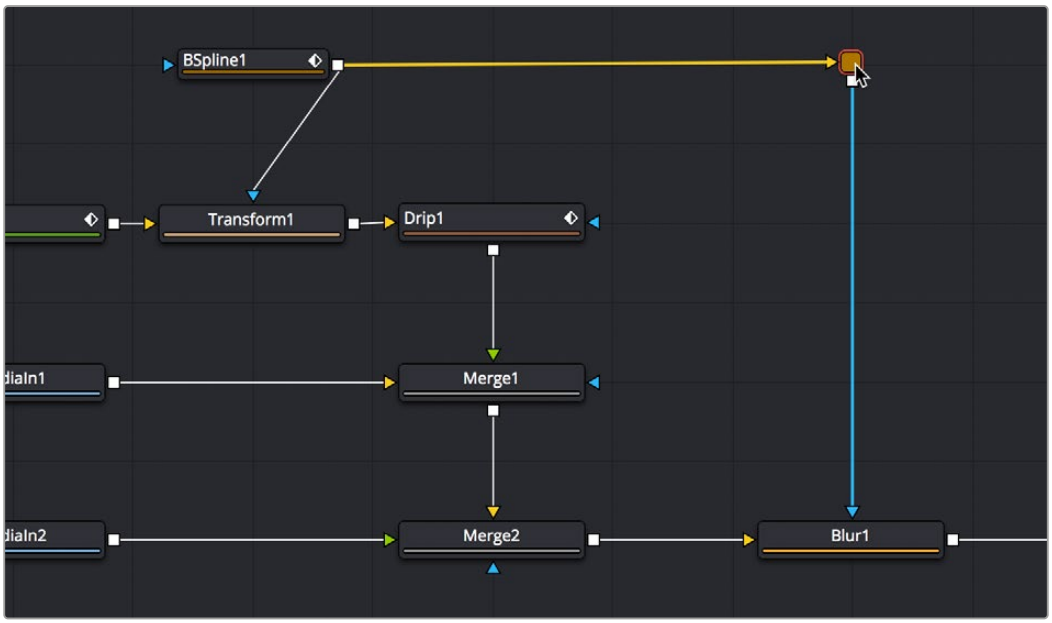
Functionally, there's no difference to your composition; this only affects how your node tree appears.

To change how connections are drawn in the Node Editor:

- Right-click the Node Editor background and choose one of the following from the contextual menu.
  - Options > Direct Pipes
  - Options > Orthogonal Pipes

Using Routers to Reshape and Branch Connections

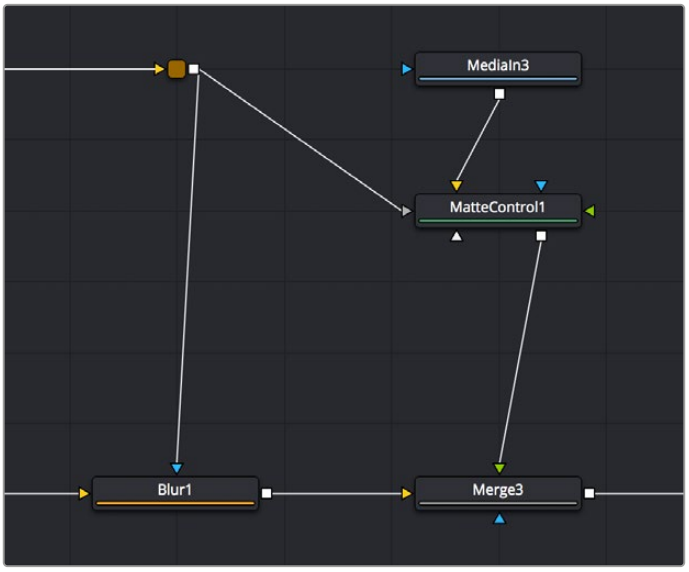
If you want to force a particular connection to be drawn at an angle to keep your node tree tidy, you can add a router to either linear or orthogonal connections to force an angle so it will be drawn however you like.



A router added to force a connection to be drawn at an angle.

Routers are tiny nodes with a single input and an output, but with no parameters except for a comments field (available in the Inspector), which you can use to add notes about what's happening in that part of the composition.

Even more usefully, you can branch a router's output to multiple nodes, which makes routers even more useful for keeping node trees neat in situations where you want to branch the output of a node in one part of your node tree to other nodes that are all the way on the opposite end of that same node tree.



A router branching its output to multiple nodes.

Methods of using routers:

- **To add a router to a connection:** Option-click anywhere on a connection.
- **To move a router:** Drag the router to a new location, and the connection will reshape itself as necessary.
- **To branch a router's output:** Drag a connection from the router output to the input of another node. You can branch a router's output as many times as you need to.
- **To remove a router:** Select any router and press the Delete key, or right-click a router and choose Delete from the contextual menu.

Swapping Node Inputs

For multiple-input nodes such as the Merge, Merge 3D, and Dissolve nodes, there's a quick method of swapping the Primary and Secondary inputs, such as the foreground and background inputs of a Merge tool, when you find you've accidentally connected them in the wrong order. If a node has more than two of its inputs connected, only the foreground and background inputs will be swapped.

To swap the primary inputs of a multi-input node, do one of the following:

- Select a node and press Command-T to reverse its inputs.
- Right-click a node and choose Swap Inputs from the contextual menu.



Before swapping node inputs (top), and after swapping node inputs (bottom), the connections don't move but the colors change.

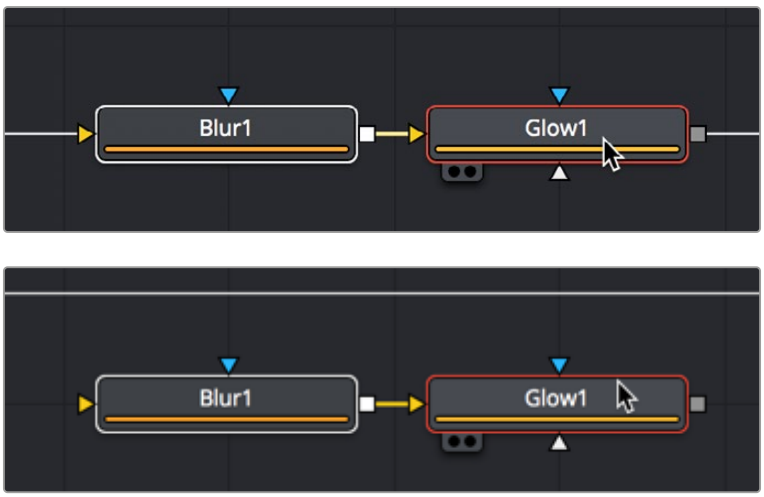
Inputs can move freely around the node, so swapping two inputs doesn't move the connection lines; instead, the inputs change color to indicate you've reversed the background (orange) and foreground (green) connections.

Extracting and Inserting Nodes

When building a composition, you'll often find that you need to rearrange nodes that you've already added, in order to connect them in different ways to obtain a better result. Happily, this is easy to do by extracting one or more nodes from one part of a node tree and inserting them at another part of the node tree.

To extract one or more nodes from their position in the node tree:

- **To extract a single node:** Hold down the Shift key, drag a node from the node tree up or down to disconnect it, and then drop the node before releasing the Shift key. That node is now detached, and the output of the next upstream node is automatically connected to the input of the next downstream node to fill the gap in the node tree.
- **To extract multiple nodes:** Select the nodes you want to extract, hold down the Shift key, drag one of the selected nodes up or down to disconnect them, and then drop the node before releasing the Shift key. Those nodes are now detached (although they remain connected to one another), and the output of the next upstream node is automatically connected to the input of the next downstream node to fill the gap in the node tree.

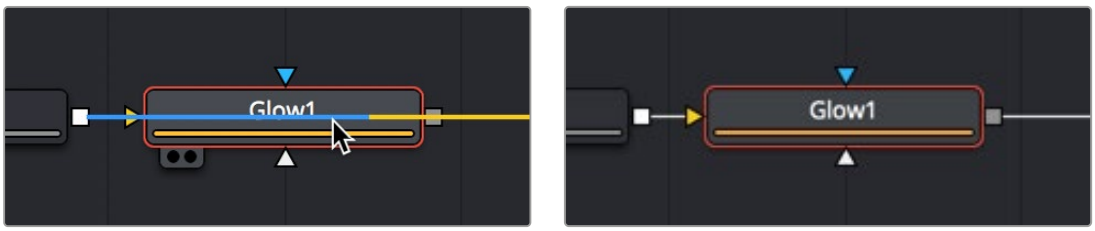


Before extracting a pair of nodes (left), and after extracting a pair of nodes (right).

After you've extracted a node, you can re-insert it into another connection somewhere else. You can only insert one node at a time.

To insert a disconnected node in the Node Editor between two compatible nodes:

1. Hold down the Shift key and drag a disconnected node directly over a connection between two other nodes.
2. Once the connection highlights, drop the node, and then release the Shift key. That node is now attached to the nodes coming before and after it.



Before inserting a node (left), and after inserting a node (right).

TIP: If you hold down the Shift key, you can extract a node and re-insert it somewhere else with a single drag.

Cut, Copy, and Paste Nodes

The standard operations of cut, copy, and paste are also available in the Node Editor. You can use them to temporarily remove nodes from the Node Editor, create duplicate nodes, or even copy the settings from one node and paste those settings into another node with compatible settings.

Cut, Copy, and Paste in the Node Editor

The standard commands all work, but with some special features specific to the Node Editor.

To copy one or more selected nodes, do one of the following:

- Right-click a node and choose Copy from the contextual menu.
- Choose Edit > Copy from the Edit menu (Command-C).

To cut one or more selected nodes, do one of the following:

- Right-click over the node and choose Cut from the contextual menu.
- Choose Edit > Cut from the Edit menu (Command-X).

When you paste into the Node Editor, you create a copy of the last node or nodes you've cut or copied. When pasting, there are a few different things you can do to control where pasted nodes appear.

To paste one or more selected nodes, do one of the following:

- **To paste nodes to be inserted after another node:** Select the node in the node tree you want to insert the pasted node(s) to, and choose Edit > Paste (Command-V).
- **To paste nodes to be disconnected from the rest of the node tree:** Deselect all nodes, and then choose Edit > Paste (Command-V), or right-click anywhere in the Node Editor and choose Paste from the contextual menu.
- **To paste disconnected nodes in a specific area of the Node Editor:** Deselect all nodes, and then click the place in the Node Editor where you want pasted node(s) to appear, and choose Edit > Paste (Command-V), or right-click anywhere in the Node Editor and choose Paste from the contextual menu.
- **To paste a node to replace an existing node in the Node Editor:** Right-click a node in the Node Editor that you want to replace, choose Paste from the contextual menu, and when a dialog appears asking if you want to replace that node, click OK. This only works when you use the contextual menu command.

TIP: When you paste a MediaIn, Loader, or Generator node so it will be inserted after a selected node in the node tree, a Merge tool is automatically created and used to composite the pasted node by connecting it to the foreground input. While this can save you a few steps, some artists may prefer to perform these sorts of merges manually, so this can be changed using the Defaults panel in the Auto tools section of the Fusion > Fusion Settings.

Pasting Node Settings

Instead of pasting a node, you can choose to paste just the parameter settings that you copied from another node. This can be useful if you've carefully set or animated parameters in one node that you want to also use in another node.

Note that you can paste settings between two nodes of the same type, or between two entirely different kinds of nodes that happen to have one or more of the same parameters in the Inspector. When copying settings from one type of node to another, only the settings that match between two nodes will be copied. A common example is to copy an animated Center parameter from a Transform node to the Center parameter of a Mask node.

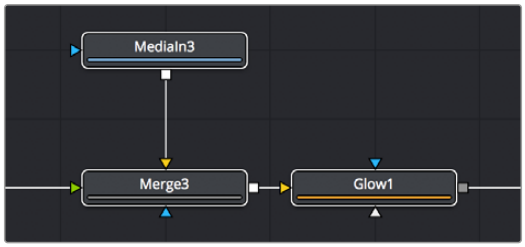
To Paste settings from one node to another:

1. Select a node that has settings you want to copy, and choose Copy from the Edit menu (Command-C).
2. Right-click a node you want to paste those settings to, and choose Paste Settings from the contextual menu.

Copying and Pasting Nodes to and from Any Text Editor

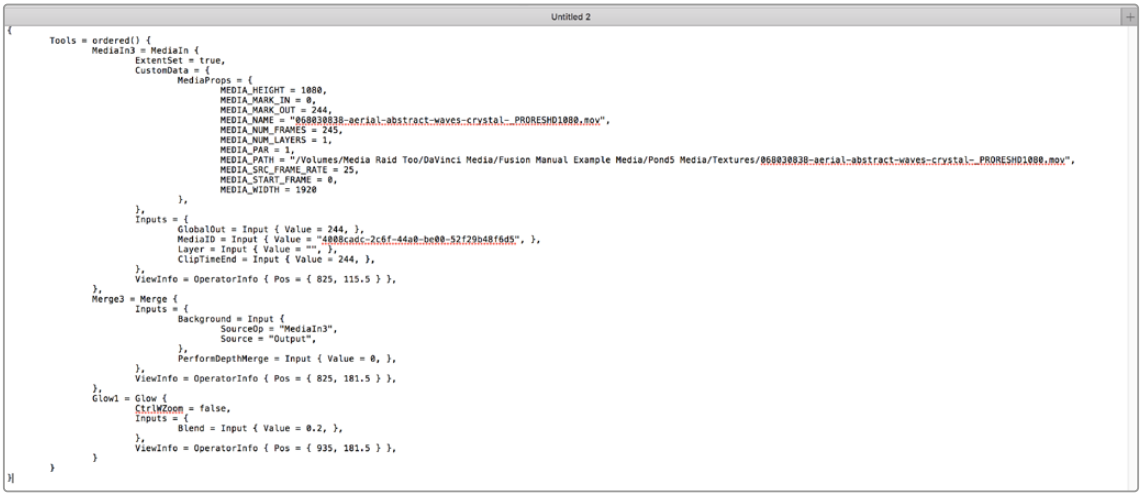
The format of nodes in the Node Editor is not binary, but is in fact a simple text format. The implications of that may not be obvious, but one example benefit is clear when you start dealing with nodes.

One or more nodes can be copied from the Node Editor and pasted directly into a text editor or email. This pastes the selection in text format, just as it's saved internally in Fusion. For example, if you copy the following set of three nodes:



A set of three nodes being copied.

And you then paste into a new text editing document, you get the following:



The same three nodes pasted into a text editor.

At this point, you have the option of editing the text (if you know what you're doing), emailing it to colleagues, or storing it in a digital notepad of some sort for future use. To use this script in Fusion again, you need only copy it and paste it back into the Node Editor.

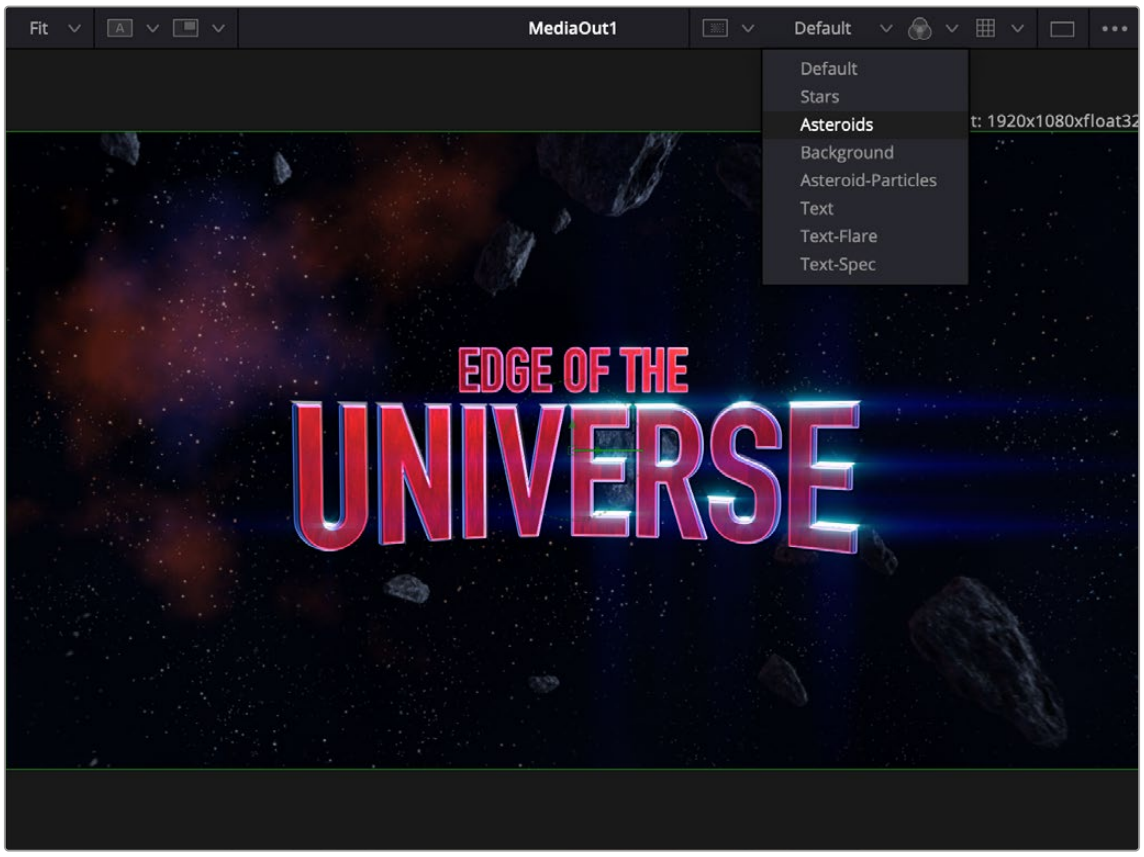
TIP: This is a very easy way to pass specific node settings back and forth between artists who may not be in the same room, city, or country.

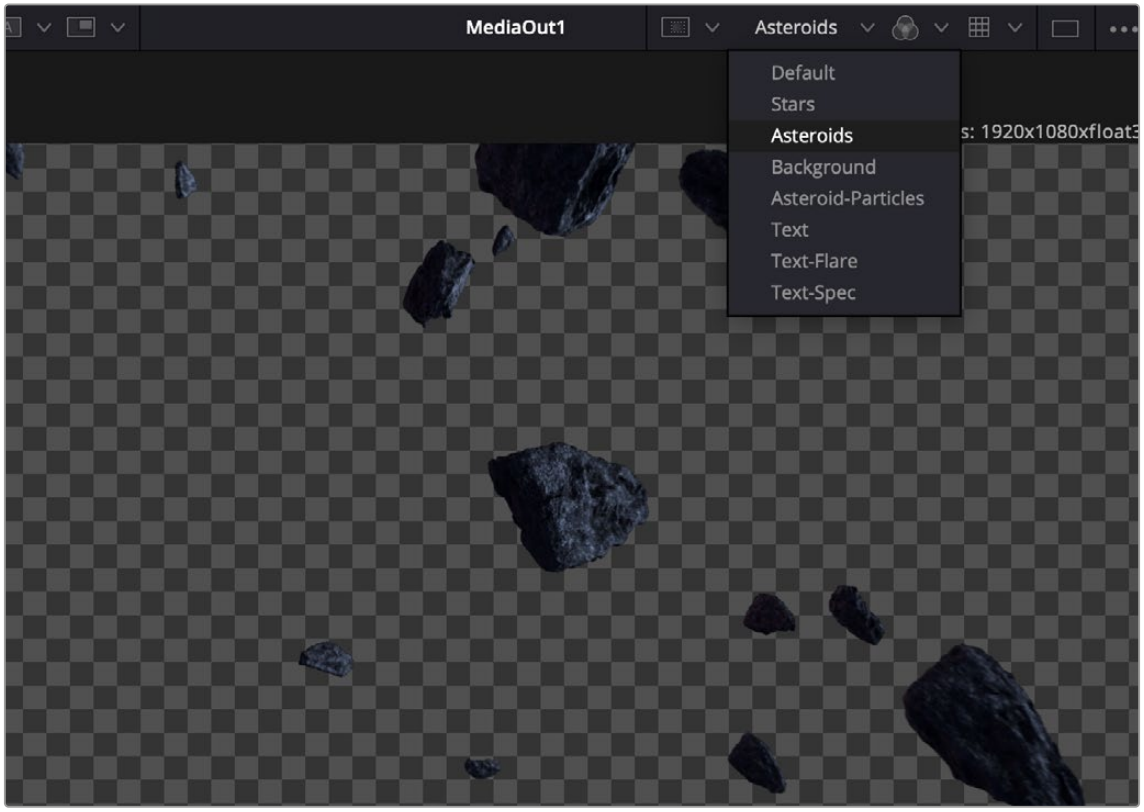
Using Multilayer Nodes (Studio Version Only)

Fusion can intelligently handle multilayer image workflows with support for EXR, Stereoscopic 3D, and PSD assets. Users can see a list of layers on the Fusion viewers and interact with layers in each node in the Settings tab in the Inspector. As this layer structure is retained throughout the toolset, this means you can isolate a tool to work on an individual layer, while still passing through the other layers unmodified.

Viewer

When loading a multilayer file, you can preview specific layers in the Viewer and select a layer from the drop-down list. Files without layer support will show only a 'default' option from the drop-down menu.





Using the Layer drop-down in the Viewer to select just the Asteroids layer from a multilayer file for viewing

Nodes and Inspector

In the Settings tab you'll find a Layer category where you have access to a number of controls to affect certain layers and to control layers from each tool input.

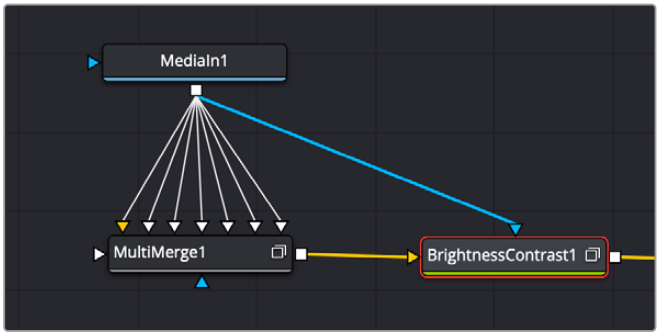
- **Process Layers:** Select the layer to process for effects from:

  Default Layer (default): To apply the effect on the base/unnamed layer, not affecting other layers.

  All Layers: To apply the effect on all available layers.

  Custom: Choice of other available layers to apply the effect on, not affecting other layers.

- **Input Layer:** The Input Layer lists all layers available from the upstream node of the default input. The selected layer can be used by the current tool or further down the node tree in isolation.
- **Effect Mask Layer:** If the tool has an Effect Mask input, this control lists the available layers that can provide mattes for the tool.



Connecting the multilayer file, both directly to the BrightnessContrast Effect Mask (blue arrow) in addition to its input, allows you to quickly create perfect masks of any layer in the file.

Tools with other named inputs (like **Image**) can show additional layer controls with a corresponding name (like **Image Layer**).

Input Layer Controls

For optimal performance, Fusion tools request only the processed layers from each input. The Input Layer, Effect Mask Layer, Image Layer and other input layer controls allow the user to customize this behavior. In each input control, you can select Auto or Match or select an explicit layer value to map:

- **Auto (Default):** Requests the layer selected in the Process Layers selection. If the layer does not exist, the tool falls back to the default layer.
- **Match:** Requests the layer selected in the Process Layers selection. If the layer does not exist, the tool shows a failure.
- **Explicit Layer Selection:** Requests the explicit layer selection from one of the listed layers. This isolates the specific layer for use in the current tool or further down the node tree.

For Fusion plugins, the Process Layer and Input Layer controls are automatically generated in the Inspector's Settings tab. Creators of plugins can decide to define and place the controls explicitly in another place, similar to how the MultiMerge places the controls in the main controls section.

Some Fusion tools include controls to handle additional layer-based decisions.

For example:

- The Renderer3D shows a new Eye control > Layers option when using a stereoscopic camera, allowing the left and right eyes to be output as individual layers.
- When combining multilayered sources using composite nodes (e.g., MultiMerge), Channel Booleans, or a MatteControl, the Settings tab has additional controls to select Background or Foreground options to isolate and combine the selected layers.
- The Combiner tool can be used to create custom layers. Set the Mode to Layer. This exposes layer name fields allowing you to name the new layers.

Instancing Nodes

Normally, when you use copy and paste to create a duplicate of a node, the new node is completely independent from the original node, so that changes made to one aren't rippled to the other. However, there are times when two nodes must have identical settings at all times. For example, when you're making identical color corrections to two or more images, you don't want to constantly have to adjust one color correction node and then manually adjust the other to match. It's a hassle, and you risk forgetting to keep them in sync if you're working in a hurry.

While there are ways to publish controls in one node and connect them to matching controls in another node, this becomes prohibitively complex and time consuming for nodes in which you're making adjustments to several controls. In these cases, creating "instanced" nodes is a real time-saver, as well as an obvious visual cue in your node tree as to what's going on.

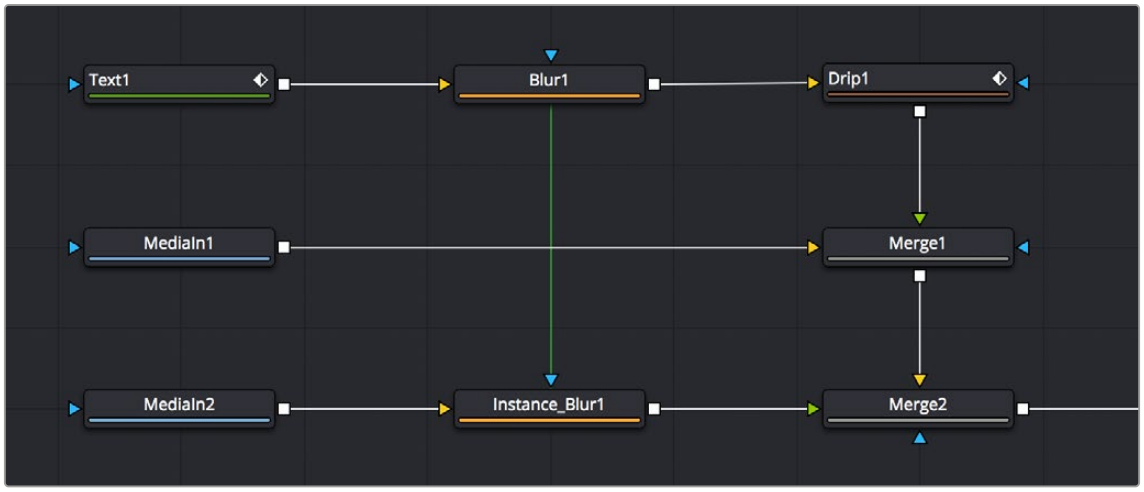
Using Instanced Nodes

Instanced nodes are nodes that have been created using the Paste Instance command, and which share settings with the original node so that a change made to one instanced node is also automatically applied to all other instances of that node (as well as the original node you copied).

To create an Instance, do the following:

1. Select a node you want to instance, and copy it (Command-C).
2. Do one of the following:
   - **To create a disconnected instance of a node:** Right-click in the background of the Node Editor, and choose Paste Instance from the contextual menu (Command-Shift-V).
   - **To insert an instanced node between two other nodes:** Select a node that's upstream of where you want to insert the instanced node, and press Command-Shift-V. Alternatively, you can right-click directly on a connection line, and choose Paste Instance from the contextual menu.

However you paste an instance, the name of that instanced node takes the form "Instance_NameOfNode." If you paste multiple instances, each instance is numbered "Instance_NameOfNode_01."



A green link line shows an instanced Blur node's relationship to the original Blur node it was copied from.

When a node tree contains instanced nodes, a green line shows the link between the original node and its instances. You have the option to hide these green link lines to reduce visual clutter in the Node Editor.

To toggle the visibility of green instance link lines in the Node Editor:

1. Right-click anywhere in the background of the Node Editor.
2. Choose Options > Show Instance Links from the contextual menu.

If you've been using an instance of a node and you later discover you need to use it to apply separate adjustments, you can "de-instance" the node.

To de-instance a node, making it independent:

1. Right-click an instanced node.
2. Choose Deinstance from the contextual menu. That node is now independent from the original node. Once you de-instance a node, you cannot re-instance it, but you can undo the operation.

NOTE: If you've de-instanced a node and you cannot undo the operation because you've restarted DaVinci Resolve, you can only recreate an instance by copying the original and pasting an instance again.

De-Instancing and Re-Instancing Specific Parameters

By default, every parameter in an instanced node is linked to the original node, so that any change you make is rippled across. However, from time to time you'll find the need to independently adjust just one or two parameters while keeping the rest of that node's parameters linked. For this reason, instead of de-instancing the entire tool, you can de-instance individual parameters.

To de-instance a single parameter:

- Right-click on a parameter's name or value in the Inspector, and choose Deinstance from the contextual menu.

If you've only de-instanced individual parameters, you can re-instance those parameters later on if you change your mind.

To re-instance a single parameter:

Right-click on a parameter's name or value in the Inspector, and choose Reinstance from the contextual menu. That parameter immediately inherits the setting of the original node.

Keeping Node Trees Organized

Similar to working with files on your desktop, even the simplest of composites require you to do some amount of organization. In this section we'll look at some basic node operations, some of which you may already be familiar with just from using your computer's operating system or other applications.

Moving Nodes

Selecting one or more nodes and dragging them moves them to a new location, which is one of the simplest ways of organizing a node tree, by grouping nodes spatially according to the role they play in the overall composition.

Keep in mind that the location of nodes in the Node Editor is purely aesthetic, and does nothing to impact the output of a composition. Node tree organization is purely for your own peace of mind, as well as that of your collaborators.

TIP: Once you've arranged the nodes in a composition in some rational way, you can use the Sticky Note and Underlay tools to add information about what's going on and to visually associate collections of nodes more definitively. These tools are covered later in this section.

Snapping Nodes to the Grid

By default, you can position nodes freely wherever you want them to be. However, keeping nodes and connection lines straight and aligned can make them easier to read. To help keep them aligned, you can have nodes you're dragging automatically snap to the grid.

To have nodes snap to the grid as they're dragged:

- Right-click over an empty area of the Node Editor, and choose Arrange Tools > To Grid from the contextual menu. All nodes you drag now snap to the nearest grid coordinate.
- Right-click over an empty area of the Node Editor, and choose Arrange Tools > To Connected from the contextual menu. All nodes you drag now snap to the horizontal or vertical position of the nodes they're attached to.

TIP: You can set "Arrange to Grid" or "Arrange to Connected" as the default for new compositions by choosing Fusion > Fusion Settings in DaVinci Resolve or File > Preferences in Fusion Studio, and turning the Fusion > Node Editor > Arrange To Grid or Arrange to Connected checkboxes on.

Commands to "Clean Up" a Node Tree

The grid in the background of the Node Editor can be used to align nodes, either by eye or automatically.

To "clean up" an unruly node tree:

Right-click in an empty section of the Node Editor, and choose Line Up All Tools to Grid from the contextual menu. All nodes in the Node Editor will move to align and center themselves along the nearest grid lines.

To "clean up" only one or more selected nodes:

Right-click one of the selected nodes and choose Line Up to Grid from the contextual menu. All selected nodes will move to align and center themselves along the nearest grid lines, while all unselected nodes will be left as they are.

Renaming Nodes

Each node that's created is automatically assigned a name (based on its function) and a number (based on how many of that type of node have been created already). For example, the first Blur node added to a composition will be called Blur1, the second will be Blur2, and so on. Although initially helpful, larger compositions may benefit from important nodes having more descriptive names to make it easier to identify what they're actually doing, or to make it easier to reference those nodes in expressions.

To rename a node:

1 Do one of the following:
   - Right-click a node and choose Rename from the contextual menu.
   - Select a node and press F2.

2 When the Rename dialog appears, type a new name, and then click OK or press Return.

> **NOTE:** If multiple nodes are selected, multiple dialogs will appear asking for a name for each tool.

Since Fusion can be scripted and use expressions, the names of nodes must adhere to a scriptable syntax. Only use alphanumeric characters (no special characters), and do not use any spaces.

Also, you cannot start a node name with a number. If you accidentally create a name that doesn't exactly follow the guidelines, spaces and invalid characters will be automatically deleted.

If you want to see the original node types instead of the node names, press and hold Command-Shift-E.

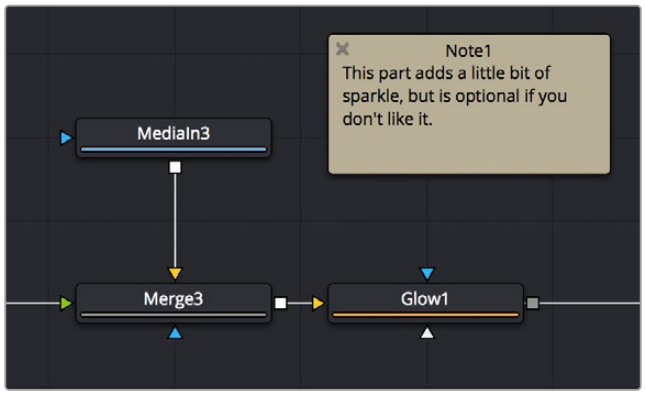
Changing Node Colors

You can change the color of any node by selecting it, opening the Inspector, and choosing a new color from the Node Color pop-up in the Inspector header for that node. Alternatively, you can right-click a node and choose a color from the Set Color submenu.

To return a node to its regular color, right-click it and choose Set Color > Clear Color from the contextual menu, or open the Node Color pop-up for a node in the Inspector, and choose Clear Color.

Using Sticky Notes

A good way to add notes about different parts of a composition, client feedback about various details, and other information you want to keep track of, is to add Sticky Notes to the Node Editor.



A Sticky Note in the Node Editor.

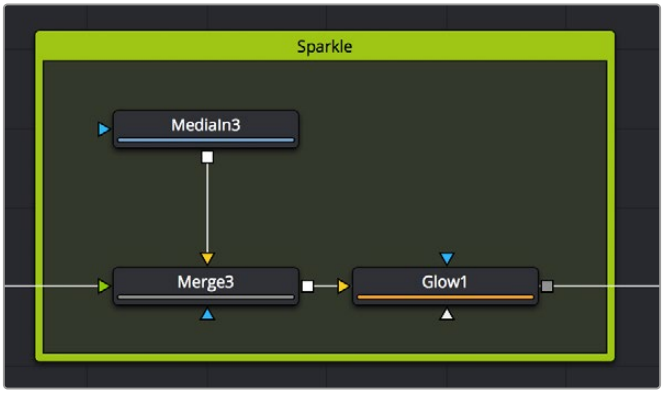
Sticky Notes are yellow boxes in which you can type whatever text you want. They can be resized, moved, and collapsed when they're not being edited, but once created they remain attached to the background of the Node Editor where you placed them until you either move them or delete them.

Methods of working with Sticky Notes:

- **To create a Sticky Note:** Click somewhere in the Node Editor where you want a Sticky Note to appear. Then, press Shift-Spacebar, type sticky, and press the Return key when the Sticky Note appears in the Select Tool window. Alternatively, you can open the Effects Library, open the Tools > Node Editor category, and click or drag the Sticky Notes node to create a new one.
- **To open a Sticky Note to full size:** Double-click a minimized Sticky Note and it expands to a larger, resizable yellow box.
- **To edit a Sticky Note:** If necessary, double-click a Sticky Note to open it to full size, and then click once in the body of the note to place a text cursor. You can edit text within the Sticky Note just like any other text editor.
- **To rename a Sticky Note:** Right-click a Sticky Note, choose Rename, type a new name into the Rename dialog, and click OK. Alternatively, you can select a Sticky Note, press F2 to open the Rename dialog, and press Return to close it when you're done.
- **To resize a Sticky Note:** Double-click a Sticky Note to open it to full size, and then drag any of the edges or corners to make it larger or smaller.
- **To minimize a Sticky Note:** Click the close box at the upper left-hand corner of the Sticky Note, and it collapses to a small tile.
- **To delete a Sticky Note:** Right-click any Sticky Note and choose Delete from the contextual menu or select the Sticky note in the Node Editor and press the Delete key.

Using Underlay Boxes

Underlay Boxes are a good way of associating a collection of nodes that work together to perform a specific task in your composition. They're simply colored rectangles that you can put nodes inside of. Once you place nodes inside an Underlay, you can move the Underlay, and all the nodes within move along with it.



An Underlay in the Node Editor.

Underlay Boxes can be named to identify the purpose of that collection of nodes, and they can be colored to be distinct from other Underlay Boxes or to adhere to some sort of color code for your compositions.

Methods of working with Underlay Boxes:

- **To create an Underlay Box:** Click somewhere in the Node Editor where you want the Underlay Box to appear. Then, press Shift-Spacebar, type under, and press the Return key when the Underlay Box appears in the Select Tool window. Alternatively, you can open the Effects Library, open the Tools > Node Editor category, and click or drag the Underlay Box node to create a new one.

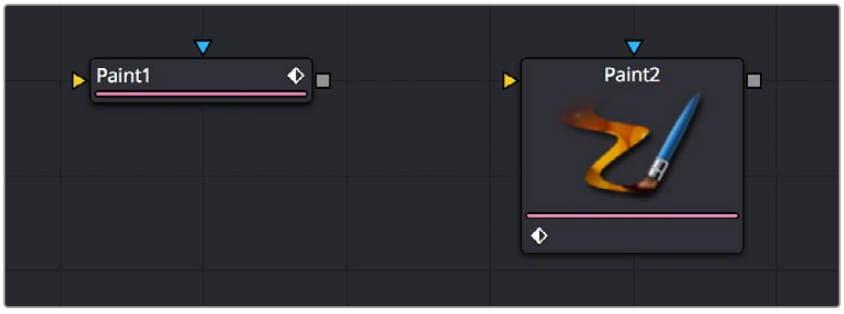
- **To create an Underlay Box around specific nodes:** Select the nodes in the Node Editor that you want surrounded by an Underlay Box. Then, press Shift-Spacebar, type under, and press the Return key when the Underlay Box appears in the Select Tool window. Alternatively, you can open the Effects Library, open the Tools > Node Editor category, and click the Underlay Box node to have it added and sized to encompass all the selected nodes.
- **To resize an Underlay Box:** Drag any of the edges or corners to make it larger or smaller.
- **To rename an Underlay Box:** Option-click the Underlay Box to select just the box and not the contents, and then right-click it and choose Rename (or press F2). Type a new name into the Rename dialog and click OK or press Return.
- **To change the color of an Underlay Box:** Option-click the Underlay Box to select just the box and not the contents, and then right-click it and choose a color from the Set Color submenu.
- **To put nodes inside of an Underlay Box:** Select the nodes you want to place inside an Underlay Box, and then drag them to fit inside. The Underlay Box must be big enough to fit all the nodes. Alternatively, you can place an Underlay Box near a collection of nodes you want to put inside it, and then resize the Underlay Box to encompass all those nodes.
- **To move an Underlay Box and all its nodes:** Once nodes have been placed inside an Underlay Box and have been deselected, you can move the entire collection of nodes together by dragging the Underlay Box by its title bar.
- **To remove nodes from an Underlay Box:** There are two ways you can remove nodes from an Underlay Box.

  With both the Underlay Box and nodes deselected, drag a bounding box or Command-click to select all nodes in the box you want to remove, and drag them out.

  Resize the Underlay Box so that it's smaller than the collection of nodes it originally encompassed. Once an Underlay Box is so small that even the last node sticks out beyond its edge, those nodes are automatically removed from the Underlay Box, and you can move or delete the Underlay Box without moving those nodes.
- **To delete an Underlay Box and all nodes within:** Select an Underlay Box and press the Delete key to delete both the Underlay Box and all nodes found inside it. If you don't also want to delete the nodes, first drag the nodes out of the box.
- **To delete an Underlay Box but keep all nodes within:** Option-click the Underlay Box to select it and not the nodes, and then press the Delete key. The nodes within remain where they were.

Node Thumbnails

Once a source or an effect has been added to the Node Editor, it's represented by a node. By default, nodes are rectangular and thin, making it easier to fit reasonably complicated grades within a relatively small area. However, if you like, you can also display node thumbnails.



A node in the Node Editor shown without and with a thumbnail.

Nodes can be displayed as a small rectangle or as a larger square. The rectangular form displays the node's name in the center, while the square form shows either the tool's icon or a thumbnail of the image it is outputting.

TIP: Even if you're not displaying node thumbnails, you can quickly obtain detailed information about a node and the data it's processing by hovering your pointer over it in the Node Editor and viewing the tooltip bar below.

Choosing Which Nodes Show Thumbnails

If you want to use node thumbnails to help visually identify media and operations in your node trees, there are a variety options for which nodes should display thumbnails in the contextual menu that appears when you right-click anywhere in the background of the Node Editor.

Force All Tile Pictures

This option shows thumbnails for every single node in the Node Editor. This can make simple node trees easier to read, but it'll make all node trees take up considerably more room.

NOTE: If Show Thumbnails is enabled, nodes may not update until the playhead is moved in the Time Ruler.

Force Active Tile Pictures

You may also choose to only show thumbnails for nodes that are currently selected, which can make it easier to see which node you're working on. When nodes become deselected, the thumbnails will be hidden again.

Force Source Tile Pictures

This enables thumbnails for all MediaIn and Loader nodes in the Node Editor, as well as all generators, and is a great way to be able to quickly see where all the clips are in a composition.

Force Mask Tile Pictures

This enables thumbnails for all Mask nodes in a composition, which can make them easier to distinguish when you're building complex shapes made from multiple Mask nodes.

Manually Showing Tile Pictures and Node Options

You also have the option of manually choosing which nodes you'd like to show thumbnails. For example, there may be certain key points of the node tree where you'd like to see a small visual representation of what's happening in the composition.

To toggle thumbnails for one or more specific nodes:

1. Select one or more nodes in the Node Editor.
2. Right-click one of the selected nodes, and choose one of the following from the contextual menu:
   - Show > Show Tile Pictures
   - Show > Show Modes/Options

When you've manually enabled thumbnails for different nodes, they'll remain visible whether or not those nodes are selected.

Switching Thumbnails between Images and Icons

Whenever you enable node thumbnails, you have the choice of having these thumbnails either display an image of the state of the image at that node, or you can instead choose to display the icon for that particular node. The setting for this affects all nodes at once.

To display icons instead of thumbnails:

Right-click anywhere in the background of the Node Editor and deselect Show Thumbnails in the contextual menu.

Sometimes Nodes Only Show Icons

As you add more and more nodes to a composition, you'll notice that some nodes never display an image in their thumbnail. In these cases, the default icon for that node is displayed instead of an image.

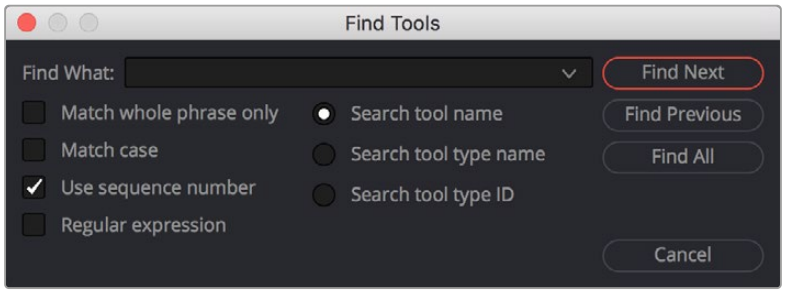
Most nodes in the Particle and 3D categories fall into this group. The exceptions are the pRender node and the Render 3D node. These two nodes are capable of displaying a rendered thumbnail if you have the menu options set for Thumbnails to be displayed.

In other cases, whether nodes display images in their thumbnail is more situational. Some Transform nodes are able to concatenate their results with one another, passing the actual processing downstream to another node later in the node tree. In this case, upstream Transform nodes don't actually process the image, so they don't produce a thumbnail.

In other situations where the Loader is not reading in a clip or the clip is trimmed in the Keyframes Editor to be out of range, it can cause the node not to process the image, so it will not produce a rendered Thumbnail image. Also, nodes that have been set to Pass Through mode are disabled and do not display a rendered Thumbnail image.

Finding Nodes

Modern visual effects require detailed work that often results in compositions with hundreds of nodes. For such large node trees, finding things visually would have you panning around the Node Editor for a long, long time. Happily, you can quickly locate nodes in the Node Editor using the Find dialog.



The Find dialog lets you quickly locate nodes wherever they are in the Node Editor.

Performing Simple Searches

To do simple searches using node names is easy.

To search for a node in the Node Editor:

1. Press Command-F, or right-click in an empty area of the Node Editor and choose Find from the contextual menu.
2. When the Find dialog appears, do the following:
   - Enter a search term in the Find field.
   - Choose search options, such as whether to match the whole phrase in the Find field, whether to match the case, whether to use a sequence number, or whether to use a regular expression in the Find field.
   - Choose what to search. Options include tool name, tool type name, or tool type ID.
3. To perform the find, do one of the following:
   - Click Find Next to try to select a downstream node matching the criteria.
   - Click Find Previous to try to select an upstream node matching the criteria.
   - Click Find All to try to select all nodes in the Node Editor that match the criteria.

   The Find window closes. If either the Find Next, Find Previous, or Find All operations are successful, the found node or nodes are selected. If not, a dialog appears letting you know that the string could not be found.

TIP: Finding all the nodes of a particular type can be very useful if you want, for example, to disable all Resize nodes. Find All will select all the nodes based on the search term, and you can temporarily disable them by pressing the shortcut for Bypass, Command-P.

Using Regular Expressions

If you need to do more complicated searches, you can turn on the Regular Expression checkbox, which lets you enter some simple expressions with which to create more complex find operations. Some useful examples of regular expressions that are valuable include the use of Character Sets.

Character Sets

Any characters typed between two brackets [ ] will be searched for. Here are some examples of character set searches that work in Fusion.

`[a-z]`
Finds: Every node using a lower caps letter

`[a-d]`
Finds: Every lower caps letter from a to d, and will find nodes with a, b, c, or d

`[Tt]`
Finds: Every node with an upper case T or a lower case t

`[aeiou]`
Finds: Every vowel

`[0-9]`
Finds: Every numeral

`[5-7]`
Finds: Every numeral from five to seven, and will find nodes numbered with 5, 6, or 7

Custom Node Settings

When a node is added to the Node Editor, its parameters are set to the default values for that type of node. If you find yourself constantly readjusting the parameters of a node to a preferred starting point as soon as it's added to the node tree, you can override the default node settings with your own custom settings.

To save new default settings for a particular type of node:

1. Create a new node.
2. Open the Inspector and customize that node's settings to the new defaults you want it to have.
3. Right-click that node in the Node Editor, or right-click that node's control header in the Inspector, and choose Settings > Save Default from the contextual menu.

TIP: You can also save six different settings for a node in the Node Editor using the Version buttons at the top of the Inspector. For more information, *see Chapter 8, "Editing Parameters in the Inspector,"* in the Fusion Reference Manual.

Managing Saved Settings

Custom node default settings are saved to a folder on your hard drive that's based on the Path Map > Defaults preference in the Fusion Settings. This path is customizable for facilities where multiple compositing artists use a common set of facility defaults, stored somewhere commonly accessible.

The default paths are:

For macOS systems, this path defaults to: /UserName/Library/Application Support/Blackmagic Design/DaVinci Resolve/Fusion/Defaults.

For Windows systems, this path defaults to C: \Users\<username>\AppData\Roaming\Blackmagic Design\DaVinci Resolve\Fusion\Defaults.

For Linux systems, this path defaults to: ~/.fusion/BlackmagicDesign/DaVinci Resolve/Fusion/Defaults.

If you browse this directory, the settings for each node are saved using a name taking the form INTERNALNAME_PUBLICNAME.settings, where INTERNALNAME is the internal name of the Fusion tool, and PUBLICNAME is the name of the Node that's derived from the internal Fusion tool. For example, the default setting for a Blur node would be called Blur_Blur.setting. This naming convention is partly to ensure that third-party plugin nodes don't overwrite the defaults for built-in Fusion nodes that happen to have the same name.

Resetting Defaults

Even if you've created new default settings for new nodes, you can always reset individual parameters to the original default setting. In addition, it's easy to restore the original default settings for new nodes you create.

To reset a single parameter to the original default settings:

1. Create a new node.
2. Open the Inspector and customize a parameter to the new default value you want it to have.
3. Right-click that parameter in the Inspector, and choose Set to Default from the contextual menu.

To reset every parameter in a node to the original defaults, do one of the following:

- Right-click on the node and choose Settings > Reset Default.
- Right-click that node's control header in the Inspector, and choose Settings > Reset Default.
- Delete the .setting file from the Defaults folder.

> **NOTE:** When you use the Settings > Reset Default command, the default .setting file is deleted. If you want to save a node's settings as alternate settings, you should use the Settings > Save As command.

Saving and Loading Alternate Node Settings

Once you change parameter values for a node using the Inspector, those values can also be saved as an alternate setting for that node, which can be reused at a later time.

To save alternate settings for a node:

1. Right-click on a tool, and then choose Settings > Save As from the contextual menu.
2. When the Save File dialog appears, enter a name for the Setting and save it to your hard drive. Unlike saved defaults, the .settings files can be saved anywhere on the file system. They do not need to be in the Default Settings folder.

To load a saved setting for one or more nodes:

1. Right-click a node and choose Settings > Load from the contextual menu.
2. Use the Open File dialog to select the settings you want to load into that node, and then click Open. Those settings are now applied to that node.

Adding Saved Settings from the File System

Saved settings in your File system can also be used to create new nodes by dragging the .setting file into the Node Editor from a standard file browser. Once dropped, that setting turns into a new node.

> **TIP:** If you drop a setting directly onto a connection line, the new node will be inserted onto that connection.

Node Modes Including Disable and Lock

Right-clicking one or more nodes and opening the contextual menu reveals a series of commands in the Modes submenu, some with accompanying keyboard shortcuts, that let you set control visibility, disable, lock, update, and cache nodes.

- **Show Controls:** Sets whether that node reveals its parameters in the Inspector when it's selected and whether its onscreen controls appear in viewers. On by default.
- **Pass Through:** (Command-P) Identical to the toggle switch in the Inspector that turns nodes off and on. Disabled nodes are ignored as image data is passed from the next previous upstream node to the next downstream node. On by default.
- **Locked:** (Command-L) Identical to the lock button in the Inspector that prevents a node from being edited in the Inspector. Off by default.
- **Update:** (Command-U) On by default. While this option is enabled, all changes to the node will cause it to re-render. When Update is disabled, you can still change the node's parameters, but those changes will not process or update the image until Update is re-enabled. While disabled, the last processed image for that node will be displayed as a freeze frame. One example of when this is useful is when you have a large or processor-intensive composition (such as a particularly intense particle system), and disabling this option temporarily will let you quickly make several quick parameter adjustments to different nodes without forcing you to wait for the node tree to re-render after every adjustment. Another example of when this is useful is when you want to quickly see the effect of animated downstream nodes while keeping upstream nodes that are too processor-intensive to play in real time from rendering additional frames.
- **Force Cache:** When enabled, this node's output for the current frame has an extremely high cache priority, essentially forcing it to stay cached in memory. Off by default.

Toggling any one of these node modes displays a badge within that node indicating its state.

Node Editor Options

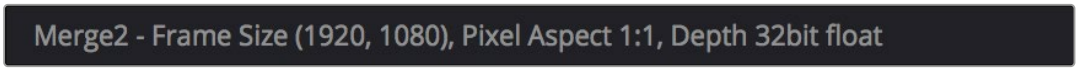
Right-clicking in an empty area of the Node Editor will bring up the contextual menu and the Options submenu. The Options submenu contains several choices that can be used to customize how the Node Editor looks and behaves.

- **Pipes Always Visible:** Enabling this option causes a connection to cross over a node instead of beneath it, sometimes making it easier to follow the connection's path.
- **Show Hidden Pipes:** When enabled, the Inspector option to Hide Incoming Connections in every node is overridden and all connections are displayed in the Node Editor.
- **Aspect Correct Tile Pictures:** Aspect Correct Tile Pictures forces the display of thumbnails to be aspect corrected, which is slower but visually more accurate. This option is enabled by default.
- **Full Tile Render Indicators:** Enabling this option causes the thumbnail to flash green when rendering, which makes it easier to identify which node is processing in a large, complex node tree.
- **Show Grid:** This option can be used to enable or disable the Node Editor's background grid.
- **Show Instance Links:** When enabled, the Node Editor draws a green connection between an instanced node and its parent.
- **Auto Remove Routers:** If routers are disconnected from a tool, they are automatically deleted from the Node Editor. This option is enabled by default to eliminate the need to delete orphaned routers.

- **Show Navigator:** Enabling this option displays a small overview window of the entire node tree in the Node Editor's top-right corner. For more information, see the Navigator section in this chapter.
- **Auto Navigator:** The Navigator only appears when one or more nodes is outside the visible area of the Node Editor. For more information, see the Navigator section in this chapter.
- **Build Flow Vertically/Horizontally:** Node trees can be built either horizontally from left to right or vertically from top to bottom. Enabling one of these options determines whether new nodes are added beneath the current node or to the right of the current tool.
- **Orthogonal/Direct Pipes:** Use these two options to decide whether connections between nodes are drawn as Direct (straight) lines or Orthogonal (bent) lines.

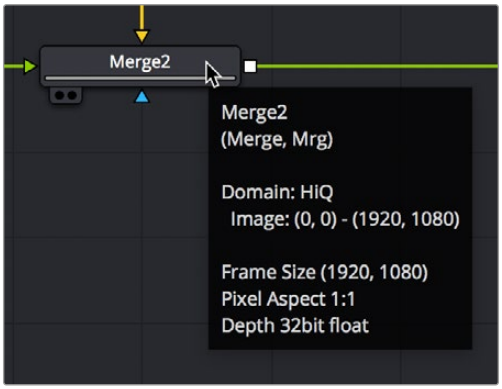
Node Tooltips and the Status Bar

Even in simple node trees, it's easy to forget some essential detail about the nodes in your comp. To help you figure out what everything's for, you can hover the pointer over any node in the Node Editor to display information in the Status bar at the bottom of the Node Editor consisting of that node's name, frame size, pixel aspect, resolution, and color depth.



The Status bar located beneath the Node Editor.

If you wait a few moments later, a more elaborate presentation of the same information appears within a floating tooltip in the Inspector. This tooltip gives you additional information about the Domain (Image and DoD) and the data range used by that clip.



The floating tooltip showing node information that appears within the Node Editor.

Chapter 6

Node Groups, Macros, and Fusion Templates

This chapter reveals how to use groups, macros, and templates in Fusion so working with complex effects becomes more organized, more efficient, and easier.

Contents

Groups 164
Creating Groups 164
Deleting Groups 165
Expanding and Collapsing Groups 165
Panning and Scaling within Open Group Windows 165
Ungrouping Nodes 165
Saving and Reusing Groups 166
Macros 166
Creating Macros 167
Using Macros 168
Re-Editing Macros 168
Other Macro Examples 169
Creating Fusion Templates 169
Getting Started with a Fusion Title Template 169
Saving a Title Macro 169
Using Your New Title Template 173
Getting Started with a Fusion Transition Template 174
Creating a Fusion Transition Template 174
Using Your New Transition Template 177
Getting Started with a Fusion Generator Template 178
Creating a Fusion Generator Template 178

Using Your New Generator Template 181
Creating a Fusion Effect Template 181
Changing Durations of a Template 182
Creating Media Drop Zones in Templates 183
Creating a Custom Template Icon 184
Using Fusion Template Bundles 184

Groups

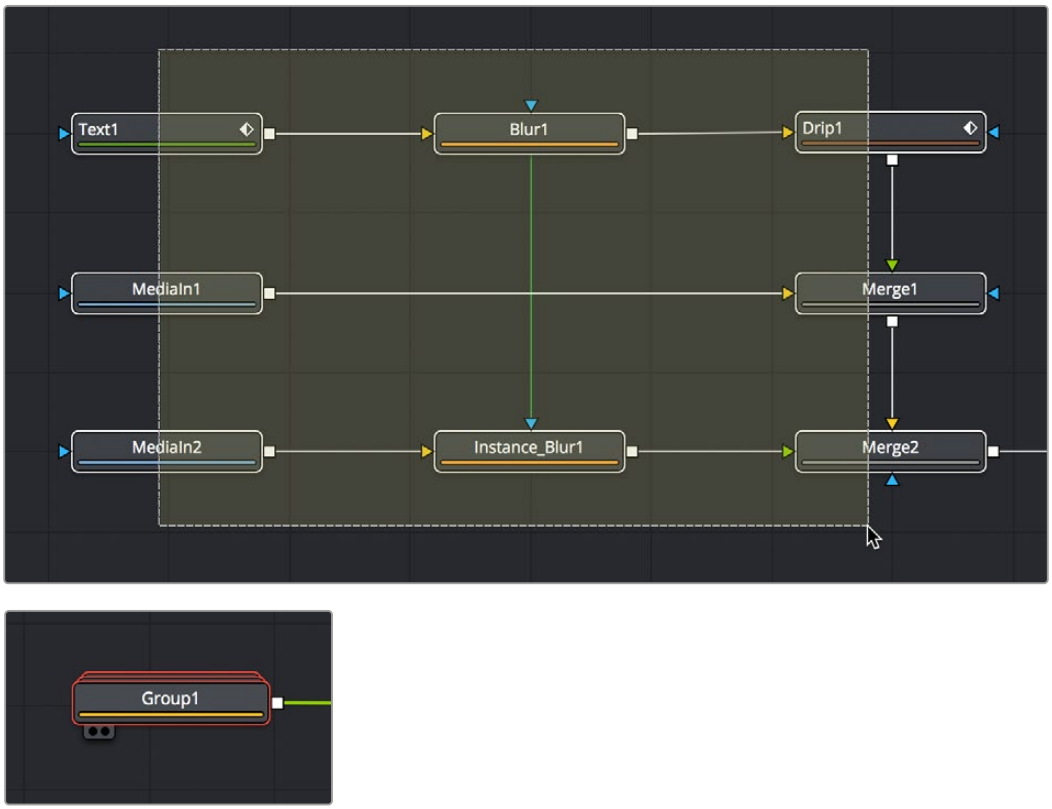
When you work on complex visual effects, node trees can become sprawling and unwieldy, so grouping tools together can help you better organize all the nodes and connections. Groups are containers in your node tree that can hold multiple nodes, similar to the way a folder on your Desktop holds multiple files. There is no limit to the number of nodes that can be contained within a group, and you can even create subgroups within a group.

Creating Groups

Creating a group is as simple as selecting the nodes you want to group together and using the Group command.

To create a group:

1. Select the nodes you want grouped together.
2. Right-click one of the selected nodes and choose Group from the contextual menu (Command-G).



Several nodes selected in preparation for making a group (top), and the resulting group (bottom).

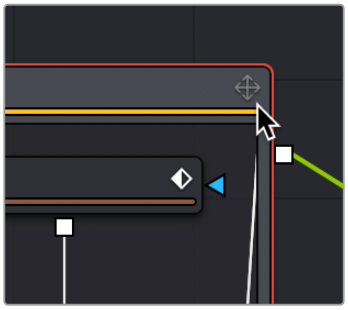
The selected nodes are collapsed into a group, which is displayed as a single node in the Node Editor. The Group node can have inputs and outputs, depending on the connections of the nodes within the group. The Group node only displays inputs for nodes that are already connected to nodes outside the group. Unconnected inputs inside the group will not have an Input knot displayed on the Group node.

Deleting Groups

Deleting a group is no different from deleting any other node in the Node Editor. Select a group and press Delete, Backspace, or Forward-Delete, and the group along with all nodes contained within it are removed from the node tree.

Expanding and Collapsing Groups

A collapsed group is represented by a single “stack” node in the node tree. If you want to modify any of the nodes inside the group, you can open the group by double-clicking it or by selecting the group node and pressing Command-E.



An open group window showing the minimize button.

When you open a group, a floating window shows the nodes within that group. This floating window is its own Node Editor that can be resized, zoomed, and panned independently of the main Node Editor. Within the group window, you can select and adjust any node you want to, and even add, insert, and delete nodes while it is open. When you’re ready to collapse the group again, click the minimize button at the top left corner of the floating window, or use the keyboard shortcut (Cmd-E).

Panning and Scaling within Open Group Windows

You can pan and scale an open group window using the same mouse buttons you use to pan and scale the main Node Editor. However, when you’re working in an expanded group and simultaneously making changes to the main node tree, you may want to prevent the expanded group from being individually panned or scaled. Turning off the Position button at the right of the group title bar locks the group nodes to the size of the nodes in the rest of the overall node tree. Turning on this Position button lets you size group nodes independently of the rest of the node tree.

Ungrouping Nodes

If you decide you no longer need a particular group, or you simply find it easier to have constant access to all the nodes in the group at once, you can decompose or “ungroup” the group without deleting the nodes within it to eliminate the group but keep the contents in the Node Editor.

To ungroup nodes, do the following:

1. Right-click on the group.
2. Choose Ungroup from the contextual menu. The nodes inside the group are placed back in the main node tree.

Saving and Reusing Groups

One of the best features of groups is that every group and its settings can be saved for later use in other shots or projects. Groups and their settings can be recalled in various ways.

A good example of when you might want to Save and Load a group is in a studio with two or more compositing artists. A lead artist in your studio can set up the master comp and create a group specifically for keying greenscreen. That key group can then be passed to another artist who refines the key, builds the mattes, and cleans up the clips. The setting can then be saved out and loaded back into the master comp. As versions are improved, these settings can be reloaded, updating the master comp.

Methods of saving and reusing groups:

- **To save a group:** Right-click a group and choose Settings > Save As from the contextual menu.
- **To reuse a group:** Drag it from your computer's file browser directly into the Node Editor. This creates a new group node in the node tree with all the same nodes as the group you saved.
- **To load the settings from a saved group to another group with the same nodes:** Right-click a group in the Node Editor and choose Settings > Load from the contextual menu.

In Fusion Studio, you can also save and reuse groups from the Bins window:

- **To save a group:** Drag the group from the Node Editor into the opened Bin window. A dialog will appear to name the group setting file and the location where it should be saved on disk. The .setting file will be saved in the specified location and placed in the bins for easy access in the future.

Macros

Some effects aren't built with one tool, but with an entire series of operations, sometimes in complex branches with interconnected parameter controls. Fusion provides many individual effects nodes for you to work with but gives users the ability to repackage them in different combinations as self-contained "bundles" that are either macros or groups. These "bundles" have several advantages:

- They reduce visual clutter in your node tree.
- They ensure proper user interaction by allowing you to restrict which controls from each node of the macro are available to the user.
- They improve productivity by allowing artists to quickly leverage solutions to common compositing challenges and creative adjustments that have already been built and saved.

Macros and groups are functionally similar, but they differ slightly in how they're created and presented to the user. Groups can be thought of as a quick way of organizing a composition by reducing the visual complexity of a node tree. Macros, on the other hand, take longer to create because of how customizable they are, but they're easier to reuse in other comps.

Creating Macros

While macros let you save complex functions for future use in very customized ways, they're actually pretty easy to create.

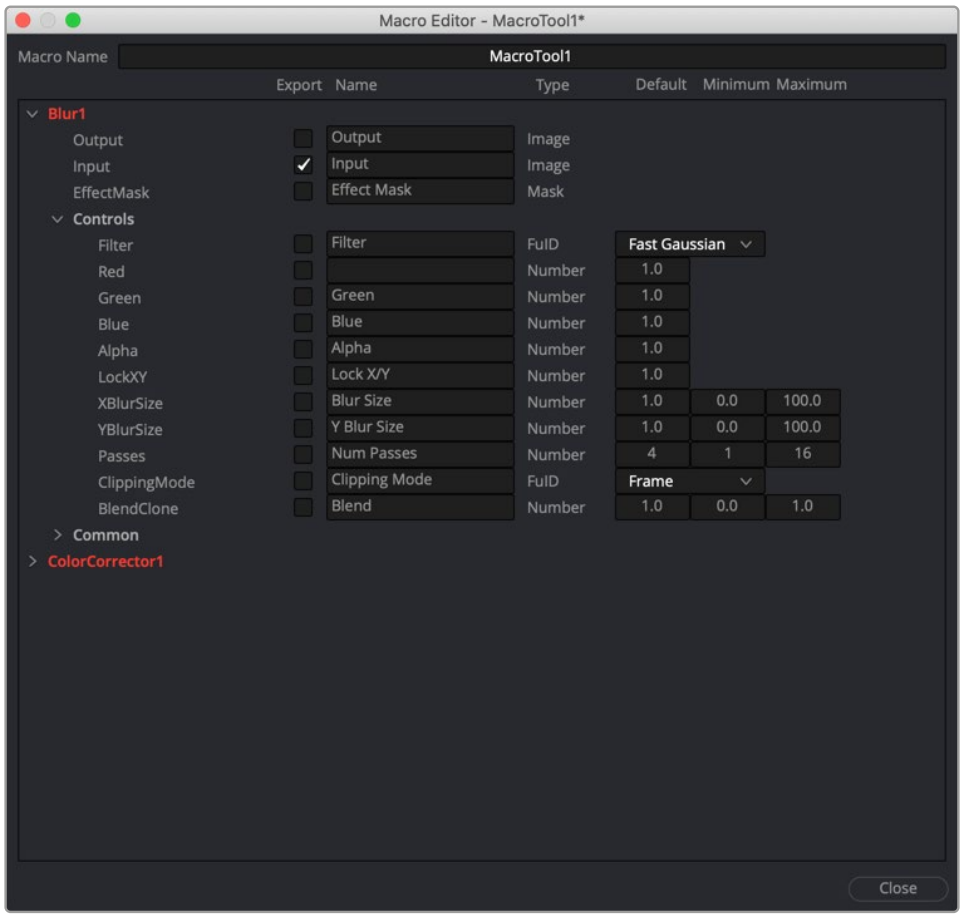
To make a macro from nodes in the Node Editor:

1 Select the nodes you want to include in the macro you're creating. Because the macro you're creating will be for a specific purpose, the nodes you select should be connected together to produce a particular output from a specific set of inputs.

> **TIP:** If you want to control the order in which each node's controls will appear in the macro you're creating, Command-click each node in the order in which you want it to appear.

2 Right-click one of the selected nodes and choose Macro > Create Macro from the contextual menu.

A Macro Editor window appears, showing each node you selected as a list, in the order in which each node was selected.



The macro editor with a Blur node and Color Corrector node.

3 First, enter a name for the macro in the field at the top of the Macro Editor. This name should be short but descriptive of the macro's purpose. No spaces are allowed, and you should avoid special characters.

4 Next, open the disclosure control to the left of each node that has controls you want to expose to the user and click the checkbox to the right of each node output, node input, and node control that you want to expose.

   The controls you check will be exposed to users in the order in which they appear in this list, so you can see how controlling the order in which you select nodes in Step 1, before you start editing your macro, is useful. Additionally, the inputs and outputs that were connected in your node tree are already checked, so if you like these becoming the inputs and outputs of the macro you're creating, that part is done for you.

   For each control's checkbox that you turn on, a series of fields to the left of that control's row lets you edit the default value of that control as well as the minimum and maximum values that control will initially allow.

5 When you're finished choosing controls, click Close.

6 A dialog prompts you to save the macro. Click Yes.

7 A Save Macro As dialog appears in which you can re-edit the Macro Name (if necessary), and choose a location for your macro.

   To have a macro appear in the Fusion page Effects Library Tools > Macros category, save it in the following locations:

   - **On macOS:** Macintosh HD/Users/username/Library/Application Support/Blackmagic Design/DaVinci Resolve/Fusion/Macros/
   - **On Windows:** C:\Users\username\AppData\Roaming\Blackmagic Design\DaVinci Resolve\Support\Fusion\Macros
   - **On Linux:** home/username/.local/share/DaVinciResolve/Fusion/Macros

   To have a macro appear in the Fusion Studio Effects Library Tools > Macros category, save it in the following locations:

   - **On macOS:** Macintosh HD/Users/username/Library/Application Support/Blackmagic Design/Fusion/Macros/
   - **On Windows:** C:\Users\username\AppData\Roaming\Blackmagic Design\Fusion\Macros
   - **On Linux:** home/username/.fusion/BlackmagicDesign/Fusion/Macros

8 When you're done, click Save.

Using Macros

Macros can be added to a node tree using the Add Tool > Macros or Replace Tool > Macros submenus of the Node Editor contextual menu.

Re-Editing Macros

To re-edit an existing macro, just right-click anywhere within the Node Editor and choose the macro you want to edit from the Macro submenu of the same contextual menu. The Macro Editor appears, and you can make your changes and save the result.

Groups Can Be Accessed Like Macros

Groups can also be loaded from the Insert Tool > Macros submenu if you save a group's setting file to the Macros folder in your file system. For example, on macOS, the Macintosh HD/Library/Application Support/Blackmagic Design/DaVinci Resolve/Fusion/Macros/ directory.

Other Macro Examples

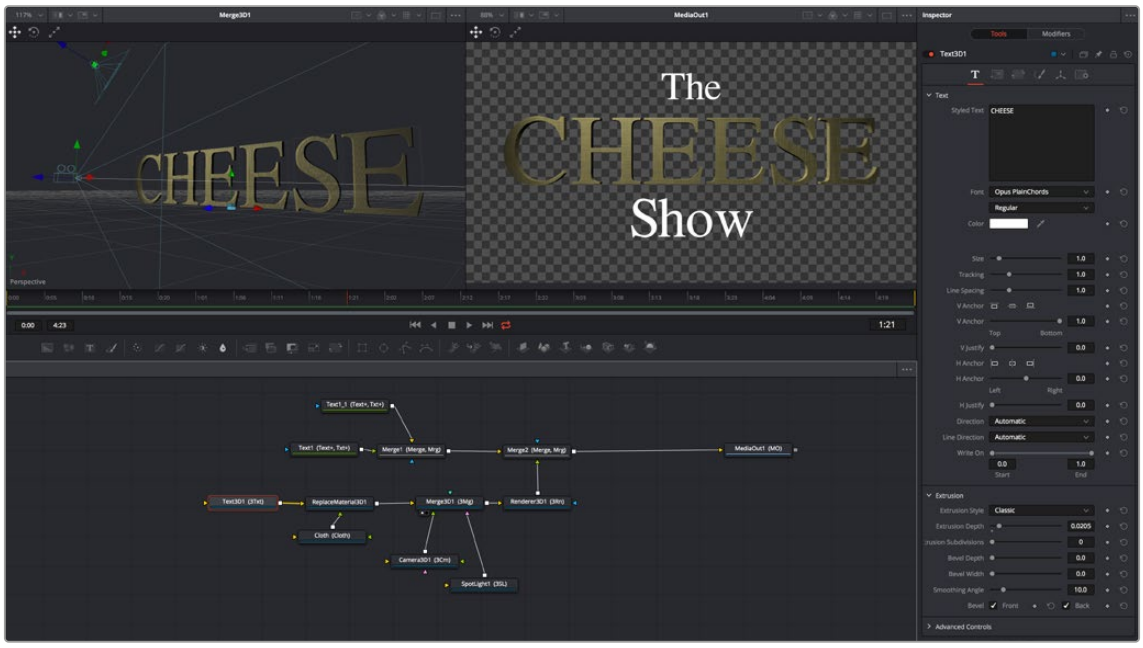
A macro can also be used as a custom LUT. Just copy the macro's .setting file to the LUTs: folder, and the macro will be selectable in the viewers as a LUT. These LUT macros can be used for more than just a color adjustment; you could make a macro that does YUV 4:2:2 resampling, a resize, a sharpening filter, or just watermarking.

Creating Fusion Templates

The integration of Fusion into DaVinci Resolve has enabled the ability to create Fusion Titles, Transitions, Effects, and Generators templates for use in the Edit page. You can create these templates in the Fusion page or within Fusion Studio and then copy them into DaVinci Resolve. Fusion Titles, Generators, and Transition templates are essentially comps created in Fusion but editable in the Timeline of the Edit page with custom controls. This section shows you how it's done.

Getting Started with a Fusion Title Template

The first part of creating a Fusion title template is to create a Fusion composition consisting of Fusion-generated objects assembled to create nearly any kind of title or generator you can imagine. If you're really ambitious, it can include animation. In this example, 3D titles and 2D titles have been combined into a show opener.

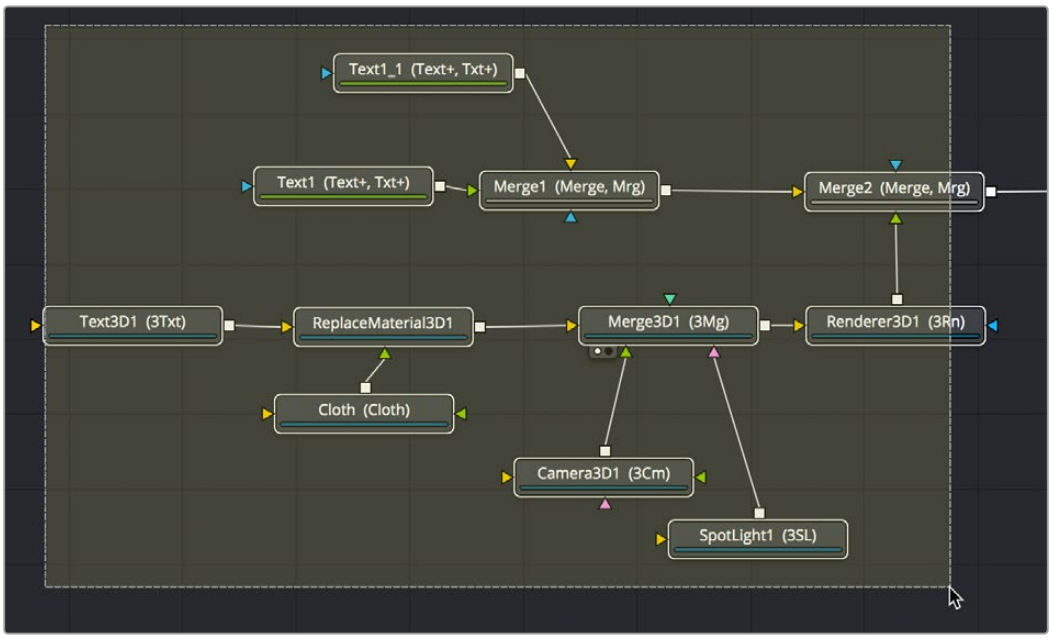


Building a composition to turn into a title template.

Saving a Title Macro

Macros are basically Fusion compositions that have been turned into self-contained nodes. Ordinarily, these nodes are used as building blocks inside of Fusion so that you can turn frequently-made compositing tricks that you use all the time into your own nodes. However, you can also use this macro functionality to build title templates for the Edit page.

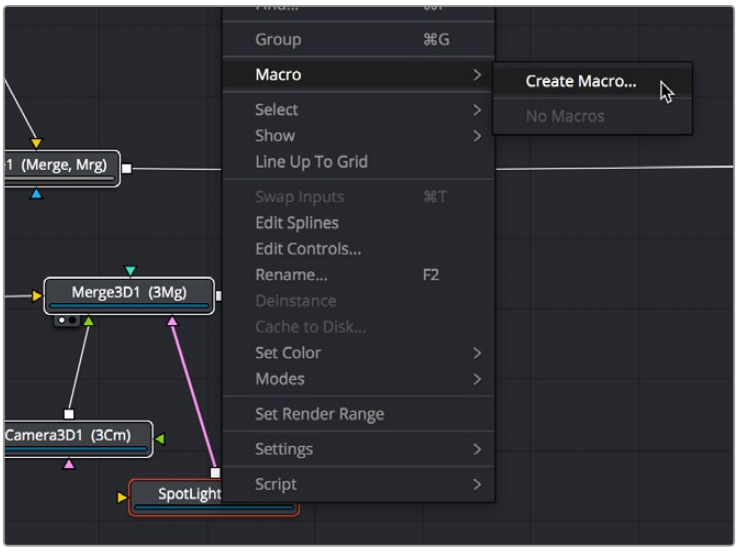
Having built your composition, select every single node you want to include in that template except for the MediaIn and MediaOut nodes in DaVinci Resolve or Loader and Saver nodes in Fusion Studio.



Selecting the nodes you want to turn into a title template.

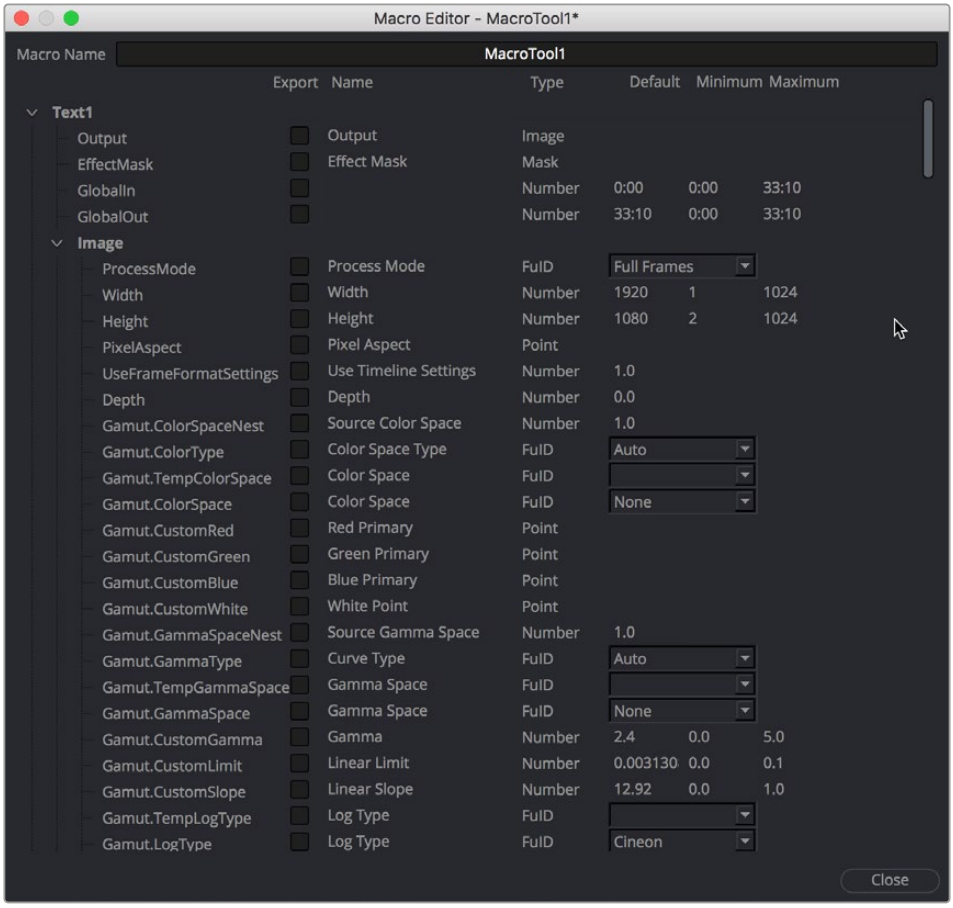
TIP: If you want to control the order in which node controls will be displayed later on, you can Command-click each node you want to include in the macro, one by one, in the order in which you want controls from those nodes to appear. This is an extra step, but it keeps things better organized later on.

Having made this selection, right-click one of the selected nodes and choose Macro > Create Macro from the contextual menu.



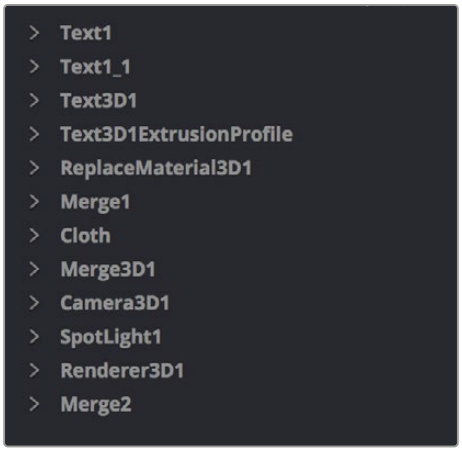
Creating a macro from the selected nodes.

The Macro Editor window appears, filled to the brim with a hierarchical list of every parameter in the composition you've just selected.



The Macro Editor populated with the parameters of all the nodes you selected.

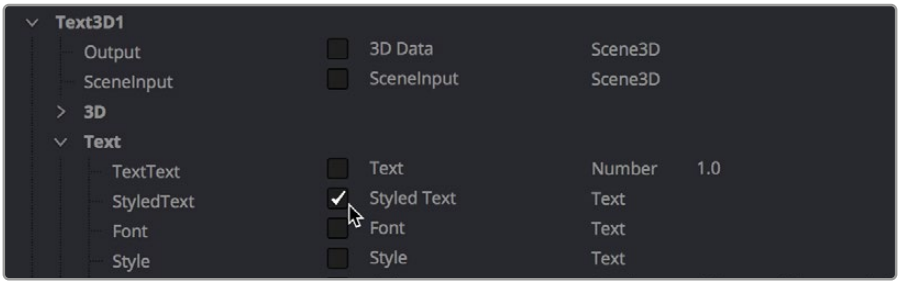
This list may look intimidating, but closing the disclosure control of the top Text1 node shows us what's really going on.



A simple list of all the nodes we've selected.

Closing the top node's parameters reveals a simple list of all the nodes we've selected. The Macro Editor is designed to let you choose which parameters you want to expose as custom editable controls for that macro. Whichever controls you choose will appear in the Inspector whenever you select that macro, or the node or clip that macro will become.

So all we have to do now is to turn on the checkboxes of all the parameters we'd like to be able to customize. For this example, we'll check the Text3D node's Styled Text checkbox, the Cloth node's Diffuse Color, Green, and Blue checkboxes, and the SpotLight node's Z Rotation checkbox, so that only the middle word of the template is editable, but we can also change its color and tilt its lighting (making a "swing-on" effect possible).



Selecting the checkboxes of parameters we'd like to edit when using this as a template.

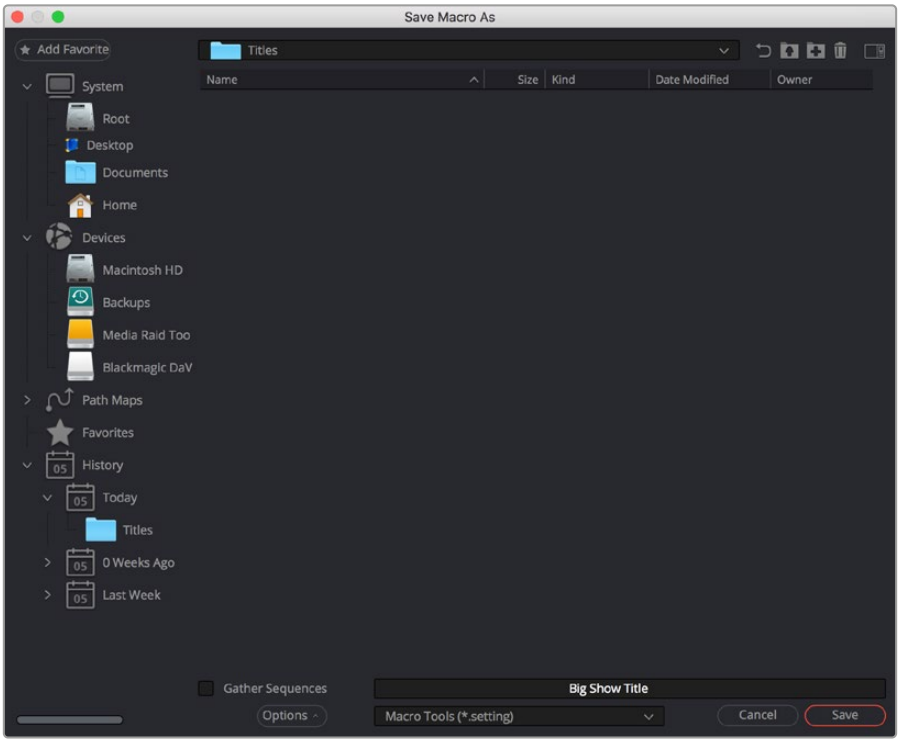
Once we've turned on all the parameters we'd like to use in the eventual template, we click the Close button, and a Save Macro As dialog appears.

To have the Title template appear in the Effects Library > Titles category of DaVinci Resolve, save the macro in the following locations:

On macOS: Macintosh HD/Users/username/Library/Application Support/Blackmagic Design/DaVinci Resolve/Fusion/Templates/Edit/Titles

On Windows: C:\Users\username\AppData\Roaming\Blackmagic Design\DaVinci Resolve\Support\Fusion\Templates\Edit\Titles

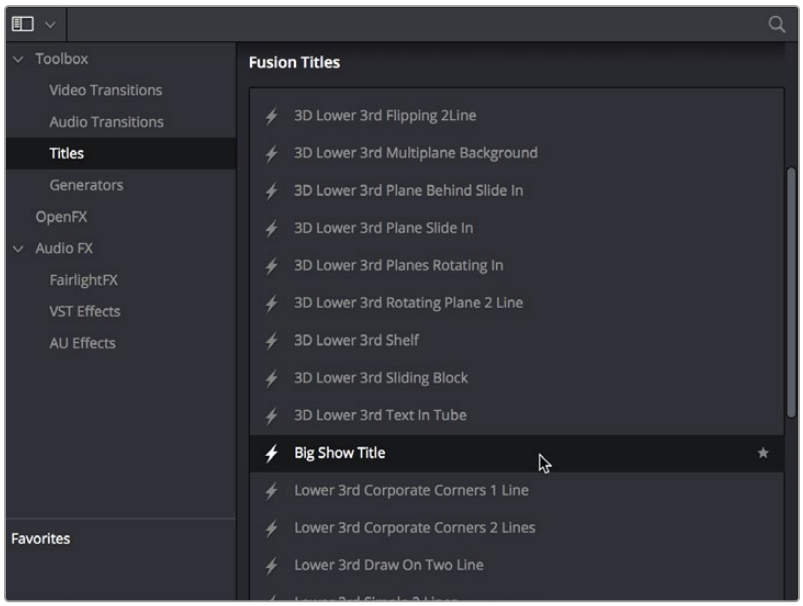
On Linux: home/username/.local/share/DaVinciResolve/Fusion/Templates/Edit/Titles



Choosing where to save a title template for the Edit page in DaVinci Resolve.

Using Your New Title Template

After you've saved your macro, you'll need to quit and reopen DaVinci Resolve. When you open the Effects Library of the Edit page, you should see your new template inside the Titles category, ready to go in the Fusion Titles list.



Custom titles appear in the Fusion Titles section of the Effects Library.

Editing this template into the Timeline and opening the Inspector, we can see the parameters we enabled for editing, and we can use these to customize the template for our own purposes.

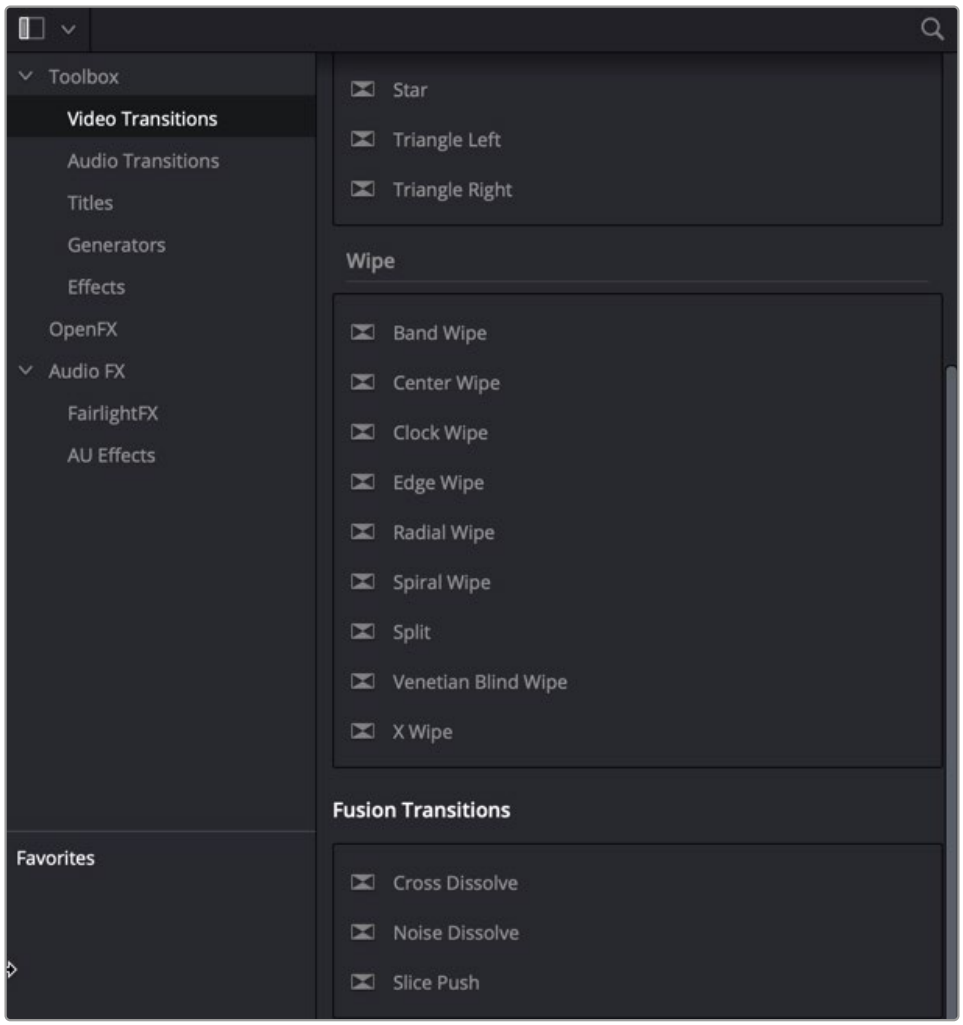


Customizing the template we made.

And that's it!

Getting Started with a Fusion Transition Template

When creating a Fusion transition template, it's easiest to start with an existing transition template and build off that. Three transitions are located in the Fusion Transitions category of the DaVinci Resolve Effects Library. The simplest transition is the Cross Dissolve, while the most complex example is the Slice Push.

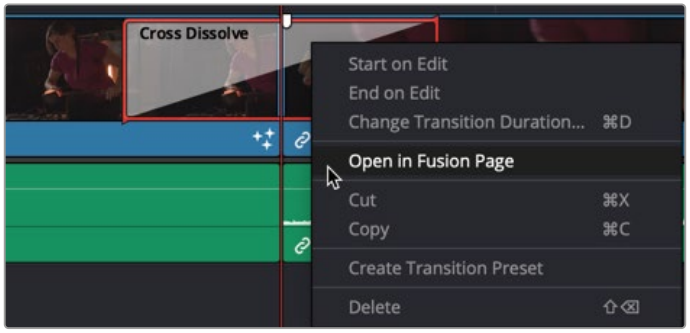


The Fusion transition templates located in DaVinci Resolve's Effects Library.

Creating a Fusion Transition Template

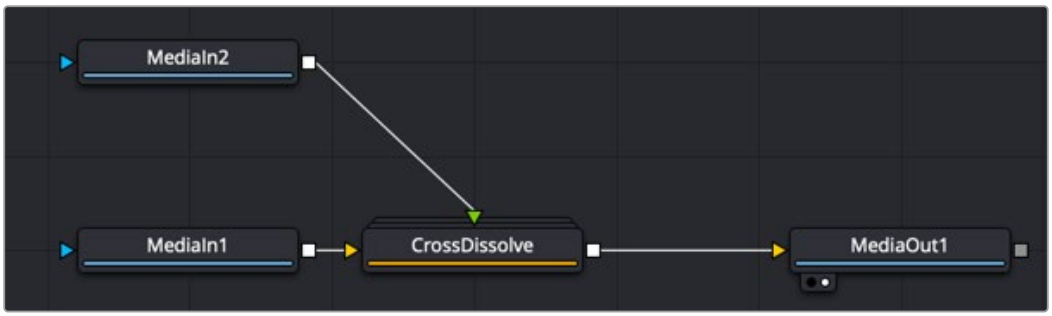
The three Fusion Transitions located in the DaVinci Resolve Effects Library are basically Fusion compositions that have been turned into macros. Ordinarily, macros are used as building blocks inside of Fusion so that you can turn frequently-made compositing tricks that you use all the time into your own nodes. However, you can also use this macro functionality to build transition templates for the DaVinci Resolve Edit page.

Once you apply a Fusion Transition in the Edit page, you can right-click it and choose Open in Fusion Page.



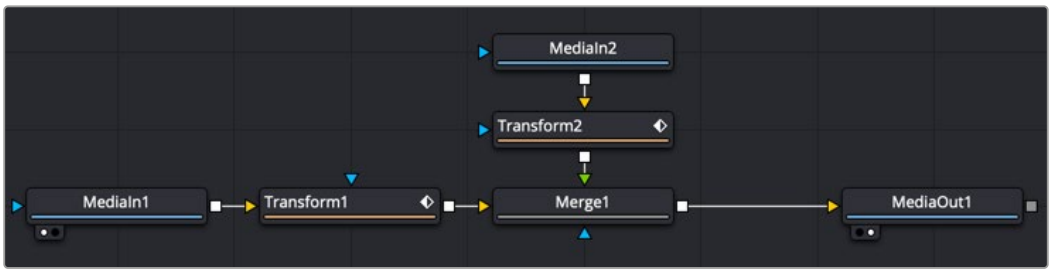
Right-clicking over a Fusion Transition in the DaVinci Resolve Edit page.

The Fusion page opens, displaying the node tree used to create the Fusion transition.



The Cross Dissolve node tree in Fusion.

The MediaIn 1 node represents the outgoing clip in the Edit page Timeline. The MediaIn 2 clip represents the incoming clip. You can modify or completely change the Cross Dissolve effect to create your own custom transition using any of Fusion's nodes.



The Fusion Cross Dissolve node tree replaced with Transforms and a Merge node.

TIP: To modify the duration of the Fusion transition from the Edit page Timeline, you must apply the Resolve Parameter Modifier to any animated parameter. In place of keyframing the transition, you create the transition using the Scale and Offset parameters of the Resolve parameter modifier.

Updating a Fusion Transition

After modifying the transition in Fusion, you can choose to update the transition in the Timeline or create a new transition, which you save into the Edit page Effects Library. To update the transition, just return to the Edit page. The transition in the Edit page Timeline reflects the changes you make in the Fusion page.

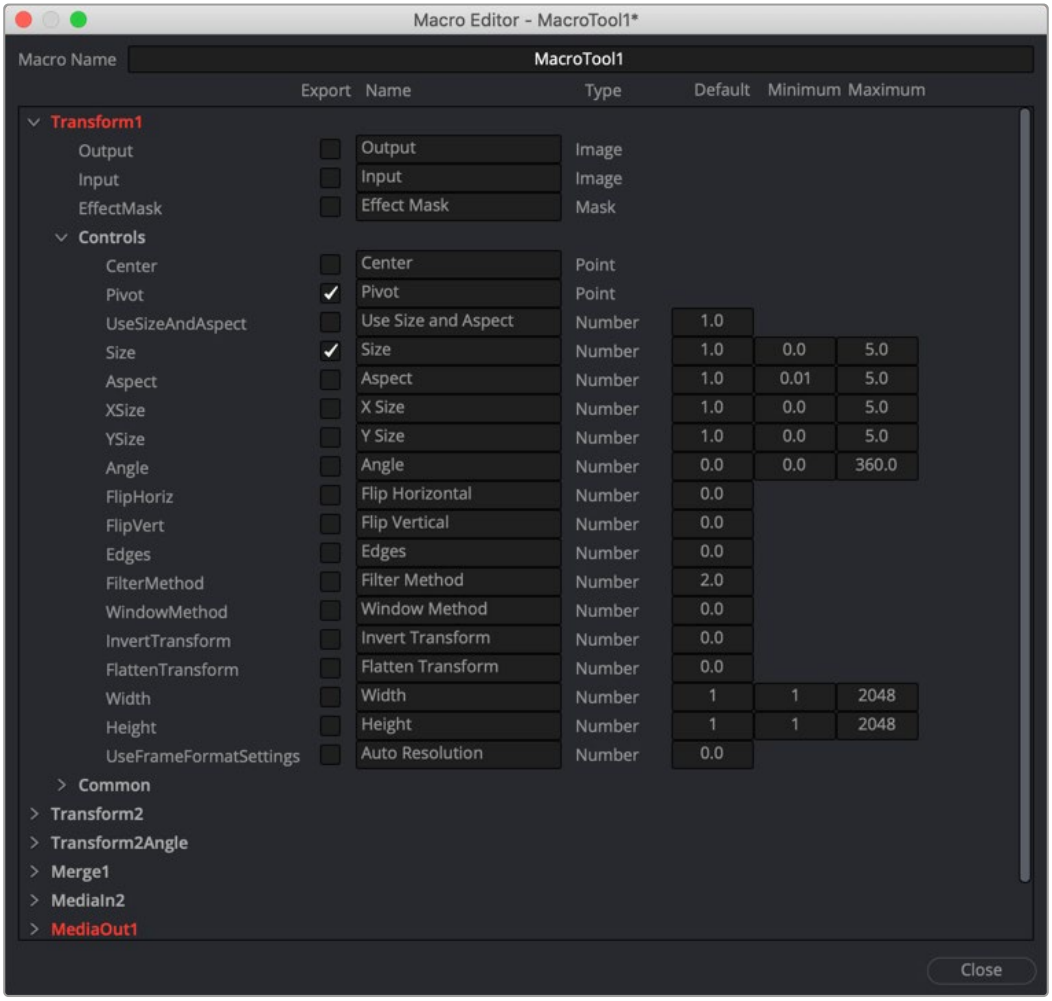
Saving a New Fusion Transition

If after modifying the transition in Fusion, you need to save it to the Effects Library to reuse it on other transitions or other projects, you must make a macro and save it to the Transitions folder.

Start by selecting every node in the Node Editor that you want to include in the transition template, including the two MediaIn nodes and the MediaOut node.

> **TIP:** Since the transition template must include the MediaIn and MediaOut nodes, the final steps for saving a transition template must be performed in DaVinci Resolve's Fusion page and cannot be performed in Fusion Studio.

Having made this selection, right-click one of the selected nodes and choose Macro > Create Macro from the contextual menu.



The Macro Editor displaying the parameters of all the nodes you selected.

The Macro Editor window appears, displaying a hierarchical list of every parameter in the composition you've just selected. The order of nodes is based on the order they were selected in the Node Editor prior to creating the macro.

The Macro Editor is designed to let you choose which parameters you want to display as custom controls in the Edit page Inspector when the transition is applied.

For transitions, you can choose not to display any controls in the Inspector, allowing only duration adjustments in the Timeline. However, you can choose a simplified set of parameters for customization by enabling the checkboxes next to any parameter name.

Once you enable all the parameters you want to use in the eventual template, click the Close button, and a Save Macro As dialog appears. Here, you can enter the name of the transition, as it should appear in the Edit page Effects Library.

To have the transition template appear in the Effects Library > Fusion Transitions category of DaVinci Resolve, save the macro in the following locations:

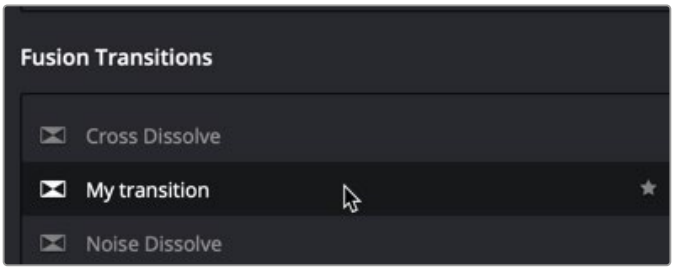
On macOS: Macintosh HD/Users/username/Library/Application Support/Blackmagic Design/DaVinci Resolve/Fusion/Templates/Edit/Transitions

On Windows: C:\Users\username\AppData\Roaming\Blackmagic Design\DaVinci Resolve\Support\Fusion\Templates\Edit\Transitions

On Linux: home/username/.local/share/DaVinciResolve/Fusion/Templates/Edit/Transitions

Using Your New Transition Template

After you've saved your macro, you'll need to quit and reopen DaVinci Resolve. When you open the Effects Library on the Edit page, the new transition template is listed in the Video Transitions category, in the Fusion Transitions list.

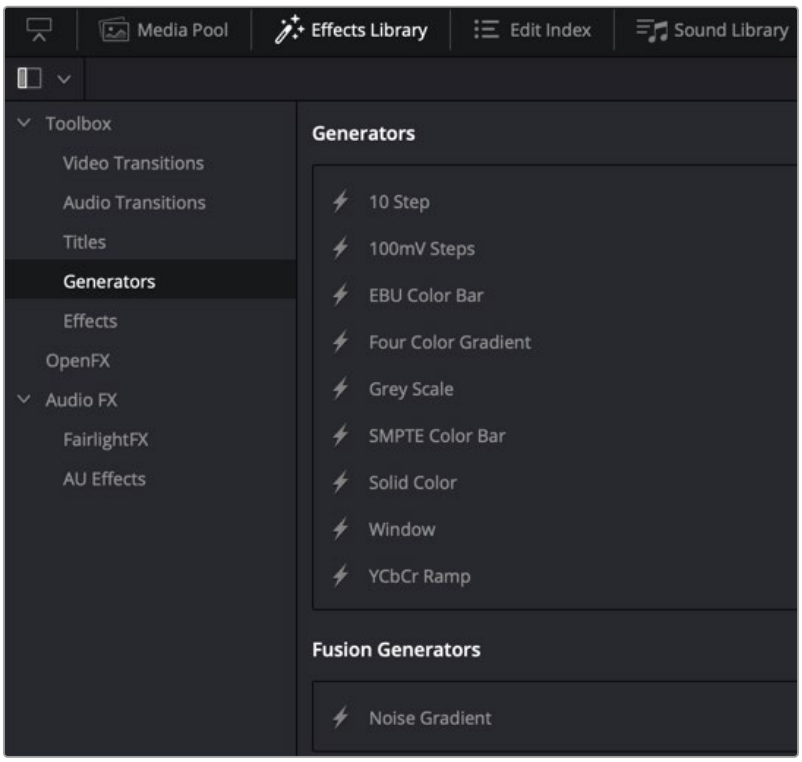


A custom Fusion Transition saved in the Edit page Effects Library.

Applying this transition to a cut in the Timeline and opening the Inspector shows the parameters you enabled for editing, if any.

Getting Started with a Fusion Generator Template

There is one simple Noise Gradient Fusion Generator located in the Effects Library that you can use as a starting point for creating your own generators.



The Fusion Generator templates located in DaVinci Resolve's Effects Library.

To open the Fusion Noise Gradient Generator in the Fusion page, do the following:

1. On the Edit page, drag the Fusion Noise Gradient Generator from the Effects Library to the Timeline.
2. Right-click over the Noise Gradient Generator and choose Open in Fusion Page from the pop-up menu.

The Fusion page opens, displaying the node tree that is used to create the Fusion Generator.

Creating a Fusion Generator Template

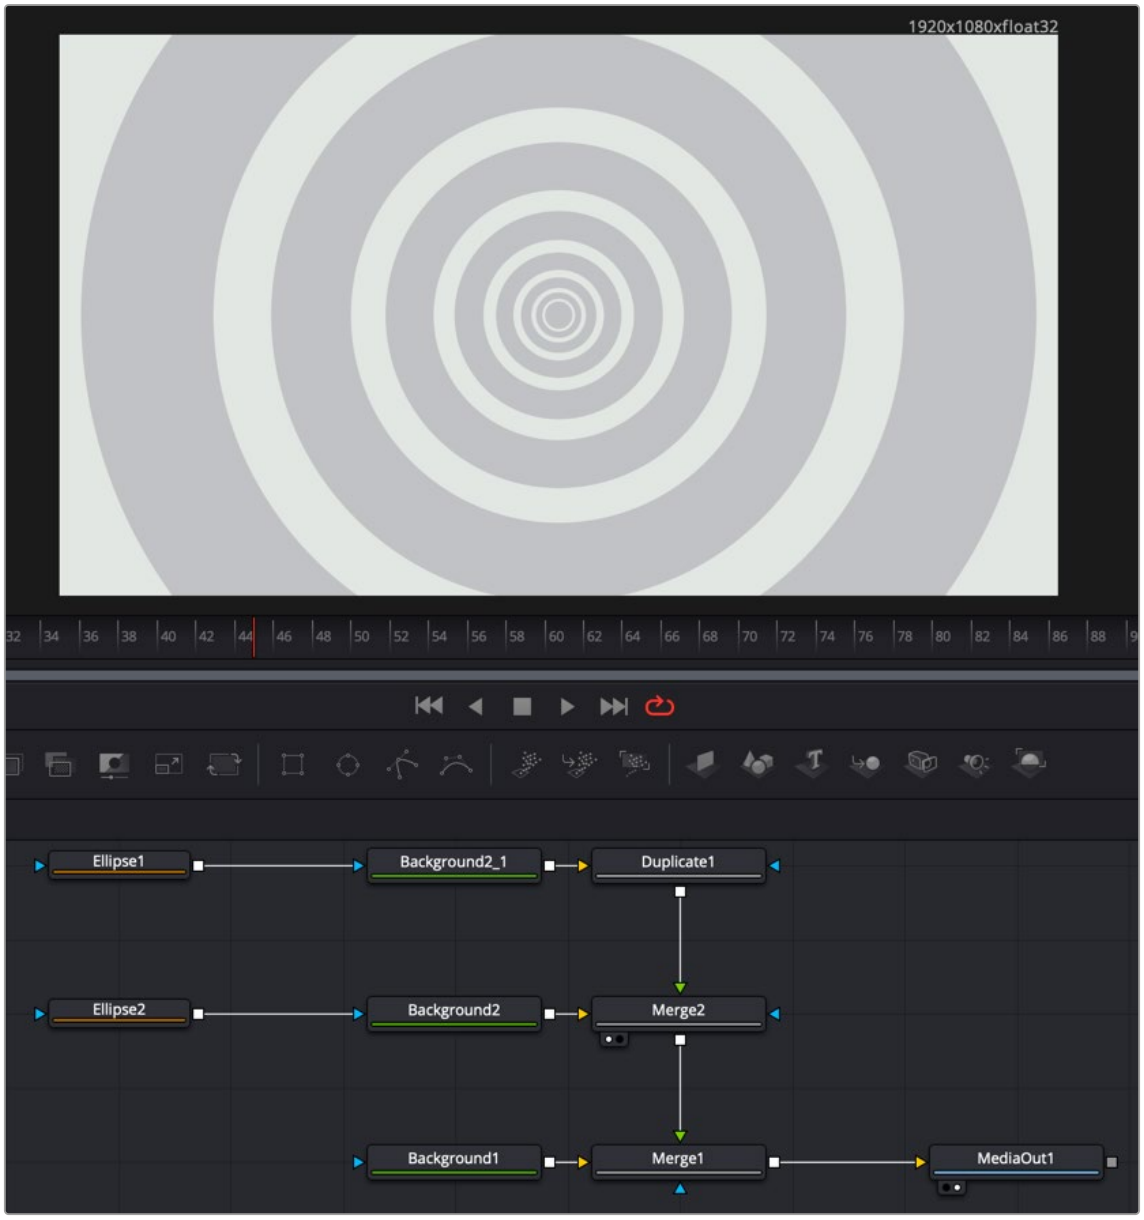
As easy as it is to begin with the Noise Gradient Generator template, you can just as easily start by adding a Fusion Composition Effect to a Timeline in the Edit page.

To begin creating a Fusion Generator Template with an empty Fusion composition, do the following:

1. On the Edit page, drag the Fusion Composition Effect from the Effects Library to the Timeline.
2. Right-click over the Composition Effect and choose Open in Fusion Page from the pop-up menu.

An empty Fusion page with a single MediaOut node opens, ready for you to create a Fusion Generator.

The Fusion Generator is a solid image generated from any number of tools combined to create a static or animated background. You can choose to combine gradient colors, masks, paint strokes, or particles in 2D or 3D to create the background generator you want.



The Fusion Generator node tree creating concentric circles.

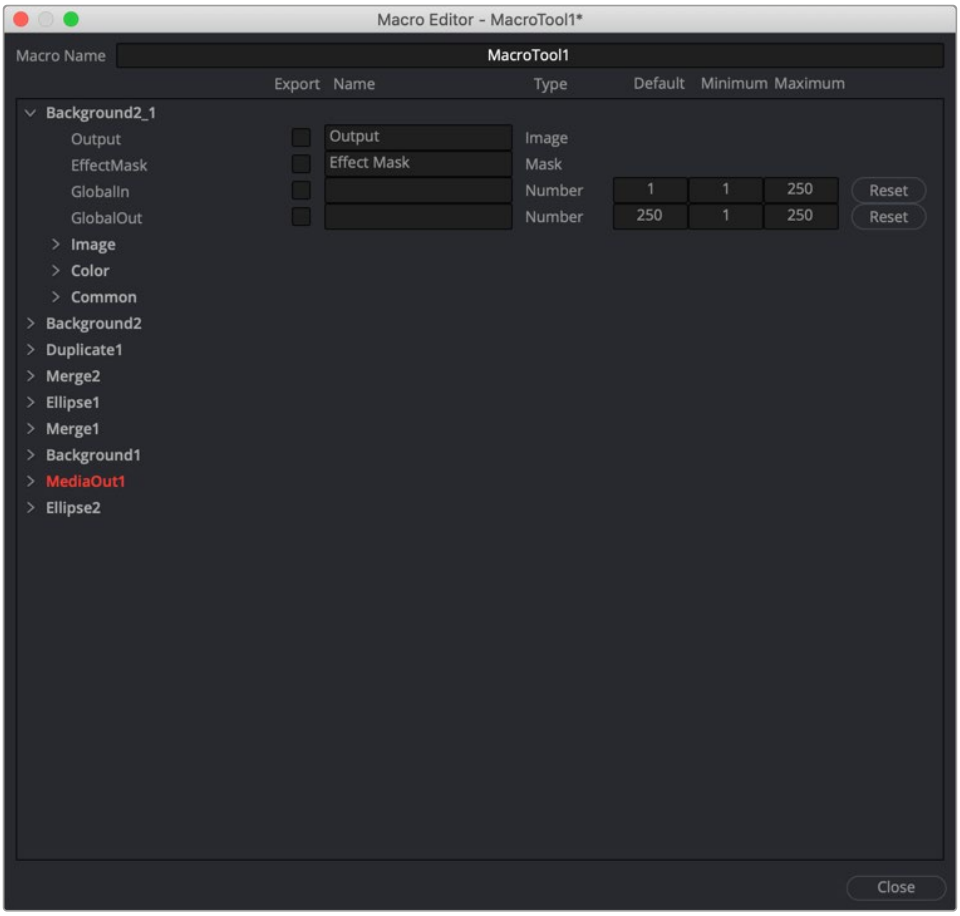
Saving a New Fusion Generator

After creating the generator you want in Fusion, you need to save it to the Effects Library to reuse it in other projects from the Edit page. To do this, you must create a Macro and save it to the Generator folder.

Ordinarily, macros are used as building blocks inside of Fusion so that you can turn frequently-made compositing tricks that you use all the time into your own nodes. However, you also use this macro functionality to build Generator templates for the DaVinci Resolve Edit page.

Start by selecting every node in the Node Editor that you want to include in the Generator template including the MediaOut node.

Having made this selection, right-click one of the selected nodes and choose Macro > Create Macro from the contextual menu.



The Macro Editor displaying the parameters of all the nodes you selected.

The Macro Editor window appears, displaying a hierarchical list of every parameter in the composition you've just selected. The order of nodes is based on the order they were selected in the Node Editor prior to creating the macro.

The Macro Editor is designed to let you choose which parameters you want to display as custom controls in the Edit page Inspector when the Generator is applied. You can choose a simplified set of parameters for customization by enabling the checkboxes next to any parameter name.

Once you enable all the parameters you want to use in the eventual template, click the Close button, and a Save Macro As dialog appears. Here, you can enter the name of the Transition, as it should appear in the Edit page Effects Library.

To have the Generator template appear in the Effects Library > Fusion Generators category of DaVinci Resolve, save the macro in the following locations:

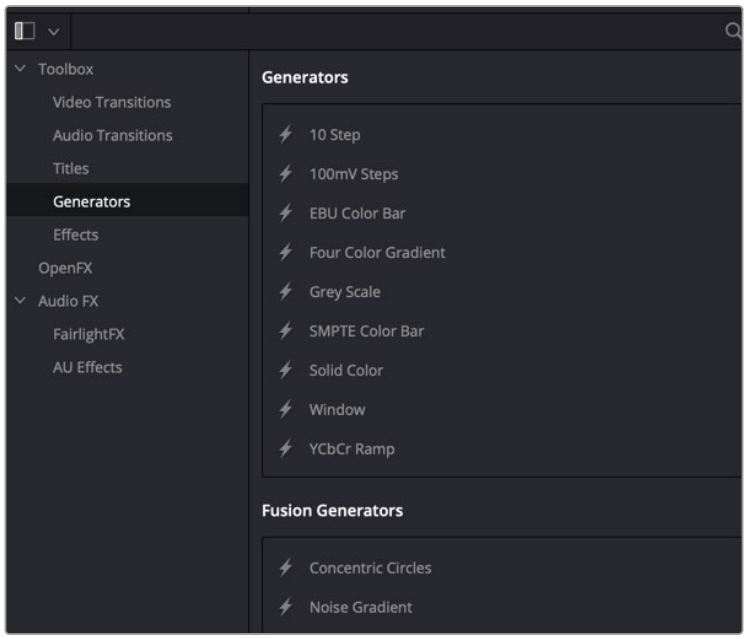
On macOS: Macintosh HD/Users/username/Library/Application Support/Blackmagic Design/DaVinci Resolve/Fusion/Templates/Edit/Generators

On Windows: C:\Users\username\AppData\Roaming\Blackmagic Design\DaVinci Resolve\Support\Fusion\Templates\Edit\Generators

On Linux: home/username/.local/share/DaVinciResolve/Fusion/Templates/Edit/Generators

Using Your New Generator Template

After you've saved your macro, you'll need to quit and reopen DaVinci Resolve. When you open the Effects Library on the Edit page, the new Generator template is listed in the Generators category, in the Fusion Generators list.



Custom Fusion Generator saved in the Edit page Effects Library.

Applying this Generator to the Timeline and opening the Inspector shows the parameters you enabled for editing, if any.

Creating a Fusion Effect Template

You start building a Fusion Effect template by bringing a clip from the Edit page Timeline into Fusion. This clip is only used for creating the template and will not be saved with the effect.

Once inside Fusion, you use Fusion's nodes to create the effect you want. You can use a single node or a hundred, depending on the effect you want to create. For instance, using Fusion's Color Correction nodes, you can create a simple color corrector you can use on the Edit page.

To create a simple Color Corrector effect, do the following:

1. Insert the Color Corrector node between the MediaIn and MediaOut nodes.
2. Select the Color Corrector node in the Node Editor, and then press Cmd-A to select the remaining nodes.
3. Right-click over any of the selected node and choose Macro > Create Macros from the contextual menu. Enabling the checkboxes in this window determines the parameters that appear in the Edit page Inspector.
4. The Macro Editor window opens. Here, you can enable the checkboxes for any parameters you want to be shown in the Edit page Inspector.
5. Enter the name of your effect at the top of the Macro Editor window.
6. To save the Macro, click Close at the bottom of the window, then click Yes in the dialog that appears asking you to save the changes.

The Macros must be saved into the correct folder for DaVinci Resolve to recognize the Macro as an effect.

In the save dialog, save the Macro in the following locations:

On macOS: Macintosh HD/Users/username/Library/Application Support/Blackmagic Design/DaVinci Resolve/Fusion/Templates/Edit/Effects

On Windows: C:\Users\username\AppData\Roaming\Blackmagic Design\DaVinci Resolve\Support\Fusion\Templates\Edit\Effects

On Linux: home/username/.local/share/DaVinciResolve/Fusion/Templates/Edit/Effects

You can save and organize your Fusion Effects into separate subfolders underneath the paths above. These subfolders will show up in the Effects section in the Edit page.

To see the effect in the Edit page Effects Library, you'll need to quit DaVinci Resolve and relaunch the application.

Creating a Fusion Effect Template for Two or More Layers

If the effect you want to create requires multiple images like a video wall, you start by creating a Fusion clip on the Edit page Timeline that includes the number of layers you want the effect to have. The clips are only used to create the number of image inputs for the template and will not be saved with the effect.

Once inside Fusion, use Fusion's nodes to create the effect you want.

Save the Macro following the same steps you use for single clip effects. Enable any of the parameters you want to control in the Edit page. To be able to switch the order of video layers within the effect, make sure you have the Layer checkbox enabled for all the MediaIn nodes.

Once you've saved the Macro and relaunched DaVinci Resolve, to use the effect on multiple timeline layers, you must create a Fusion clip. The Fusion clip should contain the same number of layers the effect requires. The order of the Timeline layers, going from the bottom track to the top, matches the MediaIn numbers. For instance, video track 1 will match the position and appearance of MediaIn1, video track 2 matches MediaIn2 and so on. If you want to change how tracks map to MediaIn nodes, you can change the Layer number in the Inspector, assuming you enabled the MediaIn Layer checkbox when creating the Macro.

Changing Durations of a Template

After you make a template in Fusion, you may want to change its duration in the Edit or Cut page Timeline. Changing the duration when animation is involved can be complicated, so there are two Modifiers in Fusion that can help determine how keyframes react when the duration is updated in the Edit or Cut page Timeline.

Anim Curves Modifier

The Animation Curves Modifier (Anim Curves) is used to dynamically adjust the timing, values, and acceleration of an animation, even if you decide to change the duration of a Comp. Using this Modifier, it becomes infinitely easier to stretch or squish animations, create smooth motion, add bouncing properties, or mirror animations curves without the complexity of manually adjusting splines.

When creating Fusion templates for the Edit or Cut page in DaVinci Resolve, the Anim Curves Modifier allows the keyframed animation you've created in Fusion to stretch and squish appropriately as the transition, title, or effect's duration changes on the Edit and Cut page Timelines.

Keyframe Stretcher Modifier

The Keyframe Stretcher modifier is primarily used when creating title templates in Fusion for use in DaVinci Resolve's Edit or Cut page. The Keyframe Stretcher modifier is added to any animated parameter so that the Hold keyframes between the initial animation on screen and the final animation off-screen stretch when the template is trimmed in the Timeline. This allows the animation to retain its timing while the static portion of the title stretches to meet the new duration requirements.

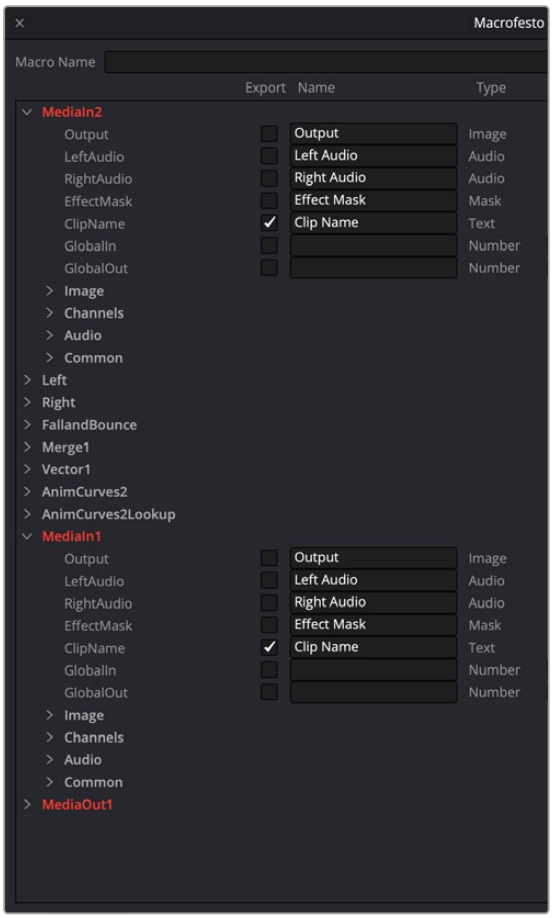
Creating Media Drop Zones in Templates

For additional ease of use in the Edit page, Fusion effects, plugins, and transitions can have media drop zones, allowing the user to simply drag clips from the Media Pool directly to the template. For example, you can instantly replace the background of a Fusion effect with a star-field, just by dragging that clip from the Media Pool to the media drop zone in the Inspector.

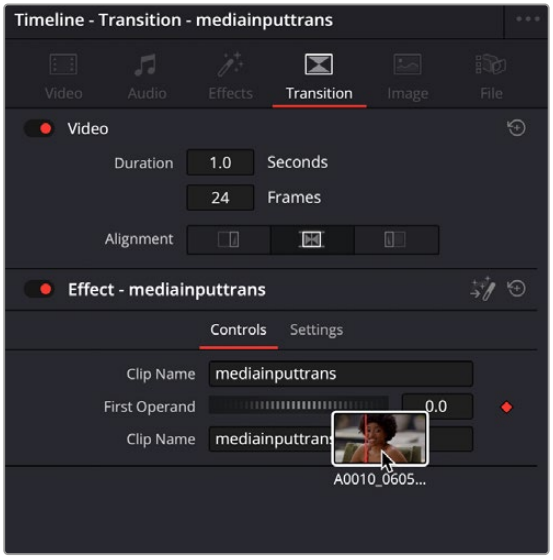
To create a media drop zone in a Fusion template:

1. Create a Macro of your current Fusion template.
2. For any MediaIn node that you want to make a media drop zone, make sure the ClipName parameter is checked.
3. Place your template in the appropriate Fusion Edit template folder (Transition, Effect, etc.), and relaunch DaVinci Resolve.

   When you use this template in the Edit page, now you can drag clips from the Edit page Media Pool and drop them in the Transition Inspector's Clip Name field. They will then instantly update the template with the chosen media.



The Clip Name parameters are checked for MediaIn2 and MediaIn1 nodes in the Macro Editor, allowing media drop zones for both incoming and outgoing sides of this Fusion transition.



A clip from the Media Pool being dropped on a Clip Name field in the resulting Transition Inspector in the Edit page

Creating a Custom Template Icon

It is possible to create a custom icon that is embedded with your template thumbnail in the Effects Library, instead of the default first three letters of the template name.

To create a simple Custom Template Icon:

1. Create a .png file of the icon you want to use for the template. The recommended size is 104 x 58 pixels, but any image will be resized to fit.
2. Name the file exactly the same name as your template, just with a .png extension rather than a .setting extension.
3. Place the .png image in the same directory as your template.

When you restart DaVinci Resolve, the icon you created will be embedded in the template thumbnail across all the Effects Libraries in the program.



A Custom Icon added to a fisheye template, before (above) and after (below)

Using Fusion Template Bundles

For ease in distributing templates to other Fusion users, multiple templates can now be bundled together into a single .drfx file. This file can then be imported back into another Fusion workstation easily to ensure that all custom templates are the same between computers.

Creating a Fusion Template Bundle requires using a specific directory structure and using your operating systems file browser and .zip compression utility. The directory structure is listed below, and you can always find a specific folder from within Fusion by right-clicking on any bin in the Effect Library and selecting Show Folder.

Folder structure for Templates used in the Edit page:

- Edit
- Effects
- Generators
- Titles
- Transitions

Folder structure for Templates used in the Fusion page:

- Fusion

To create a Fusion Template Bundle:

1. In your OS, create a folder structure above that includes the specific folder for the type of templates you want to be in the bundle. For example, if you have a transition and two effects templates, you would create an Edit folder, and two subfolders inside it named Transitions and Effects.
2. Copy your template (.setting) files into the appropriate directories. You can also include icon files and any associated assets as well.
3. Use your OS zip compression utility to create a .zip file of the directory structure.
4. Rename the .zip file in your OS with the ".drfx" extension instead of .zip. The file icon should change to reflect the new extension.

To import a Fusion Template Bundle:

1. Double-click on a .drfx file in your OS. DaVinci Resolve will launch and a dialogue box will appear asking if you want to install the template bundle.
2. Drag to the .drfx file from your OS directly into the Fusion page in DaVinci Resolve. A dialogue box will appear asking if you want to install the template bundle.

To delete a Fusion Template Bundle:

1. Navigate to the appropriate template directory in your file browser.
2. Delete the .drfx file.

IMPORTANT: The Fusion Template Bundle contains all the templates in one file. It does not uncompress them into separate template files again. Therefore if you delete the .drfx file, all associated templates inside that bundle will be removed as well.

Chapter 7

Using Viewers

This chapter covers working with viewers in Fusion, including using onscreen controls and toolbars, creating groups and subviews, managing viewer Lookup Tables (LUTs), working with the 3D viewer, and setting up viewer preferences and options.

Contents

Viewer Overview 188
Single vs. Dual Viewers 188
Floating Viewers in Fusion Studio 189
Video Output 189
Clean Feed 189
Loading Nodes into Viewers 189
Clearing Viewers 190
Position and Layout 190
The Viewer Divider 190
Zooming and Panning into Viewers 191
Flipbook Previews 191
Creating Flipbook Previews 191
Playing Flipbook Previews 192
Removing Flipbook Previews 193
Flipbook Preview Render Settings 193
Onscreen Controls 194
Showing and Hiding Onscreen Controls 195
Making Fine Adjustments to Onscreen Controls 195
Toolbars 195
Viewer Toolbar 195
Node Toolbars 196
A/B Buffers 196
Flipping between Buffers 196
Split Wipes between Buffers 197
Moving the Wipe Divider 198

Subviews 198
Showing and Hiding Subviews 199
Changing the Subview Type 199
Swapping the Subview with the Main View 199
Viewer and Subview Types 200
Viewing Selective Channels 204
Viewing Color Channels 204
Viewing Auxiliary Channels 205
Viewing Multilayer Files (Studio Version Only) 205
The 3D Viewer 207
Panning, Scaling, and Rotating a 3D Viewer 207
Viewing Objects via Wireframe 207
Changing the POV of a 3D Viewer 208
Copying a Viewer's POV to a Camera 208
Lighting and Shadows in 3D Viewers 209
Transparency in 3D Viewers 210
Grid 211
Vertex Normals 211
Quad View 212
Quad View Layouts 213
Using Quad Views for 2D Scenes 213
Guides 213
Frame Format Settings 214
Domain of Definition and Region of Interest 215
Domain of Definition (DoD) 215
Region of Interest (RoI) 216
Enabling RoI Controls 216
Managing Viewer Lookup Tables (LUTs) 217
How Lookup Tables Work in Fusion 218
Types of Viewer LUTs 218
Using Viewer LUTs 220
Editing Viewer LUTs 221
LUT Processing Order 223
Applying Multiple LUTs 223
Saving Custom LUTs 224
LUT Files 225
Viewer Preferences and Settings 226
The Viewer Options Menu 226
Locking the Viewer (Command-L) 227
Additional Viewer Options 227
Status Bar Information 228

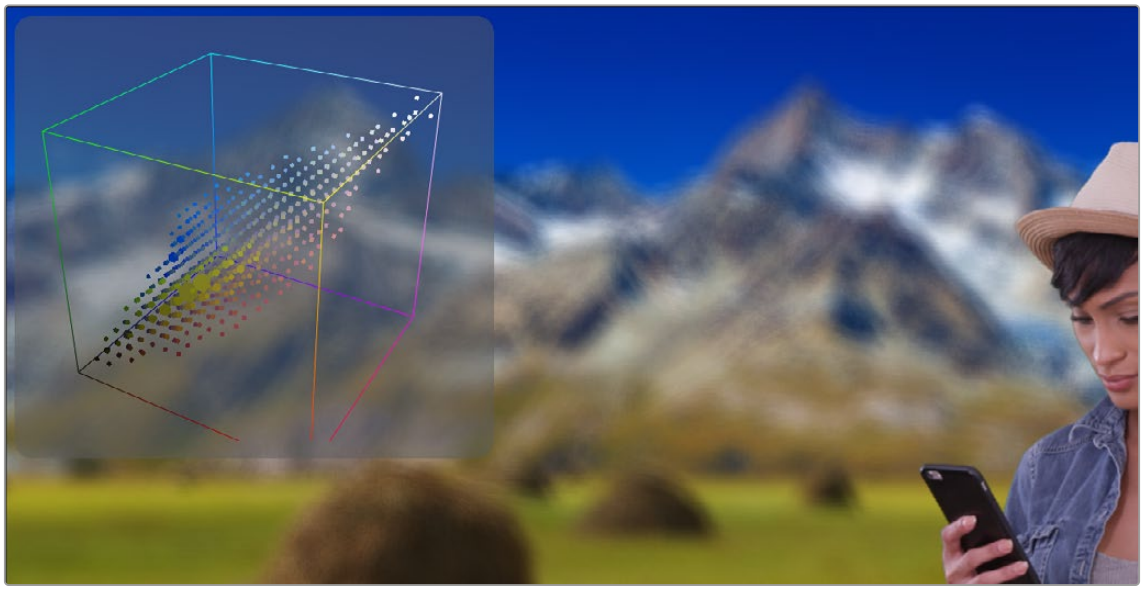
Viewer Overview

Viewers in Fusion display the current frame of the current composition in a variety of ways to help you see what you're doing and evaluate the final result of your compositing artistry. Viewers display 2D images, but they can also display a 3D environment using a 3D View as well as a special Quad viewer to help you effectively work in three dimensions.



Side-by-side dual viewers: a 3D viewer (left), and a 2D viewer (right)

Additionally, you can expose "subviews" including color inspectors, magnifiers, waveforms, histograms, and vectorscopes to help you analyze the image as you work.



Viewer with a 3D Histogram subview at the upper left-hand corner

Single vs. Dual Viewers

By default, there are two viewers positioned side by side across the top of the window. However, you can use the Single/Dual Viewer button to toggle between displaying a single viewer or two viewers displayed side by side.



The Single/Dual Viewer toggle button

Floating Viewers in Fusion Studio

In addition to the dual views above the Node Editor, Fusion Studio provides the option to use an unlimited number of floating viewers. These floating viewers are excellent for taking full advantage of a dual monitor configuration. Floating viewers can also be set to full-screen mode to make the best use of screen real estate.

To create a new floating display view, select Window > New View from the menu bar at the top of the screen. The position and configuration of the new view can be saved in the layout preferences, if required.

Video Output

When using DaVinci Resolve or Fusion Studio, if Blackmagic video hardware is present in the computer, then you can select a node to preview directly on that display. While video output can't be used for manipulating onscreen controls such as center crosshairs or spline control points, they're extremely valuable for evaluating your composition via the output format, and for determining image accuracy using a properly calibrated display.

The video hardware is configured from the DaVinci Resolve and Fusion Studio preferences.

Clean Feed

When using DaVinci Resolve with dual computer monitors, a full-screen viewer can be displayed on the secondary monitor from the Fusion page. This displays a third view indicator button under each node to control what is shown on the second display. To activate this monitor, make sure you do not have Dual Screen enabled under the Workspace menu and then select Workspace > Video Clean Feed and select your second computer display from the submenu.

Loading Nodes into Viewers

The Fusion page in DaVinci Resolve and Fusion Studio show two different things when you first open each application. When you first open the Fusion page, the output of the current empty composition (the MediaOut1 node) is usually showing in viewer 2. If you're in dual-viewer mode, viewer 1 remains empty until you assign a node to one of them. In Fusion Studio, because there are no nodes when you first begin a comp, nothing is displayed in the viewers.

To load specific nodes into specific viewers, do one of the following:

- Hover the pointer over a node, and click one of two buttons that appear at the bottom-left of the node.
- Click once to select a node, and press 1 (for the left viewer) or 2 (for the right viewer).
- Right-click a node and choose View On > None/Left View/Right View in the contextual menu.
- Right-click the control header of a node in the Inspector, and choose View On > None/Left View/Right View from the contextual menu.
- Drag a node and drop it over the viewer you'd like to load it into (this is great for tablet users).

When a node is being viewed, a View Indicator button appears at the bottom left. This is the same control that appears when you hover the pointer over a node. Not only does this control let you know which nodes are loaded into which viewer, but they also expose little round buttons for changing which viewer they appear in.



Viewer assignment buttons at the bottom left of nodes indicate when they're being viewed, and which dot is highlighted indicates which viewer that node is loaded into.

Clearing Viewers

To clear an image from a viewer, click in the viewer to make it active; a light purple outline is displayed around the active panel. With the viewer active, press the ` (accent) key. This key is usually found to the left of the 1 key on U.S. keyboards. The fastest way to remove all images from all viewers is to make sure none of the viewers is the active panel, and then press the accent key.

Position and Layout

In Fusion Studio, when you resize and change the layout of viewers in the composition, that configuration is always saved with the composition. So each time you open the composition, the size and layout is remembered. You can prevent this behavior by disabling the Recall Layout checkbox in the Fusion Global Layout preferences.

If you want all new compositions to open with a certain viewer layout, you can configure the layout of the two primary viewers, and then use the Grab Document Layout button in the Fusion Global Layout preferences to remember the layout for any new compositions. To save the position and size of floating viewers, you use the Grab Program Layout button. Finally, if you want to have the floating viewers opened automatically when you open Fusion, enable the Create Floating Views checkbox.

The Viewer Divider

You can change the relative sizes of the left and right viewers using the horizontal viewer divider that runs between them. Drag the viewer divider to increase or decrease the amount of space used by one viewer. The adjacent viewer will adjust to accommodate the new layout.

The amount of vertical space available for both viewers can be adjusted by dragging the horizontal scrollbar between the viewers and the work area below them.



The viewer divider bar

Zooming and Panning into Viewers

There are standardized methods of zooming into and panning around viewers when you need a closer look at the situation. These methods also work with the Node Editor, Spline Editor, and Keyframes Editor.

Methods of panning viewers:

- Middle-click and drag to pan around the viewer.
- Hold down Shift and Command and drag the viewer to pan.
- Drag two fingers on a track pad to pan.

Methods of scaling viewers:

- Click a viewer and press the Equals key (=) to zoom in, and the Minus key (-) to zoom out.
- Press the middle and left mouse buttons simultaneously and drag left or right to resize the viewer.
- Hold down the Command key and use your pointer's scroll control to resize the viewer.
- Hold down the Command key and drag two fingers on a track pad to resize the viewer.
- Hold down the middle mouse button, and then click the left mouse button to zoom in, or click the right button to zoom out. The scaling uses a fixed amount, centered on the position of the cursor.
- Click a viewer and press Command-1 to resize the image in the viewer to 100 percent.
- Click a viewer and press Command-2 to resize the image in the viewer to 200 percent.
- Click a viewer and press Command-F or Command-1 to reset the image in the viewer to fit the viewer.
- Click the Scale Viewer menu and choose Fit or a percentage.
- Right-click on a viewer and choose an option from the Scale submenu of the contextual menu. This includes a Custom Scale command that lets you type your own scale percentage

Methods of spinning 3D viewers:

- In 3D Perspective view, hold down the middle mouse button and the right mouse button and drag to spin the stage around.

Flipbook Previews

As you build increasingly complex compositions, and you find yourself needing to preview specific branches of your node tree to get a sense of how various details you're working on are looking, you may find it useful to create targeted RAM previews at various levels of quality right in the viewer by creating a RAM Flipbook. RAM Flipbook Previews are preview renders that exist entirely within RAM and allow you to render a node's output at differing levels of quality for quick processing in order to watch a real-time preview.

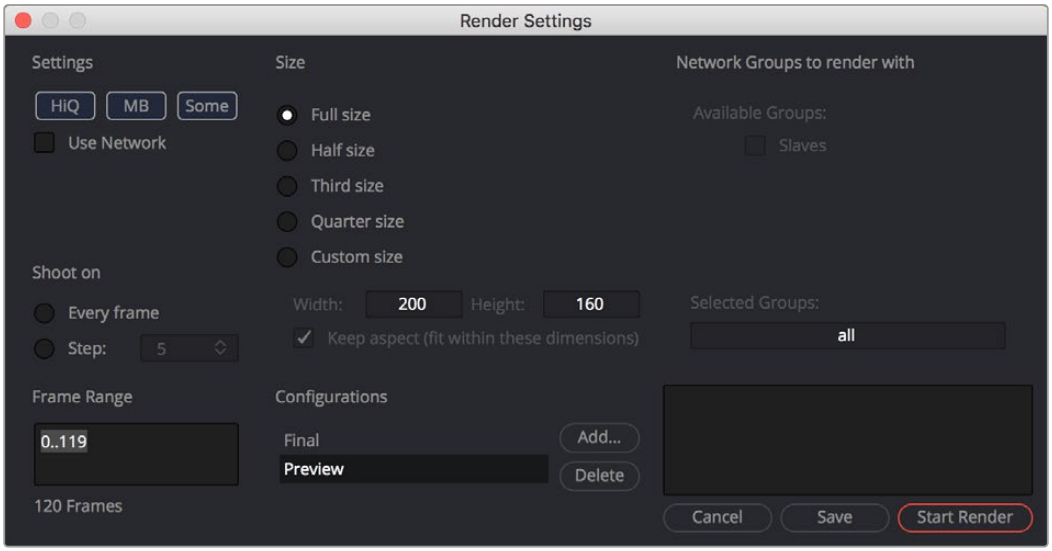
Creating Flipbook Previews

Creating a Flipbook Preview is relatively fast, once you know where to look.

To create a Flipbook Preview:

1. Choose the node in your node tree that you want to preview by doing one of the following:
   - Hold down the Option key while dragging a node into the viewer.
   - Right-click a node and choose an option from the Create/Play Preview submenu in the contextual menu.

2 When the Preview Render dialog opens, choose the quality, resolution, and motion blur settings you want to use for the Flipbook Preview.



The Flipbook Preview Render dialog

3 When you've chosen the settings you want to use, click Start Render.

The current frame range of the Time Ruler is rendered using the settings you've selected, and the result is viewable in the viewer you selected or dragged into.

Once you've created a Flipbook Preview within a particular viewer, right-clicking that viewer presents Flipbook-specific commands and options to Play, Loop, or Ping-Pong the Flipbook, to open it Full Screen, to Show Frame Numbers, and to eliminate it.

TIP: If you want to create a Flipbook Preview and bypass the Render Settings dialog by just using either the default setting or the settings that were chosen last, hold down Shift-Option while you drag a node into the viewer. The Settings dialog will not appear, and rendering the preview will start right away.

Playing Flipbook Previews

While the Flipbook Preview is loaded into a viewer, or open in full-screen mode, you can play or scrub through it using the mouse and the keyboard.

To play back a Flipbook using the mouse, do the following:

- Double-click in the viewer to start playback.

To scrub through a Flipbook using the mouse, do the following:

- Hold down the right mouse button down and drag left or right to scrub through frames.

To play back a Flipbook using the keyboard, do one of the following:

- Press the Spacebar to start or stop playback.
- Hold Shift and press the Spacebar to play in reverse.

To scrub through a Flipbook frame-by-frame using the keyboard, do one of the following:

- Press the Left or Right Arrow keys to move to the previous or next frame.
- Hold Shift and press the Left or Right Arrow keys to jump back or forward 10 frames.
- Press Command-Left Arrow to jump to the first frame.
- Press Command-Right Arrow to jump to the last frame.

TIP: The mouse and keyboard shortcuts work in full-screen mode as well.

Removing Flipbook Previews

Once you create a Flipbook Preview, you need to know how to clear it from RAM.

To eliminate a Flipbook you've created:

- Right-click within a viewer containing a Flipbook Preview, and choose Remove Preview.

Flipbook Preview Render Settings

This section covers all the settings available for rendering Flipbook Previews to RAM.

Settings

The Settings section of the Preview Render dialog includes three buttons that determine the overall quality and appearance of your Flipbook Preview. These buttons also have a significant impact on render times.

HiQ: When enabled, this setting renders the preview in full image quality. If you need to see what the final output of a node would look like, then you would enable the HiQ setting. If you are producing a rough preview to test animation, you can save yourself time by disabling this setting.

MB: The MB in this setting stands for Motion Blur. When enabled, this setting renders with motion blur applied if any node is set to produce motion blur. If you are generating a rough preview and you aren't concerned with the motion blur for animated elements, then you can save yourself time by disabling this setting.

Some: When Some is enabled, only the nodes specifically needed to produce the image of the node you're previewing are rendered.

Size

Since RAM Flipbook Previews use RAM, it's helpful to know how many frames you can render into RAM before you run out of memory. The Flipbook Preview dialog calculates the currently available memory and displays how many frames will fit into RAM. If you have a small amount of RAM in your computer and you cannot render the entire range of frames you want, you can choose to lower the resolution to a setting that delivers the best quality/duration ratio for your preview.

Network

Network rendering is only available in Fusion Studio.

For more information on network rendering, *see Chapter 4, "Rendering Using Saver Nodes,"* in the Fusion Reference Manual.

Shoot On

Sometimes you may not want to render every single frame, but instead every second, third, or fourth frame to save render time and get faster feedback. You can use the Step parameter to determine the interval at which frames are rendered.

Frame Range

This field defaults to the current Render Range In/Out set in the Time Ruler to determine the start and end frames for rendering. You can modify the range to render more or fewer frames.

Configurations

Once you've created a useful preview configuration, you can save it for later use by clicking the Add button, giving it a name, and clicking OK.

Updating a Preview

This option is designed for the interactive frame-by-frame work of rotoscoping and painting. Right-click over a preview in the viewer and choose Update from its contextual menu. When active, any frames that are modified on the previewed node are automatically updated in the preview's playback. This lets you reserve the RAM for playback. You can keep it playing on a loop or ping-pong while you work in another viewer.

Onscreen Controls

When it comes to adjusting images, the Control Panel provides very precise numerical values, but sometimes visually positioning an element using onscreen controls can get you where you want to go with less tweaking.



The Angle preview control

The viewers show onscreen controls for manipulating the parameters of the currently selected node. Common onscreen controls include crosshairs, angle indicators, polylines, and paint strokes. Each of these controls can be manipulated directly in the viewer using the mouse or keyboard.

The controls shown in viewers are determined by which nodes are selected, not by the node displayed in the viewer. For example, a downstream blur is easily viewed while manipulating the controls for a selected polygon mask or merge. If multiple nodes are selected, the controls for every selected node are shown simultaneously.

Showing and Hiding Onscreen Controls

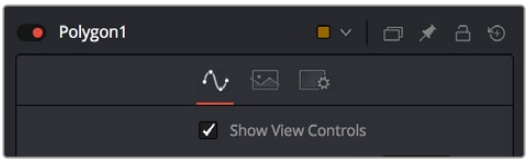
The onscreen controls for a viewer can be hidden so they don't interfere with viewing the image.

To toggle the visibility of onscreen controls, do one of the following:

- Click a viewer's Option menu and choose Show Controls to toggle the controls on or off.
- Right-click in a viewer and choose Options > Show Controls from the contextual menu.
- Select a viewer and press Command-K.

Enabling/Disabling Onscreen Controls in Specific Nodes

Some nodes, like masks, allow disabling of their onscreen controls on a per-node basis, since you often use multiple Polygon nodes to organize and animate masks.



You can disable some nodes, like the Polygon node, on a per-node basis.

Making Fine Adjustments to Onscreen Controls

If you want the visual guidance of onscreen controls with the precision of the Inspector, you can use different keyboard modifiers.

- Up and Down Arrow keys can be used to adjust the vertical position of an onscreen control by small steps.
- Holding down the Command key while using the Up and Down Arrow keys reduces the scale of each step by a factor of ten. Holding Shift increases the scale of each step by a factor of ten.

Toolbars

There are two toolbars in the viewer: a viewer toolbar, which always appears at the top of each viewer and gives you control over what that viewer shows, and an optional node toolbar that appears underneath that gives you contextual controls based on the node you've selected in the Node Editor.

Viewer Toolbar

A viewer toolbar runs across the top of each viewer, providing access to many of the most commonly used viewer-related settings, as well as an indication of the status of many of the most important settings. Most of the menus and buttons found on this toolbar are described in detail throughout this chapter.



The viewer toolbar

Node Toolbars

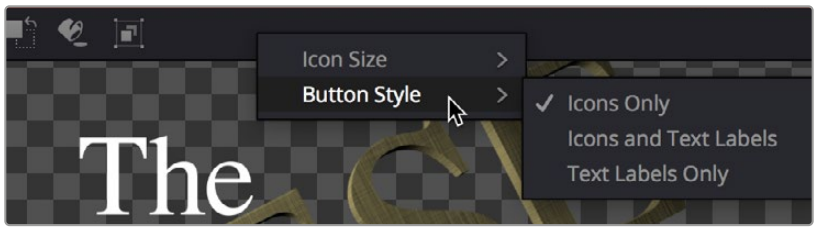
In addition to the viewer toolbar, a node toolbar is displayed underneath, at the top of the viewer display area, whenever you select a node that exposes special nodes. Examples of nodes that expose a toolbar include the text, masks, paths, paint strokes, and the 3D environment.



The node toolbar shown for the Paint node

Customizing the Node Toolbar

If you want to change the size of the buttons that appear in the Node toolbar, or turn on text names for each node, you can right-click anywhere in the empty area of the toolbar and choose new settings from the Icon Size and Button Style submenus in the contextual menu.



The contextual menu for the node toolbar

A/B Buffers

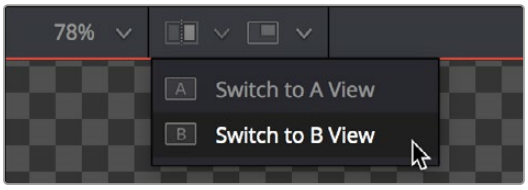
Each viewer has two buffers, each of which can contain images from different nodes, enabling easy comparison of two different nodes within the same viewer by either toggling between buffers, or via an adjustable split-wipe. Each buffer can be considered a complete and separate viewer within the same viewer pane. The A buffer is always shown by default, so when you first load a node into a viewer, the image loads into the A buffer.

Flipping between Buffers

Switching between buffers is easy, either to view a different image while keeping another image handy, or to flip between the original image and the affected image for comparison.

To switch between buffers, do one of the following:

- Select a viewer and press comma (,) to select the A buffer or press period (.) to select the B buffer.
- Click the Buffer menu and choose either Switch to A View or Switch to B View.



The Buffer menu lets you switch between buffers

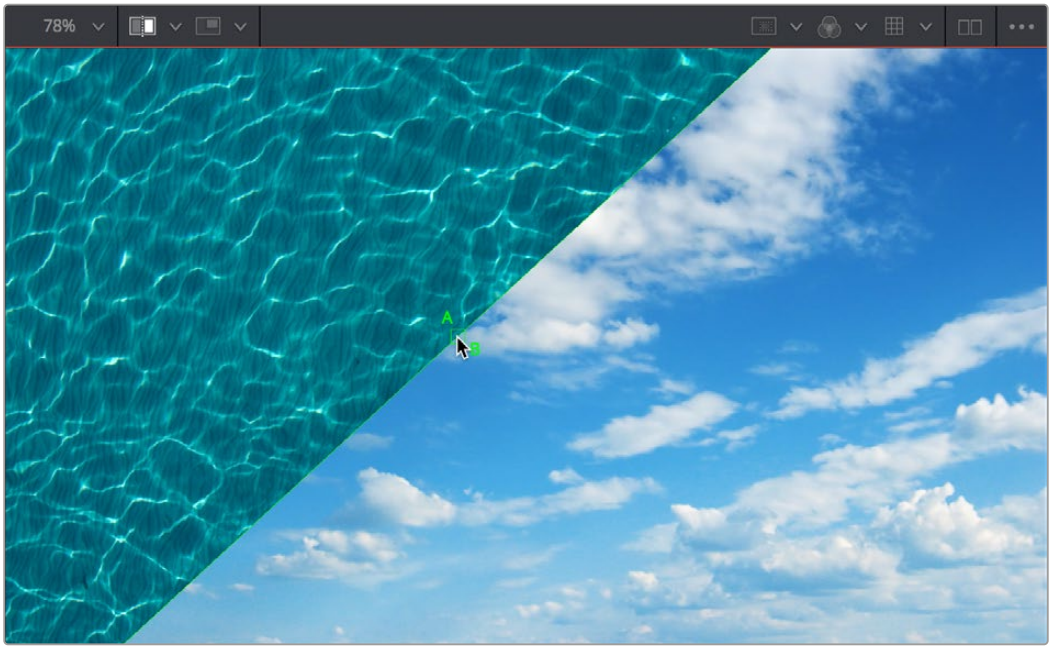
TIP: Each buffer can be set to different display settings—for example, showing different channels or different viewing LUTs, either applied to different nodes or applied to two buffered versions of the same node.

Split Wipes between Buffers

You can also wipe between both buffers, providing a more direct means of comparison.

To wipe between buffers, do one of the following:

1. Prepare to wipe between two images by loading different nodes into each buffer, or load the same node with different viewer options into each buffer.
2. To toggle the split wipe on or off, do one of the following:
   a) Click the Switch to Split Wipe View button.
   b) Press Forward Slash (/).
3. To adjust the wipe, do one of the following:
   a) Move the center of the wipe by dragging the center handle of the wipe divider.
   b) Press Command-Option and click anywhere in the viewer to jump the wipe divider to that location.
   c) Change the angle or the wipe by dragging the wipe divider. Dragging the wipe divider while holding the Shift key snaps it to the nearest 45-degree angle.
   d) Panning or zooming the viewer pans and zooms both buffers together.
4. (Optional) If you want to change the image that's displayed on that side of the split, you can drag new nodes onto either half of the viewer.
5. To turn off the wipe, click the Switch to Split Wipe View button again (or press /).



The wipe divider can be adjusted for comparing different areas of the A and B images

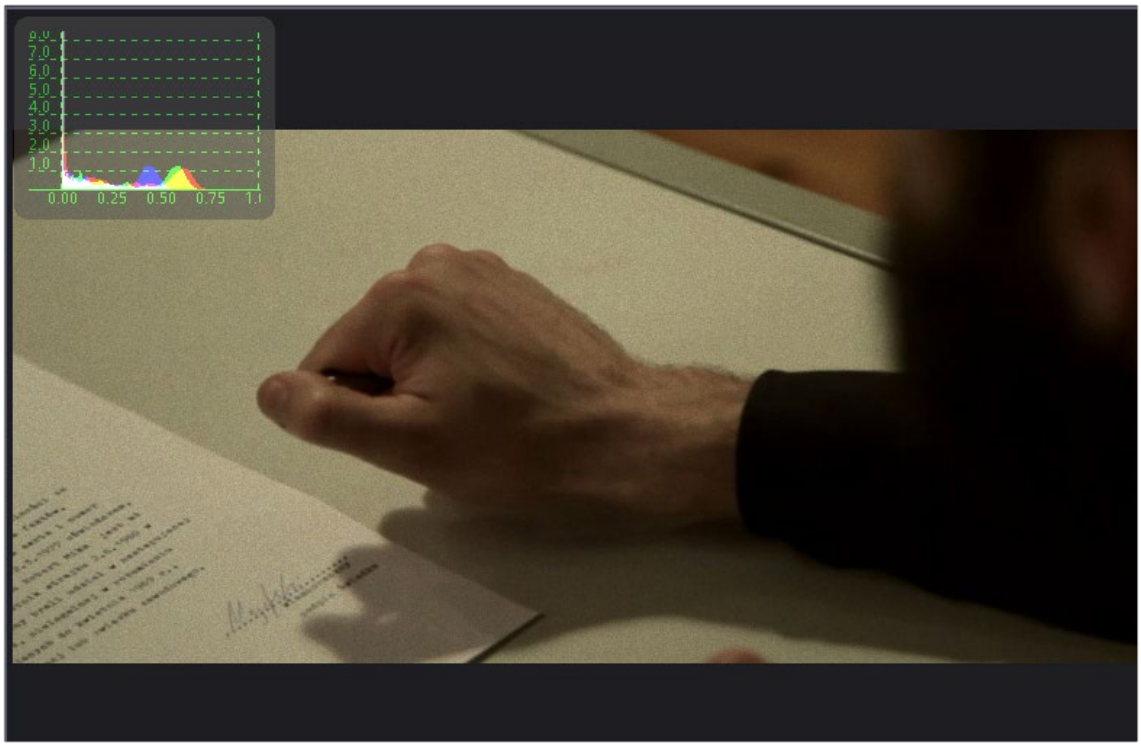
Even when you wipe, you can choose different display channels, view LUTs, or other display options for each buffer individually by clicking on the half of the wipe you want to alter, and then choosing the options you want that buffer to use. This allows easy comparison of different channels, LUTs, or other viewer settings while wiping the same image, or different images.

Moving the Wipe Divider

Occasionally, you will have either zoomed in or panned so far from the viewer divider that it's no longer visible in the viewer. Holding down Command-Option and clicking anywhere in the image will cause the viewer divider to jump to the current position of the pointer.

Subviews

A subview is a "mini" viewer that appears within the main viewer. A subview is usually used to show different information about the image.



The Subview menu with the Histogram subview displayed

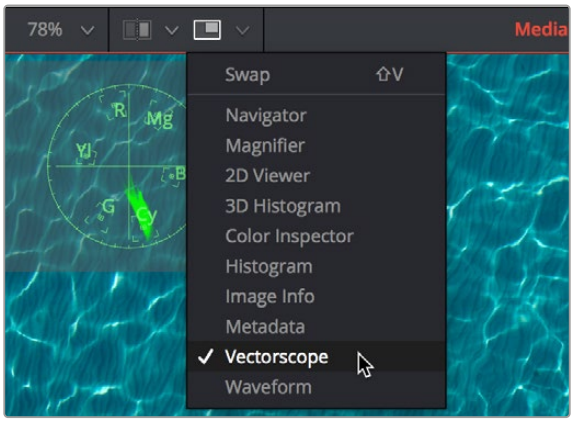
For example, the RGB channels can be viewed in the main viewer, while the alpha channel is displayed in a subview. For the most part, the subview is a fully functional miniature viewer, with its own contextual menu and options. It responds to the same keyboard shortcuts and navigation controls as any other viewer. However, there are several view types designed for use only in the subview, including the Navigator, Magnifier, Color Inspector, and Image Info.

Showing and Hiding Subviews

Subviews are easily shown and hidden.

To enable the currently selected subview in the Subview menu of a viewer, do one of the following:

- Click the Subview button in the View toolbar.
- Choose Views > Subview > Enabled from the contextual menu.
- Click a viewer, and press the V key.



The Subview button in the viewer toolbar

Changing the Subview Type

The Subview button enables and disables the subview, which usually shows the last subview you chose. You can change this at any time.

To change which subview type is displayed, do one of the following:

- Click the small arrow to the right of the Subview button to open its menu and choose which subview you want.
- Right-click within a subview to bring up the subview's contextual menu.

The Subview drop-down menu and contextual menu show all the available subview types. Once you choose an option from the list, that view will be displayed in the subview, and the Subview button will show and hide it as you wish.

Swapping the Subview with the Main View

It's possible to swap the contents of the main viewer and the subview for select subview types. However, certain view types, such as the Color Inspector and Magnifier, can only be used in the subview. In these cases, the swap will not take place.

To swap the contents of the subview with the main view, do one of the following:

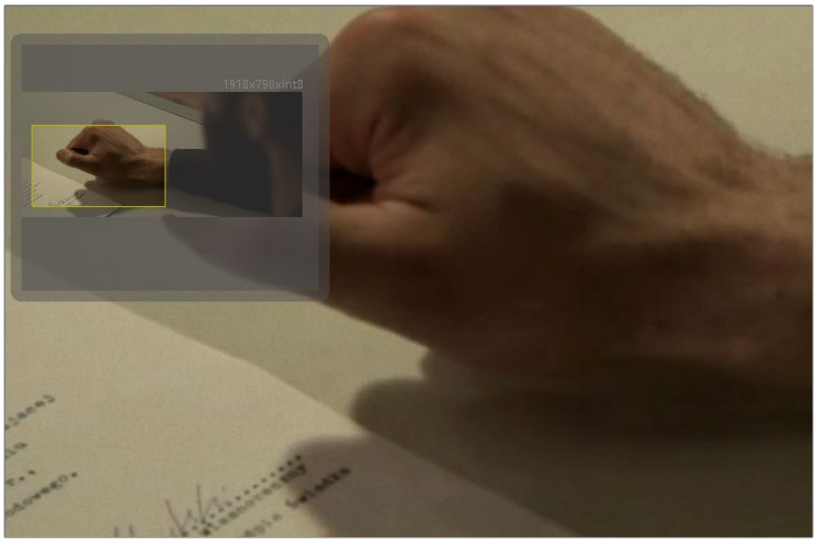
- Press Shift-V.
- Right-click in a viewer and choose Views > SubView > Swap from the contextual menu.

Viewer and Subview Types

Viewers can be changed to show a variety of different information about the image, but not all view types are available at all times. For example, the 3D Viewer is not available for a 2D node, and some of the measurement viewers are available only as subviews. Below is detailed information about the different view types available.

Navigator

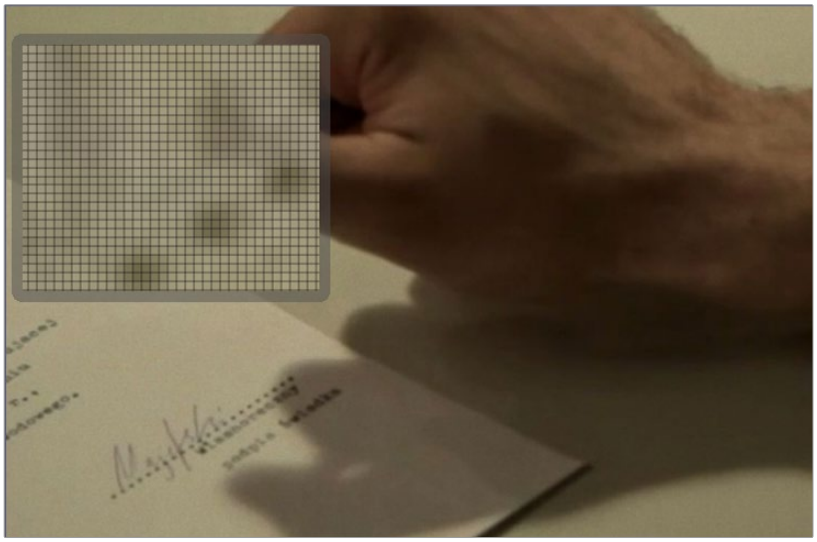
The Navigator can only be used in a subview. It provides a small overview of the entire image, with a rectangle that indicates the portion of the image that is actually visible in the main viewer. This is useful when zooming in on an image in the main view.



The Navigator subview for panning the image while zoomed in

Magnifier

The Magnifier can be used only in a subview. It shows a zoomed-in version of the pixels under the cursor in the main viewer.

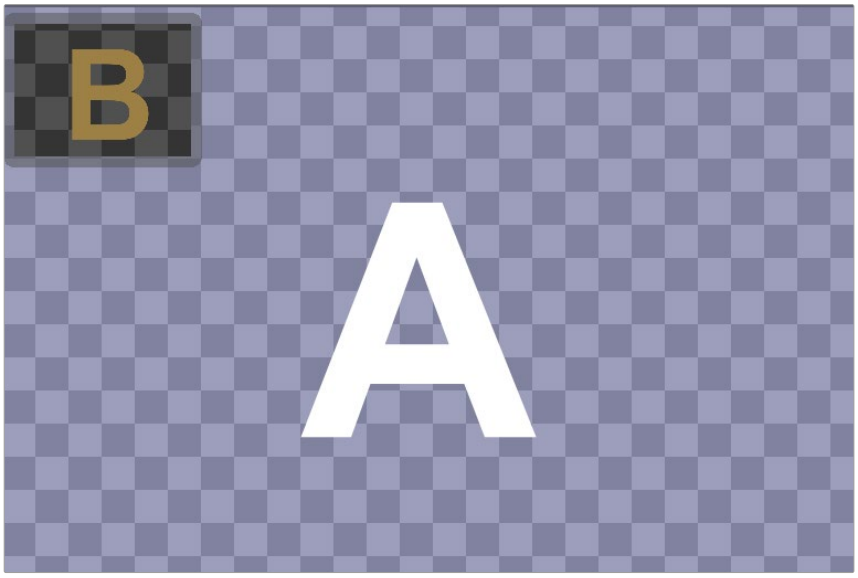


The Magnifier subview used to view a zoomed-in version of the image

2D Viewer

The 2D Viewer is the default type for showing images. When used as a subview, a different node than the one used in the main viewer can be displayed by dragging the node into the subview.

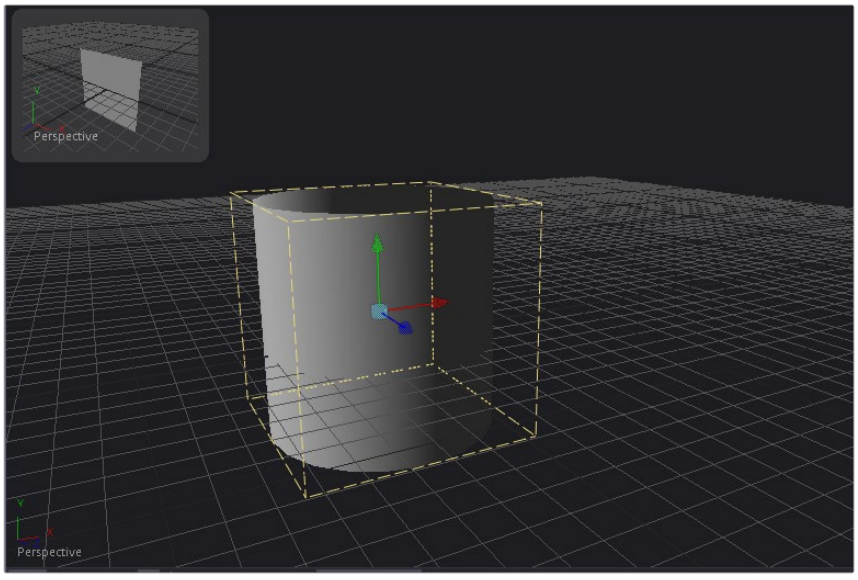
This is the only subview type that is not just a different view of the same node in the main viewer.



The subview used as another viewer

3D Image Viewer

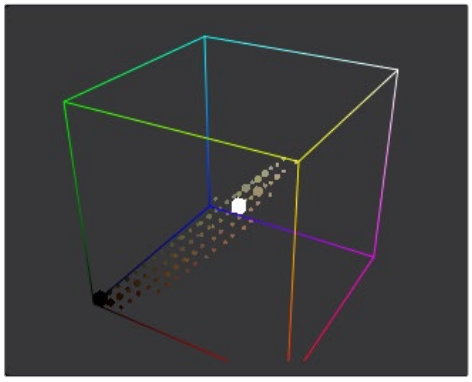
The 3D Image Viewer is available when viewing a node from the 3D category.



A 3D Image Viewer as a subview

3D Histogram

The more advanced 3D Histogram Viewer shows the color distribution in an image within a 3D cube. One advantage to a 3D Histogram is that it can accurately represent the out-of-range colors commonly found in floating-point and high-dynamic-range images. It can also be used to look at vector images like position, normal, velocity, and so on.



The 3D Histogram Viewer type for evaluating out-of-range colors

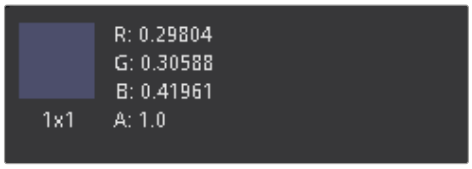
To rotate within a 3D Histogram, do one of the following:

Hold down the Option key, and drag left or right using the middle mouse button.

Hold down the middle and right mouse buttons while dragging.

Color Inspector

The Color Inspector can only be used in a subview. The Color Inspector shows information about the color channels of the pixel under the cursor. It will show all channels present, even the auxiliary channels such as Z buffer, XYZ normals, and UV mapping channels.

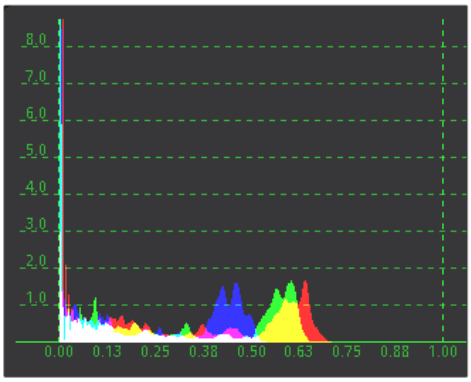


The Color Inspector subview for evaluating colors in an image

Histogram

The Histogram Viewer is an analysis node that can be used to identify problems with the contrast and dynamic range in an image. The graph shows the frequency distribution of colors in the image, including out-of-range colors in floating-point images. The horizontal axis shows the colors from shadows to highlights. The vertical axis shows the number of pixels in the image that occur at each level.

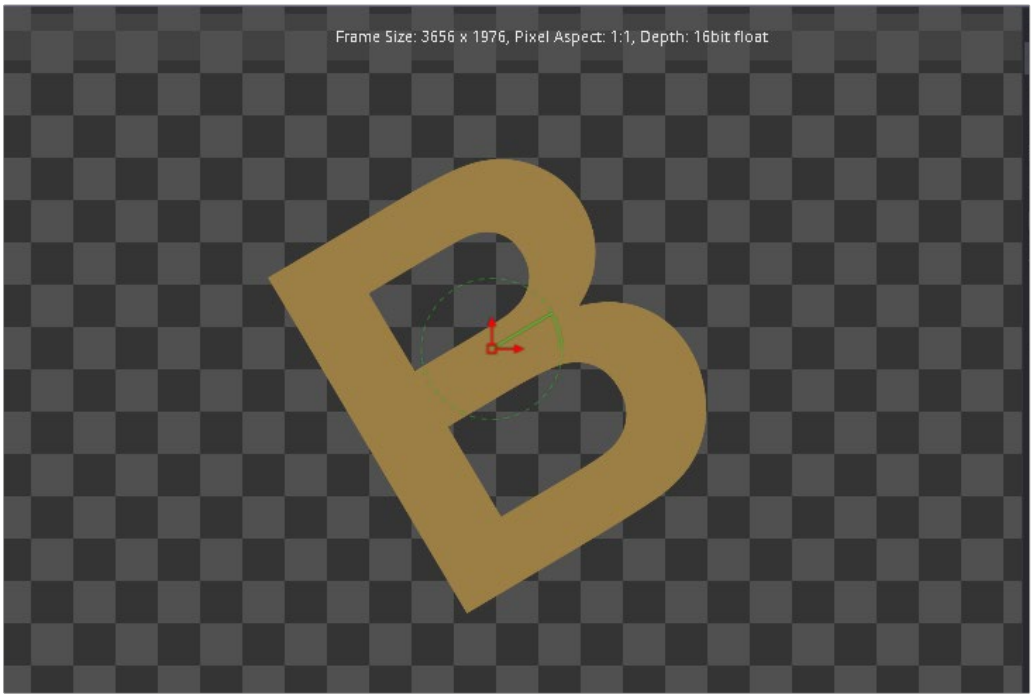
The Histogram Viewer will also display gradient information. You can use the From Image and Perturb modifiers to output gradients. If you need to see the gradient represented in a histogram, drag the modifier's title bar into the viewer.



The Histogram Viewer type for evaluating the contrast and color cast in an image

Image Info

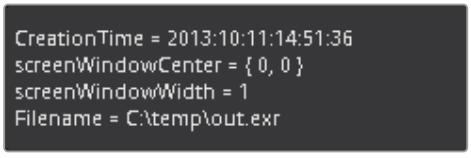
The Image Info view can only be used in a subview. The Image Info tab shows a horizontal bar across the top of the image with information about the frame size, pixel aspect, and color depth of the viewed image.



The Image Info subview for viewing size, pixel aspect, and color depth information

Metadata

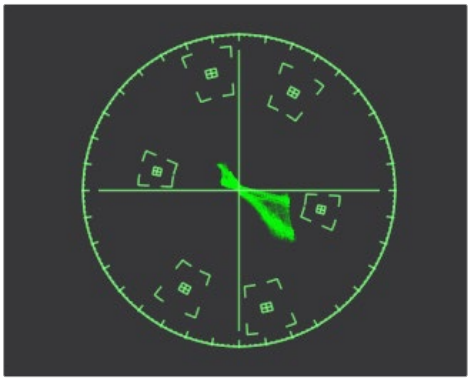
The content of this subview is based entirely on the amount of metadata in your image. Most Loaders will give the color space and file path for the image. Much more information can be displayed if it exists in the image.



The Metadata subview for viewing embedded metadata

Vectorscope

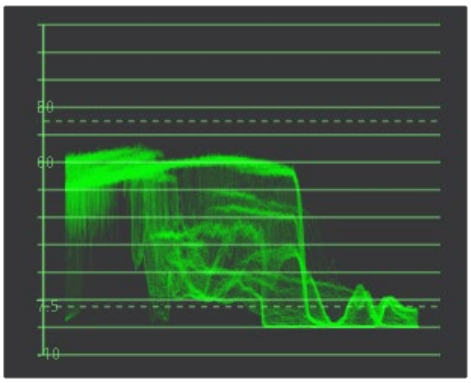
The Vectorscope Viewer duplicates the behavior of a specific type of video test equipment, displaying a circular graph that helps to visualize the intensity of chrominance signals.



The Vectorscope Viewer type for evaluating chrominance in an image

Waveform

The Waveform Viewer duplicates the behavior of a specific type of video test equipment, displaying a line or bar graph that helps to visualize the voltage or luminance of a broadcast signal.



The Waveform Viewer type for evaluating luminance in an image

Viewing Selective Channels

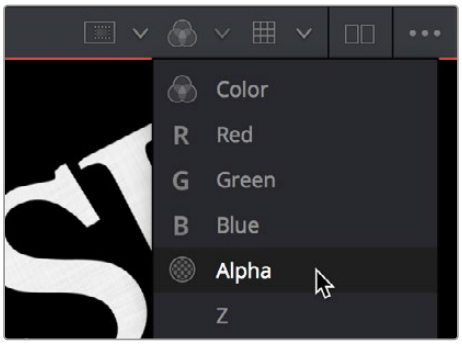
When compositing, you often deal with individual color components or channels in an image as much as you deal with the full RGB color of the entire image. The viewers and subviews can display the isolated color, alpha, depth channels, and even auxiliary channels that make up the image.

Viewing Color Channels

The default view is the full RGB color channel, but to change the channel that is displayed you can use the Channel toolbar button, the viewer's contextual menu, or keyboard shortcuts.

To toggle between RGB and alpha channels in the active viewer:

- Click the Color button in the viewer toolbar to toggle between full RGB color and that image's alpha channel.



The Channel toolbar button for switching between color, alpha, and depth channels

To toggle the channel that's displayed in the active viewer:

- Click the arrow to the right of the Color button to choose a specific channel to view from the list of available channels in the current image.
- Click the viewer you want to toggle, and press one of the following keyboard shortcuts;

  C - Full RGB color display

  R - Display red channel

  G - Display green channel

  B - Display blue channel

  A - Display alpha channel

  Z - Display Z-buffer channel

Viewing Auxiliary Channels

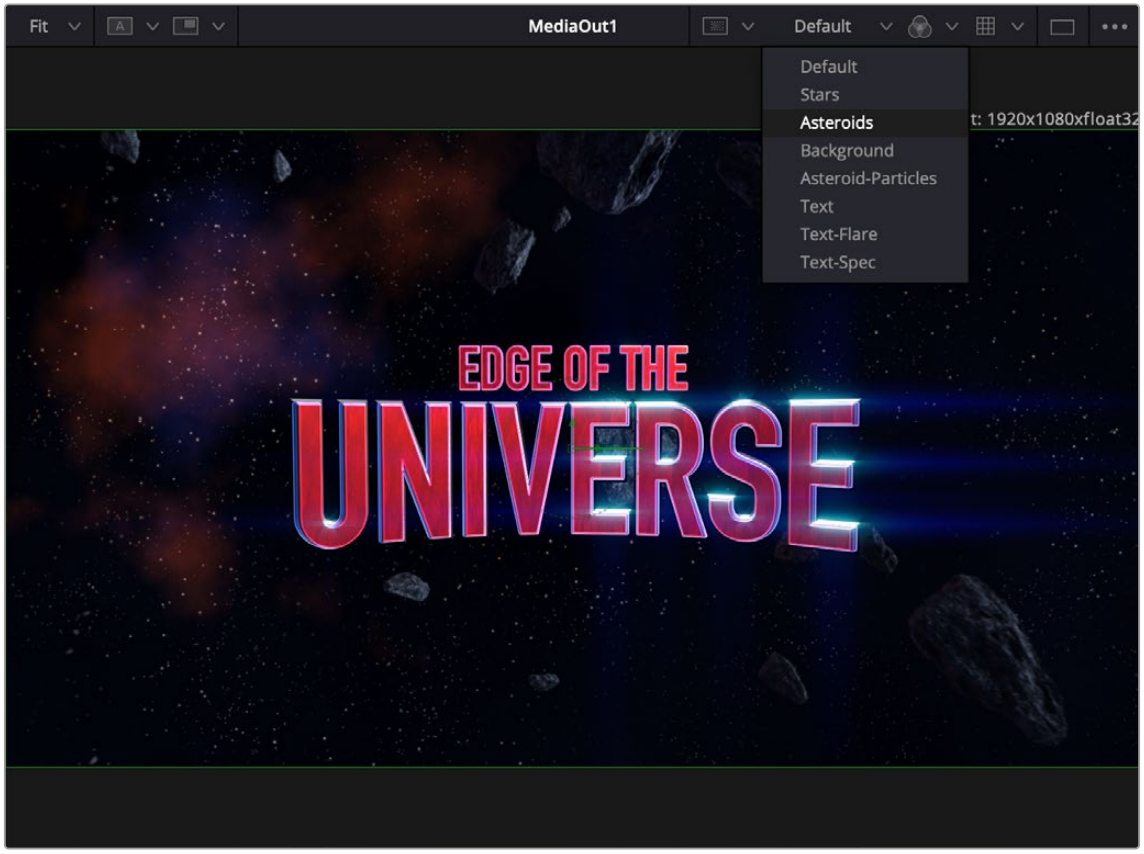
The viewers support RGBA and Z channels using keyboard shortcuts, but they support other channels as well. File formats such as OpenEXR often include auxiliary image data that provide more control and compositing options when working with rendered 3D images. To view auxiliary image data in a viewer, click the arrow to the right of the RGB button to display the drop-down menu or right-click in the viewer and choose an option from the Channels submenu of the contextual menu.

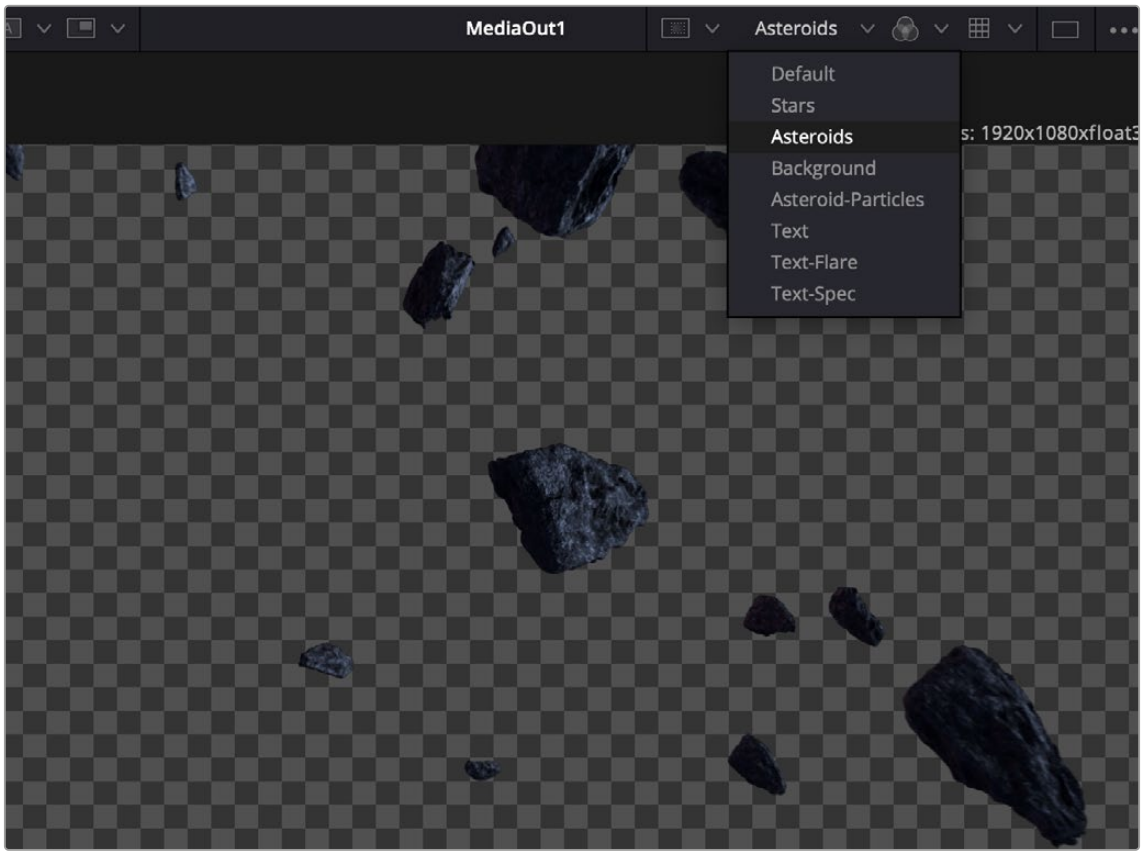
Viewing Multilayer Files (Studio Version Only)

Fusion can intelligently handle multilayer image workflows with support for EXR, Stereoscopic 3D, and PSD assets. Users can see a list of layers on the Fusion viewers and interact with layers in each node in the Settings tab in the Inspector. As this layer structure is retained throughout the toolset, this means you can isolate a tool to work on an individual layer, while still passing through the other layers unmodified.

Viewer

When loading a multilayer file, you can preview specific layers in the Viewer and selecting a layer from the drop-down list. Files without layer support will show only a 'default' option from the drop-down menu.





Using the Layer drop-down in the Viewer to select just the Asteroids layer from a multilayer file for viewing

The 3D Viewer

Building a composite in 3D space has different requirements from traditional 2D compositing. When a node from the 3D category or some particle systems is selected, a 3D Viewer is used to display the scene. The 3D Viewer shows a representation of a composite in a true GPU-accelerated 3D environment.

For more information on 3D controls, *see Chapter 25, "3D Compositing Basics,"* in the DaVinci Resolve Reference Manual or Chapter 25 in the Fusion Reference Manual.

Panning, Scaling, and Rotating a 3D Viewer

For the most part, navigation in the 3D Viewer is similar to the navigation in the 2D Viewer. Panning and zooming work with the same controls even though you're moving within a 3D space. However, when viewing a 3D scene, panning changes the point of view and thus the center point for scaling and rotation, too. A combination of panning and rotation will allow you to move the point of view anywhere in the scene.

Another small change is that there's a lower limit to the scale of a 3D scene. Continuing to zoom in past this limit will instead move ("dolly") the point of view forward. The mouse wheel will move forward slowly, and the keyboard will move more quickly.

Critically, the 3D Viewer gives you additional control to rotate the viewer within the three dimensions of the scene to better see your scene from different angles as you work.

To rotate within a 3D Viewer, do one of the following:

- Hold down the Option key and drag left or right using the middle mouse button.
- Hold down the middle and right mouse buttons while dragging.

The rotation is centered on the middle of the view.

TIP: These rotation controls can be used with the 3D Histogram subview as well.

Viewing Objects via Wireframe

3D composites not only work with 2D images on image planes but can also integrate true geometry, such as that generated by the Particle system, Text 3D node, imported FBX meshes, and basic primitives from the 3D toolset. Using a Wireframe view helps to see through a mesh or see the density of the geometry. It is much easier to see a change in the Subdivision level of an ImagePlane3D in wireframe than viewing the rendered image.

To display 3D geometry in wireframe, do the following:

- Right-click the 3D Viewer and choose 3D Options > Wireframe from the contextual menu.

Changing the POV of a 3D Viewer

Compositing a 3D scene often requires that you view the scene from different angles to make sure the position and animation are what you want. While the 3D Viewer uses a perspective camera that allows you to look at the 3D scene from any angle, you can also switch the 3D Viewer to view from the front, top, left, or right side of the scene; these are collectively called Orthographic views.

Additionally, if you have a camera or spotlight in your scene, you can switch the viewer to face the scene from the point of view of those objects.

To change the 3D viewpoints:

Right-click the viewer and choose an option from the Camera submenu of the contextual menu. The choices include Perspective, Front, Top, Left, and Right.

Changing Cameras in a 3D Viewer

If you have one or more camera objects in the 3D scene, they will be listed as options in the contextual menu. Spotlights and other lights or objects in the scene will be listed in the Other submenu. If you choose any one of these objects, the 3D Viewer jumps to display the scene from the point of view of the chosen object. While looking "through" an object, rotating, panning, or zooming, the viewer will instead change the rotation, position, and scale of the camera, light, or other object.

Copying a Viewer's POV to a Camera

There are many times you will have used the default perspective view to set up your scene and then want that point of view (POV) to become your main camera's point of view. You can easily position and orient a camera, light, or other object to match the current point of view shown in the viewer using the Copy PoV To command.

To copy the point of view in the viewer to a camera, do the following:

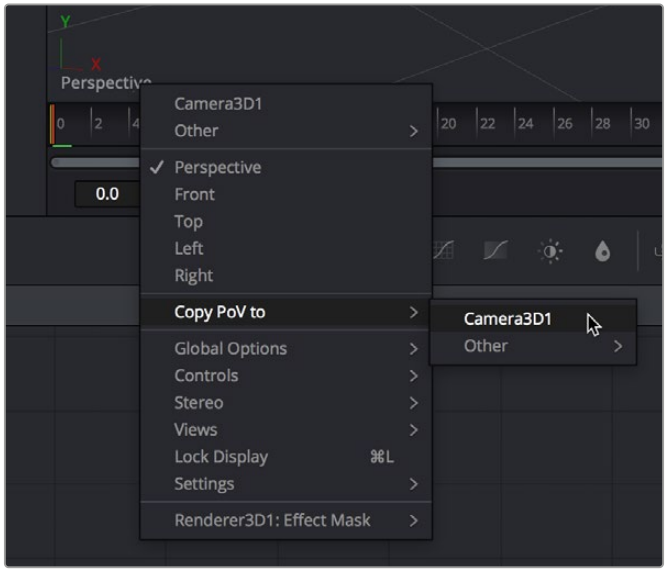
1. Set up a 3D Viewer with the point of view you want by zooming, panning, and rotating the viewer.
2. Add a camera to your 3D scene.
3. Right-click anywhere within the 3D Viewer and choose Camera > Copy PoV To > Camera3DNameOfCamera from the contextual menu.

The Camera3D's controls will inherit the viewer's position and angle values.

TIP: The Copy PoV To command uses the object's own coordinate space; any transformations performed downstream by another node are not taken into account.

POV Labels

As you switch the POV of the viewer, you can keep track of which POV is currently displayed via a text label at the bottom-left corner of the viewer. Right-clicking directly on this label, or on the axis control above it, acts as a shortcut to the Camera submenu, allowing you to easily choose another viewpoint.



The Viewpoint label and Axis control in a 3D Viewer

Lighting and Shadows in 3D Viewers

Before you add lights to a 3D scene, default lighting is provided. This basic, flat lighting allows you to see the shading on objects without requiring you to add and set up lights as you work in the 3D Viewer. Additionally, shadows are hidden by default. Once you start adding lights of your own, you need to switch modes to see what they affect as you work.

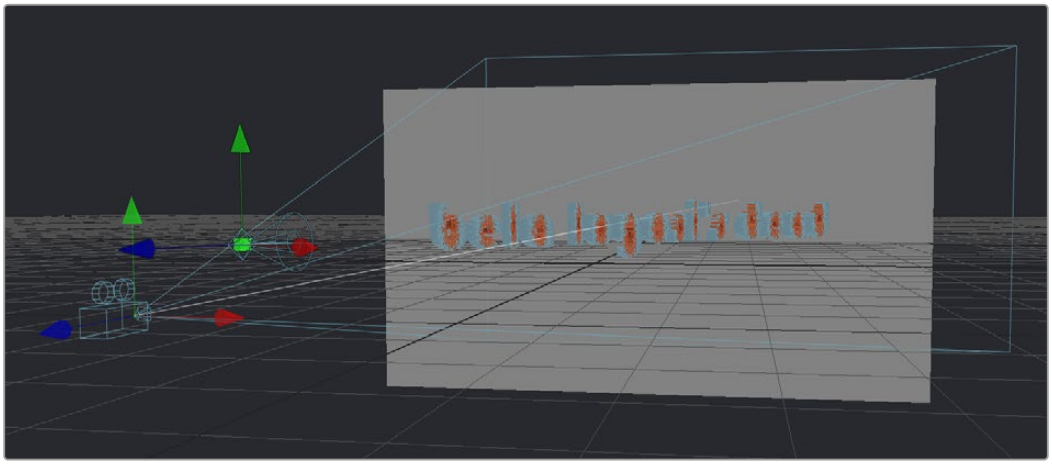
To see the effects of the default light on the scene:

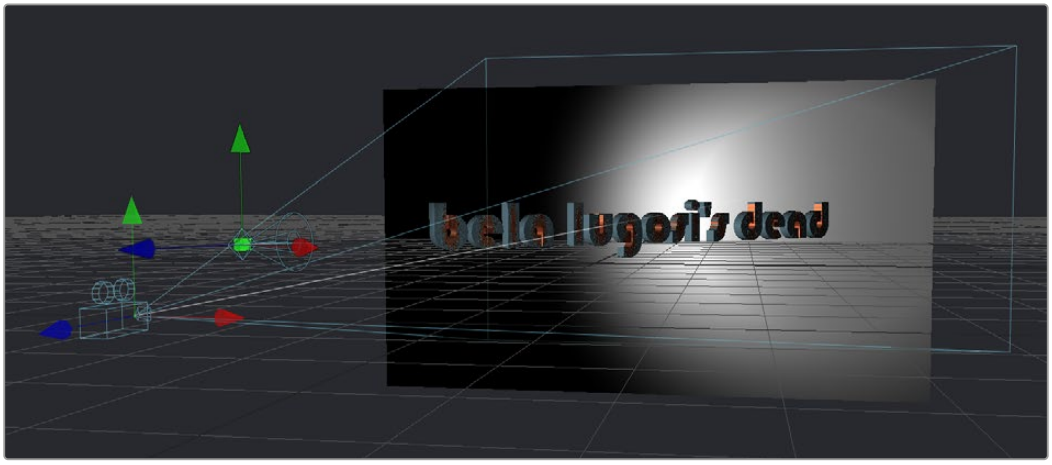
- Right-click within the 3D Viewer and choose 3D Options > Default Lights from the contextual menu.

  When you're ready to add your own lighting to a scene, you can connect light nodes in various ways to a Merge 3D node for the scene you're working on. Once you connect a light to a Merge 3D node, you need to switch the 3D Viewer over to showing the new, proper lighting.

To toggle lighting rendering within a 3D scene:

- Right-click within the 3D Viewer and choose 3D Options > Lighting from the contextual menu.





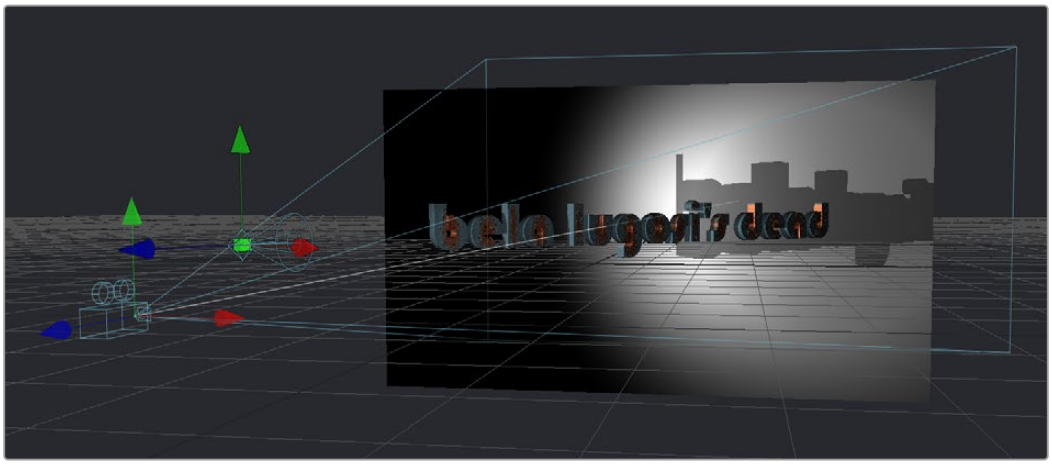
A 3D scene using default lights (top), and the same scene with lighting turned on (bottom)

TIP: Attempting to load a Light node into a viewer all by itself will result in an empty scene, with nothing illuminated. To see the effects of lights, you must view the Merge 3D node the light is connected to.

Similar to lights, the default 3D Viewer has shadows turned off. To see shadows cast from the lighting you've created, you must turn them on.

To toggle shadows rendering within a 3D scene:

- Right-click within the 3D Viewer and choose 3D Options > Shadows from the contextual menu. Enabling shadows will automatically turn on lighting, if it is not already turned on.



A 3D scene with shadows enabled along with the lights

Transparency in 3D Viewers

Image planes and 3D objects are obscured by other objects in a scene depending on the X, Y, and Z position coordinates of each object in 3D space. The default method used to determine which polygons are hidden and which are shown based on these coordinates is called Z-buffering.

Z-buffering is extremely fast but not always accurate when dealing with multiple transparent layers in a scene. Fortunately, there is another option for more complex 3D scenes with transparency: Sorted.

The Sorted method can be significantly slower in some scenes but will provide more accurate results no matter how many layers of transparency happen to be in a scene.

The default behavior in the viewer is to use Z-buffering, but if your scene requires the Sorted method, you can easily change this.

To choose a Sorted method of 3D compositing:

- Right-click anywhere within the 3D Viewer and choose one of the options in the Transparency submenu of the contextual menu;

  Full Sort: Renders every polygon in Z order to produce the most accurate rendering of transparency.

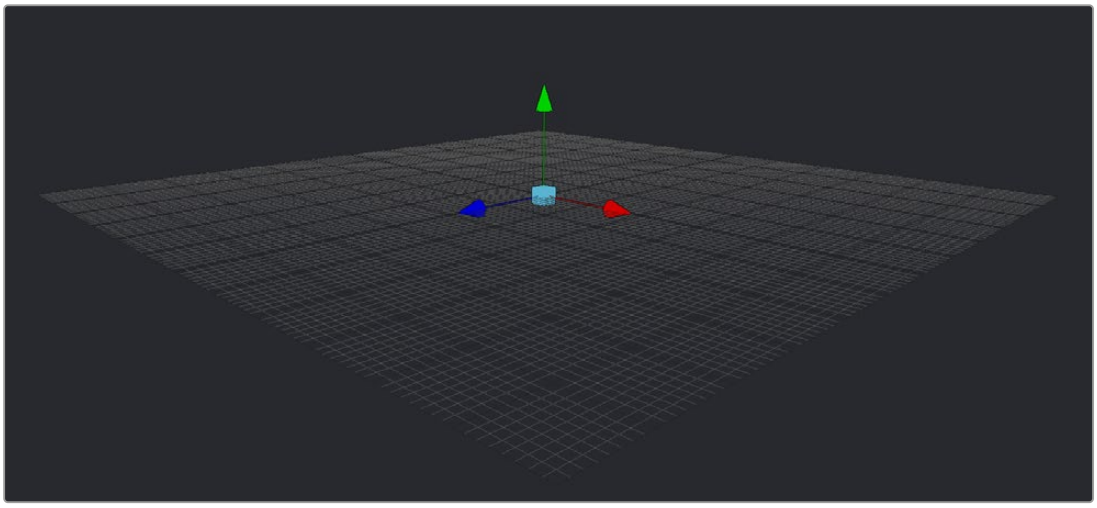
  Quick Sort: Reorders the polygons in the scene serially, from back to front, to produce a reasonably accurate rendering of transparency.

Grid

The 3D Viewer displays a grid that's used to provide a plane of reference in the 3D scene. By default, the grid is 24 x 24 units in size, centered on the origin at (0,0,0), and subdivided into large squares of 2 units with small squares of 0.25 units each. These defaults can be altered in the 3D View panel of the Fusion Settings window, available from the Fusion menu.

To toggle the grid on and off:

- Right-click anywhere within the 3D Viewer and choose 3D Options > Grid from the contextual menu.



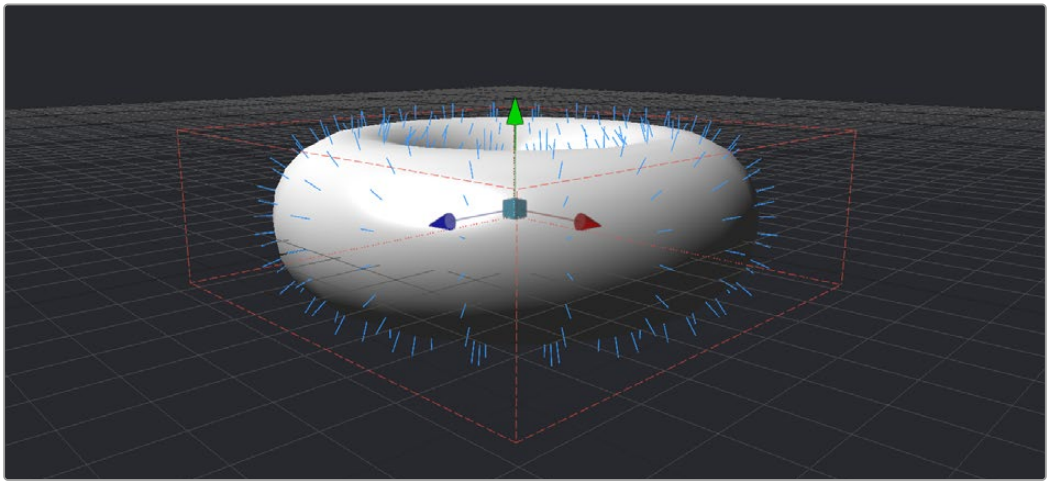
The default grid of the 3D Viewer grid with its origin at x = 0, y = 0 and z = 0

Vertex Normals

Normals indicate what direction each vertex of 3D geometry is facing, and they are used when calculating lighting and texturing on an object. When viewing any kind of 3D geometry, including an image plane or a full FBX mesh, you can display the normals for each object in a scene.

To view the normals in a scene:

- Right-click anywhere within the viewer and choose 3D Options > Vertex Normals from the contextual menu.



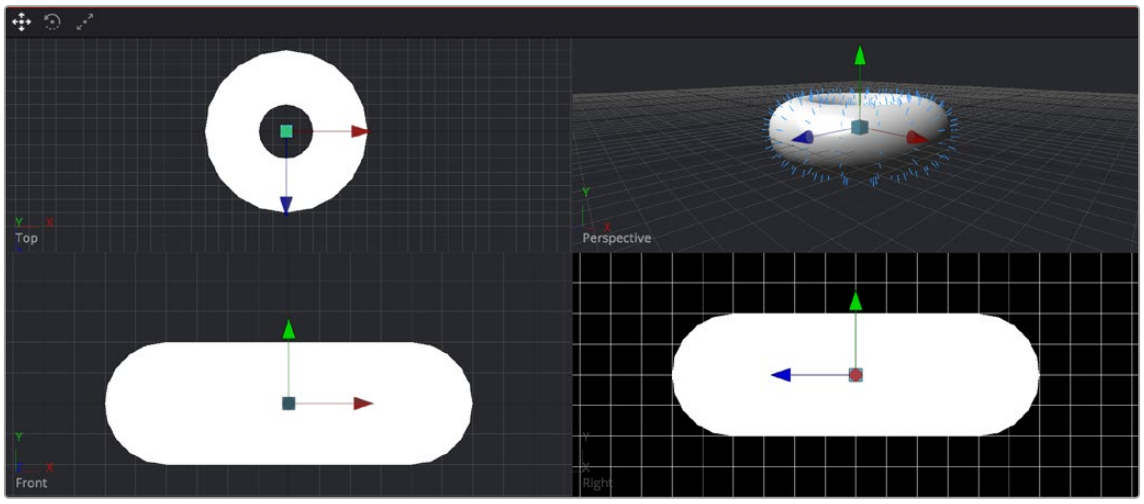
The normals viewed in a 3D scene

Quad View

3D compositing often requires you to view the scene from different points of view to better control transformations in three dimensions. While you can switch the 3D Viewer to different points of view, doing so frequently can become cumbersome. Happily, you can instead enable a Quad view, which divides the viewer into four panes. These panes can then display four different angles of the scene at one time.

To toggle the display of the Quad view, do one of the following:

- Right-click anywhere within the viewer and choose Views > Quad View from the contextual menu.
- Press Shift-Q.



A Quad view of a 3D scene

While there are four panes in the Quad view, they all show the same scene. When assigning views within a Quad view, you can choose between displaying Front, Left, Top, Bottom, and Perspective orthographic views, or you can choose the view through any camera or spotlight that's present in the scene.

To assign different views to panes of a Quad view, do one of the following:

- Right-click directly on the POV label at the bottom left of the pane you want to reassign, and choose another camera, light, or Point of View from the contextual menu.

Quad View Layouts

There are a variety of Quad view layouts, ranging from four equally sized panels to having three small panels across the bottom of a larger single panel.

To switch to a different Quad view layout, do the following:

1. Enable the Quad view.
2. Right-click anywhere within the viewer and choose an option from the Views > Quad Layouts submenu of the contextual menu.

Using Quad Views for 2D Scenes

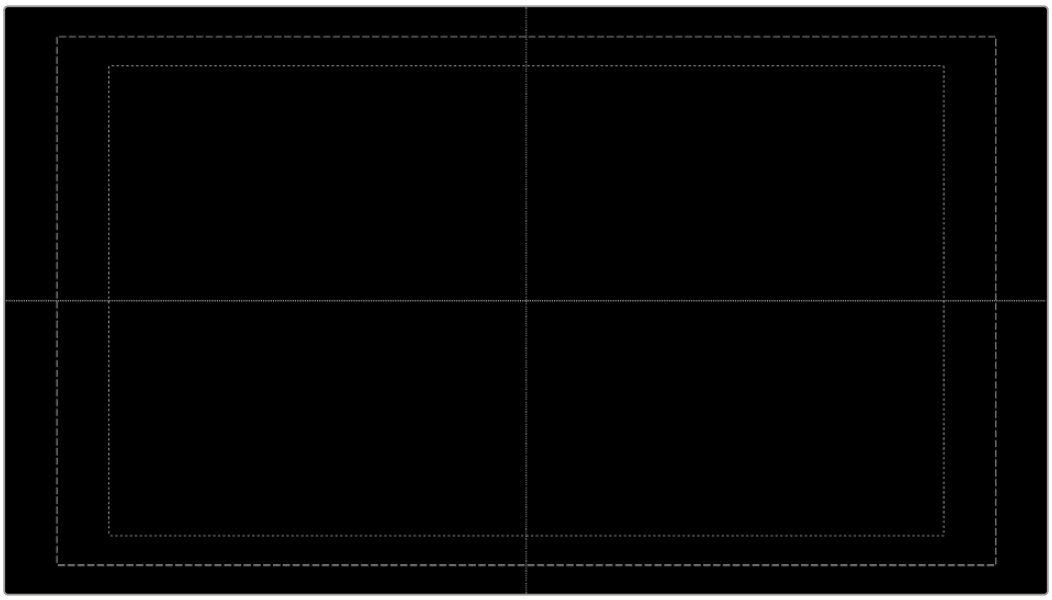
Quad views aren't only useful for 3D scenes. They can also be used with 2D scenes, with each pane showing a different image channel or subview type. For example, one pane can show the image while the other panes show the alpha channel, a vectorscope, and a histogram.

To assign different channels or subviews to panes of a Quad view for a 2D scene:

1. Right-click in a viewer and choose Views > Quad View.
2. Click once in the pane you want to reassign.
3. Do one of the following:
   a) Choose a channel from the Channel Viewer menu.
   b) Right-click in the viewer and choose Views, and then choose a Subview from the submenu.

Guides

Guides are onscreen overlays used to help you compose elements within a boundary or along the center vertical and horizontal axes. While guides are displayed in the viewer, they're not rendered into the scene. There are four commonly used guides that can be displayed, including Monitor Safety, Safe Title, Center, and Film.



The Guides submenu in the viewer's contextual menu

Methods of using guides:

- **To display guides in a viewer:** Right-click in the viewer and then choose Guides > Show Guides from the contextual menu, or press Command-G.
- **To change the aspect ratio of the displayed guides:** Right-click in the viewer and then choose an option from the Guides > Frame Aspect submenu. The frame aspect is usually set to Default, which forces the frame aspect to the same resolution as the image that's displayed in the view. However, when the frame aspect is set to a specific value, the guides will conform to the exact boundaries of the specified format and any image area outside of that will be dark gray.
- **To show or hide specific guides:** Right-click in the viewer and then choose an option from the Guides submenu. A variety of specific guides are provided, each of which can be individually enabled and disabled.

  Monitor Safety: Monitor Safety indicates the safe action area viewable on most monitors and TV screens.

  Safe Title: Safe Title indicates the safe area for titles viewable on all TV and monitor screens.

  Center: Center shows a crosshair for the center point and x- and y-axis of the view.

  Film: Some frame formats include film guides preset for you, whereas some will require customization. The film guides can be customized in the Preferences > Frame Format window.

Frame Format Settings

In the Frame Format panel of the Fusion Settings window (available in the Fusion menu), there are two film guide settings that you can use to customize these guides.

- Guide 1 contains four fields that specify the offset from the edges of the image for the left, top, right, and bottom guides, in that order. As with all offsets in Fusion, this is a resolution-independent number where 1 is the width of the full image and 0.5 is half the width of the image.
- Guide 2's text box is used to set the aspect ratio of the projection area.



The Frame Format Guides settings

Domain of Definition and Region of Interest

As a compositing environment, the Fusion page uses the standard compositing conventions of Region of Interest (RoI) and Domain of Definition (DoD) to dramatically improve performance.

Domain of Definition (DoD)

In compositing, the Domain of Definition, frequently abbreviated to DoD, refers to a rectangular region that defines what part of an image actually contains data. DoD makes the concept of an image's actual frame somewhat flexible, since rendering is no longer limited to the actual width and height of the image. This has two effects on the way Fusion renders images.

Firstly, nodes will no longer be required to render portions of the image that will not be affected by the node. This helps the renderer to optimize its performance. Secondly, Fusion can now keep track of and apply a node's effect to pixels that lie outside the visible portion of the image.

For example, consider the output of a Text+ node rendered against a transparent background. The text occupies only a portion of the pixels in the image. Without Domain of Definition, you would be required to process every pixel in the image needlessly. With a DoD, you are able to optimize effects applied to the image, producing faster results and consuming less memory in the process.

The following image shows an image with the DoD outlined.



The DoD is shown as two XY coordinates indicating the corners of an axis-aligned bounding box (in pixels)

For the most part, the DoD is calculated automatically and without the need for manual intervention. For example, all the nodes in the Generator category automatically generate the correct DoD. For nodes like Fast Noise, Mandelbrot, and Background, this is usually the full dimensions of the image. In the case of Text+ and virtually all of the Mask nodes, the DoD will often be much smaller or larger.

The OpenEXR format is capable of storing the data window of the image, and Fusion will apply this as the DoD when loading such an image through a Loader node and will write out the DoD through the Saver node.

When using the Fusion page in DaVinci Resolve, clips from the Edit page timeline or Media Pool will typically have the DoD default to the full image width of the source media. The exception is media stored in OpenEXR format.

The DoD is established as soon as the image is created or loaded into the composition. From there, it passes downstream, where viewers combine it with their Region of Interest in order to determine exactly what pixels should be affected by the node. As you work, different nodes will automatically shrink, expand, or move the DoD as they apply their effect to an image, causing the DoD to change from node to node.

Showing the DoD

If the current DoD for a node is different from the frame size of that image, it's shown in the tooltip that appears when the pointer hovers over a node in the Node Editor. The DoD is also visible in the viewer when you right-click in a viewer and choose Region > Show DoD from the contextual menu.

Setting the DoD Manually in the Node Editor

It is also possible to set the DoD for an image manually using the Tools > Miscellaneous > Auto Domain node in the Effects Library. This node can be useful when dealing with pre-created media that does not occupy the full image dimensions. For example, a rendering of a 3D character that walks toward the camera will frequently occupy only a portion of the image. The Auto Domain node can be used to animate a DoD that covers the character and ignores the rest of the image, making image processing more efficient.

Region of Interest (RoI)

The Region of Interest, frequently abbreviated to RoI, is a rectangular region similar to the Domain of Definition. However, unlike the DoD, which tells the node what pixels are actually present in the image, the RoI tells the node which pixels actually need to be rendered. When a node renders, it intersects the current RoI with the current DoD to determine what pixels should be affected.

Enabling RoI Controls

You can turn on the RoI controls to restrict rendering to a small region of the image to significantly improve performance when you're only working on a small part of a high-resolution or complex composition. For example, if you're using paint to clean up some holes in a matte on the floor of a composition with many, many high-resolution layers, 3D, and Lighting operations, you can use the RoI controls to isolate the part of the floor you're working on, which makes caching that part of the composition much faster.

To enable the RoI controls, do one of the following:

- Click the RoI button in the 2D Viewer toolbar.
- Right-click in a viewer and choose Region > Show Region from the contextual menu.

When RoI is enabled and Show Region is selected from the menu, a rectangular RoI control appears in the viewer. If this is the first time RoI has been enabled, it will be set to the full width and height of the image. Otherwise, the last known position of the RoI for that view is used. However, if you want to set the RoI to a custom area within the frame, you can do one of the following.

To adjust the RoI controls, do one of the following:

- Drag any edge of the RoI rectangle to adjust one side of the RoI.
- Drag a corner to adjust the size of the RoI rectangle from that corner.
- Drag the small circle found at the top left corner of the RoI rectangle to move the RoI without adjusting its dimensions.

  Sometimes, it's faster to simply draw a rectangle where you want the RoI to be.

To quickly draw the RoI at the desired size:

1. Choose Set from the viewer menu next to the RoI button, or right-click anywhere within the viewer and choose Region > Set Region.
2. When the pointer turns into an RoI drawing cursor, drag within the viewer to set a RoI rectangle.

Alternatively, an Auto command sets the RoI to fit whichever pixels are visible at the current zoom/pan level in the viewer. This lets you quickly limit the RoI to whatever part of the composition you've zoomed into.

To automatically draw the RoI:

- Choose Auto from the viewer menu next to the RoI button.
- Right-click anywhere within the viewer and choose Region > Auto Region.
- When you no longer need to use the RoI, you can reset it.

To reset the RoI to the full width and height of the current image, do one of the following:

- Choose Reset from the viewer menu next to the RoI button.
- Right-click anywhere within the viewer and choose Region > Reset Region from the contextual menu or from the toolbar button menu.
- Disable the ROI control, which will also reset it.

While the RoI Is Active

The RoI is only used for previewing your composition while you work, not for output from Fusion. While the RoI is active, Fusion will only request rendering of the pixels inside the region when it displays an image in that viewer. Flipbook Previews that you create in that viewer will also respect the current RoI. MediaOut and Saver nodes will always use the full image dimensions when writing the image to disk, ignoring any RoI you've set in the viewers.

The RoI improves not only rendering speed and memory use, but it can also reduce file I/O, since Loaders and MediaIn nodes only load pixels from within the RoI, if one is specified. This does require that the file format used supports direct pixel access. Cineon, DPX, and many uncompressed file formats support this feature, as well as OpenEXR and TIFF in limited cases.

Please note that changes to the viewed image size or color depth will cause the pixels outside the RoI to be reset to the image's canvas color. This also happens when switching in and out of Proxy mode, as well as during Proxy mode switching with Auto Proxy enabled. When the image size is maintained, so are the last rendered pixel values outside the RoI. This can be useful for comparing changes made within the RoI with a previous node state.

TIP: Right-clicking in a viewer and choosing Options > Show Controls for showing onscreen controls will override the RoI, forcing renders of pixels for the entire image.

Managing Viewer Lookup Tables (LUTs)

Lookup Tables, or LUTs, can be used to help match the appearance of a viewer to its eventual output destination. They're essentially image-processing operations that affect only the image being previewed in the viewer, not the image data itself. There are two basic ways that LUTs can calculate color transformations: The first is a simple 1D LUT, and the second is a more sophisticated 3D LUT.

- The simplest form of a LUT is a 1D LUT. It accounts for one color channel at a time, so it can make overall tonality changes but not very specific color changes.
- A 3D LUT looks at each possible color value (red, green, and blue) independently. A 3D LUT allows for large global changes as well as very specific color changes to be applied to images very quickly.

How Lookup Tables Work in Fusion

A Lookup Table (LUT) is a table of values used to transform the color and luminance of an image. A 1D LUT uses a two-column table for input color and output color, while a 3D LUT uses more of a matrix. A LUT is used primarily to correct for variances in the monitor or the source color space of the image. You can choose to apply a LUT to all the viewers or apply different LUTs to each viewer.

Image LUTs

Image LUTs can be applied to each viewer. In fact, you can even apply separate Image LUTs for the A and B buffers of a single viewer. These LUTs can only be applied to 2D images and not to 3D scenes. Image LUTs are routinely used to get from one scene-referred color space to another. For example, if you're working with log-encoded media but want to see how the image will look in the final color space, you can choose a LUT to make the image transform as a preview.

Buffer LUTs

The Buffer LUT is applied to the viewers regardless of contents, including 3D scenes, 3D materials, and subview types. Only one Buffer LUT can be applied. If a 2D image is being displayed with an Image LUT applied, then the Buffer LUT is applied to the result of the image LUT. Buffer LUTs are typically used to simulate another output color space that's unique to the display you're using—for instance, making a DCI-P3 projector show the image as it would look on an sRGB monitor.

To use a Buffer LUT:

1. Disable the LUT button above the viewer.
2. Right-click in the viewer and choose Global Options > Buffer LUT > Enable.
3. Right-click in the viewer and choose Global Options > Buffer LUT > Type of LUT you want to apply.

When dealing with nonlinear files from many of today's digital cinema cameras, a modern workflow would be to convert everything to linear at the beginning of the node tree, then create your composite, and then apply an Image LUT or Buffer LUT that matches the color space you want it to be in for either grading in the Color page or for final output.

However, in more elaborate production pipelines, you may need to apply multiple LUTs consecutively.

Types of Viewer LUTs

Aside from the industry standard 1D and 3D LUTs, other types of LUTs are supported, including script-based Fuse node LUTs, OCIO files, and macros assembled from standard nodes. Generally, LUT processing is performed on the graphics card's GPU in real time, although the performance of macro-based LUTs is based on the nodes they contain.

Fusion View LUT

The Fusion View LUT is the default and is a frequently used LUT type. It provides an RGBA curve that can be used to assign IN/OUT value pairs. This control is identical to that provided by the Color Curve node.

Since the purpose of the View LUT is to provide an unchanging correction for the monitor or the file's color space, however, these splines cannot be animated.

Log-Lin View LUT

The Log-Lin LUT converts logarithmic data to linear, and vice versa. This can be particularly useful when used in conjunction with supplied LUT files that expect logarithmic data. It is similar to the Cineon Log node.

Gamut View LUT

The Gamut LUT converts a source color space to an output color space, with options to deal with gamma settings, alpha channels, and premultiplication. The Gamut LUT is a frequently used LUT type to correct the viewer when working with Linear Gamma in the Node editor.

Macro LUTs

Any macro node can also be used as a viewer LUT simply by saving the macro's .setting file to the correct Fusion directory.

In DaVinci Resolve, LUTs are saved in the following locations:

On macOS: Macintosh HD/Users/username/Library/Application Support/Blackmagic Design/Fusion/LUTs/

On Windows: C:\Program Files\Blackmagic Design\Fusion\LUTs

On Linux: home/username/.local/share/DaVinciResolve/Fusion/LUTs

In Fusion Studio, LUTs are saved in the following locations:

On macOS: Macintosh HD/Users/username/Library/Application Support/Blackmagic Design/Fusion/LUTs/

On Windows: C:\Users\username\AppData\Roaming\Blackmagic Design\Fusion\LUTs

On Linux: home/username/.fusion/BlackmagicDesign/Fusion/LUTs

For this to work, the macro must have one image input and one image output. Any controls exposed on the macro will be available when the Edit option is selected for the LUT.

For more information about creating macros, *see Chapter 6, "Node Groups, Macros, and Fusion Templates,"* in the Fusion Reference Manual.

LUT Presets

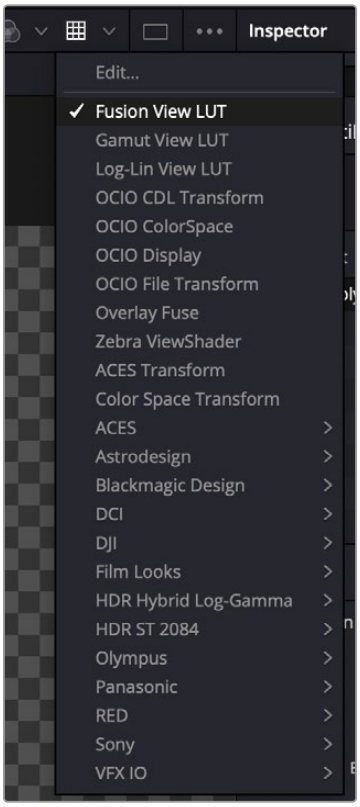
All LUTs available to DaVinci Resolve are also accessible to the Fusion page, which includes custom LUTs you've installed, as well as preset LUTs that come installed with DaVinci Resolve, such as the highly useful VFX IO category that includes a wide variety of miscellaneous to Linear and Linear to miscellaneous transforms. All of these LUTs appear by category in the viewer LUT menu.

Fuse LUTs

Fuses are scriptable plugins that are installed with the application or that you create in Fusion. A fuse named CT_ViewLUTPlugin can be applied as a LUT to a viewer. You can also script fuses that use graphics hardware shaders embedded into the LUT for real-time processing. Since fuse LUTs require shader-capable graphics hardware, they cannot be applied in software. For more information about Fuses, see the Fusion Scripting Guide located on the Blackmagic Design website.

Using Viewer LUTs

Viewer LUTs can be enabled, edited, and turned off using the viewer LUT button and menu, as well as by using the viewer contextual menu. This menu shows all LUTs available to Fusion, including custom LUTs you've installed yourself.



The viewer LUT button and menu from the Fusion page in DaVinci Resolve

To turn the current viewer LUT on and off:

- Click the LUT button in the viewer toolbar to toggle the viewer LUT on and off.
- The LUT menu can also be found as a submenu in the viewer's contextual menu.

To choose another viewer LUT:

- Open the menu to the right of the viewer LUT button and choose an option from the viewer LUT menu.

To apply a Buffer LUT:

1. Right-click anywhere within the viewer and choose Global Options > Buffer LUT > Enable.
2. To choose a specific Buffer LUT, right-click again and choose a LUT from the Global Options > Buffer LUT submenu.

   Buffer LUTs are often useful for applying monitor corrections, which do not usually change between projects.

To remove a Buffer LUT:

- Right-click anywhere within a viewer and choose Global Options > Buffer LUT > Enable to uncheck it.

Editing Viewer LUTs

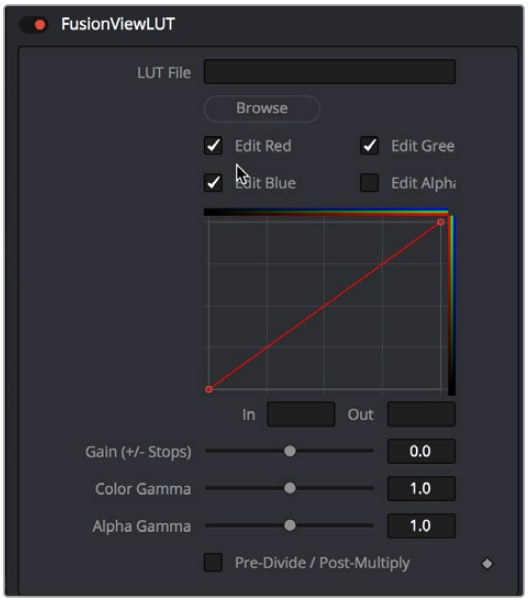
The viewers are the primary area where composites are assessed, so it's crucial that they provide an accurate representation of what the content will look like when it's played for an audience. The LUT Editor allows you to customize your viewer's output to match the gamma and color characteristics of your eventual playback device, or to test how the current image looks in a completely different color space, or how it holds up over a range of different color spaces.

To open any editable viewer LUT option's Editor:

1. Click the LUT button in the viewer toolbar to enable it.
2. Do one of the following:
   - Choose Edit from the top of the viewer LUT menu.
   - Right-click in the viewer and then choose LUT > Edit from the contextual menu.

Editing the Fusion View Lookup Table

Similarly to the Color Curves node, the Fusion View LUT Editor uses spline-based color correction. In addition to the ability to modify the separate color channels, the LUT has Gain and Gamma sliders. The Gain slider is helpful for temporarily brightening or darkening the viewed image, allowing easier examination of shadow or highlight detail. The Color Gamma and Alpha Gamma sliders are used to duplicate the gamma values of the eventual output device. Video monitors, for example, commonly have a gamma of 1.7, while computer monitors can range anywhere from 1.6 to 2.2. Alpha Gamma is applied only when viewing the alpha channel of an image, or when viewing masks.



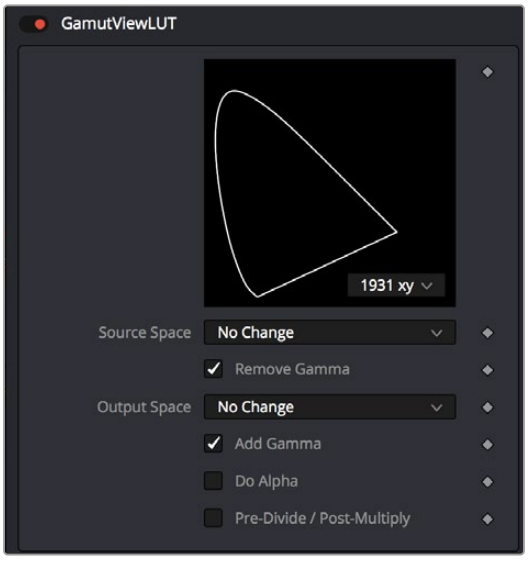
The LUT Editor for the default Fusion View LUT

Editing the Gamut View LUT

The Gamut View LUT Editor lets you choose a Source and Output color space to guide the viewer transform.

The Remove and Add Gamma checkboxes let you choose to do the gamut conversion with linear or nonlinear gamma, or they let you simply remove or add the appropriate gamma values without changing the color space.

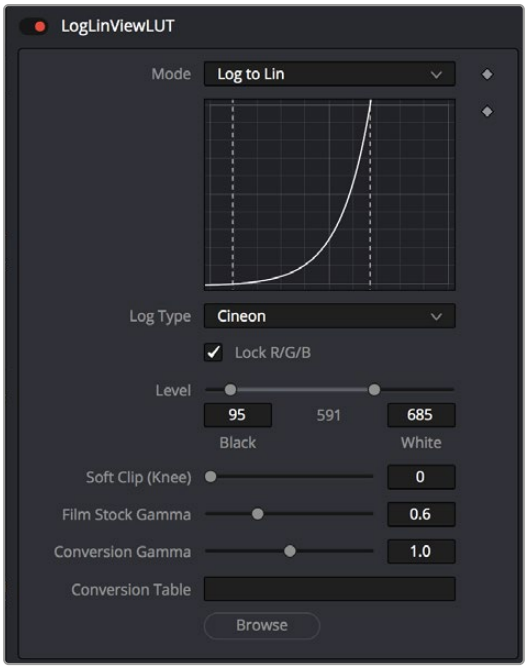
Selecting the Pre-Divide/Post-Multiply checkbox will cause the image's pixel values to be divided by the alpha values prior to this conversion, and then re-multiplied by the alpha value after this conversion. This helps to avoid the creation of illegally additive images, particularly around the edges of a blue/green key or when working with 3D rendered objects.



The Gamut View LUT Editor

Editing the Log-Lin View LUT

The Log-Lin LUT lets you apply a Log to Lin or Lin to Log operation using the Mode pop-up menu. You can choose the type of log-encoding to process from the Log Type drop-down, and choose whether to lock the R, G, and B channels together. A level adjustment lets you redefine the digital range of values used for the output, while Soft Clip (Knee), Film Stock Gamma, and Conversion Gamma sliders let you further customize the color transform. Lastly, a Conversion Table field and Browse button let you add an additional LUT as part of this operation.



The Log-Lin LUT Editor

LUT Processing Order

In elaborate workflows, facilities may apply multiple LUTs in a row before the image is seen. The order of these is important since each LUT delivers different outputs. For instance, for a Cineon file in Log color space you may often apply three LUTs. First a Log to Lin conversion, followed by a Fusion View LUT to apply a color calibration, and a third one to correct it for display on an sRGB monitor, or replace the last with a 3D DCP LUT if you are viewing on a projector.

When you select a node to be displayed, the image produced is processed before it is shown in the viewers. The processing order is slightly different for 2D images and 3D scenes.

2D images first have the image LUT applied, and the result is composited over the checker underlay. 3D scenes are instead rendered with OpenGL.



The order of processing for 2D images and 3D scenes

For either 2D or 3D, the result may be drawn to an offscreen buffer where a Buffer LUT can be applied, along with dithering, a full view checker underlay, and any stereo processing. The final result is then drawn to the viewer and any onscreen controls are drawn on top.

Applying Multiple LUTs

The viewer contextual menu can be used to apply multiple image LUTs into a processing chain.

To apply an additional LUT, do the following:

1. Right-click anywhere within the viewer.
2. From the viewer's contextual menu, choose LUT > Add New.
3. From the Add New submenu, choose a LUT to add.

To remove a LUT other than the first LUT, do the following:

1. Right-click anywhere within the viewer.
2. From the viewer's contextual menu, choose LUT > Delete.
3. From the Delete submenu, choose a LUT to remove.

A complete stacked LUT configuration can be saved to and loaded from a .viewlut file, as described below.

Saving Custom LUTs

There are a variety of ways to create and use different viewer LUTs in Fusion. You can save LUTs when you save viewer settings, you can import LUTs that have been exported from Fusion or other applications, and you can open any one of the various supported LUT file types. In addition, you can use the standard nodes in Fusion to create macros, which can then be saved and used as a LUT.

LUT Settings

The most straightforward way to save a LUT you have created using the Fusion View LUT Editor is to use the LUT > Save menu found in the viewer contextual menu. The settings are saved as an ASCII file with the extension .viewlut in the LUTs folder. Any files with this extension found in that folder will appear in the Image LUT menus for ease of loading. You can also load the settings that are not found in the menu by choosing LUT > Load from the viewer's contextual menu.

Using Viewer Settings

If you've modified a LUT, choosing Settings > Save New from the viewer's contextual menu will save all the viewer's settings, including all LUT curves and gain/gamma values. You can save these under different names, and each settings file can be reloaded at any time by choosing Settings > filename from the viewer's contextual menu. Choosing Save Default from the same menu will make these settings the standard for all new comps.

Using LUT Curves

The Viewer LUT Edit dialog can be used to import and export LUT curves. You can export the LUT curves as either ASCII or Saved format. The ASCII (.alut) file format is useful for sharing LUT curves with other software, whereas the Saved (.lut) file format is preferred for Fusion, as it is more compact, accurate, and allows further editing.

To export a LUT, do the following:

1. Click the viewer LUT button to enable it.
2. Click the viewer LUT menu, and then choose Edit.
3. Right-click on the LUT Curve Editor, and then choose Export LUT.
4. Select a LUT format at the bottom of the file browser window.
5. Enter a name for the LUT and click Save.

The Import LUT option will load LUT files back into the Curve Editor, or alternatively, if the file has been saved in Fusion's LUTs folder, it will appear in the LUT drop-down menu list.

> **TIP:** This is one way to move LUTs between viewers or to and from the Color Curves node or any other LUT Editor in Fusion.

LUT Files

Any supported LUT files in the LUTs folder can be used by choosing them either from the LUT drop-down menu or the viewer's contextual menu. This includes 1D and 3D LUTs such as Fusion's .lut, .alut and .alut3 formats, as well as .cube, .shlut, .look, .3dl, and .itx formats. This is a convenient way to access standard format LUT files for different projects.

Settings and Macros

Since LUTs are a form of color correction, you can also use any node, macro, or group of nodes as a viewer LUT.

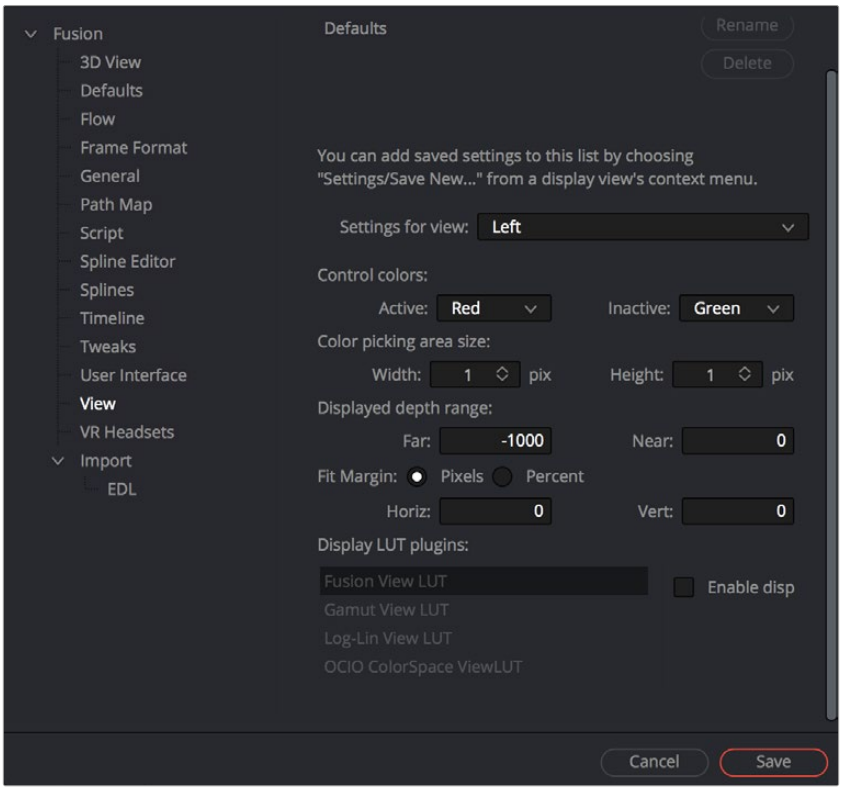
To use a node, group, or macro as a viewer LUT, do the following:

1. Select the node, group, or macro.
2. Right-click over the selected node, and then choose Settings > Save As from the menu.
3. In the file browser, go to the LUTs folder as set in Preferences > Global > Path Map > LUTS.
4. Click Save to save the .settings file.

This allows almost any combination of nodes to be used as a viewer LUT. This is the most flexible approach but is also potentially the slowest. The LUT nodes must be rendered solely on the CPU, whereas other methods are GPU-accelerated.

Setting a Default LUT

The default LUT applied when a new composition is created can be assigned in the Viewer panel of the Fusion Settings window. Clicking the Enable Display LUT checkbox allows you to select a LUT from the Display LUT plugins list.



The LUT default settings found in the View panel of the Fusion Settings window

Viewer Preferences and Settings

The default settings for each viewer can be changed using the Viewer panel in the Preferences. The position and size of each floating viewer can also be saved using the Layout menu in the Preferences.

Viewer Settings

It is often preferable to switch between entirely different viewer configurations while working. For example, while keying, the image may be in the main viewer, and the alpha channel may be in a subview. Viewer settings toward the end of a project may consist of the histogram, vectorscope, and waveform, as well as the image in a view set to Quad view.

Fusion provides the ability to quickly load and save viewer settings to help reduce the amount of effort required to change from one configuration to another.

To save a viewer setting, do the following:

1. Right-click over the viewer you want to save.
2. From the contextual menu, choose Setting > Save New.
3. Enter a name for the settings and click Save.

To load a viewer setting, do the following:

1. Right-click over the viewer you want to load a setting into.
2. From the contextual menu, choose Settings > filename.

Loading and Saving Defaults for a Viewer

The viewer can save new defaults and be returned to its defaults using the Load Defaults and the Save Defaults options in the Settings portion of the View contextual menu.

The Viewer Options Menu

The Options menu of the viewer contains several ways you can customize the look and behavior of the viewer. Many of these options are also in the viewer contextual menu.

Show Controls

When onscreen controls are not necessary or are getting in the way of evaluating the image, you can temporarily hide them using the Show Controls option. This option is toggled using Command-K.

Checker Underlay

The Checker Underlay shows a checkerboard beneath transparent pixels to make it easier to identify transparent areas. This is the default option for 2D viewers. Disabling this option replaces the checkerboard with black.

Show Pixel Grid

Enabling this option will show a light black grid that outlines the exact boundaries of pixels in the image when the image is scaled past a certain threshold. The default is Off.

Smooth Resize

The Smooth Resize option uses a smoother bilinear interpolated resizing method when zooming into an image in the viewer. When Smooth Resize is disabled, scaling uses the nearest neighbor method and shows noticeable aliasing artifacts but is more useful for seeing the actual pixels of the viewed image when you zoom all the way down to a pixel level since there is no interpolation. This option is enabled by default and can be toggled by clicking on the SmR button in the viewer toolbar.

Show Square Pixels

Depending on the frame format preferences and the type of footage loaded, many images may have pixels that are rectangular instead of square. Both the NTSC and PAL video standards, as well as some anamorphic film formats, use rectangular pixels. A computer monitor uses perfectly square pixels. To compensate for this, aspect correction is automatically performed when viewing non-square pixels. This prevents non-square pixel images from appearing squashed or stretched in the viewer.

You can enable the Show Square Pixels option to override the aspect correction. Show Square Pixels can also be toggled on and off using the 1:1 button in the viewer toolbar.

Gain/Gamma

Exposes or hides a simple pair of Gain and Gamma sliders that let you adjust the viewed image. Especially useful for “gamma slamming” a composite to see how well it holds up with a variety of gamma settings. Defaults to no change.

360˚ View

Sets the Fusion page viewer to properly display spherical imagery in a variety of formats, selectable from this submenu. Disable toggles 360 viewing on or off, while Auto, LatLong, Vert Cross, Horiz Cross, Vert Strip, and Horiz Strip let you properly display different formats of 360° video.

Locking the Viewer (Command-L)

You can lock a viewer to prevent it from updating. The node that’s loaded into that viewer still processes and the new image is queued for display in the viewer, but until you unlock it, the viewer does not update. By default, the viewer is unlocked.

Additional Viewer Options

There are additional commands when you right-click anywhere within a viewer and choose from the generically named Options submenu.

Alpha Overlay

When you enable the alpha overlay, the viewer will show the alpha channel overlaid on top of the color channels. This can be helpful when trying to see where one image stops and another begins in a composite. This option is disabled by default.

Overlay Color

When you turn the alpha overlay on, the default color is to show white for the area the alpha covers. There are times when white does not show clearly enough, depending on the colors in the image. You can change the color by choosing a color from the list of Overlay Color options.

Follow Active

Enabling the Follow Active option will cause the viewer to always display the currently active node in the Node Editor. This option is disabled by default, so you can view a different node than what you control in the Control Panel.

Show Controls

When onscreen controls are not necessary or are getting in the way of evaluating the image, you can temporarily hide them using the Show Controls option. This option is toggled using Command-K.

Show Full Color Range

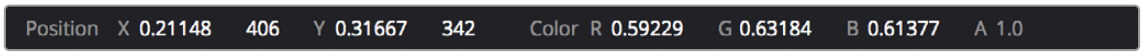
When working with floating-point images, you will occasionally need to visualize the values that fall outside the normal luminance range. Enabling the Show Full Color Range option using the toolbar button automatically normalize any image displayed in the viewer. Normalization causes the brightest pixel in a color channel to be mapped to a value of 1.0 (white) and the darkest pixel to be mapped to a value of 0.0 (black). Midrange values are scaled appropriately to fit within that range. It is also useful when viewing Z-buffer or other auxiliary channels, which often use value ranges far different from those in the color channels.

Show Labels

The Show Labels option lets you toggle the display of the text that sometimes accompanies onscreen controls in the viewer without disabling the functions that are showing those overlays, and without hiding the onscreen controls themselves.

Status Bar Information

The status bar at the bottom of the Fusion window provides the exact RGBA and Z values for the pixel beneath the pointer when it's hovering within one of the viewers. Additional information about the X and Y coordinates of the cursor and the exact pixel position are also displayed.



The status bar showing coordinates and color information

Chapter 8

Editing Parameters in the Inspector

The Inspector is where you adjust the parameters of each node to do what needs to be done. This chapter covers the various node parameters and methods for working with the available controls.

Contents

Overview of the Inspector 230
The Tools and Modifiers Panels 231
Customizing the Inspector 231
Inspector Height 231
Inspector Display Preferences 232
Opening Nodes in the Inspector 233
Pinning Multiple Nodes in the Inspector 234
Hiding Inspector Controls 235
Using the Inspector Header 235
Selecting and Viewing Nodes in the Inspector 235
Using Header Controls 236
Versioning Nodes 236
Parameter Tabs 237
The Settings Tab 237
Inspector Controls Explained 241
Fusion Slider Controls 241
Thumbwheel 241
Range Controls 242
Checkboxes 242
Drop-Down Menus 242
Button Arrays 243
Color Chooser and Picker 243
Modifiers 248

Animating Parameters in the Inspector 249
Removing Animation From a Parameter 250
Attaching a Parameter to an Existing Animation Curve 250
Connecting Parameters 250
Connecting Parameters by Publishing 250
Connecting Parameters by Pick Whipping 251
Contextual Menus 251
Customizing Node Parameters with User Controls 252
An Example of Customizing Directional Blur 253

Overview of the Inspector

While the creation and connection of nodes in the Node Editor determines the tools and order of operations that make up a composition, the Inspector is where you adjust the various parameters inside each node to do what needs to be done.



Inspector displays the Brightness Contrast controls

This chapter covers methods for opening node parameters in the Inspector to edit them in different ways according to the type of available controls.

To display the Inspector:

Click the Inspector button on the UI toolbar.

The Tools and Modifiers Panels

The Inspector is divided into two overall panels.

- The Tools panel is where the parameters of selected nodes appear so you can edit them.
- The Modifiers panel is where you edit optional extensions to the tool's standard toolset as well as automated expressions that you can attach to individual parameters to create animated effects. Additionally, certain nodes such as the Paint node generate data such as Strokes, which are saved in the Modifiers panel.



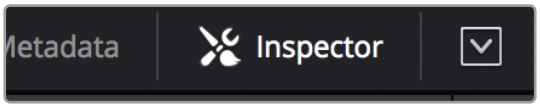
Modifiers displayed in the Modifiers panel

Customizing the Inspector

You can customize how the Inspector is presented in a variety of ways.

Inspector Height

A small arrow button at the far right of the UI toolbar lets you toggle the Inspector between full-height and half-height views, depending on how much room you need for editing parameters.

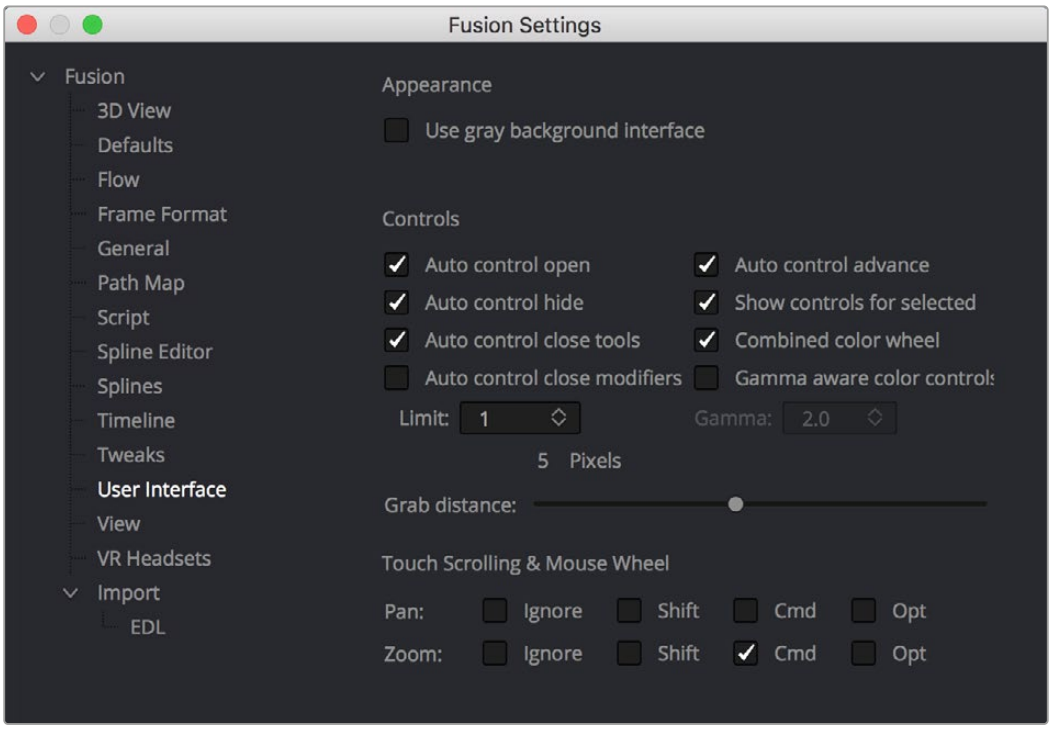


The Maximize button on the left side of the Inspector

In maximized height mode, the Inspector takes up up the entire right side of the UI, letting you see every control that a node has available, or creating enough room to see the parameters of two or three pinned nodes all at once. In half-height mode, the top of the Inspector is aligned with the tops of the viewers, expanding the horizontal space that's available for the Node Editor.

Inspector Display Preferences

By default, you see only selected nodes in the Inspector, and only the Active node is expanded to show its controls. You can change this behavior by choosing Fusion > Fusion Settings in the Fusion page or File > Preferences in Fusion Studio and opening the User Interface panel. In the User Interface, checkboxes manage the display of controls.



Control preferences in the User Interface category

- **Auto Control Open:** When enabled (the default), whichever node is active automatically opens its controls in the Inspector. When disabled, selecting an active node opens that node's Inspector header in the Inspector, but the parameters remain hidden unless you click the Inspector header.
- **Auto Control Hide:** When enabled (the default), only selected nodes are visible in the Inspector, and all deselected nodes are automatically removed from the Inspector to reduce clutter. When disabled, parameters from selected nodes remain in the Inspector, even when those nodes are deselected, so that the Inspector accumulates the parameters of every node you select over time.
- **Auto Control Close Tools:** When enabled (the default), only the parameters for the active node can be exposed. When disabled, you can open the parameters of multiple nodes in the Inspector if you want.
- **Auto Controls for Selected:** When enabled (the default), selecting multiple nodes opens multiple control headers for those nodes in the Inspector. When disabled, only the active node appears in the Inspector; multi-selected nodes highlighted in white do not appear.

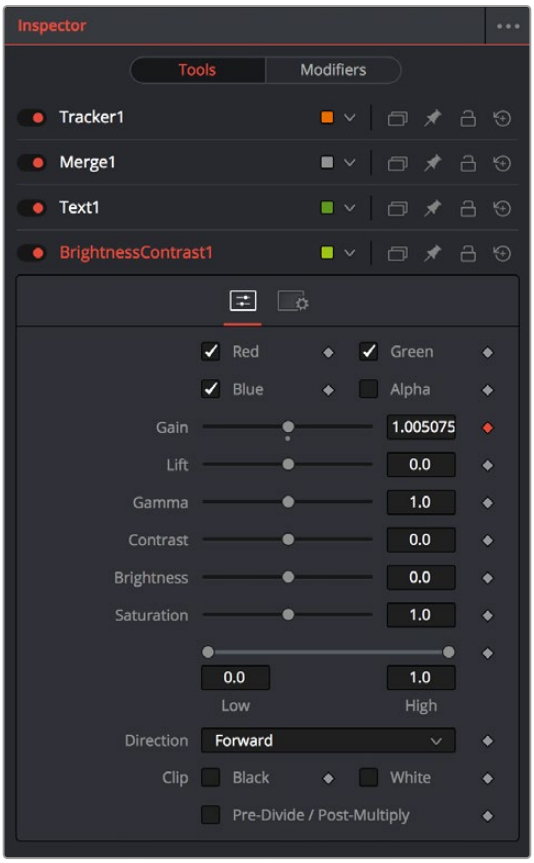
Opening Nodes in the Inspector

Before you can edit a node's parameters, you need to open it in the Inspector.

To display a node's controls in the Inspector:

Select one or more nodes from the Node Editor, Keyframes Editor, or Spline Editor.

When you select a single node so that it's highlighted orange in the Node Editor, all of its parameters appear in the Inspector. If you select multiple nodes at once, Inspector headers appear for each selected node (highlighted in white in the Node Editor), but the parameters for the active node (highlighted in orange) are exposed for editing.

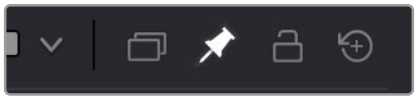


Opening multiple nodes in the Inspector

Only one node's parameters can be edited at a time, so clicking another node's Inspector header opens that node's parameters and closes the parameters of the previous node you were working on. This also makes the newly opened node the active node, highlighting it orange in the Inspector.

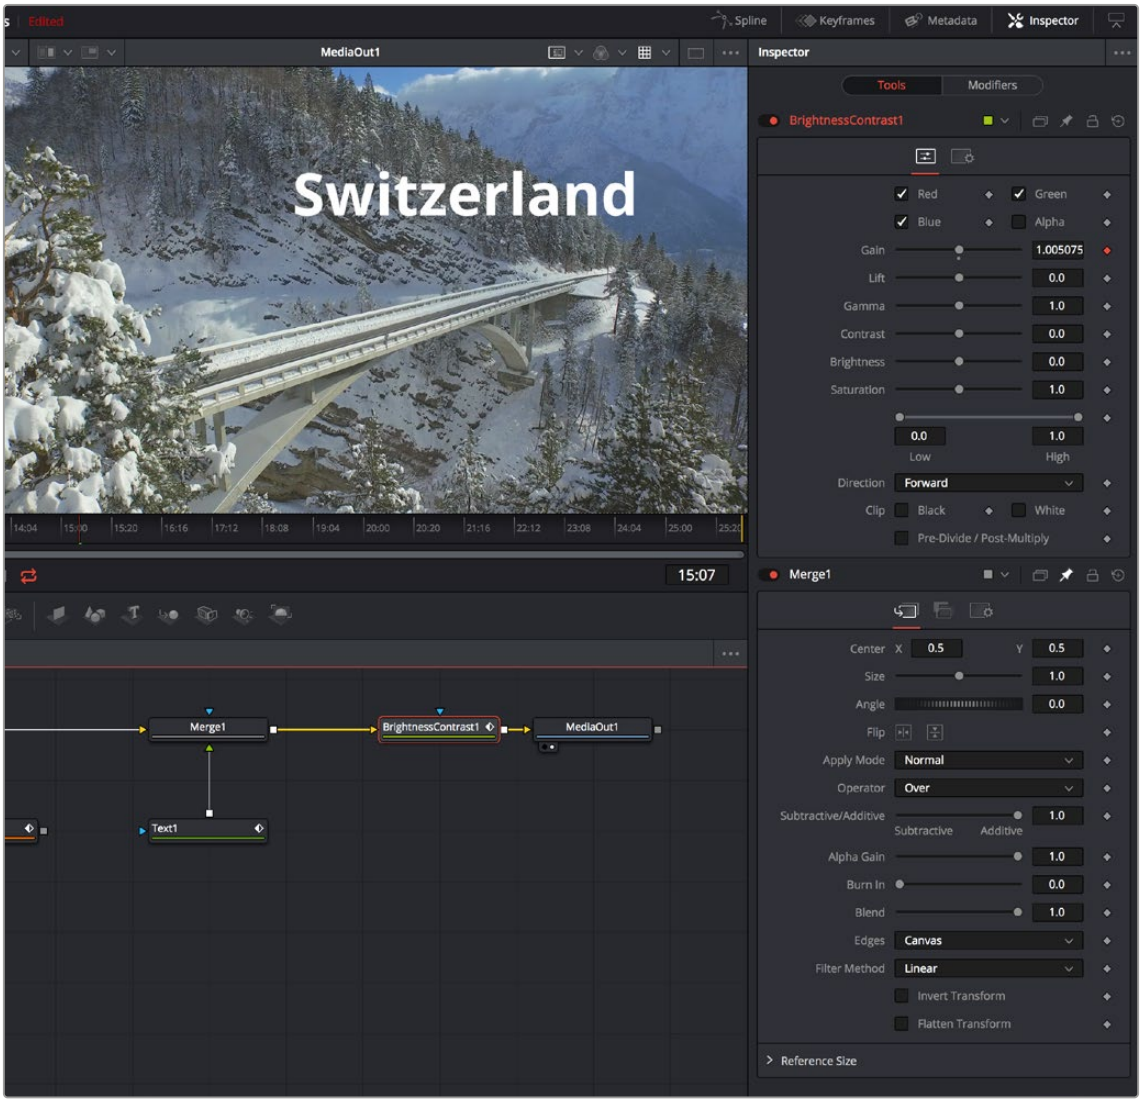
Pinning Multiple Nodes in the Inspector

For instances where you need to work quickly by editing the parameters of multiple nodes at the same time, you can use the Pin button in the Inspector header of nodes in the Inspector to keep those parameters exposed in the Inspector, regardless of whether that node is selected and active.



The Pin button of a node's Inspector header in the Inspector

While the Pin button is on, that node's parameters remain open in the Inspector. If you select another node in the Node Editor, that node's parameters appear beneath any pinned nodes.



A pinned node on the bottom, with a selected node at the top

You can have as many pinned nodes in the Inspector as you like, but the more you have, the more likely you'll need to scroll up or down in the Inspector to get to all the parameters you want to edit. To remove a pinned node from the Inspector, just turn off its Pin button in the Inspector header.

Hiding Inspector Controls

If you like, Inspector parameters for specific nodes can be hidden so they never appear, even when that node is selected. This can be useful for preventing accidental changes by you or other compositors who may be working on a composition in situations where you don't want to lock the node.

To Toggle the Inspector controls for a node or or off:

Right-click on the node in the Node Editor, or on the Inspector header, and choose Modes > Show Controls from the contextual menu.

Using the Inspector Header

When you select a node, it populates the Inspector with a title bar, or Inspector header, that displays that node's name as well as other controls that govern that node. A node's Inspector header itself has a variety of controls, but clicking (or double-clicking) on an Inspector header also exposes that node's parameters.



A node's Inspector header

When you select multiple nodes at once, you'll see multiple headers in the Inspector. By default, only the parameters for the active node (highlighted orange in the Node Editor) can be opened at any given time, although you can change this behavior in Fusion's Preferences.

Selecting and Viewing Nodes in the Inspector

Inspector headers are click targets for selecting nodes, opening and closing node parameters, and other things.

Methods of using headers:

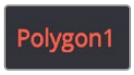
- **To select a node using the Inspector header:** When multiple nodes are selected, you can make a node the active node by clicking its Inspector header in the Inspector. As the actively selected node, the Inspector header and the corresponding node in the Node Editor are highlighted orange, and its parameters are exposed.
- **To load a node into the viewer using the Inspector header:** You can view a node by dragging its header into one of the viewers.
- **To view a node's splines with the control header:** If you want to view the animated curves of a node in the Spline Editor, you can add them by dragging the Inspector header into the Spline Editor. All animated splines for the parameters of that node will automatically be displayed.

Using Header Controls

The controls found in each node's Inspector header makes it fast to do simple things.



To turn nodes off and on: Each Inspector header has a toggle switch to the left of its name, which can be used to enable or disable that node. Disabled nodes pass image data from the previous upstream node to the next downstream node without alteration.



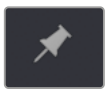
To change the Inspector header name: The name of the node corresponding to that Inspector header is displayed next. You can change the name by right-clicking the Inspector header to expose contextual menu commands similar to those found when you right-click a node in the Node Editor and choosing Rename. Alternatively, you can click an Inspector header and press F2 to edit its name. A Rename dialog appears, where you can enter a new name and click OK (or press Return).



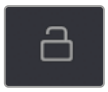
To color-code nodes: A color pop-up menu lets you color code with one of 16 colors. Choose Clear Color if you want to return that node to the default color.



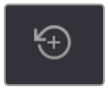
To version nodes: Turning on the Versions button displays a Version bar with six buttons. Versioning is described in the following section.



To pin Inspector controls: Clicking the Pin button "pins" that node's parameters in the Inspector so they remain in place, even if you deselect that node. You can have as many pinned nodes as you like in the Inspector, but the more you have, the more likely you'll be scrolling up and down the Inspector to navigate all the available parameters.



To lock nodes: Clicking the Lock button locks that node so no changes can be made to it.



To reset Inspector controls: The rightmost button in the Inspector header is a Reset button that resets the entire node to the default settings for that node.

Versioning Nodes

Each button is capable of containing separate parameter settings for that node, making it easy to save and compare up to six different versions of settings for each node. All versions are saved along with the node in the Node Editor for future use.

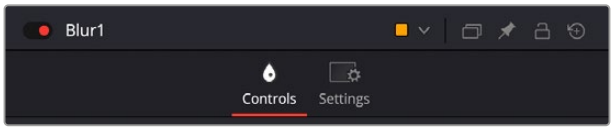


The Version bar, underneath an Inspector header, with versions enabled

An orange underline indicates the currently selected version, which is the version that's currently being used by your composition. To clear a version you don't want to use any more, right-click that version number and choose Clear from the contextual menu.

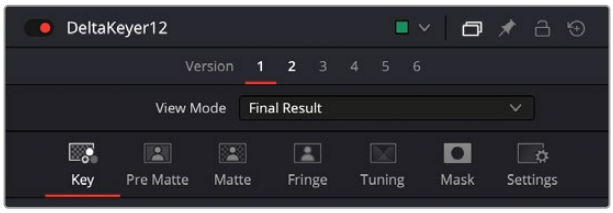
Parameter Tabs

Underneath the Inspector header is a series of panel tabs, displayed as thematic icons. Clicking one of these icons opens a separate tab of parameters, which are usually grouped by function. Simple nodes, such as the Blur node, consist of two tabs where the first contains all of the parameters relating to blurring the image, and the second is the Settings tab.



The parameter tabs of the Blur node

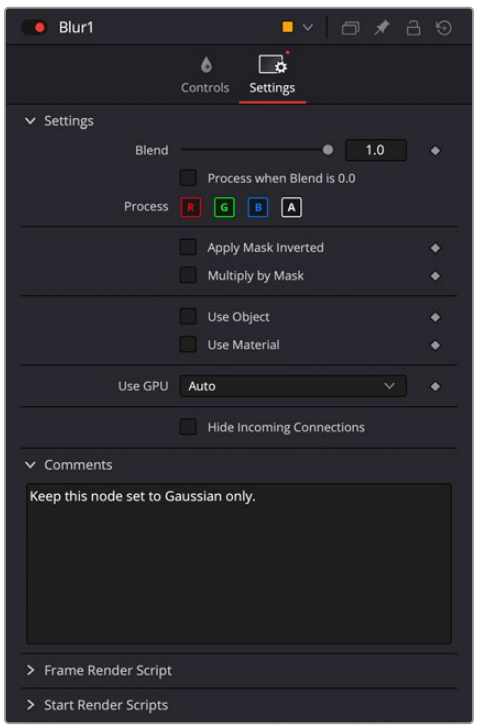
More complicated nodes have more tabs containing more groups of parameters. For example, the Delta Keyer has seven tabs: separating Key, Pre-Matte, Matte, Fringe, Tuning, and Mask parameters, along with the obligatory Settings tab. These tabs keep the Delta Keyer from being a giant scrolling list of settings and make it easy to keep track of which part of the keying process you're finessing as you work.



The parameter tabs of the Delta Keyer node

The Settings Tab

Every node that comes with Fusion has a Settings tab. This tab includes a set of standard controls that appear for nearly every node, although some nodes have special Settings tab controls that others lack.



The Settings tab in the Inspector

The following controls are common to most nodes, although some are node-specific. For example, Motion Blur settings have no purpose in a Color Space node.

Blend

The Blend control is found in all nodes, except the Loader, MediaIn, and Generator nodes. It is used to blend between the node's unaltered image input and the node's final processed output. When the blend value is 0.0, the outgoing image is identical to the incoming image. Ordinarily, this will cause the node to skip processing entirely, copying the input straight to the output. The default for this node is 1.0, meaning the node will output the modified image 100%.

Process When Blend is 0.0

This checkbox forces the node to process even when the input value is zero and the image output is identical to the image input. This can be useful on certain nodes or third-party plugins that store values from one frame to the next. If this checkbox is disabled on nodes that operate in this manner, the node will skip being processed when the Blend is set to 0, producing incorrect results on subsequent frames.

Red/Green/Blue/Alpha Channel Checkboxes

Most nodes have a set of RGBA boxes in the Settings tab. These selectable boxes let you exclude any combination of these channels from being affected by that node.



The channel limiting boxes in the Settings panel of a Transform node set so that only the green channel is affected.

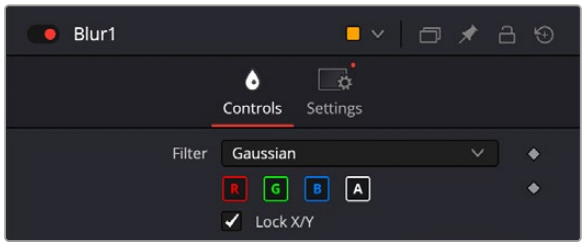
For example, if you wanted to use the Transform node to affect only the green channel of an image, you can turn off the Red, Blue, and Alpha checkboxes. As a result, the green channel is processed by this operation, and the red, blue, and alpha channels are copied straight from the node's input to the node's output, skipping that node's processing to remain unaffected.



Transforming only the green color channel of the image with a Transform effect

Skipping Channel Processing

Under the hood, most nodes actually process all channels first, but afterward copy the input image to the output for channels that have been unchecked. Modern workstations are so fast that this isn't usually noticeable, but there are some nodes where deselecting a channel actually causes that node to skip processing that channel entirely. Nodes that operate this way have a linked set of Red, Green, Blue, and Alpha boxes on another tab in the node.



Channel boxes on the Controls tab of the Blur node indicate that disabled channels won't be processed at all, to save rendering time..

In these cases, the Common Control channel boxes are instanced to the channel boxes found elsewhere in the node. Blur, Brightness/Contrast, Erode/Dilate, and Filter are examples of nodes that all have RGBY checkboxes in the main Controls tab of the Inspector, in addition to the Settings tab.

Apply Mask Inverted

When the Apply Mask Inverted checkbox is enabled, masks attached to the Effect Mask input of that node are inverted.

> **TIP:** The Apply Mask Inverted checkbox option operates only on effects masks, not on garbage masks.

Multiply By Mask

Selecting this option will cause the RGB values of the masked image to be multiplied by the Mask channel's values. This will cause all pixels of the image not included in the mask (i.e., those set to 0) to become black. This creates a premultiplied image.

Use Object/Use Material (For Masking)

Some 3D animation and rendering software can output to file formats that support auxiliary channels. Notably, the OpenEXR file format supports Object ID and Material ID channels, either of which can be used as a mask for an effect. This checkbox determines whether the channels will be used if they are available. The specific Material ID or Object ID affected is chosen using the next set of controls.

- **Sample Controls:** The Sample Controls are only displayed once the Use Object or Use Material checkbox is enabled. These controls select which ID is used to create a mask from the Object or Material channels saved in the image. You use the Sample button to grab IDs from the image in the viewer, the same way you use the Color Picker to select a color, by holding down the left mouse button on the Sample button, then dragging over to the viewer to the part of the image you want to select. The image or sequence must have been rendered from a 3D software package with those channels included.
- **Correct Edges:** The Correct Edges checkbox is only displayed once the Use Object or Use Material checkbox is enabled. When the Correct Edges checkbox is enabled, the Coverage and

Background Color channels are used to separate and improve the effect around the edge of the object. When disabled (or no Coverage or Background Color channels are available), aliasing may occur on the edge of the mask.

Motion Blur

For nodes that are capable of introducing motion, such as Transform nodes, Warp nodes, and so on, the Motion Blur checkbox toggles the rendering of motion blur on or off for that node. When this checkbox is enabled, the node's predicted motion is used to produce the blur caused by a virtual camera shutter. When the control is disabled, no motion blur is created.

When Motion Blur is disabled, no additional controls are displayed. However, turning on Motion Blur reveals four additional sliders with which you can customize the look of the motion blur you're adding to that node.

- **Quality:** Determines the number of samples used to create the blur. The default quality setting of 2 will create two samples on either side of an object's actual motion. Larger values produce smoother results but will increase the render time.
- **Shutter Angle:** Controls the angle of the virtual shutter used to produce the Motion Blur effect. Larger angles create more blur but increase the render times. A value of 360 is the equivalent of having the shutter open for one whole frame exposure. Higher values are possible and can be used to create interesting effects. The default value for this slider is 100.
- **Center Bias:** Modifies the position of the center of the motion blur. Adjusting the value allows for the creation of trail-type effects.
- **Sample Spread:** Adjusting Sample Spread modifies the weight given to each sample. This affects the brightness of the samples set with the Quality slider.

Scripting

Scripting fields are present on every node and contain one or more editable text fields that can be used to add scripts that process when that node is rendering. For more information on the contents of this tab, please consult the Scripting documentation.

Comments

A Comments field is found on every node and contains a single text field that is used to add comments and notes to that node. To enter text, simply click within the field to place a cursor, and begin typing.

When a note is added to a node, the comments icon appears in the Control Header and can be seen in a node's tooltip when the cursor is placed over the node in the Node Editor. The contents of the Comments tab can be animated over time, if required.

Additional controls appear under this tab if the node is a Loader. For more information, *see Chapter 44, "Generator Nodes."* in the Fusion Reference Manual.

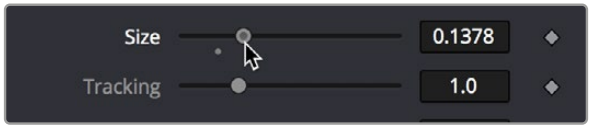
Inspector Controls Explained

Although a few nodes use fully customized interface elements that are unique to only that node, the vast majority of nodes use a mix of sliders, angle wheels, and checkboxes. This section explains how to use these controls.

Fusion Slider Controls

Slider Controls are used to select a single value from a range of values. You change the value by dragging the slider or entering a value into the edit box. This is fairly standard behavior for sliders. However, there is additional functionality that can increase your productivity when making changes with sliders.

Clicking on the gutter to the left or right of the handle will increase or decrease the value. Holding Command while clicking on the gutter will adjust the values in smaller increments. Holding Shift while clicking will adjust the value in larger increments.



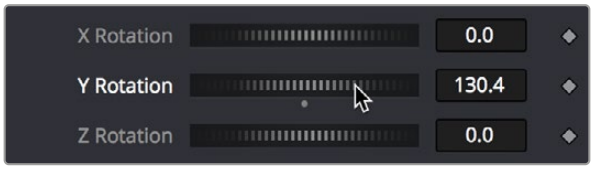
Hold Command while clicking in the gutter
to move in smaller increments

While slider controls use a minimum and maximum value range, entering a value in the edit box outside that range will often expand the range of the slider to accommodate the new value. For example, it is possible to enter 500 in a Blur Size control, even though the Blur Size sliders default maximum value is 100. The slider will automatically adjust its maximum displayed value to allow entry of these larger values.

If the slider has been altered from its default value, a small circular indicator will appear below the gutter. Clicking on this circle will reset the slider to its default.

Thumbwheel

A Thumbwheel control is identical to a slider except it does not have a maximum or minimum value. To make an adjustment, you drag the center portion left or right or enter a value directly into the edit box. Thumbwheel controls are typically used on angle parameters, although they do have other uses as well.



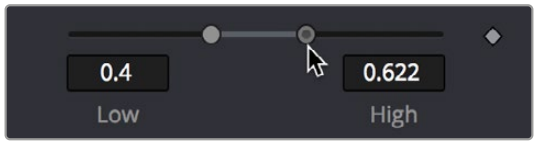
Thumbwheel controls for X,Y, and Z rotation with
arrows on either end for fine-tuning adjustments

Once the thumbwheel has been selected, you can use the Up and Down Arrows on your keyboard to further adjust the values. As with the slider control, the Command and Shift keys can be used to increase or decrease the change in value in smaller or larger increments.

If the thumbwheel has been altered from its default value, a small circular indicator will appear below above the thumbwheel. Clicking on this circle will reset the thumbwheel to its default.

Range Controls

The Range controls are actually two separate controls, one for setting the Low Range value and one for the High Range value. To adjust the values, drag the handles on either end of the Range bar. To slide the high and low values of the range simultaneously, drag from the center of the Range bar. You can also expand or contract the range symmetrically by holding Command and dragging either end of the Range bar. You find Range controls on parameters that require a high and low threshold, like the Matte Control, Chroma Keyer, and Ultra Keyer nodes.

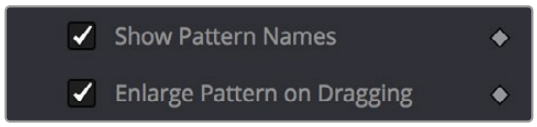


A Matte Threshold Range control

TIP: You can enter floating-point values in the Range controls by typing the values in using the Low and High numeric entry boxes.

Checkboxes

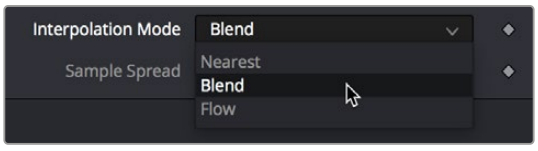
Checkboxes are controls that have either an On or Off value. Clicking on the checkbox control will toggle the state between selected and not selected. Checkboxes can be animated, with a value of 0 for Off and a value of 1.0 or greater for On.



Checkboxes used to select options for tracking

Drop-Down Menus

Drop-down menus are used to select one option from a menu. Once the menu is open, choosing one of the items will select that entry. When the menu is closed, the selection is displayed in the Inspector.



Drop-down menu in the TimeSpeed node

Drop-down menu selections can be animated, with a value of 0 representing the first item in the list, 1 representing the second, and so forth.

Button Arrays

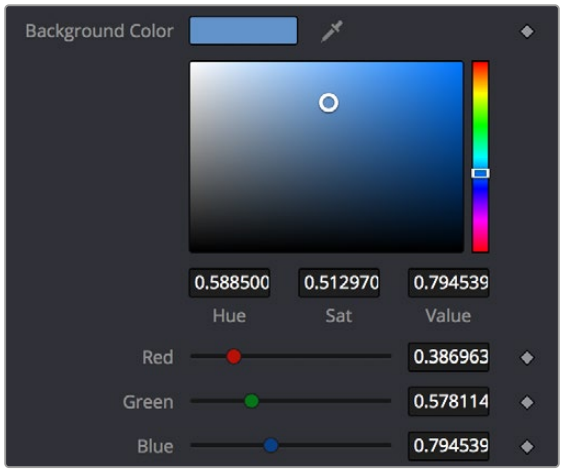
Button arrays are groups of buttons that allow you to select from a range of options. They are almost identical in function to drop-down menu controls, except that in the case of a button array it is possible to see all of the available options at a glance. Often button arrays use icons to make the options more immediately comprehensible.



The Lens Type button array in the Defocus node

Color Chooser and Picker

The Color panel is displayed wherever a parameter requires a color as its value, such as the Fill or Outline color in the Text+ node. The selected color is shown in a swatch with an Eyedropper to its right, and below the swatch is the Color Chooser.



The Color panel with transparency preview

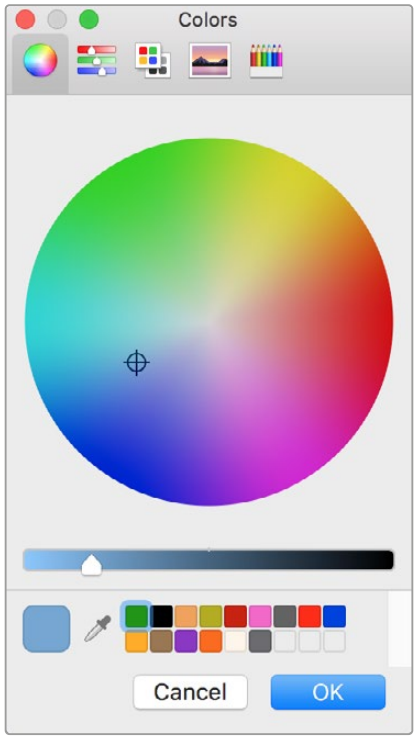
The Color panel is extremely flexible and has four different techniques for selecting and displaying colors.

TIP: Color can be represented by 0–1, 0.255, or 0–65000 by setting the range you want in the Preferences > General panel.

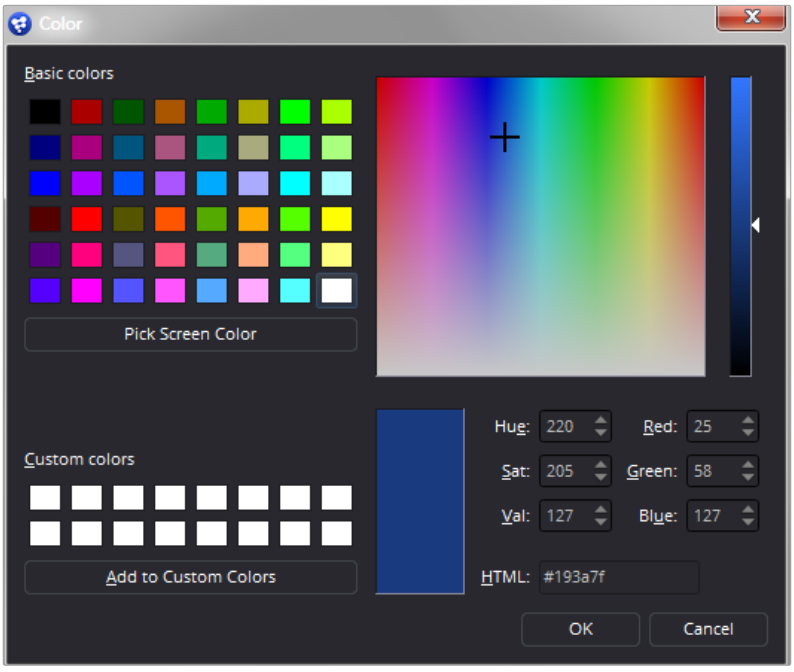
macOS and Windows Color Nodes

Clicking on the color swatch will display the operating system's standard Color Selection node.

Each operating system has a slightly different layout, but the general idea is the same. You can choose a color from the swatches provided—the color wheel on macOS, or the color palette on Windows. However you choose your color, you must click OK for the selection to be applied.



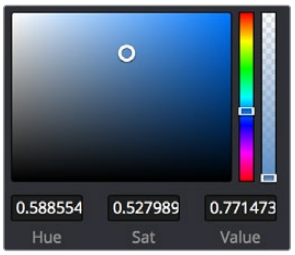
macOS Colors panel



Windows Color dialog

The Color Chooser

You also have access to the built-in color chooser, which includes sections for choosing grayscale values, as well as the currently chosen hue with different ranges of saturation and value. A hue bar and alpha bar (depending on the node) let you choose different values.



The color chooser in the Background node

Picking Colors from an Image

If you are trying to match the color from an image in the viewer, you can hold down the cursor over the Eyedropper, and then drag the pointer into the viewer. The pointer will change to an Eyedropper, and a pop-up swatch will appear above the cursor with the color you are hovering over and its values. When you are over the color you want, release the mouse button to set the color.



The Eyedropper with color swatch

The Color Picker normally selects from a single pixel in the image, but you can adjust the size of the selection by dragging into the viewer with the Eyedropper, and then holding Command and dragging out a rectangle for the sample size you want. The size change applies to all Color Pickers until the size is changed again.

Gradients

The Gradient Control bar is used to create a gradual blend between colors. The Gradient bar displays a preview of the colors used from start to end. By default, there are two triangular color stops: one on the left that determines the start color, and one on the right that determines the end color.

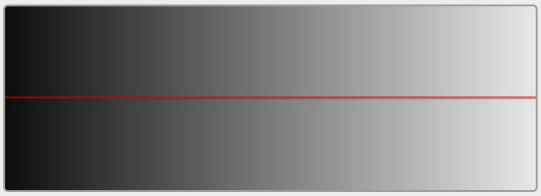


The default Gradient controls

Gradient Type

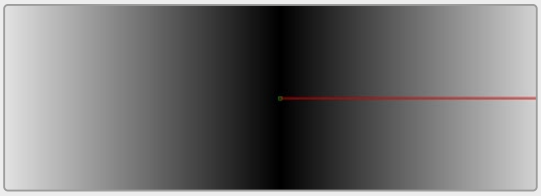
The Gradient Type button array is used to select the form used to draw the gradient.

Linear gradient: Linear draws the gradient along a straight line from the starting color stop to the ending color stop.



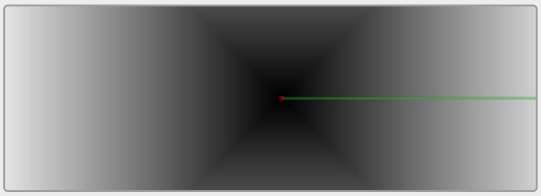
Linear gradient

Reflect gradient: Reflect draws the gradient by mirroring the linear gradient on either side of the starting point.



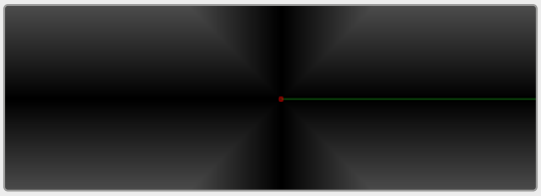
Reflect gradient

Square gradient: Square draws the gradient by using a square pattern when the starting point is at the center of the image.



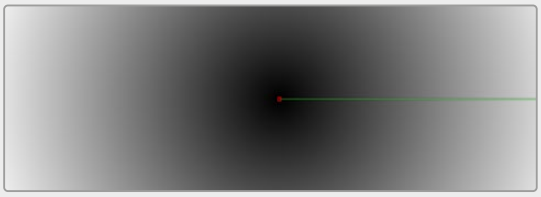
Square gradient

Cross gradient: Cross draws the gradient using a cross pattern when the starting point is at the center of the image.



Cross gradient

Radial gradient: Radial draws the gradient in a circular pattern when the starting point is at the center of the image.



Radial gradient

Angle gradient: Angle draws the gradient in a counter-clockwise sweep when the starting point is at the center of the image.



Angle gradient

Start and End Position

The Start and End Position controls have a set of X and Y edit boxes that are useful for fine-tuning the start and end position of the gradient. The position settings are also represented by two crosshair onscreen controls in the viewer, which may be more practical for initial positioning.

Gradient Colors Bar

The Gradient Colors bar is used to select the blending colors for the gradient. The default two color stops set the start and end colors. You can change the colors used in the gradient by selecting the color stop, and then using the Eyedropper or color wheel to set the new color.

You can add, move, copy, and delete colors from the gradient using the Colors bar.

To add a color stop to the Gradient Colors bar:

1. Click anywhere along the bottom of the Gradient Colors bar.
2. Use the Eyedropper or color wheel to set the color for the color stop.

To move a color stop on the Colors bar:

- Drag a color stop left or right along the Gradient Color bar.

To copy a color stop on the Colors bar:

- Hold Command while you drag a color stop.

To delete a color stop from the Colors bar, do one of the following:

- Drag the color stop up past the Gradient Colors bar.
- Select the color stop, then click the red X button to delete it.

Interpolation Space

The Gradient Interpolation Method pop-up menu lets you select what color space is used to calculate the colors between color stops.

Offset

When you adjust the Offset control, the position of the gradient is moved relative to the start and end markers. This control is most useful when used in conjunction with the repeat and ping-pong modes described below.

Once/Repeat/Ping-Pong

These three buttons are used to set the behavior of the gradient when the Offset control scrolls the gradient past its start and end positions. The Once button is the default behavior, which keeps the color continuous for offset. Repeat loops around to the start color when the offset goes beyond the end color. Ping-pong repeats the color pattern in reverse.

1x1, 2x2, 3x3, 4x4, 5x5

These buttons control the amount of sub-pixel precision used when the edges of the gradient become visible in Repeat mode, or when the gradient is animated. Higher settings will take significantly longer to render but will be more precise.

Gradient Contextual Menu

Gradients have their own contextual menu that you can bring up by right-clicking on the Gradient bar. In the Gradient contextual menu are options for animating, publishing, and connecting one gradient to another. There is also a gradient-specific modifier that builds a custom gradient by sampling colors from the output of a node in the Node Editor.

Modifiers

Modifiers are expressions, calculations, trackers, paths, and other mathematical components that you attach to a parameter to extend its functionality. When a modifier is attached to a parameter, its controls will appear separately in the Inspector Modifiers tab.

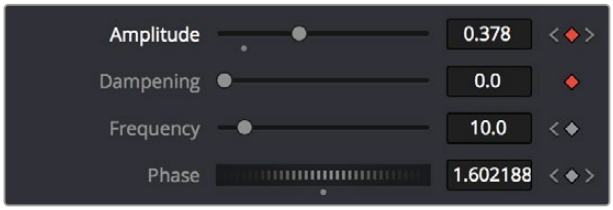
To attach a modifier:

1. Right-click over the parameter to which you want to attach a modifier.
2. Make a selection from the Modifier submenu in the contextual menu.

Animating Parameters in the Inspector

Fusion can keyframe most parameters in most nodes, in order to create animated effects such as animated transforms, rotoscoping with splines, dynamically altering warping behaviors, and so on; the list is endless. For convenience, a set of keyframing controls are available within the Inspector next to each keyframable parameter. These controls are:

- A gray Keyframe button to the right each keyframable parameter. Clicking this gray button creates a keyframe at the current position of the playhead, and turns the button orange.
- When you add a keyframe to a parameter, moving to a new frame and changing the parameter will automatically add a keyframe at the current position.
- Whenever the playhead is sitting right on top of a keyframe, this button turns orange. Clicking an orange Keyframe button deletes the keyframe at that frame and turns the button gray again.
- Small navigation arrows appear to the right and left if there are more keyframes in those directions. Clicking on navigation arrows to the right and left of keyframes jumps the playhead to those keyframes.



Orange Keyframe buttons in the Inspector show there's a keyframe at that frame

Once you've keyframed one or more parameters, If Show Modes/Options has been enabled, the node containing the parameters you keyframed displays a Keyframe badge, to show that node has been animated.



A keyframed node displays a Keyframe badge in the Node Editor

Once you've started keyframing node parameters, you can edit their timing in the Keyframes Editor and/or Spline Editor.

For more information about keyframing in Fusion, *see Chapter 9, "Animating in Fusion's Keyframes Editor,"* in the Fusion Reference Manual.

FUSION

Removing Animation From a Parameter

To remove all keyframes from a parameter:

1. Right-click over the name of the keyframed parameter in the Inspector.
2. Choose Remove “node name:parameter name” from the contextual menu.

> **TIP:** If you change the default spline type from Bézier, the contextual menu will display the name of the current spline type.

Attaching a Parameter to an Existing Animation Curve

Multiple parameters can be connected to the same animation curve. This can be an invaluable timesaver if you are identically animating different parameters in a node.

To connect a second parameter to the same animation curve:

1. Right-click on the second parameter you want to attach.
2. In the contextual menu, hover over the Connect To submenu.
3. In the Connect To submenu, choose the name of the animated parameter.

Connecting Parameters

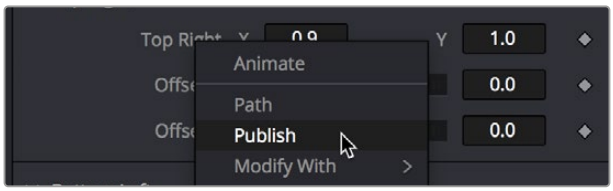
It is often useful to connect two parameters together even without an animation curve. There are two methods you can use.

Connecting Parameters by Publishing

If you want to tie two parameters together so adjusting one adjusts the other, you must connect them together using the Publish menu command on the first parameter and the Connect menu command on the second parameter.

To Publish and Connect parameters:

1. Right-click the name of the parameter you want to publish, and choose Publish from the contextual menu.
2. Right-click on the second parameter you want to attach, and choose the name of the parameter you just published from the Connect To submenu.



The Publish contextual menu

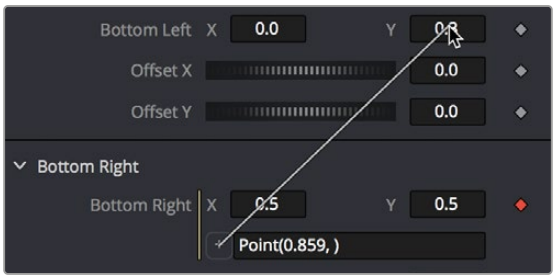
Connecting Parameters by Pick Whipping

You can also use simple expressions to link two parameters together. By using simple expressions via pick whipping, values can be connected and combined visually without the need to publish a value first. The pick whip is a temporary line drawn from one parameter to another in order to create a link between the two.

To link two parameters using a pick whip:

1. Double-click the field of a parameter you want to pick whip to another parameter, type =, and then press the Return key.
2. When Pick Whip controls appear underneath the parameter, drag a “whip” from the Add button to the target parameter.

   Now, adjusting the target parameter automatically adjusts the original parameter.



Pick whipping one parameter to another

TIP: Disabling the Auto Control Close node’s General preference, and then selecting two nodes in the Node Editor will allow you to pick whip two parameters from different nodes.

The Expression field can further be used to add mathematical formulas to the value received from the target parameter.

For more information on Pick Whipping and Expressions, *see Chapter 12, “Using Modifiers, Expressions, and Custom Controls,”* in the Fusion Reference Manual.

Contextual Menus

There are two types of contextual menus you can invoke within the Inspector.

Node Contextual Menus

To display the Node Context menu from the Inspector, right-click on the Inspector header. The node’s contextual menu includes the same menu options that are accessed by right-clicking on a node in the Node Editor.

For more information *see Chapter 5, "Working in the Node Editor."* in the Fusion Reference Manual for more information on these options.

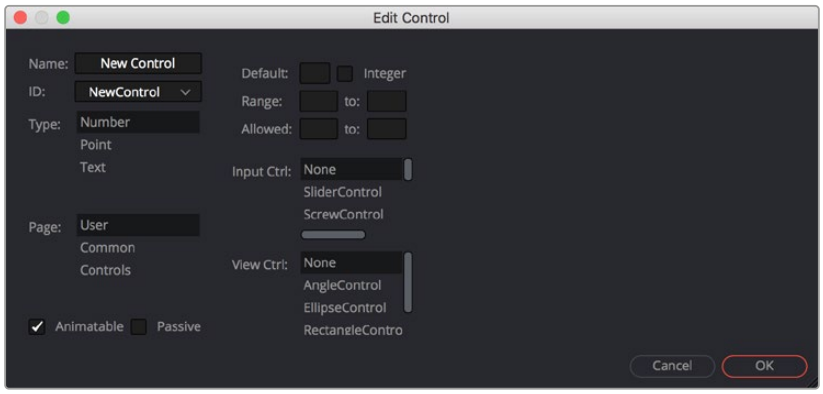
Parameter Contextual Menus

The contextual menu for individual parameters is accessed by right-clicking over the parameter's name, slider, thumbwheel, range control, button array, or other control type. For example, right-clicking on a slider will provide the slider's contextual menu, with options to animate the control or add additional modifiers. Many of these options were described in this chapter.

Customizing Node Parameters with User Controls

The user interface for each node in Fusion is designed to provide access to the parameters in a logical manner. Sometimes, though, you may want to add, hide, or change the controls. This is commonly done for simple expressions and macros, but it can be done for usability and aesthetic reasons for favorites and presets.

User custom controls can be added or edited via the Edit Control dialog. Right-click the name of a node in the Inspector (in the header bar) and choose Edit Control from the contextual menu. A new window will appear, titled Edit Control.



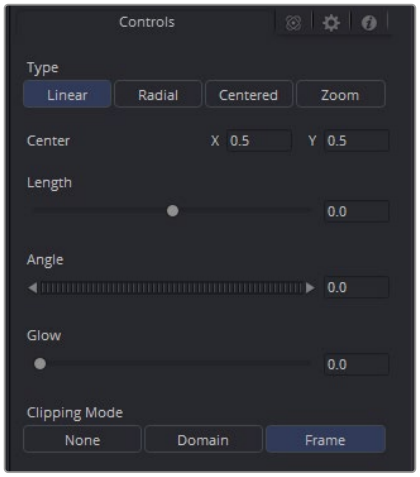
The Edit Control window

In the Input attributes, you can select an existing control or create a new one, name it, define the type, and assign it to a tab. In the Type attributes, you define the input controls, the defaults and ranges, and whether it has an onscreen preview control. The Input Ctrl attributes box contains settings specific to the selected node control, and the View Ctrl attributes box contains settings for the preview control, if any.

All changes made using UserControls are stored in the node instance itself, so they can be copy/pasted, saved to a setting, added to the Bins, or added to your favorites.

An Example of Customizing Directional Blur

In the following example, let's suppose we wanted to create a more intuitive way of controlling a linear blur than using the Length and Angle sliders independently.



Default Directional Blur controls in the Inspector

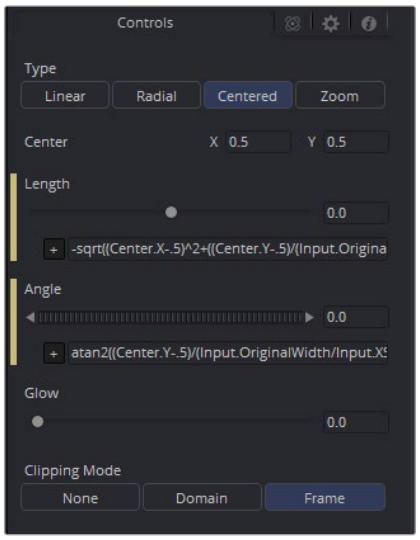
We could use the Center input control, along with its preview control, to set an angle and distance from directly within the viewer using expressions.

1 Right-click the label for the Length parameter, choose Expression from the contextual menu, and then paste the following expression into the Expression field that appears:

```
-sqrt(((Center.X-.5)*(Input.XScale))^2+((Center.Y-.5)*(Input.YScale)*(Input.
Height/Input. Width))^2)
```

2 Next, right-click the label for the Angle parameter, choose Expression from the contextual menu, and then paste the following expression into the Expression field that appears:

```
atan2(Center.Y-.5)/(Input.OriginalWidth/Input.X , .5-Center.X) * 180 / pi
```



Directional Blur controlled by the Center's position

This functions fine, but the controls are confusing. The Center control doesn't work as the center anymore, and it should be named "Blur Vector" instead. The controls for the Length and Angle aren't meant to be edited, so they should be hidden away, and we're only doing a linear blur, so we don't need the buttons for Radial or Zoom. We just need to choose between Linear and Centered.

Adding Another Control

For the first task, let's rename the Center. From the Add Control window, select Center from the ID list. A dialog will appear asking if you would like to Replace, Hide, or Change ID. We'll choose Replace. Now we are editing the Center input. We'll change the Name to Blur Vector, set the Type to Point, and the Page to Controls, which is the first tab where the controls are normally. Press OK, and our new input will appear on our node in the Node Editor. The ID of the control is still Center, so our SimpleExpressions did not change.

To hide the Length and Angle, we'll run the UserControls script again. This time when we select the Length and Angle IDs, we'll choose Hide in the dialog. Press OK for each.

Finally, to change the options available in the Type, we have two options. We can hide the buttons and use a checkbox instead, or we can change the MultiButton from four entries to two. Let's try both.

- To add the checkbox, run UserControls again, but this time instead of selecting an existing ID, we'll type Centered into the Name. This will set the name and the ID of our input to Centered. The Type is set to Number, and the Page is set to Controls. Now in the Type Attributes, set the Input Ctrl to be CheckboxControl. Press OK, and now we have our checkbox. To make the new control affect the Type, add a SimpleExpression to the Type:

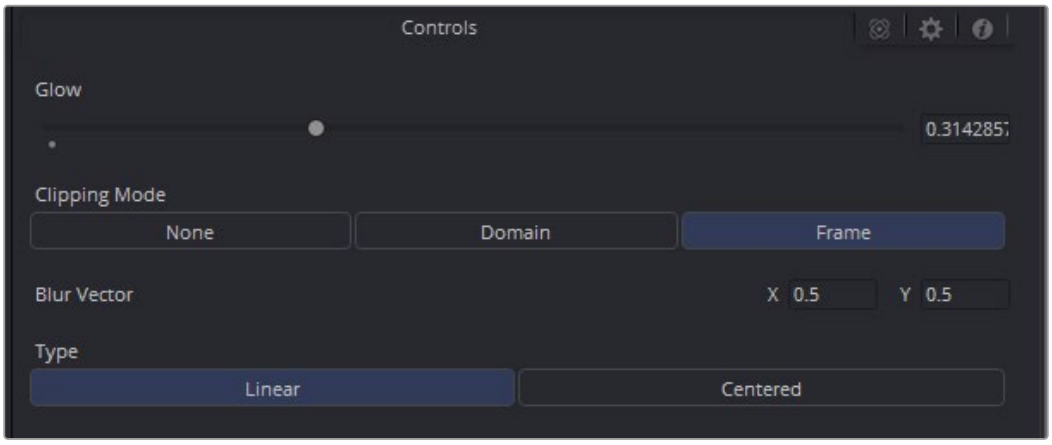
  ```
  iif(Centered==1, 2, 0).
  ```

  Once that's done, we can use the UserControls to hide the Type control.

- To make a new MultiButton, run the UserControl script, and add a new control ID, TypeNew. You can set the Name to be Type, as the Names do not need to be unique, just the IDs. Set the Type to Number, the Page to Controls, and the Input Ctrl to MultiButtonControl. In the Input Ctrl attributes, we can enter the names of our buttons. Let's do Linear and Centered. Type them in and hit Add for each. Press OK, and we have our new buttons with the unneeded options removed. To make this new control affect the original Type, add a SimpleExpression to the Type:

  ```
  iif(TypeNew==0, 0, 2).
  ```

  Once that's done, we can use the UserControls to hide the original Type control.



Directional Blurs with UserControls applied

Chapter 9

Animating in Fusion's Keyframes Editor

This chapter covers how you can keyframe effects in Fusion's Inspector and how you can edit clips, effects, and keyframes in the Keyframes Editor.

Contents

Keyframing in the Inspector 256
Removing Animation in the Inspector 257
Attaching a Parameter to an Existing Animation Curve 257
Keyframes Editor Overview 257
Keyframes Editor Tracks 258
The Timeline Header 258
The Playhead 259
Spreadsheet 259
Scaling and Panning the Timeline 260
Working with Segments in the Timeline 260
Moving Segments in the Timeline 261
Trimming Segments 261
Holding the First or Last Frame 261
Working with Keyframes in the Timeline 262
Drag and Drop Keyframe Editing 262
Keyframe Editing Using the Time Editor 262
The Keyframe Spreadsheet 263
Duplicating Spline Keyframes 263
Time Stretching Keyframes 263
Showing Keyframe Values 264
Timeline Filters 264
Selected Filtering 265

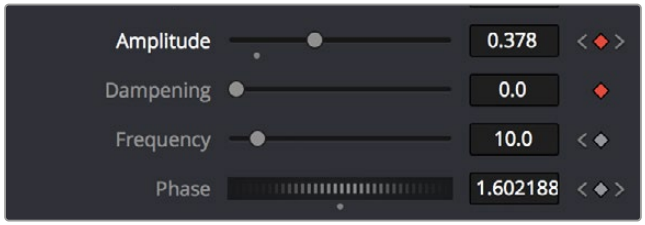
Sorting in the Timeline 266
The Tree Item Order Menu 266
The Sort Menu 266
Markers 267
Jumping to Markers 268
Renaming Markers 268
Show Marker List 268
Deleting Markers 268
Autosnap 269
Autosnap Points 269
Autosnap Markers 269
The Spreadsheet Editor 269
Selecting a Node to Edit 270
Inserting Keyframes 270
Selecting Multiple Nodes to Edit 270
Customizing the Keyframes Editor 271
Line Size 271
Display Point Values 271

Keyframing in the Inspector

Most parameters in most effects nodes can be keyframed in order to create animated effects such as animated transforms, rotoscoping with splines, dynamically altering warping behaviors, and more; the list is endless.

For convenience, a set of keyframing controls is available within the Inspector next to each keyframable parameter. These controls are:

- A gray Keyframe button to the right each keyframable parameter. Clicking this gray button creates a keyframe at the current position of the playhead, and turns the button orange.
- Whenever the playhead is sitting right on top of a keyframe, this button turns orange. Clicking an orange Keyframe button deletes the keyframe at that frame and turns the button gray again.
- Small navigation arrows appear to the right and left if there are more keyframes in those directions. Clicking on navigation arrows to the right and left of keyframes jumps the playhead to those keyframes.



Orange Keyframe buttons in the Inspector show there's a keyframe at that frame

Once you've keyframed one or more parameters, if Show Modes/Options has been enabled, the node containing the parameters you keyframed displays a Keyframe badge to show that node has been animated.



A keyframed node displays a Keyframe badge in the Node Editor

Once you've started keyframing node parameters, you can edit their timing in the Keyframes Editor and/or Spline Editor.

Removing Animation in the Inspector

To remove a keyframed spline from a parameter:

1. Right-click the keyframe control of the parameter you want to remove animation from.
2. Choose Remove [Name of parameter] from the contextual menu.

Attaching a Parameter to an Existing Animation Curve

Multiple parameters can be connected to the same animation curve. This can be an invaluable timesaver if you are identically animating different parameters in a node.

To connect a second parameter to the same animation curve:

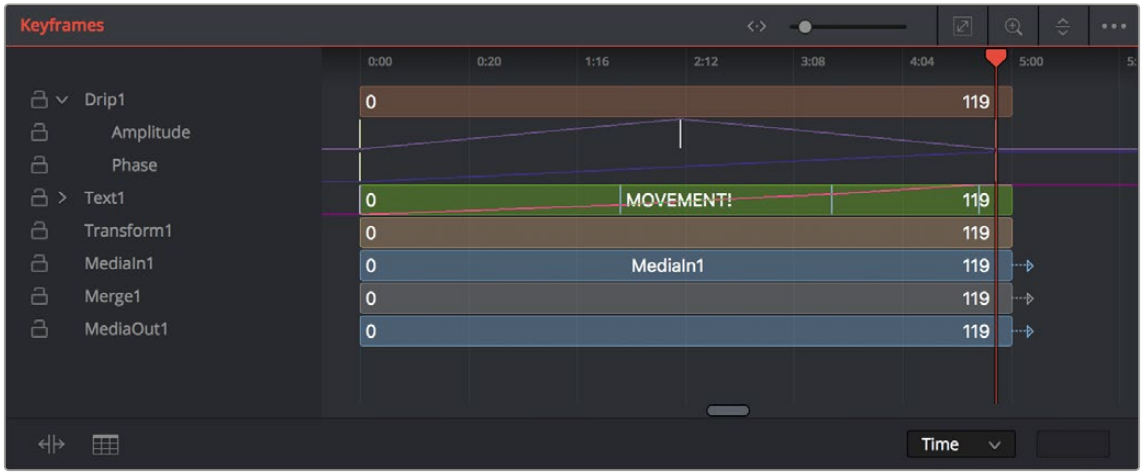
1. Right-click on the second parameter you want to attach.
2. In the contextual menu, hover over the Connect To submenu.
3. In the Connect To submenu, choose the name of the animated parameter.

Keyframes Editor Overview

The Keyframes Editor is essentially a timeline view of your composition, within which each clip and effect node in your composition is represented by a track. These tracks have the same color coding as the nodes they represent and are labeled where appropriate. A Time Ruler at the top indicates the timing of your composition, while numerous controls let you control the contents of the Keyframes Editor.

The Keyframes Editor can be used for one of two things:

- To adjust the timing of elements in a project, whether they're clips or effects. You can trim, slide, and extend clips, adjust the timing of an animation spline, or trim the duration of an effects node. You can freely rearrange the order of nodes in the Timeline without affecting the layering order of your composition. All compositing operations are handled in the Node Editor, while the Keyframes Editor manages the timing of your composition.
- To create and/or edit keyframes that you've applied to effects in a track-based manner, you can retime keyframes, add and delete keyframes, and even edit keyframe values



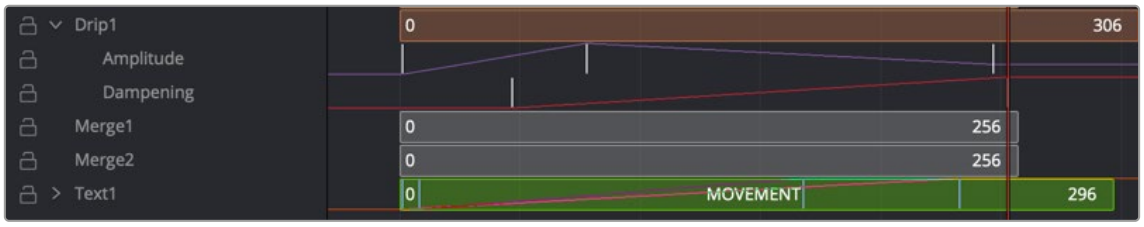
The Keyframes Editor

To show the Keyframes Editor, do one of the following:

- Click the Keyframes Editor button in the UI toolbar to toggle visibility of the Keyframes Editor on and off.
- Press F7 on the keyboard.

Keyframes Editor Tracks

While each clip and effect node in your composition is represented by a track, keyframed parameters are exposed either as keyframes superimposed upon the track to which they're applied (as seen on the MOVEMENT track), or they can be opened up onto their own tracks for more precise editing, one keyframe track per keyframed parameter, by clicking a disclosure control to the left of that track's name in the Timeline header (as seen under the "Drip1" track).



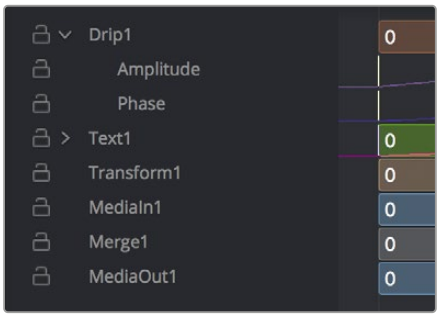
The Timeline tracks

The Timeline Header

The Timeline header area on the left side of the Timeline is a hierarchical list of all tracks in a composition. Each track displays the name of its corresponding node, a lock button, and a disclosure control for revealing keyframe tracks for each keyframe animation, modifier, and mask that's attached to it.

Collapse/Open All

A quick way to open or close all available keyframe tracks at once is to use the Expand/Collapse Tool Controls commands in the Keyframe Timeline Option menu.



The Timeline header area

The Playhead

As elsewhere in Fusion, the playhead is a red vertical bar that runs through the Timeline view to indicate the position of the current frame or time. The Keyframes Editor playhead is locked to the viewer playhead, so the image you're viewing is in sync.

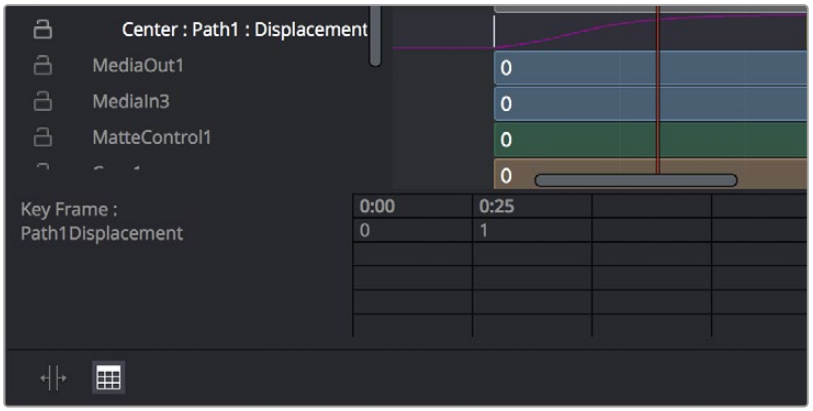
You must click on the playhead directly to drag it, even within the Timeline ruler (clicking and dragging anywhere else in the Timeline ruler scales the Timeline). Additionally, you can jump the playhead to a new location by holding down the Command-Option keys and clicking in the track area (not the Timeline ruler).



The playhead about to be dragged by the pointer

Spreadsheet

If you turn on the Spreadsheet and then click on the name of a layer in the keyframe track, the numeric time position and value (or values if it's a multi-dimensional parameter) of each keyframe appear as entries in the cells of the Spreadsheet. Each column represents one keyframe, while each row represents a single aspect of each keyframe.

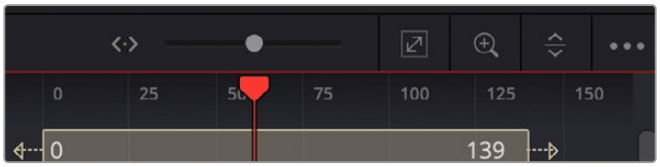


Editing keyframes in the Spreadsheet

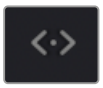
For example, if you're animating a blur, then the Key Frame row shows the frame each keyframe is positioned at, and the Blur1BlurSize row shows the blur size at each keyframe. If you change the Key Frame value of any keyframe, you'll move that keyframe to a new frame of the Timeline.

Scaling and Panning the Timeline

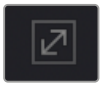
At the top, a series of zoom and framing controls let you adjust the work area containing the layers.



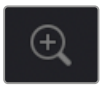
The Keyframe Editor framing controls



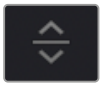
A Horizontal zoom control lets you scale the size of the editor.



A Zoom to Fit button fits the width of all tracks to the current width of the Keyframes Editor.



A Zoom to Rect tool lets you draw a rectangle to define an area of the Keyframes Editor to zoom into.



A Sort pop-up menu lets you sort or filter the tracks in various ways.



An Option menu provides access to many other ways of filtering tracks and controlling visible options.

Working with Segments in the Timeline

Most of the work in the Timeline involves trimming and aligning clip segments.

To select a single segment in the Timeline, do one of the following:

- Click the node's name in the header.
- Click the node's segment in the Timeline.

To add another segment to the selection, do one of the following:

- Hold Command and click additional segments to select discontiguous selections.
- Select a segment, and then hold Shift and click another segment to make a contiguous selection of all segments in between.

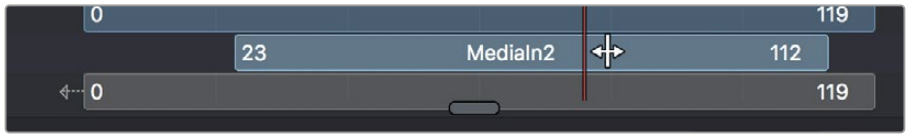
To remove a segment from the selection, do the following:

- Hold Command and click a selected segment to deselect it.

TIP: Selecting a node's name from the Timeline header also selects the node's tile in the Node Editor, with its controls displayed in the Inspector.

Moving Segments in the Timeline

To move the position of a segment, drag on the node's segment in the Keyframes Editor. The cursor will resemble a bar with two arrows pointing in either direction. Moving a segment changes where that clip begins and ends in the composition.



The Move cursor

Trimming Segments

Trimming segments has different effects on Loaders, MediaIn and Effect nodes:

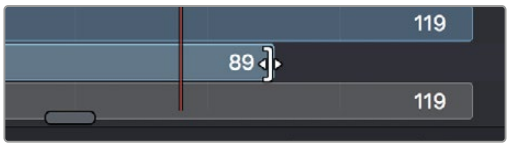
- Trimming a Loader or MediaIn node is similar to trimming clips in an editing application, in that you're changing the in and out points of the range of media that clip makes available to your composition.
- Trimming the segments of effect nodes instead modifies the range of that node's effect in the composition. Outside of the trimmed region, that effect node will behave as if it were disabled.

TIP: Shortening the duration of effects nodes can optimize processing. Imagine a Loader or MediaIn node that represents a clip that's 100 frames long and is connected to a Defocus node that's animated from frames 80–100. There is little to no point in processing the defocus node between frames 0–79, so trimming the defocus segment to start at frame 80 in the Timeline will effectively prevent it from rendering and consuming either memory or processor time until needed.

To trim a segment in the Timeline, do the following:

- Drag on either end of the node's segment in the Timeline.

  The cursor changes to a vertical bar with a single arrow when the cursor is in the right location to trim.



The Trim cursor

Holding the First or Last Frame

If you want to hold a Loader's first or last frame of a clip for a certain number of frames, also called a freeze frame, you can hold Command while you drag beyond the first or last of the segment in the Timeline.

Working with Keyframes in the Timeline

Keyframes can be drawn in one of two ways. When keyframe tracks are closed, they're drawn over the node's segment. Clicking on the disclosure icon to the left of the node's name in the track header expands the display so each keyframed parameter has its own track in the Timeline, enabling precise editing.

Furthermore, each keyframe track, whether open or closed, exposes a miniature curve overlay that provides a visual representation of the rise and fall of keyframed values. This little overlay isn't directly editable.



The Drip1 segment has its keyframe tracks exposed, while the Text1 segment has its keyframe tracks collapsed so they're displayed within the segment.

Drag and Drop Keyframe Editing

Here are pointer-based keyframe editing methods that will get you started.

Methods of selecting keyframes:

- Click a single keyframe to select it.
- Drag a bounding box over a series of keyframes to select them all.
- Command-click to select discontiguous keyframes.
- Shift-click the first and last of a range of keyframes to select a contiguous range.

Methods of adjusting keyframes:

- You can drag keyframes left and right to reposition them in time.
- You can right-click one or more selected keyframes and use contextual menu commands to change keyframe interpolation, copy/paste keyframes, or even create new keyframes.

Keyframe Editing Using the Time Editor

A drop-down and editing field at the bottom right of the Keyframes Editor lets you numerically edit the timing, in frames, of any selected keyframe, making it easy to make precise adjustments.

To change the position of a keyframe using the toolbar, do one of the following:

- Select a keyframe, and then enter a new frame number in the Time Edit box.
- Choose T Offset from the Time Editor drop-down, select one or more keyframes, and enter a frame offset.
- Choose T Scale from the Time Editor drop-down, select one or more keyframes, and enter a frame offset.



The Time button can switch to Time Offset or Time Scale for moving keyframes.

The Keyframe Spreadsheet

If you turn on the Spreadsheet and then click on the name of a layer in the keyframe track, the numeric time position and value (or values if it's a multi-dimensional parameter) of each keyframe appear as entries in the cells of the Spreadsheet. Each column represents one keyframe, while each row represents a single aspect of each keyframe.



Editing keyframes in the Spreadsheet

For example, if you're animating a blur, then the Key Frame row shows the frame each keyframe is positioned at, and the Blur1BlurSize row shows the blur size at each keyframe. If you change the Key Frame value of any keyframe, you'll move that keyframe to a new frame of the Timeline.

Duplicating Spline Keyframes

Keyframes can be duplicated, either onto the same keyframe track or onto different tracks. This can save you time if you need to repeat a keyframe sequence at another time on the same segment, or even just create identically-timed keyframes on two different segments.

To duplicate keyframes, do the following:

1. Select one or more keyframes you want to duplicate.
2. Hold Command and drag one of the selected keyframes to a new position.

Time Stretching Keyframes

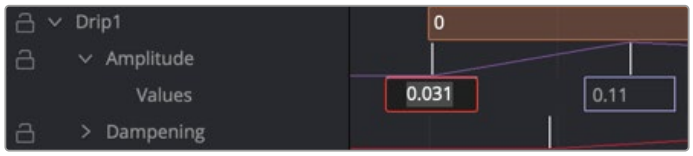
If you select a range of keyframes in a keyframe track, you can turn on the Time Stretch tool in the lower left of the Keyframes Editor, to show a box you can use to squeeze and stretch the entire range of keyframes relative to one another, to change the overall timing of a sequence of keyframes without losing the relative timing from one keyframe to the next. Alternatively, you can turn on Time Stretch and draw a bounding box around the keyframes you want to adjust to create a time-stretching boundary that way. Click the Time Stretch tool again to turn it off.



Time stretching keyframes

Showing Keyframe Values

When a node and its accompanying segment have animated parameters, keyframes appear as colored tick marks in keyframe tracks to indicate when animated changes occur. If the tracks and splines are open on a parameter, choosing Show Values from the Keyframes Editor Option menu shows editable fields beneath each keyframe. These fields show each keyframe's current value and allow you to edit them simply by entering a new number.



Show values from the Keyframes Editor Option menu

Timeline Filters

When a composition grows to include hundreds of nodes, locating specific node layers can quickly become difficult. Timeline filters can be created and applied to sift out nodes that are not necessary to the current operation. The Global Timeline preferences include a number of pre-made filters that you can enable, or you can create new ones as needed.

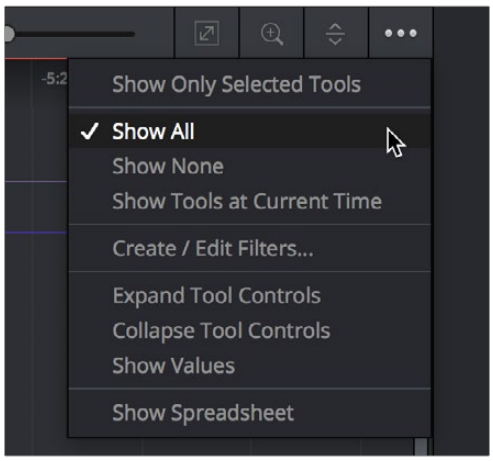
To use a Timeline filter:

Open the Keyframes Editor Option menu and choose an item from the top of the menu. Default Timeline filters include;

- Show All, which shows all node layers in the current composition.
- Show None, which hides all layers.
- Show Tools at Current Time, which only displays node layers under the playhead.
- If you've created custom filters, they appear here as well, in alphabetical order.

To go back to showing everything:

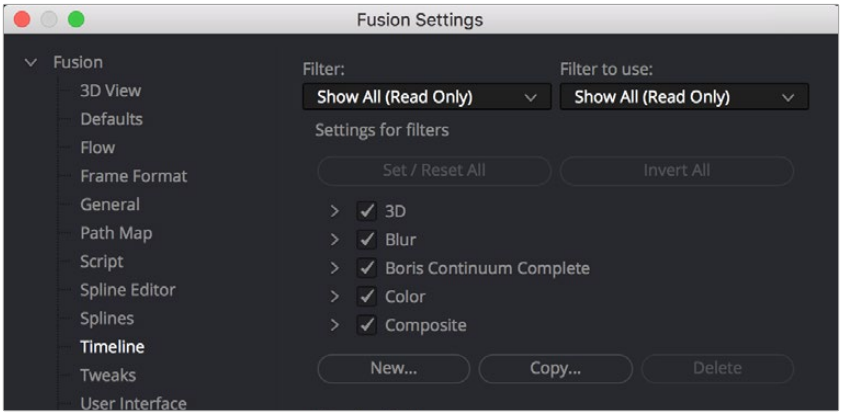
- Choose Show All from the Keyframes Editor Option menu. All layers will reappear.



Choosing a Timeline filter

To create a Timeline filter:

1 Choose Create/Edit Filters from the Keyframes Editor Option menu to open the Timeline panel of the Fusion Settings window. This is where you can create new Timeline filters.



The Global Timeline preferences for enabling filters

2 Click the New button, enter a name for your new filter setting, and click OK. The filter you created is now selected in the Filter pop-up menu at the top.

3 Use the “Settings for filters” list to turn on the checkboxes of nodes you want to be seen and turn off the checkboxes of nodes you want to filter out. Each category of node can be turned on and off, or you can open up a category’s disclosure control to turn individual nodes on and off. Clicking Invert All immediately turns off all node categories.

4 When you’re finished creating filters, click the Save button to hide the Fusion Settings window.

Filters that you’ve created in the Timeline panel of the Fusion Settings window appear in the Keyframes Editor Option menu.

To delete a filter:

1 Choose Create/Edit Filters from the Keyframes Editor Option menu to open the Timeline panel of the Fusion Settings window. This is where you can delete Timeline filters.

2 Choose the filter you want to delete from the Filter pop-up menu.

3 Click the Delete button, and when a dialog asks if you really want to do that, click OK.

Selected Filtering

Choosing “Show only selected tools” from the Keyframes Editor Option menu filters out all segments except for layers corresponding to selected nodes. This option can be turned on or off.

TIP: When “Show only selected tools” is enabled, you can continue to select nodes in the Node Editor to update what’s displayed in the Keyframes Editor.

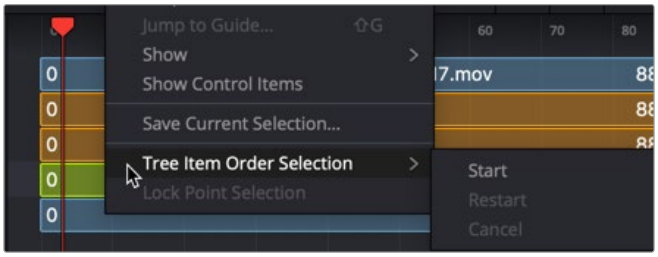
Sorting in the Timeline

You can change the order in which the nodes are displayed from top to bottom in the Timeline.

- You can use the Tree Item Order Selection menu to sort the tracks by an assigned number.
- You can use the Sort pop-up menu.

The Tree Item Order Menu

Right-clicking over any track on the Keyframes Editor will display a contextual menu that contains the Tree Item Order Selection submenu. Choosing Start from the submenu allows you to start numbering each item in the track header by clicking on them. The first item you click will be #1, the second item #2, the third #3, and so on. Once you have selected all the items in the order you want them to be organized in the Keyframes Editor, right-click over a track, and from the Tree Item Order Selection submenu, choose End. The items will be ordered using the assigned numbers, with #1 appearing above #2, which appears above #3, and so on. The first items in the Keyframe track list will always be the nodes that are the root of the node tree. The numbered nodes will appear in order after the root nodes. For example, if the node tree starts with a background node, and then connects to a Fast Noise, Blur, and Color Corrector, the background node will always appear at the top of the Keyframes Editor track list because it is the root node.

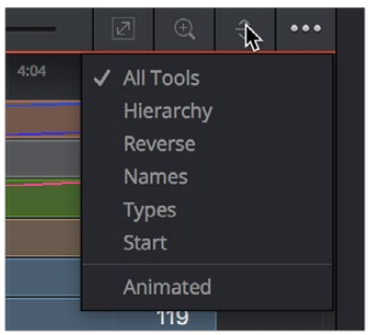


The Keyframes Editor Tree Item Order Selection menu

If you begin numbering nodes in the track header and change your mind or decide on a different order, you can choose Restart to begin numbering again or choose Cancel to keep the current order.

The Sort Menu

The Sort menu reorders how the layers of each node appear in the Keyframes Editor. Setting the menu back to All Tools will display them in a linear order, scanning the Node Editor from left to right and top to bottom. This is the default setting.



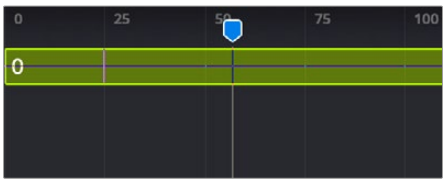
The Timeline Sort Order menu

- **All Tools:** Forces all tools currently in the Node Editor to be displayed in the Keyframes Editor.

- **Hierarchy:** Sorts with the most background layers at the top of the header, through to the most foreground layers at the bottom, following the connections of the nodes in the Node Editor.
- **Reverse:** The opposite of Hierarchy, working backward from the last node in the Node Editor toward the most background source node.
- **Names:** Sorts by the alphabetical order of the nodes, starting at the top with the beginning of the alphabet.
- **Start:** Orders layers based on their starting point in the composition. Nodes that start earlier in the Global project time are listed at the top of the header, while nodes that start later are at the bottom.
- **Animated:** Restricts the Timeline to showing animated layers only. This is an excellent mode to use when adjusting the timing of animations on several nodes at once.

Markers

Markers help identify important frames in a project that might affect how you keyframe animation. They may indicate the frame where a dragon breathes fire at a protagonist, the moment that someone passes through a portal, or any other important frame in a composition that you need to keep track of. Markers added to the Timeline in the Cut, Edit, Fairlight, or Color page will appear in the Keyframes Editor and Spline Editor of the Fusion page. They can also be added from the Keyframes Editor or the Spline Editor while working in Fusion Studio or the Fusion page. Markers in Fusion appear as a small handle with a line extending vertically through the graph view when selected.



A marker being moved in the Keyframed Editor

To create a marker, do the following:

Right-click at a frame in the Timeline Ruler of the Keyframes Editor and choose Add Marker from the contextual menu.

The most important attribute of a marker is its position. For it to add value, a marker must be placed on the frame you intended it to be on. Hovering the cursor over a marker displays a tooltip with its current frame position. If it is on the wrong frame, you can drag it along the Time Ruler to reposition it.

Markers added to the Time Ruler are editable in the Fusion page, and the changes appear back in the other DaVinci Resolve pages. Time Ruler markers can be added, moved, deleted, renamed, and given descriptive notes from within Fusion's Keyframes or Spline Editor.

NOTE: Markers attached to clips in the Edit page Timeline are visible on MediaIn nodes in Fusion's Keyframes Editor but not editable. They are not visible in the Spline Editor.

Jumping to Markers

Double-clicking a marker jumps the playhead to that marker's position.

Renaming Markers

By default, a marker uses the frame number in its name, but you can give it a more descriptive name to go along with the frame number, making it easier to identify. To rename a marker in Fusion, right-click over the marker and choose Rename Guide from the contextual menu. Enter a name in the dialog and click OK.



The Marker contextual menu is accessed by right-clicking over a marker or Keyframe Editor Time Ruler.

Show Marker List

Markers can be used to jump to specific locations in a composition using the Marker List. If you right-click over a marker or within the Keyframe Editor Time Ruler to bring up the contextual menu, you can choose Show Marker List, or press Shift-G, to display the Marker List dialog. The Marker List is a floating dialog that will remain on top of the main window until closed.

The Marker List shows all the current markers in the composition, listed according to their position in time along with any custom name you've given them. If you double-click a marker's name from the list, the playhead jumps to the marker's location.



The Marker List dialog allows you to navigate through a composition using markers.

There is a pair of checkboxes beside the names of each marker. One is for the Spline Editor, and one is for the Keyframes Editor. By default, markers are shown in both the Spline Editor and Keyframes Editor, but you can deselect the appropriate checkbox to hide the markers in that view.

Deleting Markers

You can delete a marker by dragging it up beyond the Time Ruler and releasing the mouse. You can also use the marker's contextual menu to choose Delete Marker.

Autosnap

To help with precisely positioning keyframes and the start and end of segments as you drag in the Timeline, you can have them snap to a field, a frame, or to markers. The Autosnap option is accessed through the Options section in the Keyframes Editor's contextual menu. There are two submenu options for autosnapping. One option controls the snapping behavior when you drag keyframes, control points, or the starting and ending edges of segments. The other option controls the snapping behavior of markers.

Autosnap Points

When you drag keyframes or the edges of segments, often you want them to fall on a specific frame. Autosnap restricts the placement of keyframes and segment edges to frame boundaries by default, but you have other options found in the contextual menu. To configure autosnapping on keyframes and segment edges, right-click anywhere within the Keyframes Editor and choose Options > Autosnap Points from the contextual menu. This will display the Autosnap Points submenu with options for the snapping behavior. The options are:

None: None allows free positioning of keyframes and segment edges with subframe accuracy.

Frame: Frame forces keyframes and segment edges to snap to the nearest frame.

Field: Field forces keyframes and segment edges to snap to the nearest field, which is 0.5 of a frame.

Guides: When enabled, the keyframes and segment edges snap to markers.

Autosnap Markers

When you click to create a new marker, the default behavior is that it will snap to the closest frame. If you reposition the marker, it also snaps to the nearest frame as you drag. This behavior can be changed in the Keyframes Editor's contextual menu by choosing from the Options > Autosnap Markers submenu. The options are:

None: Markers can be placed anywhere with subframe accuracy.

Frame: Frame forces all markers to snap to the nearest frame.

Field: Field forces all markers to snap to the nearest field.

The Spreadsheet Editor

The Spreadsheet Editor is a separate panel that can be displayed beneath the Keyframes Editor. It is used to compactly show the numeric values of the keyframes for selected parameters in the Keyframes Editor's header, via a table with rows and columns, showing time and value.

To reveal the Spreadsheet Editor, click on the Spreadsheet button in the toolbar. The Spreadsheet will split the Work Area panel and appear below the Keyframes Editor's interface.

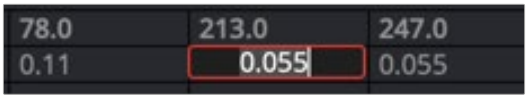


The Spreadsheet Editor showing editable data for six keyframes

Selecting a Node to Edit

To display a node's timing in the Spreadsheet, select the node's name in the Keyframes Editor header. The Start and End points of the selected node will appear in the keyframe's line of the Spreadsheet.

To edit an animation parameter in the Spreadsheet Editor, select the parameter in the Keyframes Editor header. The keyframe row includes a box for each frame number that contains a keyframe. The value of the keyframe is displayed in the cell below the frame number. Clicking on a cell allows you to change the frame number the keyframe is on or the parameter's value for that keyframe.



Clicking on the parameter's keyframe value allows you to change it

TIP: Entering a frame number using a decimal point (e.g., 10.25 or 15.75) allows you to set keyframes on a subframe level to create more natural animations.

Inserting Keyframes

You can also add new keyframes to an animation by clicking in an empty keyframe cell and entering the desired time for the new keyframe. Using the cell under the new keyframe, you can enter a value for the parameter.

Selecting Multiple Nodes to Edit

Multiple splines and nodes can be edited together in the Spreadsheet. By default, selecting a new parameter in the Timeline header will replace the parameter and keyframes currently listed in the Spreadsheet Editor. Holding down Command, you can click on additional parameters on different nodes to add them to the Spreadsheet.

Customizing the Keyframes Editor

There are a few ways you can change the appearance of the Keyframes Editor to better fit your needs. All these options are found by right-clicking anywhere within the Keyframes Editor and choosing an option from the contextual menu that appears.

Line Size

The Line Size option controls the height of each Timeline segment individually. It is often useful to increase the height of a Timeline bar, especially when editing or manipulating complex splines.

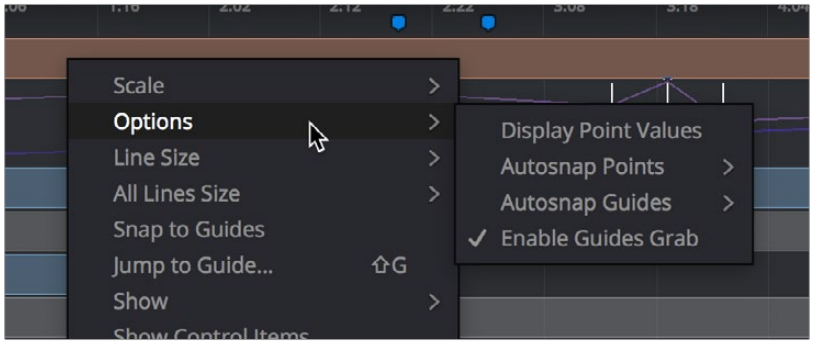
Methods of increasing or decreasing the height of segments:

To change the hight of just one segment: Right-click anywhere within the Keyframes Editor and choose a size from the Line Size submenu. The options are Minimum, Small, Medium, Large, and Huge.

To change the height of all segments: Right-click anywhere within the Keyframes Editor and choose a size from the All Line Size submenu. The options are Minimum, Small, Medium, Large, and Huge.

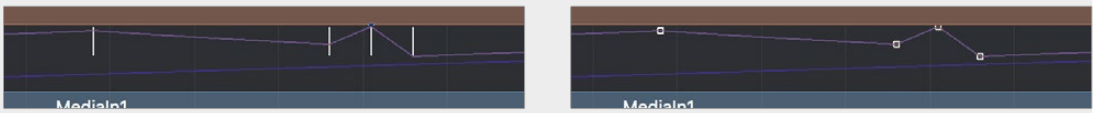
Display Point Values

A more traditional view of keyframes is to view them as control points instead of vertical bars, making them easier to select for some people. From the Timeline contextual menu, you can right-click anywhere within the Keyframes Editor and choose Options > Display Point Values to change how keyframes look.



The Options submenu for changing Display Point Values

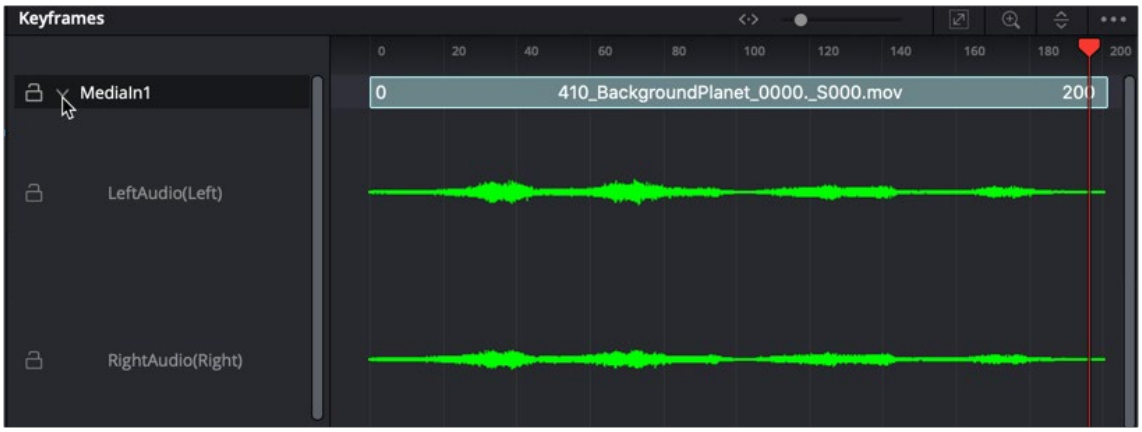
Here are the two options, compared.



Keyframes displayed as bars (left), and keyframes displayed as Point Values (right).

Displaying Audio Waveforms

You can display a MediaIn node's audio waveform in the Keyframes Editor and use it as a guide as you add and move keyframes.



Waveforms are displayed in the Keyframes Editor for all MediaIn nodes

To display the audio waveform in the Keyframes Editor:

1. Open the Keyframes Editor.
2. Click the disclosure arrow next to the MediaIn node to view the audio waveform for that clip.

To change the size of an audio waveform display:

1. Open the Keyframes Editor.
2. In the Keyframes Editor, select the audio track of the waveform you want to modify.
3. Right-click over the audio waveform and choose Line Size > Minimum/Small/Medium/Large/Huge.
4. When using Fusion Studio, you can view the audio waveforms in the Keyframes Editor by displaying the Saver node.

To view the Audio Waveform in Fusion Studio, do the following:

1. Open the Keyframes Editor.
2. Expand the Saver track to view the audio waveform.

   When you want to find the precise location of an audio beat, transient, or cue, you can slowly drag over the audio waveform to hear the audio. If you need to see more resolution in the waveform display, you can increase the size.

To change the size of an audio waveform display:

1. Open the Keyframes Editor.
2. Right-click over the audio waveform and choose Line Size > Minimum/Small/Medium/Large/Huge.

TIP: Right-clicking a track in the Keyframes Editor and choosing All Line Size > Minimum/Small/Medium/Large/Huge changes all the tracks and audio waveforms in the Keyframes Editor.

Chapter 10

Animating in Fusion's Spline Editor

This chapter covers how you can keyframe effects and control animations in Fusion's Spline Editor.

Contents

Spline Editor Overview 274
Spline Editor Interface 275
The Graph, Header, and Toolbar 275
Renaming Splines 277
Changing Spline Colors 277
Navigating Around the Spline Editor 277
Markers 278
Autosnap 281
Creating Animation Splines 281
Animating with Different Spline Types 282
Working with Keyframes and Splines 284
Adding Keyframes 284
Locked and Unlocked Controls Points 285
Selecting, Moving, and Deleting Keyframes 286
Showing Key Markers 287
Copying and Pasting Keyframes 287
Time and Value Editors 289
Modifying Spline Handles 290
Reducing Points 291
Filtering the Spline Editor 292
Working with Filters 293
Changing a Spline's Status 294
Selection States 295

Reshaping Splines Using the Toolbar 296
Interpolation 296
Reversing Splines 298
Looping Splines 298
Time Stretching 301
Shape Box 302
Ease In/Out 303
Importing and Exporting Splines 304

Spline Editor Overview

The Spline Editor is the main area where animation is manipulated and refined. You primarily use the Spline Editor to show the changing values of parameters over time in the form of splines. Whereas Keyframes explicitly set the value of a parameter on a given frame, Splines are lines or curves that interpolate the values between keyframes. Once you set a keyframe on a parameter, a spline is created and can be displayed in the Spline Editor so you can make further refinements to the animation. However, the Spline Editor is more advanced than a standard curve editor since it can also display functions, which may not be splines, like changing characters within a text string or mathematical expressions that drive your animations.



The Spline Editor with three animated parameters

What Are Splines?

All animation uses splines that describe the value of a parameter at any given point in time. The Spline Editor graph shows the time of the comp along the horizontal axis and the value of the parameter along the vertical axis.

The advantage of using splines to represent animation instead of keyframes as in the Keyframes Editor is that splines allow you to manipulate the interpolation between keyframes. For example, if a keyframe is set at a value of 1.0 on a parameter for frame 1, followed by a keyframe value of 10.0 for frame 10, the values between keyframes are smoothly interpolated or calculated based on the shape of the spline. Using the functions and controls in the Spline Editor, you have a fantastic amount of control over that interpolation.

Spline Editor Interface

The Spline Editor is not visible by default, but you can show it at any time by clicking the Spline button in the user interface toolbar. You can also display the Spline Editor by right-clicking on a node in the Node Editor or a segment in the Keyframes Editor and choosing Edit Splines from the drop-down menu.

The Spline Editor can be open alongside the Node Editor or Keyframes Editor, or displayed separately in order to take up the entire work area.

The Graph, Header, and Toolbar

The Spline Editor has three main working areas: the graph, header, and toolbar. On the left side of the Spline Editor is the header, which shows a list of animated parameters. The majority of the panel is taken up by the splines displayed in the graph area, and a toolbar runs along the bottom, providing a variety of ways to manipulate the splines.

Graph

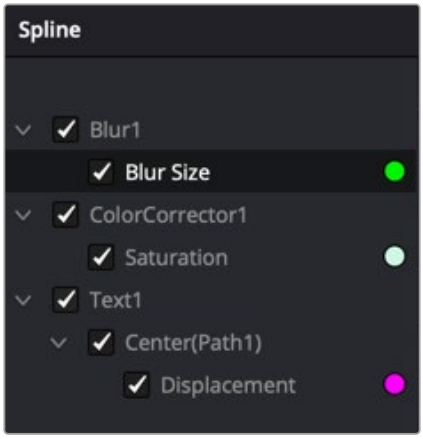
The graph is the largest area of the interface. It is here that you see and edit the animation splines. There are two axes in the graph. The horizontal axis represents time, and the vertical axis represents the spline's value. A thin bar, called the playhead, runs vertically through the graph to represent the current time as it does in the Timeline Editor. You can drag the playhead to change the current time, updating the frame displayed in the viewers.



Removing a spline from a parameter removes the animation

Spline Editor Header

The header provides a mechanism for determining what splines are visible in the graph. It shows the name of each spline in the project beneath the tool that contains that parameter. The checkbox beside each name shows whether that spline is currently displayed in the graph and whether the spline can be edited.



The Spline header

Spline Editor Toolbar

The toolbar across the bottom of the Spline Editor represents the most common operations applied to an animation spline. The various operations represented in the toolbar are all accessible from the graph's context menu as well, but the following buttons provide a faster shortcut.



Spline toolbar buttons

Playhead

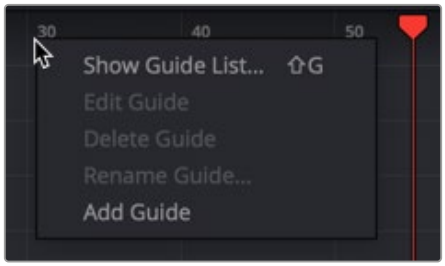
The playhead is the thin red vertical bar that runs vertically through the Spline Editor graph and represents the current time of the comp. You can drag the playhead to change the current time.

Status Bar

The status bar in the lower-right corner of the Fusion window regularly displays information about the position of the pointer, along with the time and value axes.

Contextual Menus

There are two contextual menus accessible from the Spline Editor. The Spline contextual menu is displayed by right-clicking over the graph, while the Guide contextual menu is displayed by right-clicking on the Time Ruler above the graph.



Right-click on the horizontal axis Time Ruler for the Guide menu

Renaming Splines

The name of a spline in the header is based on the parameter it animates. You can change the name of a spline by right-clicking on it in the header and choosing Rename Spline from the contextual menu.

Changing Spline Colors

Each spline in the graph is assigned a different color, making individual splines easier to identify when multiple splines are visible at once. When the spline is active, a round color swatch is displayed next to the spline's name in the header.

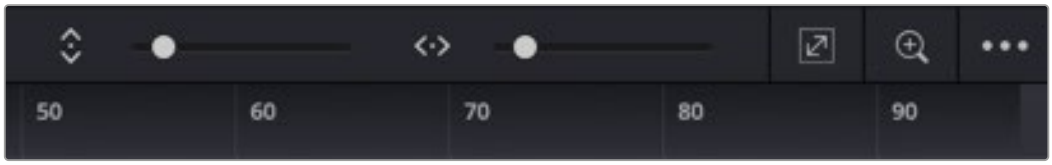
To change the color of a spline:

1. Click on the color circle next to the spline name in the header.
2. Or right-click on the name of the spline in the header and choose Change Color from the contextual menu.
3. Select the new color from the dialog box that appears and click OK.

Navigating Around the Spline Editor

It is often necessary to magnify and pan around the graph area to ensure that the splines you want to work on are visible. In general, scaling and panning the Spline Editor works the same as in all navigable parts of the Fusion interface. However, there are several unique functions to the Spline Editor for controlling your view based on the height, width, and selection of multiple animation splines.

The most obvious navigation methods to use are the scale sliders and buttons located in the upper left of the Spline Editor panel.



The Zoom Height and Zoom Width sliders, Fit button, and Zoom to Rectangle button can be used to navigate around the graph

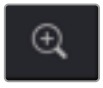
To scale and pan using the sliders and buttons:



Zoom Height and Zoom Width sliders let you change the height and width of the graph area.



The Fit button attempts to rescale the view so that all currently active splines fit within the graph.



The Zoom to Rectangle button (Command-R) allows you to draw a bounding box around the area of the graph you want centered and scaled.

To scale using the axis labels:

- Place the mouse pointer over the rulers for the horizontal or vertical axis and drag to resize the graph on that axis only. The view is scaled centered on the original position of the pointer on the ruler.



Dragging in the Time Ruler scales the graph horizontally

To scale and pan with the mouse and/or keyboard:

- With the Spline Editor active, press the + and - keys on your keyboard to zoom in and out of the graph.
- You can also zoom to a specific control point by holding down the Command key and scrolling the middle mouse wheel. The mouse pointer location determines the area that gets magnified.
- Position the mouse pointer over the graph, and then hold down the middle mouse button. With the middle mouse button pressed down, click once on the left mouse button to zoom in and once on the right button to zoom out.

To pan the graph area:

- Drag left or right using the middle mouse button or use the scroll bar along the bottom and right side of the graph.

Using the Graph contextual menu for navigation:

There are several ways to navigate the graph area using the Spline Editor contextual menu as well.

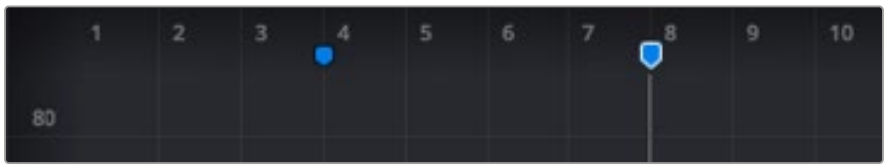
- Choose Scale > Scale to Fit (Command-F) to fit all active splines into the graph area.
- Choose Scale > Scale to Rectangle (Command-R) to draw a bounding box around the area of the graph you want centered and scaled. This has the same effect as clicking the Zoom to Rectangle button.
- Choose Scale > Default to reset the scaling of the graph area to default values.
- Choose Scale > Zoom In/Zoom Out to scale the graph area. This performs the same functions as pressing the + and - keys on the keyboard.
- Choose Scale > Auto Fit to scale the graph to fit all splines dynamically as you make splines visible and hidden. If the scaling is changed with Auto-Fit enabled, the graph area will scroll as you play the comp to view all the keyframes.
- Choose Scale > Auto Scroll to scroll the graph area if the splines fall outside the graph horizontally as you play.
- Choose Scale > Manual to disable all automatic attempts at showing splines in the graph.
- Choose Options > Fit Times to automatically scale along the X-axis to fit the selected spline. All visible splines are taken into account, not just the newly selected spline. With this option off, activating a new spline will not change the horizontal scale.
- Choose Options > Fit Values to automatically scale along the Y-axis to fit the selected spline. All visible splines are taken into account, not just the newly selected spline. With this option off, activating a new spline will not change the vertical scale.

Markers

Markers help identify important frames in a project. They may indicate a frame where a ray gun shoots a beam in the scene, the moment that someone passes through a portal in the image, or any other important event in the composite.

Markers added to the Timeline in the Cut, Edit, Fairlight, or Color page will appear in the Keyframes Editor and Spline Editor of the Fusion page. They can also be added from the Keyframes Editor or the Spline Editor while working in Fusion Studio or the Fusion page. Markers appear along the top of the horizontal axis Spline Editor's Time Ruler. They are displayed as small blue shapes, and when selected, a line extends from each guide down vertically through the graph.

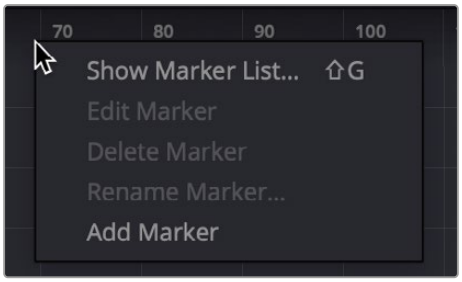
NOTE: Markers attached to clips in the Cut, Edit, Color, or Fairlight pages Timeline are not visible in Fusion's Spline Editor.



Unselected markers appear as blue shapes along the top, while selected markers display a vertical line running through the graph

Working with Markers

Markers call attention to a particular frame within a comp. They can be named, displayed in a list, and edited. After you add markers, you can easily jump the playhead between them, change their position, or delete them altogether.



Markers can be added by right-clicking in the horizontal time axis

To create a marker:

- Right-click in the horizontal axis Time Ruler and choose Add Marker.

To delete a marker, do one of the following:

- Drag the marker up outside the Spline Editor panel.
- Right-click on the marker and then choose Delete Marker from the menu.
- Select the marker and then press Delete or Backspace on the keyboard.
- From the Marker List, select a guide in the list and click the Del button.

To move a marker to a new frame, do one of the following:

- Drag the marker handle along the time axis.
- Right-click in the marker area and choose Options > Enable Marker Grab, and then drag the marker's vertical line to move the guide.

To move the Playhead to a marker:

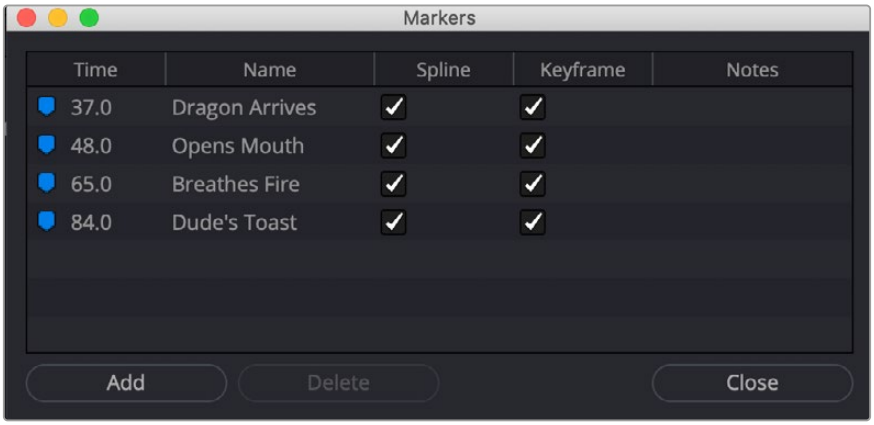
- Right-click on a marker and choose Set Current Time To [Frame number].

Using the Marker List

The Marker List is a list of all markers in the current comp. It can display the markers from either the Keyframes Editor, the Spline Editor, or both panels simultaneously. Clicking on the frame number or name of a guide causes the current time to change to that marker's frame. Since the Marker List is a floating window, it can remain open, allowing you to quickly jump to different markers while you work in the Spline Editor.

To show the Marker List:

- Right-click in the horizontal axis and choose Show Marker List, or press Shift-G.



The Marker List above shows markers in the current comp.

If markers currently exist in the comp, they are automatically displayed in the Marker List, regardless of whether they were added in the Keyframes Editor or the Spline Editor or any other page in DaVinci Resolve. You can also add markers directly from the Marker List, which can be helpful if you have multiple markers you need to add, and you know the rough timing.

To add a guide from the Marker List:

1. Click the Add button in the Marker List window.
2. Enter a frame number in the Time field.
3. Press Tab or click the Close button to close the Marker List.

To name a marker, do one of the following:

- In the Marker List, double-click in the Name column, to the right of the frame number, and enter a name for the marker.
- Right-click over a marker in the horizontal axis and choose Rename Marker. In the dialog that opens, enter a name for the marker.

Displaying Marker with the Timeline

The Marker List window includes checkboxes next to each marker that determines whether a marker displays in the Spline Editor, the Keyframes Editor, both, or neither. By default, when you create markers, they are active in both panels. To hide a marker from appearing in either panel, deselect the appropriate checkbox.

Autosnap

To assist in precisely positioning keyframe control points along the horizontal (time) axis, you can enable the Spline Editor's Autosnap function. Right-clicking over a spline and choosing Options > Autosnap Points provides a submenu with four options.

None: Allows free, sub-frame positioning of the keyframes.

Frame: Keyframes snap to the nearest frame.

Fields: Keyframes snap to the nearest field.

Markers: Keyframes snap to the nearest marker.

Autosnapping and Markers

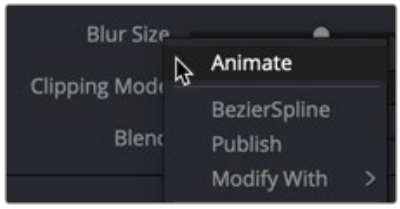
By default, a newly-created marker snaps to the closest frame. Moving markers with the mouse also snaps them to the current frame. You can change this behavior by selecting Options > Autosnap Markers > None or by selecting Options > Autosnap Markers > Field from the contextual menu.

Creating Animation Splines

Animation splines are created automatically when you keyframe a parameter in the Inspector or the Keyframes Editor. However, you can create an animation spline without first having to add a keyframe.

To create a spline:

- Right-click on the parameter to be animated in the Inspector, and choose Animate from the contextual menu.

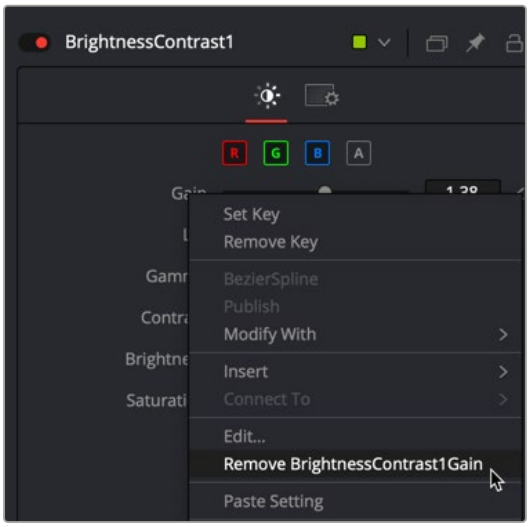


Choosing Animation from the menu displays a spline in the Spline Editor

Selecting Animate from the contextual menu connects the parameter to the default spline type. This is usually a Bézier Spline unless you change the default spline in the Defaults panel of the Fusion Preferences.

Deleting Animation Splines

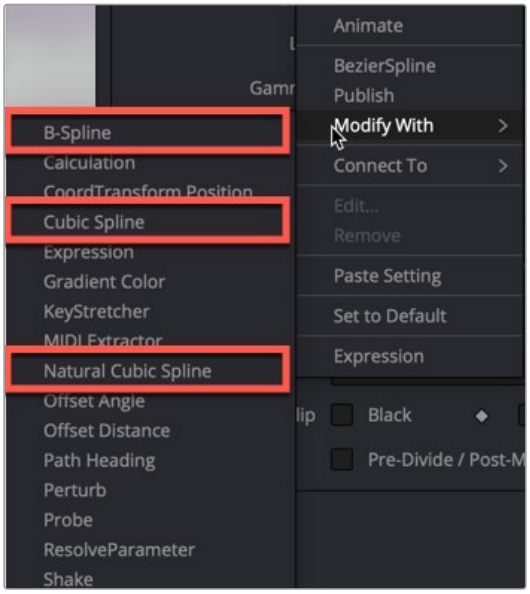
To remove an animation spline from a parameter, right-click on the control in the Inspector and select Remove [tool parameter's name] from the contextual menu. Removing a spline from a parameter only deletes the spline if no other tool in the composition is connected to the same spline at that time.



Removing a spline from a parameter removes the animation

Animating with Different Spline Types

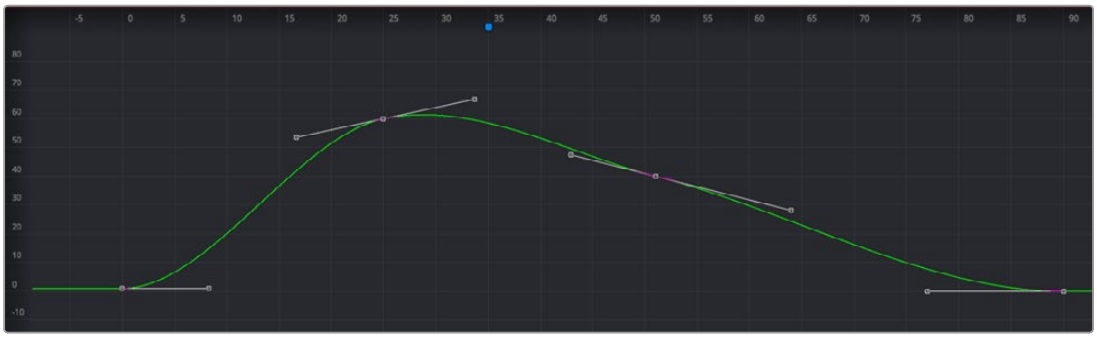
A Bézier spline is the default spline unless changed in the Preferences. However, if you want to use a spline type other than Bézier for the animation curve, you can choose the spline type from the Modify With contextual submenu before creating any keyframes.



Three spline types in the Modify With menu

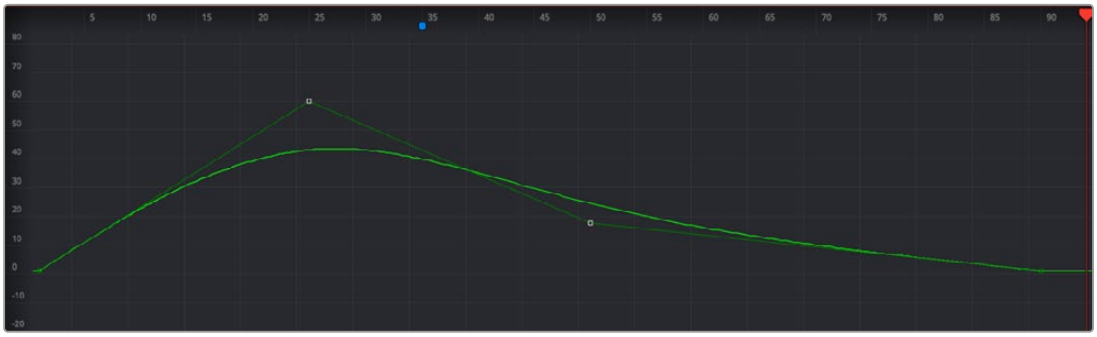
- **Bézier Spline:** Bézier splines are the default curve type. Three points for each keyframe on the spline determine the smoothness of the curve. The first point is the actual keyframe, representing the value at a given time. The other two points represent handles that determine how smoothly

the curve for the segments leading in and out of the keyframe are drawn. Bézier is the most used spline type because Bézier splines allow you to create combinations of curves and straight lines.



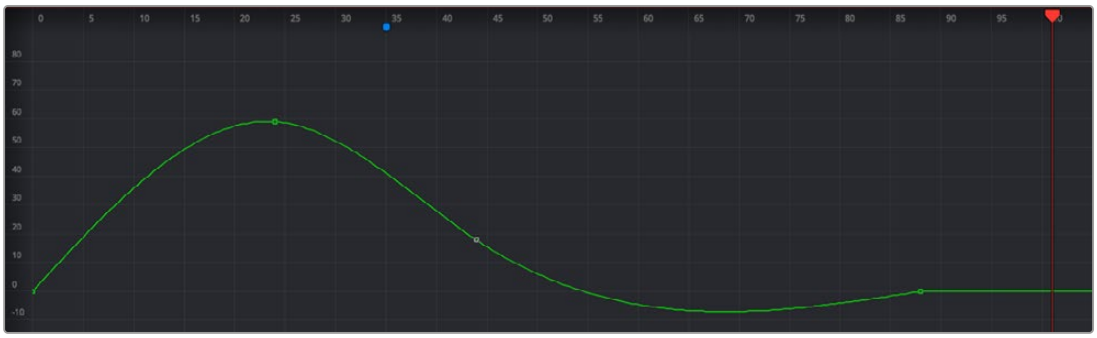
Bézier spline

- **Modify with > B-Spline:** B-splines use a single point to determine the smoothness of the curve. Instead of using handles, a single control point determines the value as well as the smoothness of the curve. Holding down the W key while dragging left or right on the control point adjusts the tension of the curve.



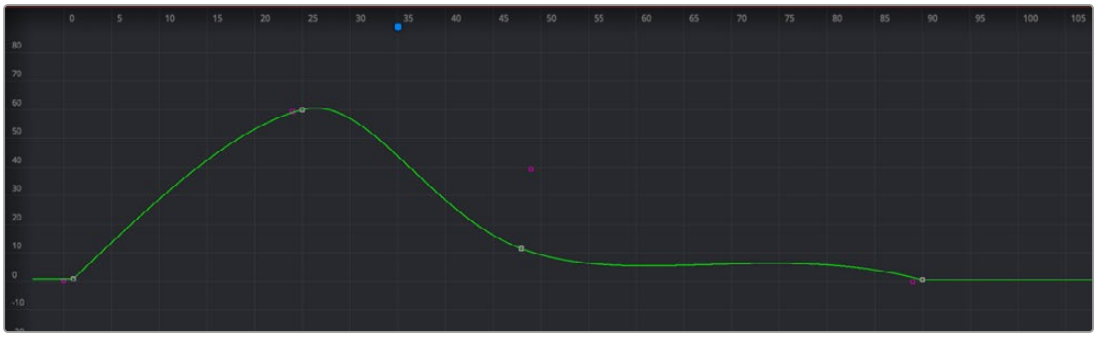
B-spline

- **Modify with > Cubic Spline:** Cubic splines are similar to Bézier splines, in that the spline passes through the control point. However, Cubic splines do not display handles and always make the smoothest possible curve. In this way, they are similar to B-splines. This spline type is almost never used.



Cubic spline

- **Modify with > Natural Cubic Spline:** Natural Cubic splines are similar to Cubic splines, except that they change in a more localized area. Changing one control point does not affect other tangents beyond the next or previous control points.



Natural Cubic spline

Working with Keyframes and Splines

Once you animate a parameter and display the Spline Editor, you can manipulate the spline's keyframes (and thus the animation) in a variety of ways. By selecting Keyframe control points, you can move, copy, and change the interpolation of your animation.

Adding Keyframes

Once you create one keyframe, additional keyframes are automatically added to a spline whenever you move the playhead and change the value of that spline's parameter. For example, if you change the strength of an animated glow at frame 15, a keyframe with the new value occurs on frame 15.

In the Spline Editor, control points can also be added directly to a spline by clicking on the spline where you want to add the new keyframe.

Adding Keyframes at the Playhead

If you want to add a new keyframe at the current playhead location, pressing Command-K on the keyboard or right-clicking in the graph and choosing Set Key adds a keyframe under the playhead.

Adding Equal Keyframes

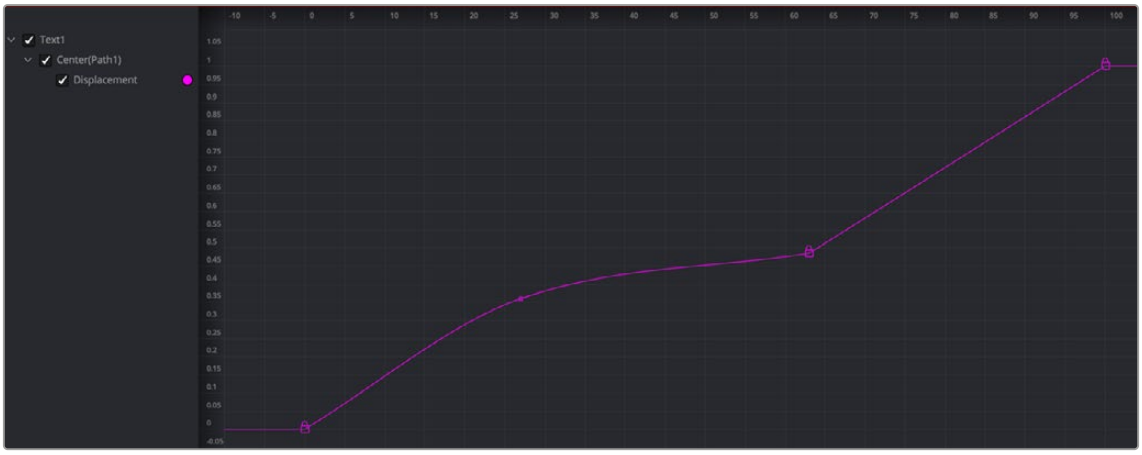
If you want to hold a value over several frames, right-clicking in the graph area and choosing Set Key Equal To displays a submenu to add a new keyframe with a value equal to the next or the previous keyframe.



To hold a value between two keyframes choose from the Set Key Equal To submenu

Locked and Unlocked Controls Points

When animating the Center X/Y or Pivot X/Y parameters on any tool, you create a displacement spline in the Spline Editor. The displacement spline represents the relative offset position of the animated object along its path. Since the displacement spline is relative, keyframes use a value between 0.0 and 1.0. A displacement value of 0.0 in the Spline Editor indicates that the object is at the very beginning of a path. A value of 1.0 indicates that the object is positioned at the end of the path.



The displacement spline represents the relative position along a motion path

Displacement paths are composed of locked and unlocked points. Whether a point is locked is determined by how you added it to the polyline. Locked points on the spline have an associated point in the viewer's motion path; unlocked points do not have a corresponding point in the viewer's motion path. Each has a distinct behavior, as described below.

TIP: You can convert displacement splines to X and Y coordinates by right-clicking over the motion path in the viewer and choosing Path#: Polyline > Convert to X/Y Path.

Locked Points

Locked points are the motion path equivalents of keyframes. They are created by moving the playhead position and changing the parameter value. These points indicate that the animated object must be in a specified position on a specified frame. Since these keyframes are only related to position along the path, they can only be moved horizontally along the spline's time axis.

The locked points appear as larger-sized lock icons in the Spline Editor. Each locked key has an associated point on the motion path in the viewer.

Deleting a locked point changes the overall timing of the motion.

Unlocked Points

Unlocked points are created by clicking directly on the spline in the Spline Editor. These points give additional control over the acceleration along the motion path without adjusting the path itself. Conversely, you can add unlocked points in the viewer to control the shape of the motion path without changing the timing.

You can change an unlocked point into a locked point, and vice versa, by selecting the point(s), right-clicking, and choosing Lock Point from the contextual menu.

> For more information on motion paths and locked keyframes, see Chapters 71 and 73 in the DaVinci Resolve manual or Chapters 8 and 10 in the Fusion Studio manual.

Selecting, Moving, and Deleting Keyframes

The placement of keyframes greatly affects the style of the animation. Using the graph, you can select keyframes and move them up or down to change their value or move them left and right to change the timing. Keyframes can be copied and pasted between splines and parameters.

Methods of selecting keyframes:

- Click directly on a keyframe on the spline, or drag a bounding box around the keyframe.
- Drag a bounding box that encompasses multiple keyframes to select more than one.
- To add or remove a keyframe from the current selection, hold down the Command key while selecting the keyframes. This will remove currently selected keyframes and add currently unselected keyframes.
- Press Command-A or right-click in the graph area and choose Select Points > Select All from the contextual menu to select all keyframes from the active splines.

Moving Keyframes

You can freely move keyframes with the mouse, keyboard, or the edit point controls. Keyframes can even pass over existing points as you move them. For instance, if a keyframe exists on frame 5 and frame 10, the keyframe at frame 5 can be repositioned to frame 15.

To move keyframes with the mouse:

- Drag the selected keyframe to its new position in the graph. If more than one keyframe is selected, all selected keyframes will be moved simultaneously.
- Hold down the Option key before dragging the keyframe to constrain its motion to a single axis.

To move keyframes with the keyboard:

- The Up and Down Arrow keys will adjust the value of the keyframes by a small amount.
- The current scale of the graph determines the degree of vertical movement applied to the value with each key press. The closer the zoom in the spline, the finer the adjustment.
- Hold down the Shift key while pressing the Up or Down Arrow keys to increase the value adjustment in larger increments.

To move keyframes using edit fields:

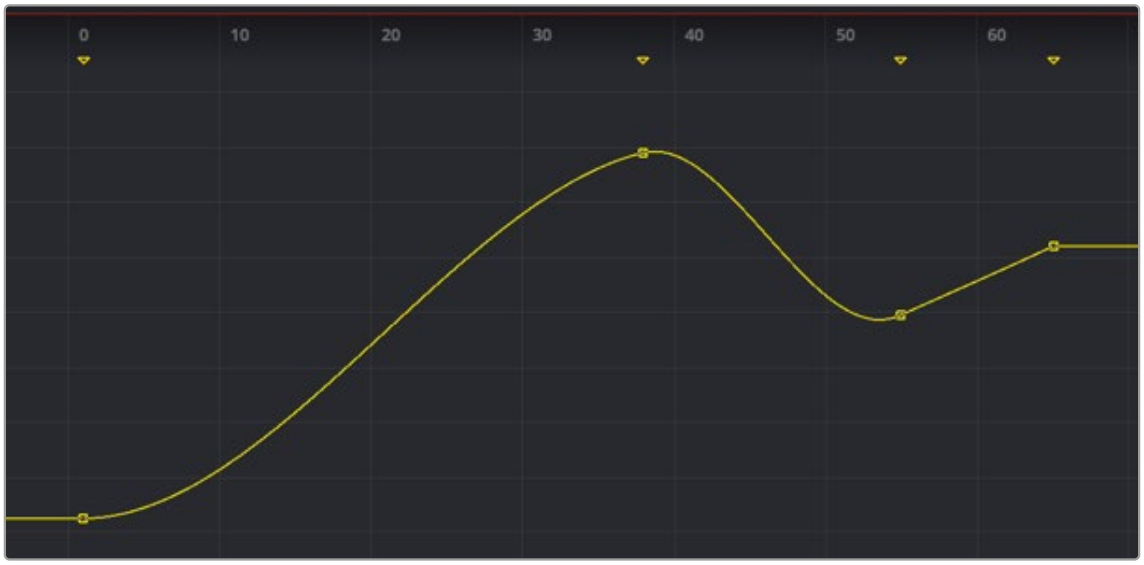
- The Value and Time Editors are found on the far right of the toolbar. These number fields allow explicit values to be entered for selected keyframes. These controls are explained in more depth later in this chapter.

To delete one or more keyframes:

- Select one or more keyframes and press the Delete or Backspace key on the keyboard. This only removes the keyframes; it does not remove the spline, even if there are no keyframes on the spline. To remove the spline, right-click over the parameter in the Inspector and choose Remove [parameter name].

Showing Key Markers

You can adjust the position of the keyframes in time, without worrying about manipulating splines, by using the key markers. The horizontal time axis can show markers that indicate the position of each keyframe. The display of these markers is enabled by right-clicking in the graph and choosing Show > Key Markers from the contextual menu, or by clicking on the Show Key Markers button in the toolbar.



The key markers show keyframes in the horizontal axis using the same color as the splines

Copying and Pasting Keyframes

To precisely match animation, Keyframes can be copied to a new location on the same spline or onto completely different splines and different tools.

There are two options in the graph's contextual menu for copying keyframes. Choosing Copy Points (Command-C) copies all selected points. Choosing Copy Value copies a single point identified by the pointer from multiple selected points. This does not deselect your selection set, and you can pick out numbers as needed.

To copy and paste points to a new location on the same spline:

1. Select the desired keyframes on the spline.
2. Right-click over the spline and choose Copy Points from the contextual menu or press Command-C.
3. Click in an empty area of the graph to deselect all the copied points.
4. Move the playhead to the area of the spline where you want the points pasted and press Command-V.

   Or, move the pointer over the spline where you want the points pasted, and when the spline highlights, right-click and choose Paste Points/Value.

Alternatively, you can copy and paste keyframes by dragging them with the mouse. After you select the points, hold down the Command key and drag the points along the spline to where you want them pasted.

To copy and paste keyframes from one spline to another:

1. Make one spline the active visible spline and select the desired keyframes on the spline.
2. Right-click over the spline and choose Copy Points from the contextual menu or press Command-C.
3. Set the spline to viewed or disabled using the status checkbox next to the spline's name in the header.
4. Make the destination spline the active visible spline and select the keyframe on the spline where the new keyframes should be pasted.
5. Right-click and choose Paste Points/Value or press Command-V.

You can copy a single point's value from a group of selected points. Since this process does not deselect the selected set, you can continue picking out values as needed without having to reselect points.

To copy and paste a keyframe value:

1. Make one spline the active visible spline and select all the keyframes on the spline.
2. Right-click over a single point and choose Copy Value from the contextual menu.
3. Set the spline to viewed or disabled using the status checkbox next to the spline's name in the header.
4. Make the destination spline the active visible spline and select the keyframe on the spline where the new keyframe should be pasted.
5. Right-click and choose Paste Points/Value or press Command-V.

Keyframes can also be pasted with an offset, allowing you to duplicate a spline shape but increase the values or shift the timing using an offset to X or Y.

To paste keyframes points and values with an X or Y offset:

1. Make one spline the active visible spline and select the desired keyframes on the spline.
2. Right-click over the spline and choose Copy Points from the contextual menu or press Command-C.
3. Set the spline to viewed or disabled using the status checkbox next to the spline's name in the header.

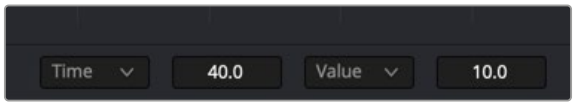
4 Make the destination spline the active visible spline and select the keyframe on the spline where the new keyframes should be pasted.

5 Right-click and choose Paste with Offset. In the Offset dialog, enter the Y value, which will be added to the values of the pasted keyframes.

TIP: You cannot copy and paste between different spline types. For instance, you cannot copy from a Bézier spline and paste into a B-spline.

Time and Value Editors

The Time and Value Editors in the lower-right corner of the Spline Editor are used to change the position and parameter value of a keyframe by entering a number into the number field for each button. Each field can switch between three modes that help modify the time and value of a keyframe in three precise but distinct ways. The default mode for each field takes the explicit frame number or parameter value at which you want the keyframe set. The other modes offset and scale the keyframe's position or value.



Use the number fields to enter in a value or a specific time to change the selected keyframe

Time Editor

The Time Editor is used to modify the current time of the selected keyframe. You can change the Time mode to enter a specific frame number, an offset from the current frame, or spread the keyframes based on the distance (scale) from the playhead. You can select one of the three modes from the Time mode drop-down menu.



Three time editing modes are selectable from the Time mode drop-down menu

Time

The number field shows the current frame number of the selected control point. Entering a new frame number into the field moves the selected control point to the specified frame. If no keyframes are selected or if multiple keyframes are selected, the field is empty, and you cannot enter a time.

Time Offset

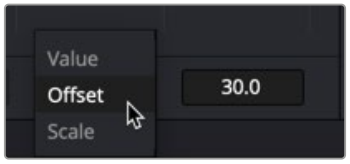
Selecting T Offset from the drop-down menu changes the mode of the number field to Time Offset. In this mode, the number field offsets the selected keyframes positively or negatively in time. An offset of either positive or negative values can be entered. For example, entering an offset of 2 moves a selected keyframe from frame 10 to 12. If multiple keyframes were selected in the previous example, all the keyframes would move two frames forward from their current positions.

Time Scale

Selecting T Scale from the drop-down menu changes the mode of the number field to Time Scale. In this mode, the selected keyframes' positions are scaled based on the position of the playhead. For example, if a keyframe is on frame 10 and the playhead is on frame 5, entering a scale of 2 moves the keyframe 10 frames forward from the playhead's position, to frame 15. Keyframes on the left side of the playhead would be scaled using negative values.

Value Editor

The Value Editor is used to modify the selected keyframe's parameter value using one of three Value modes. You can change the Value mode to enter a specific value for a parameter, an offset from the value, or to spread the values. The mode is chosen from the Value mode drop-down menu.



Three value editing modes are selectable from the Value mode drop-down menu

Value

The number field shows the value of the currently selected keyframes. Entering a new number into the field changes the value of the selected keyframe. If more than one keyframe is selected, the displayed value is an average of the keyframes, but entering a new value will cause all keyframes to adopt that value.

Value Offset

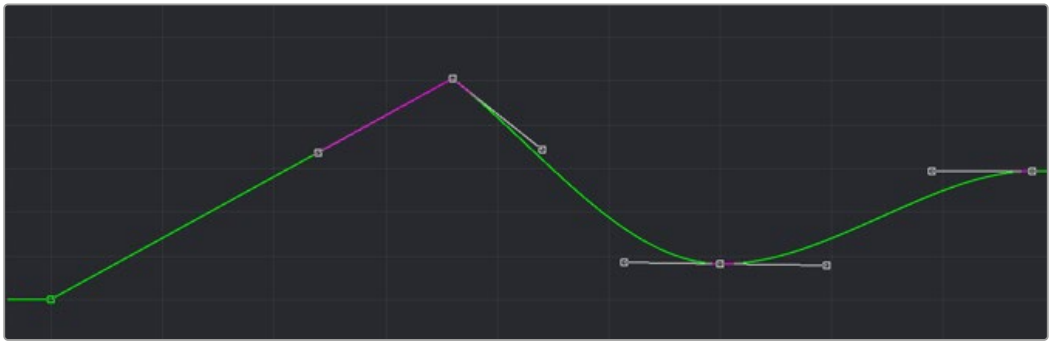
Choosing Offset from the drop-down menu sets the Value Editor to the Offset mode. In this mode, the value for the selected keyframes are offset positively or negatively. An offset of either positive or negative values can be entered. For example, entering a value of -2 changes a value from 10 to 8. If multiple keyframes are selected, all the keyframes have their values modified by -2.

Value Scale

Choosing Offset from the drop-down menu sets the Value Editor to the Scale mode. Entering a new value causes the selected keyframes' values to be scaled or multiplied by the specified amount. For example, entering a value of 0.5 changes a keyframe's value from 10 to 5.

Modifying Spline Handles

All Bézier spline keyframes have a pair of control handles to shape the spline as it passes through the key point. These handles are only displayed when the keyframe is selected. Initially, these handles are set to linear, creating straight line changes between keyframes. However, any control point can be made smooth by right-clicking over it and choosing Smooth or pressing Shift-S.



Bézier splines can mix linear and smooth curves

Dragging on a keyframe's handles adjusts the slope of the segments passing through the spline. By default, the two control handles on a control point are locked together so that if one moves, the one on the other side moves with it. This maintains a constant tension through the keyframe. There are situations, however, when it is desirable to modify these control handles separately for a more pronounced curve or effect.

To temporarily break the control handles on a Bézier spline, moving one independently of the other:

1. Select the control point to be modified.
2. Hold down the Command key and drag one of the control handles. They will now move independently of each other, as long as the Command key is held down.

To treat all Bézier handles as independent in the Spline Editor:

Right-click in the graph and choose Independent Handles from the Options contextual menu.

Enabling this option causes all the Bézier handles to be independent. This is the same as using the Command key when moving a handle, except it is applied to all control points until it is disabled.

Reducing Points

When there are too many control points too close together on a spline, you can choose Reduce Points to decrease their number, making it easier to modify the remaining points. The overall shape of the spline is maintained as closely as possible while eliminating redundant points from the path.

To reduce the number of points on a spline:

1. Select the range of keyframes you want to thin out.
2. Right-click in the graph area and choose Reduce Points from the contextual menu.
3. When the Reduce Points dialog appears, drag the slider to a lower value.

You can set the slider value as low as possible as long as the spline still closely resembles the shape of your original spline.

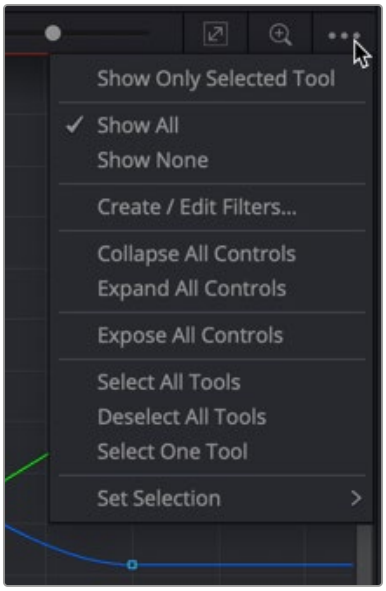
TIP: When the value is 100, no points will be removed from the spline. Use smaller values to eliminate more points.

Filtering the Spline Editor

The animation splines for multiple parameters can be displayed within the Spline Editor simultaneously, and Fusion offers several ways for you to choose which spline to view and which to edit.

A complex composition can easily contain dozens, if not hundreds, of animation curves. As a composition grows, locating a specific spline can become more difficult. There are two ways to filter the splines shown in the Spline Editor: display selected tools only or create a filter to show only certain tools.

The Spline Editor includes different ways to control which splines are displayed. The majority of these options are available in the Options menu, located in the upper-right corner of the Spline Editor panel.



The Options menu is used to control which splines are displayed in the Spline Editor

Show Only Selected Tool: You can choose to limit the splines displayed in the Spline Editor by showing only the splines from selected tools. Choosing this option at the top of the Options menu displays only the splines for tools currently selected in the Node Editor.

Show All/None: The default behavior of the Spline Editor displays all the splines for all the nodes with animated parameters. You can override this by enabling Show Only Selected Tools in the Options menu. You can also disable the Show All setting by choosing Show None, in which case the Spline Editor remains empty.

Expose All Controls: The Expose All Controls option is a way of not filtering the parameters. Choosing this option displays all parameters in the Spline Editor header for all nodes in the Node Editor. It can be a fast way of activating one of the parameters and automatically adding an animation spline for it if one does not exist.

With a large number of nodes displayed, which themselves might have a large number of parameters, this might lead to cluttering and slowing down the interface. This option is most effective when used in conjunction with the Show Only Selected Tool option to limit the number of nodes and parameters displayed and yield optimum performance.

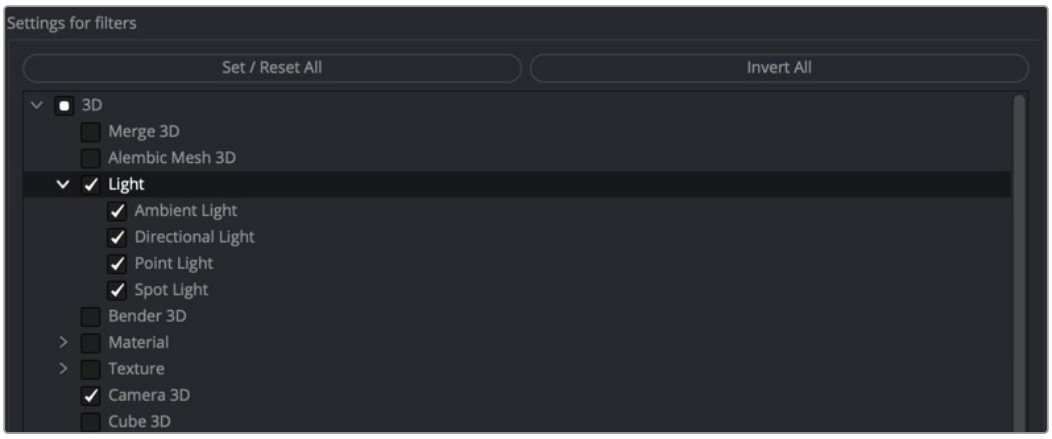
Follow Active: The Follow Active option is located by right-clicking in the graph and choosing Options > Follow Active. This option provides a way to filter the splines in the graph while not filtering the header list of tools. Where the Show Only Selected Tool option hides other tools in the header, the Follow Active option leaves the header displaying all the tools but automatically enables only the splines of the Active tool.

Working with Filters

Filters allow you to select the specific types and classes of tools shown in the Spline Editor and Keyframes Editor. For example, you can make a filter that shows only particle nodes or one that only shows color correction and brightness/contrast tools.

To create a filter:

1. From the Options menu, choose Create/Edit Filters.
2. Click the New button to create a new filter and name the new filter in the dialog box.
3. Enable a checkbox next to the entire category or the individual tools in each category to determine the tools included in the filter.



Enable each tool you want to keep in the Spline Editor when the filter is selected

The Invert All and Set/Reset All buttons can apply global changes to all the checkboxes, toggling the selected states as described.

To switch the selection state of the categories when creating a filter list:

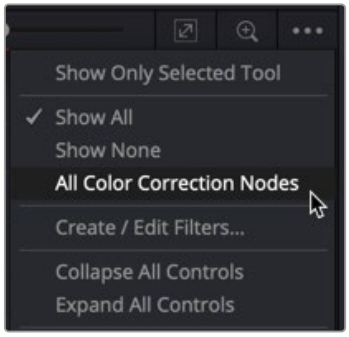
1. Click the Invert All button.
2. After configuring the custom filter, click the Save button to close the Settings dialog and save the filter.

To enable all checkboxes or disable all checkboxes:

1. Click the Set/Reset All button as many times as needed until all categories are either checked or unchecked.
2. After configuring the custom filter, click the Save button to close the Settings dialog and save the filter.

To apply a filter to the Spline Editor:

- Choose the desired filter by name from the Options menu. The filter applies to both the Spline Editor and the Timeline.



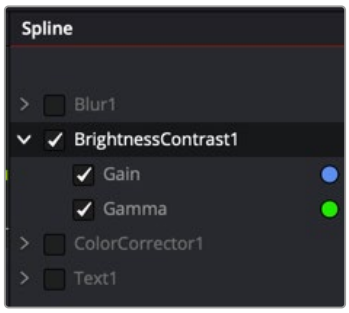
Each filter you create is listed in the Options menu

To disable a filter and show all tools in the Spline Editor again:

- Choose Show All from the Options menu.

Changing a Spline's Status

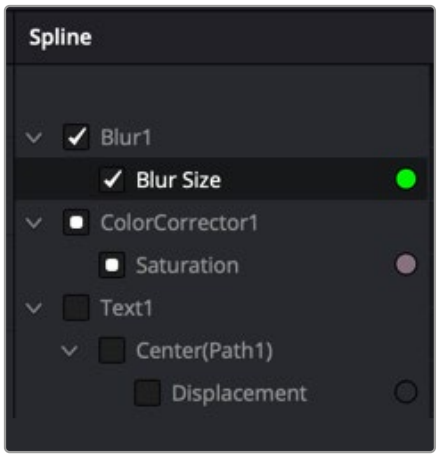
The Spline header is a hierarchical list of animated parameters and their parent nodes. Clicking the disclosure arrow next to a tool's name reveals all the names of the animated parameters on that tool. Clicking directly on the parameter name in the Spline header activates that spline for display and editing.



The Spline header with the Brightness Contrast tool and its animated Gain and Gamma parameters

Tool Status Checkbox

Next to the name of each spline is a checkbox that indicates the spline's status. When you select a parameter name, the checkbox becomes active, allowing you to see and edit the spline in the graph. There are three selection modes for each checkbox: active, viewed, and disabled. Clicking directly on the checkbox will toggle it between these three states. Changing the state of the parent node checkbox sets the state for all splines for that node.



Active: When the checkbox is enabled with a check mark, the spline is displayed in the graph and can be edited.

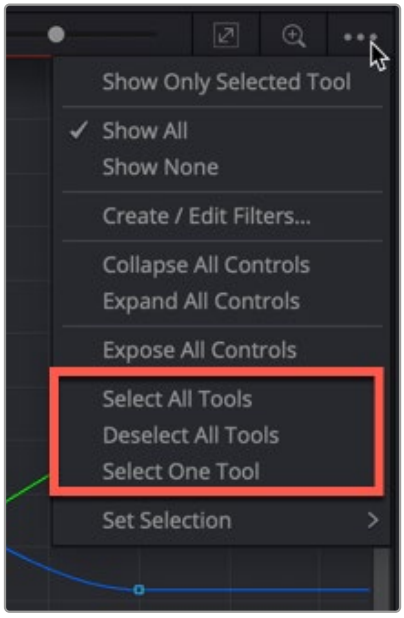
Viewed: When the checkbox is enabled with a solid gray box, the spline is visible in the graph but cannot be edited. It is read-only.

Disabled: When the checkbox is clear, the spline is not visible in the graph and cannot be edited.

Checkboxes determine which splines are visible and editable

Selection States

There are three selection options, labeled Select All Tools, Deselect All Tools, and Select One Tool, that determine how the items in the Spline Editor header behave when a checkbox or label is selected to activate a spline. These states are located in the Options menu in the upper-right corner of the Spline Editor.



Select All Tools: Choosing this option activates all splines for editing.

Deselect All Tools: Choosing this option sets all spline checkboxes to disabled.

Select One Tool: This option is a toggle. When Select One Tool is chosen from the menu, only one spline in the header is active and visible at a time. Clicking on any spline's checkbox will set it to active, and all other splines will be cleared. When disabled, multiple splines can be active in the header.

The Options menu selection states make it easier to select and deselect all parameters in the header

Selection Groups

It is possible to save the current selection state of the splines in the header, making selection groups that can easily be reapplied when needed. To create a selection group, right-click over any parameter in the header or in an empty area of the graph and choose Save Current Selection from the contextual menu. A dialog will appear to name the new selection.

To reapply the selection group, choose the selection group by name from the Set Selection menu in the same contextual menu. Other context menu options allow selection groups to be renamed or deleted.

Reshaping Splines Using the Toolbar

There are several ways to manipulate the shape of a spline, thereby altering the animation that spline generates. Other than manually adjusting Bézier handles, you can quickly squish, stretch, loop, and reverse a spline. You can also quickly change the interpolation between keyframes from the default linear motion to more natural smooth motion. All these options are provided at the bottom of the Spline Editor in the toolbar. The toolbar is divided into different groups for setting interpolation, reversing splines, looping splines, time stretching, and reshaping splines.

Interpolation

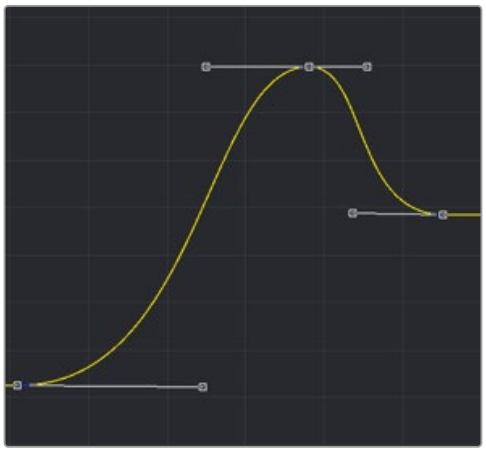
Keyframes are specific frames in an animation where control points are set to exact values on a given parameter. Interpolation is the method used to fill in the unknown values between two keyframes. Fusion automatically interpolates between two keyframes. However, you may want to modify the interpolation to achieve a specific style of animation. The Spline Editor includes several interpolation methods you can choose from using the toolbar.



Interpolation buttons in the toolbar: Smooth, Linear, Invert, Step In, and Step Out

Smooth

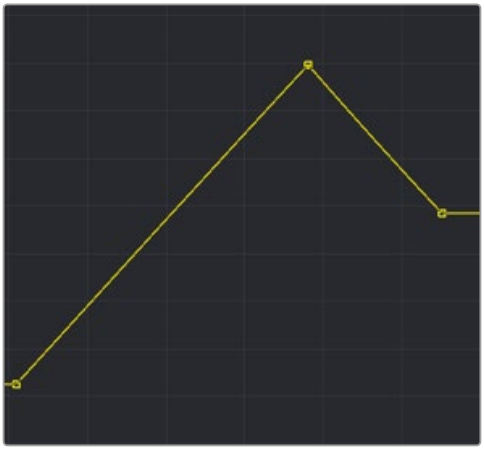
A smoothed segment provides a gentle keyframe transition in and out of the keyframe by slightly extending the direction handles on the curve. This slows down the animation as you pass through the keyframe. To smooth the selected keyframe(s), press Shift-S or click the toolbar's Smooth button.



Smooth interpolation between keyframes

Linear

A linear segment effectively takes the shortest route between two control points, which is a straight line. To make the selected keyframe(s) linear, press Shift-L or click the Linear button in the toolbar.



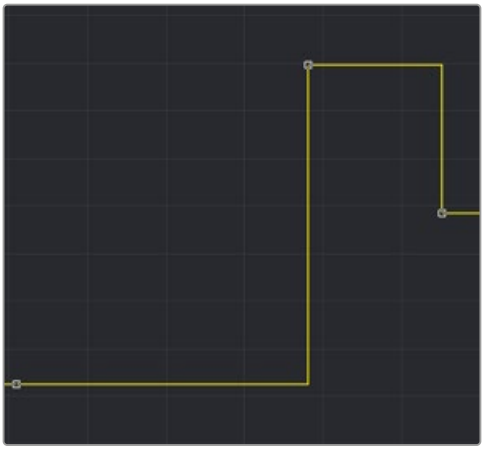
Linear interpolation between keyframes

TIP: Invert is used only for non-animated LUT splines, which are currently only available in the LUT Editor window.

Step In/Step Out

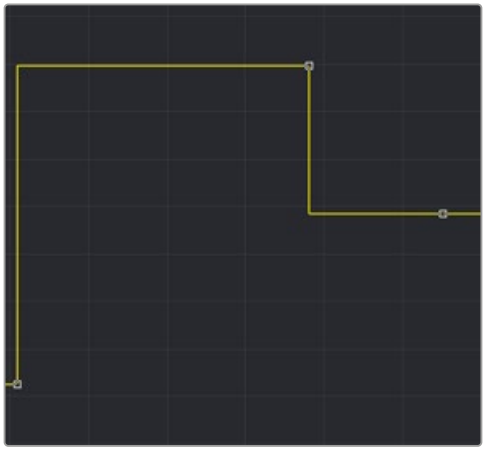
On occasion, it is not desirable to have any interpolation between two keyframes. Instead, the value of one keyframe may hold its value until another keyframe changes it. For these cases, use the Step In or Step Out mode.

Step In causes the value of the previous keyframe to hold, then jump straight to the value of the next keyframe.



Step In holds a value until the next keyframe is reached in the comp

Step Out causes the value of the selected keyframe to hold right up to the next keyframe.



Step Out switches immediately to the next keyframe value in a comp

Step In and Step Out modes can be set for selected keyframes by clicking on the toolbar buttons for each mode, or by right-clicking and choosing the appropriate option from the contextual menu. The keyboard shortcuts I and O can also be used to enable Step In and Step Out on selected keyframes.

Reversing Splines

Reverse inverts the horizontal direction of a segment of an animation spline. To apply reverse, choose a group of points in a spline and click the Reverse button, or right-click and choose Reverse from the contextual menu, or press the V key. The group of points is immediately mirrored horizontally in the graph. Points surrounding the reversed selection may also be affected.



The Reverse button in the toolbar

Looping Splines

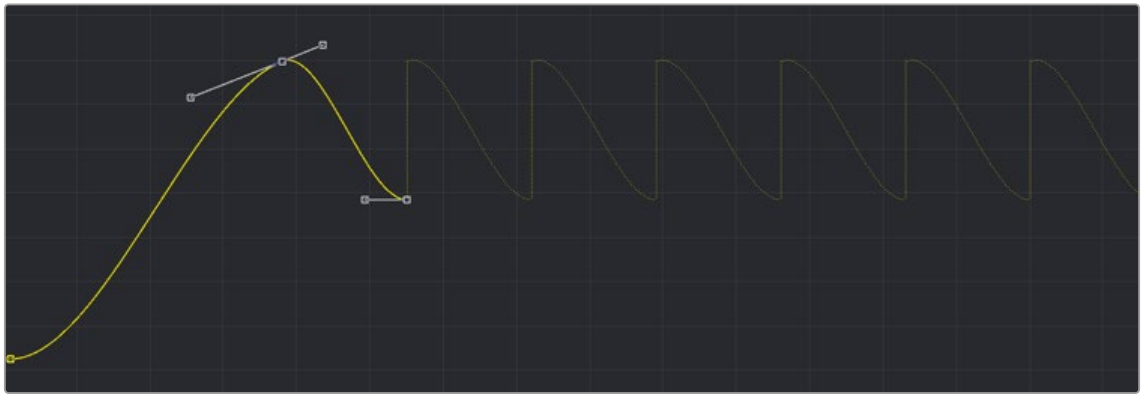
It is often useful to repeat an animated section, either infinitely or for a specified number of times, such as is required to create a strobing light or a spinning wheel. Fusion offers a variety of ways to repeat a selected segment.



The various Loop buttons in the toolbar

Set Loop

To repeat or loop a selected spline segment, select the keyframes to be looped. Select Set Loop from the contextual menu or click on the Set Loop button in the toolbar. The selected section of the spline repeats forward in time until the end of the global range, or until another keyframe ends the repeating segment.



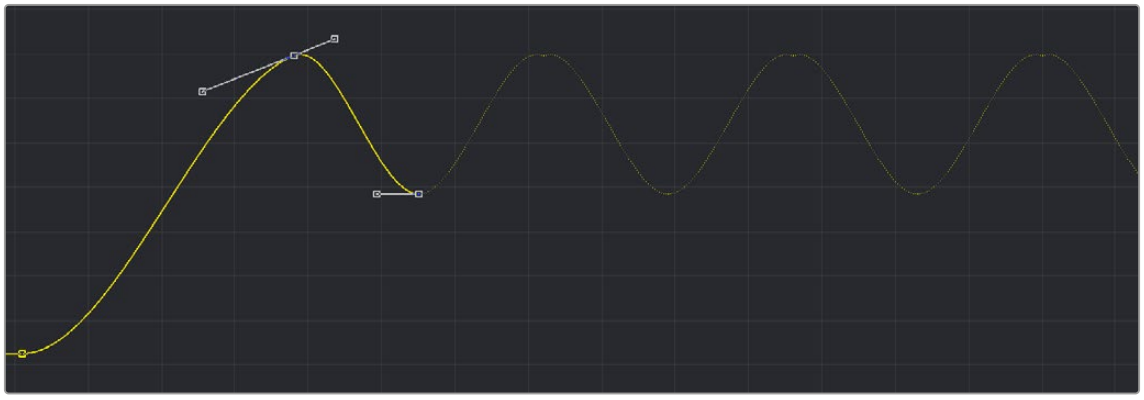
A looped section in the graph

Changing and Removing the Loop

You can change the looped segment by modifying any of the keyframes or control points originally used to create the loop. Simply select one of the originating key points, make any necessary modifications, and the looped segment updates. To remove the loop, select the keyframes you used to create the loop, and then click the Loop button in the toolbar.

Ping-Pong

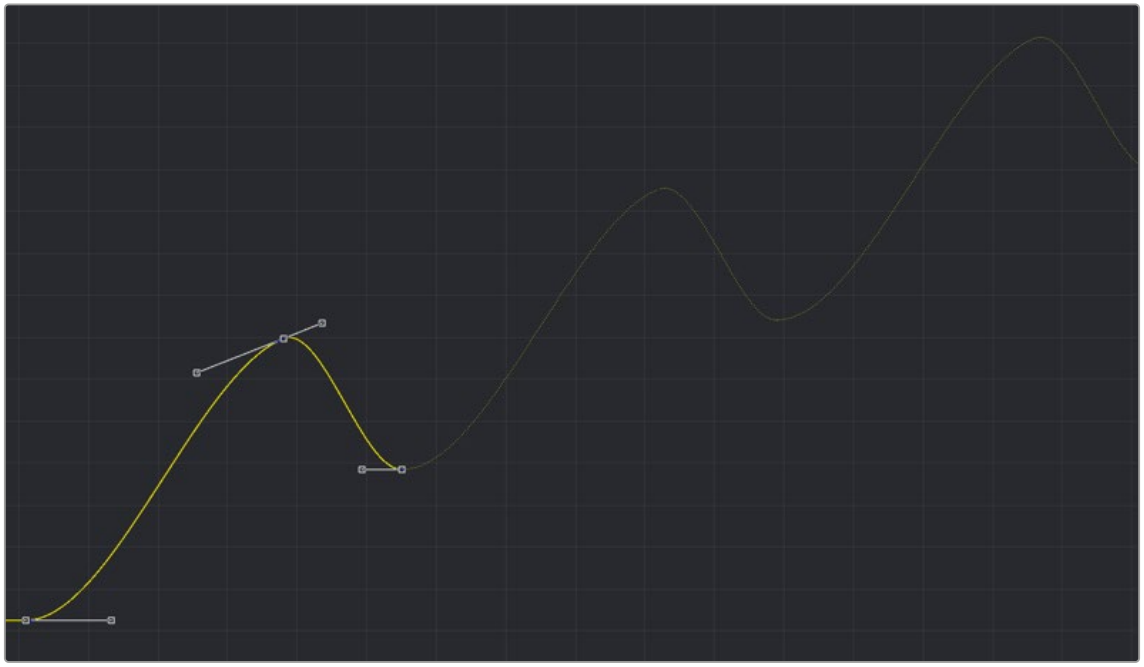
The Ping-Pong Loop mode repeats the selected segment, reverses each successive loop, and then repeats. Ping-pong looping can be enabled on the selected segments from the context menu or the toolbar.



A ping-pong section in the graph

Relative Loop

The Relative Loop mode repeats the segment like the Loop, but each repetition adds upon the last point of the previous loop so that the values increase steadily over time.



A Relative Loop section in the graph

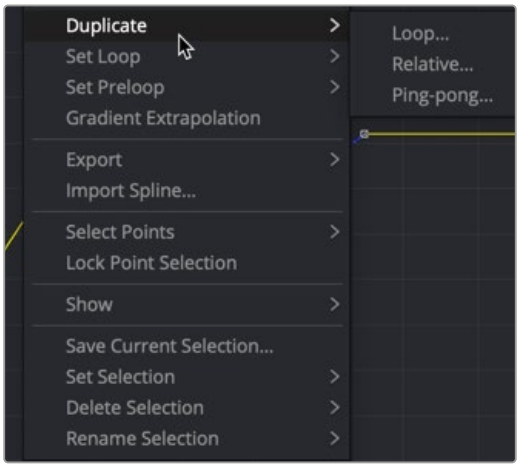
Looping Backward

You can choose Set Pre-Loop by right-clicking in the graph area and choosing it from the contextual menu. This option contains the same options for looping as the Loop option buttons in the toolbar, except that the selected segment is repeated backward in time rather than forward.

Repeating Splines X Number of Times

You can duplicate splines and repeat them a set number of times by right-clicking in the graph area and choosing Duplicate from the contextual menu. Duplicated splines are like looped splines, except that the selected segment repeats only a specified number of times, and each repetition is a copy rather than an instance. Adjustments to the original segment do not alter the shape of its repetitions.

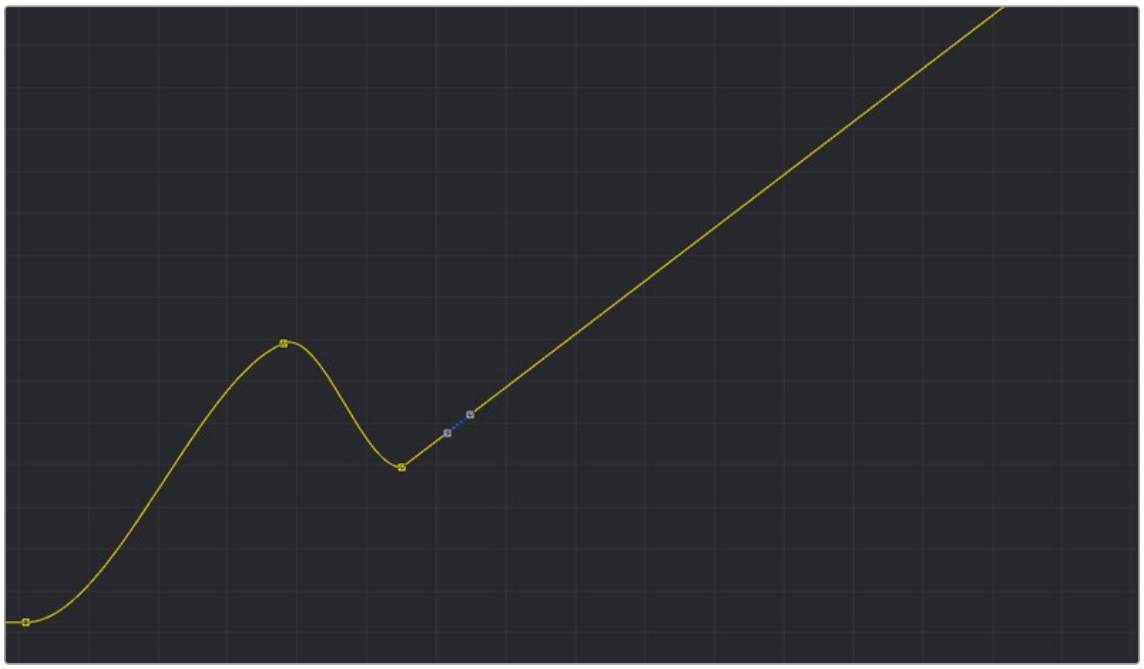
The Duplicate modes are only accessed from the Duplicate contextual menu, which reveals a submenu with all the looping modes described above. Selecting any of these modes opens a dialog in which the number of repetitions can be entered.



The Duplicate contextual submenu

Gradient Extrapolation

You can choose Gradient Extrapolation by right-clicking in the graph area and choosing it from the contextual menu. This option continues the trajectory of the last two keyframes.



The Gradient Extrapolation applied to the spline

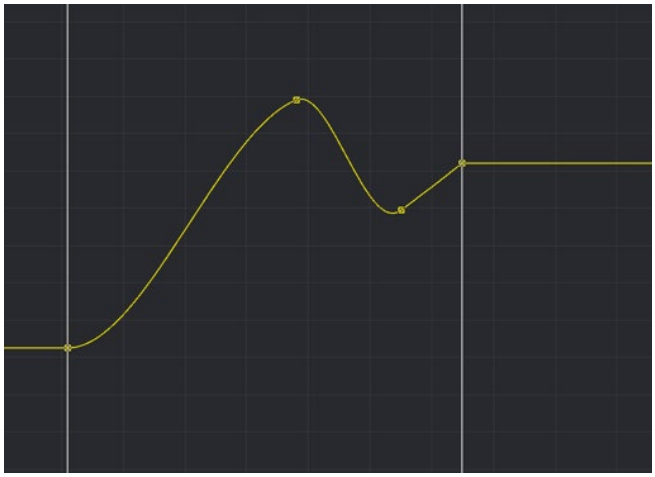
Time Stretching

Time Stretching allows for a selected group of keyframes to be proportionally stretched or squashed. This allows you to change the duration of the animation while keeping the relative distance between each keyframe. To enable spline stretching, select the group of keyframes that you want to time stretch, and then choose Modes > Time Stretching from the graph's contextual menu or click the Time Stretch button in the toolbar.



The Time Stretch button in the toolbar

When you have more than one keyframe selected on the spline, enabling Time Stretch surrounds the outer keyframes with two vertical white bars. Drag on the white vertical bars to stretch or shrink the timing of the spline segments within the bars. Drag these bars back and forth to stretch or squash the spline segment.



The Time Stretch bars in the graph

TIP: If no keyframes are selected when you enable Time Stretch, drag a rectangle to set the boundaries of the Time Stretch.

To disable the Time Stretching mode:

- Click on the Time Stretch button in the toolbar again or reselect Modes > Time Stretching from the contextual menu.

Shape Box

The Shape Box transform mode is similar to Time Stretching; however, it can adjust the vertical scaling of keyframe values as well as time.

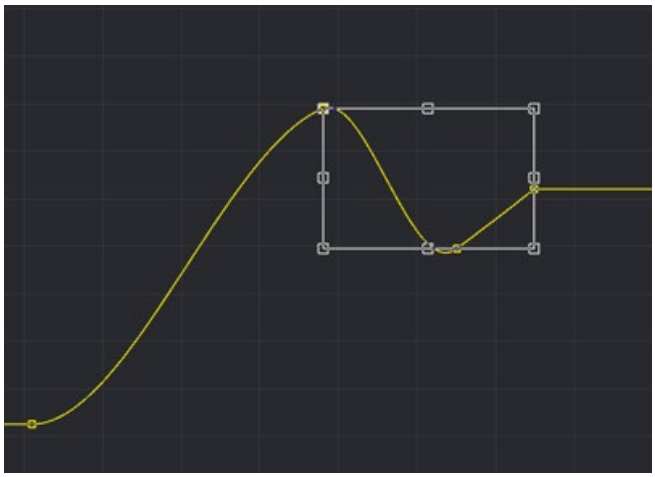
To enable the Shape Box, do one of the following:

- Select a group of keyframes, and then choose Modes > Shape Box from the contextual menu.
- Select a group of keyframes, and then click the Shape Box button in the toolbar.
- Select a group of keyframes, and then press Shift-B to enable or disable the Shape Box mode.



The Shape Box button in the toolbar

A white rectangle outlines the selected points when the mode is enabled. To scale, skew, or stretch the spline, drag on any of the control points located around the box. To move all the keyframes, drag on the box edges.



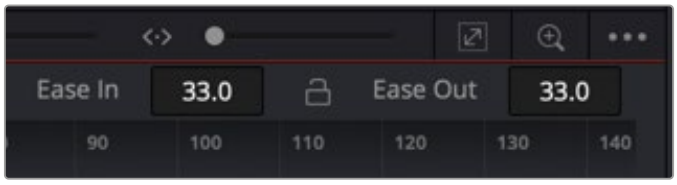
The Shape Box in the graph

TIP: If no points are selected, or if you want to select a new group of keyframes, you can drag out a new rectangle at any time.

Ease In/Out

For a more precise way to adjust the length of Bézier direction handles attached to selected keyframes, you can use the Spline Ease dialog. To show the dialog, select a keyframe in the graph, and then choose Edit > Ease In/Out from the graph's contextual menu or press T on the keyboard.

The Ease In/Out controls appear above the graph area. You can drag over the number fields to adjust the length of the direction handles or enter a value in the fields.



The Ease In/Out controls above the graph area

Clicking the Lock In/Out button will collapse the two sliders into one, so any adjustments apply to both direction handles.

Importing and Exporting Splines

Spline shapes can be imported and exported from or to an ASCII text file. This makes it easier to save complex spline curves for later reuse, or to transfer tracking, path, and animation data from one application to another. Exported splines are assigned the file extension .spl for easy identification.

To export a spline:

1. Select the active spline in the Spline Editor.
2. Right-click on the spline in the graph area to display the contextual menu and select Export.
3. Choose from three format options in the submenu.
4. Enter a name and location in the file browser dialog, and then click Save.

Exporting a spline gives you three options. You can export the Samples, Key Points, or All Points. Samples adds a control point at every frame to create an accurate representation of the spline. Key Points replicates the control point positions and values on the spline using linear interpolation. All Points exports the spline as you see it in the Spline Editor, using the same position, value, and interpolation.

To import a spline:

1. Add an animation spline for the parameter.
2. In the Spline Editor, right-click on the animation spline and select Import Spline from the contextual menu.
3. In the File Browser dialog, select the spline curve .spl file, and then click Open.

Importing a new curve will replace any existing animation on the selected spline.

Chapter 11

Animating with Motion Paths

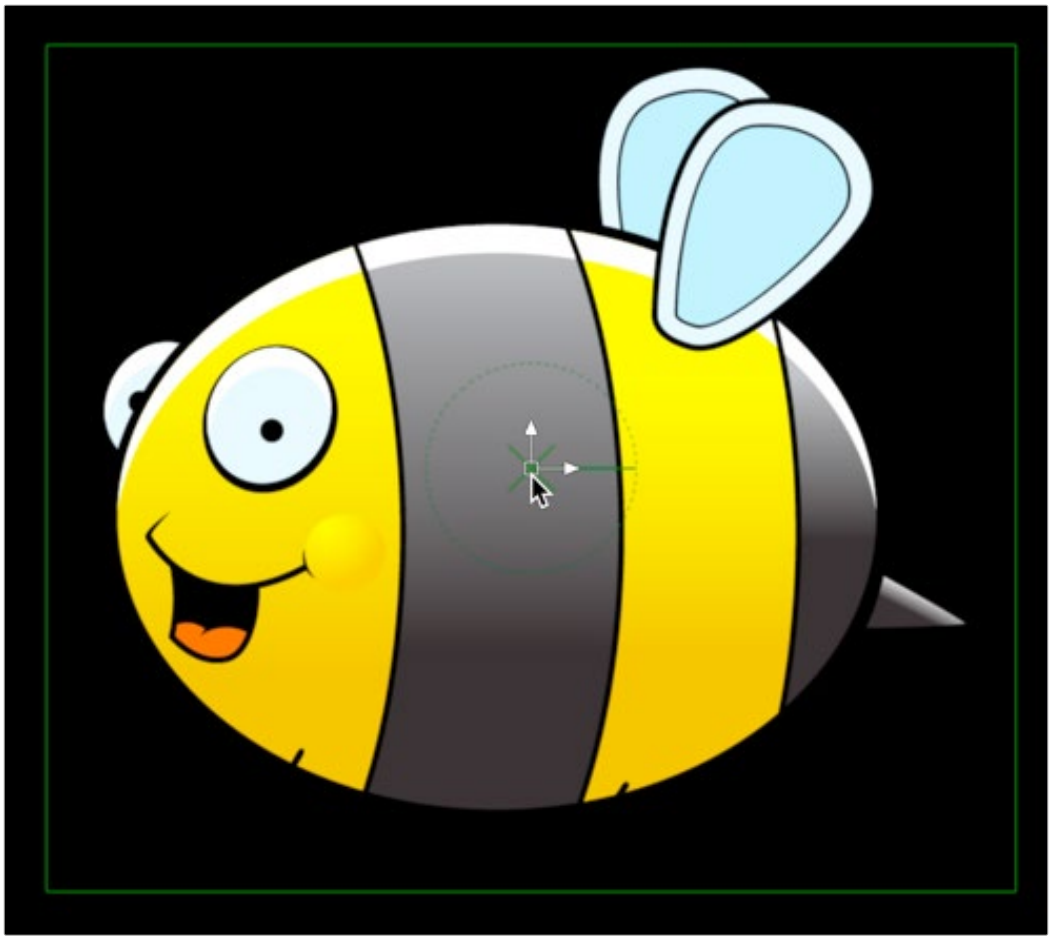
Layers and 3D objects can move along a designated spline shape to create motion path animations. This chapter discusses how you can create, edit, and use motion paths in Fusion.

Contents

Animating Using Motion Paths 306
Types of Motion Paths 307
Polyline Path 307
Path Modifier 311
Controlling Speed and Orientation along a Path 311
XY Path 314
Types of Control Points 316
Locked Points 316
Unlocked Points 319
Locking and Unlocking Points 320
Tips for Manipulating Motion Paths 320
Compound Motion Paths Using Path Centers 320
Copying and Pasting Motion Paths 321
Removing Motion Paths 321
Recording Motion Paths 322
Importing and Exporting Polylines 322
Native Format 322

Animating Using Motion Paths

Motion paths are created from splines (polylines) including paint strokes and masks that have a path modifier applied to them. The path modifier defines the movement that transforms ordinary spline shapes into motion paths. You apply the path modifier either explicitly to the Center X/Y or Pivot X/Y parameters or implicitly by keyframing the Center or Pivot parameters. In either case, the spline path is displayed and then visually adjusted in the viewers. For all motion paths, the coordinate control in the viewer represents the position of an object or effect, such as a Transform or Rays node's center that moves along the path. Coordinate controls are represented onscreen with vertical and horizontal positioning arrows or an X.



A Center Offset onscreen control for an object's Transform node

The following nodes have parameters that can be animated using path modifiers to move an image around the composition. These include, but are not limited to:

- **Transform:** Center X/Y can be animated to move an image around.
- **DVE:** Center X/Y can be animated to move an image around.
- **Merge:** Center X/Y can be animated to move the Foreground connected image around.
- **Paint:** Stroke Controls > Center X/Y can be animated to move a stroke around.
- **Camera 3D node:** Translation X/Y/Z
- **Shape 3D node:** Translation X/Y/Z

The following nodes have parameters that can be animated using paths to alter the direction of a visual effect. These include, but are not limited to:

- **Directional Blur:** Center X/Y can be animated to change the direction of the blur.
- **Hot Spot:** Primary Center X/Y can be animated to move the hot spot around.
- **Rays:** Center X/Y can be animated to change the angle at which rays are emitted.
- **Polygon/BSpline/Ellipse/Rectangle/Triangle mask:** Center X/Y can be animated to move the mask.
- **Corner Positioner:** Top Left/Top Right/Bottom Left/Bottom Right X/Y can be animated to move each corner of the corner-pinned effect.
- **Vortex:** Center X/Y can be animated to move the center of the warping effect.

NOTE: It's not possible to add a motion path to a one-dimensional value, such as blur strength or merge angle. However, you can use the Spline Editor to edit these kinds of values in a visual way.

Types of Motion Paths

There are three types of motion paths: Polyline paths, XY paths, and 3D motion paths for 3D scenes.

- A Polyline path is generated by applying the path modifier. It uses two splines to define the path; one for the shape of the path displayed in the viewer, and a Displacement spline for the speed of the object along the path, displayed in the Spline Editor. The Polyline path is the default type of path modifier, and most documentation in this chapter assumes that this type is used.
- The XY path modifier employs a spline for the X position of the point and another for the Y position. The XY path modifier is explained in detail later in this chapter.
- 3D motion paths pertain only to positional controls within 3D scenes.

Polyline Path

Polyline paths are the easiest motion paths to work with. You can use the spline shape in the viewer to control the shape of the path, while a single Displacement curve in the Spline Editor is used to control the acceleration along the path. The most obvious way to create a Polyline motion path is by keyframing the Center X/Y parameter of a Transform node in the Inspector.

To create a Polyline motion path using the Center X/Y parameter in the Inspector:

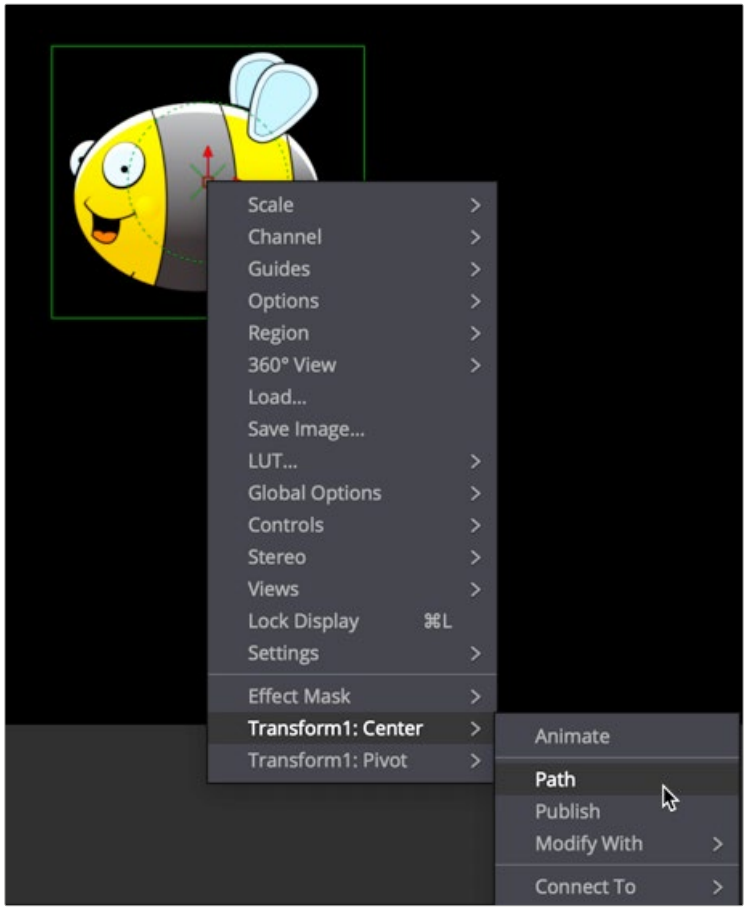
1. Position the playhead on the frame where the motion will begin.
2. In the Inspector, click the gray Keyframe button to the right of the Center X and Y parameters.

   This action applies the path modifier in the Modifiers tab in the Inspector.
3. Adjust the Center X and Y for the first keyframe position.
4. Position the playhead on the frame where the motion should change or stop.
5. In the Inspector, change the Center X and Y parameters to set a keyframe for the new location automatically.
6. In the viewer, modify and refine the motion path by selecting a control point and using any of the spline controls in the viewer toolbar.
7. Open the Spline Editor and adjust the Displacement spline to control the speed of the object along the path.

Keyframing Center X/Y is not the only way to apply the path modifier. An alternative method is to apply the path modifier to the Center X/Y parameter either in the Inspector or using the coordinate control in the viewer.

To apply a Polyline path to the Center X/Y parameter:

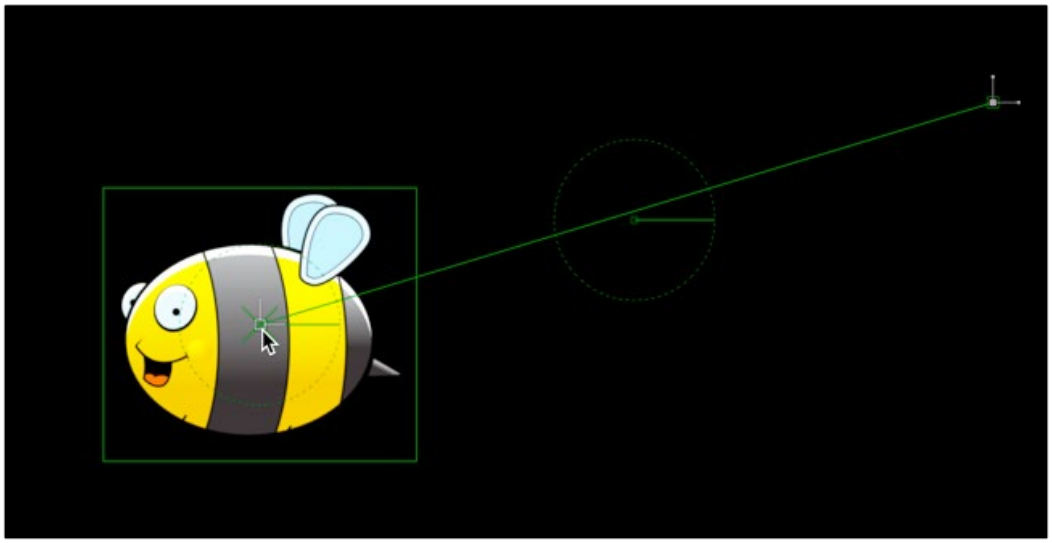
1. Position the playhead on the frame where the motion will begin.
2. Position the center coordinate control for the object, or effect, at its starting position.
3. Right-click the onscreen center coordinate control in the viewer (or over the Center X and Y parameter in the Inspector), and choose Path from the contextual menu for that control.



Applying a path modifier to the center of a Transform node

The object now has a path modifier applied, so without setting a keyframe you can drag the object to begin creating a motion path in the viewer.

4. Move the playhead to a new frame.
5. Drag the onscreen coordinate control or adjust the Offset or Center values in the Inspector. A keyframe is automatically created on the motion path, and a polyline is drawn from the original keyframe to the new one.



Two keyframes spaced several frames apart display a motion path showing the direction of animation.

6 The position of the center control is interpolated between the two keyframes. You can continue adding points by moving the playhead and adjusting the object's position until the entire motion path you need is created. For motion paths, there's no need to close the spline path; you can leave it open.

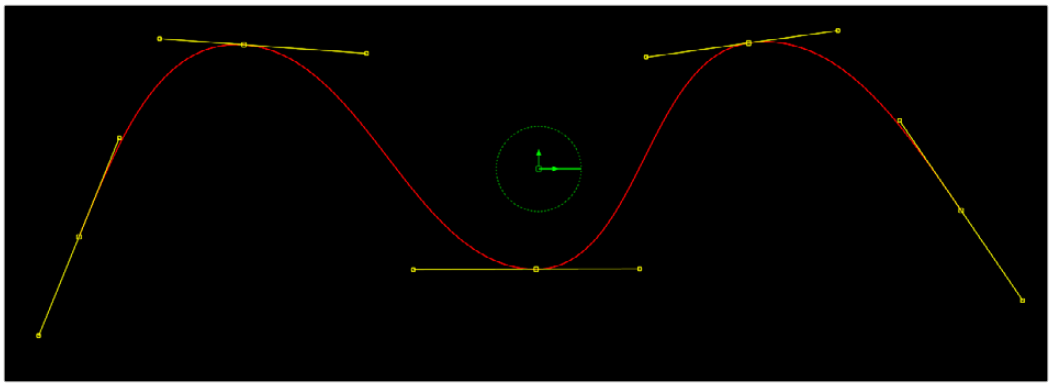
7 Upon completion, set the polyline to Insert and Modify mode by selecting a point on the path and pressing Command-I or clicking the Insert and Modify button on the toolbar. Don't worry too much about the overall shape of the motion path at this point. The shape can be refined further by adding additional points to the polyline and modifying the existing points.

A final alternative method for creating a motion path is to draw a spline shape first and then connect a path modifier to the spline. Using any of Fusion's spline tools, you can draw the shape of the path and then connect the path modifier to the published spline. Once the path modifier and the published spline are connected, you can keyframe the Displace parameter to move an image along the path. This method is useful when you want to use a paint stroke or mask shape as a motion path.

To use a mask shape or paint stroke as a path:

1 Draw a polygon mask or paint stroke in the shape of the path you want to use.

You do not have to close the mask shape.



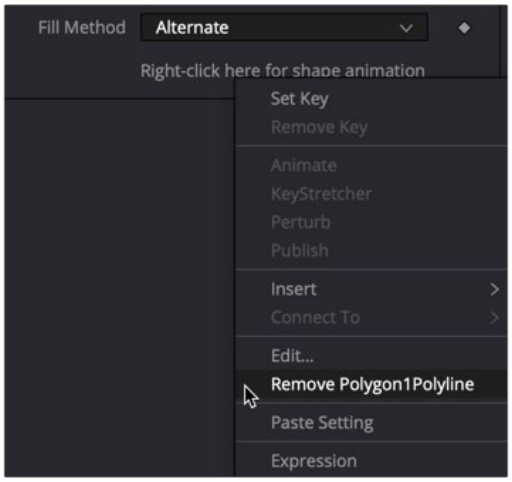
An open Polyline mask shape

2 When done drawing the shape, click the Insert and Modify button in the viewer toolbar to leave the mask shape as an open spline.

3 At this point you can select any of the control points along the spline and press Shift-S to make them smooth or Shift-L to make them linear.

All mask polylines have animation enabled by default, but that is usually not desirable for a motion path. You will need to remove this keyframe animation if you are using a mask shape.

4 At the bottom of the Inspector, right-click on the “Right-click here for shape animation” label and choose Remove Polygon1Polyline.



Right-clicking at the bottom of the Inspector to remove auto-keyframing from a mask shape

5 Right-click at the bottom of the Inspector again and select Publish to give other nodes access to this spline shape. (For a paint stroke, you will need to make the Stroke editable first by clicking the Make Editable button in the Stroke Controls.)

This enables the Modifiers tab with the Published Polyline modifier. This published spline can be used to define the shape of splines in other nodes.

6 Connect a Transform node to the image you want to have follow the path.

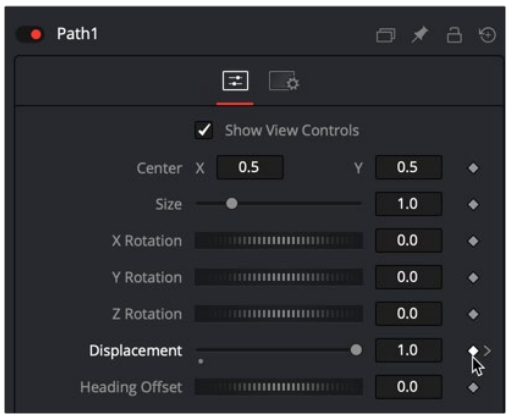
7 Right-click over the Center X/Y parameter in the Inspector and choose Path.

This adds a path modifier into the Modifiers tab.

8 In the Inspector, click the Modifiers tab and double-click the Path1 heading to open its parameters.

The Displacement parameter already has a keyframe on it automatically. You’ll want to remove that so you can set your own.

9 Click the red keyframe button to the right of the Displacement parameter to remove it.



Remove the automatic displacement keyframe prior to creating your own keyframes.

10 At the bottom of the Modifiers tab, right-click on “Right-click here for shape animation” and choose Connect To > Polygon1Polyline.

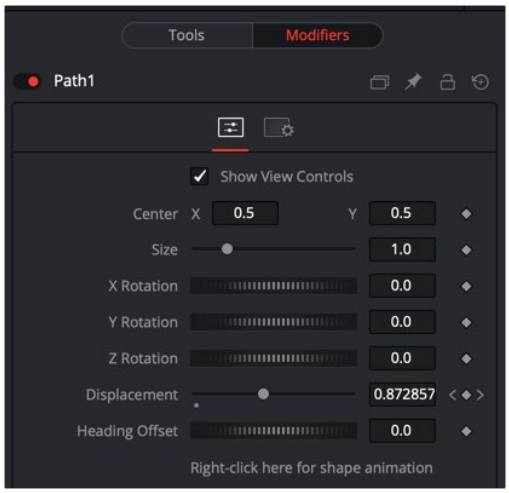
11 To quickly see where your object has gone, drag the Displacement slider back and forth.

12 You may want to use the Size parameter to adjust the size of the overall path.

The Displacement slider is meant to be keyframed for animating the object along the path.

Path Modifier

In terms of functionality, it makes no difference which method you use to generate the path modifier. All the above methods are just different ways to get to the same point. Whichever way you decide to add the path modifier, the Modifiers tab contains controls for the path.

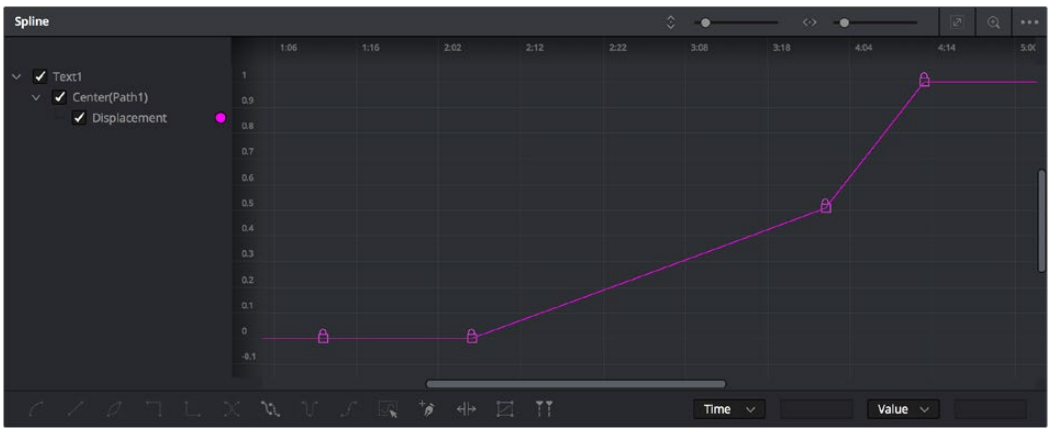


Creating a path adds controls to the Modifiers tab in the Inspector.

You can use the path modifier controls in the Inspector to change the position, size, and rotation of the entire path shape. The Displacement parameter is represented as a spline in the Spline Editor, which determines the object’s position along the path and the Heading Offset is used for the orientation along the path.

Controlling Speed and Orientation along a Path

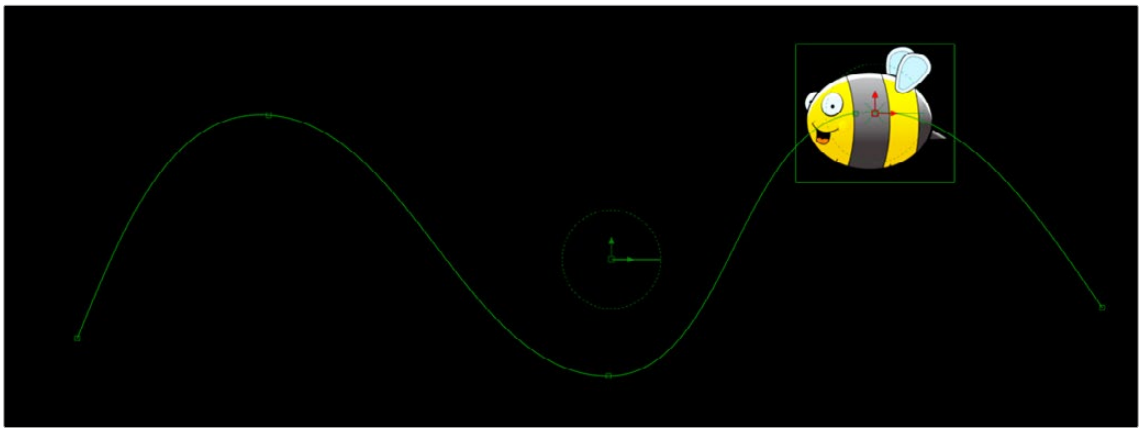
Every Polyline path has an associated Displacement spline in the Spline Editor. The Displacement spline represents acceleration, or the position of the animated object along its path, represented as a value between 0.0 and 1.0.



The Displacement curve of a Poly path represents the acceleration of an object on a path.

Smaller values are closer to the beginning of a path, while larger values are increasingly closer to the end of the path.

For instance, let's say you have a bumblebee that bobs up and down as it moves across the screen. To have the bee accelerate as it moves up and down but slow down as it reaches its peaks and valleys you use the Displacement curve.



A curvy path defined by a spline shape

The curved shape path does not define how fast the bee moves. The speed of the bee at any point along the path is a function of the Displacement parameter. You can control the Displacement parameter either in the Modifiers tab or in the Spline Editor.

To use the Displacement parameter to control the speed of an object on a path, do the following.

1. Position the playhead at the start of the animation.
2. In the Modifiers tab, drag the Displacement parameter to 0.0. This positions the object at the start of the path.
3. Click the Keyframe button to the right of the Displacement parameter.
4. Position the playhead somewhere further into the comp and drag the Displacement parameter until the object is where you want it to be based on the current frame.
5. Continue updating the playhead and the Displacement parameters at key points in the comp until you have reached the end of the path.

After the initial animation is set, you can use the Displacement curve in the Spline Editor to adjust the timing.

To adjust the timing of an object along a path:

1. Open the Spline Editor and enable the Displacement spline in the header.
2. Move control points horizontally closer together to increase the speed between the two points while maintaining the location of the object along the path.

   A longer, flatter curve between two points indicates a slower rate of change.
3. Drag a control point up or down to change its location on the path while maintaining the timing between two points.



A Displacement curve in the Spline Editor

TIP: Holding down the Option key while clicking on the spline path in the viewer will add a new point to the spline path without adding a Displacement keyframe in the Spline Editor. This allows you to refine the shape of the path without changing the speed along the path.

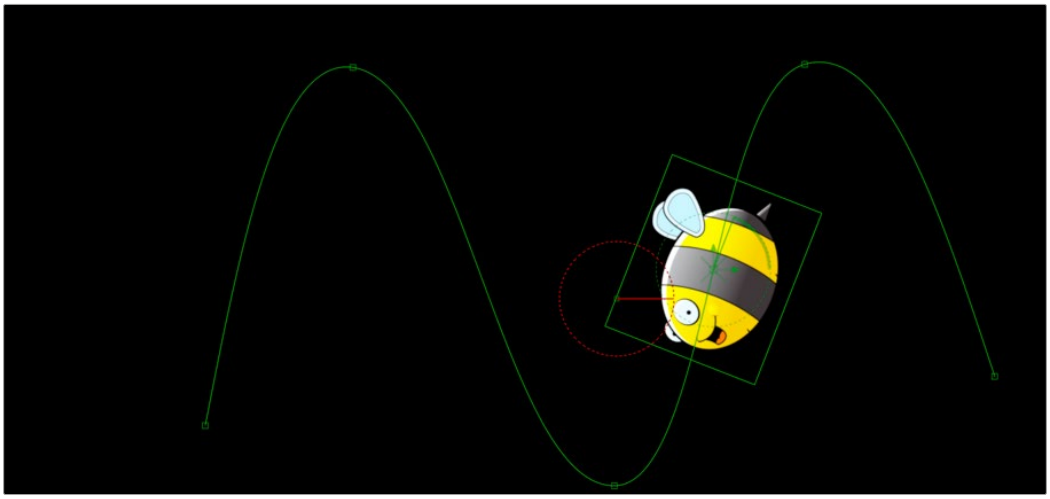
Using a Path Modifier to Adjust Orientation

The Heading parameter is used to adjust the orientation of the object along the path. For instance, if you want the bee in our animation to auto-orient based on the direction of the path, you can connect the bee's angle to the Heading parameter.

To connect the Heading to an object's angle:

1. Right-click over an object's angle parameter in a Transform node.
2. Choose Connect To > Path > Heading.

In our example comp, the bee now auto-orients as the path descends and rises.



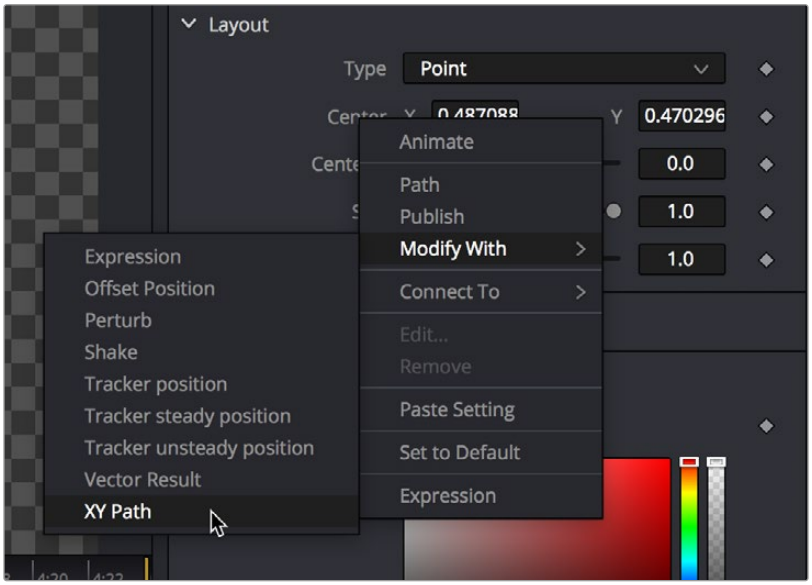
The Transform's angle parameter connected to the path modifier's Heading parameter

XY Path

Unlike a Polyline path, the XY path modifier uses separate splines in the Spline Editor to calculate position along the X-axis and along the Y-axis.

To animate a coordinate control using an XY path modifier:

- Right-click on the center coordinate control in the viewer or the Center X/ Y parameter in the Inspector, and then choose Modify With > XY Path from the contextual menu.



Adding the XY Path modifier to a Center parameter in the Inspector

At first glance, XY paths work like Polyline paths. To create the path once the modifier is applied, position the playhead and drag the onscreen control where you want it. Position the playhead again and move the onscreen control to its new position. The difference is that the control points are only there for spatial positioning. There is no Displacement parameter for controlling temporal positioning.



XY path modifier controls in the Modifiers tab of the Inspector

Instead of dragging in the viewer, you can use the controls in the Modifiers tab to create a motion path, while using the object's original Inspector controls as an offset to this motion path. You can use the XYZ parameters to position the object, the Center X/Y parameters to position the entire path, the Size and Angle to scale and rotate the path, and the Heading Offset control to adjust the orientation.



Using an XY path modifier to animate a piece of text

Using the XY Paths in the Spline Editor

The Spline Editor for the XY path displays the X and Y channel splines. Changes to the path in the viewer or the Inspector will be displayed as keyframes on these splines in the Spline Editor. Unlike a Polyline path, XY paths do not include a Displacement curve. The speed of the object along the path is tied to the path itself and cannot be separated from the timing of the keyframes that define that path.

> **TIP:** XY path and Poly path can be converted between each other from the contextual menu. This gives you the ability to change methods to suit your current needs without having to redo animation.

The advantage of the XY path modifier is that you can explicitly set an XY coordinate at a specific time for more control.



XY Path displays X and Y curves in the Spline Editor but does not include a Displacement control.

Switching Default Paths

If you want to change the default path type to XY path, you can choose Fusion > Preferences > Globals in Fusion Studio or Fusion > Fusion Settings in DaVinci Resolve. In the Settings window, select the Defaults category and choose XY Path for the Point With drop-down menu. The next time you keyframe the Center X/Y parameter or choose Animate from CenterX/Y's contextual menu, an XY path modifier will be used instead of a Polyline path.

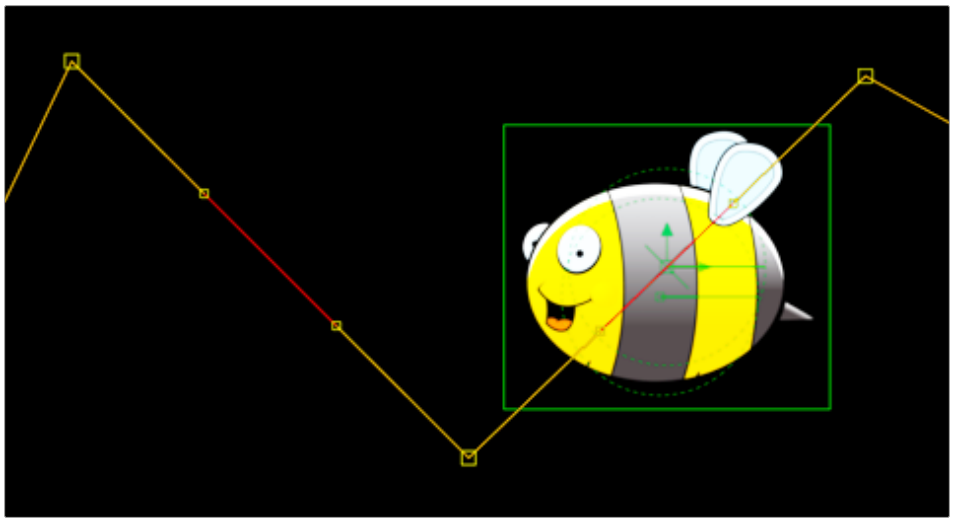
Types of Control Points

The control points along an XY path in the viewer are locked to the control points on the X and Y curve in the Spline Editor. The number of points are identical, and adding a control point in one place adds it to the other. That is not the case with a Polyline path. Polyline paths are composed of locked and unlocked points. Whether a point is locked is determined by how it was added to the Polyline. Locked points on the motion path in the viewer will have an associated point on the Displacement spline in the Spline Editor; unlocked points will not have corresponding points. Each has a distinct behavior, as described below.

Locked Points

Locked points are the motion path equivalents of keyframes. They are created by changing the playhead position and moving the animated control. These points indicate that the animated control must be in the specified position at the specified frame.

The locked points are displayed as larger-sized hollow squares in the viewer. Each locked key has an associated point on the path's Displacement curve in the Spline Editor.

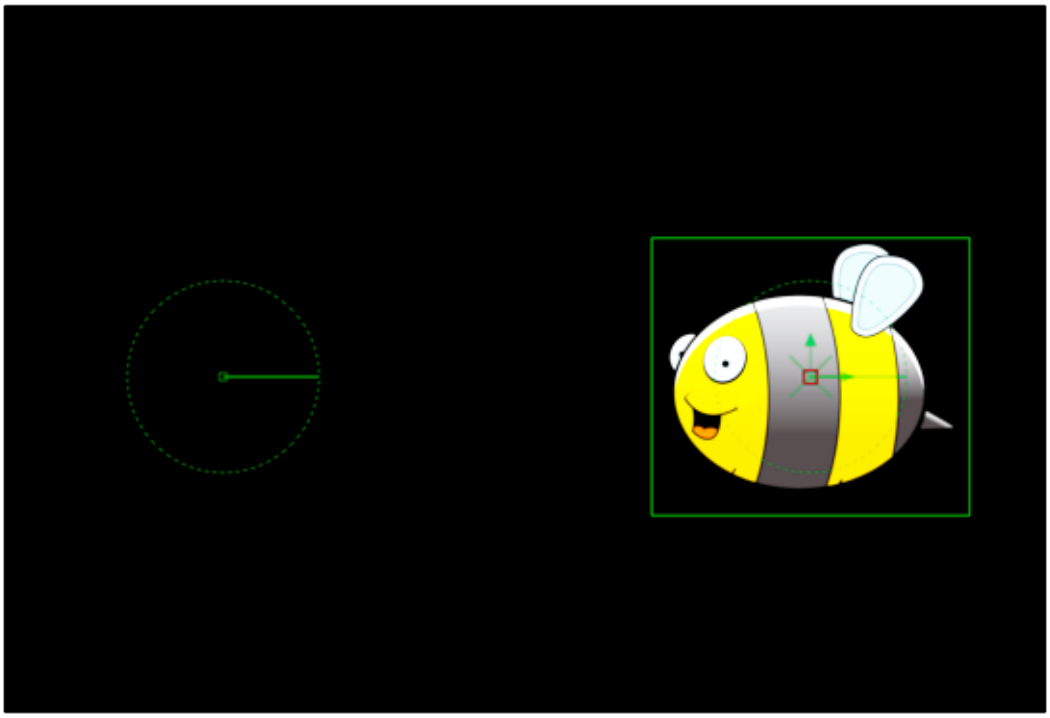


Locked points on a path in the viewer

Deleting a locked point from the motion path will change the overall timing of the motion.

To change the duration of a path using locked points:

1. Connect an object to a Transform node.
2. Position the where you want to start the motion path.



A graphic placed on the right side of the frame

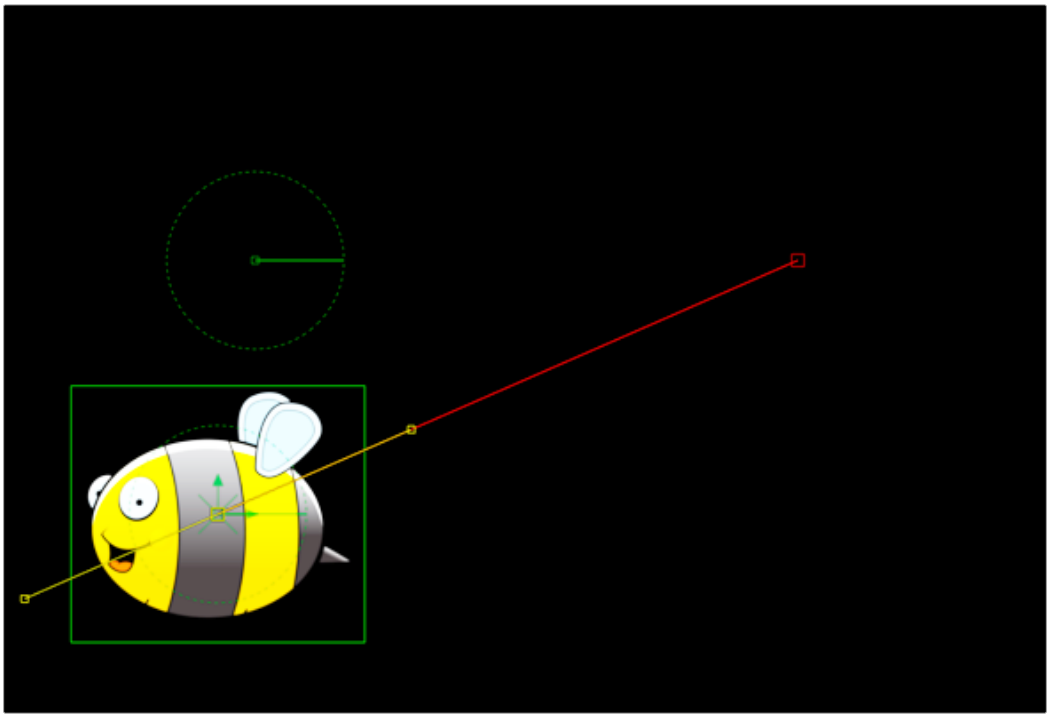
3 Set the playhead at frame 0.

4 Select the Transform node and in the Inspector, click the Keyframe button to the right of the Center X/Y parameter.

This adds the path modifier and creates the first locked point of the path.

5 Position the playhead at frame 45.

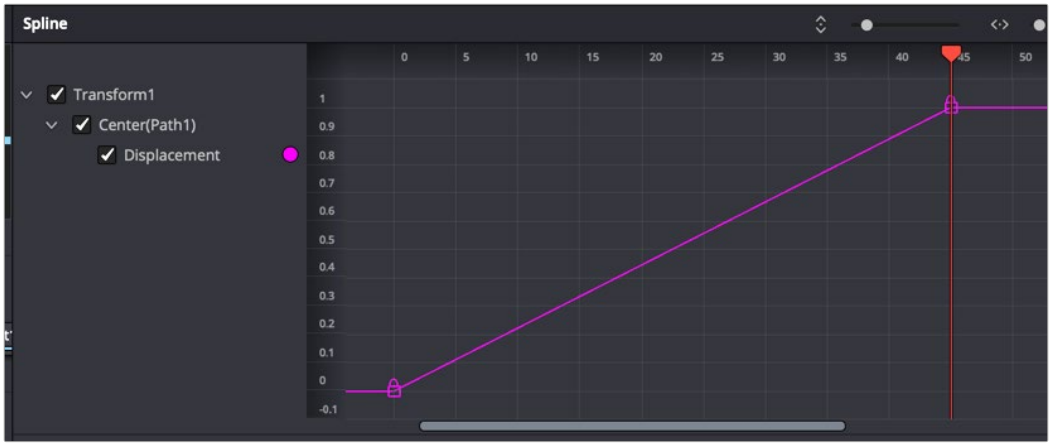
6 Move the object's center to the lower center of the screen.



Moving the playhead and repositioning the bee adds a second locked point

This sets the second locked point.

7 View the Spline Editor and display Path1's Displacement spline.



The path's Displacement spline with locked keyframes

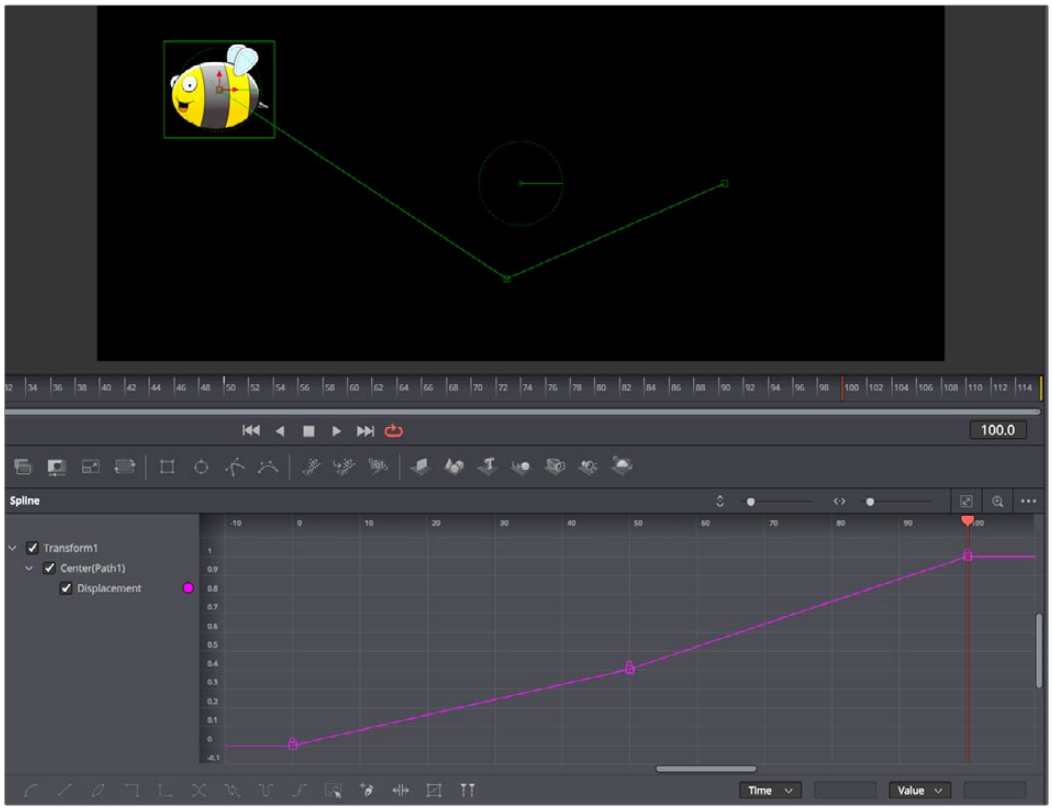
At a value of 0.0, the control will be located at the beginning of the path. When the value of the Displacement spline is 1.0, the control is located at the end of the path.

8 Select the keyframe at frame 45 in the Displacement spline and drag it to frame 50.

The motion path is now 50 frames long, without making any changes to the motion path's shape.

If you try to change this point's value from 1.0 to 0.75, it cannot be done because the point is the last in the animation, so the value must be 1.0 in the Displacement spline.

9 Position the playhead on frame 100 and move the bee center to the upper-left corner of the screen.



Moving locked points changes the duration of a motion path without changing its shape

This will create an additional locked point and set a new ending for the path.

Unlocked Points

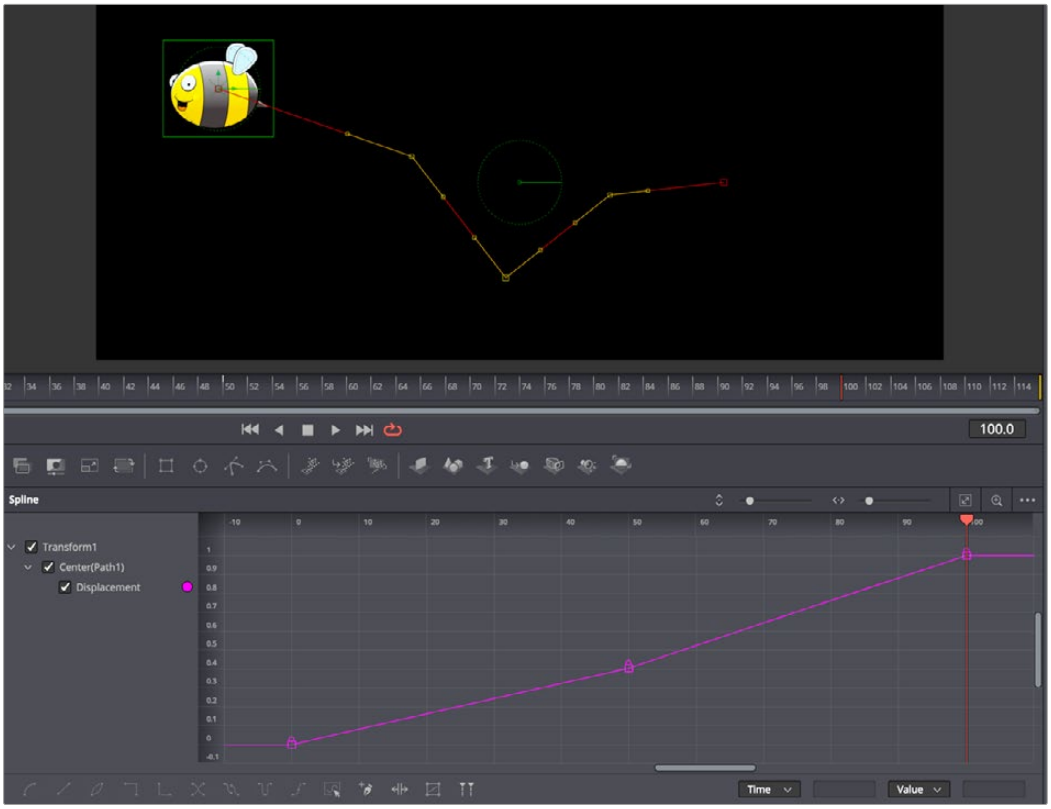
Unlocked points are created when additional points are added to the motion path while in Insert and Modify modes. These points are used to adjust the overall shape of the motion path, without directly affecting the timing of the motion. This means you can add whatever unlocked points you want to reshape a motion path, without that having any effect on the timing of the animation on that path. This makes it vastly easier to fine-tune a motion path spatially after you've perfected the timing of the animation temporally.

Unlocked points *do not* have corresponding points on the path's Displacement spline. They appear in the viewer as smaller, solid square points.

To add unlocked points to a motion path, do the following:

1. Select a motion path spline by using the Tab key to cycle controls until the path is selected.
2. To insert points along a path, click the Insert and Modify button in the toolbar.
3. Click on the path and create two new points: one half-way between the first and the second points, and the other half-way between the second and the third points.

   The two points just added are not present in the motion path's Displacement spline. These are unlocked points, used to shape the motion but unrelated to the timing of the path.

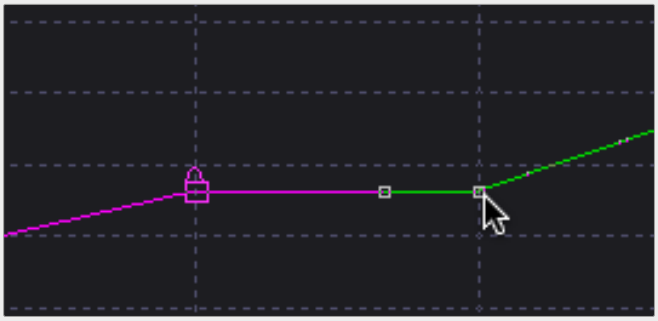


Unlocked points added to the motion path are not displayed on the Displacement spline

You can add unlocked points to the Displacement spline as well. Additional unlocked points in the Spline Editor can be used to make the object's motion pause briefly.

To pause motion along a motion path using an unlocked point:

Select a locked point on the Displacement spline and then hold down the Command key while dragging the point horizontally to another frame. The point is copied as an unlocked point to the new frame.



Create a pause in the motion by copying locked points

Knowing the difference between locked and unlocked points gives you independent control over the spatial and temporal aspects of motion paths.

Locking and Unlocking Points

You can change an unlocked point into a locked point, and vice versa, by selecting the point(s) and choosing the Lock Point option from the contextual menu.

Tips for Manipulating Motion Paths

There are a variety of ways you can create and edit motion paths in the viewer.

Compound Motion Paths Using Path Centers

Every motion path has a defined center represented by a crosshair. Path centers allow paths to be connected to other controls and behave in a hierarchical manner, which is an exceptionally powerful way of creating complex motion by combining relatively simple paths.

A useful example of this technique would be animating the path of a bee in flight. A bee often flies in a constant figure eight pattern while simultaneously moving forward. The easy way of making this happen involves two paths working together.

In the following example, the bee would be connected to a first path in a Transform node, which would be a figure eight of the bee moving in place. This first path's center would then be connected to another path defining the forward motion of the bee through the scene via a second Transform node.

To create a compound motion path:

1. Create a figure-eight motion path by keyframing an object or drawing a polyline mask.

   (If using a polygon mask, you'll need to remove the auto-keyframing and publish the mask prior to the next step.)

2. Add a polyline mask and create a smooth curve spline that travels across the screen.

3 At the bottom of the Inspector, right-click over the “Right-click here for shape animation” label and choose Remove Polygon Polyline to remove the auto-animation behavior.

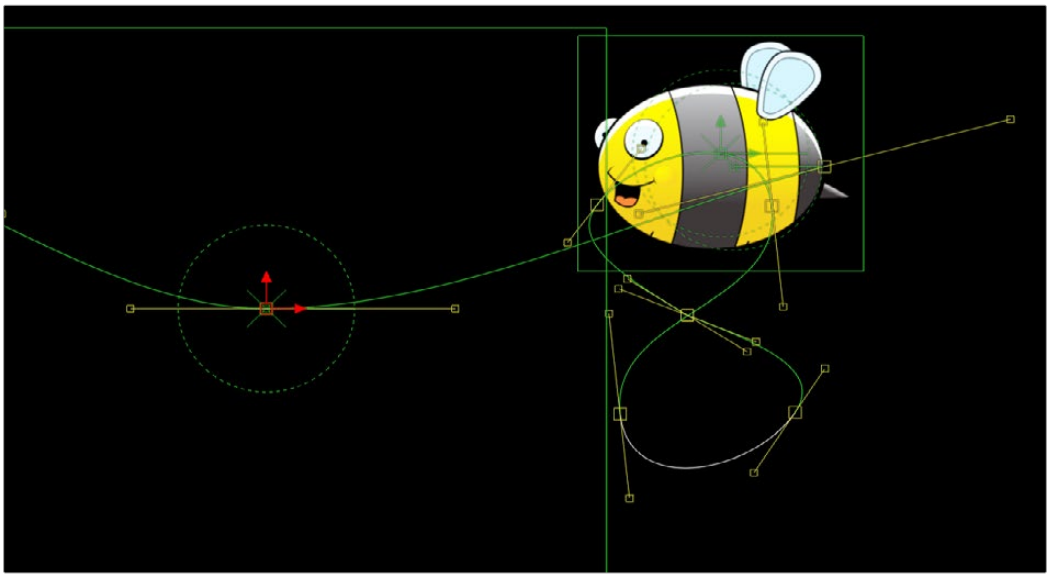
4 Right-click over the label again and choose Publish.

5 Select the object’s Transform node and click the Modifiers tab.

6 Right-click over the Path1 Center X/Y parameter and choose Path.

7 At the bottom of Path2, choose Connect To > Polygon: Polyline.

8 Keyframe the Path2 Displacement to move the object along the second path.



Two motion paths working together

Copying and Pasting Motion Paths

It is possible to copy an entire motion path to the clipboard and then paste it onto another node or path or composition.

Methods of copying and pasting motion paths:

- **To copy a motion path:** In the Inspector’s Modifiers tab, right-click on the path’s control header and choose Copy from the contextual menu.
- **To cut a motion path out of a node:** In the Inspector, right-click on the path’s control header and choose Cut from the contextual menu.
- **To paste the copied path over another path:** In the Inspector, right-click on the path’s control header and choose Paste from the contextual menu.

When pasting a path, the old motion path will be overwritten with the one from the clipboard.

Removing Motion Paths

There are multiple ways to remove or delete a path, and they all involve a right-click contextual menu. Removing a path modifier does not remove the object or the spline shape; it only removes the animation from the object and the modifier from the Modifiers tab in the Inspector.

To remove a path, do one of the following:

- In the Modifiers tab, right-click over the Path1 header and choose Delete from the menu.
- In the Inspector, right-click over the Center XY parameter and choose Remove Path1 from the menu.
- Right-click the center coordinate control in the viewer for the object you're animating, and choose Remove Path1 from the submenu of the NameOfObject; Center submenu.



Removing an entire motion path at once

Recording Motion Paths

You can animate both of the control's spatial and temporal information at the same time using the Record mode. This is useful when both position and speed are crucial to achieve the desired result.

Right-click on the desired path and select Record from the contextual menu. This displays a submenu of available data that may be recorded.

Use the Record Time option in conjunction with the Draw Append mode to create complex motion paths that will recreate the motion precisely as the path is drawn.

The time used to record the animation may not suit the needs of a project precisely. Adjust the path's Displacement spline in the Spline Editor to more correctly match the required motion.

Importing and Exporting Polylines

You can import and export polyline shapes into a common editable ASCII text file or its native format. These methods are used to save a particularly useful or generic mask or path for future use or for use in another application, such as Maya or LightWave. You can also import FXF, SSF, or Nuke shape files.

Native Format

To save a polyline shape in Fusion's native ASCII format, right-click on the header of the Mask node in the Inspector and select Settings > Save As from the contextual menu. Provide a name and path for the saved file and select OK to write a file with the .setting extension. This file will save the shape of a mask or path, as well as any animation splines applied to its points or controls.

To load the saved setting back into Fusion, you first create a new polyline of the same type, and then select Settings > Load from the mask's context menu or drag the .setting file directly into the Node Editor.

If you want to move a polyline from one composition to another, you can also copy the node to the clipboard, open your second composition, and paste it from the clipboard into the new composition.

Chapter 12

Using Modifiers, Expressions, and Custom Controls

Some of the most powerful aspects of Fusion are the different ways it allows you to go beyond the standard tools delivered with the application.

This chapter provides an introduction to a variety of advanced features, including Modifiers, Expressions, and Scripting, which can help you extend the functionality and better integrate Fusion into your studio.

FUSION

Contents

The Contextual Menu for Parameters in the Inspector 324
Using Modifiers 324
Adding the Right Modifier for the Job 324
Adding Modifiers to Individual Parameters 324
Combining Modifiers and Keyframes 325
Publishing a Parameter 325
Connecting Multiple Parameters to One Modifier 326
Adding and Inserting Multiple Modifiers 326
Performing Calculations in Parameter Fields 328
Using SimpleExpressions 328
Pick Whipping to Create an Expression 330
Removing SimpleExpressions 330
Customizing User Controls 331
FusionScript 334

The Contextual Menu for Parameters in the Inspector

Most of the features in this chapter are accessed via a contextual menu that appears when you right-click most parameters in the Inspector. Different contextual menus are available based on where in the Inspector you right-click. Specifically, right-clicking over parameter names or sliders displays a feature-rich contextual menu that can add animation, expression fields, modifiers to extend functionality, as well as publishing and linking capabilities, allowing you to adjust multiple controls simultaneously.

Using Modifiers

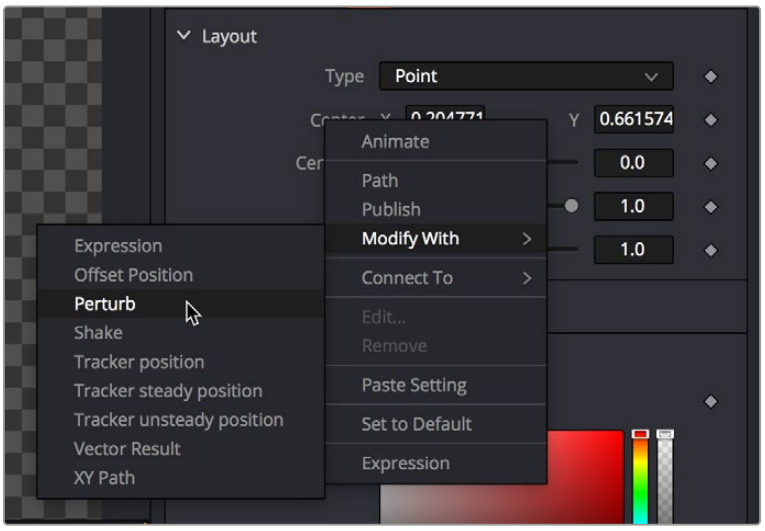
Parameters can be controlled with modifiers, which are extensions to a node's toolset. Many modifiers can automatically create animation that would be difficult to achieve manually. Modifiers can be as simple as keyframe animation or linking the parameters to other nodes, or modifiers can be complex expressions, procedural functions, external data, third-party plugins, or fuses.

Adding the Right Modifier for the Job

Which modifiers are available depends on the type of parameter you're right-clicking over. Numeric values, text, polylines, gradients, and points each have different sets of modifiers that will work with them, so the Modify With menu is filtered based on each parameter.

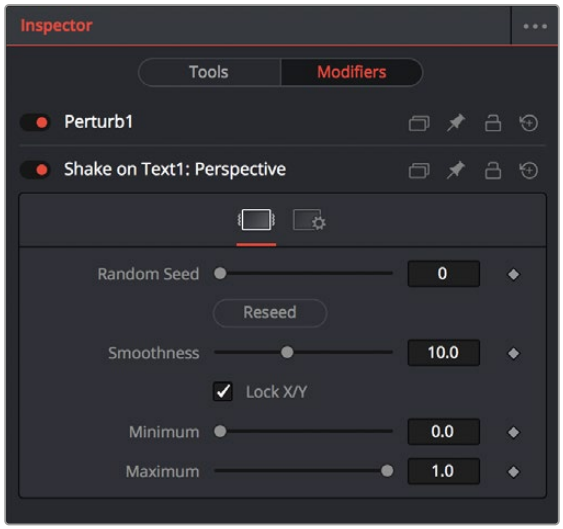
Adding Modifiers to Individual Parameters

You add modifiers to a parameter using the Inspector's contextual menu, or by right-clicking the onscreen control for a parameter in the viewer. Either way, a dynamic list of modifiers that are appropriate for the selected parameters is displayed. For instance, a Perturb modifier can be added to any slider to auto-animate the parameter by randomly wiggling the value. Once the modifier is added, controls are displayed in the Modifiers tab to adjust the speed and amplitude of the random animation.



The Inspector contextual menu is used to add a Perturb modifier to the Center X and Y parameters

A modifier's controls are displayed in the Modifiers tab of the Inspector. When a selected node has a modifier applied, the Modifiers tab will become highlighted as an indication. The tab remains grayed out if no modifier is applied.



The Modifiers tab with two modifiers applied

Modifiers appear with header bars and header controls just like the tools for nodes. A modifier's title bar can also be dragged into a viewer to see its output.

Combining Modifiers and Keyframes

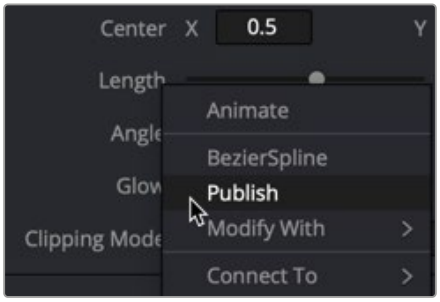
Modifiers that auto-animate parameters like Perturb and Shake can be combined with keyframes to create more natural, organic looking animations. For instance, you can create the general motion path for an element by keyframing the Center X and Y parameters, and then apply a modifier to the same parameters to create a secondary wiggling motion.

To combine a keyframed motion path with a Perturb modifier:

1. Add a Transform to an image like a butterfly or spaceship.
2. Select the Transform node in the Node Editor.
3. In the Inspector, click the Keyframe button to the right of the Center X and Y parameters.
4. Position the butterfly or spaceship image where you want the start of the animation to begin.
5. Continue to move the playhead in the render range and reposition the image until you create a figure-8 motion path.
6. Right-click over the Center X label in the Inspector and choose Modify With > Perturb.
7. Click the Modifiers tab at the top of the Inspector and adjust random, wiggling motion by setting the Strength, Wobble, and Speed parameters while the animation plays.

Publishing a Parameter

The Publish modifier makes the value of a parameter available, so that other parameters can connect to it. This allows you to simultaneously use one slider to adjust other parameters on the same or different nodes. For instance, publishing a motion path allows you to connect multiple objects to the same path.



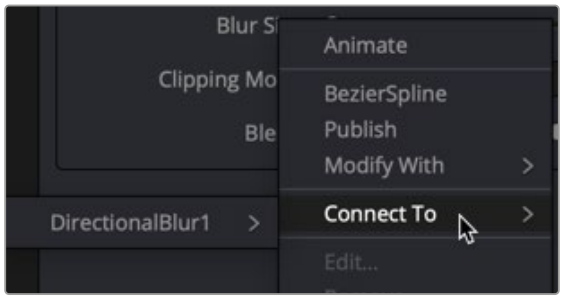
Publish a parameter in order to link other parameters to it

Once a parameter is published, you can right-click another parameter and choose Connect To > [published parameter name] from the contextual menu. The two values are linked, and changing the parameter value of one in the Modifiers tab changes the other.

Using the pick whip between two parameters provides similar linking behavior with more flexibility. Pick whipping between parameters is covered later in this chapter.

Connecting Multiple Parameters to One Modifier

A single modifier or published parameter can be applied to multiple parameters so they all act as one. This is handled through the Connect To submenu. As with modifier assignment, the list is filtered to show only options that are suitable for the parameter you've right-clicked on. When you do this, the connection is bidirectional; editing either parameter will cause the other parameter to change.



The Connect To menu connects modifiers to multiple parameters

Adding and Inserting Multiple Modifiers

Modifiers can be connected to each other and branched, just like nodes in the Node Editor. For example, the Calculation modifier outputs a number, but has two Number parameters, both of which can have modifiers added to them. If you want to insert a modifier between the existing modifier and the modified parameter, use the Insert submenu of the parameter's contextual menu.

Available Modifiers in Fusion:

- **Anim Curves:** Adds an animation curves modifier that allows you to dynamically adjust the timing, scaling, and acceleration of an animation.
- **Bézier Spline:** Adds a Bézier spline to the Spline Editor for animating the selected parameter.
- **B-Spline Spline:** Adds a B-Spline spline to the Spline Editor for animating the selected parameter.
- **Calculation:** Creates an indirect link that includes a mathematical expression between two parameters.

- **CoordTransform Position:** Calculates the current 3D position of a given object even after multiple 3D transforms have repositioned the object through the node tree hierarchy.
- **Cubic Spline:** Adds a Cubic spline to the Spline Editor for animating the selected parameter.
- **Expression:** Allows you to add a variable or a mathematical calculation to a parameter, rather than a straight numeric value. The Expression modifier provides controls in the Modifiers tab, giving you more room and parameters than the SimpleExpression
- **From Image:** This modifier takes color samples of an image along a user-definable line and creates a gradient from those samples.
- **Gradient Color Modifier:** Creates a customized gradient and maps it into a specified time range to animate a value.
- **KeyStretcher:** Used to stretch keyframes in a Fusion Title template when trimming the template in the Edit page or Cut page timeline.
- **MIDI Extractor:** Modifies the value of a parameter using the values stored in a MIDI file.
- **Natural Cubic Spline:** Adds a Natural Cubic spline to the Spline Editor for animating the selected parameter.
- **Offset (Angle, Distance, Position):** The three Offset modifiers are used to create variances, or offsets, between two positional values. For instance, when this modifier is added to a size parameter, you can change the size of an object using the distance between two onscreen controls (position and offset).
- **Path:** Produces two splines to control the animation of of an object: An onscreen motion path (spacial) and a Time spline visible in the Spline Editor (temporal).
- **Perturb:** Generates smoothly varying random animation for a given parameter.
- **Probe:** Auto-animates a parameter by sampling the color or luminosity of a specific pixel or rectangular region of an image.
- **Publish:** The first step in linking two non-animated parameters is to use the Publish modifier to publish a parameter. That allows other parameters to use the Connect To submenu and link to the published parameter.
- **Resolve Parameter:** Allows you to modify the duration of a Fusion transition template from the Edit page Timeline. The Resolve Parameter Modifier is applied to any animated parameter instead of keyframing the transition.
- **Shake:** Similar to Perturb, Shake generates smoothly varying random animation for a given parameter.
- **Track:** Attaches a single point tracker to the selected parameter. The tracker can then track an object onscreen to animate the parameter. This is quicker and more direct than using the normal Tracker node; however, it offers less flexibility since the resulting tracker is only a single point and can only be used for the selected parameter.
- **Vector Result:** Similar to the Offset modifier, Vector Result is used to offset position parameters using origin, distance, and angle controls to create a vector. This vector can then be used to adjust any other parameter.
- **XY Path:** Produces an X and Y spline in the Spline Editor to animate the position of an object..

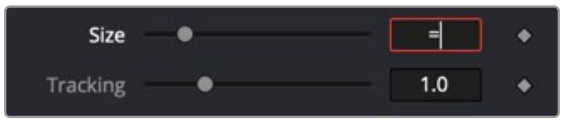
For more information on all modifiers available in Fusion, *see Chapter 64, "Modifiers,"* in the Fusion Reference Manual.

Performing Calculations in Parameter Fields

You can enter simple mathematical equations directly in a number field to calculate a desired value. For example, typing 2.0 + 4.0 in most number fields will result in a value of 6.0. This can be helpful when you want a parameter to be the sum of two other parameters or a fraction of the screen resolution.

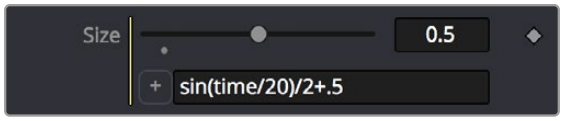
Using SimpleExpressions

Simple Expressions are a special type of script that can be placed alongside the parameter it is controlling. These are useful for setting simple calculations, building unidirectional parameter connections, or a combination of both. You add a SimpleExpression by entering an equals sign directly in the number field of the parameter and then pressing Return.



Entering an equals sign opens the SimpleExpression field with Pick Whip control

An empty field will appear below the parameter, and a yellow indicator will appear to the left. The current value of the parameter will be entered into the number field. Using Simple Expressions, you can enter a mathematical formula that drives the value of a parameter or even links two different parameters. This helps when you want to create an animation that is too difficult or impossible to set up with keyframing. For instance, to create a pulsating object, you can use the sine and time functions on a Size parameter. Dividing the time function can slow down the pulsing while multiplying it can increase the rate.



A SimpleExpression using the sine and time functions

Inside the SimpleExpression text box, you can enter one-line scripts in Lua with some Fusion-specific shorthand. Some examples of Simple Expressions and their syntax include:

Expression	Description
`time`	This returns the current frame number.
`Merge1.Blend`	This returns the value of another input, Blend, from another node, Merge1.
`Merge1:GetValue("Blend", time-5)`	This returns the value from another input, but sampled at a different frame, in this case five frames before the current one.
`sin(time/20)/2+.5`	This returns a sine wave between 0 and 1.

Expression	Description
`iif(Merge1.Blend == 0, 0, 1)`	This returns 0 if the Blend value is 0, and returns 1 if it is not. The iff() function is a shorthand conditional statement, if-then-else.
`iif(Input.Metadata.ColorSpaceID == "sRGB", 0, 1)`	This returns 0 if the image connected to the current node's Input is tagged with the sRGB colorspace. When no other node name is supplied, the expression assumes the Input is coming from the current node. It is equivalent to self.Input. The Input in most, but not all, Fusion nodes is the main image input shown in the Node Editor as an orange triangle. Images have members that you can read, such as Depth, Width, Metadata, and so on.
`Point(Text1.Center.X, Text1.Center.Y-.1)`	Unlike the previous examples, this returns a Point, not a Number. Point inputs use two members, X and Y. In this example, the Point returned is 1/10 of the image height below the Text1's Center. This can be useful for making unidirectional parameter links, like offsetting one Text from another.
`Text1.Center - Point(0,.1)`	This is similar to the previous expression. This SimpleExpression returns Text instead of a Number or Point.
`Text("Colorspace: "..(Merge1.Background.Metadata.ColorSpaceID)`	The string inside the quotes is concatenated with the metadata string, perhaps returning: Colorspace: sRGB
`Text("Rendered "..os.date("%b %d, %Y").. " at "..os.date("%H:%M").."\n on the computer "..os.getenv("COMPUTERNAME").. " running "..os. getenv("OS").."\n from the comp "..ToUNC(comp.Filename))`	This returns a much larger Text, perhaps something like: Rendered Nov 12, 2019 at 15:43 on the computer Rn309 running Windows_NT from the comp \\ SRVR\Proj\Am109\SlateGenerator_A01.comp
`os.date("%H:%M")`	The OS library can pull various information about the computer. In the previous example, os.date gets the date and time in hours:minutes.
`"..os.getenv("COMPUTERNAME").. " running "..os.`	Any environment variable can be read by os.getenv, in this case the computer name and the operating system.
`"\n from the comp "..ToUNC(comp.Filename)`	To get a new line in the Text, \n is used. Various attributes from the comp can be accessed with the comp variable, like the filename, expressed as a UNC path.

TIP: When working with long SimpleExpressions, it may be helpful to drag the Inspector panel out to make it wider or to copy/paste from a text editor or the Console.

After setting an expression that generates animation, you can open the Spline Editor to view the values plotted out over time. This is a good way to check how your SimpleExpression evaluates over time.



A sine wave in the Spline Editor, generated by the expression used for Text1: Size

For more information about writing Simple Expressions, see the Fusion Studio Scripting Guide, and the official Lua documentation.

Pick Whipping to Create an Expression

With a SimpleExpression field open, a + button is displayed on the left side. Dragging the + button onto another control, or "pick whipping," links the two parameters, similar to the Connect To menu. However, unlike a Connect To parameter link, the pick whip allows you to modify the connection by further editing the expression.



Pick whipping to connect one parameter to another quickly

SimpleExpressions can also be created and edited within the Spline Editor. Right-click on the parameter in the Spline Editor and select Set SimpleExpression from the contextual menu. The SimpleExpression will be plotted in the Spline Editor, allowing you to see the result over time.

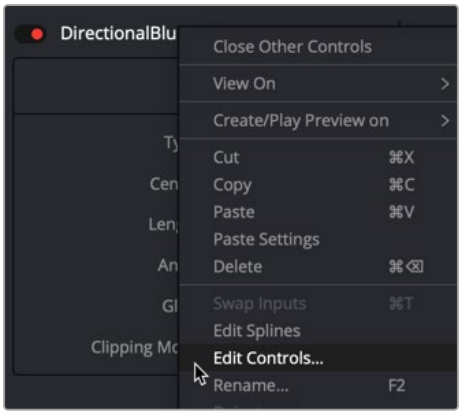
Removing SimpleExpressions

To remove a SimpleExpression, right-click the name of the parameter, and choose Remove Expression from the contextual menu.

Customizing User Controls

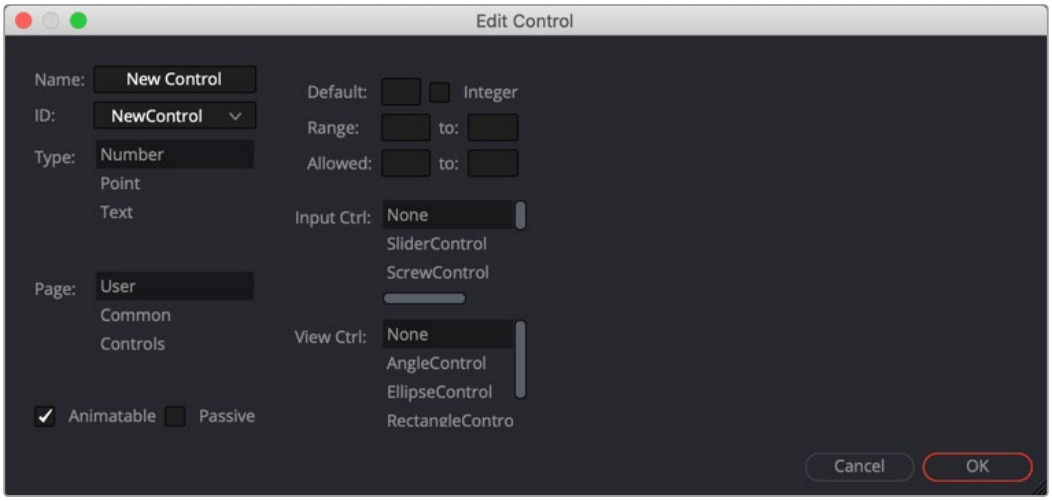
Each tool's parameters are organized in a logical order in the Inspector. The most used controls are closer to the top, and the more subtle refinement controls are lower in the list. Sometimes, though, you may want to add, hide, or change the controls. You often need to do this for SimpleExpressions and macros, but you may also do this for usability and aesthetic reasons for favorites and presets.

Custom controls can be added or edited via the Edit Control dialog, which you access by right-clicking over the node's name in the Inspector and choosing Edit Controls from the menu.



Choose Edit Controls to create and customize parameters

In the Edit Control dialog, you use the ID menu to select an existing parameter or create a new one. You can name the control and define whether it is a text field, number field, or a point using the Type attributes list.

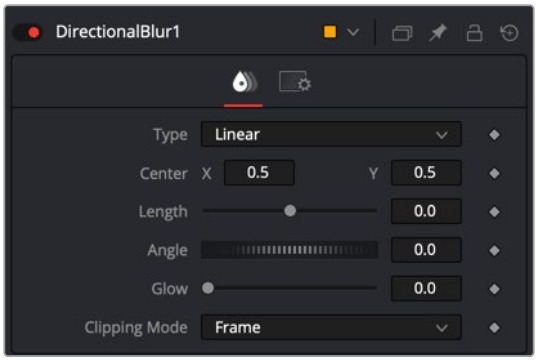


Edit Control dialog

You use the page list to assign the new control to one of the tabs in the Inspector. There are also settings to determine the defaults and ranges, and whether it has an onscreen preview control. The Input Ctrl box contains settings specific to the selected Type, and the View Ctrl attributes box contains a list of onscreen preview controls to be displayed, if any.

All changes made using the Edit Controls dialog get stored in the current tool instance, so they can be copied/ pasted to other nodes in the comp. However, to keep these changes for other comps, you must save the node settings, and add them to the Bins in Fusion Studio or to your favorites.

As an example, we'll customize the controls for a DirectionalBlur:



The Inspector for a default direction blur

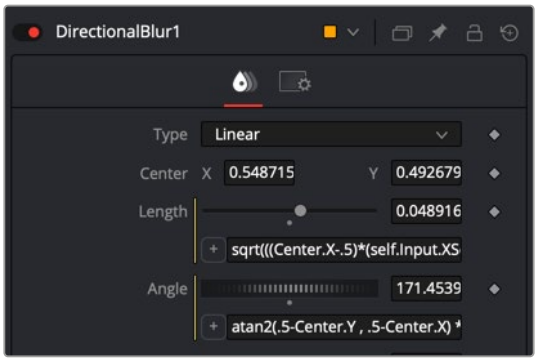
Let's say we wanted a more interactive way of controlling a linear blur in the viewer, rather than using the Length and Angle sliders in the Inspector. Using a SimpleExpression, we'll control the length and angle parameters with the Center parameter's onscreen control in the viewer. The SimpleExpression would look something like this:

For Length:

```
sqrt(((Center.X-.5)*(self.Input.XScale))^2+((Center.Y-.5)*(self.Input.YScale)*(self.Input.Height/self.Input.Width))^2)
```

For Angle:

```
atan2(.5-Center.Y , .5-Center.X) * 180 / pi
```

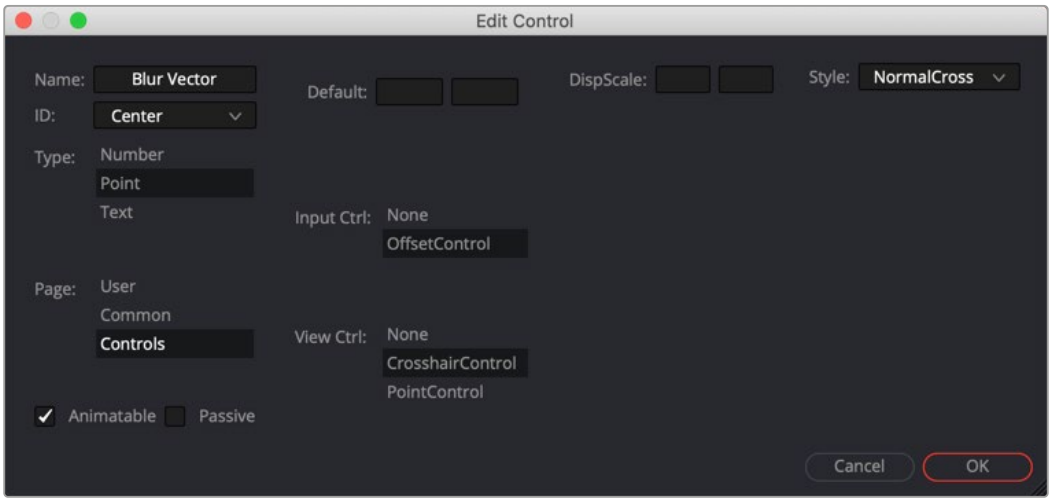


DirectionalBlur controlled by the Center's position.

This admittedly somewhat advanced function does the job fine. Dragging the onscreen control adjusts the angle and length for the directional blur. However, now the names of the parameters are confusing. The Center parameter doesn't function as the center anymore. It is the direction and length of the blur. It should be named "Blur Vector" instead. You no longer need to edit the Length and Angle controls, so they should be hidden away, and since this is only for a linear blur, we don't need the Type menu to include Radial or Zoom. We only need to choose between Linear and Centered. These changes can easily be made in the Edit Controls dialog.

To change the Inspector parameters for the example above:

1. In the Edit Control dialog, select the Center from the ID list.
2. A dialog will appear asking if you would like to Replace, Hide, or Change ID. We'll choose Replace.
3. Change the Name to Blur Vector.
4. Set the Type to Point.
5. Select Controls in the Page list. (Controls is the first tab in the Inspector, normally.)
6. Click OK.



DirectionalBlur Center parameter name changed to Blur Vector.

The new Blur Vector parameter now appears in the Inspector. The internal ID of the control is still Center, so our SimpleExpressions did not change.

To hide the Length and Angle parameters in the Inspector:

1. In the Edit Control dialog, select the Length from the ID list.
2. Select Controls from the Page list.
3. In the input Ctrl list, select Node.
4. Click OK.
5. In the Edit Control dialog, select the Angle from the ID list.
6. In the input Ctrl list, select Node.
7. Click OK.

Finally, to remove Radial and Zoom options from the Type menu:

1. In the Edit Control dialog, select the Type from the ID list.
2. Select Controls from the Page list.
3. Select Radial from the Items list and click Del to remove it.
4. Select Zoom from the Items list and click Del to remove it.
5. Click OK.

The Type menu now includes only two options.

If you want to replace the Type menu with a new checkbox control, you can do that by creating a new control and a very short expression.

To create a new control:

1. In the Edit Control dialog, enter Center Blur in the Name field.
2. Select the New Control from the ID list.
3. Set the Type to Number, and set the Page to Controls.
4. Set the Input Ctrl to CheckboxControl.
5. Click OK.

To make this new checkbox affect the original Type menu, you'll need to add a SimpleExpression to the Type:

```
iif(TypeNew==0, 0, 2)
```

The "iif" operator is known as a conditional expression in Lua script. It evaluates something based on a condition being true or false.

FusionScript

Scripting is an essential means of increasing productivity. Scripts can create new capabilities or automate repetitive tasks, especially those specific to your projects and workflows. Inside Fusion, scripts can rearrange nodes in a comp, manage caches, and generate multiple output files for delivery. They can connect Fusion with other apps to log artist time, send emails, or update webpages and databases.

FusionScript is the blanket term for the scripting environment in Fusion. It includes support for Lua as well as Python 2 and 3 for some contexts. FusionScript also includes libraries to make certain common tasks easier to do with Lua and Python within Fusion.

You can run interactive scripts in various situations. Common scripts include:

- Utility Scripts, using the Fusion application context, are found under the File > Scripts menu.
- Comp Scripts, using the composition context, are found under the Script menu or entered into the Console.
- Tool Scripts, using the tool context, are found in the Tool's context menu > Scripts.

Other script types are available as well, such as Startup Scripts, Scriptlibs, Bin Scripts, Event Suites, Hotkey Scripts, Intool Scripts, and SimpleExpressions. Fusion Studio allows external and command-line scripting as well and network rendering Job and Render node scripting.

FusionScript also forms the basis for Fuses and ViewShaders, which are special scripting-based plugins for tools and viewers that can be used in both Fusion and Fusion Studio.

For more information about scripting, see the Fusion Scripting Documentation, accessible from the Documentation submenu of the Help menu.

Chapter 13

Bins

This chapter covers the bin system in Fusion Studio. Bins allow for the storage and organization of clips, compositions, tool settings, and macros, similar to the Media Pool and Effects Library in DaVinci Resolve. It includes a built-in Studio Player for creating a playlist of multiple shots and their versions.

Bins can be used in a server configuration for organizing shots and collaborating with other team members across the studio.

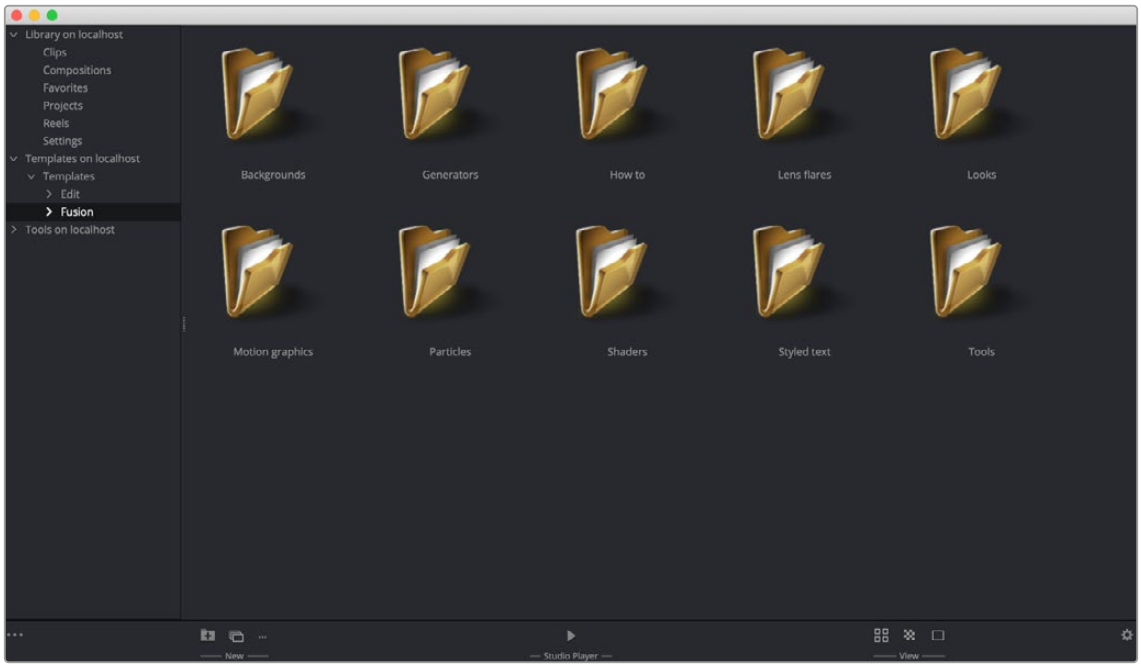
Contents

Bins Overview 336
Bins Interface 336
Viewing and Sorting Bins 337
Organizing Bins 339
Adding and Using Content 340
File Type Details 341
Using Content from Bins 342
Jog and Shuttle 342
Stamp Files 343
Using the Studio Player 343
Playing a Single Clip 344
Creating a Reel 345
Connecting Bins Over a Network 352
Adding a Remote Bin Entry 353
Accessing Remote Bins 353
Permissions 353
Studio Player and Bin Server 354

Bins Overview

Bins are folders that provide an easy way of accessing commonly used tools, settings, macros, compositions, and footage. They can keep all your custom content and resources close at hand, so you can use them without searching through your hard drives. Bins can also be shared over a network to improve a collaborative workflow with other Fusion artists.

Bins are only available in Fusion Studio.



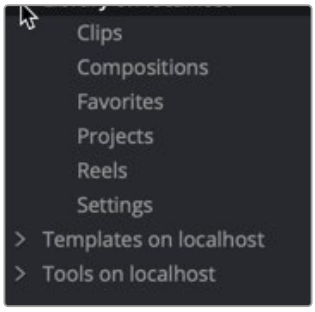
The Bins window

Bins Interface

The Bins window is separated into two panels. The sidebar on the left is a list of the bins, while the Content panel on the right displays the selected bin's contents.

To open the Bins window, do one of the following:

- Choose File > Bins from the menu bar.
- Press Command-B.

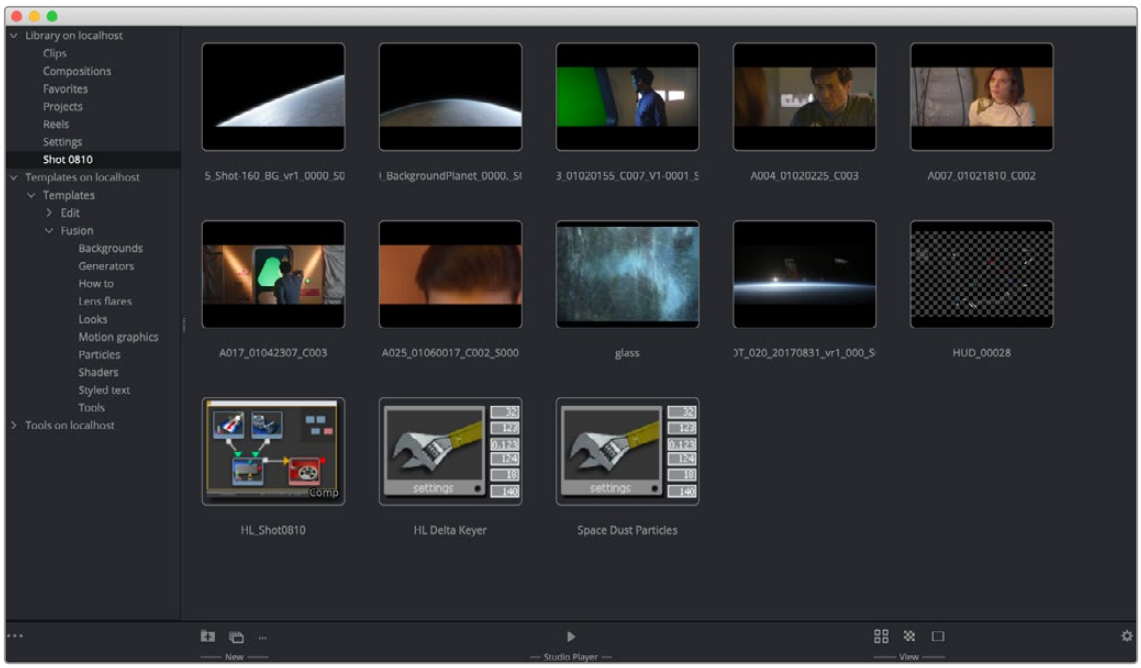


The Bins sidebar

TIP: When opening the Bins window on macOS, the window may open behind the current Fusion Studio window. Check in the dock for the Bins window icon or move the Fusion Studio window to locate the Bins window if you are working on a single monitor.

The sidebar organizes content into bins, or folders, using a hierarchical list view. These folders can be organized however they suit your workflow, but standard folders are provided for Clips, Compositions, Favorites, Projects, Reels, Settings, Templates, and Tools. The Tools category is a duplicate of all the tools found in the Effects library. The Tools bin is a parent folder, and parent folders contain subfolders that hold the content. For instance, Blurs is a subfolder of the Tools parent folder. Parent folders can be identified by the disclosure arrow to the left of their name.

When you select a folder from the sidebar, the contents of the folder are displayed in the Contents panel as thumbnail icons.



The Bins icon view

Viewing and Sorting Bins

A contextual menu is used to access most of a bin's features. You show the contextual menu by right-clicking in an empty area in the Contents panel. Right-clicking on an item will show the same menu with additional options for renaming, playing, or deleting the item.

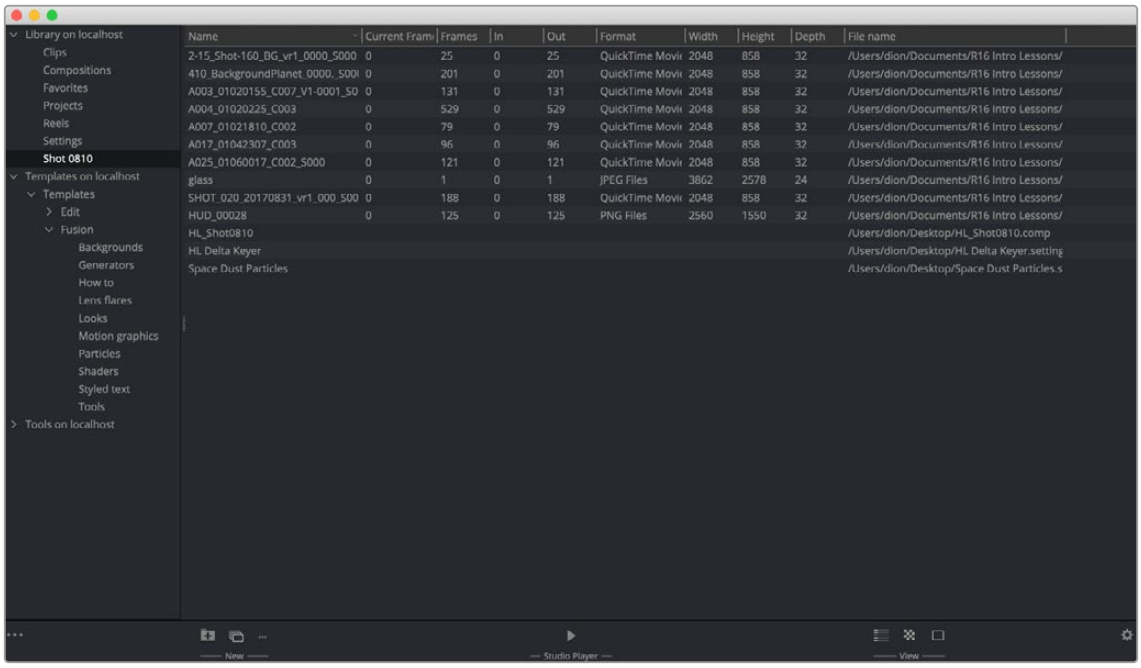
Icon or List View

One use of the contextual menu is to switch between viewing the contents as thumbnail icons or as a list.

To view a bin's contents in List view:

1. Right-click in an empty area of the Contents panel.
2. From the contextual menu, choose View > Details.

Each bin in the sidebar can be set to List view or Icon view independently of each other. So while you may have some bins you want to see as a list, others may be easier to view as icons.



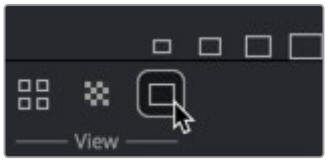
The Bins List view

Sort Order in List View

Clicking on the heading of a column in List view will sort the list in ascending order, and clicking it again will reverse the sort order.

Icon Size in Icon View

The icons can be adjusted to small, medium, large, or huge by clicking the Size button in the bottom toolbar or right-clicking in an empty area of the Contents panel to bring up the contextual menu and choosing a size from the Icon Size submenu.



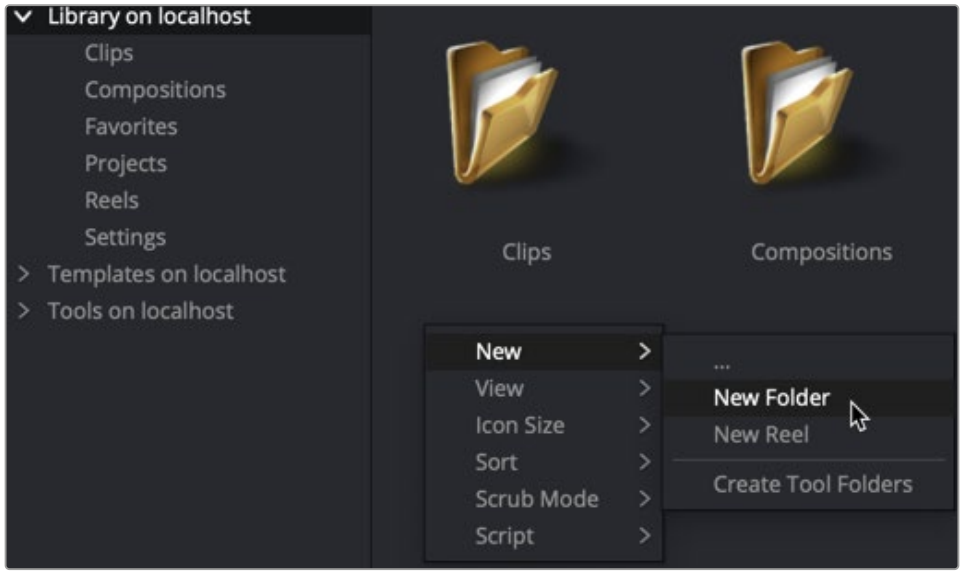
Use the Size button to select the icon size in the bin.

Organizing Bins

Once you begin adding your own categories and content, you can have hundreds of items that need to be organized. To keep all of your elements accessible, you'll want to use some basic organization, just like keeping files and documents organized on your computer.

To create a new folder in the sidebar:

1. In the sidebar, select the parent folder under which the new folder will be listed.
2. Right-click in an empty area of the Contents panel.
3. From the contextual menu, choose New > New Folder.
4. Enter a name for the new folder, and then click OK on the dialog.



The New Folder menu

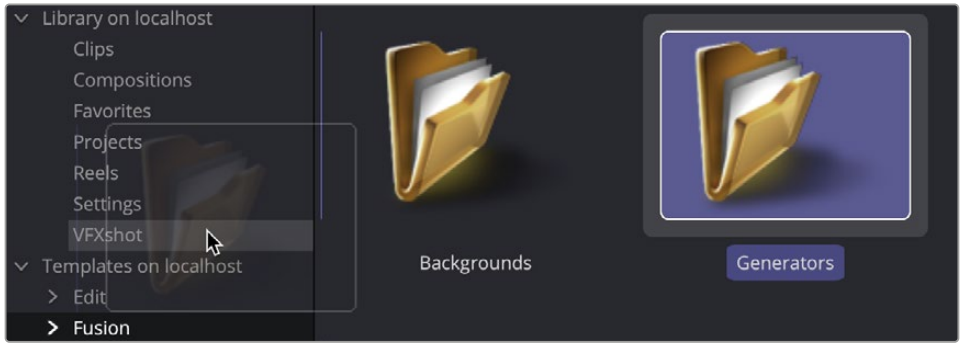
You can also click the New Folder icon in the toolbar.

To rename a bin folder:

1. Right-click on the folder icon in the Contents panel.
2. Choose Rename from the contextual menu or press F2 on the keyboard.

To move a folder into or out of a parent folder:

1. Select the parent folder that contains the folder you want to move.
2. In the Contents panel, drag the folder into the sidebar where you want it moved.



Drag a folder from the Contents panel to move it in the sidebar.

When you drag a folder onto another folder in the sidebar, you create a hierarchical subfolder. Dragging it to the Library parent folder at the top of the sidebar will add it to the top level of the Bins window.

To remove a folder from the bins:

1. Select the folder in the Contents panel.
2. Press Command-Delete (Backspace on Windows and Linux).

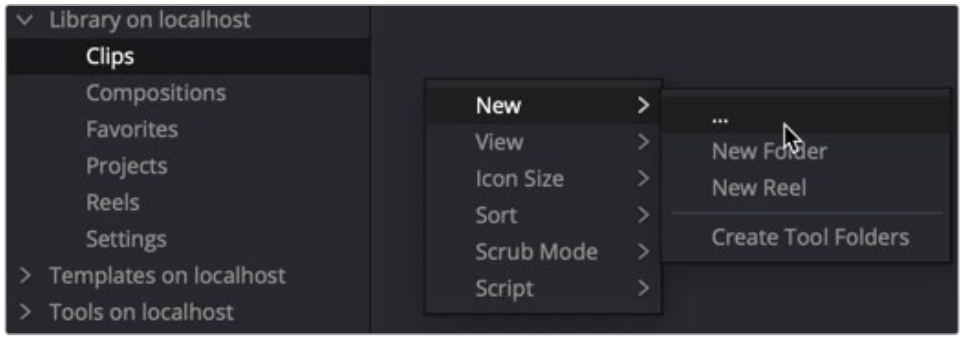
> **TIP:** You cannot undo removing a folder from the Bins window.

Adding and Using Content

You can add and use different types of content with bins. Fusion Studio compositions, tools, saved tool settings, macros, tool groups, and any file format that is supported in Fusion can be added to bins and then used in compositions at a later time.

To add an item to a bin:

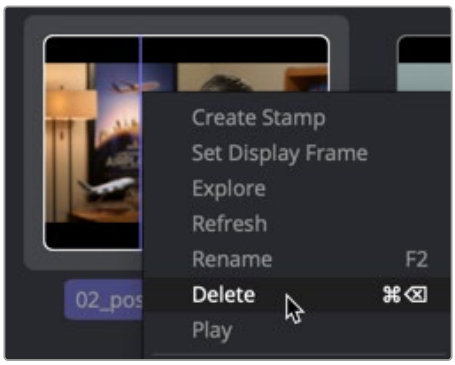
1. Select a bin in the sidebar where you want to add the content.
2. Right-click in the Contents panel.
3. Choose New > ... from the contextual menu
4. Select the media, settings, macros, or comps in the file browser, and then click Open.



Add Item in the contextual menu

To remove an item from the bins:

1. Select the folder in the Contents panel.
2. Right-click on the item and choose Delete from the contextual menu or press Command-Delete.



The Delete function in the contextual menu

TIP: Unsupported files like PDFs can also be saved in the bin, and the application that supports it will launch when you double-click the file. It's a useful location for scripts and notes.

If you have an operating system file browser window open, you can drag files directly into a bin as well. When adding an item to the bins, the files are not copied. A link is created between the content and the bins, but the files remain in their original location.

File Type Details

Some types of content have additional methods of being added to bins. Some additional methods are because of the file type and some are because of where they are located.

Projects and Media

In addition to using the New... contextual menu as explained earlier, Fusion Studio project files with the extension ".comp" and media can also be added to bins by dragging them to the Contents panel from a file browser.

Tool Settings

If you want to add a node with custom settings to a bin, first save the node as a setting by right-clicking over it in the Node Editor and choosing Settings > Save As. Once the setting is saved, you can add it to a bin by dragging it from a File Browser into the Bins window.

Image Sequences and Stills

Image sequences are automatically identified on disk and loaded as clips rather than stills, so it is not necessary to select more than one frame from an image sequence when dragging it into a bin.

To ignore the image sequence and import only a single frame, hold Shift when you drag the frame into a bin. This can be useful when trying to import a single still image from a series of still shots with a DSLR. The numbers maybe sequential, but you just need one still image from the series.

Using Content from Bins

Once you have content in your bins, you'll want to add them to a composition. In general, you can either drag the content directly into the Node Editor or double-click it to add it; however, each type of content behaves a little differently when added to the Node Editor.

Media

Dragging media into the Node Editor from a bin window creates a new Loader that points to the media on disk. Still files or photos are automatically set to loop.

Compositions

To add a composition, you must right-click and choose Open. Dragging a comp item onto an open composition will have no effect. When a composition is added, it is opened in a new window. It is not added to the existing composition.

Tools, Tool Settings, and Macros

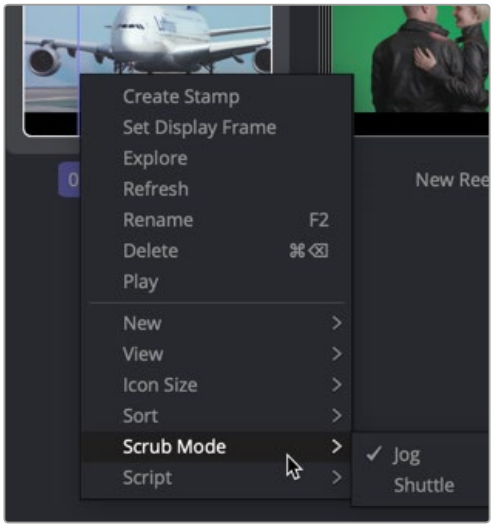
When you add tools from the Tools categories to a composition, the methods you use and results you get are similar to adding tools using the toolbar buttons or the Effects Library. Dragging a tool allows you to place it anywhere in the Node Editor, unconnected, or, if you drag it over a connection line, inserted between two existing tools. Double-clicking a tool in the bin will insert it after the Active tool in the Node Editor. Dragging a tool from a bin into a viewer will insert that tool after the currently viewed tool.

Settings and macros work a bit differently than tools. They can only be added to the Node Editor by dragging and droping. Dragging a setting or macro allows you to place it in the Node Editor, unconnected, or, if you drag it over a connection line, inserted between two existing tools.

Jog and Shuttle

You can scrub clips in Icon view in the bin using one of two modes. Jog mode is the default mode. It moves the clip forward and backward as long as you are dragging the mouse. Once the mouse stops, the clip pauses.

You can choose Shuttle mode by right-clicking over the clip's thumbnail in the bin and choosing Scrub Mode > Shuttle.



Jog and Shuttle contextual menu

Shuttle mode begins playing the clip forward or backward once you press the right mouse button and drag either left or right. The clip continues to play until the mouse button is released or you reach the end of the clip.

Stamp Files

Stamp files are low-resolution, local proxies of clips, which are used for playback of clips stored on a network server, or for very large clips.

To create a stamp file for a clip:

1. Right-click on the clip in the bin.
2. Choose Create Stamp from the contextual menu.

The Status bar at the top of the Bins window shows the progress of the stamp creation. Since the stamp creation is a background process, you can queue other stamps and can continue working with the composition.

Using the Studio Player

The Studio Player is a timeline-based playback interface built into the Bins window. It allows you to play and organize versions of compositions, make notes, and collaborate on shots and projects. The resolution-independent player uses any format that Fusion Studio can ingest, like EXR, ProRes, BMD RAW, QuickTime, and others. Clips use a RAM cache for playback, so even large formats can loop playback as long as there is available memory. Clips can contain audio and are output to a video monitor using Blackmagic Design Decklink or UltraStudio devices, allowing you to screen dailies and review shots. You can add annotation notes to the project, and setting up a Bin Server allows multiple artists to access the Studio Player.

Here are some Studio Player highlights:

- Clips can be played as single events, looped, or using a ping-pong playback, with a definable loop range.
- Clip metadata can be viewed, with live update during scrub/play.
- The Timeline interface can create a playlist for reviewing multiple clips.
- Per-shot color adjustment controls allow for consistent display of clips from different formats.
- Annotation notes can be typed on each clip and version, as well as the entire project.
- Audio Scratch track can be enabled for each clip during playback.
- Clip versions are stored in the same project to allow for quick access to previous work and for comparison of progress.
- Guide overlays can be customized to show monitor/title safety and show crops to various output formats.
- Blackmagic Design UltraStudio and DeckLink playback devices are supported for reviewing clips on video monitors and projectors.
- The fully collaborative workflow automatically synchronizes reel changes, annotations, and color adjustments across multiple workstations, allowing multiple artists or supervisors to access the same projects simultaneously.

- Remote sync allows multiple Studio Players to follow the master. Actions performed on the master, such as playback and scrubbing, will also be executed on the slaves, allowing the reel to be reviewed across multiple workstations or sites.
- Studios can automate tasks using the Fusion scripting engine to control the Studio Player.

Playing a Single Clip

Clips created from image sequences, MOV files, and some RAW formats can be previewed using the Studio Player without having to first add the clip to a node tree.

To play a clip in the Studio Player, do one of the following:

- Double-click a clip in a bin to open the Studio Player.
- Select the clip and click the Play button at the bottom of the bin window.



The Studio Player includes a large viewer, a Timeline, and a toolbar along the bottom.

Once you have the clip open in the Studio Player, you can click the Play button in the toolbar at the bottom of the window.

Scrubbing the Timeline

You can quickly preview the clip by scrubbing through it rather than playing it.

To scrub a clip in Studio Player:

1. In the Timeline, drag the playhead to the area you want to scrub over.
2. Use the Left and Right Arrow keys to move one frame forward or backward.

Closing the Studio Player

After you have finished previewing in the Studio Player, you can return to the Bins to switch to a new clip or continue working in Fusion.

To close the current clip in the Studio Player and return to the bins:

- Click the three-dot Options menu in the lower-left corner and choose Close.

Creating a Reel

A reel is a playlist or clip list that is viewed either as storyboard thumbnails or a timeline. In the bin, you create a new reel item to hold and play back multiple clips. The thumbnail for a reel appears with a multi-image border around it, making it easer to identify in the Bin window.

To create a reel in the current bin:

- Right-click in an empty area of the bin and choose New > Reel
- Click the Reel button along the bottom toolbar.



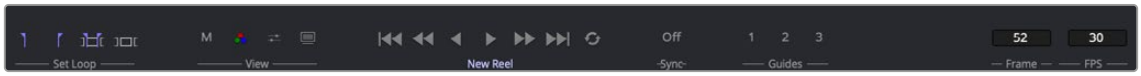
Use the New Reel button
to create a reel in the current bin.

Once created and named, the reel appears in the current bin. Double-clicking the reel opens the Studio Player interface along the bottom of the Bin window. An empty Studio Player appears in the top half of the window.

To add a clip to a reel:

1. Click the three-dot Options menu in the lower-left corner and choose Show Bins.
2. Drag a clip from one of the bins to the empty reel in the lower half of the window.

The toolbar across the bottom of the interface has various controls for setting a loop, showing and adjusting color, playback transport controls, collaboration sync, guide overlays, and frame number and playback speed in fps.



The toolbar along the bottom of the Studio Player includes controls to customize playback.

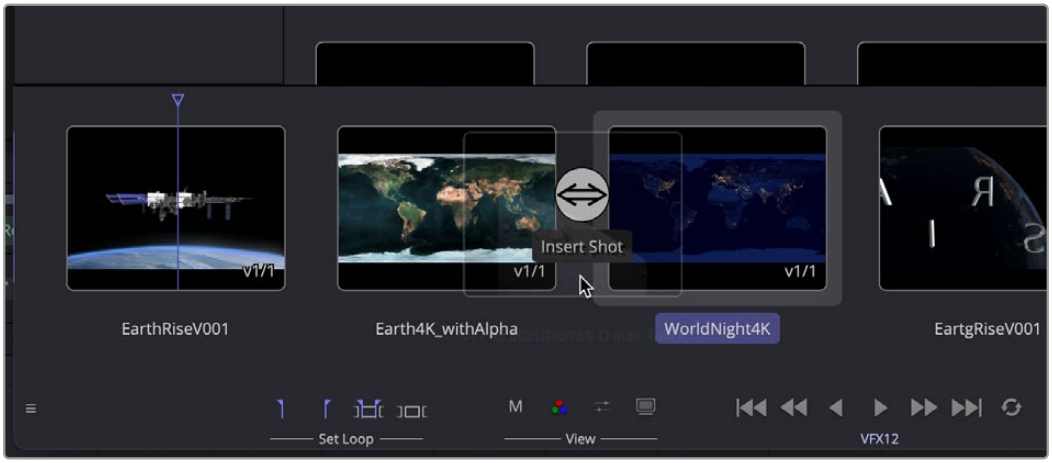
Toolbar buttons

- **Set Loop In/Out:** Sets the start and end frame for playing a section of the Timeline in a loop.
- **Shot:** Sets the loop for the entire clip.
- **Reset:** Disables the loop mode.
- **M:** Shows Metadata of the image.
- **RGB and alpha:** Toggles between displaying the color and alpha of the clip.
- **Brightness Gamma:** Adjusts the brightness and gamma of the viewer and is applied to all clips. Individual clip color controls can also be applied using another menu.

- **Video:** Outputs the image to Blackmagic Design DeckLink and UltraStudio devices.
- **Transport controls:** Used to play forward, backward, fast forward, and fast backward, as well as go to the start and end of a clip.
- **Sync:** A three-way toggle allowing the Studio Player to be controlled or to control another player over a network. The Off setting disables this functionality.
- **Guide buttons:** These three buttons control the visibility of three customizable guide settings.

Inserting Shots

Clips or comps from the bin can be dragged to the storyboard area of the reel to add and organize a playlist.



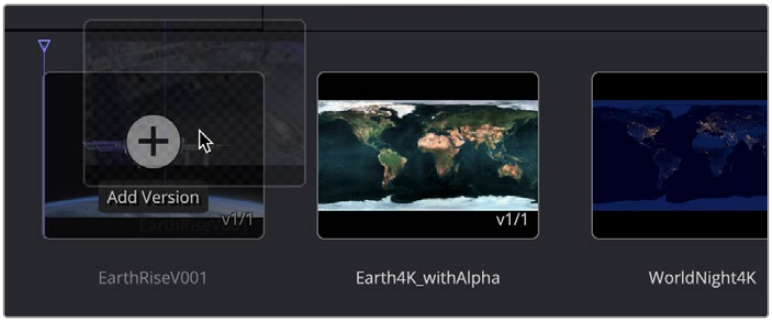
You can insert a shot between existing clips by positioning the new clip in between existing items in the reel.

To play the clips in a reel:

1. Click the three-dot Options menu in the lower-left corner and choose Show Player.
2. Position the playhead where you want to begin playing and click the Play button.

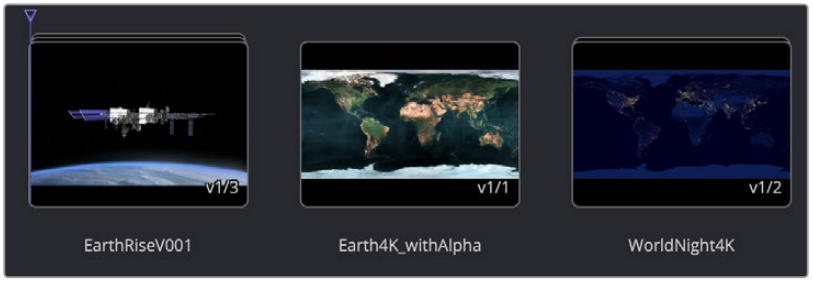
Creating Versions

Alternatively, you can add a version to an existing clip by dragging the new item on top.



You can create a new version of a clip by dragging the same clip into the player on top of itself.

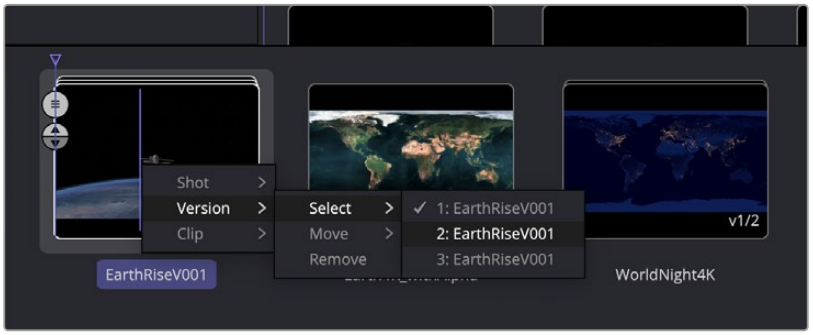
Versions of a shot will appear as stacked icons in the storyboard reel. The number of stacks in the icon indicate the number of versions included with that clip. The current version and the total versions are indicated by the number in the lower-right hand corner of the icon. In the example below, the first shot has three versions, the second shot has one version, and the last clip has two versions.



Clip versions are visible by the number in the lower right and graphically represented by the number of stacked icons.

Version Menu

You can choose which version to view by right-clicking over the clip in the storyboard and selecting it from the Version > Select menu.



Clip versions can be chosen by right-clicking on the icon in the player.

The Version menu also includes options to move or rearrange the order of the clip versions as well as remove a version, thereby deleting it from the stack.

Shot Menu

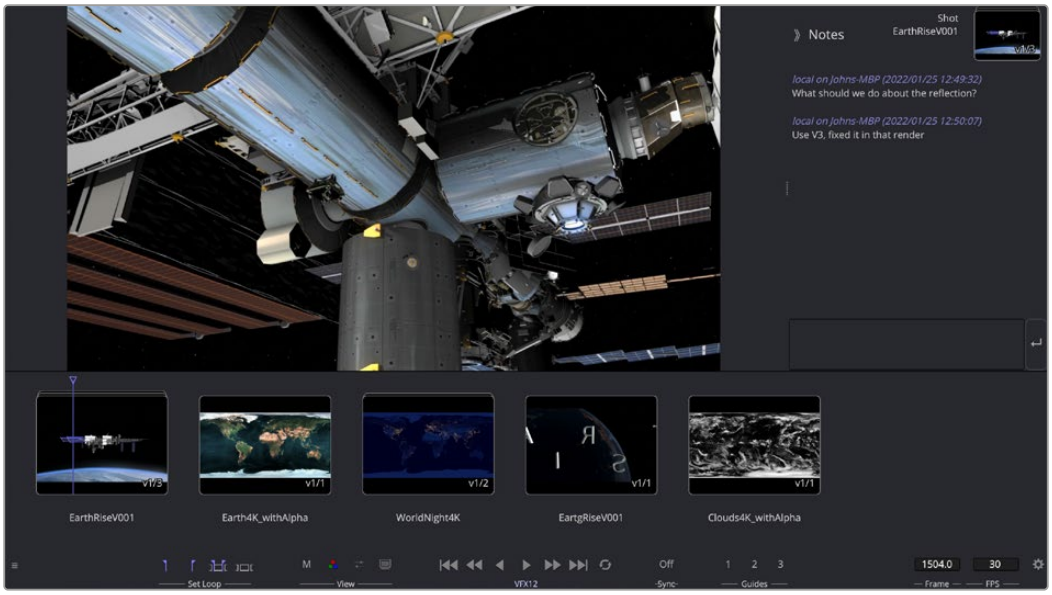
The per-clip Shot menu includes functions to Rename the shot, Remove the shot, Trim the clip's In and Out points, add notes, adjust the color, and add an audio soundtrack.



The Shot menu lets you modify the clip in the player.

- **Rename:** Allows you to change the name of the shot.
- **Remove:** Deletes the entire shot and all the versions from the project reel.

- **Trim:** Opens the Trim dialog to adjust the clip In point and Out point on the Timeline.
- **Notes:** Opens the Notes window to the right of the interface, allowing you to add a note to the specific shot.
- **Color:** Opens the Color Adjustments panel to perform Lift, Gamma, Gain, Brightness, and Contrast adjustments to clips.
- **Audio:** Allows you to attach an audio file to a clip in the Reel.



When notes are added, they are time and date stamped as well as name stamped.
The naming is from the bin login name and computer name.

Selecting Color from the Shot menu allows you to make tonal adjustments per clip using Color Decision List (CDL) style controls.

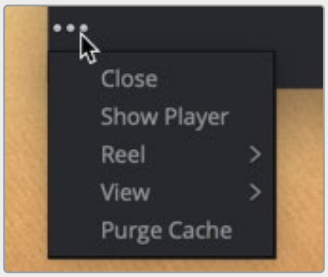


The Color Adjust controls from the Shot menu

The Audio menu option can import an audio .wav file that will play back along with the selected clip.

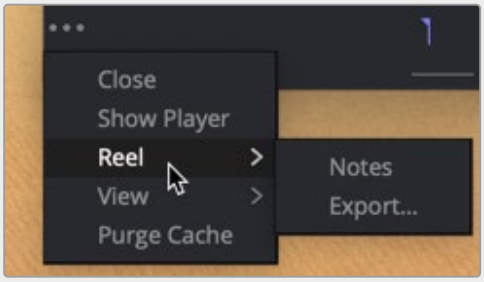
Options Menu

The three-dot Options menu in the lower left of the interface displays the menu that can be used to switch between viewer and the bin in the top half of the window. It is also used to clear the memory used in playback by selecting Purge Cache.



The Options menu includes options to switch the top half of the window between the viewer or bin contents.

Selecting Reel > Notes opens the Notes dialog to add annotations text to the entire reel project. The Reel > Export option saves the reel to disk as an ASCII readable format so it can be used elsewhere or archived.



The Reel submenu opens an area for production notes on the entire reel.

The View menu is used when you want to switch between the reel storyboard layout and a Timeline layout.

Guides

You can assign customizable guide overlays to three Guide buttons along the bottom of the Studio Player. Fusion includes four guides to choose from, but you can add your own using the XML Guide format and style information provided at the end of this chapter. You assign a customizable guide to one of the three Guide buttons by right-clicking over a button and selecting a guide from the list. To display the guide, click the assigned button.



Three toolbar Guide buttons enable guide overlays in the viewer.

Guides can be customized and placed in the Guides folder:

- **On macOS:** Macintosh HD/Users/username/Library/Application Support/Blackmagic Design/Fusion/Guides/
- **On Windows:** C:\Users\username\AppData\Roaming\Blackmagic Design\Fusion\Guides.
- **On Linux:** home/username/.fusion/BlackmagicDesign/Fusion/Guides.

Guides are a simple XML formatted text document saved with the .guide extension, as defined below. This makes it easy to create and share guides.

Custom Guide Format

The guides are files that have drawing instructions a bit like code, like this:

```
Guide
{
      Name = “10 Pixels”,
      Elements =
         {
             HLine { Y1=”10T” },
             HLine { Y1=”10B” },
             VLine { X1=”10L” },
             VLine { X1=”10R” },
         },
}
```

Or an example of safe area:

```
Guide
{
      Name = “Safe Frame”,

      Elements =
         {
                  HLine { Y1=”10%”, Pattern = 0xF0F0 },
                  HLine { Y1=”90%”, Pattern = 0xF0F0 },
                  HLine { Y1=”95%” },
                        HLine { Y1=”5%” },
                  VLine { X1=”10%”, Pattern = 0xF0F0 },
                  VLine { X1=”90%”, Pattern = 0xF0F0 },
                  VLine { X1=”95%” },
                        VLine { X1=”5%” },
                        HLine { Y1=”50%”, Pattern = 0xF0F0, Color = { R =
```

```
1.0, G = 0.75, B = 0.05, A=1.0 } },
                        VLine { X1="50%", Pattern = 0xF0F0, Color = { R =
1.0, G = 0.75, B = 0.05, A=1.0 } },
        },
}
```

Guide Styles

The style of a guide is defined by a set of properties that appear in the format shown below:

```
<HLine Y1="33%" Pattern="C0C0" Color="FFFFFFFF"/>
```

- **HLine:** Draws a horizontal line and requires a Y-value, which is measured from the top of the screen. The Y-value can be given either in percent (%) or in absolute pixels (px).
- **Vline:** Draws a vertical line and requires an X-value, which is measured from the left of the screen. The X-value can be given either in percent (%) or in absolute pixels (px).
- **Pattern:** The Pattern value is made up of four hex values and determines the visual appearance of the line.

  Examples for such patterns include:

  ```
  >>FFFF draws a solid line ________________
  >>EEEE a dashed line --------------------
  >>ECEC dash-dot line -.-.-.-.-.-.-.-.-.-.-.
  >>ECCC dash-dot-dot -..-..-..-..-..-..-..
  >>AAAA dotted line ……
  ```

- **Color:** The Color value is composed of four groups of two hex values each. The first three groups define the RGB colors; the last group defines the transparency. For instance, the hex value for pure red would be #FF000000, and pure lime green would be #00FF0000
- **Rectangle:** Draws a rectangle, which can be empty or filled, and supports the same pattern and color settings described above.

  It requires two X- and two Y-values to define the extent <Rectangle Pattern="F0F0" X1="10%"

  ```
  Y1="10%" X2="90%" Y2="90%">.
  ```

- **FillMode:** Applies to rectangles only and defines whether the inside or the outside of the rectangle should be filled with a color. Leave this value out to have just a bounding rectangle without any fill.

  ```
  >>FillMode = ("None"|"Inside"|"Outside")
  ```

- **FillColor:** Applies to rectangles only and defines the color of the filled area specified by FillMode.

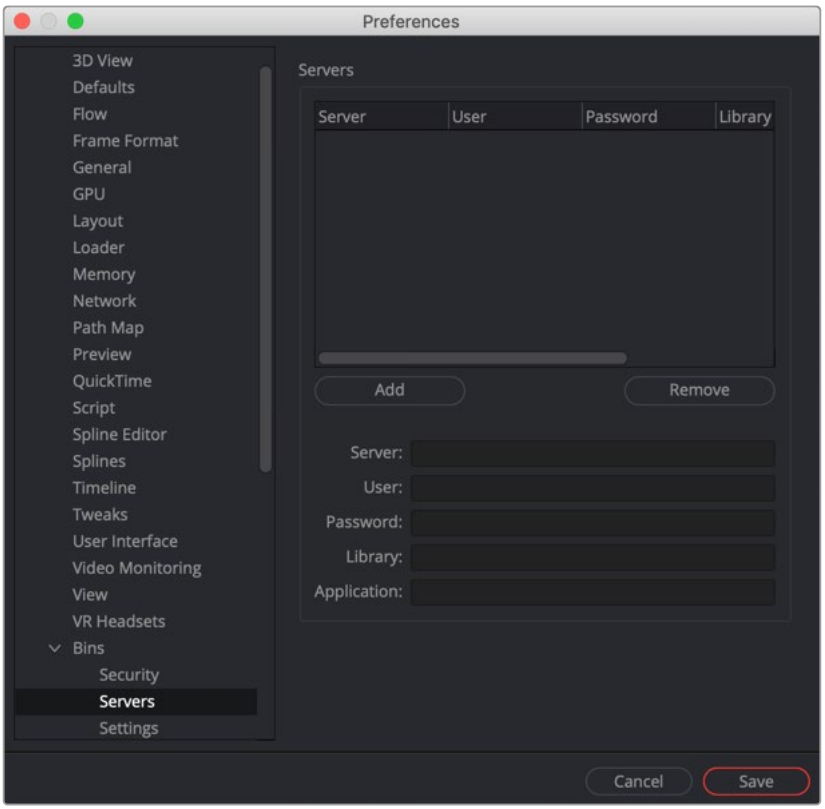
  ```
  >>FillColor="FF000020"
  ```

Connecting Bins Over a Network

You can share bins among computers running Fusion on the network so that multiple visual effects artists can share assets, presets, and even entire compositions. These shared bins are called remote bins, and everyone can share one or more remote bin in a studio.

To connect to a remote system and display its bins:

1. Choose Fusion Studio > Preferences.
2. In the Preferences dialog, select Global > Bins > Servers in the list.

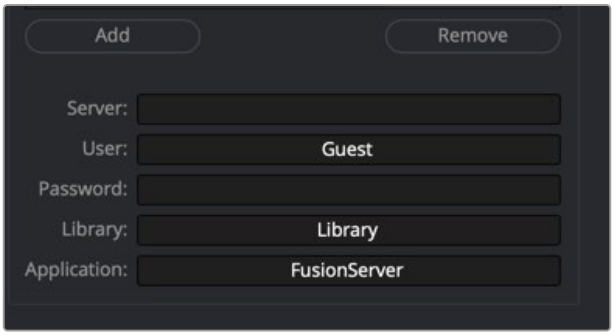


The bin servers Preferences panel

This panel shows a list of the available bin servers, with buttons below for entries to be added to or deleted from the list.

Adding a Remote Bin Entry

If you want to add another Remote bin to the list of available Remote bins, you can click the Add button in the bin servers Preferences panel. The text controls below the button will become enabled for editing. In the Server field, type the system name or IP address where the bin is hosted.



Add the IP address where the bin server is hosted.

Then add a User name and Password if one is needed to access the server.

The Library field lets you name the bins. So if you want to create a bin for individual projects, you would name it in the Library field, and each project would gets its own bin.

The Application field allows larger studios to specify some other program to serve out the bin requests.

Once you've finished setting up the bin server information and clicked Save in the Preferences window, you can open the Bins window to test your bin server. Opening the Bins window is the first time your connection to the server will be tested. If it cannot connect, the bin server will still be listed, with access denied or unavailable marked next to the name on the bins sidebar.

There is no practical limit to the number of bins that can be accessed.

Accessing Remote Bins

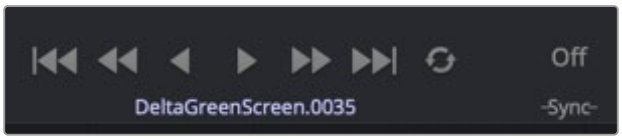
Bin servers behave just like a local bins. Any bin added in the preferences show in the Bins sidebar as another top-level item. The available bins are shown by name with a status and, if required, a password. Bins unavailable to you are marked as (unavailable).

Permissions

Unlike other directories on a server, your access to bins on a network is stored in the bin document. The bins themselves contain all the users and passwords in plain text, making it easy for someone to administer the bins.

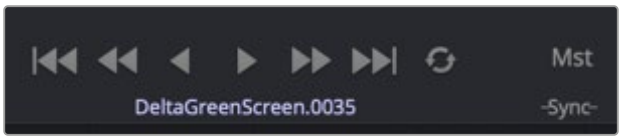
Studio Player and Bin Server

Reel projects can be shared by multiple artists across the studio via the bin server system, reviewing and adding versions and notes, all independently at the same time. With the Sync function, multiple people can collaborate together with synced playback and scrubbing.



The Sync button is a three-way toggle button: Off, Slave, and Master.

When the Sync function is On, the transport controls can be set to control the playback or follow the master controller.



Master enables controlling other slave players on the network.



The local transport controls can be disabled, causing playback to follow the master.

Chapter 14

Fusion Connect

This chapter goes into detail on how to use the Fusion Connect AVX2 plug-in with an Avid Media Composer editing system. The Fusion Connect AVX plug-in is only available with Fusion Studio.

Contents

Fusion Connect Overview 356
System Requirements 356
The Effects Palette 356
The Layer Input Dialog 357
Applying Fusion Connect to a Transition Point 357
Export Clips 357
Edit Effect 358
Browse for Location 359
Auto Render in Fusion 359
Red on Missing Frames 359
Compress Exported Frames 359
Edit Effect Also Launches Fusion 359
Versioning 360
Create New Version 360
Version 360
About RAW Images 360
About Color Depth 360
Manual vs. Auto-Render 360
Fusion/Avid Project Relationship 362
Rendering with Fusion 363
Directory Structure of Fusion Connect Media 363
Advanced Project Paths 365
Configuring Paths on macOS 365
Configuring Paths on Windows 366
Fields and Variables 366
Environment Variables 367

Fusion Connect Overview

Fusion Connect is an AVX2 plug-in for Avid Media Composer. It allows editors to create a conduit between the Timeline in Avid editing products and Fusion Studio. Fusion Connect exports clips from the Avid Timeline as image sequences and assembles Fusion compositions that allow you to work your magic on the content.

Fusion can be started automatically by the plug-in if Fusion is installed on the same system, or it can be used on remote computers to modify the composition.

System Requirements

Fusion Connect has the following requirements:

- **Supported Avid products:** Media Composer 8.x
- **Supported product:** Fusion Studio 8.1 or later
- **Installation:** Two files will be installed in your Media Composer:

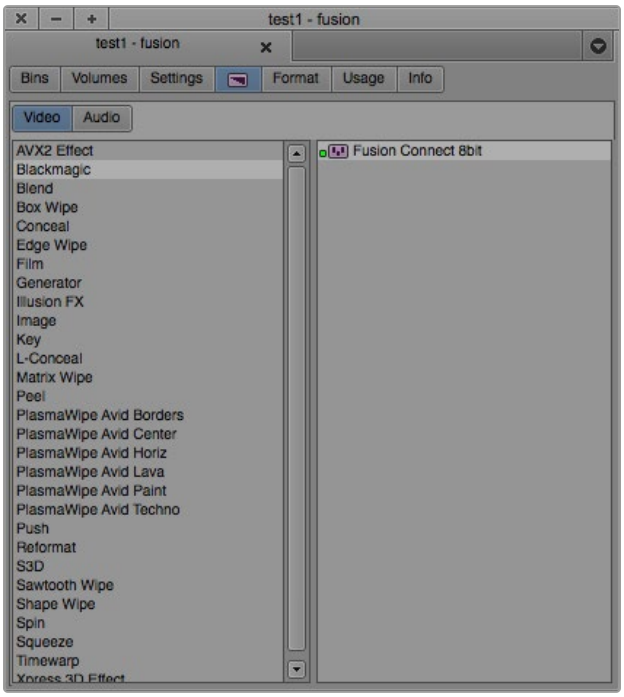
  Fusion Connect.avx

  BlackmagicFusionConnect.lua

- **Avid's default directory:** \Avid\AVX2_Plug-ins

The Effects Palette

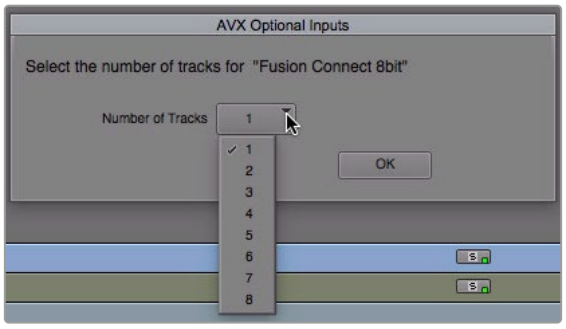
After launching Media Composer, Fusion Connect is located in the Blackmagic Effects Palette category. As with any segment or transition effect in the Effects Palette, you can apply the Fusion Connect AVX2 plug-in to any clip. This includes filler*, edit transition point, or video track layers on the Avid Timeline.



Avid Media Composer Effects palette showing the Blackmagic category and Fusion Connect AVX plug-in.

The Layer Input Dialog

When you apply the Fusion Connect AVX plug-in to a clip or a layer in the Timeline, you are presented with an AVX Optional Inputs dialog box.



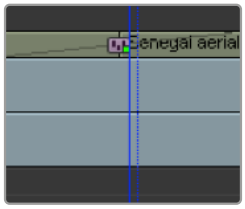
The Fusion Connect AVX plug-in can access multiple layers from the Media Composer Timeline.

Once the layer count is selected, Fusion Connect will be applied to the Timeline.

- Select the layer count equal to the number of video track layers you want to ingest into Fusion.
- Filler can be used as a layer.
- Fusion Connect will allow a maximum of eight layers.

Applying Fusion Connect to a Transition Point

If you apply Fusion Connect to a transition point, no dialog box will display, and the AVX plug-in will simply be applied to the Timeline transition point.



A Fusion Connect AVX plug-in applied as a transition to a cut in the Timeline

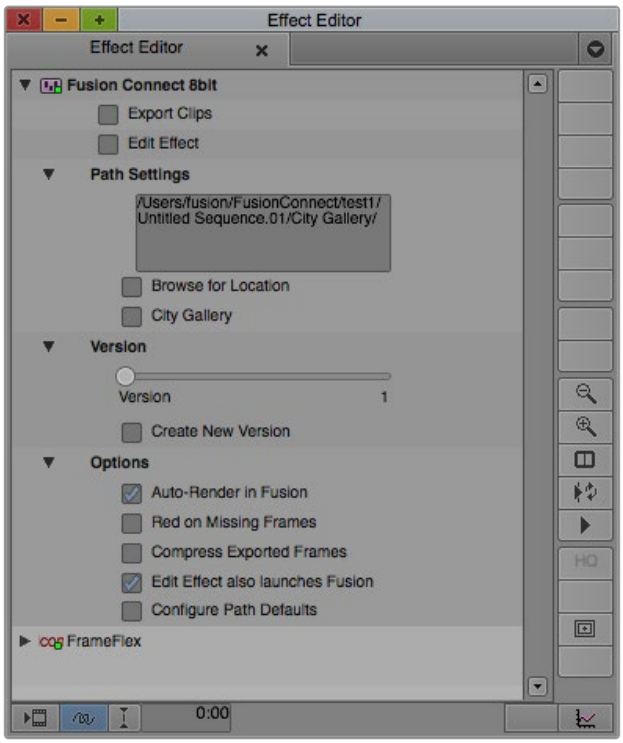
You can use the Avid dialog boxes or smart tools to adjust the length and offset of the transition to Start, Center, End, or Custom.

Export Clips

By pressing the Export Clips button in the Effects Editor, Fusion Connect exports all the associated clips as image sequences, to provide access to them in Fusion. Any previously exported existing images are overwritten, ensuring that all media needed by Fusion is accessible. Performing an export is desired if you want to use Fusion on a different computer than the one with Media Composer installed.

When using Fusion on the same computer as Media Composer, there is no need to export the clips explicitly by checking the Export Clips checkbox. Without this option enabled, Fusion Connect saves the source frames each time images are displayed, scrubbed, or played back from the Timeline. Depending on your Media Composer Timeline settings, these interactively exported images might be half-resolution based on Avid proxy settings. When scrubbing around the Timeline, only a few frames—namely, those that are fully displayed during scrubbing—might be written to disk.

TIP: Set your Timeline Video Quality button to Full Quality (green) and 10-bit color bit depth. If the Timeline resolution is set to Draft Quality (green/yellow) or Best Performance (yellow), Fusion receives subsampled, lower-resolution images.



Fusion Connect AVX plug-in controls in the Effects Editor

Edit Effect

After exporting the clips, the Edit Effect button performs three subsequent functions:

- Creates a Fusion composition, with Loaders, Savers, and Merges (for layers), or a Dissolve (for transitions). This function is only performed the first time a comp is created when the Fusion Connect AVX2 plug-in is applied.
- Launches Fusion (if installed on the machine), if it is not already launched.
- Opens the Fusion comp associated with created effects.

Browse for Location

The Fusion Connect media folders are created on the drive where the associated Avid media resides, defaulting to the root level of that drive. However, you can choose a new location for the Fusion Connect media.

To choose another location to store and access the media:

1. In the Media Composer Effects Editor, click the Browse for Location button.
2. In the File Browser, change the location and create additional folders if desired.

The path settings field in the Effects Editor updates to show the current location. If you apply Fusion Connect to another clip in the Timeline, the last location is remembered.

Auto Render in Fusion

The Auto Render button is a toggle that allows you to automatically render your Fusion comp from within Avid. Please note that this method of rendering has limitations and can be significantly slower than rendering directly in Fusion. It is mostly used for batch rendering on the Avid Timeline. Auto Render also exports the necessary media without having to manually execute the Export Clips function first.

Red on Missing Frames

The Red on Missing Frames button is a toggle used to display red images within the Avid Timeline viewer (Timeline monitor) if no rendered frames from Fusion Studio can be found, or if the rendered frames are not of a high enough resolution. When disabled, the original untouched frames will be shown instead of the red frame.

Compress Exported Frames

When enabled, this button creates Fusion RAW files with compression for, both exported and rendered frames. This creates smaller file sizes, saving disk space. As with any other compression algorithm, the compression adds time to the write process of the file sequence.

Edit Effect Also Launches Fusion

When enabled, this button opens Fusion Studio when the Edit Effect button is clicked, allowing for a more direct Avid/Fusion workflow. When disabled, this button will not launch Fusion Studio when the Edit Effect button is clicked, but a Fusion .comp file is created. This is useful when working with Fusion Studio on a separate computer than the one running the Media Composer software.

Versioning

Creating visual effects is almost always an iterative process. You'll often need to create revisions after your first pass at the effect. Built into Fusion Connect is a versioning feature that lets you create multiple revisions of an effect and switch between them from within Media Composer.

Create New Version

This checkbox creates a copy of the current comp without affecting the original.

Any changes that are rendered in the copy will be written to a new folder and become another version of the rendered result played on the Avid Timeline. Previous versions of the comp and their rendered results are accessible using the Version slider.

Version

This slider selects which version of the comp is used in the Media Composer Timeline. It can be used to interactively switch from one version to the other in order to compare the results.

About RAW Images

Fusion Connect creates a Fusion RAW file image sequence for intermediate folders between Avid Media Composer and Fusion Studio, in order to preserve all the image information. This allows the images to reside on disk and not take up space in RAM. The benefits of using Fusion RAW include:

- The ability to continue the editing process while an effect is rendering
- The ability to take advantage of network rendering
- The ability to retime footage
- The ability to run Fusion Studio remotely

About Color Depth

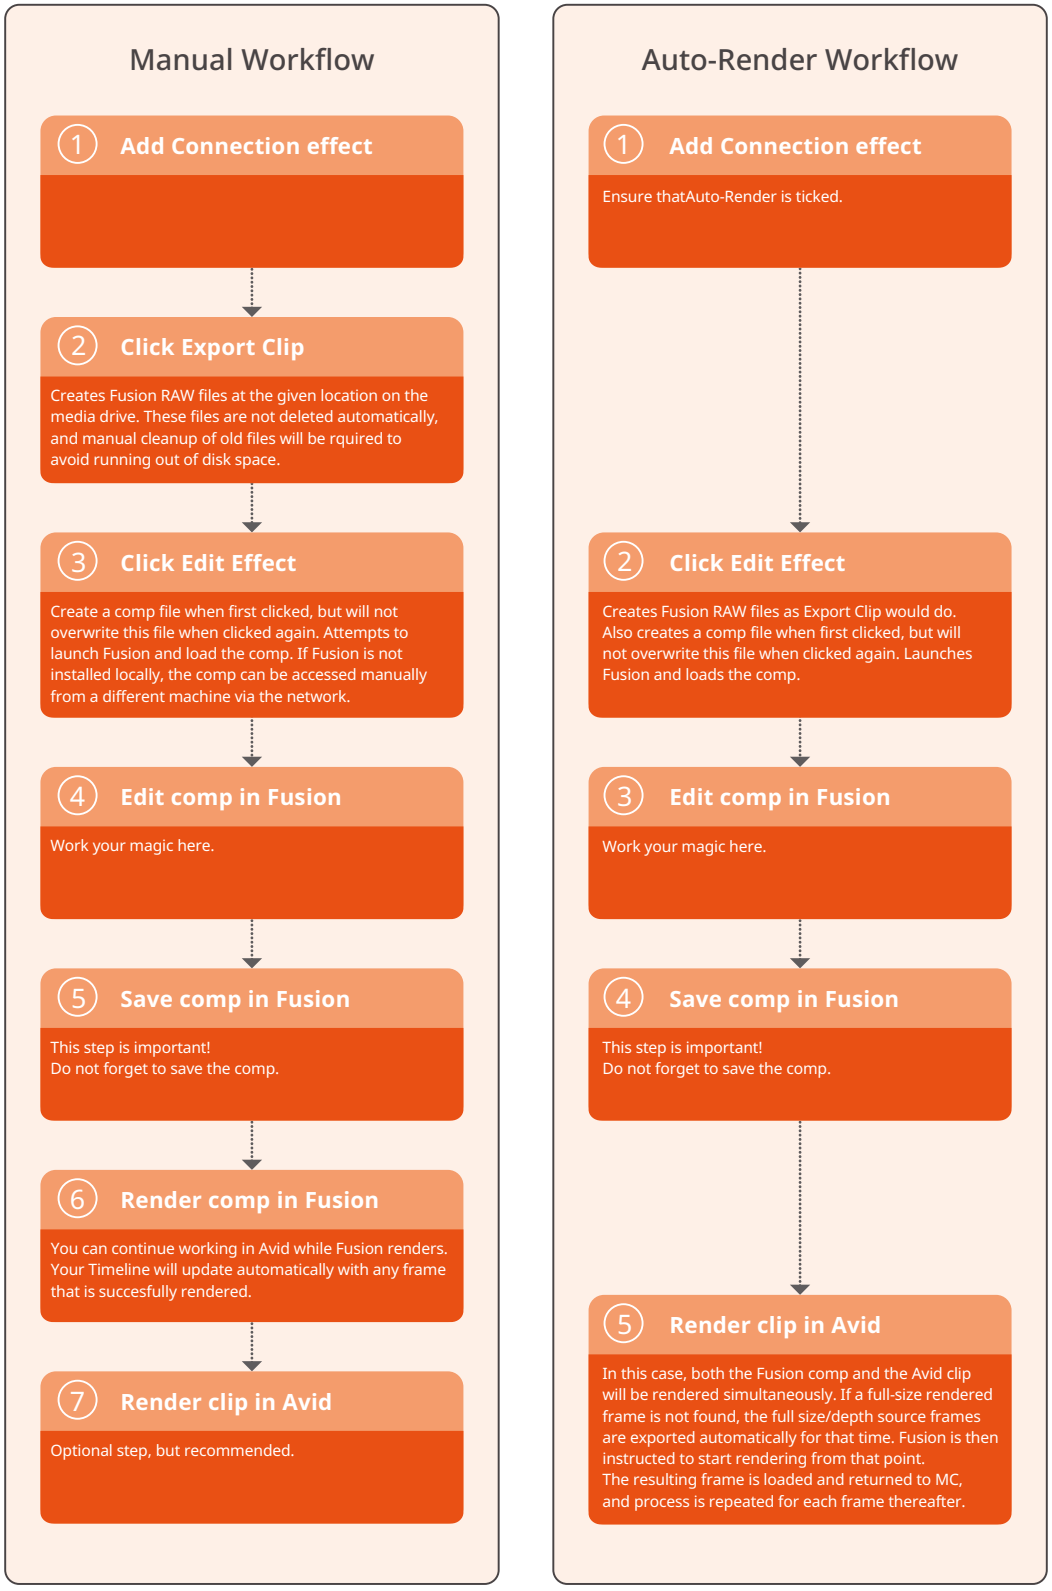
Fusion Connect derives its images directly from the RGB data within Avid Media files. This allows the images to be codec agnostic. All RAW files from Avid that begin as 8- or 10-bit images are remapped to 16-bit float in Fusion. Rendered results from Fusion are processed in 16-bit float to maintain the full color fidelity supported by Media Composer.

Manual vs. Auto-Render

While Auto-Render is the easier workflow, the manual approach offers faster renderings in Fusion and more control over the performance and memory usage on your system.

- In the manual workflow, Fusion Studio is not required to be installed on the Avid system itself but can reside on any other computer.
- For Auto-Render, Fusion Studio must be installed on the local computer.

The following diagram shows typical workflows for manual and automatic renders.



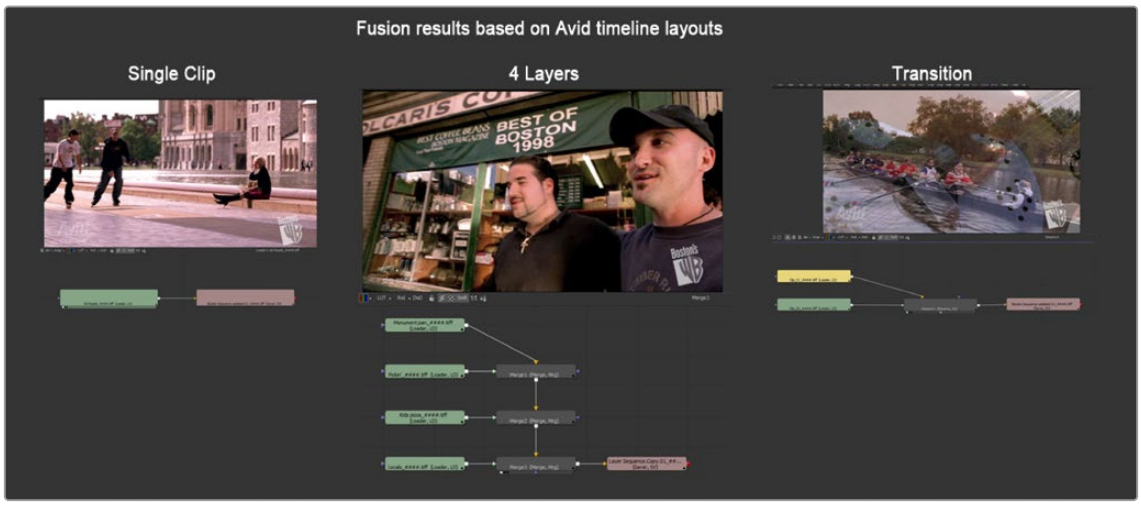
Avid/Fusion layer to comp relationship for auto and manual renders

Once the initial trip from Avid to Fusion is complete, depending on the type of clip to which you assigned the effect, in Fusion Studio you will be presented with one of the following:

- A Loader node representing a single clip.
- Two or more Loaders connected to Merge nodes representing layers.
- Two Loaders connected to a Dissolve node representing a transition.

In all three node tree layouts outlined above, there will also be a Saver node. The Saver node is automatically set to render to the directory that is connected back to the Media Composer Timeline with the correct format. If for some reason the file format or file path are changed in the Saver node, the Fusion Connect process will not render correctly.

TIP: Due to the design of the AVX2 standard, hidden embedded handles are not supported. To add handles, prior to exporting to Fusion, increase the length of your clip in the Media Composer Timeline to include the handle length.

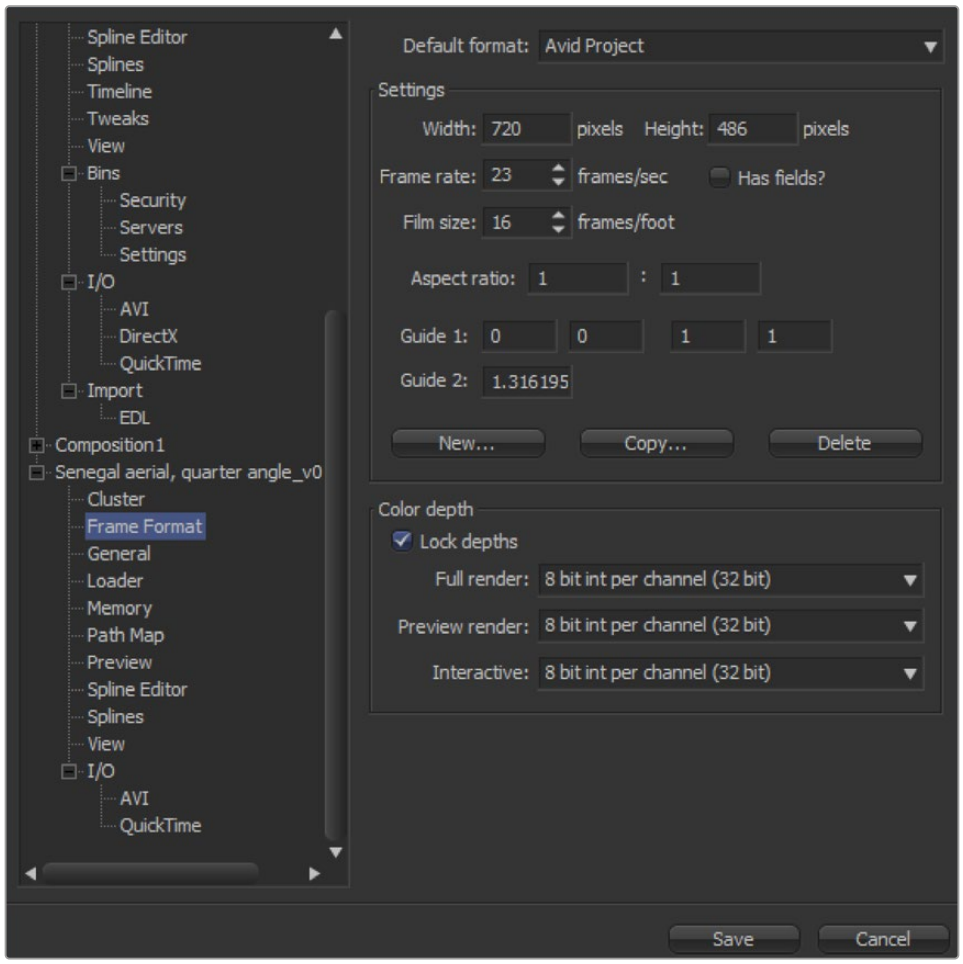


Fusion Node Editor representations of a single clip segment effect in the Media Composer Timeline (left), a multi-layered composite (center), and the transition (right)

TIP: If segments in the Avid Timeline have different durations, move the longest clip to the top layer and apply Fusion Connect to that clip. This will ensure all clips are brought into Fusion without being trimmed to the shortest clip.

Fusion/Avid Project Relationship

The frame rate and image size preferences created in Media Composer are adopted within Fusion's frame rate preferences. This allows for consistency in formats for the roundtrip process from Avid to Fusion and back to Avid. The format settings do not prevent you from using or mixing any other sized imaging within the composition as Fusion is resolution independent.



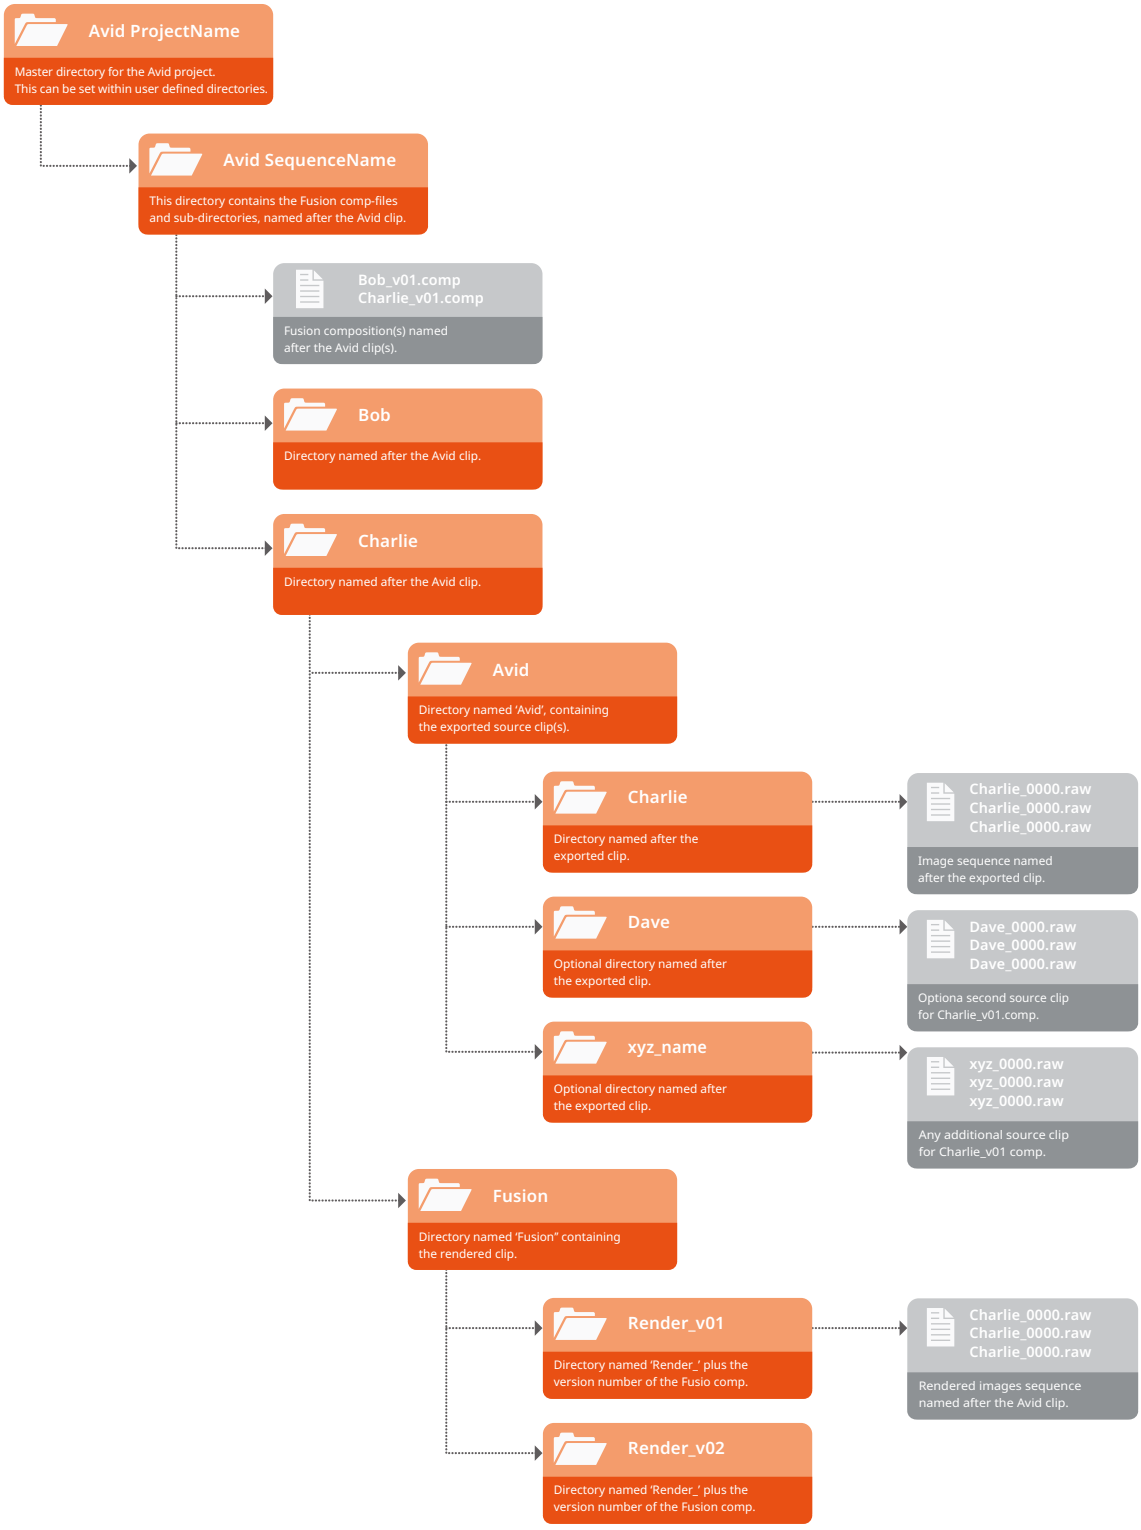
Fusion Connect AVX uses frame rate and resolution from the Media Composer Timeline.

Rendering with Fusion

When you perform a render of your comp inside Fusion Studio, the results are rendered to the output folder created by Fusion Connect during the initial application of the plug-in to the Timeline. Upon rendering, you immediately see the results of the rendered Fusion comp in the Avid Timeline. Even while Fusion is rendering, you can continue with the edit process on any clip except for the associated clip being rendered at the time.

Directory Structure of Fusion Connect Media

Fusion Connect creates a logical folder structure that is not affiliated with the Avid Media Files folder but rather the Fusion Connect AVX2 plug-in. Based on data gathered during the AVX application to the Timeline, a logical folder hierarchy is automatically created based on Avid projects, sequences, and clips. This structure allows for multiple instances of Fusion Studio to access the media and multiple instances of the AVX to relate to a single Fusion comp. In a re-edit situation, media is never duplicated but is overwritten to avoid multiple copies of identical media.



If you apply the effect to a transition, the naming behavior might be somewhat different.

By default, Media Composer refers to the two clips of a transition as “Clip_001” and “Clip_002”. Based on the naming convention, Fusion Connect will create folders with matching names. If such folders already exist, because another transition has already used up “Clip_001” and “Clip_002”, the numbers will be incremented automatically.

Likewise, "_001" will be added incrementally to the group folder name, if a folder of that name already exists. The corresponding comp file will be named accordingly.



Fusion Connect AVX creates folder structures in the OS to save media and Fusion compositions. Those names are reflected in the Timeline.

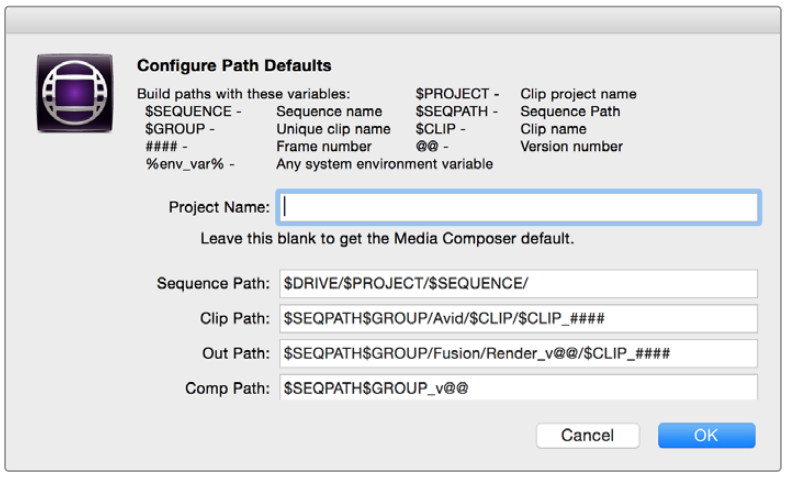
You will notice that the Fusion Connect icon is a green dot (real-time) effect. If your hardware is fast enough, the results that populate the plug-in will play in real time. However, it's recommended that you render the green dot effect, which will force an MXF precompute to be created to guarantee real-time playback.

Advanced Project Paths

The Fusion Connect AVX2 plug-in controls the pathing of Fusion Connect's .raw media and Fusion's .comp files as well as showing and hiding project and sequence level directories. This is achieved through environment variables, which are set in the operating system. This gives you the most flexible control over pathing your media, and as the name depicts, you can change variables (controls) in various application environments. This is useful for network storage and when Fusion Studio is running on other systems.

Configuring Paths on macOS

When using Fusion Connect on macOS, the Configure Path Defaults dialog looks like this:

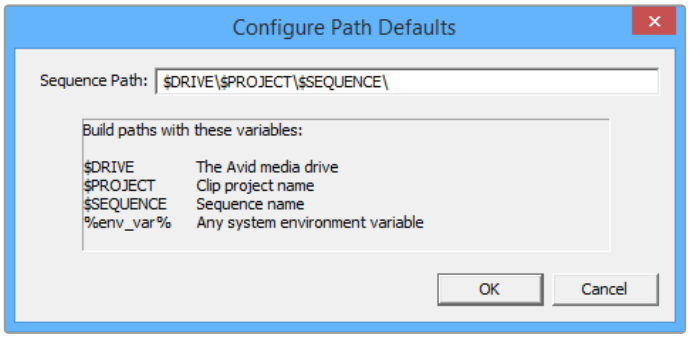


Environment variable path maps on macOS.

Default paths can be configured using variables similarly as on Windows, but for added convenience it is possible to enter any desired path defaults directly into fields in the dialog, without the need for using environment variables.

Configuring Paths on Windows

When using Fusion Connect on Windows, the Configure Path Defaults dialog looks like this:



Environment variable path maps on Windows.

Fusion Connect can define the user variables directly in the Fusion Connect plug-in. Click the Configure Path Defaults button to launch the path defaults dialog editor. In the Options section of the Fusion Connect AVX2 plug-in, click the triangle to reveal the path details.

Fields and Variables

Variables are the control title defined specifically by the application that is being controlled, while values are the instructions that tell a variable what to do. The fields and variables that can be used on macOS and Windows are described in the following table:

Field	Variable	Environment Variable	Description
Project Name	$PROJECT	CONNECT_PROJECT	Overrides the current Avid Project name.
	$DRIVE	CONNECT_DRIVE	Drive or folder for all Connect projects
	$SEQUENCE	–	Name of Avid sequence
Sequence Path	$SEQPATH	CONNECT_SEQUENCE_PATH	Folder for all Connect files in this sequence
	$GROUP	–	Unique name of this Connect instance
	$CLIP	–	Name of exported clip
Clip Path	–	CONNECT_CLIP_PATH	Folder for exported clips from Avid
Out Path	–	CONNECT_OUT_PATH	Folder for rendered results from Fusion
Comp Path	–	CONNECT_COMP_PATH	Location and name of Fusion comp file

Environment Variables

The pathing can be set in the environment variables of the system, so that IT management of the project paths can be achieved.

Accessing the Environment Variables on Windows

The quickest way to access environment variables is through your Windows control panel by searching the word “env” without the quotes. You have a choice of editing at a user level or system level.

User Variables

Click on the link that says “Edit environment variables for your account.”

System Variables

Click on the link that says “Edit the system environment variables.”

Accessing the Environment Variables on macOS

To set an environment variable on macOS, you must use the Terminal window. Environment variables on macOS are added to the .bash_profile directory for the current user.

User Variables

For system-wide operations, place the environment variable in ~/.bash_profile

TIP: System variables control the environment throughout the operating system, no matter which user is logged in.

User variables always overrule any system variable; therefore, the user variable always wins if control for a specific function is duplicated in the user and system variable.

System Variables

For system-wide operations, place the environment variable in /etc/profile

Environment Variables and What They Mean

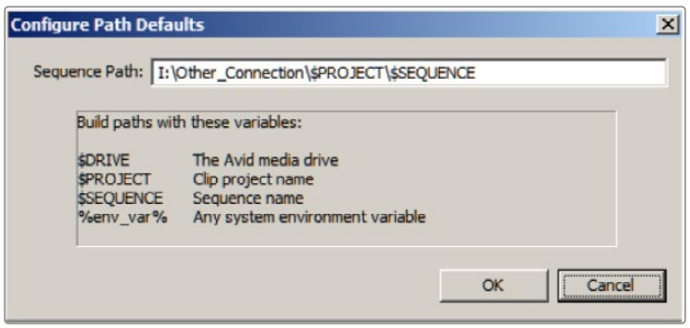
The Fusion Connect AVX plug-in can use environment variables to set different locations for certain folders or files. Each computer OS has a unique way of entering environment variables, but in every OS, the variable must be typed exactly as shown.

TIP: If you type directly in Fusion Connect’s Path Editor, you do not have to type the variable, just the value. You also can make modifications without having to restart the Media Composer! The only caveat is that in order to remove a variable, you must exit Media Composer and clear the environment variable in the Windows interface or macOS Terminal and restart the Media Composer.

Other values you can control derived from your Avid Bin include:

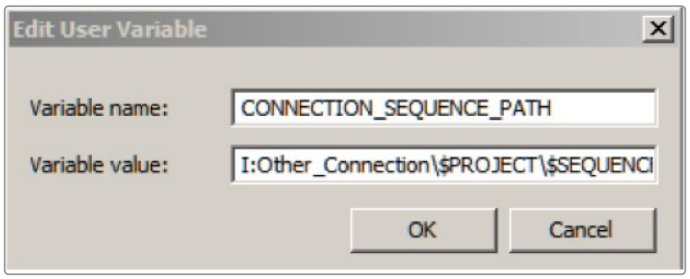
Values	Description
$DRIVE	This will force the directory to the drive where the Avid media is stored.
$PROJECT	This will force a directory based on the Avid project name for which the media was digitized/imported or AMA linked.
$SEQUENCE	This will force a directory based on the Avid SEQUENCE name for which the media was digitized/imported or AMA linked.

Here is an example of how a variable can be set up to support project and sequence names within your directory.



Environment variable project and sequence name examples on Windows

Here is the same example in the Windows environment variable editor.



Windows environment user variable editor

Chapter 15

Preferences

This chapter covers the various options that are available from the Fusion Preferences Window.

Contents

Preferences Overview 370
Categories of Preferences 371
Preferences In Depth 374
3D View 374
AVI 375
Defaults 376
Flow 377
Frame Format 379
General 381
GPU 383
Layout 384
Loader 385
Memory 387
Network 389
Path Maps 390
Preview 394
QuickTime 395
Script 396
Spline Editor 397
Splines 398
Timeline 400
Tweaks 401
User Interface 404
Video Monitoring 406
View 407
VR Headsets 409
Bins/Security 411
Bins/Server 412
Bins/Settings 413
EDL Import 414
Customization 415
Shortcuts Customization 415
Customizing Preferences 416

Preferences Overview

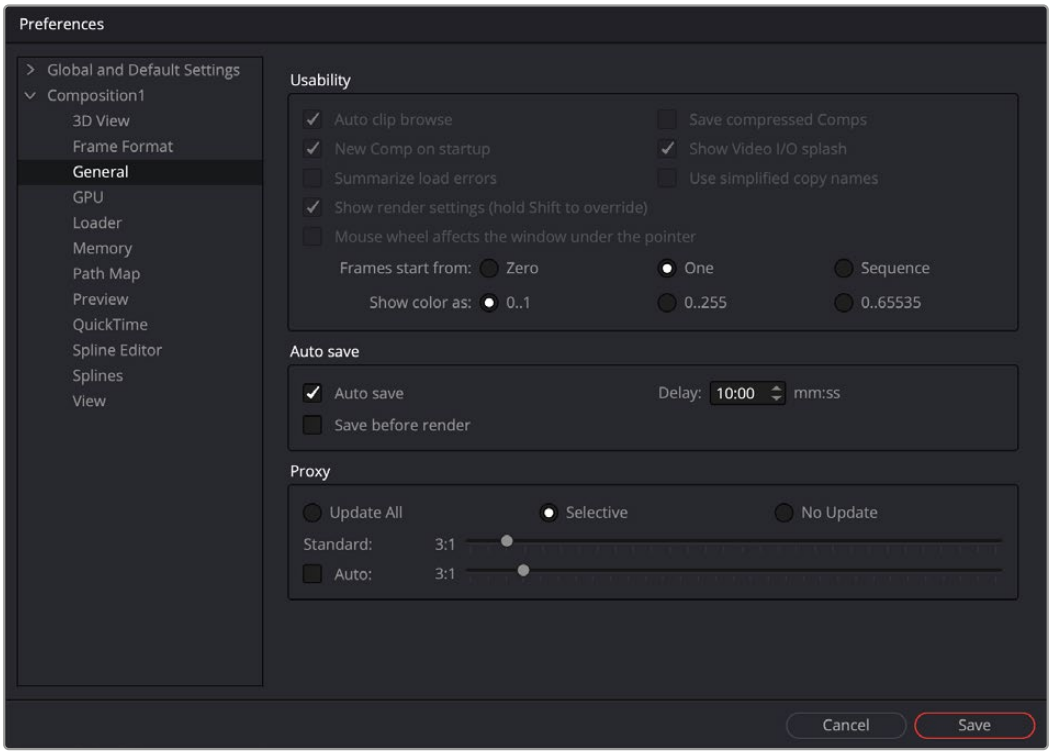
The Preferences window provides a wide variety of optional settings available for you to configure Fusion's behavior to better suit your working environment. These settings are accessed via the Preferences window. The Preferences window can be opened from a menu at the top of the interface.

In DaVinci Resolve, to open the Fusion Preferences window, do one of the following:

- On macOS, switch to the Fusion page and choose Fusion > Fusion Settings.
- On Windows, switch to the Fusion page and choose Fusion > Fusion Settings.
- On Linux, switch to the Fusion page and choose Fusion > Fusion Settings.

In Fusion Studio, to open the Fusion Preferences window, do one of the following:

- On macOS, choose Fusion Studio > Preferences.
- On Windows, choose File > Preferences.
- On Linux, choose File > Preferences.



The Preferences window includes composition settings.

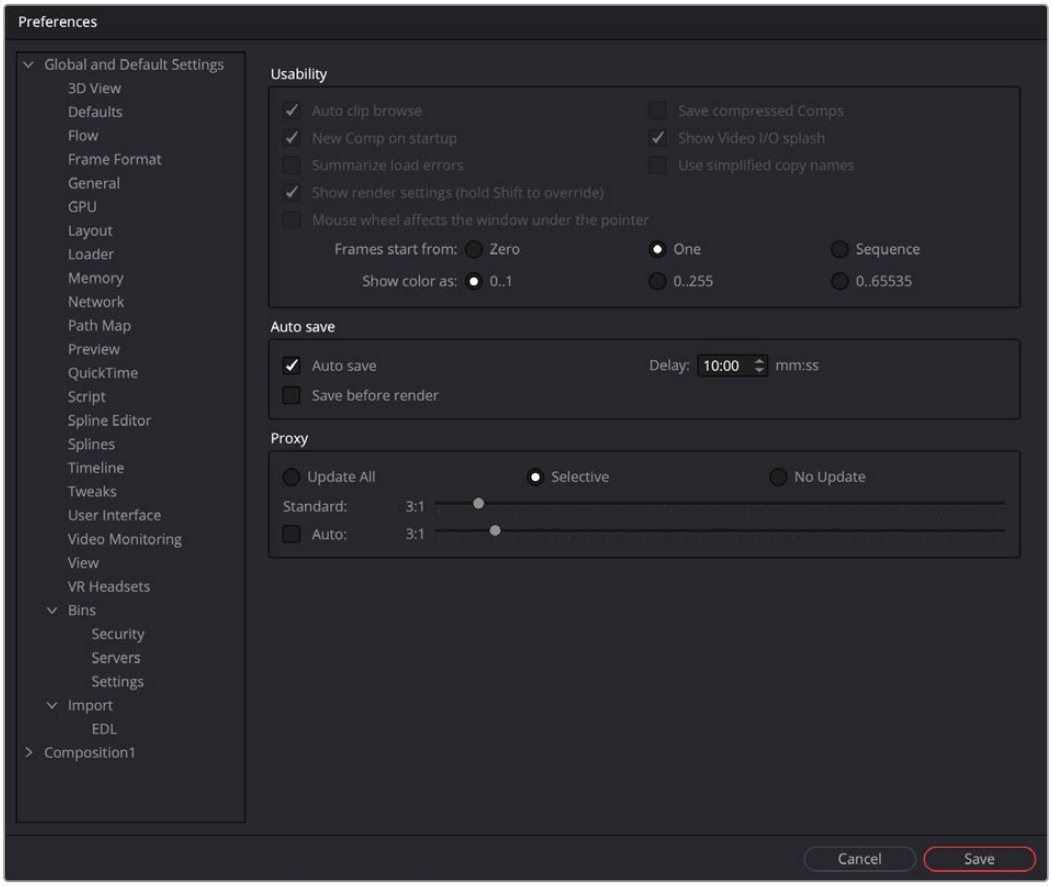
Global and Composition Preferences

The Preferences window is divided into a category sidebar on the left and the settings panel on the right. In Fusion Studio, there are two levels of preferences: Global and Composition. The Fusion page in DaVinci Resolve uses just a single Global preference that affects every project, new and existing.

The Global preferences are used to set options specific to Fusion's overall behavior as well as defaults for each new composition. The Composition preferences in Fusion Studio can further modify the currently open composition without affecting the Global preferences or any other composition that is open but not displayed.

Categories of Preferences

The first entry in the Preferences sidebar is assigned to the Global preferences. Clicking the Global and Default Settings disclosure arrow reveals the following sections.



Fusion Studio Global and Default settings

3D View

The 3D View preferences offer control over various parameters of the 3D Viewers, including grids, default ambient light setup, and Stereoscopic views.

AVI (Windows Fusion Studio Only)

The AVI preferences configure the default compression options when creating AVI files from a Saver node. These settings can be overridden on a case-by-case basis using the Format tab in the Saver's Inspector.

Defaults

The Defaults preferences are used to select default behavior for a variety of options, such as animation, global range, timecode display, and automatic tool merging.

Flow

You use the Flow preferences to set many of the same options found in the Node Editor's contextual menu, like settings for Tile picture, the Navigator, and pipe style.

Frame Format

The Frame Format preferences are used to create new frame formats as well as select the default image height and width when adding new creator tools like Background and Text+. You also set the frame rate for playback.

General

The General preferences contain options for the general operation, such as auto save, and gamma settings for color controls.

GPU (Fusion Studio Only)

The GPU preferences include options for selecting specific GPU acceleration methods based on your computer platform and hardware capabilities. It is also used for enabling caching, and debugging GPU devices and tools.

Layout (Fusion Studio Only)

You can use the Layout preferences to save the exact layout of Fusion's windows.

Loader (Fusion Studio Only)

Using the Loader preferences, you can set options for the default Loader's depth and aspect ratio as well as define the local and network LoaderCache settings.

Memory (Fusion Studio Only)

Memory management for multi-frame and simultaneous branch rendering is configured in the Memory preferences.

Network (Fusion Studio Only)

The Network rendering preferences are used to configure options such as selecting a render master, email notification, and whether the machine can be used as a render slave.

Path Map

Path Map preferences are used to configure virtual file path names used by Loaders and Savers as well as the folders used by Fusion to locate comps, macros, scripts, tool settings, disk caches, and more.

Preview (Fusion Studio Only)

The Preview preferences is where you configure the Preview creation and playback options.

QuickTime (macOS Fusion Studio Only)

This section lets you preconfigure the QuickTime codec used for rendering.

Script

The Script preferences include a field for passwords used to execute scripts externally, programs to use for editing scripts, and the default Python version to use.

Spline Editor

The Spline Editor preferences allow you to set various spline options for Autosnap behavior, handles, markers, and more.

Splines

Options for the handling and smoothing of animation splines, Tracker path defaults, onion-skinning, roto assist, and more are found in the Splines preference.

Timeline

The Timeline preferences is where you create and edit Timeline/Spline filters and set default options for the Keyframes Editor.

Tweaks

The Tweaks preferences handle miscellaneous settings for modifying the behavior when loading frames over the network and queue/network rendering.

User Interface

These preferences set the appearance of the user interface window and how the Inspector is displayed.

Video Monitoring (Fusion Studio Only)

The Video Monitoring preferences is where you can configure your Blackmagic video display hardware for monitoring on an HD, Ultra HD, or DCI 4K display.

View

The View preferences are used to manage settings for viewers, including default control colors, Z-depth channel viewing ranges, default LUTs, padding for fit, zoom, and more.

VR Headsets

The VR Headsets preferences allow configuration of any connected Virtual Reality headsets, including how stereo and 3D scenes are viewed.

Bins (Fusion Studio Only)

There are three panels as part of the Bins preferences: a Security panel where you set users and passwords for serving the local bins; a Servers panel used to select which remote Bin servers are connected; and a Settings panel for stamp rendering.

Import

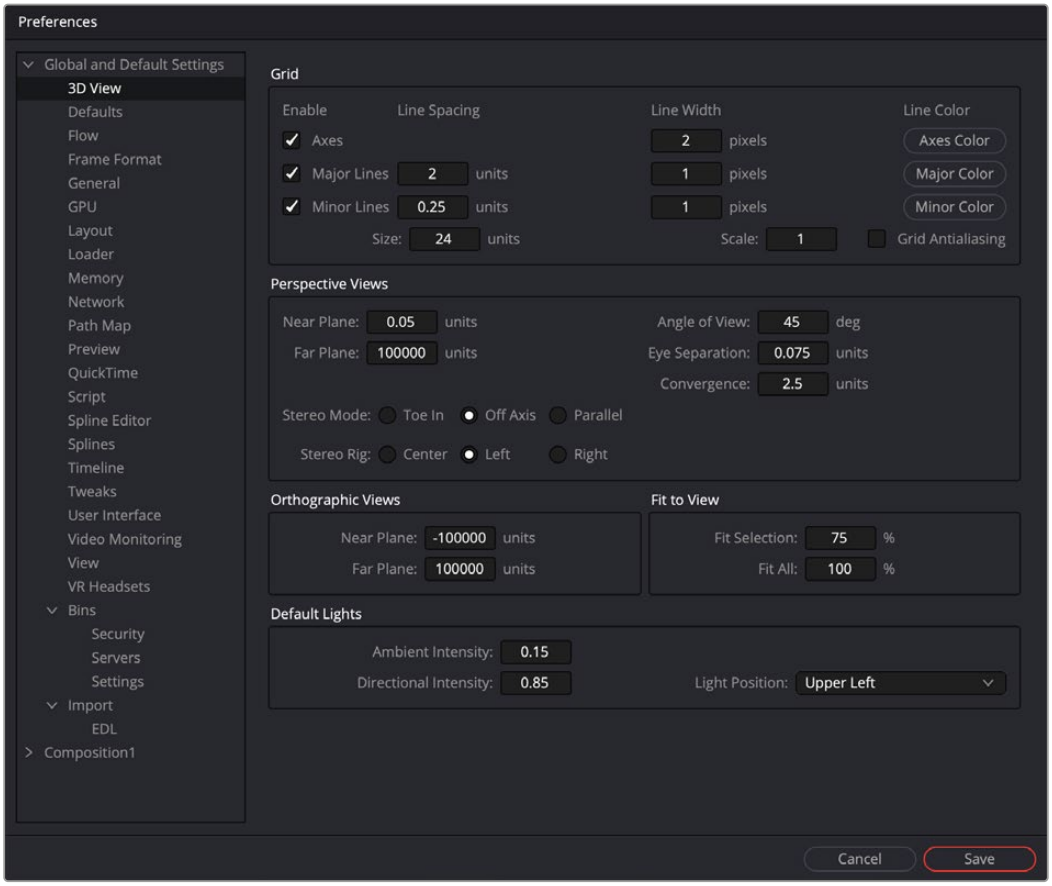
The Import settings contain options for EDL Import that affect how flows are built using the data from an EDL.

Preferences In Depth

Within each category is a deep set of controls for configuring Fusion so that it better fits your working environment. The preferences contain both software and hardware options that affect all newly created comps. The following section explains every option located in the preferences categories.

3D View

The 3D View preferences contain settings for various defaults in the 3D Viewers, including grids, default ambient light setup, and Stereoscopic views.



The 3D View preferences

Grid

The Grid section of the 3D View preferences configures how the grid in 3D Viewers are drawn.

- **Grid Antialiasing:** Some graphics hardware and drivers do not support antialiased grid lines, causing them to sort incorrectly in the 3D Viewer. Disabling this checkbox will disable antialiasing of the grid lines. To turn off the grid completely, right-click in a 3D Viewer and choose 3D Options > Grid.
- **Size:** Increasing the Size value will increase the number of grid lines drawn. The units used for the spacing between grid lines are not defined in Fusion. A “unit” is whatever you want it to be.
- **Scale:** Adjusting the overall scaling factor for the grid is useful, for example, if the area of the grid appears too small compared to the size of your geometry.

Perspective Views

The Perspective Views section handles the appearance of the perspective view in both a normal and stereoscopic project.

- **Near Plane/Far Plane:** These values set the nearest and furthest point any object can get to or from the camera before it is clipped. The minimum setting is 0.05. Setting Near Plane too low and Far Plane too far results in loss of depth precision in the viewer.
- **Eye Separation/Convergence/Stereo Mode:** This group of settings defines the defaults when stereo is turned on in the 3D Viewer.

Orthographic Views

Similar to the Perspective Views, the Orthographic Views (front, top, right, and left views) section sets the nearest and furthest point any object can get to or from the viewer before clipping occurs.

Fit to View

The Fit to View section has two value fields that manage how much empty space is left around objects in the viewer when the F key is pressed.

- **Fit Selection:** Fit Selection determines the empty space when one or more objects are selected and the F key is pressed.
- **Fit All:** Fit All determines the empty space when you press F with no objects selected.

Default Lights

These three settings control the default light setup in the 3D Viewer.

The default ambient light is used when lighting is turned on and you have not added a light to the scene. The directional light moves with the camera, so if the directional light is set to “upper left,” the light appears to come from the upper-left side of the image/camera.

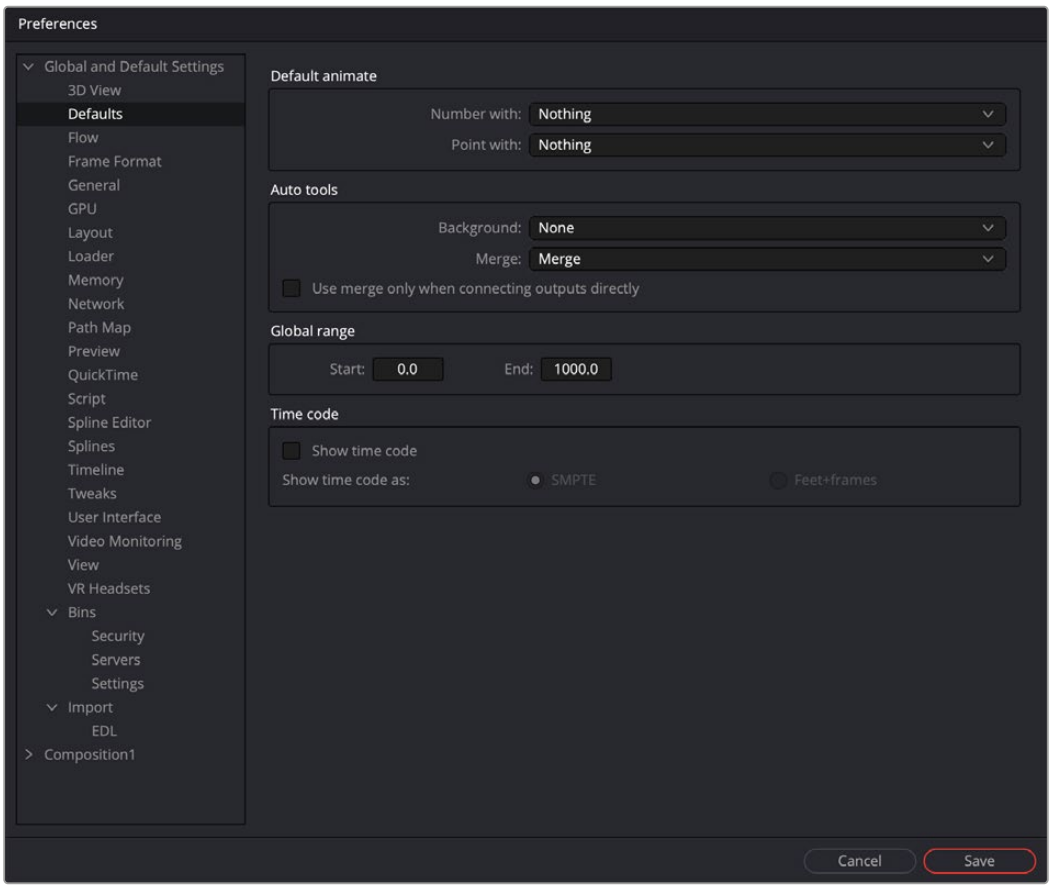
AVI

The AVI preference is only available in Fusion Studio on Windows. It configures the default AVI codec settings when you select AVI as the rendering file format in the Saver node.

- **Compressor:** This drop-down menu displays the AVI codecs available from your computer. Fusion tests each codec when the application opens; therefore, some codecs may not be available if the tests indicate that they are unsuitable for use within Fusion.
- **Quality:** This slider determines the amount of compression to be used by the codec. Higher values produce clearer images but larger files. Not all codecs support the Quality setting.
- **Key Frame Every X Frames:** When checked, the codec creates keyframes at specified intervals. Keyframes are not compressed in conjunction with previous frames and are, therefore, quicker to seek within the resulting movie. Not all codecs support the keyframe setting.
- **Limit Data Rate To X KB/Second:** When checked, the data rates of the rendered file are limited to the amount specified. Not all codecs support this option. Enter the data rate used to limit the AVI in kilobytes (kB) per second, if applicable. This control does not affect the file unless the Limit Data Rate To option is selected.

Defaults

The choices made here are used to determine Fusion's behavior when new tools are added to the Node Editor and when parameters are animated.



The Defaults preferences

Default Animate

The Default Animate section is used to change the type of modifier attached to a parameter when the Animate option is selected from its contextual menu. The default option is Nothing, which uses a Bézier spline to animate numeric parameters and a path modifier for positional controls.

- **Number With and Point With:** Drop-down lists are used to select a different modifier for the new default. For example, change the default type used to animate position by setting the Point with the drop-down menu to XY Path.

  Choices shown in this menu come from installed modifiers that are valid for that type of parameter. These include third-party plug-in modifiers, as well as native modifiers installed with Fusion.

Auto Tools

The Auto Tools section determines which tools are added automatically for the most common operations of the Background tools and Merge operations.

- **Background:** When set to None, a standard Background tool is used; however, the drop-down menu allows you to choose from a variety of tools including 2D and 3D tools to customize the operation to your workflow.

- **Merge:** When set to None, nothing happens. When set to Merge, connecting the outputs of two tools or dragging multiple clips on the Node Editor uses a standard Merge. Other valid options for this are Anaglyph, Channel Booleans, and Dissolve.
- **Use Merge Only When Connecting Outputs Directly:** When this option is active, Merges are not automatically added when you drag multiple clips from the Finder or Windows Explorer onto the Flow area.

Global Range

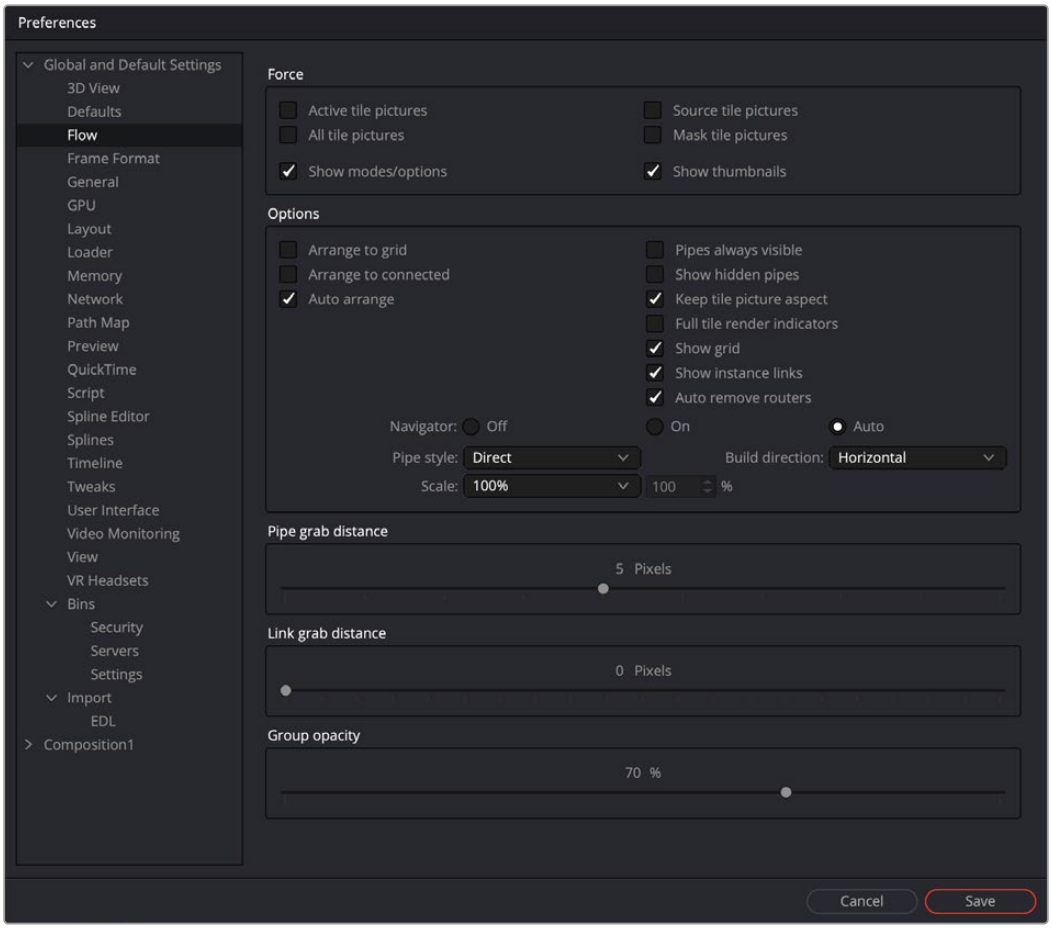
Using the Start and End fields, you can define the Global Start and End frames used when creating new compositions.

Time Code

You use this option to determine whether new compositions will default to showing SMPTE Time Code or frames (Feet + Frames) to represent time.

Flow

Many of the same options found in the Node Editor's contextual menu, like settings for Tile Picture, the Navigator, and Pipe Style, are found in this category.



The Flow preferences

Force

The Force section can set the default to display pictures in certain tool tiles in the Node Editor rather than showing plane tiles. The Active checkbox sets pictures for the actively selected tool, the All checkbox enables pictures for all tiles, and the Source and Mask checkbox enables tile pictures for just Source and Mask tools.

When All is enabled, the picture shown will either be a thumbnail of the image rendered by the tool if the tool has rendered, or if the Show Thumbnails option is disabled, the tool's default icon is used. Concatenated transforms will also show a default icon.

- **Show Modes/Options:** Enabling this option will display icons in the tool tile depicting various states, like Disk Caching or Locked.
- **Show Thumbnails:** When this checkbox is selected, tool tiles set to show tile pictures will display the rendered output of the tool. When the checkbox is cleared, the default icon for the tool is used instead.

Options

The Options section includes several settings that control or aid in the layout and alignment of tools in the Node Editor.

- **Arrange to Grid:** This enables a new node tree's Snap to Grid option to force the tool layout to align with the grid marks in the flow.
- **Arrange to Connected:** Tools snap to the vertical or horizontal positions of other tools they are connected to.
- **Auto Arrange:** This option enables the Node Editor to shift the position of tools as needed to make space when inserting new tools or auto-merging layers.
- **Show Grid:** This enables or disables the display of the Node Editor's background grid.
- **Auto Remove Routers:** Pipe Routers or "elbow nodes" in the Node Editor are considered "orphaned" if the tools connected to either the input or output are deleted. When this option is enabled, Orphaned Routers are automatically deleted.
- **Pipes Always Visible:** When enabled, the connection lines between tools are drawn over the top of the tool tiles.
- **Keep Tile Picture Aspect:** Enabling this option forces tool tile thumbnail pictures to preserve the aspect of the original image in the thumbnail.
- **Full Tile Render Indicators:** Enabling this checkbox causes the entire tile to change color when it is processing. This can make it easier to identify which tools are processing in a large composition. The coloring itself will form a progress bar to alert you to how close slower tools are to finishing their process.
- **Show Instance Links:** This option is used to select whether Instance tools will show links, displayed as green lines, between Instance tools.
- **Navigator:** The Navigator is a small square overview of the entire composition. It is used to quickly navigate to different parts of a node tree while you are zoomed in. The checkboxes in this section determine when the Navigator is displayed, if at all.

  On: The Navigator will always be visible.

  Off: The Navigator will always be hidden.

  Auto: The Navigator will only be visible when the Node Editor's contents exceed the currently visible Work area.

- **Pipe Style:** This drop-down menu selects which method is used to draw connections between tools. The Direct method uses a straight line between tools, and Orthogonal uses horizontal and vertical lines.

- **Build Direction:** When auto-building or laying out a node tree, Build Direction controls whether tools are organized horizontally or vertically.
- **Scale:** The Scale menu allows you to select the default zoom level of the Node Editor when a new composition is created.

Pipe Grab Distance

The Pipe Grab Distance slider allows you to choose how close the pointer must be (in pixels) to a connection line in the node tree when selecting them.

Link Grab Distance

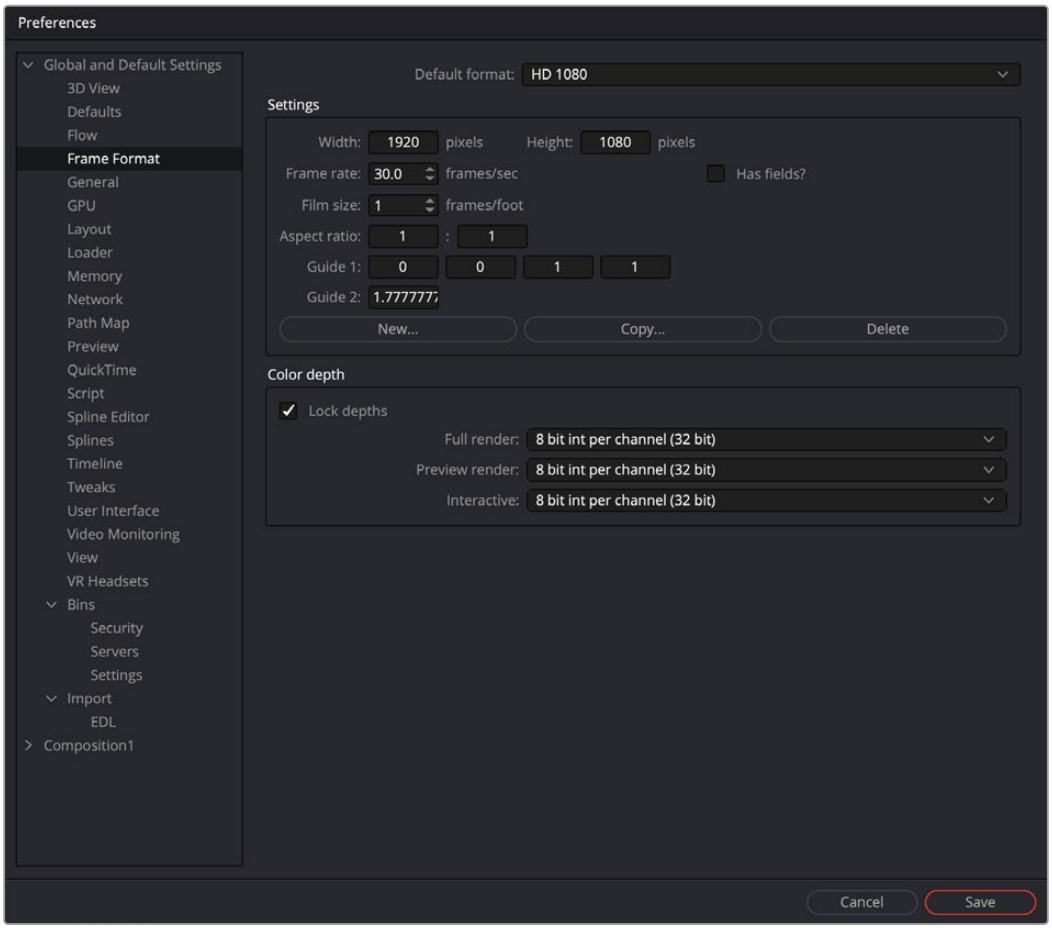
The Link Grab Distance slider allows you to choose how close the pointer must be (in pixels) to a knot on a node before a connection is made or removed. 0 pixels means you must be directly on the knot, while a maximum setting of 20 means you can be 20 pixels from a knot and still connect to it.

Group Opacity

This slider controls the opacity of an expanded group's background in the Node Editor.

Frame Format

Frame Format preferences allow you to select the resolution and frame rate for the nodes that generate images like Background, fast noise, and Text+. It also sets the color bit depth for final renders, previews, and interactive updates in the viewer. The color bit depth settings only apply to Fusion Studio. Rendering in DaVinci Resolve always use 32-bit float.



The Frame Format preferences

Default Format

This drop-down menu is used to select the default resolution for Generator tools from a list of presets. This is only a default setting; these settings can be overridden using the Resolution settings in a node's Inspector.

Use the Edit boxes to change any of the default settings. When creating a new setting, press the New button and enter a name for the setting in the dialog box that appears and enter the parameters.

Settings

The Settings section defines the format that is selected in the Default Format menu. You can modify an existing format or create a new one.

- **Width/Height:** When creating a new format for the menu or modifying an existing menu item, you specify the Width or Height in pixels of the format using these fields.
- **Frame Rate:** Enter or view the frames per second played by the format. This sets the default Frame Rate for previews and final renders from the Saver tool. It also sets the playback for the comp itself, as well as the frame to time code conversion for tools with temporal inputs.
- **Has Fields:** When this checkbox is enabled, any Creator or Loader tool added to the Node Editor will be in Fields process mode.
- **Film Size:** This field is used to define how many frames are found in one foot of film. The value is used to calculate the display of time code in Feet + Frames mode.
- **Aspect Ratio:** These two fields set the pixel aspect ratio of the chosen frame format.
- **Guide 1:** The four fields for Guide 1 define the left, top, right, and bottom guide positions for the custom guides in the viewer. To change the position of a guide, enter a value from 0 to 1. The bottom-left corner is always 0/0, the top-right corner is always 1/1. If the entered value's aspect does not conform to the frame format as defined by the Width and Height parameters, an additional guide is displayed onscreen. The dotted line represents the image aspect centered about Guide 1's Center values.
- **Guide 2:** This setting determines the image aspect ratio in respect to the entire frame format width and height. Values higher than 1 cause the height to decrease relative to the width. Values smaller than 1 cause height to increase relative to width.
- **New:** You use the New button to create a new default setting in the drop-down menu. Once you click the button, you can name the setting in the dialog box that appears.
- **Copy:** The Copy button copies the current setting to create a new one for customization.
- **Delete:** The Delete button will remove the current setting from the default drop-down list.

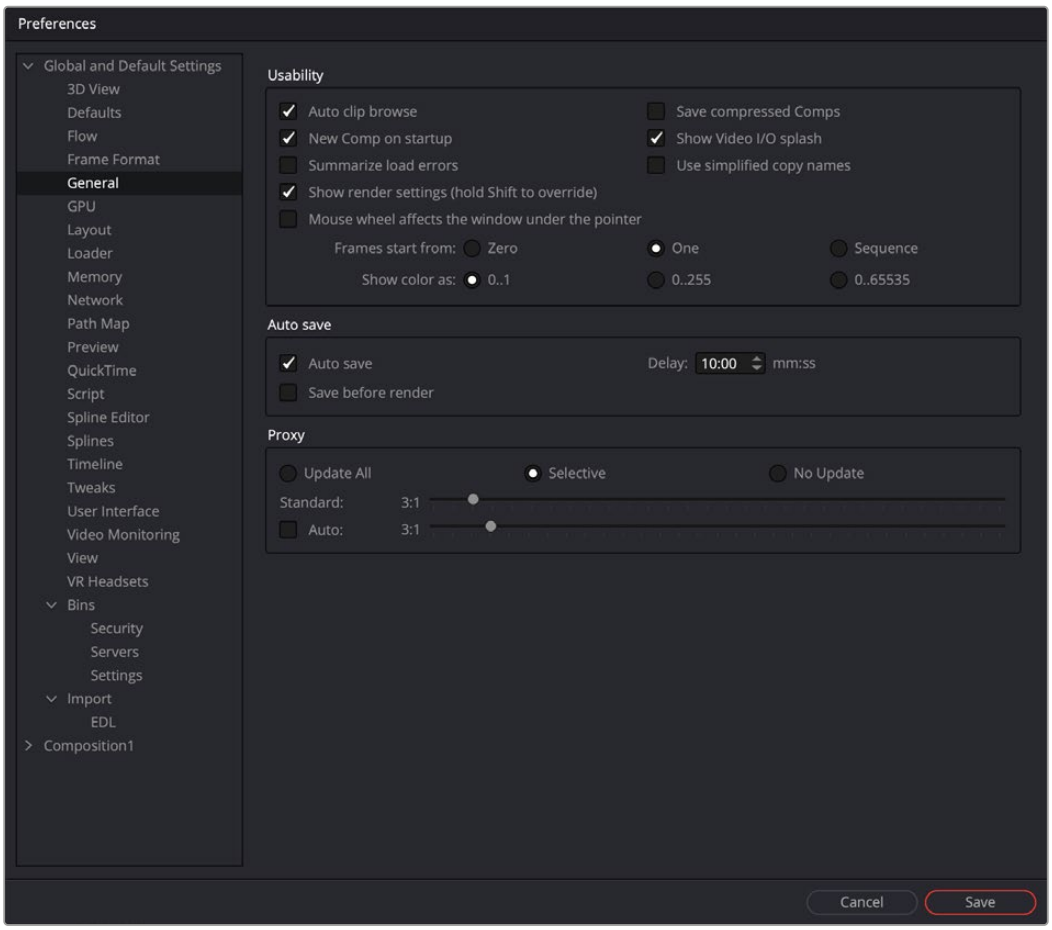
Color Depth

The three menus in the Color Depth section are used to select the color mode for processing preview renders, interactive renders, and full (final) renders. Processing images at 8-bit is the lowest color depth and is rarely sufficient for final work these days but is acceptable for fast previews. 16-bit color has much higher color fidelity but uses more system resources. 16-bit and 32-bit float per channel uses even more system resources and is best for digital film and HDR rendered images.

Generally, these options are ignored by the composition unless a Loader or Creator tool's Color Depth control is set to Default.

General

The sections contained in the General preferences affect the behavior of the Inspector as well as some other user interface elements.



The General preferences

Usability

Usability has a number of project, Node Editor, and user interface settings that can make the application easier to work with, depending on your workflow.

- **Auto Clip Browse:** When this checkbox is enabled, the File Browser is automatically displayed when a new Loader or Saver is added to the Node Editor.
- **New Comp on Startup:** When checked, a new, empty project is created each time Fusion Studio is launched. This has no effect in DaVinci Resolve's Fusion page.
- **Summarize Load Errors:** When loading node trees or "comps" that contain unknown tools (e.g., comps that have been created on other computers with plugins not installed on the current machine), the missing tools are summarized in the console rather than a dialog being presented for every missing tool.
- **Save Compressed Comps:** This option enables the saving of compressed node trees, rather than ASCII based text files. Compressed node trees take up less space on disk, although they may take a moment longer to load. Node trees containing complex spline animation and many paint strokes can grow into tens of megabytes when this option is disabled. However, compressed comps cannot be edited with a text editor unless saved again as uncompressed.

- **Show Video I/O Splash:** This toggles whether the Splash image will be displayed over the video display hardware. This is only applies to Fusion Studio.
- **Use Simplified Copy Names:** This option reduces the occurrence of underscores in tool names when copying.
- **Show Render Settings:** When this checkbox is selected, the Fusion Render Settings dialog will be displayed every time a render is started in Fusion Studio. Holding Shift while starting a render will prevent the display of the dialog for that session, using whatever settings were applied during the last render. Disabling this option reverses this behavior.
- **Mouse Wheel Affects the Window Under the Pointer:** Normally the mouse wheel or trackpad swiping works in the currently active window. With this option enabled, it will work in the window underneath the cursor, so you don't have to click into a window first to make it active.
- **Frames Start From:** This designates the starting frame number for clip times in the Loader and its Clip list.
- **Show Color As:** This setting determines the numeric scale used to represent colors. The available options are Normalized (0 to 1), 8-bit (0 to 255), and 16-bit (0 to 65,535). This does not affect the actual processing or quality of the image, but it can make the mental math sometimes used to figure out adjustments a bit easier.

Auto Save

The Auto Save settings only apply to Fusion Studio. To set auto backups for the Fusion page in DaVinci Resolve, use the DaVinci Resolve Project Load and Save Preferences.

When Auto Save is enabled in Fusion Studio, comps are automatically saved to a backup file at regular intervals defined by the Delay setting. If a backup file is found when attempting to open the comp, you are presented with the choice of loading either the backup or the original.

If the backup comp is opened from the location set in the Path Map preference, saving the backup will overwrite the original file. If the backup file is closed without saving, it is deleted without affecting the original file.

- **Save Before Render:** When enabled, the comp is automatically saved before a preview or final render is started.
- **Delay:** This preference is used to set the interval between Auto Saves. The interval is set using mm:ss notation, so entering 10 causes an Auto Save to occur every 10 seconds, whereas entering 10:00 causes an Auto Save every 10 minutes.

Proxy

- **Update All, Selective, No Update:** The Update mode button is located above the toolbar. You can use this preference to determine the default mode for all new comps. Selective is the usual default. It renders only the tools needed to display the images in the Display view. All will render all tools in the composition, whereas None prevents all rendering.
- **Standard and Auto:** These sliders designate the default ratio used to create proxies when the Proxy and Auto Proxy modes are turned on. These settings do not affect the final render quality.

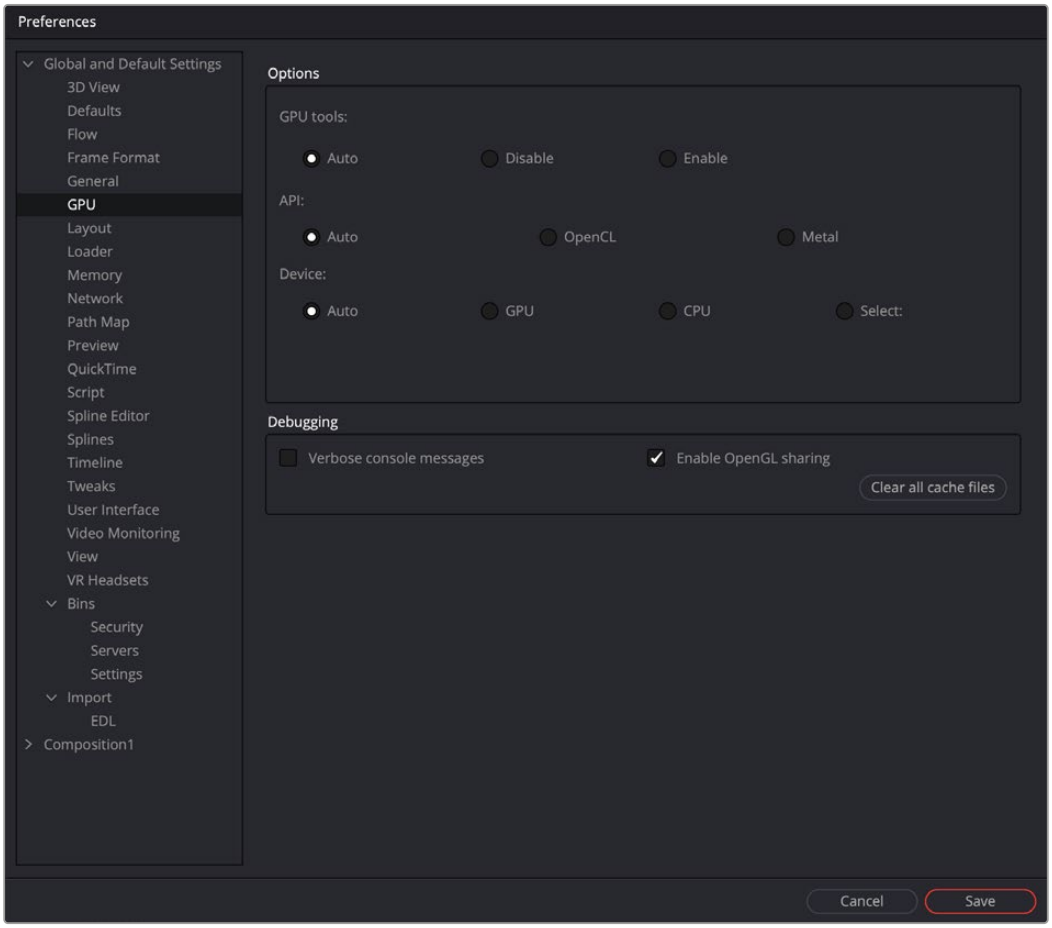
Even though the images are being processed smaller than their original size, the image viewing scales in the viewers still refer to original resolutions. Additionally, image processing performed in Proxy Scale mode may differ slightly from full-resolution rendering.

The Proxy and Auto Proxy size ratios may be changed from within the interface itself by right-clicking on the Prx and APrx buttons above the toolbar and selecting the desired value from the contextual menu.

GPU

The GPU preference is only available in Fusion Studio. In DaVinci Resolve, you can configure the GPU processing in Resolve's Memory and GPU preferences.

In Fusion Studio, the GPU preference is used to specify the GPU acceleration method used for processing, based on your computer platform and hardware capabilities. It is also used for enabling caching and debugging GPU devices and tools.



The GPU preferences

Options

The GPU options include radio buttons to select whether the GPU is used when processing and, if so, which computer framework is used for communicating with the GPU.

- **GPU Tools:** This preference has three settings: Auto, Disable, and Enable. When set to Disable, no GPU acceleration is used for tools or third-party plugins. Fuses may still require GPU acceleration. If Enable is selected, GPU acceleration is available for tools and plugins, if appropriate drivers are installed.
- **API:** The API setting selects the GPU processing method to use.
- **Device:** The Device setting determines which GPU hardware to use in the case of multiple GPUs. The Auto setting gives priority to GPU processing; however, if it is unavailable, Fusion uses the platform default. Currently, both the AMD and CPU options require either the AMD Catalyst 10.10 Accelerated Parallel Processing (APP) technology Edition driver or the ATI Stream SDK 2.1 or later to be installed. The Select setting allows you to choose the device explicitly.

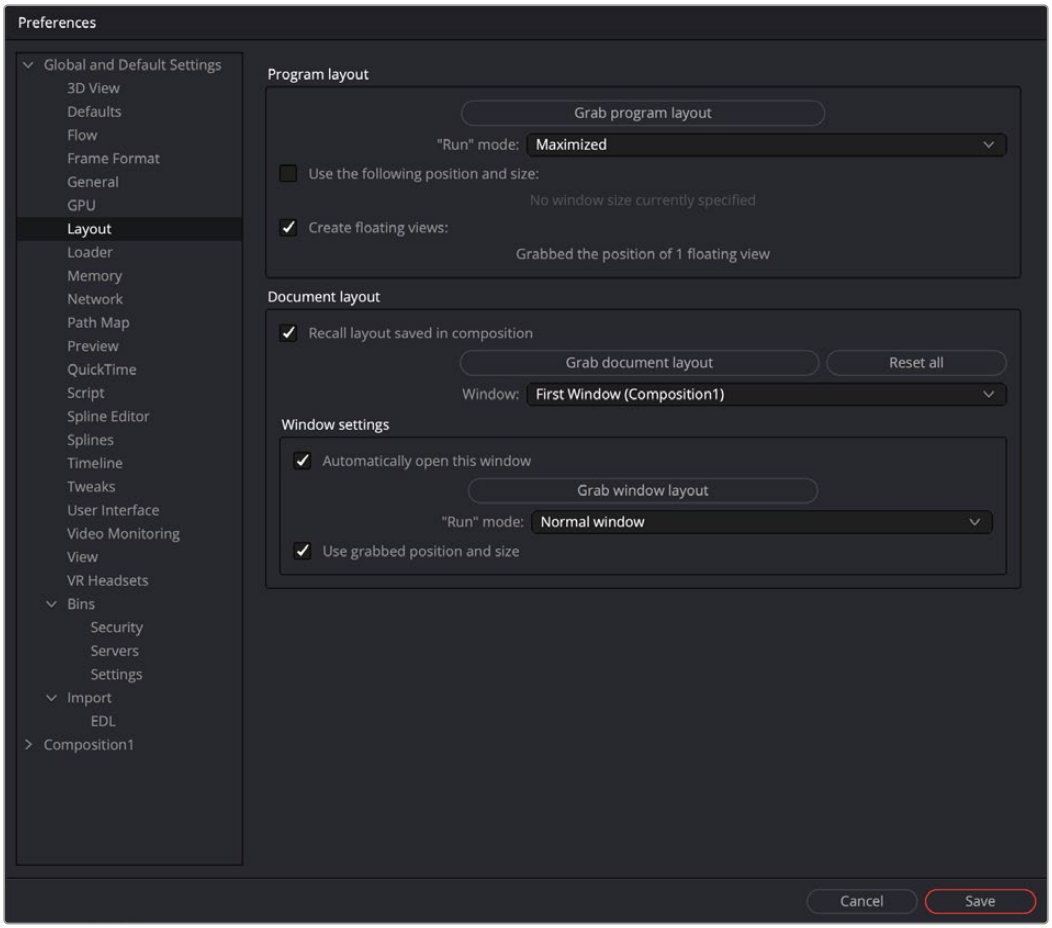
Debugging

The more advanced preferences located in this section are designed for diagnostics and analyzing GPU operations.

- **Verbose Console Messages:** Enabling this option causes information to be shown in the Console. For example, Startup Logs, Compiler Warnings, and Messages.
- **OpenGL Sharing:** Enabling this option shares system RAM with onboard GPU RAM to create a larger, but slower, OpenGL memory pool.
- **Clear Cache Files:** This option will clear already compiled GPU code and then recompile the kernels.

Layout

The Layout preferences are only available in Fusion Studio. To save a Layout in DaVinci Resolve's Fusion page, use the Workspace > Layout Presets menu. The Layout options are used to control the layout, size, and position of various windows in Fusion's interface at startup or when a comp is created.



The Layout preferences

There are a lot of options, but in practice, you simply organize the interface the way you prefer it on startup and when a new composition is created, then open this Preferences panel and click on the three buttons to grab the Program Layout, the Document Layout and the Window Settings.

Program Layout

The Program Layout is used to save the overall Fusion interface window and any open floating windows. Each new composition you open within the larger overall Fusion interface window will adhere to these preferences.

- **Grab Program Layout:** Pressing this button stores the application's overall current position and size.
- **Run Mode:** This menu is used to select the application's default mode at startup. You choose between a Maximized application window, a Minimized application, or a Normal application display.
- **Use the Following Position and Size:** When checked, the values stored when Grab Program Layout was selected will be used when starting Fusion Studio.
- **Create Floating Views:** When checked, the position and size of the floating viewers will be saved when the Grab Program Layout button is used.

Document Layout

The Document Layout is used to save the layout of panels and windows for the current Fusion comp.

- **Recall Layout Saved In Composition:** When checked, all Document Layout settings in the controls below will be recalled when a saved composition is loaded.
- **Grab Document Layout:** Pressing this button stores the entire interface setup, including all the internal positions and sizes of panels and work areas.
- **Window:** When multiple windows on the same composition are used, this menu is used to select the window to which the Window Settings will apply.

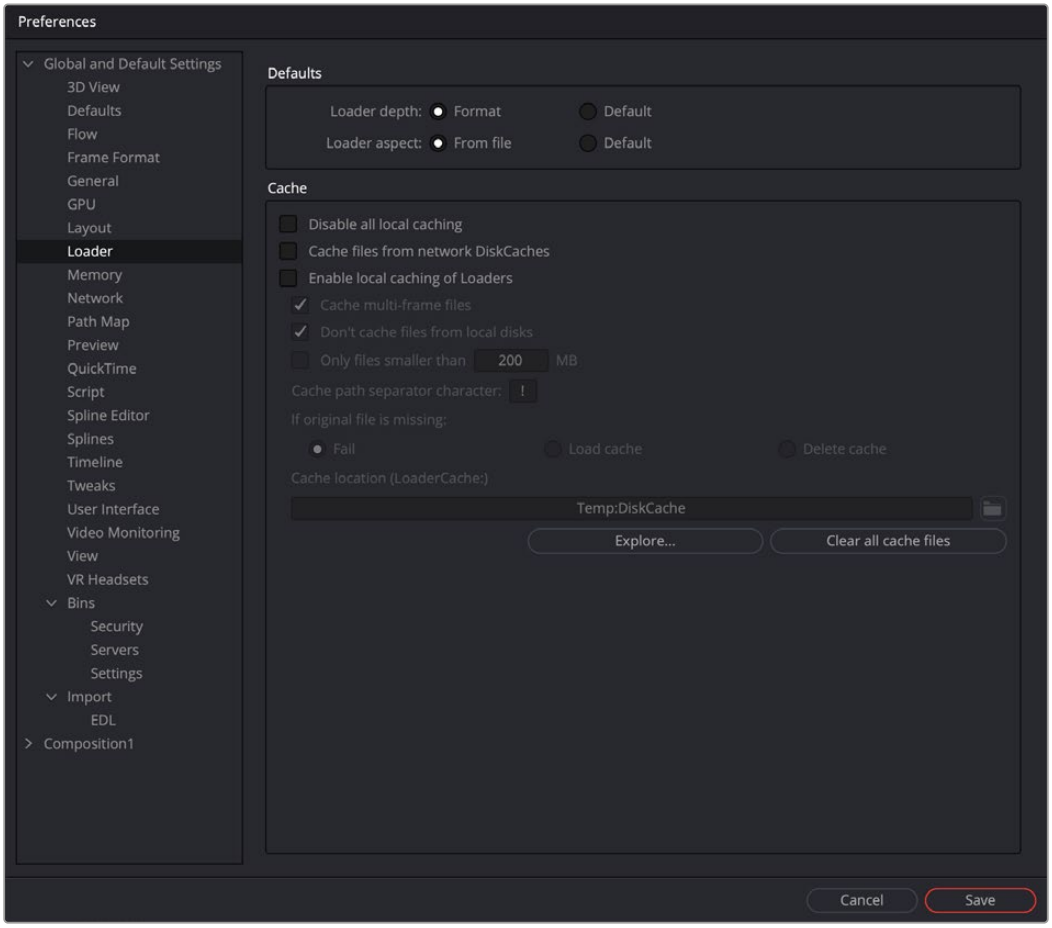
Window Settings

Rather than saving entire comp layouts, you can save position and size for individual floating windows and panels within a comp using the Window Settings.

- **Automatically Open This Window:** When checked, the selected window will automatically be opened for new flows.
- **Grab Window Layout:** Pressing this button stores the size and position of the selected window.
- **Run Mode:** Select the default run mode for the selected window. You can choose between a Maximized window, a Minimized window, or a Normal window display.
- **Use Grabbed Position and Size:** When checked, the selected window will be created using the stored position and size.

Loader

The Loader preferences are only available in Fusion Studio. Using the Loader preferences, you can set options for the default Loader's color depth and aspect ratio as well as define the local and network cache settings.



The Loader preferences

Defaults

The Defaults section includes two settings to determine how color depth and aspect ratio are handled for Loaders.

- **Loader Depth:** The Loader Depth defines how color bit depth is handled when adding a Loader. Choosing Format means that the correct bit depth is automatically selected, depending on the file format and the information in the file's header. Choosing Default sets the bit depth to the value specified in the Frame Format preferences.

Cache

The Cache preferences allow you to control how disk caching operates in Fusion. You can set how and where the cache is generated, when the cache is removed, how the cache reacts when source files are not available, as well as many other cache related options. This is not to be confused with RAM cache, which is controlled in the Memory preferences.

- **Disable All Local Caching:** This setting disables local caching.
- **Cache Files from Network DiskCaches:** If a tool has disk caching enabled, and the disk cache files are stored remotely on the network, then enabling this option will use a local copy of those cache files, similarly to the local cache on a networked Loader.
- **Enable Local Caching of Loaders:** Files will be copied into the LoaderCache path set below or in the Path Maps preferences.

- **Cache Multi-Frame Files:** Files like AVI or QuickTime will be copied into the LoaderCache path. This may take some time if the file is large.
- **Don't Cache Files from Local Disks:** Files that do not sit on a network drive will not be copied into the LoaderCache path. You can disable this option if you have, for example, a fast SSD cache drive and want to use it for local files as well to speed up file access while working interactively.
- **Only Files Smaller Than xxx MB.:** Files larger than the value set here will not be copied into the LoaderCache path.
- **Cache Path Separator Character:** When Enable Local Caching of Loaders is enabled, you can use this setting to rebuild the path of the original files in LoaderCache.

  For instance, given the default "!" character, the original path X\Project\MyShots\ Shot0815\ will be translated into X!Project!MyShots!Shot0815! in the LoaderCache path. Other separator characters may be used, including the "\" character, which will use subdirectories in LoaderCache: X\Project\ MyShots\Shot0815\.
- **If Original File Is Missing:** This setting provides three options to determine the caching behavior when the original files can't be found. The Fail option behaves exactly as the Default Loader in Fusion. The Loader will not process, which may cause the render to halt. The Load Cache option loads the cache even though no original file is present.The Delete Cache option clears missing files from the cache.
- **Cache Location:** For convenience, this is a copy of the LoaderCache path set in the Path Maps preferences.
- **Explore:** This button opens the LoaderCache path in the macOS X Finder window or a Windows Explorer window.
- **Clear All Cache Files:** This button deletes all cached files present in the LoaderCache path.

Memory

The Memory preferences are only available in Fusion Studio. To control Fusion's memory when using the Fusion page in DaVinci Resolve, open DaVinci Resolve's Memory and GPU preferences.

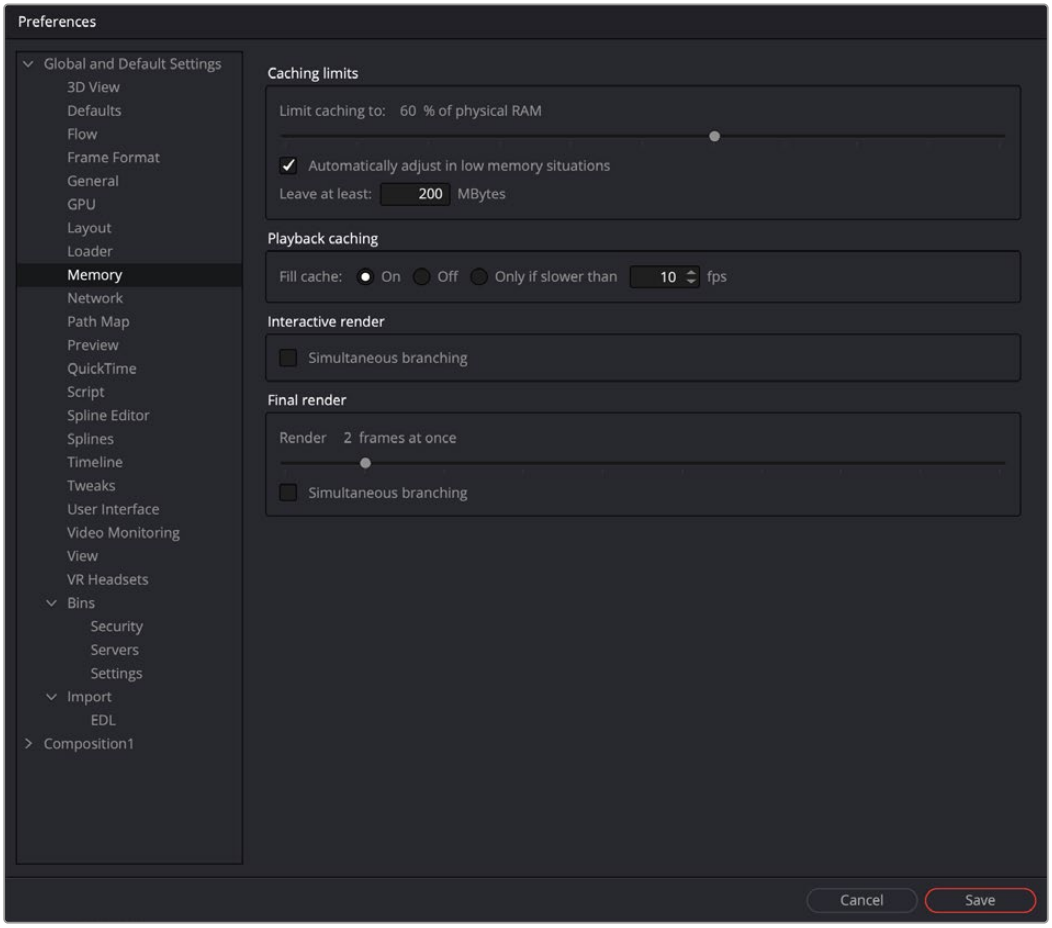
Occasionally, it will be necessary to adjust the Memory preferences in order to make the best use of available memory on the computer. For example, some people prefer a higher cache memory for faster interactive work, but for final renders the cache memory is often reduced, so there's more memory available for simultaneous processing of tools or multiple frames being rendered at once.

Caching Limits

The Caching Limits include options for Fusion's RAM cache operation. Here, you can determine how much RAM is allocated to the RAM cache for playing back comps in the viewer.

- **Limit Caching To:** This slider is used to set the percentage of available memory used for the interactive tool cache. Available memory refers to the amount of memory installed in the computer.

  When the interactive cache reaches the limit defined in this setting, it starts to remove lower priority frames in the cache to clear space for new frames.
- **Automatically Adjust In Low Memory Situations:** This checkbox will set the caching to adjust when memory is low. The console will display any cache purges.
- **Leave At Least X MBytes:** This setting is used to set the hard limit for memory usage. No matter what the setting of the Cache Limit, this setting determines the amount of physical memory available for use by other applications. Normally, this value should not be smaller than 25 MBytes.



The Memory preferences

Interactive Render

The Interactive Render option allows you to optimize Fusion's processing based on the amount of RAM you have installed in your system.

- **Simultaneous Branching:** When checked, more than one tool will be processed at the same time. Disable this checkbox if you are running out of memory frequently.

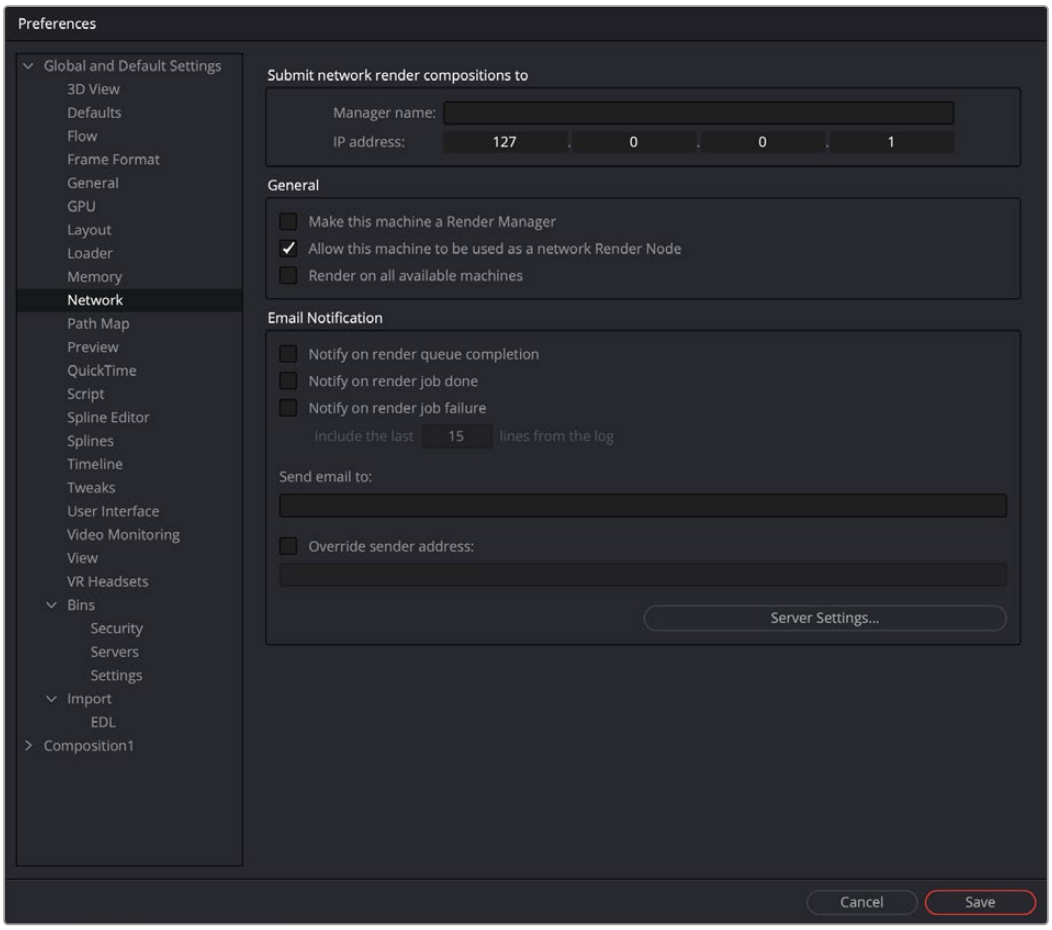
Final Render

These settings apply to memory usage during a rendering session, either preview or final, with no effect during an interactive session.

- **Render Slider:** This slider adjusts the number of frames that are rendered at the same time.
- **Simultaneous Branching:** When checked, more than one branch of a node tree will be rendered at the same time. If you are running low on memory, turn this off to increase rendering performance.

Network

The Network preferences are only available in Fusion Studio. These preferences are used to set up and control network rendering in Fusion Studio. The majority of settings are found in the Render Manager dialog.



The Network preferences

Submit Network Render Compositions

In these fields, you enter the Master Name and IP address of the computer that will manage all network renders sent from this machine. If a standalone render master is in use on the network, these fields may be pre-filled and may not be editable. This is done to prevent multiple unauthorized render masters from being created by each person in a facility.

> To re-enable editing of the master name and IP, create the environment variable FUSION_ NO_ MANAGER and set the value to True. Check your operating system user guide for how to create environment variables.

General

The General preferences are designed with the most used options at the top in the General section. These options determine in what capacity the system is used during network rendering.

- **Make This Machine a Render Master:** When enabled, Fusion will accept network render compositions from other computers and manage the render. It does not necessarily mean that this computer will be directly involved in the render, but it will submit the job to the render nodes listed in the Render Manager dialog.
- **Allows This Machine to Be Used as a Network Slave:** When enabled, this computer can be used as a Render node and will accept compositions for network rendering. Deselect it to prevent other people from submitting compositions to render on this computer.
- **Render on All Available Machines:** Enable this checkbox to ignore groups and priorities configured in the Render Manager. Compositions submitted from this computer for network rendering will always be assigned to every available slave.

Email Notification

You can use the Email Notification section to set up who gets notified with status updates regarding the render jobs and the network.

- **Notify Options:** These checkboxes cause emails to be sent when certain render events take place. The available events are Queue Completion, Job Done, and Job Failure.
- **Send Email to:** Enter the address or addresses to which notifications should be sent. You separate multiple addresses with a semicolon.
- **Override Sender Address:** Enter an email address that will be used as the sender address. If this option is not selected, no sender address is used, which may cause some spam filters to prevent the message from being delivered to the recipient.

Server Settings

This section covers Clustering and Network Rendering.

> For more information on these settings and clustering, *see Chapter 4, "Rendering Using Saver Nodes,"* in the Fusion Reference Manual.

Path Maps

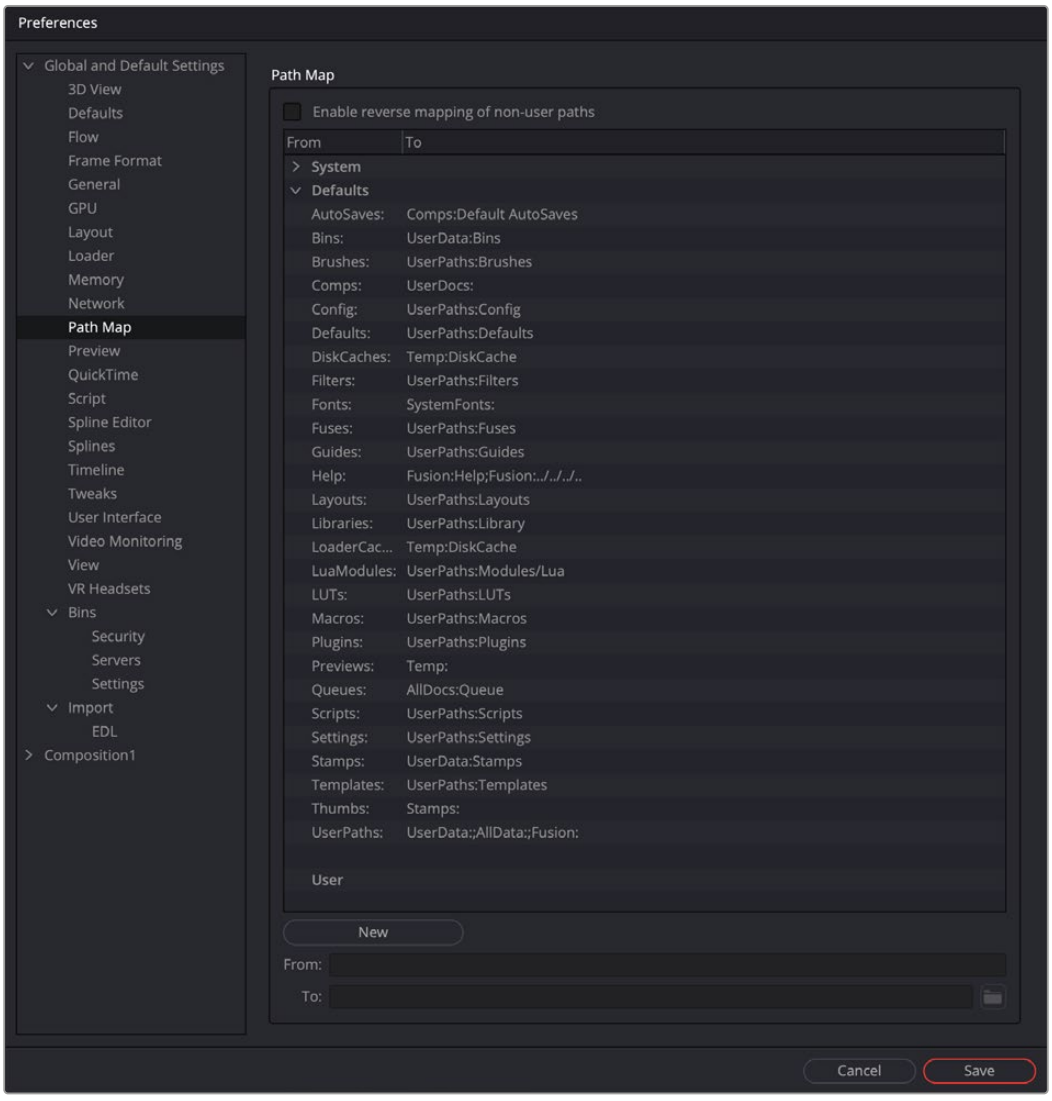
Path Maps are virtual paths used to replace segments of file paths with variables. For example, define the path 'movie_x' as actually being in X\Shows\Movie_X. Using this example, Fusion would understand the path 'movie_x\scene_5\ scan.000.cin' as actually being X:\Shows\ Movie_X\scene_5\scan.000.cin.

For Fusion Studio, there are two main advantages to virtual path maps instead of actual file paths. One is that you can easily change the path to media connected to Loaders (for example, when moving a comp from one drive to another), without needing to make any changes in the composition. The other advantage is when network rendering, you can bypass the different OS filename conventions.

- **Enable Reverse Mapping of Paths Preferences:** This checkbox is at the top of the Path Map settings. When enabled, Fusion uses the built-in path maps for entries in the path's settings when applying mapping to existing filenames. The main benefit is for Fusion Studio. Enabling this checkbox causes Loaders to automatically use paths relative to the location of the saved composition when they are added to the Node Editor.

For more information on using relative paths for Loaders, *see Chapter 45, "I/O Nodes,"* in the Fusion Reference Manual.

As with other preferences in Fusion Studio, paths maps are available in both Global and Composition preferences. Global preferences are applied to all new compositions, while Composition path maps are only saved with the current composition. Composition path maps will override Global path maps with the same name.



The Path Map preferences

The Global paths maps are divided into three sections:

- **System Path Maps:** The operating system determines system path maps, and they define Fusion's global locations. You can override specific System path maps using the Defaults or current Composition Path Map settings. If you change your mind at a later time, you are always able to return to Fusion's "factory" defaults using the System path maps. There are several top-level path maps established in the System Path Map settings.

  AllData: The folder where Fusion saves all shared application data.

AllDocs: The folder where Fusion saves the public/shared document folder.

AllLUTs: The nested LUTs path in the Defaults section, where Fusion saves LUTs.

Fusion: The folder where Fusion Studio app is installed. For example, if you open Fusion from C:\Program Files\Fusion, then the path Fusion:\Help refers to C:\Program Files\Fusion\Help. If you instead used a copy of Fusion found in \\post-server\fusion\16, then Fusion:\Help would expand to \\post-server\fusion\16\Help.

FusionLibs: The Fusion libraries used for the application.

Profile: The folder where default Fusion preferences file is saved.

Profiles: The folder where Fusion individual user preferences are saved.

Programs: The location of Fusion Studio or DaVinci Resolve.

SystemFonts: The folder where the OS saves fonts that appear for Text+ and Text 3D nodes.

Temp: The system's temporary folder.

UserData: The folder where Fusion saves all user-specific miscellaneous roaming data. The individual elements included in the roaming data are listed in the Default Path Maps section. For Fusion Studio on Windows, this is "C:\Users\username\AppData\Roaming\". On Linux, this will be "$HOME/Blackmagic/Fusion". On macOS, this is "Users/UserName/Library/Application Support/Blackmagic Design/Fusion".

UserDocs: The folder where Fusion saves the user's document folders.

- **Default Path Maps:** The Defaults are user-editable path maps. They can reference the System paths, as part of their paths. For instance. the Temp folder is defined in the System path and used by the Default DiskCache path map to refine the nested location (Temp:DiskCache). Default path maps can also redirect paths without using the Global System path maps. After you change a Default, the updated setting can be selected in the Preferences window, and a Reset button at the bottom of the Preferences window will return the modified setting to the System default.

  AutoSaves: This setting determines the Fusion Comp AutoSave document's location, set in the Fusion General preferences.

  Bins: Sets the location of Fusion Studio bins. Since the bins use pointers to the content, the content is not saved with the bin. Only the metadata and pointers are saved in the bins.

  Brushes: Points Fusion to the folder that contains custom paintbrushes.

  Comps: The folder where Fusion Studio compositions are saved. On macOS or Windows, the default location is in Users/YourUserName/Documents/Blackmagic Design/Fusion.

  Config: Stores Configuration files used by Fusion Studio during its operation.

  Defaults: Identifies the location of node default settings so they can be restored if overwritten.

  DiskCache: Sets the location for files written to disk when using the Cache to Disk feature. This location can be overridden in the Cache to Disk window.

  Edit templates: The location where Fusion macros are saved in order to appear as templates in the DaVinci Resolve Effects Library.

  Filters: Points to a folder containing Convolution filters like sharpen, which can be used for the Custom Filter node.

  Fonts: The default path map for Fonts points to the operating system fonts folders. Changing this will change the fonts that are available in the Text+ or Text 3D nodes as well as any Fusion Title Template. In DaVinci Resolve. This path map does not affect the five additional Edit page titles (L Lower 3rd, R Lower 3rd, M Lower 3rd, Scroll, and Text.)

  Fuses: Points to a folder containing Fusion Fuses plugins.

  FusionTemplates: Location where Fusion macros are saved in order to appear as templates in Fusion's Effects Library.

Guides: Location where custom viewer guide overlays are stored.

Help: Identifies where Fusion Studio PDF files are located.

Layouts: Location where Fusion Studio custom window layouts are saved.

Libraries: Points to a support folder where custom Effects Library items can be stored.

LoaderCache: The Fusion Studio Loader preferences allow the Loader to cache when reading from a slow network. This path map point to the local drive location for that cache.

LuaModules: Location for Lua Scripting modules.

LUTs: Points to a folder containing Look Up Tables (LUTs).

Macros: Points to the location for user created macros. The macros saved to this location appear in the macros category of the Effects Library and in the right-click Edit Macro contextual menu.

Plugins: This refers to user specific OpenFX plugins that you do not want loaded for all users.

Previews: Path map used for the older style, file sequence flipbook previews.

Queues: Location of the Render manager list.

Scripts: Location of Lua and Python scripts. This path can be further refined into specific scripts for tools (nodes), comps, and other specific script types.

Settings: Location where custom Node settings are saved.

Stamps: Location for preview movies generated in a Fusion Studio bin. This is an outdated path map since bins now include the Studio Player.

Templates: Location of the Templates folder. Saving Macros to the Template folder will cause them to appear in the Effects Library in a Templates category. In Fusion Studio, the Templates category does not appear until a Macros is saved into the folder.

Thumbs: Location for clip thumbnails generated in a Fusion Studio bin. This is an outdated path map since bins include now include the Studio Player.

UserPaths: Used for locations of studio- or facility-specific tools, like custom plugins and scripts located on a central server.

- **User Path Maps:** User paths are new paths that you have defined that do not currently exist in the Defaults settings.

  Comp refers to the folder where the current composition is saved. For instance, saving media folders within the same folder as your Fusion Studio comp file is a way to use relative file paths for Loaders instead of actual file paths.

Modifying a System Path Map

To modify an existing System path map, select the path map in the System section. Click the folder icon at the bottom of the Preferences window, and enter the name of the path map in the From field below. Enter the value of the path map in the To: field.

Modifying a Default Path Map

To modify an existing Default path map, select the path map in the Default section. Click the folder icon at the bottom of the Preferences window, and enter the name of the path map in the From field below. Enter the value of the path map in the To: field.

Creating a User Path Map

To create a path map, click on the New button and enter the name of the path map in the From field below. Enter the value of the path map in the To: field.

Deleting a Path Map

To delete a user-created path map, select it from the list and click the Delete button. System and Default path maps cannot be deleted; only user created path maps can be removed from the Path Maps list.

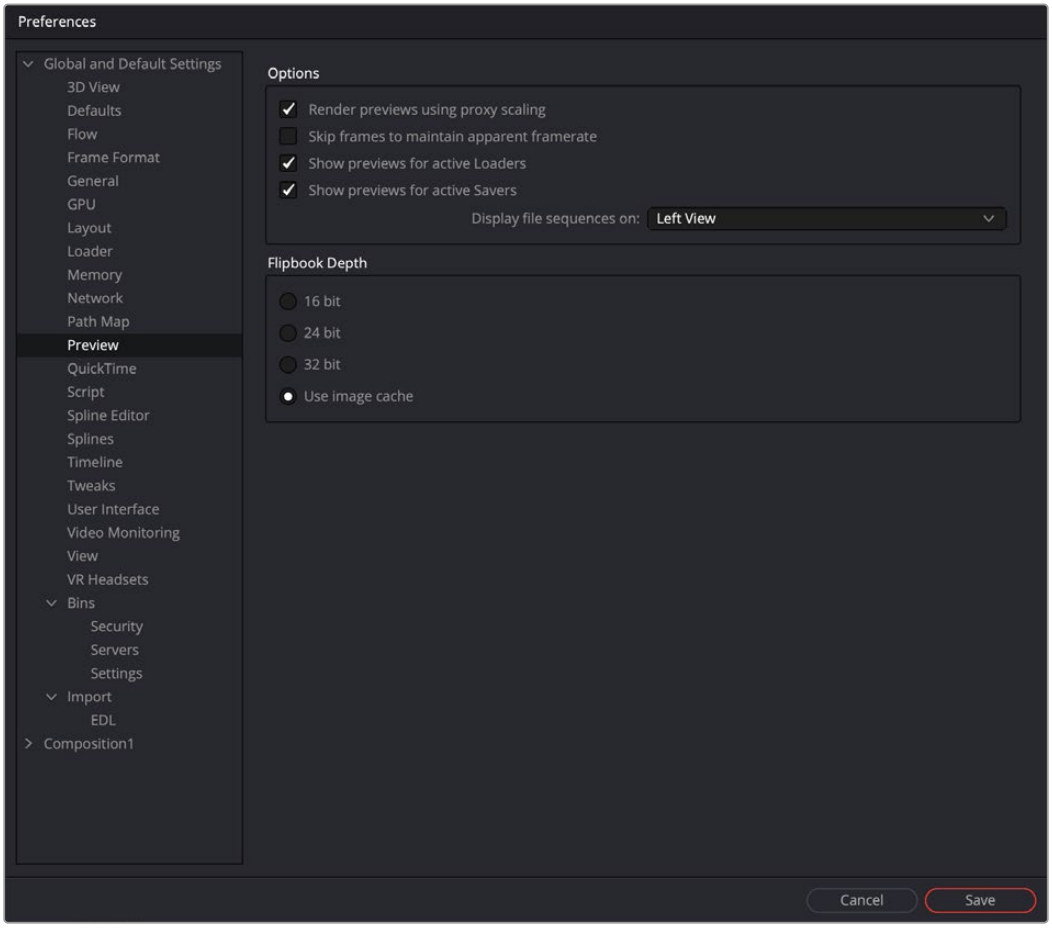
Nesting Path Maps

When defining your own path map, you can use an existing path map in the new definition. For example, define a path map called ‘Episode’ that maps to MyDrive\ Projects\Episode1. Then create new path maps called Renders and Stills that map to Episode\ Renders_v1 and Episode\Stills_v1.

Preview

Preview is only available in Fusion Studio. Previews in DaVinci Resolve use the Scratch Disk setting in the Media Storage preferences.

In the Preview preferences, you configure the creation and playback options for preview renders.



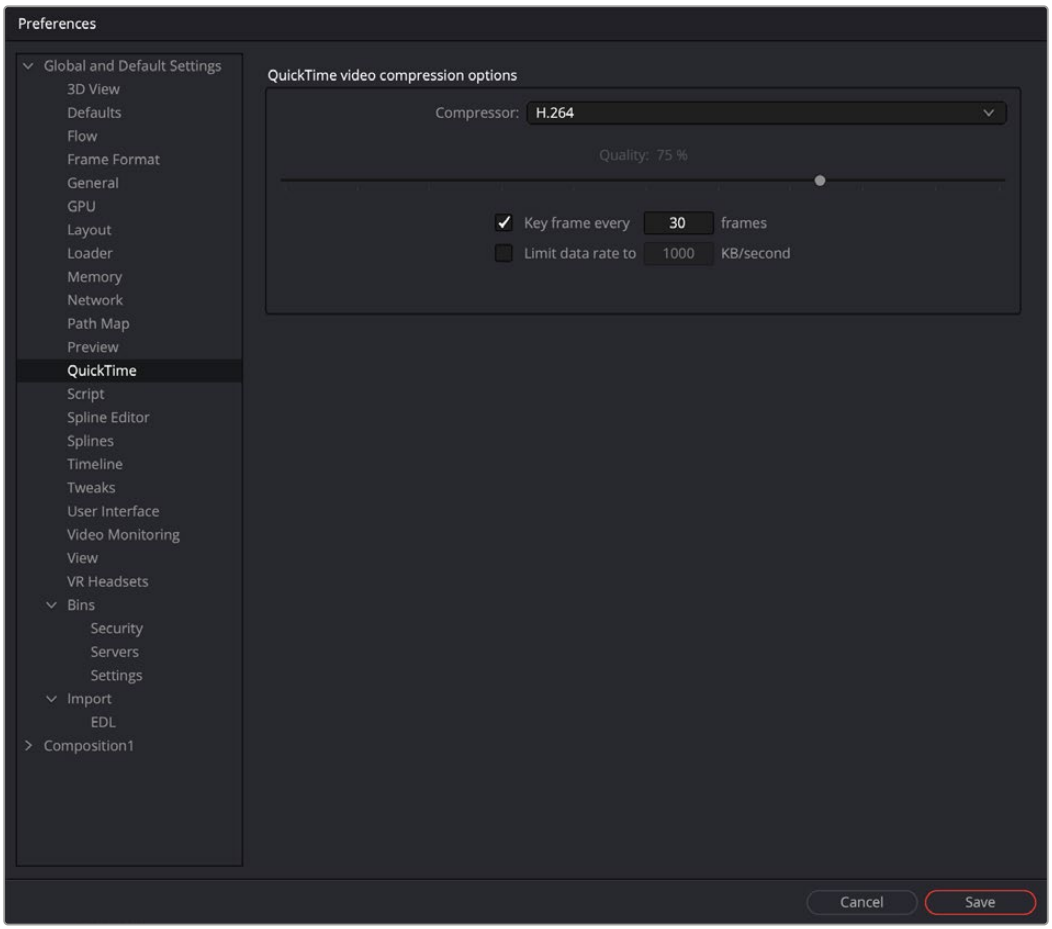
The Preview preferences

Options

- **Render Previews Using Proxy Scaling:** When checked, this option scales down the images to the preview size for the Loader and Creator tools. This causes much faster rendering. If this option is disabled, frames will be rendered at full size and are then scaled down.
- **Skip Frames to Maintain Apparent Framerate:** When checked, frames are skipped during playback of Flipbooks and file sequences to maintain the frame rate setting.
- **Show Previews for Active Loaders:** This setting determines whether the preview playback controls are shown below the Inspector when a Loader with a valid file is activated.
- **Show Previews for Active Savers:** This setting determines whether the preview playback controls below the Inspector are shown when a Saver with a valid file is activated.
- **Display File Sequences On:** This setting determines which viewer or external monitor is used for the interactive and file sequence playbacks as well as for the scrubbing function in the bins.

QuickTime

The QuickTime preferences are only available in Fusion Studio on macOS. These settings configure the default QuickTime codec settings when you select QuickTime as the rendering file format in the Saver node.



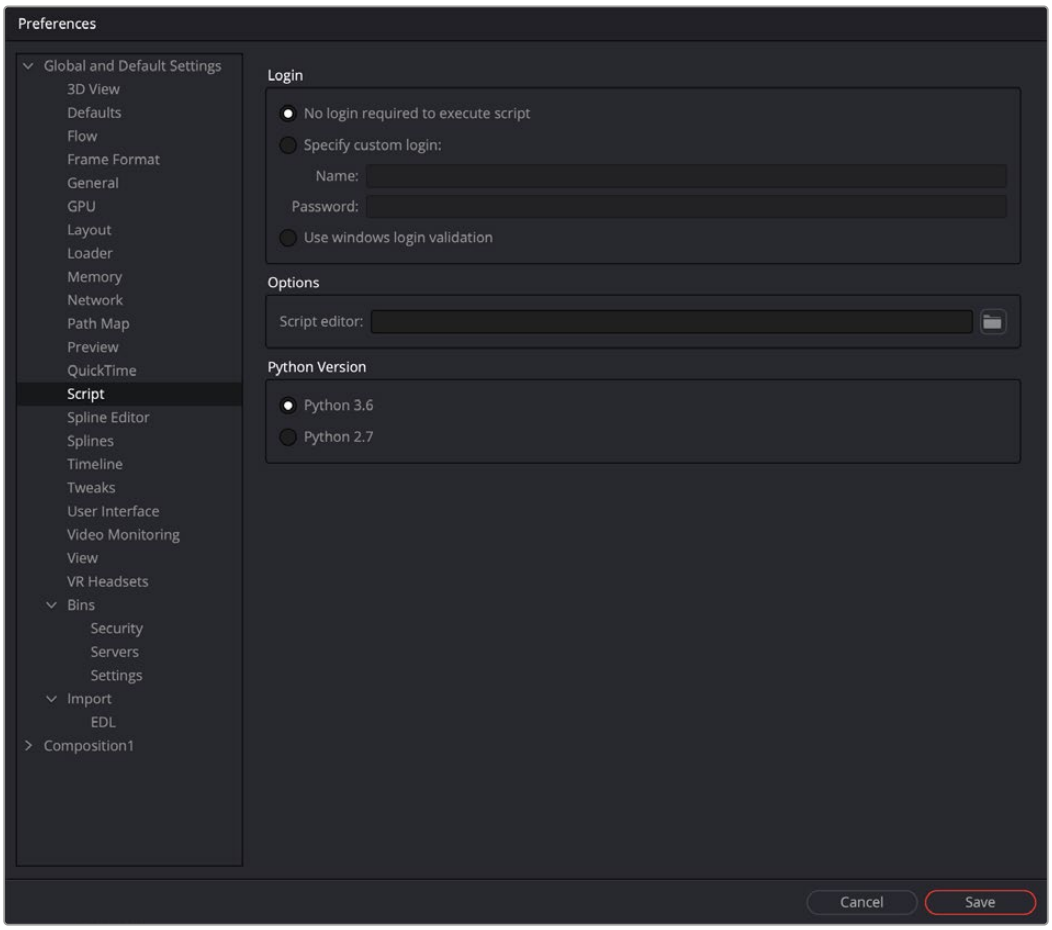
The QuickTime preferences

- **Compressor:** This drop-down menu displays the QuickTime codecs available from your computer. Fusion tests each codec when the program is started; therefore, some codecs may not be available if the tests indicate that they are unsuitable for use within Fusion.

- **Quality:** This slider is used to determine the amount of compression to be used by the codec. Higher values produce clearer images but larger files. Not all codecs support the Quality setting.
- **Key Frame Every X Frames:** When checked, the codec will create key frames at specified intervals. Key frames are not compressed in conjunction with previous frames and are, therefore, quicker to seek within the resulting movie. Not all codecs support the key frame setting.
- **Limit Data Rate To X KB/Second:** When checked, the data rates of the rendered file will be limited to the amount specified. Not all codecs support this option. Enter the data rate used to limit the QuickTime in kilobytes (kB) per second, if applicable. This control will have no effect if the Limit Data Rate To option is not selected.

Script

The preferences for Scripting include a field for passwords used to execute scripts from the command line and applications for use when editing scripts.



The Script preferences

Login

There are three login options for running scripts outside of the Fusion application.

- **No Login Required to Execute Script:** When enabled, scripts executed from the command line, or scripts that attempt to control remote copies of Fusion, do not need to log in to the workstation in order to run.

- **Specify Custom Login:** If a username and password are assigned, Fusion will refuse to process incoming external script commands (from FusionScript, for example), unless the Script first logs in to the workstation. This only affects scripts that are executed from the command line, or scripts that attempt to control remote copies of Fusion. Scripts executed from within the interface do not need to log in regardless of this setting. For more information, see the Scripting documentation.
- **Use Windows Login Validation:** When using Fusion on Windows, enabling this option verifies the user name and password (also known as credentials) with the operating system before running the script.

Options

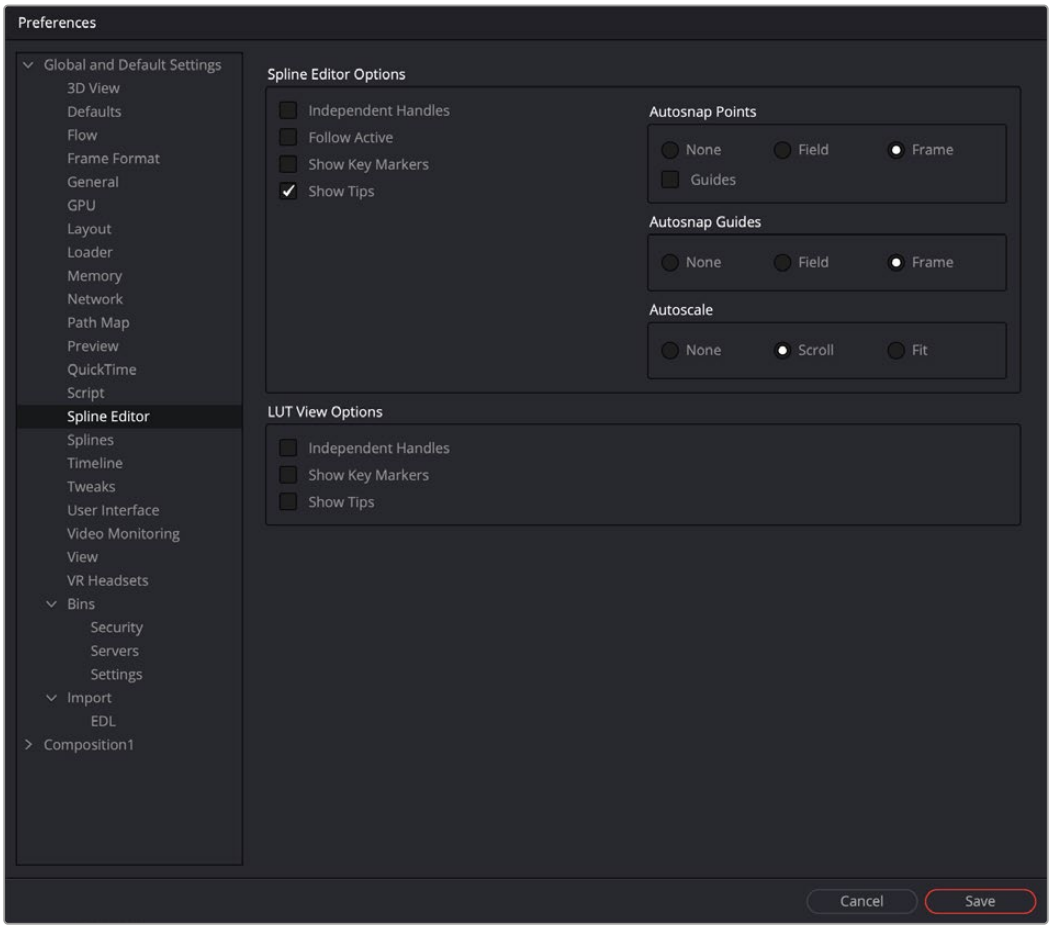
Script Editor: Use this preference to select an external editor for scripts. This preference is used when selecting Scripts > Edit.

Python Version

Two options are presented here for selecting the version of Python that you plan on using for your scripts.

Spline Editor

The Spline Editor preferences allow you to set various spline options for Autosnap behavior, handles, markers, and more. This only affects splines displayed in the Spline Editor, not splines created in the viewer using the polygon tool or paths.



The Spine Editor preferences

Spline Editor Options

These settings control the spline behavior in the Spline Editor, as well as the appearance of the graph area.

- **Independent Handles:** Enabling this option allows the In or Out direction handle on newly created key frames to be moved independently without affecting the other. This option is also available via the Options submenu when right-clicking in the Spline Editor graph.
- **Follow Active:** The Spline Editor focuses on the currently active tool. This option is also available via the Options submenu when right-clicking in the Spline Editor graph.
- **Show Key Markers:** Small colored triangles will be displayed at the top of the Spline Editor Time Ruler to indicate key frames on active splines. The colors of the triangles match the colors of the splines. This option is also available via the Show submenu when right-clicking in the Spline Editor graph.
- **Show Tips:** Toggles if tooltips are displayed or not. This option is also available via the Show submenu when right-clicking in the Spline Editor graph.
- **Autosnap Points:** When moving points in the Spline Editor, these will snap to the fields or frames or can be moved freely. This option is also available via the Options submenu when right-clicking in the Spline Editor graph.
- **Guides:** When moving points in the Spline Editor, these will snap to guides as well. This option is also available via the Options submenu when right-clicking in the Spline Editor graph.
- **Autosnap Guides:** When moving or creating guides, these will snap to the fields or frames or can be moved freely. This option is also available via the Options submenu when right-clicking in the Spline Editor graph.
- **Autoscale:** Keeps the Spline Editor scales intact on changing the editable spline content of the graph. This scale is also available via the Options submenu when right-clicking in the Spline Editor graph.
- **Scroll:** Scrolls horizontally and vertically to show all or most of the spline points. This option is also available via the Scale submenu when right-clicking in the Spline Editor graph.
- **Fit:** Zooms to fit all points within the spline graph, if necessary. This option is also available via the Scale submenu when right-clicking in the Spline Editor graph.

LUT View Options

These settings let you control how the LUT view is displayed.

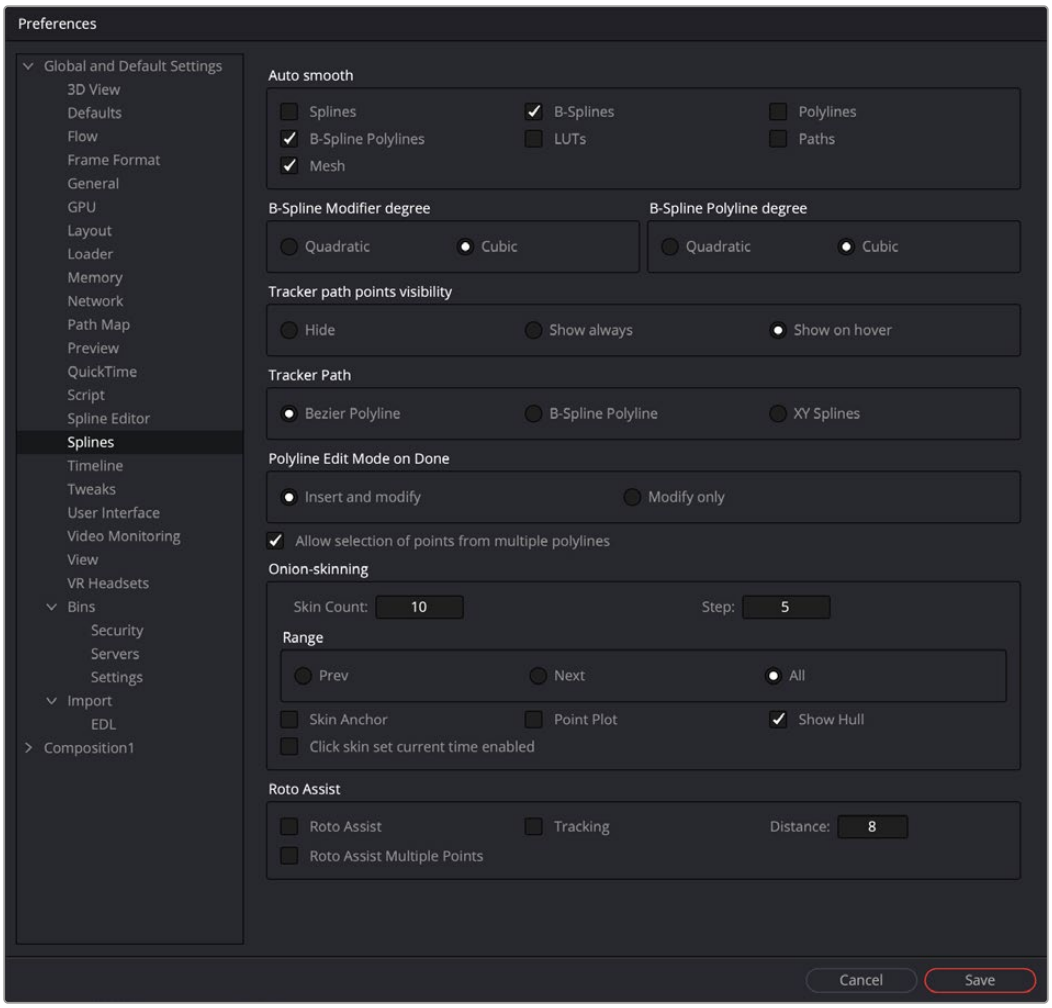
- **Independent Handles:** Enabling this option allows the In or Out direction handle on newly created key frames to be moved independently without affecting the other.
- **Show Key Markers:** Small colored triangles will be displayed at the top of the Spline Editor Time Ruler to indicate key frames on active splines. The colors of the triangles match the colors of the splines.
- **Show Tips:** Toggles whether tooltips are displayed.

Splines

Options for the handling and smoothing of animation splines, tracker path defaults, and rotoscoping are found in the Splines preferences.

- **Autosmooth:** Automatically smooths out any newly created points or key frames on the splines selected in this section. You can choose to automatically smooth animation splines, B-Splines, polyline matte shapes, LUTs, paths, and meshes.
- **B-Spline Modifier Degree:** This setting determines the degree to which the line segments influence the resulting curvature when B-Splines are used in animation. Cubic B-Splines determine a segment through two control points between the anchor points, and Quadratic B-Splines determine a segment through one control point between the anchor points.

- **B-Spline Polyline Degree:** This setting is like the one above but applies to B-Splines used for masks.

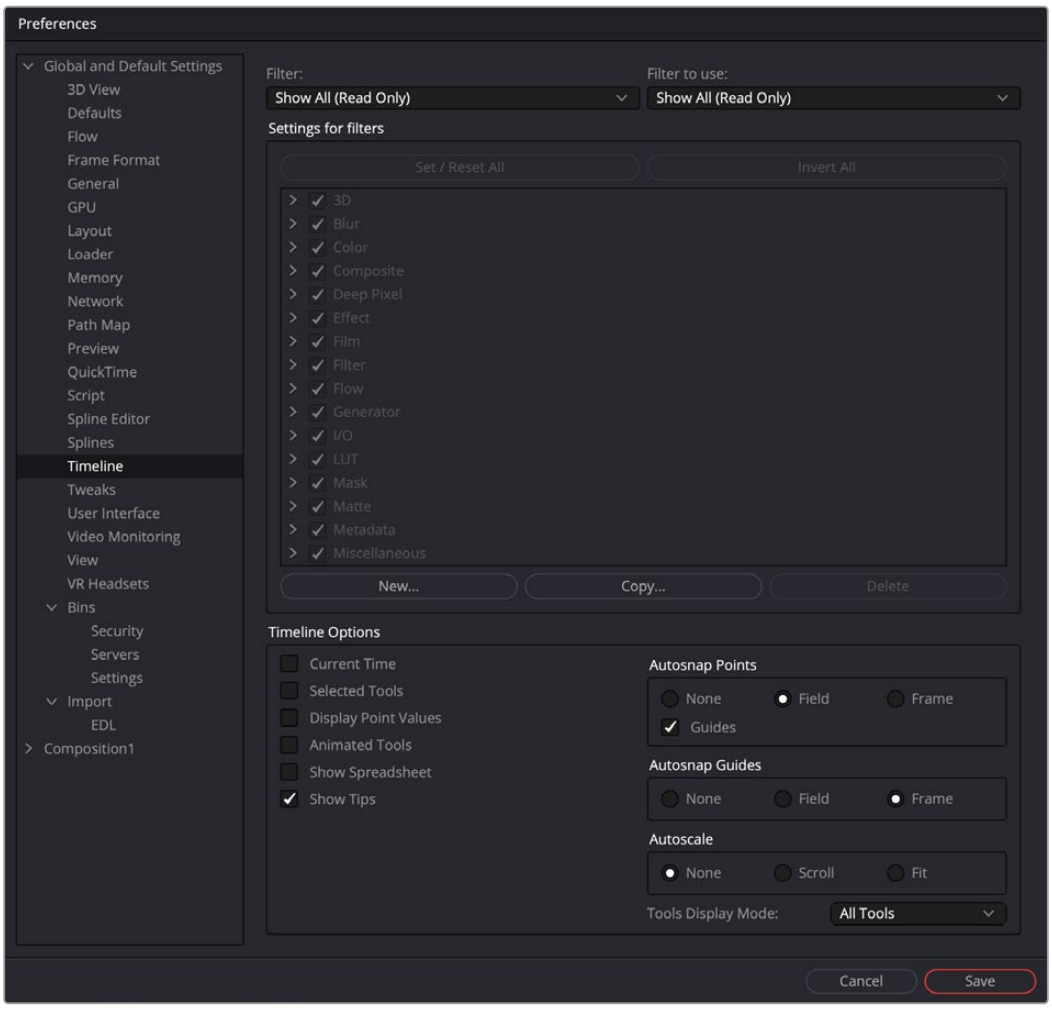


The Splines preferences

- **Tracker Path Points Visibility:** This setting determines the visibility of the control points on tracker paths. You can show them, hide them, or show them when your cursor hovers over the path, which is the default behavior.
- **Tracker Path:** The default tracker creates Bézier-style spline paths. Two other options in this setting allow you to choose B-Spline or XY Spline paths.
- **Polyline Edit Mode on Done:** This setting determines the state of the Polyline tool after you complete the drawing of a polyline. It can either be set to modify the existing control points on the spline or modify and add new control points to the spline.
- **Onion Skinning:** The Onion Skinning settings determine the number of frames displayed while rotoscoping, allowing you to preview and compare a range of frames. You can also adjust if the preview frames only from the frame prior to the current frame, after the current frames, or split between the two.

Timeline

The Timeline preferences is where you create and edit Keyframes Editor/Spline Editor filters and set default options for the Keyframes Editor.



The Timeline preferences

Filter/Filter to Use

The Filter menu populates the hierarchy area below the menu with that setting. It lets you edit the filters. The Filter to Use menu selects the default filter setting located in the Keyframes Editor Options menu.

Settings for Filters

This area is used to create a new filter and define its settings. You start by first clicking the New button and entering the name of the new Filter. You then select any of the tools that you want the filter to contain. Only tools that are checked will appear in the Keyframes Editor or Spline Editor when the filter is selected. You can also create a copy of the filter using the Copy button or remove a filter from the list by clicking the Delete button.

Timeline Options

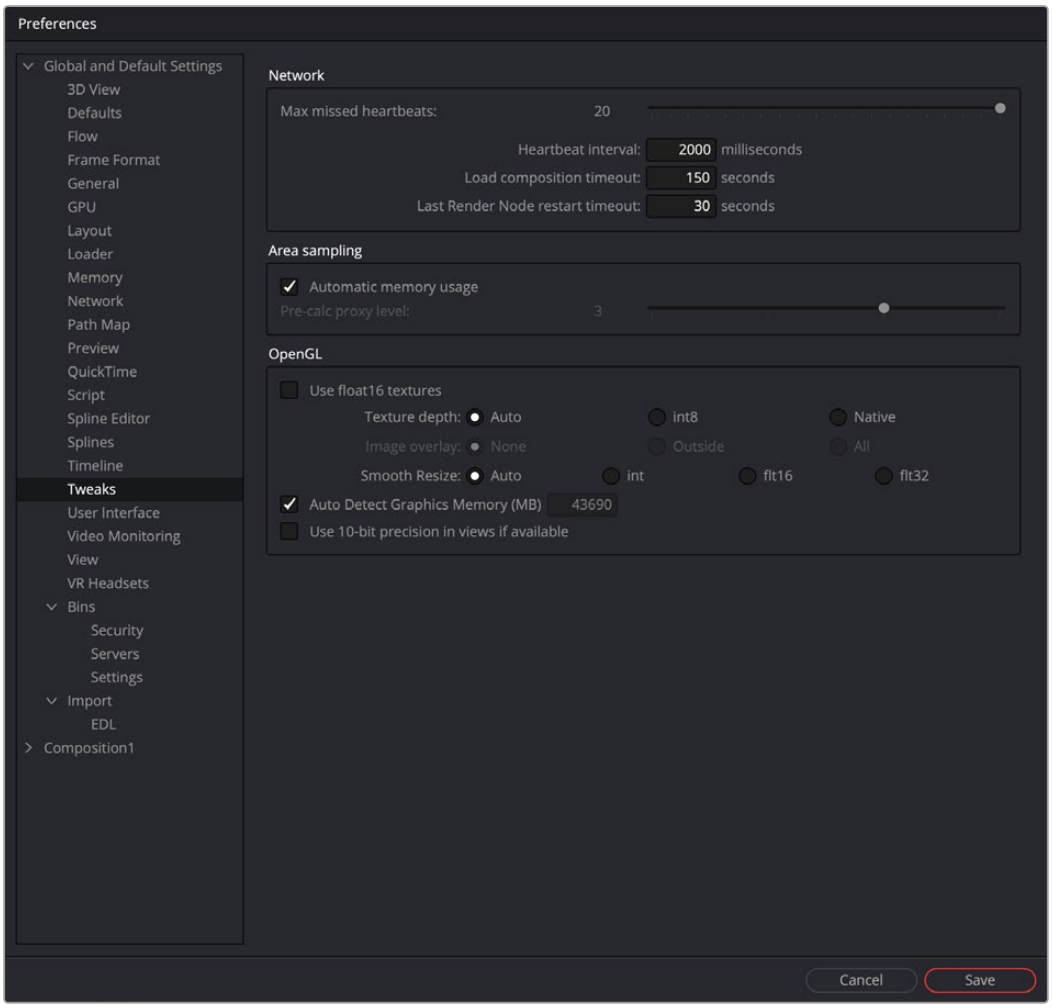
The Timeline Options configure which options in the Keyframe Editor are enabled by default. A series of checkboxes correspond to buttons located in the Timeline, allowing you to determine the states of those buttons at the time a new comp is created.

For more information on the Keyframes Editor functions, *see Chapter 9, "Animating in Fusion's Keyframes Editor,"* in the Fusion Reference Manual.

- **Autosnap Points:** When moving points in the Keyframes Editor, the points will snap to the fields or to the frames, or they can be moved freely.
- **Guides:** When moving points in the Keyframes Editor, the point will snap to the guides that are placed in the Timeline graph.
- **Autosnap Guides:** When moving or creating guides, the guides will snap to the fields or to the frames, or they can be moved freely.
- **Autoscale:** Keeps the Timeline scales intact while changing the editable spline content in the graph. When set to scroll, the Timeline scrolls horizontally and vertically to show all or most of the spline points when changing the editable spline content in the graph. When set to Fit, the Timeline zooms to fit all points within the graph, if necessary.
- **Tools Display Mode:** This menu controls the default sort order of the tools displayed in the Keyframes Editor. The default can be changed using the Sort order menu in the upper right of the Keyframes Editor.

Tweaks

The Tweaks preferences handle a collection of settings for fine-tuning Network rendering in Fusion Studio and graphics hardware behavior.



The Tweaks preferences

Network

The Network section is used to control and monitor the health of communication packets over TCP/IP when rendering over a network in Fusion Studio.

- **Maximum Missed Heartbeats:** This setting determines the maximum number of times the network is checked before terminating the communication with a Render node.
- **Heartbeat Interval:** This sets the time between network checks.
- **Load Composition Timeout:** This timeout option determines how long the Render Manger will wait for a composition to load before moving on to another task.
- **Last Slave Restart Timeout:** This timeout option determines how long the Render Manager will wait for a render salve to respond before using another render slave.

File I/O

The File I/O options are used to control the performance when reading frames or large media files from both direct and networked attached storage.

- **I/O Canceling:** This option enables a feature of the operating system that allows queued operations to be canceled when the function that requested them is stopped. This can improve the responsiveness, particularly when loading large images over a network.

  Enabling this option will specifically affect performance while loading and accessing formats that perform a large amount of seeking, such as the TIFF format.

  This option has not been tested with every hardware and OS configuration, so it is recommended to enable it only after you have thoroughly tested your hardware and OS configuration using drive loads from both local disks and network shares.
- **Enable Direct Reads:** Enabling this checkbox uses a more efficient method when loading a large chunk of contiguous data into memory by reducing I/O operations. Not every operating system employs this ability, so it may produce unknown behavior.
- **Read Ahead Buffers:** This slider determines the number of 64K buffers that are use to read ahead in a file I/O operation. The more buffers, the more efficient loading frames from disk will be, but the less responsive it will be to changes that require disk access interactively.

Area Sampling

The Area Sampling options allow you to fine-tune the RAM usage on Render nodes by trading off speed for lower RAM requirements.

- **Automatic Memory Usage:** This checkbox determines how area sampling uses available memory. Area sampling is used for Merges and Transforms. When the checkbox is enabled (default), Fusion will detect available RAM when processing the tool and determine the appropriate trade-off between speed and memory.

  If less RAM is available, Fusion will use a higher proxy level internally and take longer to render. The quality of the image is not compromised in any way, just the amount of time it takes to render. In node trees that deal with images larger than 4K, it may be desirable to override the automatic scaling and fix the proxy scale manually. This can preserve RAM for future operations.
- **Pre-Calc Proxy Level:** Deselecting the Automatic Memory will enable the Pre-Calc Proxy Scale slider. Higher values will use less RAM but take much longer to render.

Open GL

This section controls how Fusion makes use of your graphics card when compositing in 3D with the Renderer 3D node. Most settings may be left as they are, but since OpenGL hardware varies widely in capabilities and different driver revisions can sometimes introduce bugs, these tweaks can be useful if you are experiencing unwanted behavior.

- **Disable View LUT Shaders:** OpenGL shaders can often dramatically accelerate View LUTs, but this can occasionally involve small trade-offs in accuracy. This setting will force Fusion to process LUTs at full accuracy using the CPU instead. Try activating this if View LUTs do not seem to be giving the desired result.
- **Use Float16 Textures:** If your graphics hardware supports 16-bit floating-point textures, activating this option will force int16 and float32 images to be uploaded to the viewer as float16 instead, which may improve playback performance.
- **Texture Depth:** Defines in what depth images are uploaded to the viewer.

  Auto: The Auto option (recommended) lets Fusion choose the best balance of performance and capability.

  int8: Similar to the Use Float16 Textures switch, this option can be used to force images to be uploaded to the Display View as int8, which can be faster but gives less range for View LUT correction.

  Native: The Native option uploads images at their native depth, so no conversion is done.
- **Image Overlay:** The Image Overlay is a viewer control used with Merge and Transform tools to display a translucent overlay of the transformed image. This can be helpful in visualizing the transformation when it is outside the image bounds but may reduce performance when selecting the tool if cache memory is low. There are three settings to choose from: None, Outside, and All.

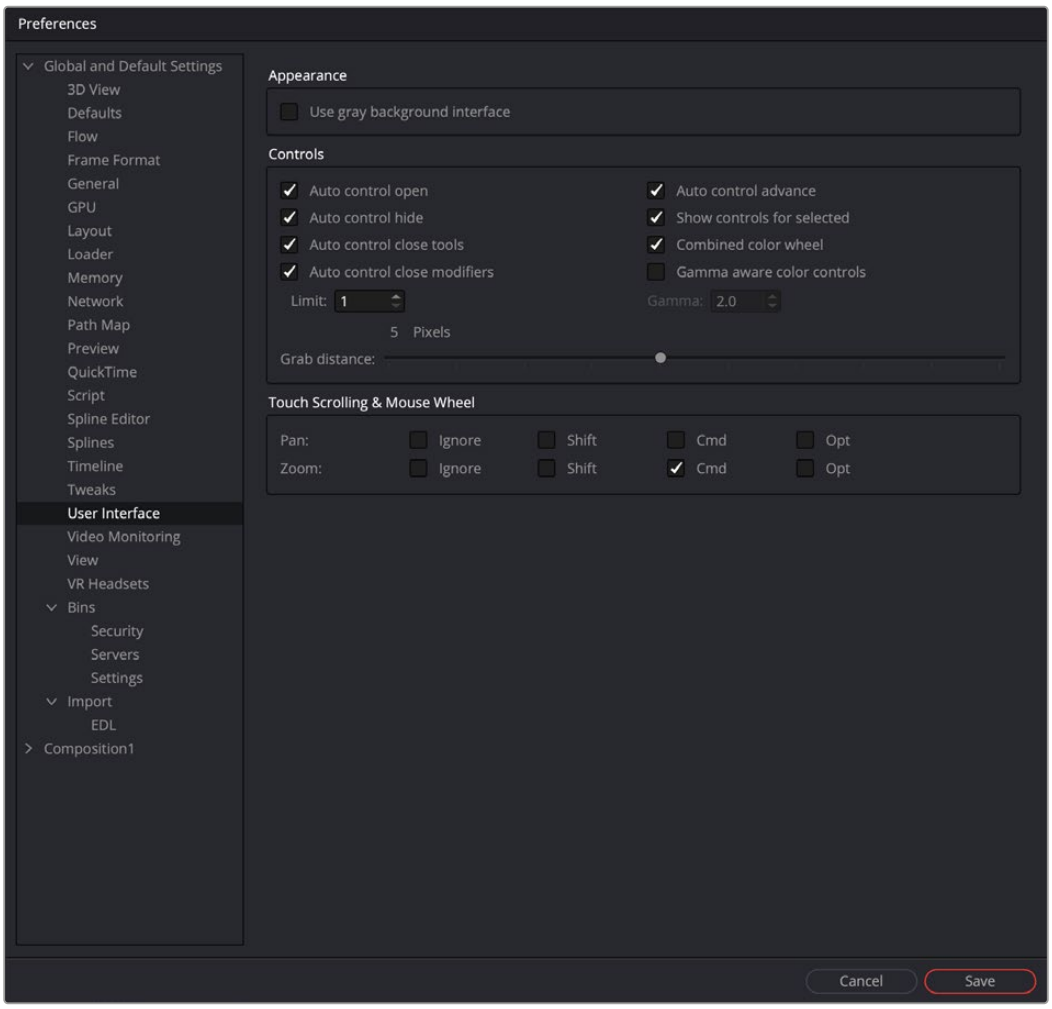
  None: This setting never displays the translucent overlay or controls, which can reduce the need for background renders, in some cases resulting in a speed up of the display.

  Outside: This will display only those areas of the control that are outside the bounds of the image, which can reduce visual confusion.

  All: Displays all overlays of all selected tools.
- **Smooth Resize:** This setting can disable the viewer's Smooth Resize behavior when displaying floating-point images. Some older graphics cards are not capable of filtering floating-point textures or may be very slow. If Smooth Resize does not work well with float images, try setting this to flt16 or int.
- **Auto Detect Graphics Memory (MB):** Having Fusion open alongside other OpenGL programs like 3D animation software can lead to a shortage of graphics memory. In those cases, you can manually reduce the amount of memory Fusion is allowed to use on the card. Setting this too low or too high may cause performance or data loss.
- **Use 10-10-10-2 Framebuffer:** If your graphics hardware and monitor support 30-bit color (Nvidia Quadro/AMD Radeon Pro, and some Nvidia GeForce/AMD Radeon), this setting will render viewers with 10 bits per primary accuracy, instead of 8 bits. Banding is greatly reduced when displaying 3D renders or images deeper than 8-bit.

User Interface

The User Interface preferences set the appearance of the user interface window and how the Inspector is displayed.



The User Interface preferences

Appearance

When enabled, the Use Gray Background Interface checkbox will change the color of the background in Fusion's panels to a lighter, more neutral shade of gray.

Controls

This group of checkboxes manages how the controls in the Inspector are displayed.

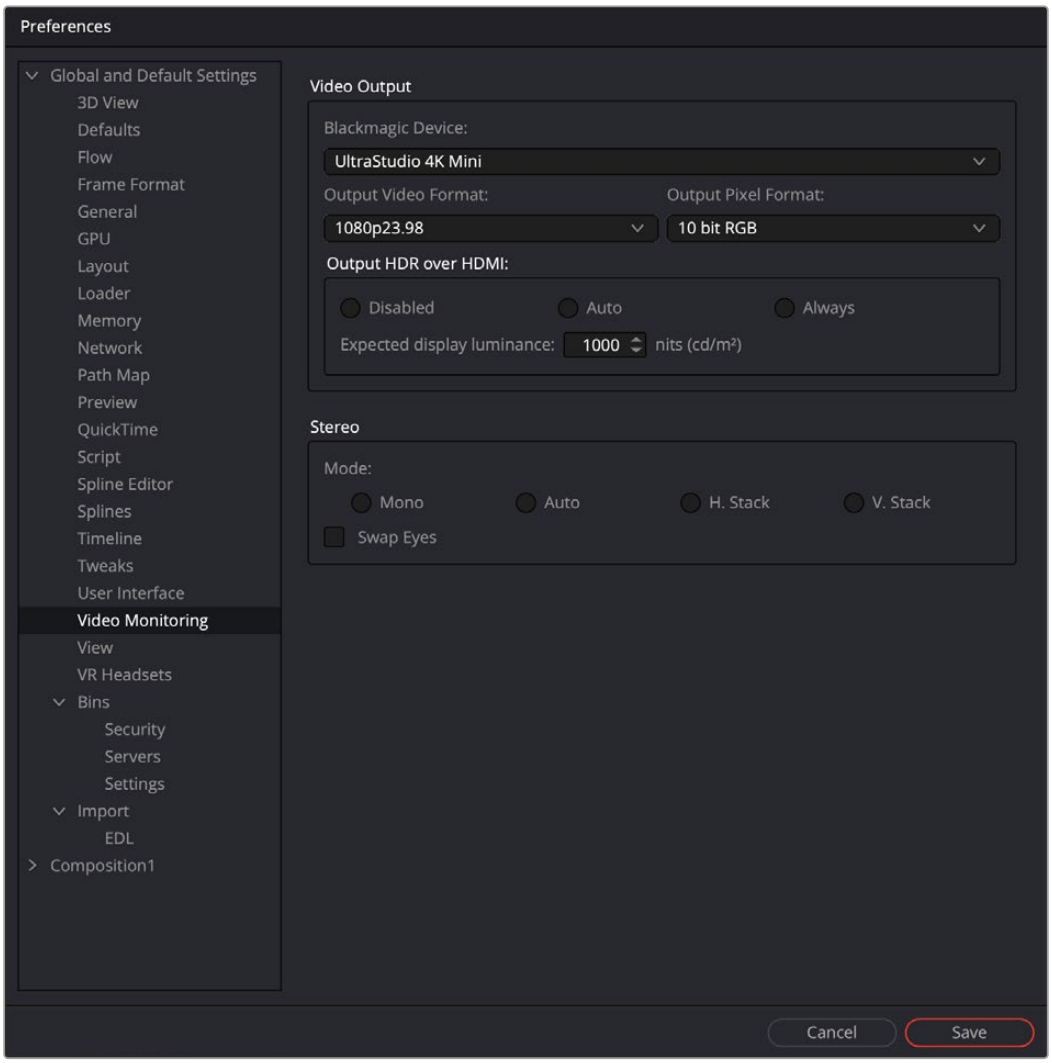
- **Auto Control Open:** When disabled, only the header of the selected node is displayed in the Inspector. You must double-click the header to display the parameters. When enabled, the parameters are automatically displayed when the node is selected.
- **Auto Control Hide:** When enabled, only the parameters for the currently active tool (red outline) will be made visible. Otherwise, all tool headers will be visible and displayed based on the Auto Control Open setting.
- **Auto Control Close Tools:** When enabled, only the active (red outlined) tool in the Node Editor will have controls displayed. Any previous active node's tools will be closed in the Inspector. When disabled, any number of tools may be opened to display parameters at the same time. This setting has no effect if the Auto Control Hide checkbox is enabled.
- **Auto Control Close Modifiers:** When enabled, only one modifier's parameters will be displayed for the active node. Any additional modifiers for the active node will show only their header.
- **Auto Control Advance:** If the Auto Control Advanced checkbox is enabled, the Tab key and Return/Enter key will cause the keyboard focus to advance to the next edit box within the Inspector. When disabled, Return/Enter will cause the value entered to be accepted, but the keyboard focus will remain in the same edit box of the control. The Tab key can still be used to advance the keyboard focus.
- **Show Controls for Selected:** When this option is disabled, only the active tool's parameters are shown in the Inspector. By default, it is enabled, showing controls for the active tool as well as all selected tools.
- **Combined Color Wheel:** When the Color Corrector tool is displayed in the Inspector, enabling this checkbox will show one color wheel with buttons to switch between the master, shadow, midtones, and highlight channels. Otherwise, four color wheels are displayed in the Inspector.
- **Gamma Aware Color Controls:** This setting adjusts color correction nodes when working with Rec. 709 images in a non-color managed project. Rec. 709 images appear correct on the computer monitor because monitors have a gamma adjustment built in. When working in the Rec. 709 color space without color management, enabling Gamma Aware color removes the gamma, applies the color correction as if it where linear, and then reapplies the gamma. For Rec. 709 images, enable the Gamma Aware setting and enter a Gamma value of 2.4. In a color managed linear project, this should be set to Off or a value of 1.0. When dealing with mixed color spaces, Fusion reads the metadata from the image and sets the Aware gamma value based on the metadata available.
- **Grab Distance:** This slider ranges from 1 to 10 and defaults to 5. It designates the active area around the mouse pointer and can be modified if you have difficulties in selecting points for modification in paths and spline curves. Smaller values will require a more accurate selection with the mouse pointer.

Touch Scrolling and Mouse Wheel

This group of settings allows you to configure which, if any, keyboard modifiers are needed to pan or zoom a panel when using a trackpad or middle mouse wheel.

Video Monitoring

This setting is only available in Fusion Studio. Control over video hardware for the Fusion Page is done in the DaVinci Resolve preferences. The Video Monitoring preferences are used to configure the settings of Blackmagic Design capture and playback products such as DeckLink PCIe cards and UltraStudio i/O units.



The Video Monitoring preferences

Video Output

This group of drop-down menus allows you to select the type of video I/O device you have installed, the output resolution, and the pixel format. These settings have nothing to do with your rendered output; it is only for your display hardware.

The Output HDR over HDMI settings are used to output the necessary metadata when sending high dynamic range signals over HMDI 2.0a and have it correctly decided by an HDR capable video display.

The Auto setting detects the image's values and outputs HDR. This will not affect non HDR images.

The Always setting forces HDR on all the time. This can can be useful when checking non HDR and HDR grades.

When Auto or Always is selected, you can then set the “nit” level (slang for cd/m^2) to whatever peak luminance level your HDMI connected HDR display is capable of.

Stereo Mode

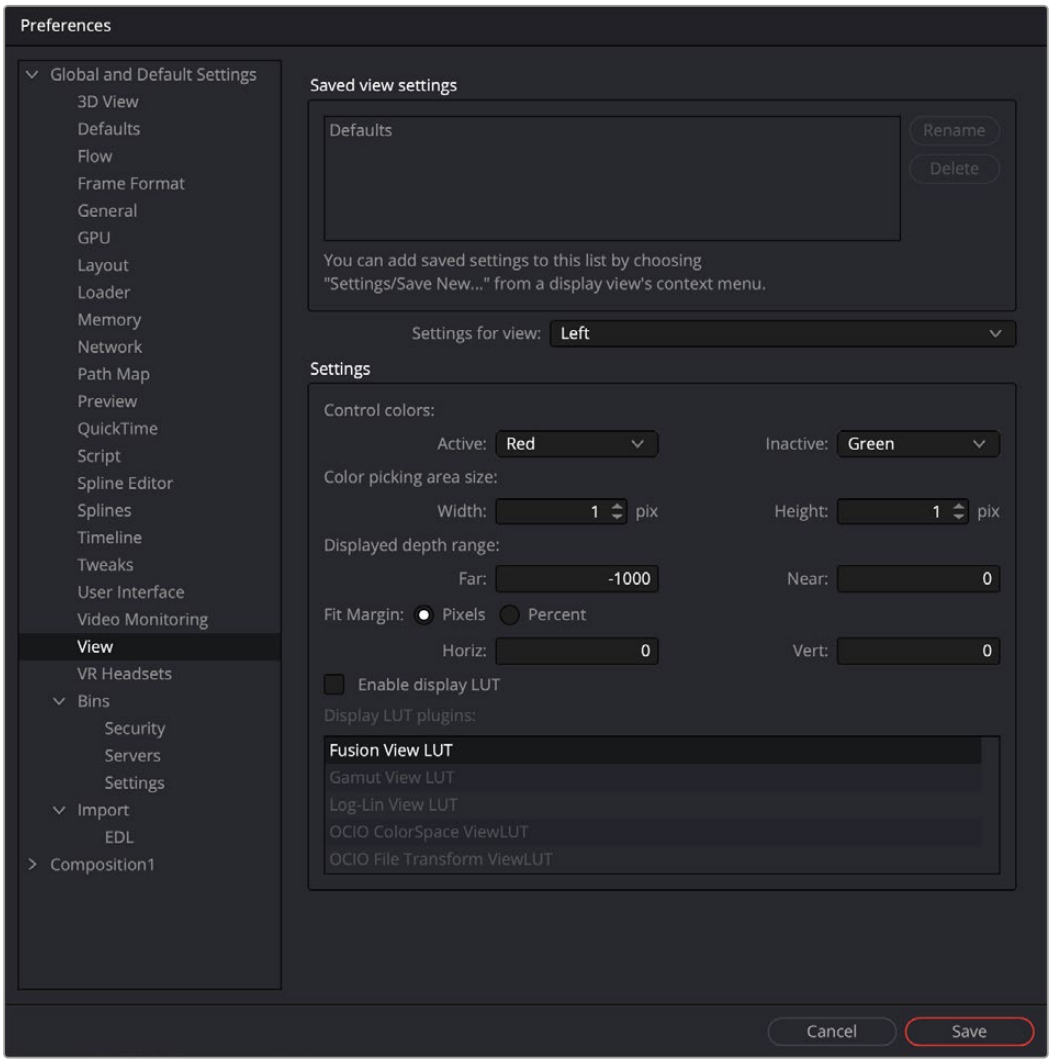
This group of settings configures the output hardware for displaying stereo 3D content.

- Mono will output a single non stereo eye.
- Auto will detect which method with which the stereo images are stacked.
- Use the Vstack option if the stereo images are stacked vertically as left on top and right at the bottom.
- Use the Hstack option if the stereo images are stacked horizontally as left and right.

The Swap eyes checkbox will swap the eyes if stereo is reversed.

View

The View preferences are used to manage settings and default controls for viewers.



The View preferences

Saved View Settings

The area at the top of the view preferences lists the currently saved settings that you create from the viewer's contextual menu. You can use the Rename and Delete buttons to manage the selected entries in the list.

> For more information on the viewer and its contextual menu, *see Chapter 7, "Using Viewers."* in the Fusion Reference Manual.

Settings for View

Each viewer has its own preferences. The Settings for View drop-down menu is used to select the viewer you want to configure.

Control Colors

The Control Colors setting allows you to determine the color of the active/inactive onscreen controls.

Color Picking Area Size

You can use these width/height controls to set the number of pixels sampled when using the Color Picker in the viewers.

Displayed Depth Range

The Displayed Depth Range setting controls the view normalization of the Z-Channel.

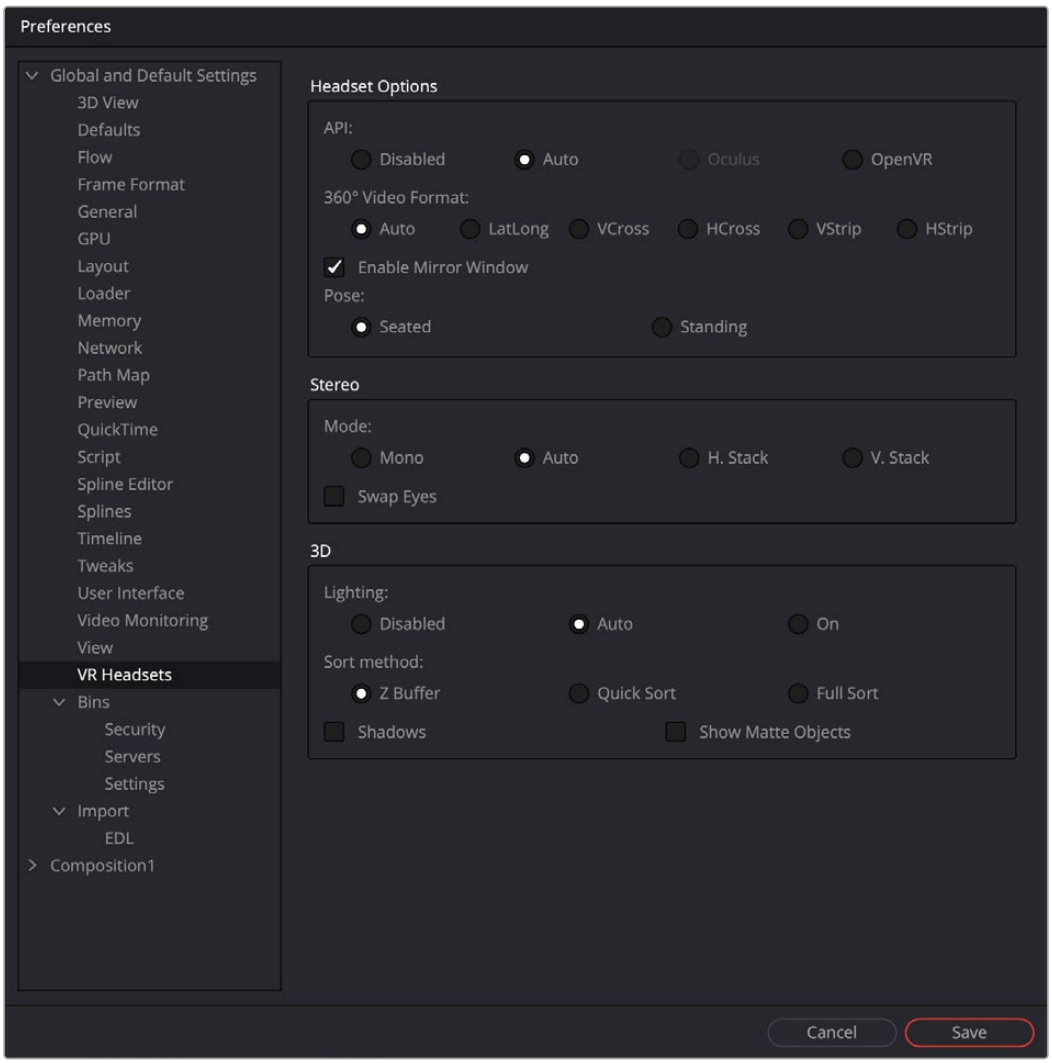
Fit Margin

The Fit Margin setting determines how much padding is left around the frame when the Fit button is pressed or Fit is selected from the viewer's contextual menu.

- **Display LUT Plugins**
  This list shows the available display LUTs and activates the selected one as default.

VR Headsets

The VR Headsets preferences allow configuration of any connected Virtual Reality headsets, including how stereo and 3D scenes are viewed.



The VR Headsets preferences

Headset Options

The Headset options are used to select the type of VR headset you are using to view the composite as well as the video layout of the 360° view.

API

- **Disabled:** Disabled turns off and hides all usage of headsets.
- **Auto:** Auto will detect which headset is plugged in.
- **Oculus:** Oculus will set the VR output to the Oculus headset.
- **OpenVR:** OpenVR will support a number of VR headsets like the HTC Vive.

360° Video Format

- **Auto:** Auto will detect the incoming image layout from the metadata and image frame aspect.
- **VCross and HCross:** VCross and HCross are the six square faces of a cube laid out in a cross, vertical or horizontal, with the forward view in the center of the cross, in a 3:4 or 4:3 image.
- **VStrip and HStrip:** VStrip and HStrip are the six square faces of a cube laid vertically or horizontally in a line, ordered as Left, Right, Up, Down, Back, Front (+X, -X, +Y, -Y, +Z, -Z), in a 1:6 or 6:1 image.
- **LatLong:** LatLong is a single 2:1 image in equirectangular mapping.
- **Enable Mirror Window:** Enable Mirror Window will show a window displaying the headset user's live view.

Stereo

Similar to normal viewer options for stereo 3D comps, these preferences control how a stereo 3D comp is displayed in a VR headset.

Mode

- **Mono:** Mono will output a single non stereo eye.
- **Auto:** Auto will detect the method with which the stereo images are stacked.
- **Vstack:** Vstack stereo images are stacked vertically as left on top and right at the bottom.
- **Hstack:** Hstack stereo images are stacked horizontally as left and right.
- **Swap Eyes:** Swap eyes will swap the eyes if stereo is reversed.

3D

Similar to normal viewer options for 3D comps, these preferences control how a 3D comp is displayed in a VR headset.

Light

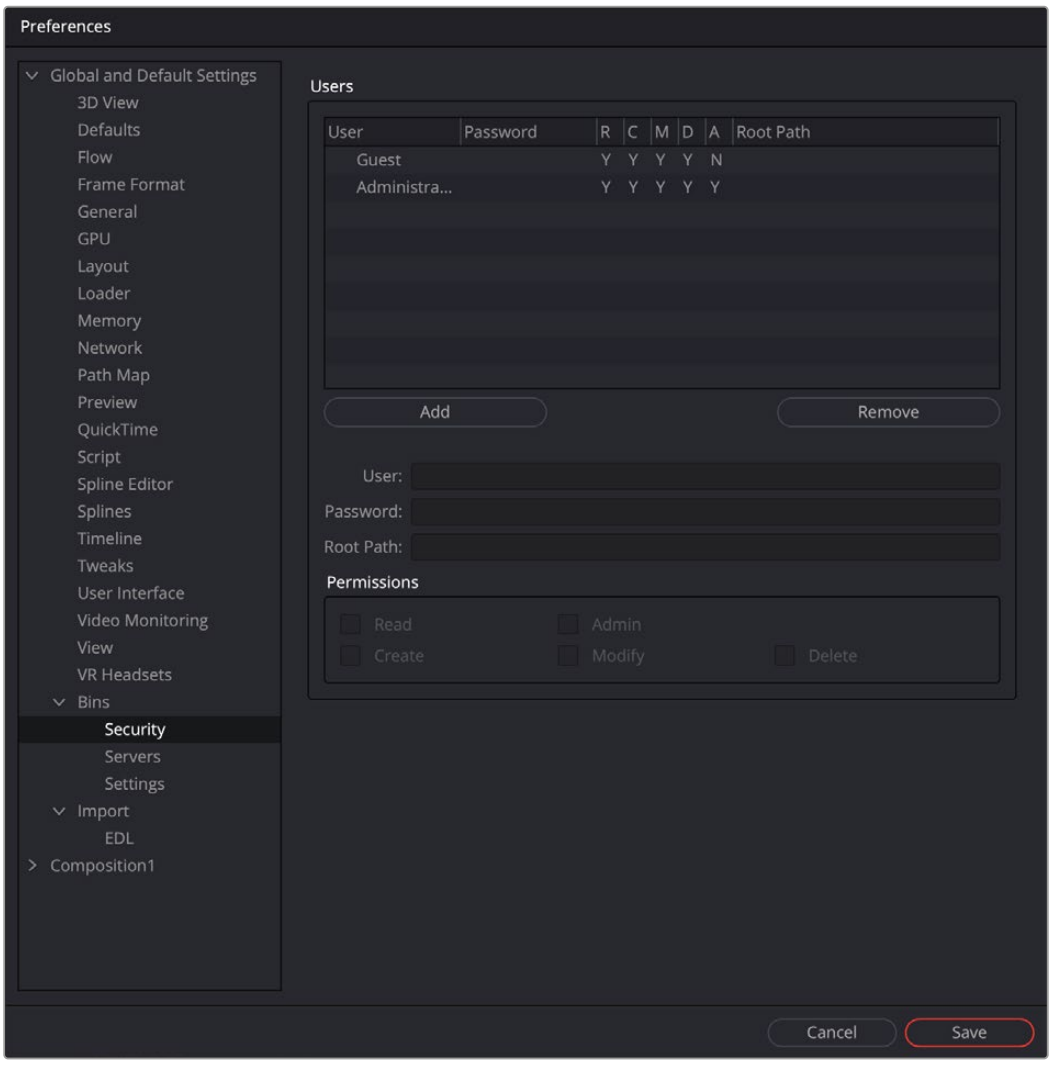
- Disabled lighting is off.
- Auto will detect if lighting is on in the view.
- On will force lighting on in the VR view.

Sort Method

- Z buffer sorting is the fast OpenGL method of sorting polygons.
- Quick Sort will sort the depth of polygons to get better transparency rendering.
- Full Sort will use a robust sort and render method to render transparency .
- Shadows can be on or off.
- Show Matte Objects will make matte objects visible in view or invisible.

Bins/Security

Bins preferences are only available in Fusion Studio. These preferences are used to manage the Bin users and their permissions.



The Bins Security preferences

Users List

The Users List is a list of the users and their permissions. You can select one of the entries to edit their settings using the User and Password edit boxes.

- **Add:** The Add button is used to add a new user to the list by entering a username and password.
- **Remove:** Click this button to remove the selected entry.

User

This editable field shows the username for the selected Bin Server item. If the username is unknown, try “Guest” with no password.

Password

Use this field to enter the password for the Bin user entered in the Users list.

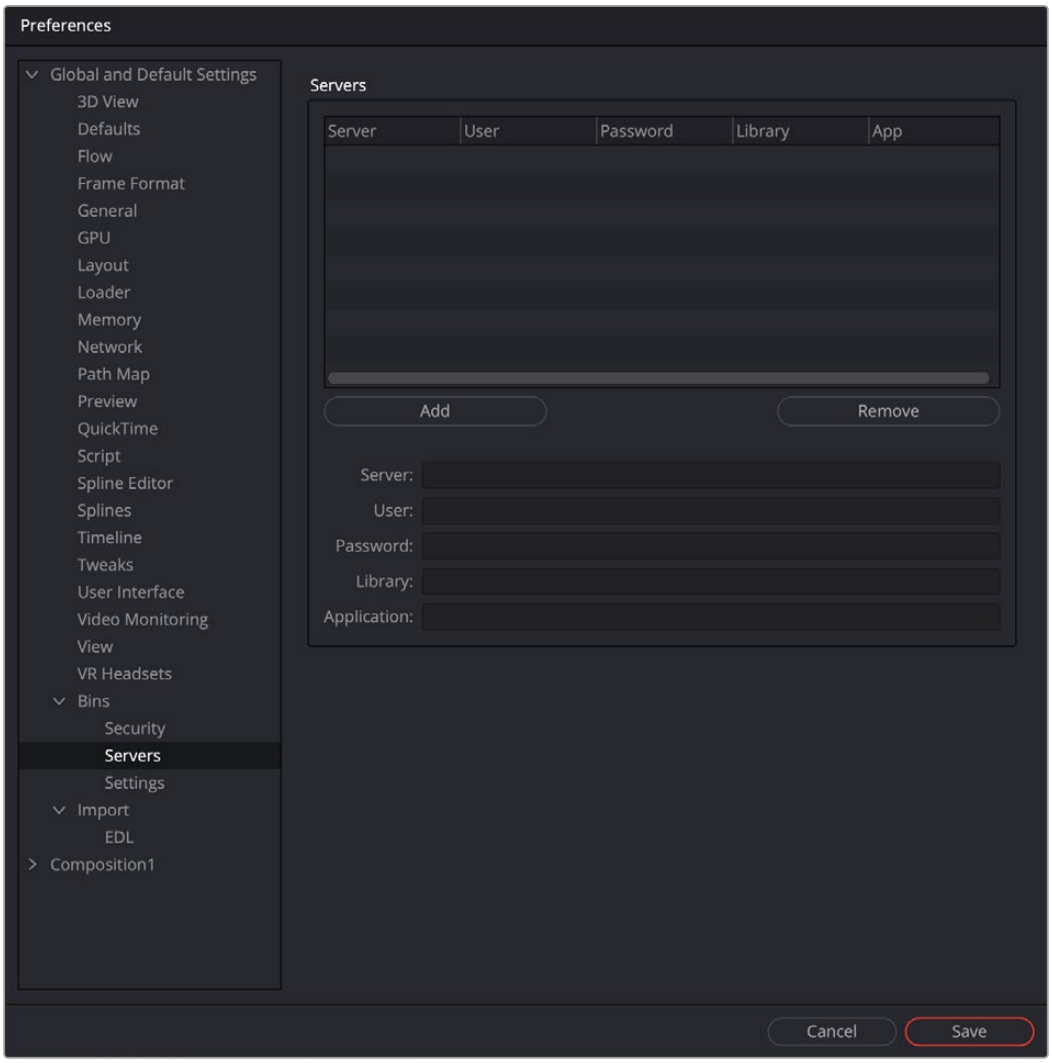
Permissions

The administrator can set up different permission types for users.

- **Read:** This will allow the user to have read-only permission for the bins.
- **Create:** This will allow the user to create new bins.
- **Admin:** This gives the user full control over the bins system.
- **Modify:** This allows the user to modify existing bins.
- **Delete:** This allows the user to remove bins.

Bins/Server

These preferences are used to add Bin Servers to the list of bins Fusion will display in the Bins dialog.



The Bin Servers preferences

Servers

This dialog lists the servers that are currently in the connection list. You can select one of the entries to edit its settings.

- **Add:** Use this button to add a new server to the list.
- **Remove:** Click this button to remove the selected entry.

Server

This editable field shows the name or IP address of the server for the selected entry in the list.

User

This editable dialog shows the username for the selected Bin Server item.

Password

Use this field to enter the password for the server entered in the Server list.

Library

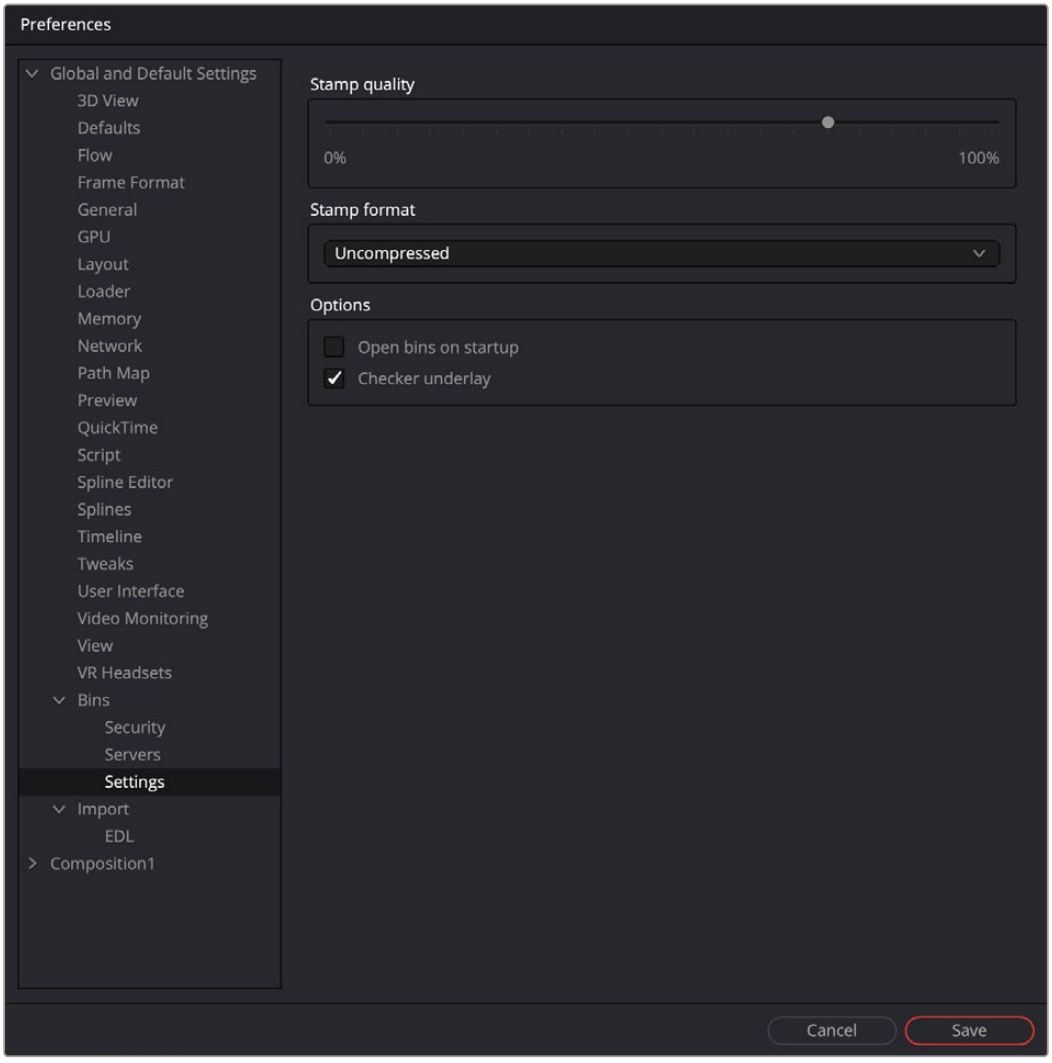
The Library field lets you name the bins. If you wanted to create a bin for individual projects, you would name it in the Library field and each project would gets its own bin.

Application

The Application field allows larger studios to specify some other program to serve out the Bin requests.

Bins/Settings

These preferences are used to control the default behavior of bins.



The Bins Settings preferences

Stamp Quality

The Stamp Quality is a percentage slider that determines the compression ratio used for Stamp thumbnail creation. Higher values offer better quality but take up more space.

Stamp Format

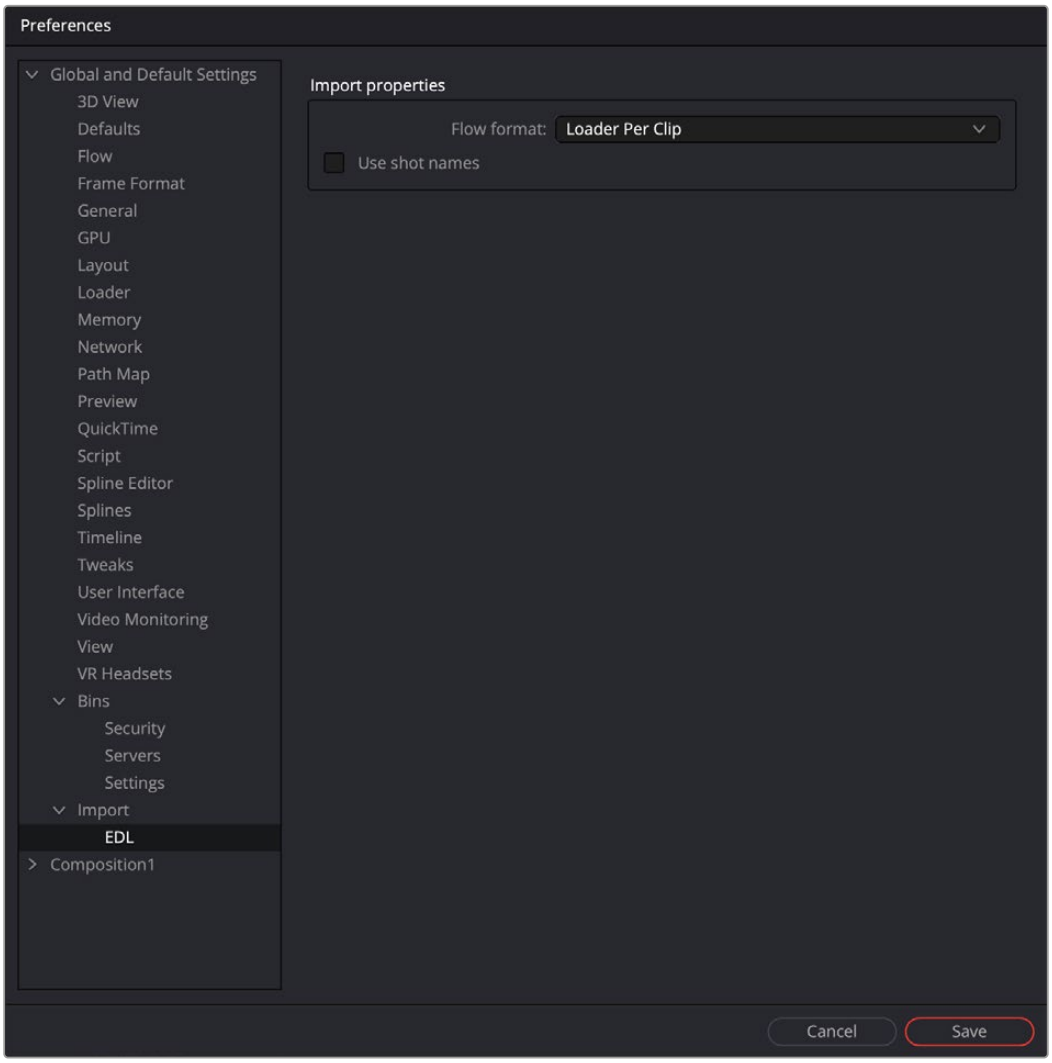
This drop-down list determines whether the Stamp thumbnails will be saved as compressed or uncompressed.

Options

- **Open Bins on Startup:** When Open Bins on Startup is checked, the bins will open automatically when Fusion is launched.
- **Checker Underlay:** When the Checker Underlay is enabled, a checkerboard background is used for clips with alpha channels. When disabled, a gray background matching the Bin window is used as the clip's background.

EDL Import

The EDL Import options are used to determine how compositions are created from imported CMX-formatted EDL files.



The EDL Import preferences

Flow Format

This drop-down menu provides three options that determine how the node tree is constructed for the imported EDL file.

- **Loader Per Clip:** A Loader will be created for each clip in the EDL file.
- **A-B Roll:** A node tree with a Dissolve tool will be created automatically.
- **Loader Per Transition:** A Loader with a Clip list will be created, representing the imported EDL list.

Use Shot Names

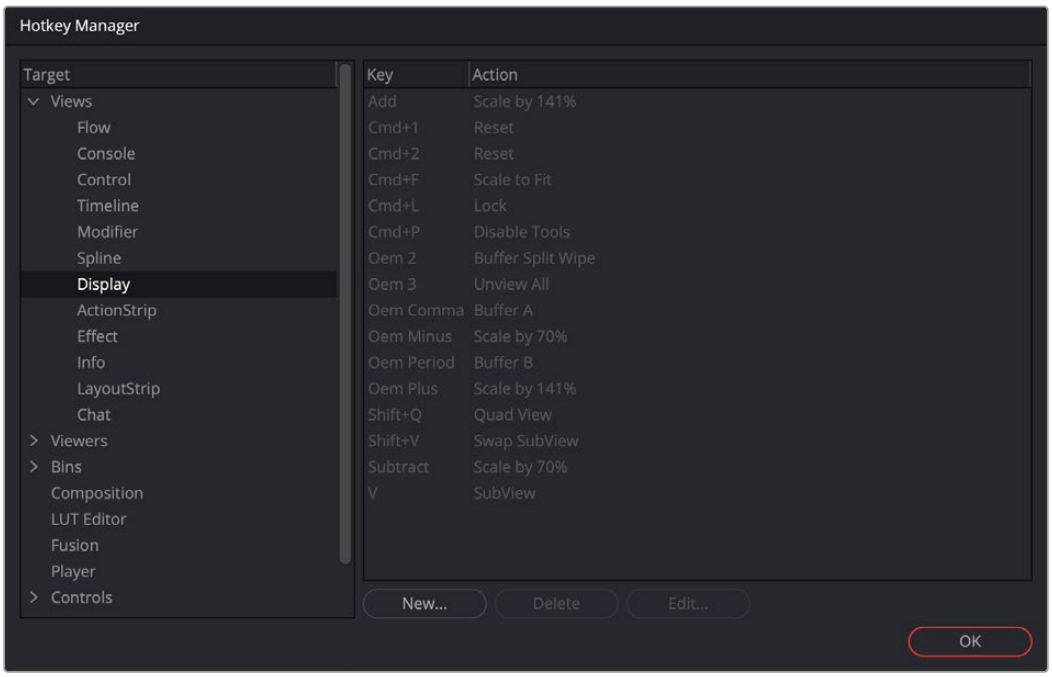
When checked, shot names stored in the EDL file are used to locate the footage.

Customization

The following section covers the customization of preferences that are not technically part of the Preferences window. Using Fusion Studio's Hotkey Manager window, you can customize the keyboard shortcuts, making the entire process of working in Fusion not only faster but potentially more familiar if you are migrating from another software application. You can also customize Fusion with environment variables to switch between different preferences files, allowing different working setups based on different users or job types. Both of these customization options are only available in Fusion Studio.

Shortcuts Customization

Keyboard shortcuts can be customized in Fusion Studio. You can access the Hotkey Manager by choosing Customize HotKeys from the View menu.

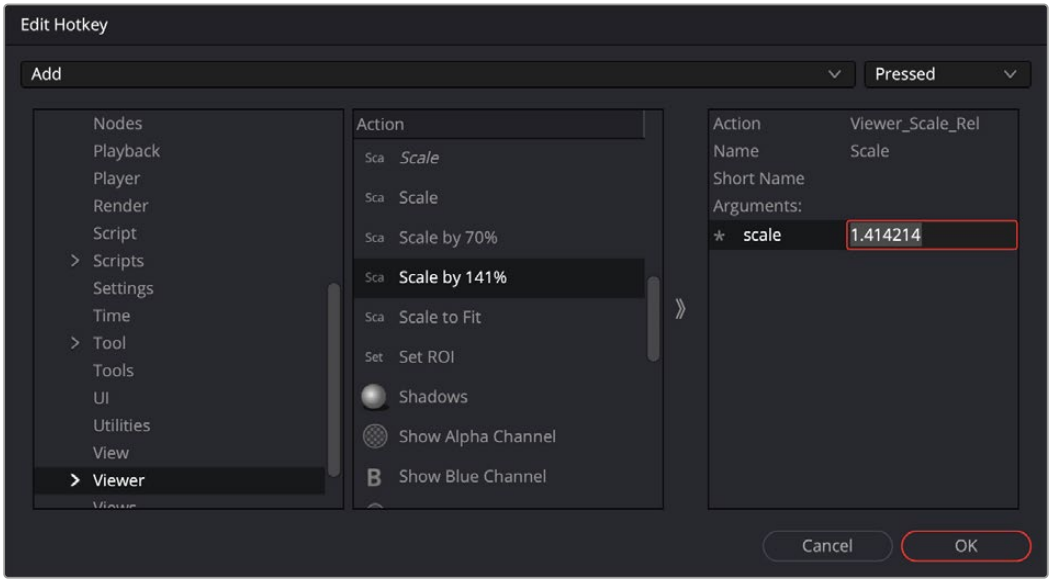


The Hotkey Manager

Fusion has active windows to focus attention on those areas of the interface, like the Node Editor, the viewers, and the Inspector. When selected, a gray border line will outline that section. The shortcuts for those sections will work only if the region is active. For example, Command-F in the

View will scale the image to fit the view area; in the Flow view, Command-F will open the Find tool dialog; and in the Spline editor, it will fit the splines to the window.

On the right is a hierarchy tree of each section of Fusion and a list of currently set hotkeys. By choosing New or Edit, another dialog will appear, which will give specific control over that hotkey.



The Hotkey Editor

Creating a new keyframe will give you the key combo to press, and this Edit Hotkey dialog will appear where the Action can be defined at top right: pressed, repeated, or released. The Name and abbreviated Short Name can be set, as can the Arguments of the action.

Customizing Preferences

Fusion Studio's preferences configure Fusion's overall application default settings and settings for each new composition. Although you access and set these preferences through the Preferences window, Fusion saves them in a simple text format called Fusion.prefs.

These default preferences are located in a \Profiles\Default folder and shared by all Fusion users on the computer. However, you may want to allow each user to have separate preferences and settings, and this requires saving the preferences to different locations based on a user login.

To change the saved location of the preferences file requires the use of environment variables.

Setting the Preferences Location

When you fist open Fusion, the environment variable FUSION_PROFILE_DIR defines the folder that contains the Profiles folder. If this variable defines a valid path, then the preferences are saved to this folder. If the FUSION_PROFILE_DIR does not exist, then Fusion attempts to create it. If it cannot create the path, then the preferences are stored in the default Path Map location: AllData:\Profiles.

Typically, all users share the same preferences.If you want each user to save separate preferences within their home folder, you must create another environment variable with the name FUSION_PROFILE (e.g., FUSION_PROFILE=jane). Using this second environment variable, Fusion will look for the preferences in the PROFILE_DIR of the user profile. Using a login script, you can make sure the FUSION_PROFILE is set to the name of the logged in user.

Creating a Master Preferences File

When working with multiple Fusion users in a studio, you may want to standardize on a few settings. Using the FUSION_MasterPrefs environment variable, you can create one or more site-wide preferences in addition to your local personal preferences.

FUSION_MasterPrefs must contain the full path to at least one preferences file. If you have multiple preferences paths, separate them using semicolons. Fusion does not write to these prefs files, and they may contain a subset of all available settings. You may change settings in these files and use them only where local prefs do not already exist unless you set the Locked flag.

Locking Preferences

If the line “Locked = true,” appears in the main table of a master file, all settings in that file are locked and override any other preferences. Locked preferences cannot be altered by the user.



2D Compositing

FUSION

CONTENTS

16 Controlling Image Processing and Resolution 419
17 Managing Color for Visual Effects 430
18 Understanding Image Channels 443
19 Compositing Layers in Fusion 479
20 Rotoscoping with Masks 504
21 Paint 528
22 Using the Tracker Node 556
23 Planar Tracking 592
24 Using Open FX, Resolve FX, and Fuse Plugins 599

Chapter 16

Controlling Image Processing and Resolution

This chapter covers the overall image-processing pipeline. It discusses color bit-depth and how to control the output resolution in a resolution-independent environment.

Contents

Fusion's Place in the DaVinci Resolve Image-Processing Pipeline 420
Source Media into the Fusion Page 420
Forcing Effects into the Fusion Page 421
Output from the Fusion Page to the Color Page 421
What Viewers Show in Different DaVinci Resolve Pages 421
Managing Resolution In Fusion 422
Changing the Resolution of a Clip 422
Compositing with Different-Resolution Clips 423
Sizing Between DaVinci Resolve Pages 423
Color Bit Depths 424
Understanding Integer vs. Float 424
Setting Color Depth in Fusion Studio 425
Combining Images with Different Color Depths 427
Advantages of Floating-Point Processing 427

Fusion's Place in the DaVinci Resolve Image-Processing Pipeline

When working in a single unified environment like DaVinci Resolve, it is important to understand the order of operations among the pages. DaVinci Resolve exposes some of this via the order of the page buttons at the bottom of the screen, with the Media, Cut, and Edit page at the beginning of the chain and the Color, Fairlight, and Deliver page at the end. However, this isn't the whole story, especially when it comes to the Fusion page. The following sections describe where the Fusion page fits in the image-processing chain of DaVinci Resolve.

Source Media into the Fusion Page

For ordinary, single clips coming in from the Edit or Cut page, the MediaIn node in the Fusion page represents the source media, as modified by the Clip Attributes window. Although you select the clip from the Edit or Cut page Timeline, in the Fusion page, the clip is accessed from the Media Pool.

TIP: The decoding or debayering of RAW files occurs prior to all other operations, and as such, any RAW adjustments will be displayed correctly in the Fusion page.

This means you have access to the entire source clip in the Fusion page, but the render range is set to match the duration of the clip in the Timeline. You also use the full resolution of the source clip, even if the Timeline is set to a lower resolution. However, none of the Edit or Cut page Inspector adjustments carry over into the Fusion page, with the exception of the Lens Correction adjustment.

When you make Zoom, Position, Crop, or Stabilization changes in the Edit or Cut page, they are not visible in the Fusion page. The same applies to any Resolve FX or Open FX third-party plugins. If you add these items to a clip in the Edit or Cut page, and then you open the Fusion page, you won't see them taking effect. All Edit and Cut page timeline effects and Inspector adjustments, with the exception of the Lens Correction adjustment, are computed after the Fusion page but before the Color page. If you open the Color page, you'll see the Edit and Cut page transforms and plugins applied to that clip, effectively as an operation before the grading adjustments and effects you apply in the Color page Node Editor.

With this in mind, the order of effects processing in the different pages of DaVinci Resolve can be described as follows:



TIP: Retiming applied to the clip in the Edit page Timeline is also not carried over into the Fusion page.

Forcing Effects into the Fusion Page

There is a way you can force clips with Edit page Inspector adjustments, plugins, retiming, and Color page grades into the Fusion page, and that is to turn that clip into a compound clip or a Fusion clip. When Edit page effects and Color page grading are embedded within compound or Fusion clips, MediaIn nodes corresponding to compound clips route the effected clip into the Fusion page. However, bringing a compound clip or Fusion clip into the Fusion page changes the resolution of the source clip to match the Timeline resolution. For more information, see the section "Sizing Between DaVinci Resolve Pages" in this chapter.

Output from the Fusion Page to the Color Page

The composition output from the Fusion page's MediaOut node are passed on via the Color page's source input, with the sole exception that if you've added plugins to that clip in the Edit or Cut page, then the handoff from the Fusion page to the Color page is as follows:



What Viewers Show in Different DaVinci Resolve Pages

Owing to the different needs of compositing artists, editors, and colorists, the viewers show different states of the clip.

- **The Edit page source viewer:** Always shows the source media, unless you're opening a compound clip that's been saved in the Media Pool. If Resolve Color Management is enabled, then the Edit page source viewer shows the source media at the Timeline color space and gamma.
- **The Edit page Timeline viewer:** Shows clips with all Edit page effects, Color page grades, and Fusion page effects applied, so editors see the program within the context of all effects and grading.
- **The Fusion page viewer:** Shows Media Pool source clips at the Timeline color space and gamma, including Resolve Color Management adjustments. but no Edit page Inspector adjustments or Resolve FX effects and no Color page grades.
- **The Color page viewer:** Shows clips with all Edit page effects, Color page grades, and Fusion page effects applied.

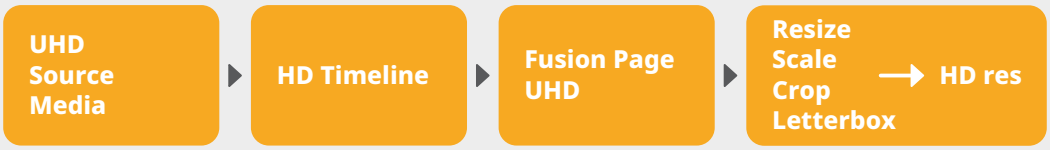
Managing Resolution In Fusion

There is no formal resolution for a comp in Fusion.

- In the Fusion page, the default resolution for Fusion-generated images, like the Background tool, Fast Noise, and Text+, etc., is determined by the Timeline resolution.
- In Fusion Studio, the resolution for generators is set in Preferences > Frame Format, where you can select a standard resolution from the drop-down menu, or set the Width and Height manually.
- The resolution of a composition in the Fusion page is initially determined by the source resolution of the input image.

For example, if you have a Timeline set to HD resolution (1920 x 1080) and have a clip on the timeline whose source resolution is UHD (3840 x 2160), opening this clip in the Fusion page will set the composition's resolution to UHD (3840 x 2160).

All Fusion-generated images in this composition will be UHD resolution.



Fusion timeline sizing for the Fusion page

TIP: The output of the Fusion page is placed back into the Edit page Timeline based on DaVinci Resolve's Image Sizing setting. By default, DaVinci Resolve uses an image sizing setting called Scale to Fit. This means that even if the Fusion page outputs a 4K composition, it conforms to 1920 x 1080 if that is what the project or a particular Timeline is set to. Changing the image sizing setting in DaVinci Resolve's Project Settings affects how Fusion compositions are integrated into the Edit page Timeline.

Changing the Resolution of a Clip

If your comp uses a single image, you can change the pixel output resolution in several ways. Three common tools that change the pixel resolution of a clip are the Resize, Scale, and Crop nodes. A fourth node, Letterbox, is less commonly used but also changes the pixel resolution of a clip.

These four nodes are located in the Transform category of the Effects library. Resize is also located in the toolbar.

- **Crop:** Sets the output resolution of the node using a combination of X and Y size along with X and Y offset to cut the frame down to the size you want. Crop removes pixels from the image, so if you later use a Transform node and try to move the image, those pixels are not available.
- **Letterbox:** Sets the output resolution of the node by adding horizontal or vertical black edges where necessary to format the frame size and aspect ratio.
- **Resize:** Sets the output resolution of the node using absolute pixels.
- **Scale:** Sets the output resolution of the node using a relative percentage of the current input image size.

TIP: To change resolution and reposition a frame without changing the pixel resolution of a clip, use the Transform node.

Compositing with Different-Resolution Clips

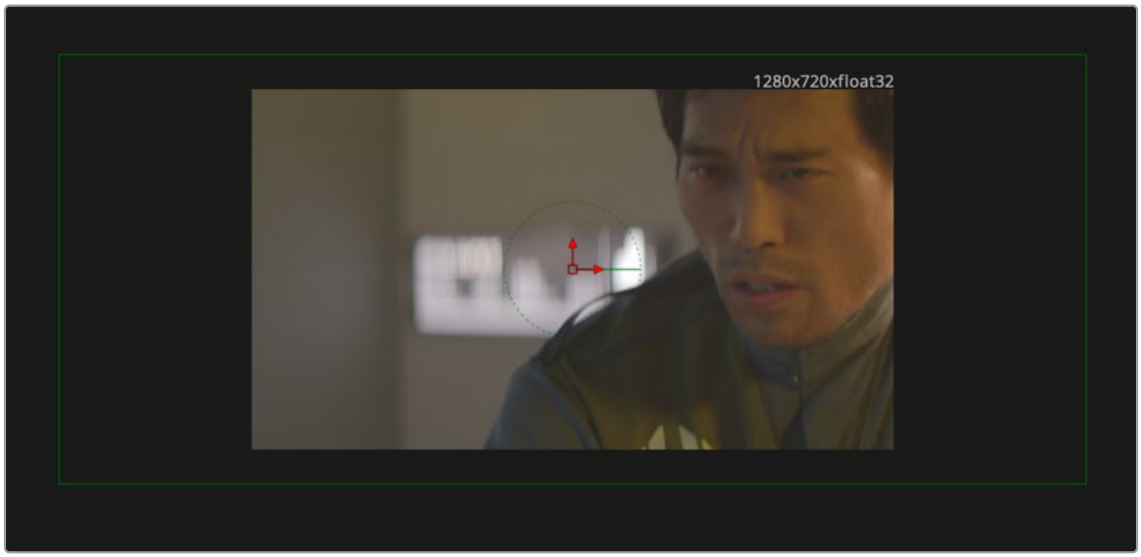
When you composite images with different resolutions using the Merge node, the image that's connected to the orange background input determines the output resolution of the Merge node.

Often, it's easiest to control the comp resolution right at the start by connecting a node with the desired output resolution you want to the orange background input on the Merge node. A Background node is often used in this situation because it consumes meager system resources.



A Background node determines the output resolution of the merge

The Background node sets the output size, and the foreground image is cropped if it is larger.



A Background node created at 1280 x 720 crops the larger foreground.
However, all the pixels of the larger foreground are available for repositioning.

Sizing Between DaVinci Resolve Pages

The order of sizing operations between DaVinci Resolve pages is a bit more nuanced. It's important to understand which sizing operations happen in the Fusion page, and which happen after, so you know which effects alter the image that's input to the Fusion page, and which effects alter the page's output.

For example, lens correction, while not strictly sizing, is nonetheless an effect that changes how the image begins in your Fusion composition. However, the Edit or Cut page stabilization function is an effect that comes after the Fusion page, so it does not appear in the composition you're creating.

The order of sizing effects in the different pages of DaVinci Resolve can be described as follows:



Sizing with Compound and Fusion Clips

Another way to modify the resolution before clips get handed off from the Edit page to the Fusion page is to create a compound clip or a Fusion clip. Both compound clips and Fusion clips change the working resolution of the individual clips to match the Timeline resolution. For instance, if two 4K clips are stacked one on top of the other in an HD timeline, creating a compound or Fusion clip resizes the clips to HD. The full resolution of the individual 4K clips is not available in Fusion and is therefore handed off to the Color page at the rescaled size. To maintain the full resolution of source clips, bring only one clip into the Fusion page from the Edit or Cut page Timeline, and then bring other clips into the Fusion composition using the Media Pool. Of course, if your clips are full HD and your timeline is full HD, then creating a Fusion clip or compound clip does not affect the resolution.



Fusion timeline sizing for Fusion and Compound clips

Color Bit Depths

The term bit depth describes how many colors are available in the color palette used to make up an image. The higher the bit depth, the greater the precision of color in the image, and therefore the greater the color reproduction. The higher precision is most apparent in gradients with subtle changes. Lower bit-depth gradients have noticeable banding artifacts, whereas higher bit-depth images can reproduce more colors, so fewer, if any, banding artifacts occur. The Fusion page within DaVinci Resolve always uses 32-bit float bits per channel precision to process images. However, in Fusion Studio you can choose to process images with 8-bit integer, 16-bit integer, 16-bit float, and 32-bit float bits per channel. Although always working at 16-bit float or 32-bit float will produce the best quality, it may be more efficient to use a lower bit depth if your images are 8-bit or 16-bit integer formats to begin with.

Understanding Integer vs. Float

Generally, 8-bit integer color processing is the lowest bit depth you'll come across for video formats. 8-bit images come from older or consumer-grade video equipment like mobile phones and camcorders. If you try to perform any significant gamma or color correction on 8-bit images, you can often see more visible banding.

16-bit integer color depth doubles the amount of precision, eliminating problems with banding. Although you can select 16-bit integer processing for an 8-bit clip, it does not reduce banding that already exists in the original file. Still, it can help when adding additional effects to the clip. This sounds like the best solution until you realize that many digital cameras like Blackmagic Design URSA Mini Pro

and others record in formats that can capture over-range values with shadow areas below 0.0 and super highlights above 1.0, which are truncated in 16-bit integer.

The 16-bit float color depth sacrifices a small amount of the precision from standard 16-bit integer color depth to allow storage of color values less than 0 and greater than 1.0. 16-bit float, sometimes called half-float, is most often found in the OpenEXR format and contains more than enough dynamic range for most film and HDR television purposes yet requires significantly less memory and processing time than is required for full float, 32-bit images.



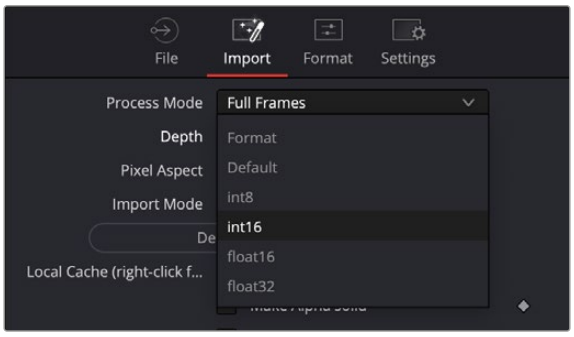
Preserving over-range values allows you to change exposure while maintaining highlights

Processing at 32-bit float can work with shadow areas below 0.0 and highlights above 1.0, similar to 16-bit float, except with a much greater range of precision but also much greater memory and processing requirements.

Setting Color Depth in Fusion Studio

As we said earlier, DaVinci Resolve always processes at 32-bit float bits per channel; however, you can use less memory and still achieve more-than-acceptable results using the Performance Mode setting located in the User > Playback Preferences panel.

Fusion Studio automatically uses the color depth that makes the most sense for each file format. For example, if you read in a JPEG file from disk, then the color depth for the Loader is set to 8 bits per channel. Since the JPEG format is an 8-bit format, loading the image at a greater color depth would generally be wasteful. If a 16-bit TIFF is loaded, the color depth is set to 16 bits. Loading a DPX file defaults to 32-bit float, whereas OpenEXR generally defaults to 16-bit float. However, you can override the automatic format color depth using the settings found in the Import tab of the Loader node's Inspector. The Loader's Inspector, as well as the Inspector for images generated in Fusion (i.e., text, gradients, fast noise, and others), has a Depth menu for 8-bit, 16-bit integer, 16-bit float, and 32-bit float.

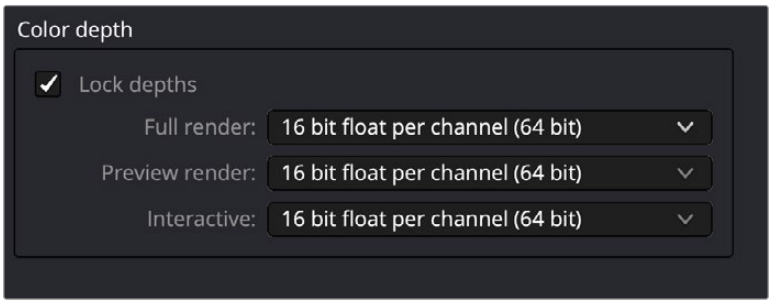


The Loader's Inspector Color Bit Depth settings

Configuring Default Color Depth Preferences

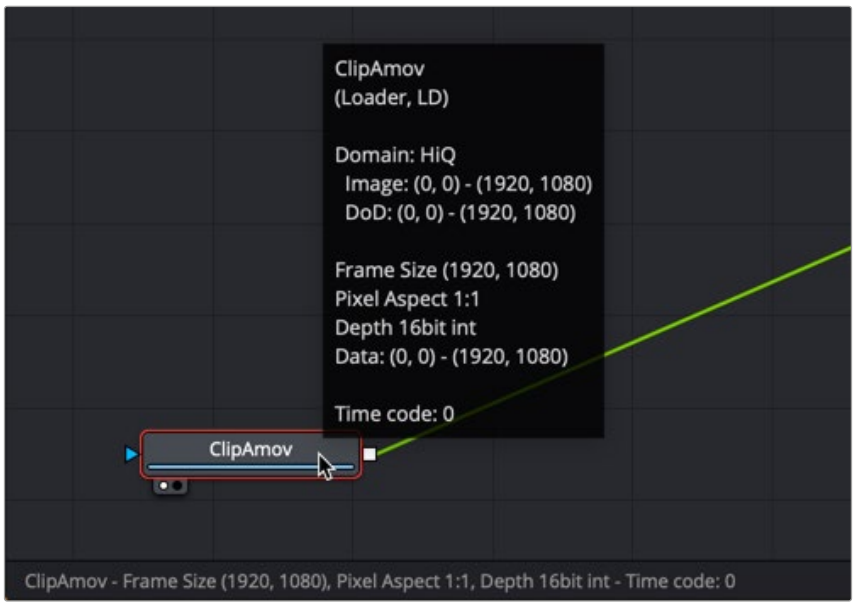
The default color depth setting forces the tool to process based on the settings configured in the Node Editor's Frame Format preferences. These are used to set a default value for color depth, applied when a Generator tool is added to the Node Editor. There are three drop-down menus to configure color depth in the preferences. They specify the different color depths for the interactive session, final renders, and preview renders.

To improve performance as you work on your comp, you can set the Interactive and Preview depth to 8-bits per channel, while final renders can be set to 16-bit integer. However, if your final render output is 16-bit float or 32-bit float, you should not use the integer options for the interactive setting. The final results may look significantly different from interactive previews set to integer options.



The Frame Format Color Depth settings

If you aren't sure what the color depth process is for a tool, you can position the pointer over the node's tile in the Node Editor, and a tooltip listing the color depth for that node will appear on the Status bar.



Hover over a node to view its Color Bit Depth setting.

TIP: When working with images that use 10-bit or 12-bit dynamic range or greater, like Blackmagic RAW or Cinema DNG files, set the Depth menu in the Inspector to 16-bit float or 32-bit float. This preserves highlight detail as you composite.

Combining Images with Different Color Depths

You can combine images with different color depths in a single composition. When images of different color depths are combined, the image from the background input of the node determines the bit depth output, and the foreground image is adjusted to match.

Advantages of Floating-Point Processing

There are two major advantages to floating-point processing that make the additional RAM requirements and longer render times worth your while. The first benefit is that floating-point values are more accurate than integer values. The second benefit is the preservation of shadow and highlight values that go beyond the normal tonal range.

Greater Accuracy

Using 16- or 32-bit floating-point processing prevents the loss of accuracy that can occur when using 8- or 16-bit integer processing. The main difference is that integer values cannot store fractional or decimal values, so rounding occurs in all image processing. Floating-point processing allows decimal or fractional values for each pixel, so it is not required to round off the values of the pixel to the closest integer. As a result, color precision remains virtually perfect, regardless of how many operations are applied to an image.

If you have an 8-bit pixel with a red value of 75 (dark red) and that pixel is halved using a Color Correction tool, the pixel's red value is now 37.5. Since you cannot store decimal or fractional values in integers, that value is rounded off to 37. Doubling the brightness of the pixel with another Color Correction tool should bring back the original pixel value of 75 but because of rounding 37 x 2 is 74. The red value lost a full point of precision due to integer rounding on a very simple example. This is a problem that can result in visible banding over several color corrections. Similar problems arise when merging images or transforming them. The more operations that are applied to an image, the more color precision is lost to rounding when using 8- or 16-bit integer processing.

Accessing Extended Highlights and Shadows

Increasingly more productions are capturing out-of-range images thanks to digital cinema cameras like the Blackmagic URSA Mini Pro and even the Pocket Cinema 6K camera. These cameras capture very high dynamic range RAW images and maintain color detail even in heavily over or underexposed frames. The extended white color detail can also give very nice, natural results when blurred, glowed, color corrected, or even just when faded or dissolved. While it is possible to work with these RAW images using integer data, doing so results in the loss of the extended range values, losing all detail in the highlights and shadows. Float processing makes working with logarithmic RAW images considerably easier by preserving highlight and shadow detail.

If you have an 8-bit pixel that has a red value of 200 (bright red) and a Color Gain tool is used to double the brightness of the red channel, the result is 200 x 2, or 400. However, 8-bit color values are limited to a range of 0 through 255. So the pixel's value is clipped to 255, or pure red. If now the brightness is halved, the result is half of 255, or 127 (rounded), instead of the original value of 200.

When processing floating-point colors, pixel values brighter than white or darker than black are maintained. There is no value clipping. The pixel is still shown in the viewer as pure red, but if float processing is used instead of 8-bit, the second operation where the gain was halved would have restored the pixel to its original value of 200.

Using Float with 8-Bit HD Video

There is also some value to using float color depths with an 8-bit HD video when the images require a lot of color correction. Using float helps maintain precision by avoiding the rounding errors common to 8-bit processing, as described above.

Detecting Extended Highlight and Shadow Values

Although floating-point processing preserves extended values below 0.0 and greater than 1.0, also called “out-of-range values,” the viewer still displays them as black or white. This can make it difficult for you to determine the overall dynamic range of an image.

To discover whether there are out-of-range values in a viewed image:

Right-click in the viewer and choose Options > Normalized Color Range.



Use the Normalized Color Range pop-up menu to detect out-of-range images.

Enabling this display mode rescales the color values in the image so that the brightest color in the image is remapped to a value of 1.0 (white), and the darkest is remapped to 0.0 (black).

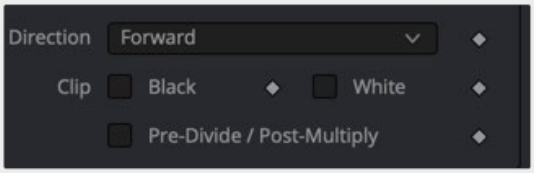
The 3D Histogram subview can also help visualize out-of-range colors in an image.

For more information, *see Chapter 7, “Using Viewers.”* in the Fusion Reference Manual.

Clipping Out-of-Range Values

When processing in floating point, there may be situations where the out-of-range values in an image need to be clipped. The Brightness/Contrast tool provides checkboxes that can be used to clip out-of-range values to 0 or 1.

For example, there may be files that contain out-of-range alpha values. Since the alpha channel represents the opacity of a pixel, it makes little sense to be more than completely transparent or more than fully opaque, and compositing such an image may lead to unexpected results. To easily clip alpha values below 0 and above 1, add a Brightness/Contrast toolset to Clip Black and Clip White, with only the Alpha checkbox selected.



Clip White and Clip Black settings in Brightness/Contrast can be used to clip mattes.

Alternatively, you can clip the range by adding a Change Depth node and switching to 8-bit or 16-bit integer color depths.

Chapter 17

Managing Color for Visual Effects

This chapter discusses LUTs, color space conversions, and the value of compositing with linear gamma while previewing the image in the viewer using the gamma of your choice.

Contents

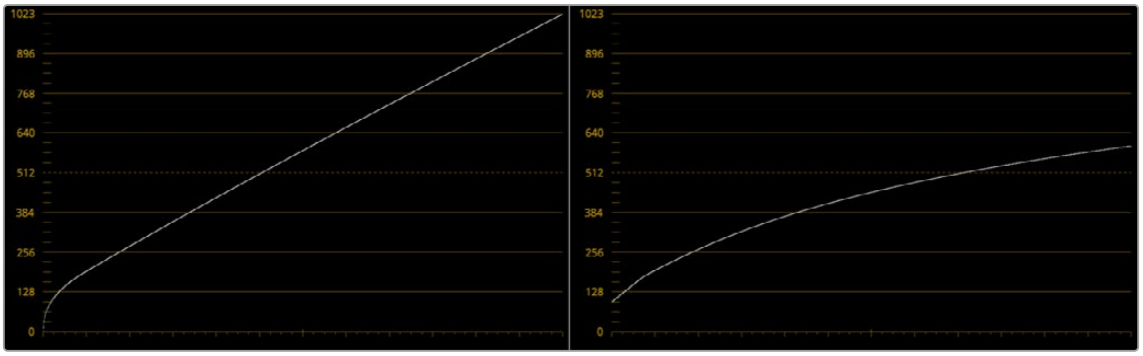
Color Management 431
All Compositing Is Math 432
Introducing Color Management in Fusion 432
Converting to Linear Gamma 433
Applying LUTs to a Viewer 436
Using Resolve Color Management 438
Using ACES Color Management in Resolve 439
Using OCIO for ACES Color Management in Fusion 440
Applying OCIO LUTs in the Viewer 441

Color Management

The simplified goal of color management is to make sure that the image you see on your computer screen is what your audience sees when they view it on a television, cinema screen, or mobile device. For such a simple goal, problems arise that you, as the compositor, must manage. These problems begin with the fact that our eyes see luminance one way, and a computer display represents luminance differently.

Each capture device records images using a nonlinear tonal curve or gamma curve to compensate for this difference. Specifically, Rec. 709 HD gamma curves are designed so that when shown on HD displays, the images have built-in compensation for the display. The result is that HD images on HD displays appear normal to us.

Digital cinema cameras have taken the concept of gamma curves further. They use gamma curves as a way to maximize the bit depth of an image and store a wider dynamic range. Digital cinema cameras' gamma curve (often collectively referred to as *log gamma*), give more attention to the darker mid-tones where the human eye is most sensitive. This allows them to save images with brighter highlights and more detail in shadows.



A Rec. 709 HD gamma curve (left) and a nonlinear, or log gamma, curve (right)

The problem is that these images do not look normal on any monitor. Clips recorded with a log gamma curve typically have a low contrast, low saturated appearance when viewed on an sRGB computer display or Rec. 709 HD video monitor. This problem is easy to fix using a LookUp Table, or LUT. A LUT is a form of gamma and color correction applied to the viewer to normalize how the image is displayed on your screen.



A clip displayed with a nonlinear, log gamma curve (left) and corrected in the viewer using a LUT (right)

All Compositing Is Math

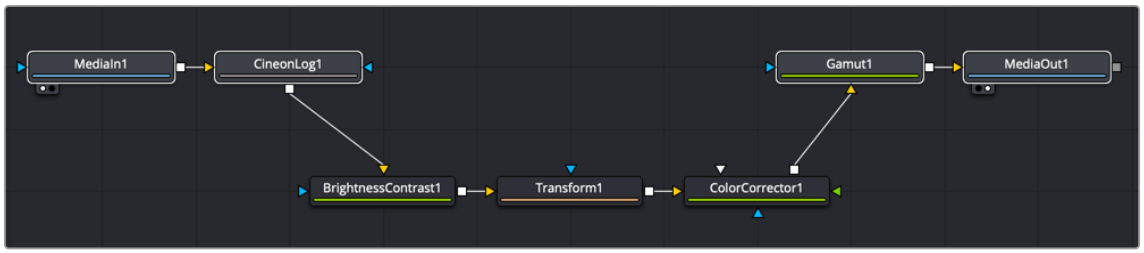
The LUT applied to the viewer only solves the problem in the viewer. Now we come to the larger problem. The image data is still using a log gamma curve. Fusion, and every other image-processing application, operates with the assumption that the image data has linear gamma. The image-processing filters you apply to images use standard math functions, like 1 + 1 = 2. Common operations such as those that add pixels like Brightness, or divide pixels (a.k.a. “unpremultiply”), or composite modes that include multiplication such as “screen,” and many other compositing tasks assume that 1 + 1 always equals 2. In other words, if you perform an operation that doubles the amount of brightness, then every pixel should be twice as bright. However, if you are starting with a nonlinear gamma curve, pixels are not being adjusted linearly, so some pixels might end up 1.2 x as bright, 1.7 x as bright, or 2.4 x as bright. Now the math is 1 + 1 = 3. The further your images are from linear gamma, the more pronounced the math error. A Rec. 709 HD clip shows less error than a log gamma clip from a digital cinema camera. However, an error is still an error, and the more compositing operations you perform on the image, the more the error is compounded.

You can see a more practical example when you apply filtering effects, such as a blur, to an image with any gamma setting. The image probably looks fine. However, if you convert the image to a linear gamma first and then apply the blur, the images (especially those with extremely bright areas) are processed with greater accuracy, and you should notice a different and superior result.

The answer to these problems is to manage your color before compositing.

Introducing Color Management in Fusion

Images loaded into Fusion by default are not color managed. The image is displayed directly from the file to the viewer without any interpretation or conversion. However, Fusion includes nodes that convert the output of each image to linear gamma at the beginning of your composite. The same nodes can convert from linear back to your desired output gamma at the end of your composite, just prior to the Saver or MediaOut node.



A log clip converted to linear and then converted for output

To manually set up a linear gamma workflow in Fusion:

1. Use a Gamut or CineonLog node after all MediaIn or Loader nodes to convert them to linear.
2. Apply a GAMUT View LUT to the viewers to correct the display of a linear image to sRGB or Rec. 709.
3. Before a Saver or MediaOut node, insert a Gamut or CineonLog node to convert from linear to your target output format.

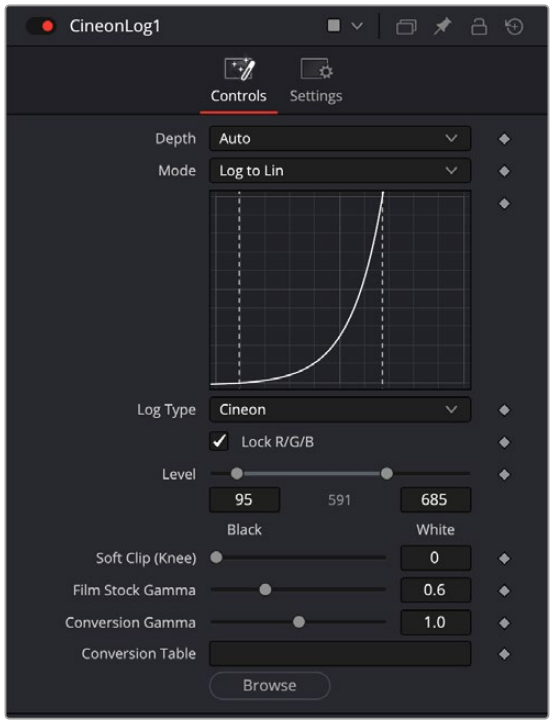
Converting to Linear Gamma

Whether an image comes from the Edit page in DaVinci Resolve, or from a Loader in Fusion Studio, the color and gamma are read directly into Fusion, with no modification. For some simple operations on sRGB or Rec 709 clips, this may be fine, but it's not always the ideal way to work, especially for log-encoded media. The ideal way to work with log-encoded media is to convert images to linear gamma, since the majority of image-processing operations in Fusion expect gamma to be linear and will produce superior results.

> **TIP:** 3D rendered CGI images are often generated as EXR files with linear gamma, and converting them is not necessary. However, you should check your specific files to make sure they are using linear gamma.

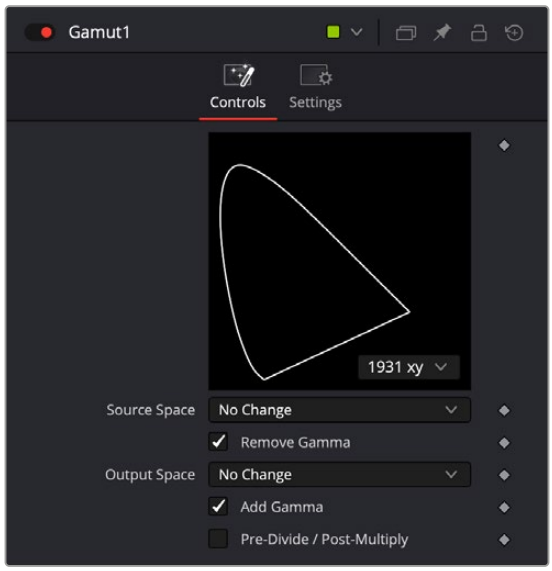
Fusion includes several kinds of nodes you can use to convert the image out of each MediaIn or Loader node to linear gamma at the beginning of your composite, and then convert from linear back to your desired output gamma at the end of your composite. These include:

- **CineonLog node:** The CineonLog node, found in the Film category of the Effects Library, performs a conversion from any of the formats in the Log Type menu to linear, and also reverses the process, adding log gamma back to a clip. This is most often used for images coming from common digital cinema cameras like BlackMagic Design, Arri, or Red. The CineonLog node is added directly after a MediaIn or Loader node. The Mode menu chooses the direction of the conversion to or from linear.



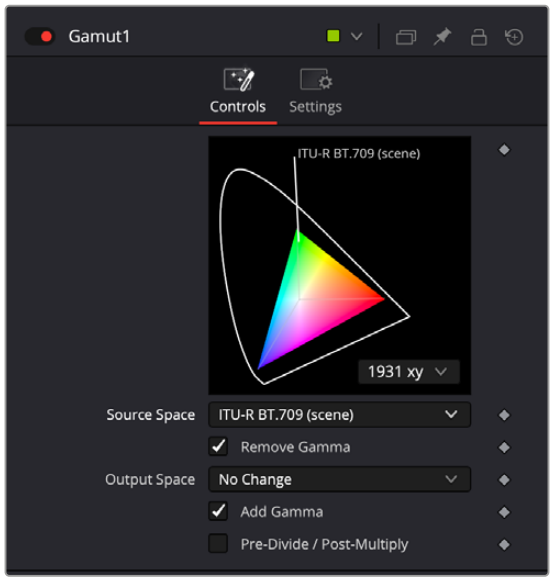
Add a CineonLog tool to convert log gamma curves to linear and vice versa.

- **Gamut node:** The Gamut node, found in the Color category of the Effects Library, lets you perform linear conversions based on color space. This node converts to linear or from linear and is often inserted after a MediaIn or Loader node or just before a MediaOut or Saver node. Depending on where you insert the node, you either choose from the Source Space controls or the Output Space controls.



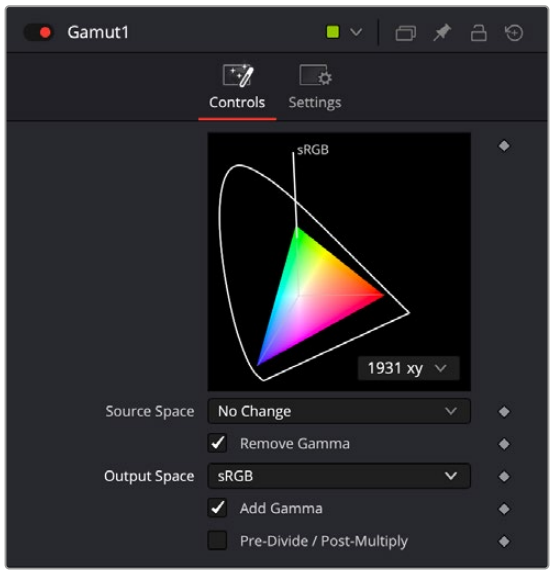
Add a Gamut tool to convert gamma curves to linear based on color space.

When converting media to linear gamma, set the Source Space menu to the color space of your source material. For instance, if your media is full 1080 HD ProRes, then choose ITU-R BT.709 (scene) for gamma of 2.4. Then, enable the Remove Gamma checkbox if it isn't already enabled, to use linear gamma.



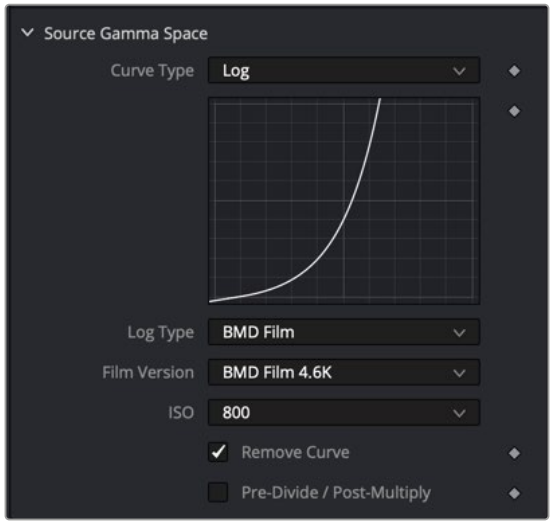
Source Space is used to convert to linear gamma.

When converting from linear gamma for output, you insert the Gamut node before your output node, which is a Saver in Fusion Studio or a MediaOut node in DaVinci Resolve's Fusion page. Make sure the Source Space menu is set to No Change, and set the Output Space to your output color space. For instance, if your desired output is full 1080 HD, then choose either sRGB or ITU-R BT.709 (scene) for gamma of 2.4. Then, enable the Add Gamma checkbox if it isn't already enabled, to format the output of the Gamut node for your final output.



Output Space is used to convert from linear gamma.

- **MediaIn and Loader nodes:** MediaIn and Loader nodes have Source Gamma Space controls in the Inspector that let you identify and remove the gamma curve without the need to add another node. If your files include gamma curve metadata like RAW files, the Auto setting for the Curve Type drop-down menu reads the metadata and uses it when removing the gamma curve. When using intermediate files or files that do not include gamma curve metadata, you can choose either a log gamma curve by choosing Log from the Curve Type menu or a specific color space using the Space option from the menu. Clicking the Remove Curve checkbox then removes the gamma curve, converting the image to linear gamma.



MediaIn and Loader nodes include a Remove Curve checkbox in the Inspector.

- **FileLUT node:** The FileLUT node, found in the LUT category of the Effects Library, lets you do a conversion using any LUT you want, giving you the option to manually load LUTs in the ALUT3, ITX, 3DL, or CUBE format to perform a gamma and gamut conversion. Although LUTs are very commonly placed at the end of a node tree for final rendering, you'll get more accurate gamma and color space conversions using the Gamut and CineonLog nodes to transform your MediaIn and Loader nodes into linear.

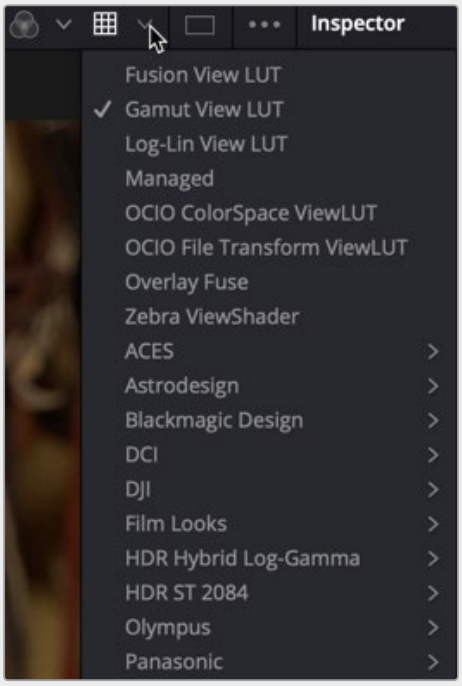
Applying LUTs to a Viewer

Images converted to a linear gamma don't look correct. They usually look very dark, with extremely bright highlights and oversaturated colors. Happily, even though the image may appear to be incorrect, the fact that Fusion can work entirely with floating-point color data means that you're not actually clipping or losing any image data. It just looks completely wrong when viewing the linear state of your image data directly.



A clip displayed with a nonlinear, log gamma curve (left) and the clip transformed to linear gamma (right)

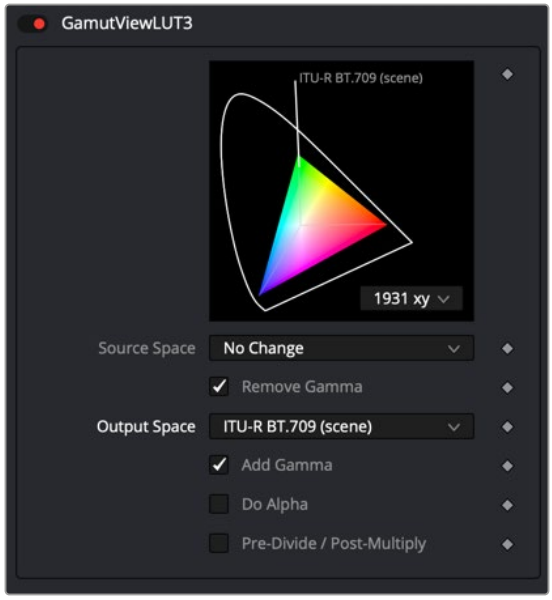
It would be impossible to work if you couldn't view the image as it's supposed to appear within the final gamut and gamma you'll be outputting. For this reason, each viewer has a LUT menu that lets you enable a "preview" color space and/or gamma conversion, while the node tree is processing correctly in linear gamma.



Apply a Gamut View LUT to preview the image in your intended output color space.

To preview the images in the viewer using sRGB or Rec. 709 color space:

1. Enable the LUT button above the viewer.
2. From the Viewer LUT drop-down menu, choose either a Gamut View LUT, or a LUT from the VFX IO category that transforms linear to Rec. 709 or sRGB.
3. If you choose the Gamut View LUT, then choose Edit from the bottom of the LUT menu to configure the LUT.
4. In the LUT Editor, set the Output Space to the target color space you want.
5. Enable the Add Gamma checkbox to apply the gamma curve based on the selected color space.



Gamut View LUT Output Space set to Rec. 709 for HD with the Add Gamma checkbox enabled

TIP: If your monitor is calibrated differently, you need to select a LUT that matches your calibration.

Whether you use the Gamut View LUT or a LUT for your specific monitor calibration, you can save the viewer setup as the default.

To Save the Gamut LUT setup as the default viewer setup:

- Right-click in the viewer, and then choose Settings > Save Defaults.

For every comp, the viewer will now be preconfigured based on the saved defaults.

For more information on Viewer LUTs, *see Chapter 7, "Using Viewers."* in the Fusion Reference Manual.

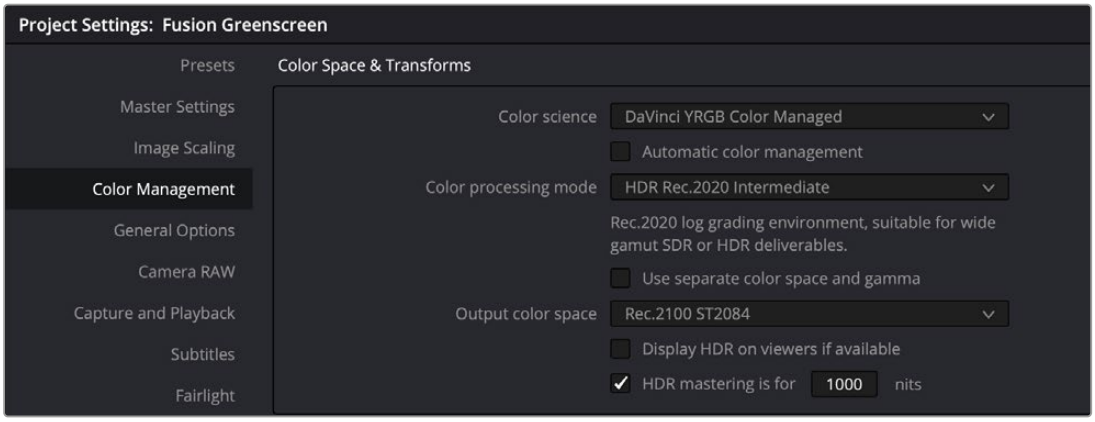
Using Resolve Color Management

If you're using the Fusion page within Davinci Resolve, you have the option of enabling DaVinci Resolve's scene-referred color management, instead of inserting Gamut and CineonLog nodes. When DaVinci YRGB Color Managed (RCM) is enabled, the color of MediaIn nodes in the Fusion page is handled differently. The RCM automatically determines the input color spaces of all files used for all MediaIn nodes in the Fusion page, which then automatically get converted to linear gamma. The MediaOut node then gets converted back into the Color Processing mode to get graded in the Color page or further edited in the Edit page Timeline.

Despite the seeming complexity of color management, using RCM is actually simple. In essence, all you have to do is (A) turn on RCM, and (B) choose the Color Processing mode and Output Color Space combination you want to use.

To enable Resolve Color Management:

1. Open the Color Management panel of the Project Settings.
2. Choose DaVinci YRGB Color Managed from the Color Science drop-down menu.
3. Check the Automatic color management box for a simplified selection, or un-check the box to adjust the Color Processing mode and Output color spaces manually.
4. Set the desired options for the Color Processing mode and Output Color Space RCM settings.



When DaVinci YRGB Color Managed is enabled, Timeline color space is used for all MediaIn nodes in the Fusion page.

To override the input color space for differently recorded clips in the Media Pool:

1. Enable DaVinci YRGB Color Management as explained above.
2. Save and close the Settings dialog.
3. In the Media Pool, select the clip or clips you want to assign a new Input Color space.
4. Right-click one of the selected clips.
5. Choose the Input Color Space that corresponds to those clips from the contextual menu.

Using RCM eliminates a few steps, since the input color space math used to transform the source preserves all wide-latitude image data, making highlights easily retrievable without any extra steps. With RCM enabled, there is no need to insert CineonLog or Gamut nodes while in the Fusion page. The transforms from and to linear are done automatically based on the RCM settings. Switching to the Fusion page converts the images to linear and enables the LUT button in the viewers with the Managed LUT selected. The Managed LUT uses the RCM settings to take a linear image and display it based on the RCM output color space.

For more information on Resolve Color Management, see Chapter 9, “Data Levels, Color Management, and ACES,” in the *DaVinci Resolve 20 Reference Manual*

Using ACES Color Management in Resolve

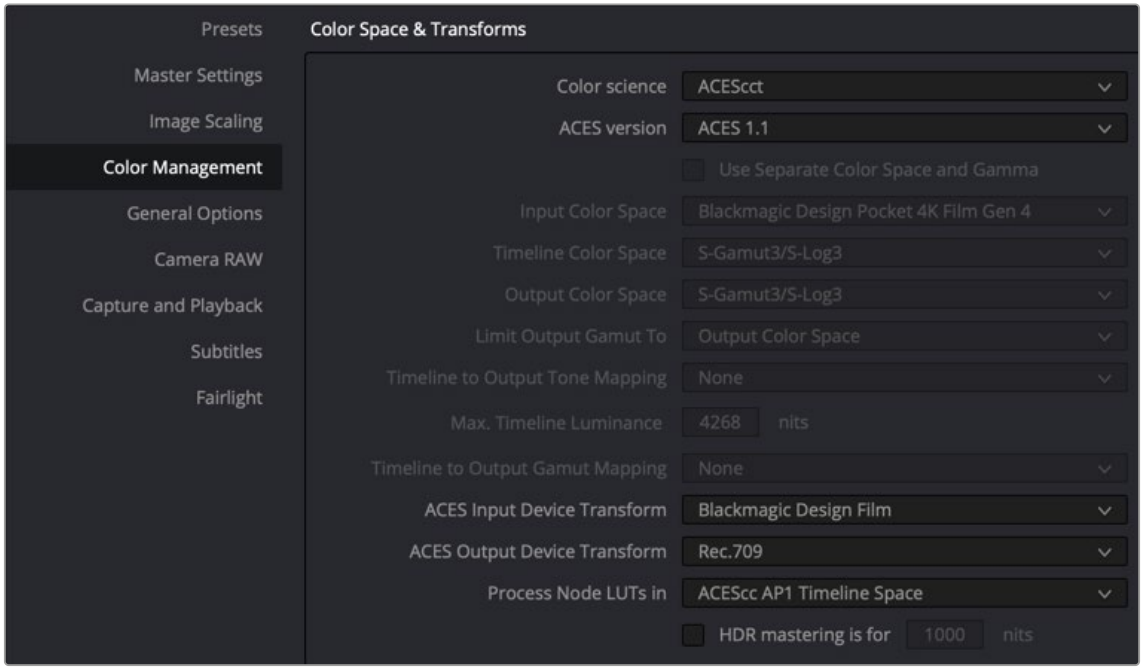
The ACES (Academy Color Encoding Specification) color space is another standard for managing color throughout an entire production. It’s designed to make start-to-finish, scene-referred color management a reality for digital cinema workflows. Just like DaVinci’s RCM, ACES makes it easier to extract high-precision, wide-latitude image data from raw camera formats, to preserve high-quality image data from acquisition through the color grading process and to output high-quality data for broadcast viewing, film printing, or digital cinema encoding.

ACES works by assigning an IDT (Input Device Transform) to every camera and acquisition device. The IDT specifies how media from that device is converted into the ACES color space. At the end of the pipeline, an ODT (Output Device Transform) is applied to convert the image data from ACES color space into the gamut of your final output.

Similar to setting up RCM, DaVinci Resolve’s color management project settings can be configured for ACES, which carries through the Edit, Fusion, and Color pages.

NOTE: When using Fusion Studio, the OpenColorIO (OCIO) framework is used for ACES color management.

The Color Science drop-down menu in the Color Management panel of the Project Settings is used to set up the ACES color management in DaVinci Resolve.



When ACES is enabled, IDT and ODT are used to identify input and output devices.

- **Color Science:** Using this drop-down menu, you can choose either ACEScct or ACEScc color science. This is primarily a personal preference since they are mostly identical, but the shadows respond differently to grading operations. In the Fusion page, images are automatically converted to linear, so whoever does the grading has more of a reason to choose one or the other.

  ACEScc: Choose ACEScc color science to apply a standard Cineon style log encoding to the ACES data before it is processed by DaVinci Resolve.

  ACEScct: This variation of ACEScc adds a roll-off at the toe of the image to make color correction lift operations "feel" more like they do with film scans and LogC encoded images.
- **ACES Version:** When you've chosen one of the ACES color science options, this menu becomes available to let you choose which version of ACES you want to use.
- **ACES Input Device Transform:** This menu lets you choose which IDT (Input Device Transform) to use for the dominant media format in use.
- **ACES Output Device Transform:** This menu lets you choose an ODT (Output Device Transform) with which to transform the image data to match your required deliverable.
- **Process Node LUTs In:** This menu lets you choose how you want to process LUTs in the Color page and does not affect the Fusion page.
- **Disable tone mapping for Fusion conversion:** Checking this box will remove any tone mapping from the ACES color management.

For more information on ACES within DaVinci Resolve, see Chapter 9, "Data Levels, Color Management, and ACES." in the *DaVinci Resolve 20 Reference Manual*.

Using OCIO for ACES Color Management in Fusion

When using Fusion Studio or not using color management in DaVinci Resolve, you have the option to use OpenColorIO nodes in Fusion to composite in an ACES color space.

OpenColorIO (OCIO) is an open-source color management framework for visual effects and computer animation. OCIO is compatible with the Academy Color Encoding Specification (ACES). Three OCIO nodes located in the Color category of the Effects Library allow you to use OCIO color space transforms in Fusion.

OCIO CDL Transform node allows you to create, save, load, and apply a Color Decision List (CDL) grade.

OCIO Color Space allows sophisticated color space conversions based on an OCIO config file.

OCIO File Transform allows you to load and apply a variety of LookUp Tables (LUTs).

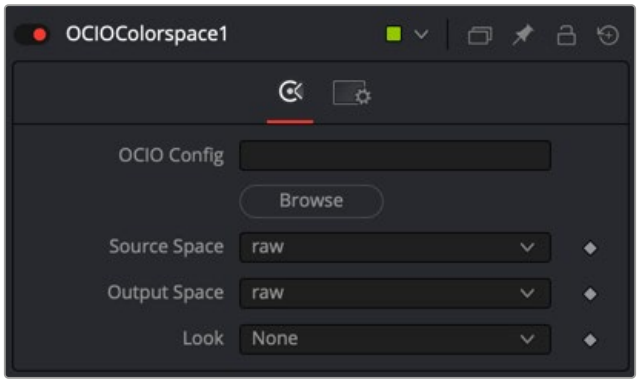
Using OCIO for converting MediaIn or Loader nodes to linear gamma is based on the OCIO Color Space node. Placing the OCIO Color Space node directly after a Loader (or MediaIn in DaVinci Resolve) displays the OCIO Source and Output controls in the Inspector.



OCIO Color Space nodes can be used to work in an ACES color managed environment.

Within the Inspector for OCIO Color Space node, Fusion includes default Source and Output transforms for standard color spaces. However, to use the full OCIO standard, you'll need to download and install the OCIO config file. You can download the config file from the OCIO website. https://opencolorio.org

Clicking the Browse button in the Inspector will allow you to navigate to the downloaded config file. From the download, locate the ACES 1.0.3 or later folder and select the file config.ocio.



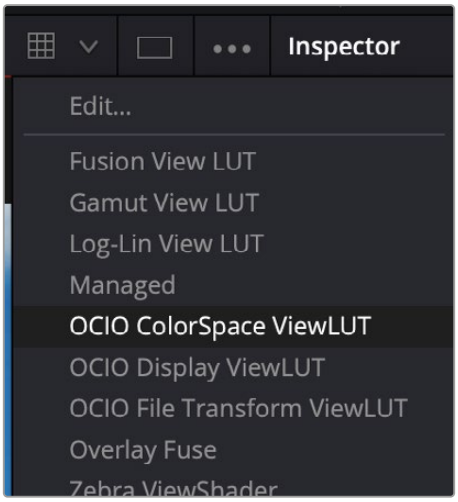
Source and Output menus are populated based on the config.ocio file that you download.

The Source menu is used to choose the color profile for your Loader or MediaIn node. The default raw setting shows an unaltered image, essentially applying no color management to the clip. The selection you make from the menu is based on the recording profile of your media.

The Output menu is set based on your deliverables. When working in Fusion Studio, typically the Output selected is ACEScg, to work in a scene linear space.

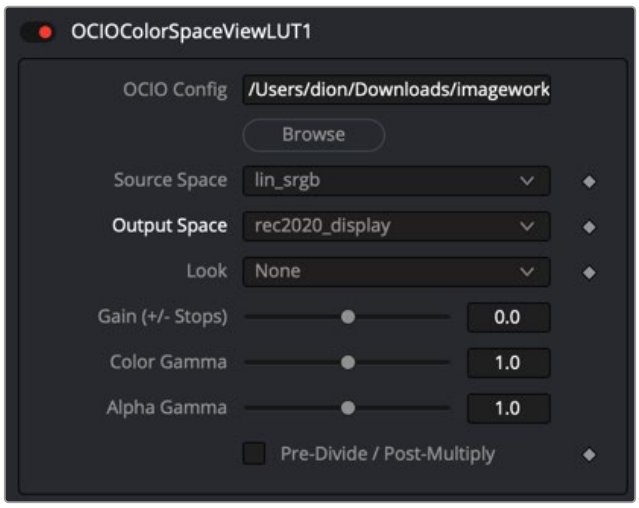
Applying OCIO LUTs in the Viewer

The viewer also includes OCIO View LUTs to calibrate the viewers. Once the OCIO Color Space View LUT is selected from the LUT menu above the viewer, choosing Edit for the same menu opens a dialog where you can load the OCIO config file.



The OCIO Color Space View LUT is located in the Viewer LUT menu in Fusion Studio.

By default, the same standard options are available in the View LUT. However, clicking the Browse button allows you to load the same config file you loaded into the OCIO Color Space node. Once loaded, all the expanded OCIO options are available. If you selected the OCIO Color Space node to output ACEScg, you use the OCIO View LUT to go from a source setting of linear sRGB to an output setting of sRGB or Rec. 709 in most cases.



The OCIO Color Space View LUT dialog is used to configure the viewer when using the OCIO Color Space node in the Node Editor.

TIP: If your monitor is calibrated differently, you will need to select a LUT that matches your calibration.

To set the LUT using OCIO:

1. Click the LUT menu and choose the OCIO Color Space View LUT.
2. From the same menu, select Edit.
3. In the View LUT editor that opens, set the source's color space to lin sRGB
4. Set the output space to sRGB or REC 709, assuming you are viewing on a standard computer monitor. You now see a normalized image in the viewer, but all color operations will be on linear images.

Whether you use the OCIO Color Space LUT or a LUT for your specific monitor calibration, you can save the viewer setup as the default.

To save the OCIO ColorSpace LUT setup as the default viewer setup:

- Right-click in the viewer, and then choose Settings > Save Defaults. Now, for every comp, the viewer is preconfigured based on the saved defaults.

Chapter 18

Understanding Image Channels

This chapter seeks to demystify how Fusion handles image channels and, in the process, show you how different nodes need to be connected to get the results you expect.

It also explains the mysteries of premultiplication, and presents a full explanation of how Fusion is capable of using and even generating auxiliary data.

Contents

Channels in Fusion 444
Types of Channels Supported by Fusion 444
Fusion Node Connections Carry Multiple Channels 445
Node Inputs and Outputs 445
Node Colors Tell You Which Nodes Go Together 449
Using Channels in a Composition 451
Channel Limiting 452
Adding Alpha Channels 453
How Channels Propagate During Compositing 454
Rearranging or Combining Channels 454
Understanding Premultiplication 455
The Rules of Premultiplication 457
Alpha Channel Status in MediaIn and Loader Nodes 459
Controlling Premultiplication in Color Correction Nodes 459
Controlling Premultiplication With Alpha Divide and Alpha Multiply 459
Multi Channel Compositing 460
Compositing with Beauty Passes 460
Working with Auxiliary Channels 466
Auxiliary Channels Explained 468
Propagating Auxiliary Channels 476
Nodes That Use Auxiliary Channels 476
Image Formats That Support Aux Channels 478
Creating Auxiliary Channels in Fusion 478

Channels in Fusion

Fusion introduces some innovative ways of working with the many different channels of image data that modern compositing workflows encompass. This chapter's introduction to color and data channels and how they're affected by different nodes and operations is a valuable way to begin the process of learning to do paint, compositing, and effects in Fusion.

If you're new to compositing, or you're new to the Fusion workflow, you ignore this chapter at your peril, as it provides a solid foundation to understanding how to predictably control image data as you work in this powerful environment.

Types of Channels Supported by Fusion

Digital images can be divided into separate streams of image data called Channels, each of which carries a specific kind of image data. Nodes that perform different image processing operations typically expect specific channels to provide predictable results. You are probably familiar with the three standard color channels of red, green, and blue, but there are many others. This section describes the different kinds of channels that Fusion supports.

RGB Color channels

The red, green, and blue channels of any still image or movie clip combine additively to represent everything we can see via visible light. Each of these three channels is a grayscale image when seen by itself. When combined additively, these channels represent a full-color image.

Alpha Channels

An alpha channel is an embedded fourth channel that defines different levels of transparency in an RGB image. Alpha channels are typically embedded in RGB images that are generated from computer graphics applications. In Fusion, white denotes solid areas, while black denotes transparent areas. Grayscale values range from more opaque (lighter) to more transparent (darker).

If you're working with an imported alpha channel from another application for which these conventions are reversed, never fear. Every node capable of using an alpha channel is also capable of inverting it.

Single-Channel Masks

While similar to alpha channels, mask channels are single channel images, external to any RGB image and typically created by Fusion within one of the available Mask nodes. Mask nodes are unique in that they propagate single-channel image data that defines which areas of an image should be solid and which should be transparent. However, masks can also define which parts of an image should be affected by a particular operation, and which should not. Mask channels are designed to be connected to specific mask inputs of nodes including Effect Mask, Garbage Mask, and Solid Mask inputs.

Auxiliary Channels

Auxiliary channels (covered in more detail later in this chapter), describe a family of special-purpose image data that typically expose 3D data in a way that can be used in 2D composites. For example, Z-Depth channels describe the depth of each pixel in an image along a Z axis, while an XYZ Normals channel describes the orientation (facing up, facing down, or facing to the left or right) of each pixel in an image. Auxiliary channel data is generated by rendering 3D images, so it usually accompanies or is embedded with RGB images generated by 3D modeling and animation applications.
These channels can also be generated from within Fusion via the Renderer 3D node, which outputs a 3D scene that you've assembled and lit as 2D RGBA channels, with optionally accompanying auxiliary channels.

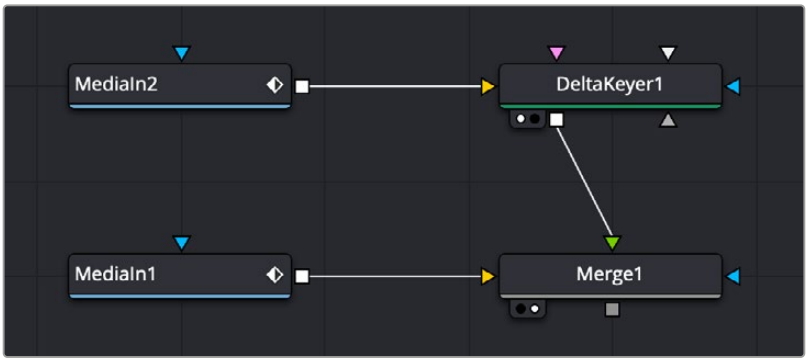
The reason to use auxiliary data is that 3D rendering is computationally expensive and time-consuming, so outputting descriptive information about a 3D image that's been rendered empowers compositing artists to make sophisticated alterations in 2D. You can add motion blur, perform relighting, and composite with depth information faster than re-rendering the 3D source material over and over.

TIP: You can view any of a node's channels in isolation using the Color drop-down menu in the viewer. Clicking the Color drop-down menu reveals a list of all channels within the currently selected node, including red, green, blue, or auxiliary channels.

Fusion Node Connections Carry Multiple Channels

The connections that pass image data from one node to the next in Fusion's Node Editor are capable of carrying multiple channels of image data. That means that a single connection may route RGB, or RGBA, or RGBAZ-Depth, or even just Z-Depth, depending on the Input you connect to and the node's function.

In the following example, the two MediaIn nodes output RGB data. However, the Delta Keyer creates an alpha channel and combined it with MediaIn2's RGB image. The RGB-A of the Delta Keyer becomes the foreground image that the Merge node can use to create a two-layer composite.



The alpha channel generated by the Delta key is used for compositing by the foreground input of the Merge node.

NOTE: Node trees shown in this chapter may display MediaIn nodes found in DaVinci Resolve's Fusion page; however, Fusion Studio Loader nodes are interchangeable unless otherwise noted.

Running multiple channels through single connection lines makes Fusion node trees simple to read, but it also means you need to keep track of which nodes process which channels to make sure that you're directing the intended image data to the correct operations.

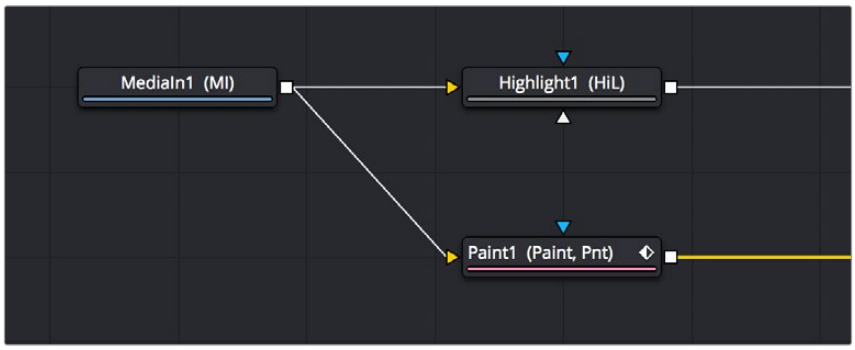
Node Inputs and Outputs

By default, MediaIn nodes in the Fusion page and Loader nodes in Fusion Studio output RGBA channels. When you connect one node's output to another node's input, the active channels are passed from the upstream node to the downstream node, which then processes the image according to that node's function. Only one node output can be connected to a node input at a time. In this simple example, a MediaIn node's output is connected to the input of a Highlight node to create a sparkly highlight effect.



MediaIn node connected to a Highlight node connected to MediaOut node in the Fusion page.

When connecting nodes, a node's output carries the same channels no matter how many times the output is "branched." You cannot send one channel out on one branch and a different channel out on another branch.

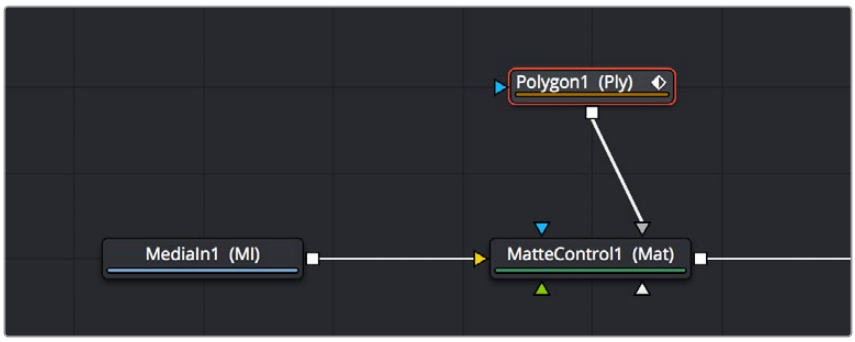


The MediaIn node's output is branched but carries the same RGB channels to both inputs.

Using Multiple Inputs

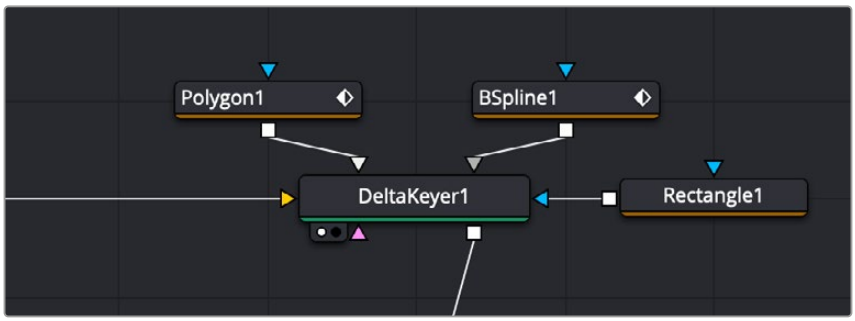
Most nodes have two inputs, one for RGBA and another for an effect mask that can be optionally used to limit the effect of that node to a particular part of the image. However, some nodes have three or even more inputs, and it's important to make sure you connect the correct image data to the appropriate input to obtain the desired result. If you connect one node's output to another node's input and nothing happens, chances are you've connected to the wrong input.

For example, the MatteControl node has a background input and a foreground input, both of which accept RGBA channels. However, it also has SolidMatte, GarbageMatte, and EffectsMask inputs that accept alpha or mask channels to modify the transparency of the Node's output. If you want to perform the extremely common operation of using a MatteControl node to create an alpha channel using a Polygon node for rotoscoping an image, you need to make sure that you connect the Polygon node to the GarbageMatte input to obtain the correct result. The GarbageMatte input is automatically set to alter the alpha channel of the foreground image. If you connect to any other input, your Polygon mask may not produce expected results.



Polygon node connected to the GarbageMatte input of a MatteControl node for rotoscoping

In another example, the DeltaKeyer node has a primary input (labeled “Input”) that accepts RGBA channels, but it also has three Matte inputs. These SolidMatte, GarbageMatte, and EffectsMask inputs on the Delta Keyer accept alpha or mask channels to modify the matte being extracted from the image in different ways.



The DeltaKeyer combines the multiple mask nodes in different ways.

If you position your pointer over any node’s input or output, a tooltip appears in the Tooltip bar at the bottom of the Fusion window, letting you know what that input or output is for, to help guide you to using the right input for the job. If you pause for a moment longer, another tooltip appears in the Node Editor itself.



(Left) The node input’s tooltip in the Tooltip bar,
(Right) The node tooltip in the Node Editor

Connecting to the Correct Input

When you’re connecting nodes, pulling a connection line from the output of one node and dropping it right on top of the body of another node makes a connection to the default input for that node, which is typically the “Input” or “Background” input.



Side by side, dropping a connection on a node’s body to connect to that node’s primary input

However, if you drop a connection line right on top of a specific input, then you'll connect to that input, so it's important to be mindful of where you drop connection lines as you wire up different node trees together.



Side by side, dropping a connection on a specific node input, note how the inputs rearrange themselves afterwards to keep the node tree tidy-looking

TIP: If you hold the Option key down while you drag a connection line from one node onto another, and you keep the Option key held down while you release the pointer's button to drop the connection, a menu appears that lets you choose which specific input you want to connect to, by name.

Inputs Require Specific Channels

Usually, you're prevented from connecting a node's output to another node or node input that's not compatible with it. For example, if you try to connect a Text3D node's output directly to the input of a regular Merge node, it won't work; 3D nodes do not produce RGB images, they produce 3D geometry data, so you must first connect to a Renderer3D node that creates the RGB output appropriate for 2D compositing operations.

In other cases, connecting the wrong image data to the wrong node input won't give you any error, it simply fails to produce the result you were expecting, necessitating you to troubleshoot the composition. If this happens to you, check the Fusion Effects section of this manual to see if the node you're trying to connect to has any limitations as to how it must be attached.

TIP: This chapter tries to cover many of the easy-to-miss exceptions to node connection that are important for you to know, so don't skim too fast.

Always Connect the Background Input First

Many nodes combine images in different ways, using "background" and "foreground" inputs. This includes the Merge node, the Matte Control node, and the Channel Booleans node, to cite common examples. The color of inputs on a node can help you make the right corrections. For instance, background inputs are always orange, and foreground inputs are always green.

When you first connect any node's output to a multi-input node, you usually want to connect the background input first. This is handled for you automatically when you first drop a connection line onto the body of a new multi-input node. The orange-colored background input is almost always connected first (the exception is Mask nodes, which always connect to the first available Mask input). This is good because you want to get into the habit of always connecting the background input first.

If you connect to only one input of a multi-input node and you don't connect to the background input, you may find that you don't get the results you wanted. This is because each multi-input node expects that the background is connected before anything else so that the internal connections and math used by that node can be predictable.

TIP: The only node to which you can safely connect the foreground input prior to the background input is the Dissolve node, which is a special node that can be used to either dissolve between two inputs, or automatically switch between two inputs of unequal duration.

Node Colors Tell You Which Nodes Go Together

Each node in Fusion accomplishes a single type of effect or operation. These single-purpose nodes make it easier to decipher a complex composition when examining its node tree. Single-purpose nodes also make it easier to focus on fine-tuning specific adjustments, one at a time, when assembling an ever-growing tree.

Because each Fusion node has a specific function, they're categorized by type to make it easier to keep track of which nodes require what types of image channels as input, and what image data you can expect each node to output. These general types are described here.



A node tree showing the main categories of node colors

Blue MediaIn and Loader Nodes, and Green Generator Nodes

Blue MediaIn nodes and blue Loader nodes add clips to a composite, and green Generator nodes create images. Both types of nodes output RGBA channels (depending on the source and generator), and may optionally output auxiliary channels for doing advanced compositing operations.

Because these are sources of images, both kinds of nodes can be attached to a wide variety of other nodes for effects creation besides just 2D nodes. For example, you can also connect MediaIn nodes to Image Plane 3D nodes for 3D compositing, or to pEmitter nodes set to “Bitmap” for creating different particle systems. Green Generator nodes can be similarly attached to many different kinds of nodes, for example, attaching a FastNoise node to a Displace 3D node to impose undulating effects to 3D shapes.

Shape nodes are also green, although they must be attached to a specialized set of gray modifier and render nodes (all of which begin with the letter “s” and appear in the Shape category of the Effects Library).

2D Processing Nodes, Color Coded by Type

These encompass most 2D processing and compositing operations in Fusion, all of which process RGBA channels and pass along auxiliary channels. These include:

- Orange Blur nodes
- Olive Color Adjustment nodes (color adjustment nodes additionally concatenate with one another)
- Pink Paint nodes
- Dark orange Tracking nodes
- Tan Transform node (transform nodes additionally concatenate with one another)
- Teal VR nodes
- Dark brown Warp nodes
- Gray, which includes Compositing nodes as well as many other types.

Additionally, some 2D nodes such as Fog and Depth Blur (in the Deep Pixel category) accept and use auxiliary channels such as Z-Depth to create different perspective effects in 2D.

> **TIP:** Two 2D nodes that specifically don’t process alpha channel data are the Color Corrector node and the Gamut node. The Color Correction node lets you color correct a foreground layer to match a background layer without affecting an alpha channel. The Gamut node lets you perform color space conversions to RGB data from one gamut to another without affecting the alpha channel.

Purple Particle System Nodes

These are nodes that connect to create different particle systems, and they’re incompatible with other kinds of nodes until you add a pRender node that outputs 2D RGBA and auxiliary data that can be composited with other 2D nodes and operations.

Dark Blue 3D Nodes

These are 3D operations, which generate and manipulate 3D data (including auxiliary channels) that are incompatible with other kinds of nodes until processed via a Renderer 3D node, which then outputs RGBA and auxiliary data.

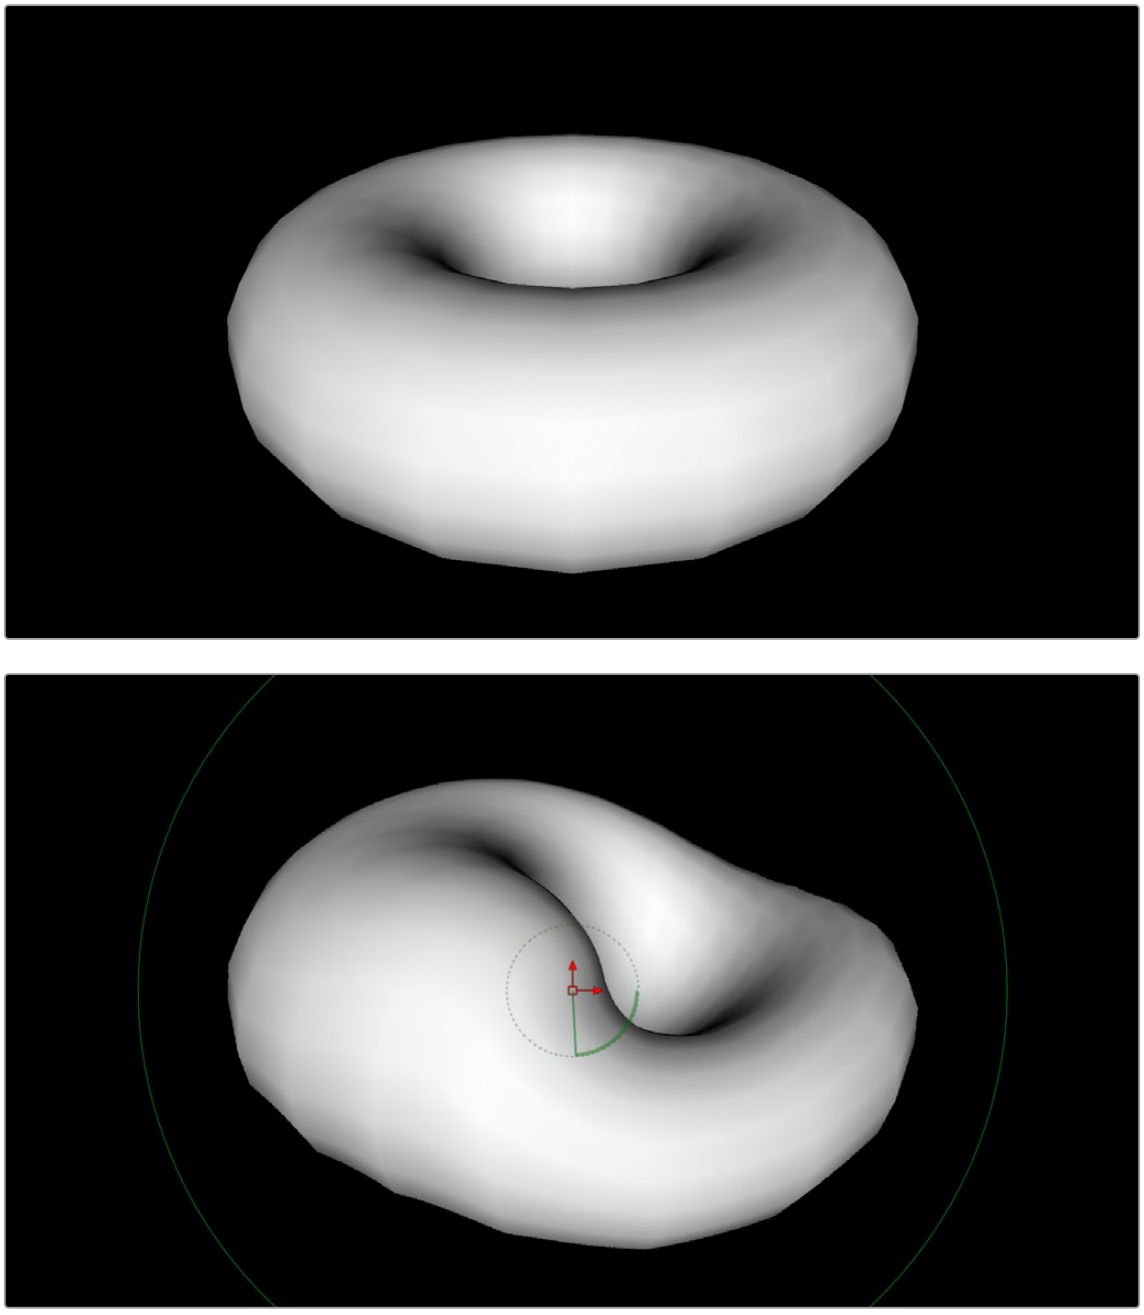
Brown Mask Nodes

Masks output single-channel images that can only be connected to one another (to combine masks) or to specified Mask inputs. Masks are useful for defining transparency (Alpha masks), defining which parts of an image should be cropped out (Garbage masks), or defining which parts of an image should be affected by a particular node operation (Effects masks).

Using Channels in a Composition

When you connect one node's Output to another node's Input, you feed all of the channels that are output from the upstream node to the downstream node. 2D nodes, which constitute most simple image processing operations in Fusion, propagate all channel data from node to node, including RGB, alpha, and auxiliary channels, regardless of whether or not that node actually uses or affects a particular channel.

2D nodes also typically operate upon all channel data routed through that node. For example, if you connect a node's output with RGBA and XYZ Normals channels to the input of a Vortex node, all channels are equally transformed by the Size, Center, and Angle parameters of this operation, including the alpha and XYZ normals channels, as seen in the following screenshot.



(Top) The Normal Z channel output by a rendered torus, (Bottom) The Normal Z channel after the output is connected to a Vortex node; note how this auxiliary channel warps along with the RGB and A channels

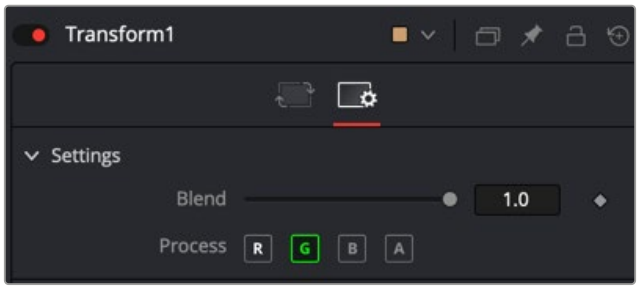
This is appropriate because in most cases, you want to make sure that all channels are transformed, warped, or adjusted together. You wouldn't want to shrink the image without also shrinking the alpha channel along with it, and the same is true for most other operations.

On the other hand, some nodes deliberately ignore specific channels when it makes sense. For example, the Color Corrector and Gamut nodes, both of which are designed to alter RGB data specifically, do not affect auxiliary channels. This makes them convenient for color-matching foreground and background layers you're compositing, without worrying that you're altering the depth information accompanying that layer.

TIP: If you're doing something exotic and you actually want to operate on a channel that's usually unaffected by a particular node, you can always use the Channel Booleans node to reassign the channel. When doing this to a single image, it's important to connect that image to the background input of the Channel Booleans node, so the alpha and auxiliary channels are properly handled.

Channel Limiting

Most nodes have a set of Red, Green, Blue, and Alpha buttons in the Settings tab of that node's Inspector. These buttons let you exclude any combination of these channels from being affected by that node.



The channel limiting buttons in the Settings panel of a Transform node, so only the Green channel is affected

For example, if you wanted to use the Transform node to affect only the green channel of an image, you can turn off the Red, Blue, and Alpha buttons. As a result, the green channel is processed by this operation, and the red, blue, and alpha channels are copied straight from the node's input to the node's output, skipping that node's processing to remain unaffected.



Transforming only the green color channel of the image with a Transform effect

Skipping Channel Processing

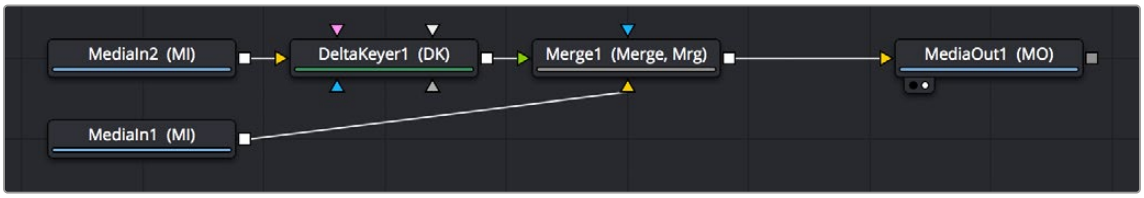
Under the hood, most nodes process all channels first, but afterward copy the input image to the output for channels that have been enabled. Modern workstations are so fast that this isn't usually noticeable, but there are some nodes where deselecting a channel actually causes that node to skip processing that channel entirely. Nodes that operate this way have a linked set of Red, Green, Blue, and Alpha buttons on another tab in the node. In these cases, the Common Control channel buttons are instanced to the channel buttons found elsewhere in the node.

Blur, Brightness/Contrast, Erode/Dilate, and Filter are examples of nodes that all have RGBA buttons in the main Controls tab of the Inspector, in addition to the Settings tab.

Adding Alpha Channels

Much of visual effects compositing has to do with placing a foreground subject over a background. Possibly the most fundamental method is through the use of an alpha or matte channel. If an alpha channel is not contained within the clip, you add one via keying or rotoscoping. While more specific methods are covered in detail in later chapters, here's an example of how this is handled within Fusion.

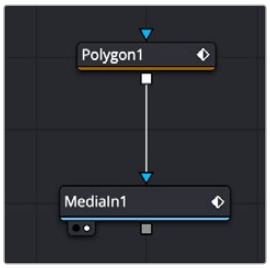
In the case of extracting an alpha matte from a green screen image, you typically connect the image's RGB output to the "Input" input of a Keyer node such as the Delta Keyer, and you then use the keyer's controls to extract the matte. The Keyer node automatically inserts the alpha channel that's generated alongside the RGB channels, so the output is automatically RGBA. Then, when you connect the keyer's output to a Merge node to composite it over another image, the Merge node automatically knows to use the embedded alpha channel coming into the foreground input to create the desired composite, as seen in the following screenshot.



A simple node tree for keying; note that only one connection links the DeltaKeyer to the Merge node

Rotoscoping, or manually drawing a mask shape using a Polygon or other Mask node is another technique used to create the matte channel. There are many ways to configure the node tree for this task, but the simplest setup is just to connect a Polygon or B-Spline mask node to the Effect Mask input of a MediaIn or Loader node.

TIP: When rotoscoping, it is best to leave the Mask node disconnected from the image while you draw the shape. This allows you to view the MediaIn node while drawing. Connect the Matte node once you have finished drawing the shape.

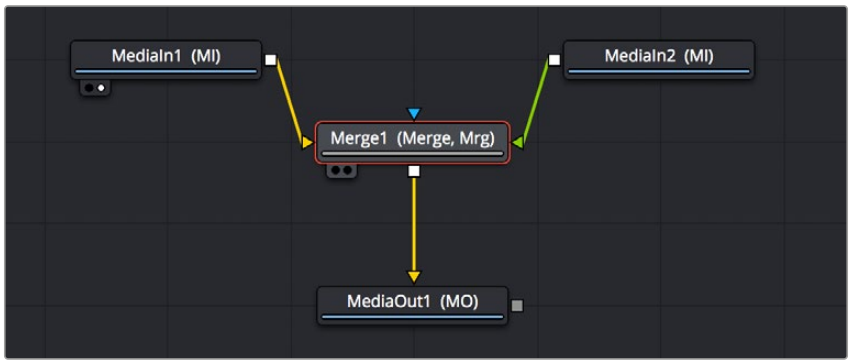


A simple rotoscoping setup using a Polygon node directly connected to the Effect Mask input of a MediaIn.

In both cases, you can see how the node tree's ability to carry a single channel or multiple channels of image data over a single connection line simplifies the compositing process.

How Channels Propagate During Compositing

Images are combined or composited together using the Merge node. The Merge node takes two RGBA inputs labeled "Foreground" (green) and "Background" (orange) and combines them into a single RGB output (or RGBA if both the foreground and background input images have alpha), where the foreground image is in front (or on top, depending on what you're working on), and the background image is, you guessed it, in back.



A simple Merge node composite

Auxiliary channels, on the other hand, are handled in a much more specific way. When you composite two image layers using the Merge node, auxiliary channels only propagate through the image that's connected to the background input. The rationale for this is that in most CGI composites, the background is most often the CG layer that contains auxiliary channels, and the foreground is a live-action green screen plate.

Since most compositions use multiple Merge nodes, it pays to be careful about how you connect the background and foreground inputs of each Merge node to make sure that the correct channels flow properly.

TIP: Merge nodes are also capable of combining the foreground and background inputs using Z-Depth channels using the "Perform Depth Merge" checkbox, in which case every pair of pixels are compared. Which one is in front depends on its Z-Depth and not the connected input.

Rearranging or Combining Channels

Last, but certainly not least, it's also possible to rearrange and re-combine channels in any way you need, using one of four different node operations. For example, you might want to combine the red channel from one image with the blue and green channels of a second image to create a completely different channel mix. Alternatively, you might want to take the alpha channel from one image and merge it with the alpha channel of a second image in different ways, adding, subtracting, or using other intersection operations to create a very specific blend of the two.

The following nodes are used to re-combine channels in different ways:

- **Channel Boolean:** This is a 3D node used to remap and modify channels of 3D materials using a variety of simple pre-defined math operations.
- **Channel Booleans:** Used to shuffle or rearrange RGBA and auxiliary channels within a single input image, or among two input images, to create a single output image. If you only connect a single image to this node, it must be connected to the background input to make sure everything works.
- **Copy Aux:** The Copy Aux node is used to remap channels between RGBA channels and auxiliary data channels in a single 2D image. The Copy Aux node is mostly a convenience node, as the copying can also be accomplished with more effort (and flexibility) using a Channel Booleans node.
- **Matte Control:** Designed to do any combination of the following: (a) re-combining mattes, masks, and alpha channels in various ways, (b) modifying alpha channels using dedicated matte controls, and (c) copying alpha channels into the RGB stream of the image connected to the background input in preparation for compositing. You can copy specific channels from the foreground input to the background input to use as an alpha channel, or you can attach masks to the garbage matte input to use as alpha channels as well.

Understanding Premultiplication

Now that you understand how to direct and recombine RGB images and alpha channels in Fusion, it's time to go more deeply into alpha channels to make sure you always combine RGB and alpha channels correctly for each operation you perform in your composite. This might seem simple, but small mistakes are easy to make and can result in unsightly artifacts. This is arguably one of the most confusing areas of visual effects compositing, so don't skip this section.

When alpha channel and RGB pixels are both contained within a media file, such as a 3D rendered animation that contains RGB and transparency, or a motion graphics movie file with transparency baked in, there are two different ways they might be combined, and it's important to know which is in use.

- **Unpremultipled (Straight):** An RGB image unaltered by the semi-transparency information in a fourth channel (alpha channel)
- **Premultiplied:** An RGB image that has each channel multiplied by its alpha channel before compositing.

The term Premultiplied alpha is a term that has historically been used by editors, visual effects artists, and motion graphics designers, but it's imprecise. The alpha channel itself is not multiplied. The R, G, and B channels are multiplied by the alpha. In the end, the alpha channel stays the same, but the values contained in the R, G, and B channels are modified.



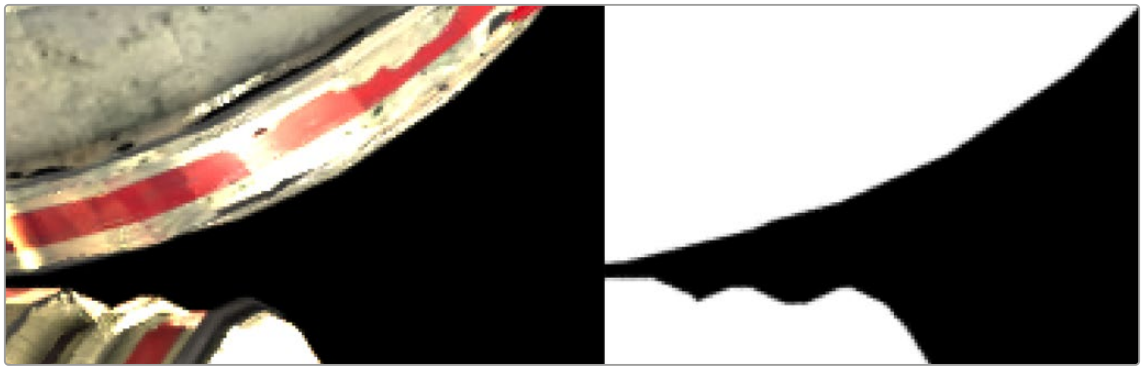
A RGB image (left) and its alpha channel (right)

Non-premultiplied images, sometimes called “straight” alpha channels, have RGB channels that are unaltered (not multiplied) by the alpha channel. The result is that the RGB image has no anti-aliased edges and no semi-transparency. It’s usually obvious where the RGB image ends and the alpha matte begins. The image below is an example of the ragged edges seen in the RGB channels when using a non-premultiplied alpha channel. But notice the smooth semi-transparent edges found in the alpha.



A detailed view of a non-premultiplied RGB image (left) and its alpha channel (right)

A premultiplied alpha channel means the RGB pixels are multiplied by the alpha channel. This method guarantees that the RGB image pixels include semi-transparency where needed, like anti-aliased edges. Most computer-generated images are premultiplied for convenience, because they’re easier to review smoothly without actually being placed inside of a composite.



A detailed view of a premultiplied image (left) and its alpha channel (right)

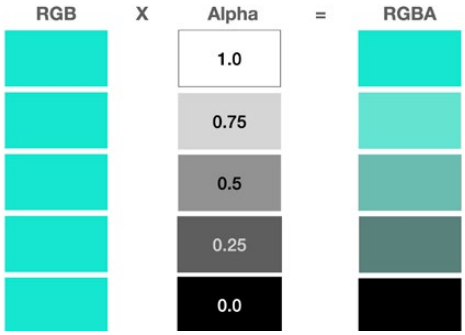
What does this mean for compositing? The edges of a premultiplied image look much smoother, making them the preferred choice when compositing foreground over the background in a merge. Consequently, all alpha channels are made to be premultiplied in compositing operations if they weren’t already.

On the other hand, it is always preferred to color correct a non-premultiplied RGBA image, because you don’t want to alter the pixel values of an image after the RGB channels have been multiplied by the alpha channel.

If you think of this from a mathematical perspective:

- **RGB pixel value x 0 = 0:** The black transparent areas of an alpha channel have a pixel value of 0. When you take the value of an RGB pixel and multiply it by 0 (n x 0 = 0) then by the laws of multiplication, the RGB value becomes 0, or fully transparent.
- **RGB Pixel value x 1 = RGB Pixel:** The solid or opaque white areas have a value of 1.0. When you take the value of an RGB pixel and multiply it by 1 (n x 1 = n), then the RGB value stays the same, fully opaque.

- **RGB Pixel value x 0.3 = A different color:** Along the edges of an alpha channel are gray pixels, indicating semi-transparency. These semi-transparent pixels have a value falling somewhere between 1.0 and 0.0. To apply the alpha channel's anti-aliased edges to the RGB channels, you multiply the pixel values. The multiplication process mixes some percentage of the transparent pixels (black) with the RGB pixels. Although this is desired to get good anti-aliased edges, you can not color correct the image because it alters the smooth semi-transparency you created once it is done.



RGB pixels are multiplied by varying degrees of transparency and the result is a different RGB value.

The Rules of Premultiplication

Following from the information presented above, when you're compositing multiple images together and one or more has a built-in alpha channel, you want to make sure you follow these general rules to avoid problems:

- Always use premultiplied images with a Merge node.
- Only color-correct images that are not premultiplied.
- Always filter and transform images that are premultiplied.
- Never double premultiply an image.

Premultiplication and the Merge Node

The foreground input of a Merge node expects a premultiplied RGBA image. It is an additive merge, meaning the semi-transparent areas of the foreground are added over the background. However, if the image is not premultiplied, the pixels that should be transparent are still added, which typically results in an unwanted bright fringe around the edges of your foreground subject.

If you are compositing with a non-premultiplied alpha, you can fix these bright edges by changing the Merge to perform a Subtractive merge in the Inspector.



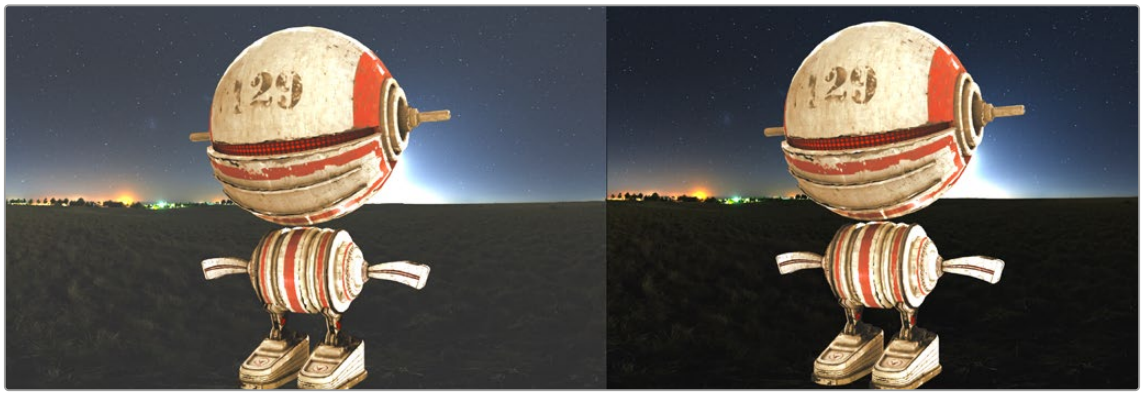
Nonpremultiplied edges in an additive Merge (left) and premultiplied edges in an additive Merge (right)

TIP: When an RGB image and a Mask node are combined using, for instance, a Matte Control node, if the RGB image is not multiplied by the mask in the Matte control, the checkerboard background pattern in the viewer will appear only semi-transparent when it should be fully transparent.

Color Correcting Premultiplied RGBA Images

When you premultiply an image, you tie any changes made to the brightness of the RGB pixels to the alpha channel pixels as well. If, for example, you raise the brightness of the foreground image in any way, you'll also raise the brightness of the alpha channel, which will likely not be desirable as this will change the transparency the alpha channel creates. The visible result of this when viewing the Merge node is that the entire background will become brighter (or darker if you'd lowered the RGB brightness) based on your color adjustment.

For this reason, the rule is always to divide the semi-transparent pixels before performing any color correction on an image with an alpha channel. You can do this turning on the Pre-Divide/Post Multiply checkbox in any node that performs color correction. Alternatively, you can use the Alpha Divide and Alpha Multiply nodes to do the same thing. These methods are covered in more detail later in this chapter.



Color correcting a premultiplied foreground incorrectly alters the background (left).
Color correcting a nonpremultiplied foreground works correctly (right).

Double Premultiplied RGBA Means Double Trouble

A common mistake made by many artists is to over-compensate for premultiplication. As important as it is to premultiply the alpha before compositing in a Merge node, it's just as important not to double premultiply the alpha. Performing a premultiply operation two times in a row can create a darken halo effect around your images. You are effectively multiplying by the gray semi-transparent pixels twice; this is not optimal.



Double premultiplied image displays dark edges (left);
premultiplied image with correct edges (right)

Premultiplied Alpha Channels and Filtering

When dealing with filtering, the state of the RGBA channels shouldn't matter for most composites. However, an exception might be if the filter algorithm you choose includes color modification. For instance, if a filter attempts to simulate a defocus by blooming highlights like a real light source, that filter might over-brighten pixels near a transparent edge, which will result in some manner of artifact when that image is composited.

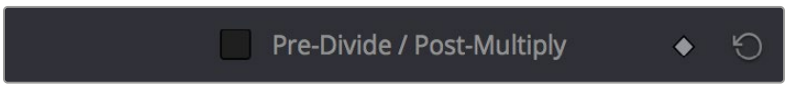
Alpha Channel Status in MediaIn and Loader Nodes

When using a Loader node to add a clip to the composite, the Import tab in the Inspector includes a group of checkboxes that let you determine how an embedded alpha channel is handled. There are checkboxes to make the alpha channel solid (ignore transparency), to Invert the alpha channel, and to Post-Multiply the RGB channels with the alpha channel, should that be necessary.

When using a MediaIn node, you can modify how an embedded alpha channel is interpreted using the Clip attributes window. The Clips Attributes window includes an Alpha Mode menu setting to choose if the alpha channel is ignored, treated as premultiplied, inverted, or treated as nonpremultiplied (straight).

Controlling Premultiplication in Color Correction Nodes

Most nodes that require you to explicitly deal with the state of premultiplication of an RGBA image have a "Pre-Divide, Post-Multiply" checkbox. This includes simple color correction nodes such as Brightness Contrast and Color Curves, as well as the Color Correct node, which has the "Pre-Divide/Post-Multiply" checkbox in the Options panel of its Inspector settings.

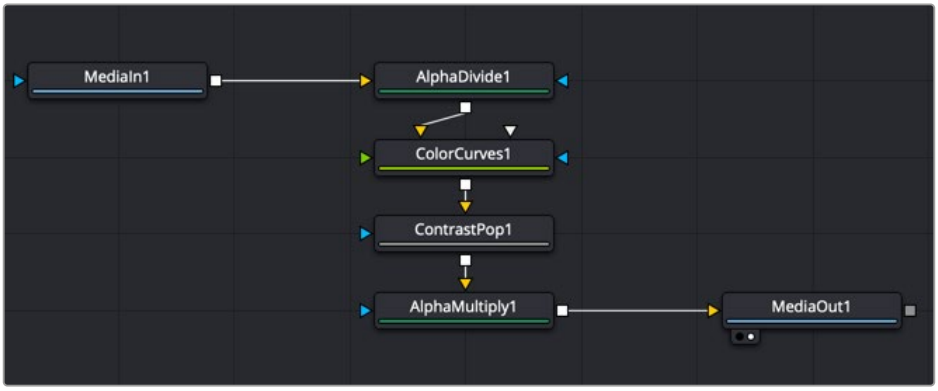


The Pre-Divide/Post-Multiply checkbox of t
he Color Curves node, seen in the Inspector

This checkbox allows you to connect an RGBA premultiplied image to the node and perform a color correction operation. It takes the RGBA image input, performs a divide operation to remove the semi-transparency and then performs a multiplication operation before outputting the color corrected image. This way, the color correction is done using a nonpremultiplied image but the resulting output is a Merge-friendly premultiplied image.

Controlling Premultiplication With Alpha Divide and Alpha Multiply

The Alpha Divide and Alpha Multiply nodes, found in the Matte category of the Effects Library, are available when multiple operations in a row expect a "straight" alpha channel. Instead of performing repetitive Pre-Divide/Post Multiply operations on each node, you can use these two nodes to book-end the other nodes. Simply add the Alpha Divide node when you want the RGBA image data not to be premultiplied, and add the Alpha Multiply node when you want the image data to be premultiplied again. For example, if you're using third-party OFX nodes that make color adjustments, you may need to control premultiplication before and after such an adjustment manually.



A node tree with explicit Alpha Divide and Alpha Multiply nodes

Multi Channel Compositing

If you've gone to the cinema and seen any recent superhero movie, you will have seen the results of sophisticated 3D rendering, combined with a whole lot of compositing. 3D applications can render very realistic images but the time it takes to render each frame of those realistic images can be measured in hours, not minutes. Any change to the 3D images, even relatively simple operations like color adjustments, changes to focus, filtering, or additional masking, means these images will need to be re-rendered entirely, which means you waiting, often many times over. For efficiencies sake, it's much faster to make iterative changes that can be accomplished with 2D image processing operations in Fusion instead.

To have the flexibility you need to make common changes to 3D images after-the-fact, the various attributes that make up the 3D scene are separated and rendered as different image sequences, often referred to as render passes. For example, render passes are often created for attributes like raw color, shadows, and reflections, that can be recombined as a 2D composite to produce the final result. Having different attributes rendered into different image sequences gives you a significant amount of flexibility, since now each image attribute can be color corrected, blurred, or further processed independently of the other attributes of the image, with fast-processing operations in Fusion.

The most common render passes that are typically generated come from the RGBA channels of the 3D scene. These are collectively called beauty passes and can consist of attributes like color, shadows, lighting, reflections, environment, and others.

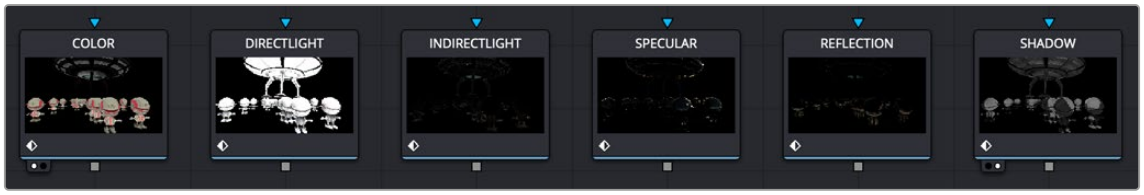
Render passes can also contain non-RGB data. Different effects applications have different names for these passes, such as Data Channels, or AOVs (Arbitrary Output Variables). In Fusion, these channels are called Auxiliary Channels, and they contain 3D data such as Depth, Normals, Motion Vectors, and UV Coordinates (to name just a few).

When compositing a 3D render consisting of multiple render passes, the beauty passes are handled using one technique, and the Auxiliary Channels are handled with another. Since Fusion nodes carry RGBA channels by default, we'll cover beauty passes first, and then explain how to work with Auxiliary Channels later in this chapter.

Compositing with Beauty Passes

Each attribute of a beauty pass can be rendered into individual image sequences, so you end up with a series of numbered images, one for the diffuse pass, one for the reflection pass, another of the shadow pass, and so on. Alternatively, all the passes can be contained in a multi-part EXR image sequence. Multi-part EXRs benefit from requiring a bit less file management, but the passes are handled similarly in Fusion using either method.

A single MediaIn or Loader node only handles a single beauty pass since only one set of RGBA channels gets output per node. Setting up your composite in Fusion requires you to use a separate MediaIn or Loader node for each pass.

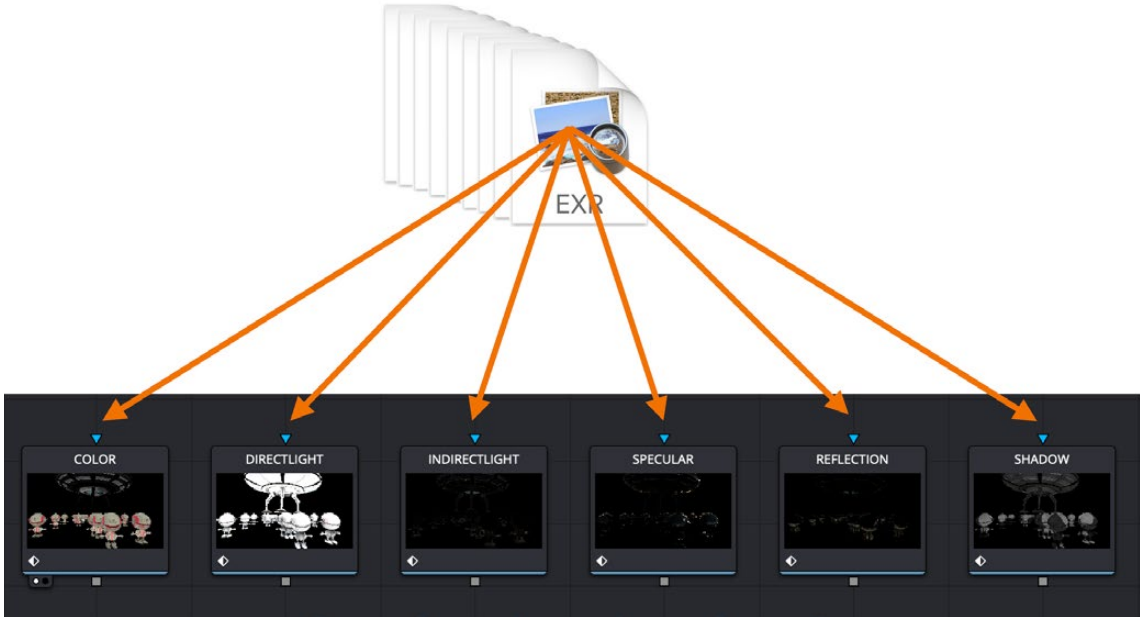


A MediaIn or Loader node is needed for each beauty pass.

TIP: The EXR format in Fusion is optimized when multiple Loaders are used to read the same EXR file. The file is only loaded once to access all the channels.

Beauty Pass Setup

Separated components of 3D renders, such as diffusion, shadows, or reflections, can be rendered individually using RGB channels. If you're provided with individual image sequences for each image component, they'll import and open in Fusion and can be displayed in the viewer, just like any other clip. If you're using multi-part EXR files, then you essentially have multiple RGB images contained within a single file. However, the MediaIn and Loader nodes can only make use of one set of RGB channels at a time. If you want to composite multiple RGB beauty passes together with one another, you must use multiple MediaIn or Loader nodes that point to the same image sequence but are assigned to different passes contained in the EXR file.

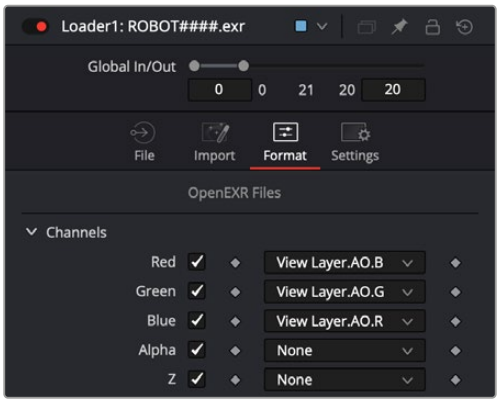


Multiple Loader or MediaIn nodes connect to a multi-part EXR image sequence

TIP: It is wise to rename each Loader or MediaIn to represent the beauty pass it contains.

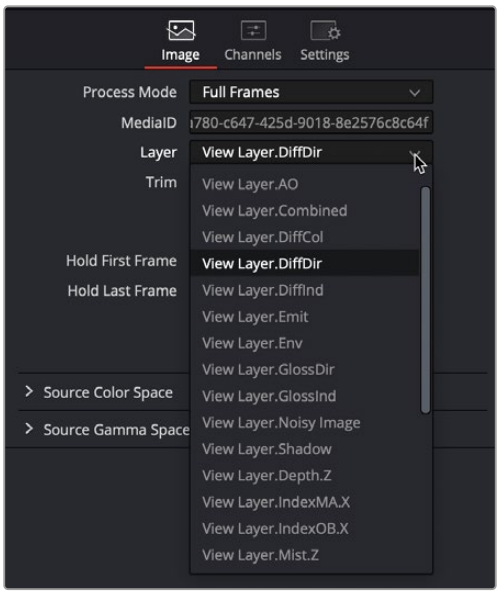
Mapping Each Set of Beauty Passes to a Particular Node

Depending on whether you're using a MediaIn node or a Loader node, beauty passes can be mapped to RGB channels using either the Image tab, the Channels tab, or the Format tab in the Inspector. When using a Loader node, the Loader's Format tab is used.



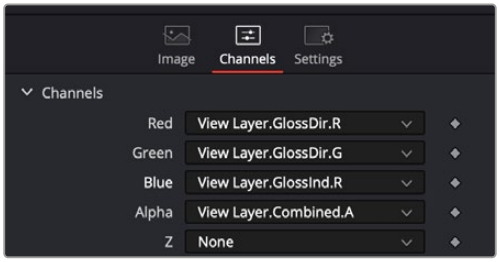
Beauty passes mapped to red, green, and blue channels in a Loader node

The MediaIn's Image tab includes a Layer menu. Any pass included in a multi-part EXR image sequence can be selected from this menu and automatically assigned to the RGBA channels.



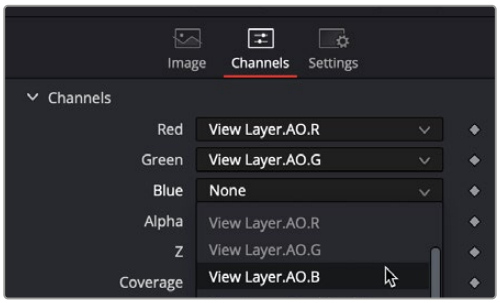
The Layer menu in a MediaIn node showing headings for combined channel passes

In most cases, the menu shows the combined channel passes, meaning the individual red, green, blue, and alpha channels cannot be selected. Because the alpha channel is not included in many beauty passes, you sometimes need to borrow the alpha channel from a different beauty pass. For this reason, it's often better to use the Channels tab for mapping the individual channels of a beauty pass to the channels of the MediaIn node.



The MediaIn node's Channels tab or the Loader's Format tab provides access to individual channels.

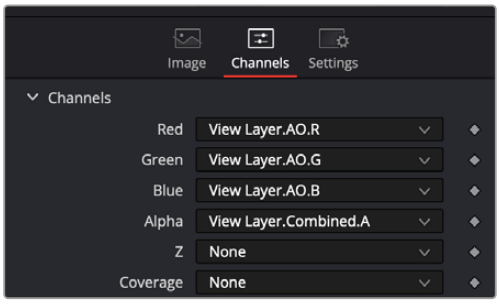
The MediaIn node in DaVinci Resolve and the Format tab in a Loader node include the same channel mapping functionality. The Channels tab and the Format tab include individual RGBA menus at the top of the tab. You can use these menus to map the RGBA channels from any pass contained in the multi-part EXR. For instance, if you want to map the Ambient Occlusion pass to the RGB channels, choose AO. R (Red) from the Red channel menu, AO. G (Green) from the Green channel menu, and AO. B (Blue) from the Blue channel menu.



A pass' individual channels mapped to the red, green, blue, and alpha channels in the MediaIn node

TIP: Different 3D applications will label beauty passes in different ways. For instance, the name for an Ambient Occlusion beauty pass may be AO, AM_OCC, or some other abbreviation.

The Ambient Occlusion beauty pass does not include an alpha channel. To composite it, you can reuse the alpha channel pass from another beauty pass. In the image below, the alpha channel is mapped using the combined render pass' alpha channel.

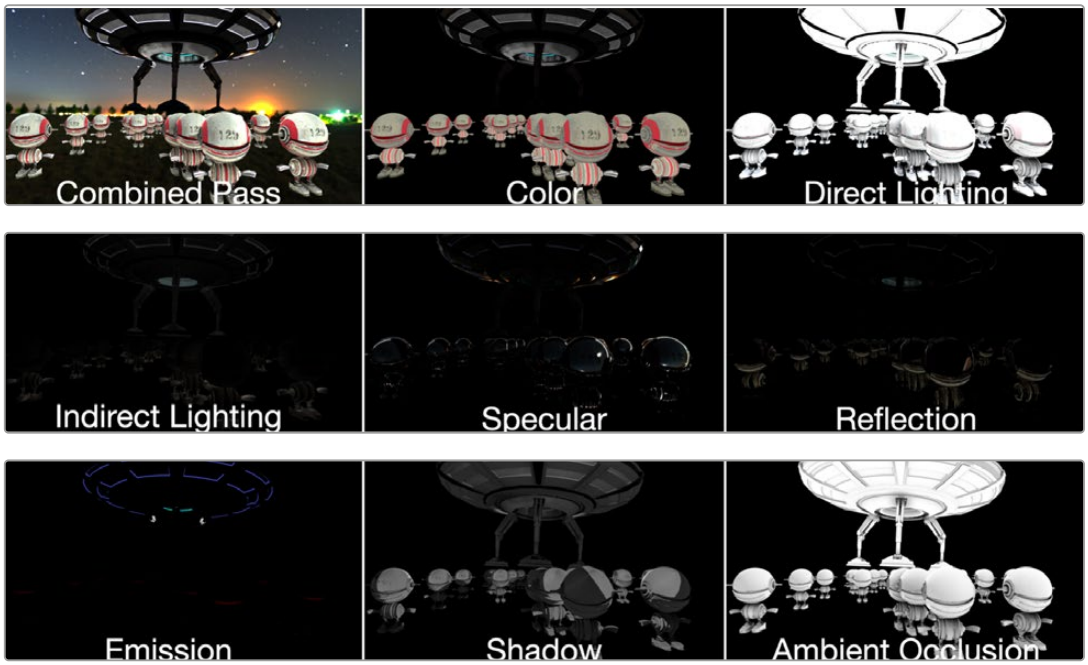


The alpha channel from a different beauty pass combined with the Ambient Occlusion pass

TIP: When using the Format tab in the Loader node, the checkbox next to each channel needs to be turned on for the corresponding channel to become available in the node's output.

Compositing Multiple Beauty Passes in the Node Editor

Once all of your passes are brought in and mapped to RGBA channels, you'll end up with a series of MediaIn or Loader nodes. How many MediaIn and Loader nodes is really up to your workflow. There is no predefined number of passes you will use. Every studio decides for themselves the standard set. However, there are common render passes that are involved in most composites. The following is a list of commonly used render passes and their generic names.



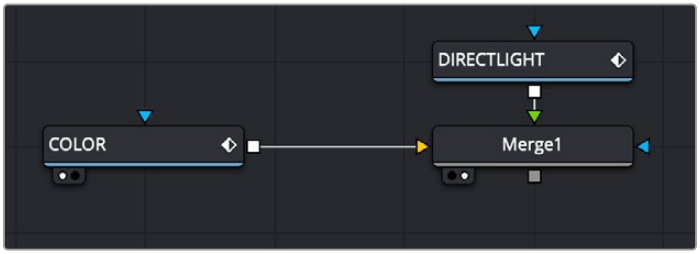
Commonly used beauty passes, compared

Compositing multiple beauty passes into a single output image is relatively straight forward. 3D rendering applications typically output linear gamma, so no Gamut or other color space conversion nodes are required if you're keeping the image in a linear color space for ease of compositing.

The basic compositing is accomplished with either a Merge node or a Channels Booleans node. Both allow for additive combining of render passes. There's no strict requirement for compositing each pass in any particular way, although in most situations a simple additive composite should work just fine.

To begin compositing render passes using a Merge node:

1. Connect a Color pass to the background input of a Merge node.
2. Connect a Direct Lighting pass to the foreground input.
3. Adjust the Alpha gain and Blend parameters to get the look you desire.

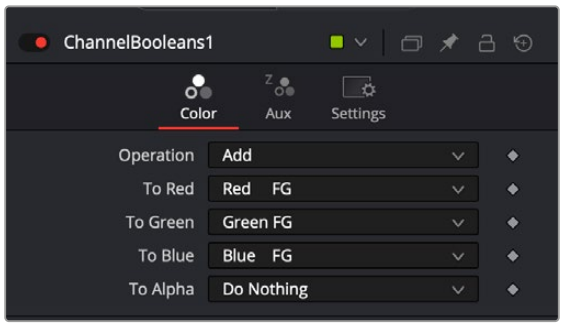


Compositing beauty passes starts by connecting a background and foreground into a Merge node

If you prefer, you can use a Channels Booleans node to make the same composite. In this case, there is no technical difference between the two nodes.

To composite render passes with a Channels Booleans node:

1. Connect a Color pass to the background input of a Channels Booleans node.
2. Connect a Direct Light pass to the foreground input.
3. Choose Add from the Operations menu
4. Choose Do Nothing from the Alpha To menu.



Channel Booleans is set to Add to combine the foreground input and the background input.

One of the exceptions to the steps above are Shadow passes, such as Ambient Occlusion. In that case, a *multiply* Apply mode is usually employed.

To composite an Ambient Occlusion render pass using a Merge node:

1. Connect last of a sequence of Merge node render passes to the background input of a Merge node.
2. Connect an Ambient Occlusion pass to the foreground input.
3. Choose Multiply from the Apply Mode menu.
4. Adjust the Gain and Blend parameters to get the look you desire.

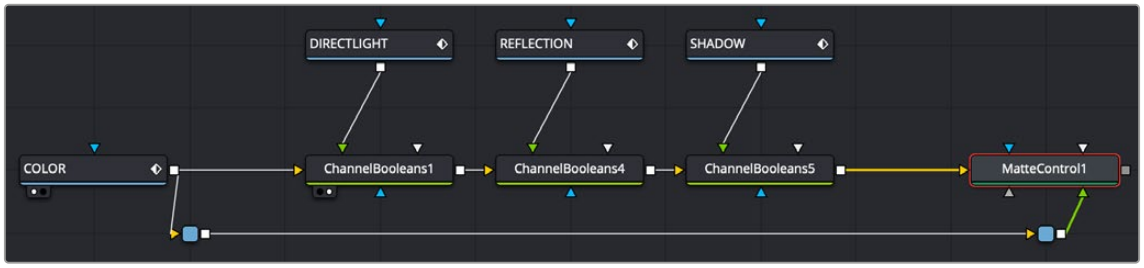
As straightforward as this sounds, compositing using a recipe doesn't always work for every shot. When using different images, you may need to experiment with varying techniques of compositing for the best results.

Embedding Alpha into Beauty Passes

Alpha channels are not included in all beauty passes. If your shot requires that you composite your assembled beauty passes over live-action or some other background, it may be up to you to add the alpha channel from a pass that does include it.

To add an alpha channel into your assembled beauty pass composite, do the following:

1. Connect the last Merge or Channel Booleans output into the background input of a Matte Control node.
2. Connect the render pass that contains the alpha into the green Foreground input of the Matte Control node.
3. In the Matte Control's Inspector, choose Combine Alpha from the Combine menu.
4. Choose Copy from the Combine Op menu.



An alpha channel from the color pass added back into the completed beauty pass node tree

TIP: Alpha channels from 3D renderings are typically premultiplied. That being the case, be sure to turn on the Pre Divide/Post Multiply checkbox on any node that performs color correction. If using more than one node in a row to perform color correction, use the Alpha Divide and Alpha Mult nodes instead.

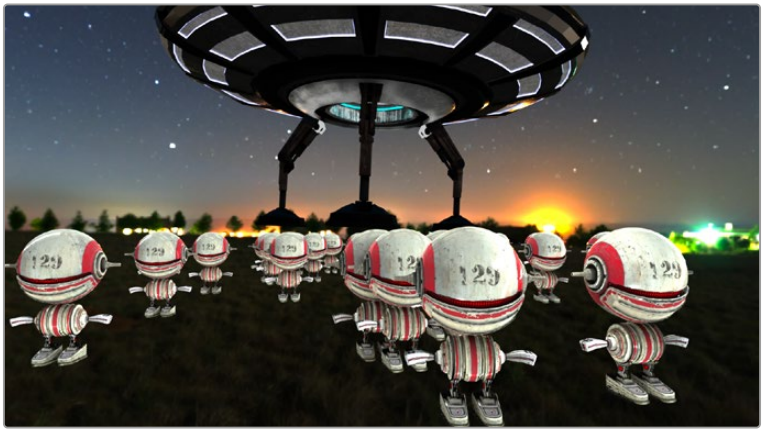
Working with Auxiliary Channels

Auxiliary channels are not RGB images; they are a family of special-purpose 3D image data that typically describes position, orientation, and object information for use in 2D composites. For example, Z-Depth channels describe the depth of each pixel of an image along a Z-axis, while an XYZ Normals channel describes the orientation (facing up, facing down, facing to the left or right) of every pixel in an image.

Similar to the use of multiple beauty passes, one of the most common reasons to use auxiliary data is to eliminate the need to re-render computationally expensive 3D imagery, by enabling even more aspects of rendered images to be manipulated after-the-fact. 3D rendering is computationally expensive and time-consuming, so outputting descriptive information about a 3D image allows sophisticated alterations to occur in 2D compositing, which is faster to perform and adjust.

There are two ways of obtaining auxiliary channel data:

- First, auxiliary data may be embedded within a clip rendered from a 3D application, most often using the EXR file format. In this case, it's best to consult your 3D application's documentation to determine which auxiliary channels can be generated and output.
- You may also obtain auxiliary channel data by generating it within Fusion, via 3D operations output by the Renderer 3D node, by the Optical Flow node, or by the Disparity node.

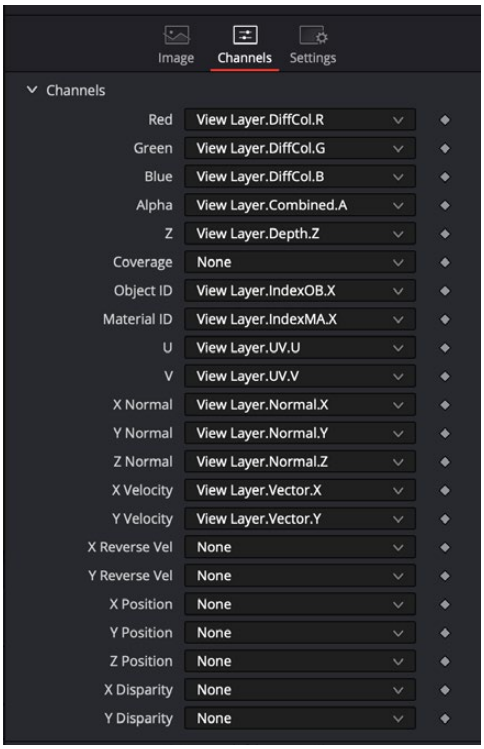


An RGBA 3D rendered scene that also contains auxiliary channels

TIP: When trying to locate information about auxiliary channels in other software, some 3D applications refer to auxiliary channels as Arbitrary Output Variables (AOVs), render elements, or secondaries.

Aux Channel Setup

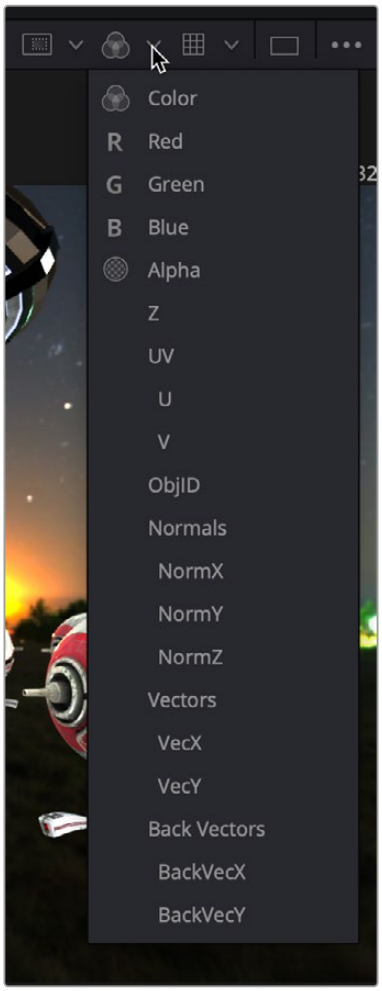
When using a MediaIn or Loader node that links to a multi-part EXR file with auxiliary channel data, the Inspector's Channels or Format tab includes a pre-defined set of auxiliary channels for mapping purposes. Each pre-defined channel includes a menu that displays every attribute included in a multi-part EXR. From the menu, you select the render pass that should be assigned to the corresponding channel. As described earlier, RGB beauty passes like diffuse, shadow, and reflection are mapped to the red, green, and blue channels. Auxiliary passes include preset mappings.



A multi-part EXR file with embedded render passes mapped to their aux channels in a MediaIn node

Displaying Channels in the Viewer

Once you map an auxiliary channel in the Inspector, you can display the data as an RGB image in the viewer. Clicking the drop down Color menu at the top of the viewer reveals a list of every active auxiliary channel for the currently viewed node.



Select an aux channel from the color drop-down menu to display the aux channel in the viewer

TIP: The Color Inspector SubView can be used to read numerical values from all of the channels.

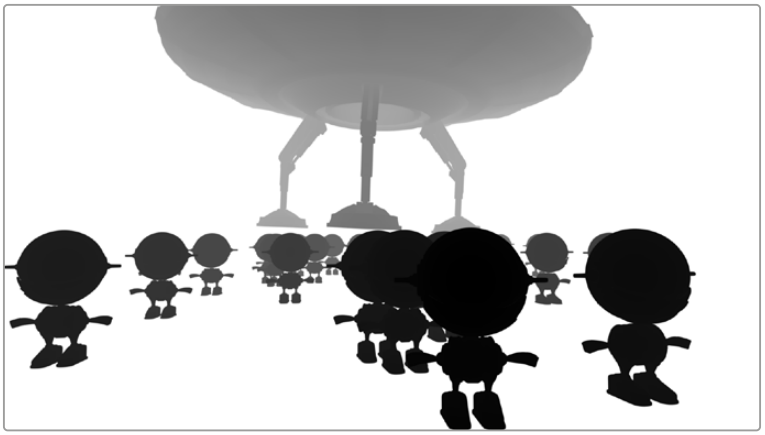
Auxiliary Channels Explained

Fusion is capable of using auxiliary channels to perform depth-based compositing, to create masks and mattes based on Object or Material IDs, and for texture replacements. Nodes that work with auxiliary channel information have been specifically developed to work with this data. The auxiliary channels that are supported in Fusion are described below.

Z-Depth

Each pixel in a Z-Depth channel contains a value that represents the relative depth of that pixel in the scene. In the case of overlapping objects in a model, most 3D applications take the depth value from the object closest to the camera when two objects are present within the same pixel since the closest object typically obscures the farther object.

When present, Z-Depth can be used to perform depth merging using the Merge node or to control simulated depth-of-field blurring using the Depth Blur node.

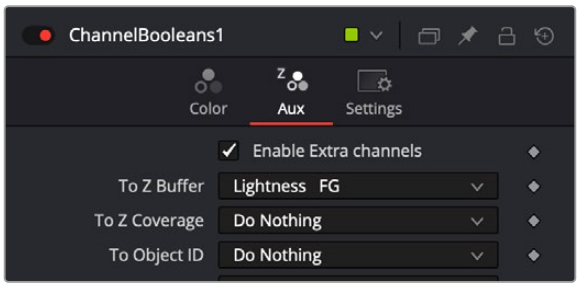


The rendered Z-Depth channel for the previous RGBA image

For this example, we'll examine the case where the Z-Depth channel is provided as a separate file. The Z- channel can often be rendered as an RGB image. You'll need to combine the beauty and Z pass using a Channel Booleans node. When the Z pass is rendered as an image in the RGB channels, the Channels Booleans node is used to re-shuffle the Lightness of the foreground RGB channel into the Z-channel.

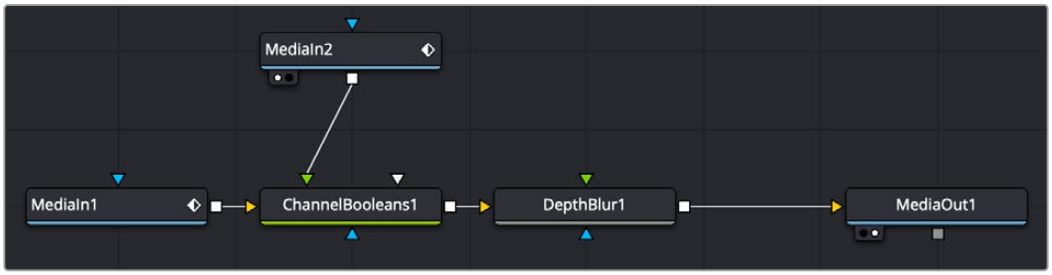
To combine the Z-pass and beauty pass:

1. Connect the MediaIn node containing the beauty pass to the background input of the Channel Booleans node.
2. Connect the MediaIn node containing the Z-Depth pass to the green foreground input of the Channel Booleans node.
3. Select the Channel Booleans node, and use the Inspector to set the To Red, To Green, To Blue, and To Alpha menus to Do Nothing.
4. Select the Aux tab, and set the To Z Buffer menu to Lightness FG.
5. Connect the output of the Channels Booleans node into the Depth Blur node.



The Aux tab configured to shuffle the Foreground Lightness to the Z-Depth channel.

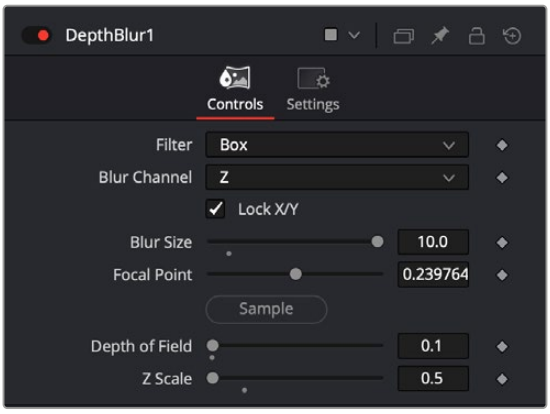
The Depth Blur node is one of the nodes that take advantage of a Z-channel in order to create blurry depth-of-field simulations. To set this up, the output of the MediaIn node connects to the background input on the Depth Blur.



The Depth Blur uses the Z-channel that is enabled in the Channel Booleans node.

The Depth Blur's controls in the Inspector are very dependent on the type of image you're using. It can be easier to begin by adjusting the controls in the Inspector to some better defaults. Start by increasing the Blur Size to 10. This will make it easier to see even the smallest of changes. Next, instead of using the Focal Point, you should pick a focal point in the image by dragging the Sample button into the viewer and selecting a pixel that determines the part of the picture to keep in focus.

The final setup steps are to lower Z Scale to somewhere around 0.2 (if you're using a floating-point image), and leave the Depth of Field alone for now. This should show you some blurring in the image.



Start by improving the defaults if your image is 16- or 32-bit floating point.

Once you see these experimental results, you can return to each parameter and refine it as needed to achieve the actual look you want.



An image using a Z-Depth channel for blurring

TIP: Z-Depth channels often contain negative values. If this causes problems, you can choose Normalize Color Range from the viewer's Options menu to apply a normalization to the viewer, keeping the image within a range from 0 to 1.

Z-Coverage

The Z-Coverage channel is a somewhat extinct render pass in most 3D applications. It was a way of restoring antialiasing to rendered color masks and Z-Depth passes. It indicated pixels in the Z-Depth that contained two objects. The value was used to indicate, as a percentage, how transparent that pixel was in the final depth composite. It can still be used today if you are rendering files from one of the few applications that can produce them.

TIP: The wide adoption of an open-source matte creation technology called Cryptomatte, has somewhat superseded mattes created from Coverage, Background, Object ID, and Material ID passes.

Background RGBA

This channel is a somewhat extinct render pass in most 3D applications. It contained the color values from the objects behind the pixels described in the Z coverage.

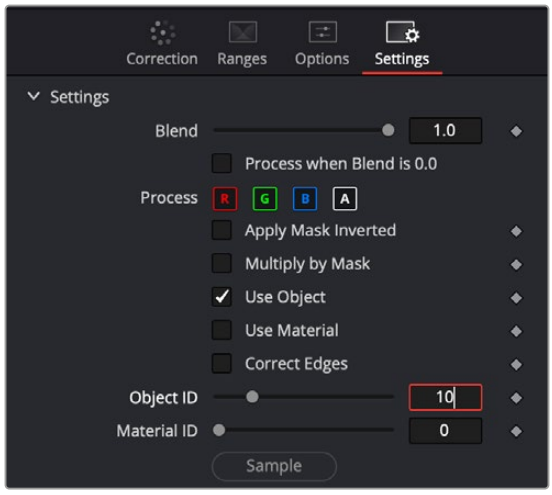
Object ID

Most 3D applications are capable of assigning ID values to objects in a scene. Each pixel in the Object ID channel will be identified by that ID number, allowing for the creation of masks.

If you want to use an Object ID in a comp, like all aux channels you must map the Object ID pass to the Object ID channel in the MediaIn or Loader Node.

To use an ObjectID pass, do the following:

1. In the MediaIn or Loader node, use the Channels or Format tab to map the Object ID pass to the Object ID aux channel.
2. In whatever node you want to have affected by the ObjectID matte, select the Settings tab, turn on the Object ID checkbox, and select the ID number assigned to the object.



The common Settings tab on most nodes contains ObjectID controls.

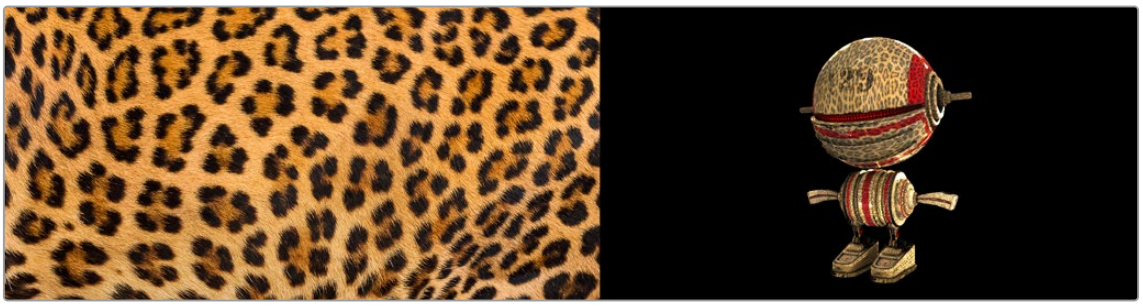
Material ID

Most 3D applications are capable of assigning ID values to materials in a scene. Each pixel in the Material ID channel will be identified by that ID number, allowing for masks based on materials.

You can set up Material IDs using the Settings tab, similarly to how ObjectIDs are set.

UV Texture

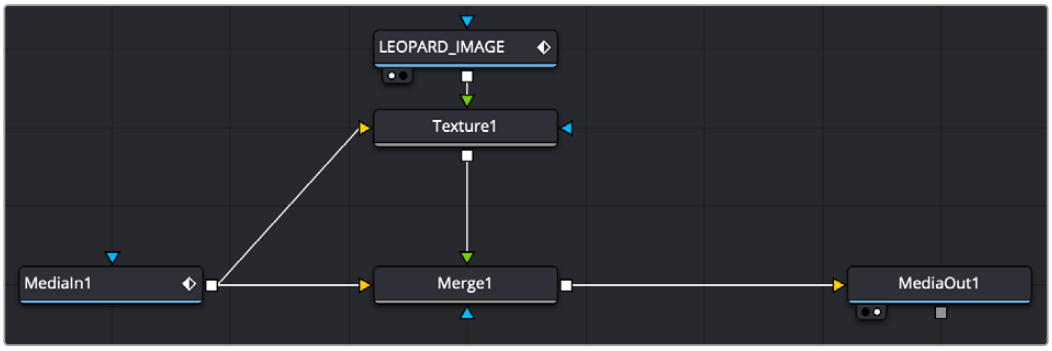
The UV Texture channels contain information about how pixels in an image map to texture coordinates. This is used to retexture an object in a 2D image. For instance, If you want to apply a logo onto a rendered object, you can use the UV aux channel with the Texture node.



Texture (left) applied to 2D image (right) using UV channels and texture node.

To use a UV pass, do the following:

1. In the MediaIn or Loader node, use the Channels or Format tab to map the U and V pass to the U and V aux channels.
2. Connect the output of the MediaIn or Loader node to the background input of the Texture node.
3. Connect the texture image you want to use to the foreground input of the Texture node.
4. If you want to combine the original texture with the new texture, use a merge with the background input from the original image and the foreground input from the Texture node.
5. Adjust Merge's Apply mode, Alpha Gain, and blend to get the desired mix of the two textures.

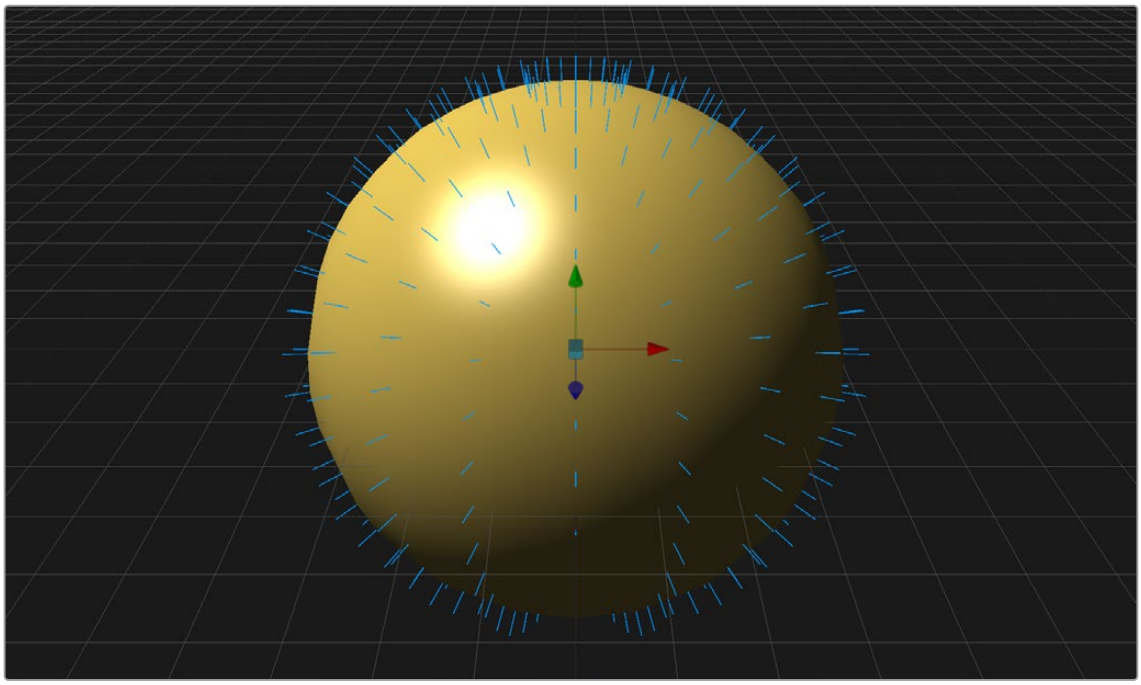


UV channels from a MediaIn node used in a Texture node and merged over the original image

TIP: If you are using a separate UV render pass with the UV data in the RGB channels, map red to U and green to V in a Channel Booleans node.

X, Y, and Z Normals

The X, Y, and Z Normal channels contain information about each pixel's orientation (the direction it faces) in 3D space. The normals are often displayed as lines coming out from your object perpendicular to the surface, letting you visualize the relationship between the surface and camera.



The Normals display the direction of the surface.

The Normals X, Y, and Z channels are often used with a Shader node to perform relighting adjustments on a 2D rendered image.

To setup a Shader node to use XYZ Normals, do the following:

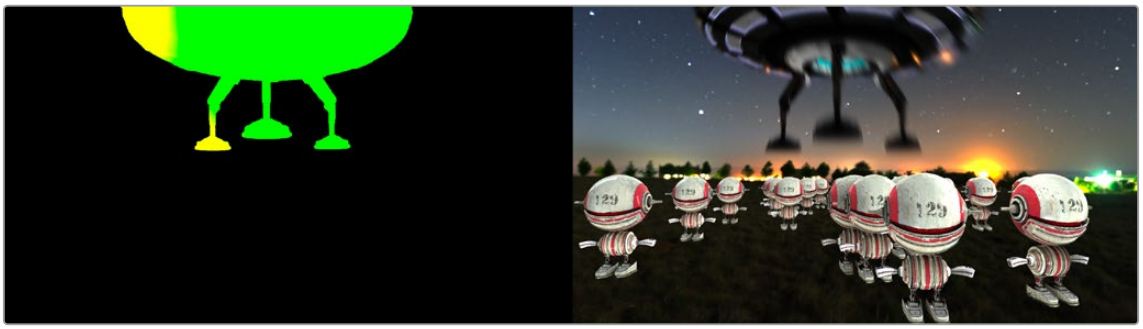
1. In the MediaIn or Loader node, use the Channels or Format tab to map the individual X, Y, and Z Normals pass to the X Normal, Y Normal, and Z Normal channels.
2. Connect the output of the MediaIn or Loader node to the background input of the Shader node.
3. Optionally, connect a floating-point EXR image to be used as a reflection image to the Shader node's reflection input.
4. Adjust the Shader controls to perform relighting.



Original 2D image (left) and Normals used for relighting (right)

XY Vector and XY BackVector

The Vector channels indicates the pixel's motion from frame to frame. It can be used to apply motion blur to an image or to generate optical flow analysis for retiming. The XY Vector points to the next frame, while the XY BackVector points to the previous frame.

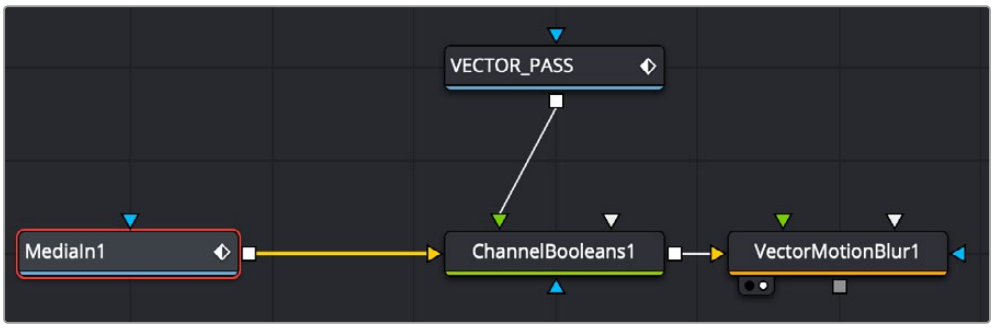


XY Vector pass (left) used with Vector Motion Blur to generate motion blur on spaceship (right)

Often the vector pass will be rendered in a separate pass as an RGB image. The X and Y vector data is located in the R and G channels. In order to place them in the vector channels, you can use a Channel Booleans node.

To use a Motion Vector pass to create motion blur, do the following:

1. Add a MediaIn or Loader node for the image and the Vector render pass.
2. Connect the output of the image into the background of a Channel Booleans node.
3. Connect the output of the Vector render pass to the Channel Boolean's foreground
4. In the Channel Booleans inspector, set the To Red, To Green, To Blue, and To Alpha all to Do Nothing.
5. Select to the Aux tab.
6. Turn on Enable Extra Channels.
7. Set the To X Vector drop-down menu to Red FG, and then set the To Y Vector drop-down menu to Green FG.
8. Connect the Channel Booleans node's output to the yellow background input on a Vector Motion Blur node.

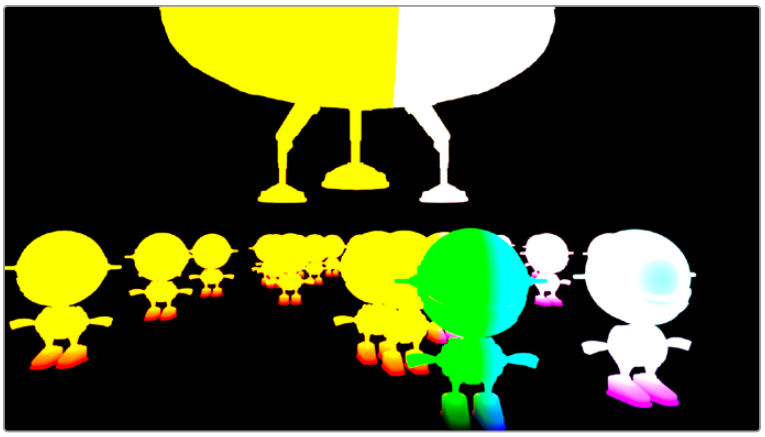


The Vector render pass is combined with the beauty image using the Channels Booleans node, which then feeds the Vector Motion Blur node.

World Position

World Position Pass (WPP) is an auxiliary channel, sometimes referred to as Point Position, XYZ pass, or WPP. It's used to represent each pixel's 3D (XYZ) position as an RGB color value. The result is data that can be viewed as a very colorful RGB image. Like Z-Depth, this can be used for compositing via depth. However, it can also be used for masking based on 3D position, regardless of camera transforms.

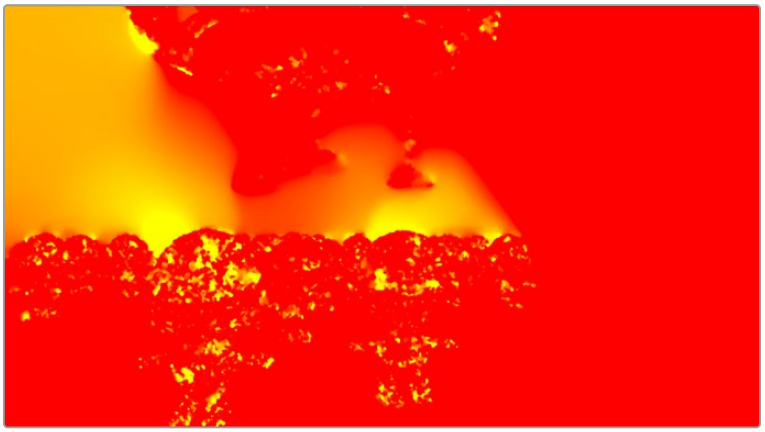
The colors correspond to a pixel's position in 3D, so if a pixel sits at 0/0/0 in a 3D scene, the resulting pixel's will have an RGB value of 0/0/0 or black. If a pixel sits at 1/0/0 in the 3D scene, the resulting pixel is fully red. Due to the huge extent, 3D scenes can have the WPP channel should always be rendered in 32-bit floating-point to provide the accuracy needed.



XYZ Position

XY Disparity

XY Disparity is the only channel listed here that is not generated in a 3D application. These channels indicate where each pixel's corresponding matte can be found in a stereo image. Each eye, left and right, will use this vector to point to where that pixel would be in the other eye. This can be used for adjusting stereo effects, or to mask pixels in stereo space.



Combined X and Y Disparity channels

Using Cryptomatte in Fusion

Cryptomatte is an open-source technology that has been widely adopted in 3D applications. Unlike Z-Depth mattes or Object IDs, Cryptomatte automatically generates anti-aliased ID mattes from 3D renders with support for motion blur, transparency, and depth of field.

Fusion does not natively support the Cryptomatte format. However, using a free plugin from third-party developers, you can use Crypotmatte render passes in Fusion.

Crypotmatte for Fusion can be downloaded and installed for free:
https://github.com/Psyop/Cryptomatte

Or to use an easier installer, you can download Reactor, which comes bundled with Cryptomatte and offers many other free, useful Fusion plugins. Reactor can be found at:
https://www.steakunderwater.com

Propagating Auxiliary Channels

Ordinarily, auxiliary channels are propagated along with RGBA image data, from node to node, among gray-colored nodes, including those in the Blur, Filter, Effect, Transform, and Warp categories. Basically, most nodes that simply manipulate channel data propagate (and potentially manipulate) auxiliary channels with no problems.

However, when you composite two image layers using the Merge node, auxiliary channels only propagate through the image that's connected to the background input. The rationale for this is that in most composites that include computer-generated imagery, the background is most often the CG layer that contains auxiliary channels, while the foreground is a live-action green screen plate with subjects or elements that are combined against the background, which lack auxiliary channels.

Nodes That Use Auxiliary Channels

The availability of auxiliary channels opens up a world of advanced compositing functionality. This section describes every Fusion node that has been designed to work with images that contain auxiliary channels.

- **Copy Aux:** The Copy Aux tool can copy auxiliary channels to RGB and then copy them back. It includes some useful options for remapping values and color depths, as well as removing auxiliary channels.
- **Channel Booleans:** The Channel Boolean tool can be used to combine or copy the values from one channel to another in a variety of ways.
- **Custom Tool, Custom Vertex 3D, pCustom:** The "Custom" tools can sample data from the auxiliary channels per pixel, vertex, or particle and use that for whatever processing you would like.
- **Depth Blur:** The Depth Blur tool is used to blur an image based on the information present in the Z-Depth. A focal point is selected from the Z-Depth values of the image and the extent of the focused region is selected using the Depth of Field control. The Scale value default is based on an 8-bit image so it is important to lower the scale value when using the Depth Blur with 16- or 32-bit float files.
- **Disparity to Z, Z to Disparity, Z to WorldPos:** These tools use the inherent relationships between depth, position, and disparity to convert from one channel to another.
- **Fog:** The Fog tool makes use of the Z-Depth to create a fog effect that is thin closer to the camera and thickens in regions farther away from the camera. You use the Pick tool to select the Depth values from the image and to define the Near and Far planes of the fog's effect.

- **Lumakeyer:** The Lumakeyer tool can be used to perform a key on the Z-Depth channel by selecting the Z-Depth in the channel drop-down list.
- **Merge:** In addition to regular compositing operations, Merge is capable of merging two or more images together using the Z-Depth, Z-Coverage, and BG RGBA buffer data. This is accomplished by enabling the Perform Depth Merge checkbox from the Channels tab.
- **New Eye:** For stereoscopic footage, New Eye uses the Disparity channels to create new viewpoints or to transfer RGBA data from one eye to the other.
- **Shader:** The Shader tool applies data from the RGBA, UV and the Normal channels to modify the lighting applied to objects in the image. Control is provided over specular highlights, ambient and diffuse lighting, and position of the light source. A second image can be applied as a reflection or refraction map.
- **Shadow:** The Shadow tool can use the Z-Depth channel for a Z-Map. This allows the shadow to fall onto the shape of the objects in the image.
- **Smooth Motion:** Smooth Motion uses Vector and Back Vector channels to blend other channels temporally. This can remove high frequency jitter from problematic channels such as Disparity.
- **SSAO:** SSAO is short for Screen Space Ambient Occlusion. Ambient Occlusion is the lighting caused when a scene is surrounded by a uniform diffuse spherical light source. In the real world, light lands on surfaces from all directions, not from just a few directional lights. Ambient Occlusion captures this low frequency lighting, but it does not capture sharp shadows or specular lighting. For this reason, Ambient Occlusion is usually combined with Specular lighting to create a full lighting solution. The SSAO tool uses the Z-Depth channel but requires a Camera3D input.
- **Stereo Align:** For stereoscopic footage, the Disparity channels can be used by Stereo Align to warp one or both of the eyes to correct misalignment or to change the convergence plane.
- **Texture:** The Texture tool uses the UV channels to apply an image from the second input as a texture. This can replace textures on a specific object when used in conjunction with the Object ID or Material ID masks.
- **Time Speed and Time Stretcher:** These tools can use the Vector and BackVector channels to retime footage.
- **Vector Distortion:** The forward XY Vector channels can be used to warp an image with this tool.
- **Vector Motion Blur:** Using the forward XY Vector channels, the Vector Motion Blur tool can apply blur in the direction of the velocity, creating a motion blur effect.
- **Volume Fog:** Volume Fog is a raymarcher that uses the World Position channels to determine ray termination and volume dataset placement. It can also use cameras and lights from a 3D scene to set the correct ray start point and Illumination parameters.
- **Volume Mask:** Volume Mask uses the World Position channels to set a mask in 3D space as opposed to screen space. This allows a mask to maintain perfect tracking through a camera move.

TIP: The Object ID and Material ID auxiliary channels can be used by some tools in Fusion to generate a mask. The “Use Object” and “Use Material” settings used to accomplish this are found in the Settings tab of that node’s controls in the Inspector.

Image Formats That Support Aux Channels

Fusion supports auxiliary channel information contained in a variety of image formats. The number of channels and methods used are different for each format.

- **OpenEXR (*.exr):** The OpenEXR file format is the primary format used to contain an arbitrary number of additional image channels. Many renderers that will write to the OpenEXR format will allow the creation of channels that contain entirely arbitrary data. For example, a channel with specular highlights might exist in an OpenEXR. In most cases, the channel will have a custom name that can be used to map the extra channel to one of the channels recognized by Fusion.
- **SoftImage PIC (*.PIC, *.ZPIC and *.Z):** The PIC image format (used by SoftImage) is an older image format that can contain Z-Depth data in a separate file marked by the ZPIC file extension. These files must be located in the same directory as the RGBA PIC files and must use the same names. Fusion will automatically detect the presence of the additional information and load the ZPIC images along with the PIC images.
- **Wavefront RLA (*.RLA), 3ds Max RLA (*.RLA) and RPF (*.RPF):** This is an older image format capable of containing any of the image channels mentioned above. All channels are contained within one file, including RGBA, as well as the auxiliary channels. These files are identified by the RLA or RPF file extension. Not all RLA or RPF files contain auxiliary channel information, but most do. RPF files have the additional capability of storing multiple samples per pixel, allowing different layers of the image to be loaded for very complex depth composites.
- **Fusion RAW (*.RAW):** Fusion's native RAW format is able to contain all of the auxiliary channels as well as other metadata used within Fusion.

Creating Auxiliary Channels in Fusion

The following nodes create auxiliary channels:

- **Renderer 3D:** Creates these channels in the same way as any other 3D application would, and you have the option of outputting every one of the auxiliary data channels that the Fusion page supports.
- **Optical Flow:** Generates Vector and Back Vector channels by analyzing pixels over consecutive frames to determine likely movements of features in the image.
- **Disparity:** Generates Disparity channels by comparing stereoscopic image pairs.

Chapter 19

Compositing Layers in Fusion

This chapter is intended to give you a solid base for making the transition from a layer-based compositing application to Fusion's node-based interface. It provides practical information about how to start structuring a node tree for simple layered composites.

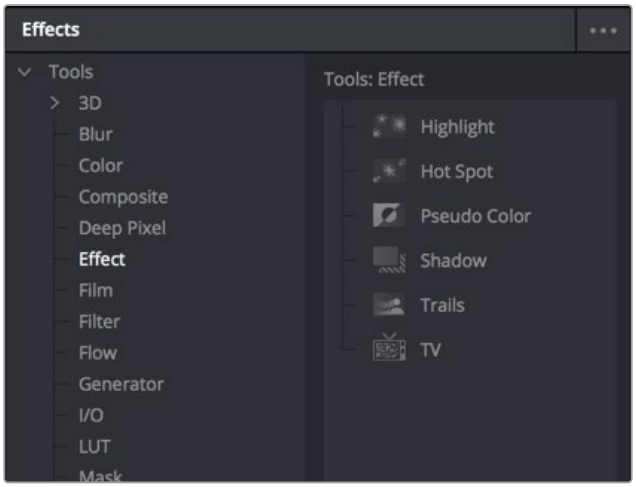
Contents

Applying Effects 480
Adding a Node to the Tree 480
Editing Parameters in the Inspector 481
Replacing Nodes 482
Adjusting Fusion Sliders 482
Compositing Two Clips Together 483
Adding Additional Media to Compositions 483
Automatically Creating Merge Nodes 484
Fixing Problem Edges in a Composite 484
Using Composite Modes in the Merge Node 486
Creating and Using Text 487
Creating Text Using the Text+ Node 487
Styling and Adjusting Text 488
Using Text as a Mask 490
Using Transform Controls in the Merge Node 493
Building a Simple Green-Screen Composite 494
Mapping Timeline Layers to Nodes in Fusion 494
Pulling a Green-Screen Key Using the Delta Keyer 496
Dealing with Spill 500
Masking a Graphic 500

Applying Effects

Before we dive into multi-layered composites, let's start by looking at some very simple effects and build up from there. Opening the Effects Library, then clicking the disclosure control to the left of Tools, reveals a list of categories containing all the nodes available in Fusion. As mentioned before, each node does one thing, and by using these nodes in concert you can create extremely complex results from humble beginnings.

Clicking the Effect category reveals its contents. In this example we'll use the TV effect.



Browsing the Effect category to find the TV node

Adding a Node to the Tree

Assuming that the MediaIn node in the Fusion page or the Loader node in Fusion Studio is the currently selected node in the Node Editor, clicking once on the TV node, for example, in the Effects Library automatically adds that node to the node tree to the right of the selected node. In the Fusion page, it immediately takes effect in the viewer thanks to the fact that the MediaOut1 node is what's loaded in the viewer, since that means that all nodes upstream of the MediaOut1 node will be processed and shown.



A new node added from the Effects Library

In Fusion Studio, you must press the 1 or 2 key on the keyboard to load the selected node in the viewer.

There are many other ways of adding nodes to your node tree, but it's good to know how to browse the Effects Library as you get started.

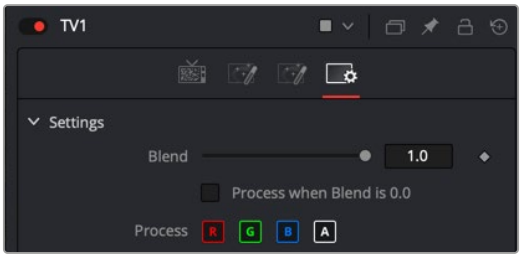
Editing Parameters in the Inspector

To improve an effect, you can make an adjustment to a node's parameters in the Inspector at the right. The selected node shows its controls in the Inspector, in which most nodes have several tabs of controls seen as little icons just underneath that node's title bar.



The Inspector showing the parameters of the TV effect

Clicking the last panel on any node opens the Settings panel. Every node has a Settings panel, and this is where the parameters that every node shares, such as the Blend slider and RGBA buttons, are found. These let you choose which image channels are affected, and let you blend between the effect and the original image.



The Settings panel, which includes channel limiting and mask handling controls that every node shares

In the case of the TV effect, for example, the resulting image has a lot of transparency because the scan lines being added are also being added to the alpha channel, creating alternating lines of transparency. Turning off the Alpha checkbox results in a more solid image, while opening the Controls tab (the first tab) and dragging the Scan Lines slider to the right to raise its value to 4 creates a more visible television effect.

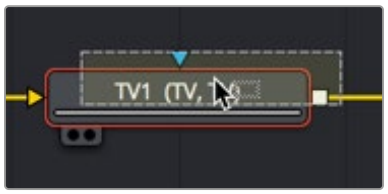


The original TV effect (left), and modifications to the TV effect to make the clip more solid (right)

Replacing Nodes

In the Effect category of the Effects Library, you'll also find a Highlight node that adds glints to the highlights of an image.

Instead of clicking the Highlight node, which would add it after the currently selected node, dragging and dropping a node from the Effects Library on top of a node in the Node Editor replaces the node in the Node Editor.



Dragging a node from the Effects Library onto a node in the Node Editor to replace it

In our example, the Highlight1 node takes the TV node's place in the node tree, and the new effect can be seen in the viewer, which in this example consists of star highlights over the lights in the image.

It's time to use the Inspector controls to customize this effect.

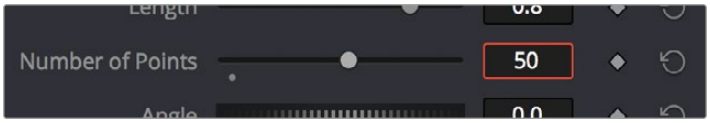
Adjusting Fusion Sliders

When you drag a slider in the Fusion Inspector—in this case, the Number of Points slider—a little dot appears underneath it. This dot indicates the position of the default value for that slider and also serves as a reset button if you click it.



Adjusting a slider reveals a reset button underneath it

Each slider is limited to a different range of minimum and maximum values that is particular to the parameter you're adjusting. In this case, the Number of Points slider maxes out at 24. However, you can remap the range of many (but not all) sliders by entering a larger value in the number field to the right of that slider. Doing so immediately repositions the slider's controls to the left as the slider's range increases to accommodate the value you just entered.



Entering a larger value to expand the range over which a slider will operate

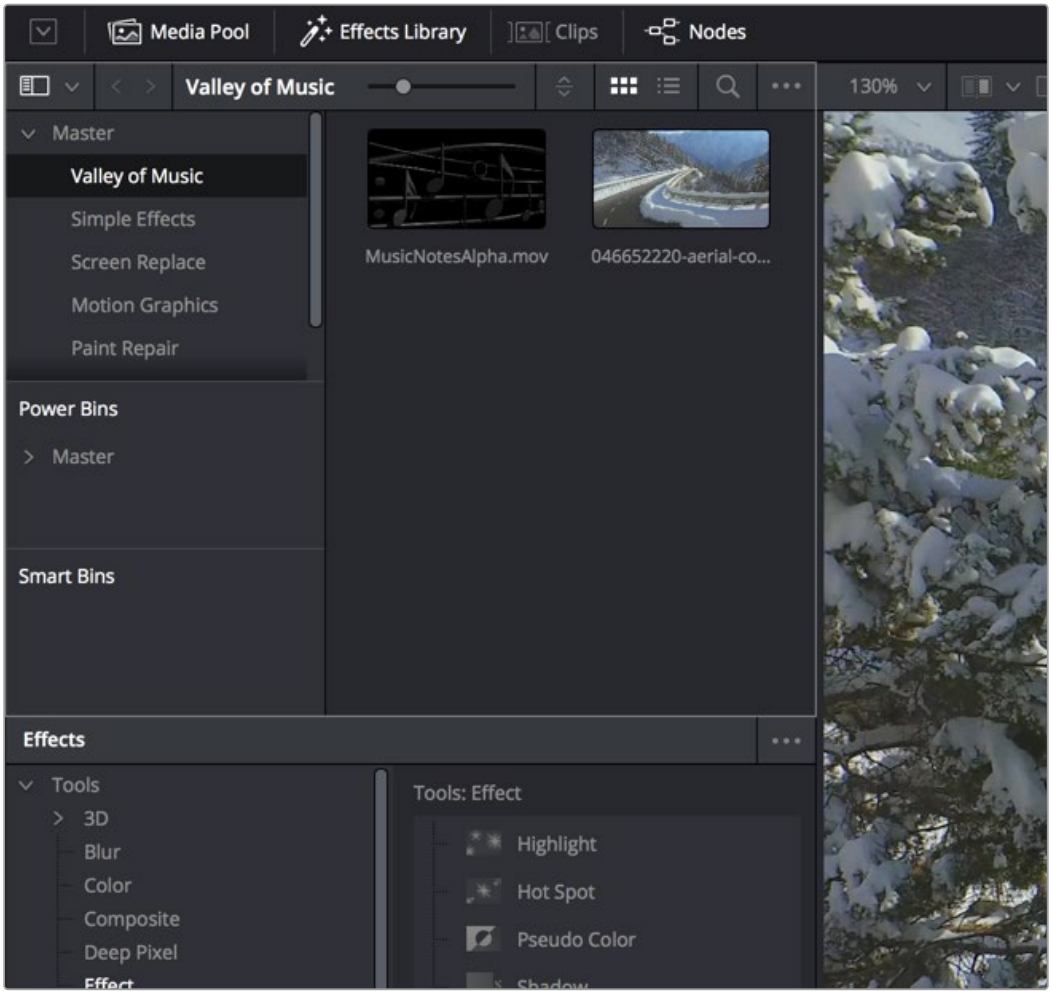
Compositing Two Clips Together

As entertaining as it is adding individual nodes to create simple effects, eventually you'll need to start adding additional layers of media in order to merge them together as composites. Let's turn our attention to another example in which we need to combine a background clip with a foreground clip that includes a built-in alpha channel, to see simple layering in action.

Adding Additional Media to Compositions

You'll often find that even though you start out wanting to do something relatively simple, you end up adding additional media to create the effect that you need.

- In Fusion Studio, you do this by adding additional Loader nodes. If you add a new Loader node to an empty area of the Node Editor, you'll add an unconnected Loader2 node (incremented to keep it unique) that you can then connect how you want.
- In the Fusion page, you can open the Media Pool and drag clips directly to the Node Editor to add them to your node tree. If you drag a clip from the Media Pool to an empty area of the Node Editor, you'll add an unconnected MediaIn2 node (incremented to keep it unique) that you can then connect in any way you want.

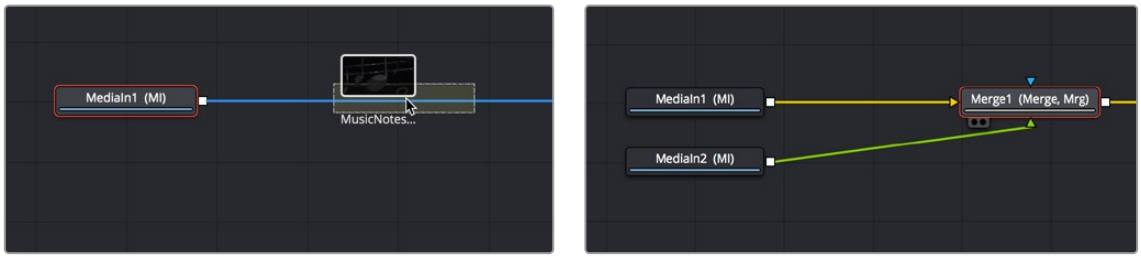


The Media Pool as seen in the Fusion page

Automatically Creating Merge Nodes

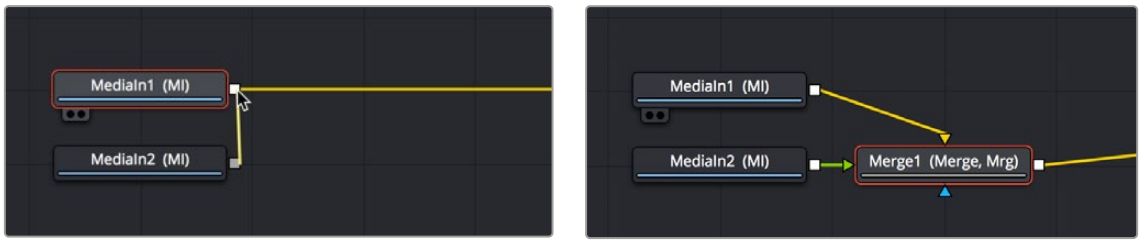
If you want to connect the incoming clip immediately to your node tree as the top layer, or foreground, of a composite, in Fusion Studio select the Loader1 node and then add a second Loader node. In the Fusion page, drag the new clip from the Media Pool right on top of any connection line.

In both cases, the new MediaIn or Loader node automatically becomes the “foreground input”.



Dragging a node from the Media Pool onto a connection (left), and dropping it to create a Merge node composite (right)

The Node Editor is filled with shortcuts like this to help you build your compositions more quickly. Here’s one for when you have a disconnected node that you want to composite against another node with a Merge node. Drag a connection from the output of the node you want to be the foreground layer, and drop it on top of the output of the node you want to be the background layer. A Merge node will be automatically created to build that composite. Remember: background inputs are orange, and foreground inputs are green.



Dragging a connection from a disconnected node to another node’s output (left), and dropping it to create a Merge node composite (right)

Adding Clips to a Fusion Composition From the File System

If you drag clips from the file system directly into the Node Editor, they’ll be added to the DaVinci Resolve Media Pool automatically. So, if you have a library of stock animated background textures and you’ve just found one you want to use using your file system’s search tools, you can simply drag it straight into the Node Editor and it’ll be added to the currently selected bin of the Media Pool.

Fixing Problem Edges in a Composite

Most often, the Merge node does a perfectly good job when handed a foreground image with premultiplied alpha transparency to composite against a solid background image. However, from time to time, you may notice a small bit of fringing at the edge of the border of a foreground element and transparent area, such as seen in the following close-up. This slight lightening at the edge is a tell-tale sign that the clip probably wasn’t premultiplied. The Merge node expects all foreground images with alpha channels to be premultiplied. But this is something that’s easily fixed.



A bit of fringing at the edge of a foreground element surrounded by transparency

Click to select the Merge node for that particular composite, and look for the Subtractive/Additive slider.



The Subtractive/Additive slider, which can be used to fix or improve fringing in composites

Drag the slider all the way to the left, to the Subtractive position, and the fringing disappears.



A clip with alpha exhibits fringing (left), and after fixing fringing by dragging the Subtractive/Additive slider to the left (right)

The Subtractive/Additive slider, which is only available when the Apply mode is set to Normal, controls whether the Normal mode performs an Additive merge, a Subtractive merge, or a blend of both. This slider defaults to Additive merging, which assumes that all input images with alpha transparency are premultiplied (which is usually the case). If you don't understand the difference between Additive and Subtractive merging, here's a quick explanation:

- An Additive merge, with the slider all the way to the right, is necessary when the foreground image is premultiplied, meaning that the pixels in the color channels have been multiplied by the pixels in the alpha channel. The result is that transparent pixels are always black, since any number multiplied by 0 is always 0. This obscures the background (by multiplying with the inverse of the foreground alpha), and then simply adds the pixels from the foreground.
- A Subtractive merge, with the slider all the way to the left, is necessary if the foreground image is not premultiplied. The compositing method is similar to an Additive merge, but the foreground image is first multiplied by its own alpha to eliminate any background pixels outside the alpha area.

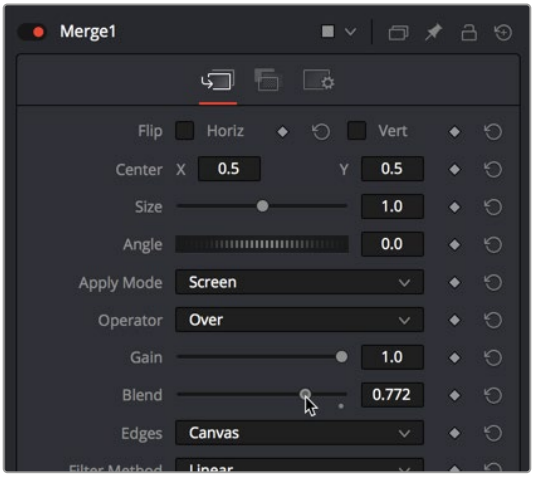
The Additive/Subtractive slider lets you blend between two versions of the merge operation, one Additive and the other Subtractive, to find the best combination for the needs of your particular composite. Blending between the two is occasionally useful for dealing with problem composites that have edges that are calling attention to themselves as either too bright or too dark.

For example, using Subtractive merging on a premultiplied image may result in darker edges, whereas using Additive merging with a non-premultiplied image will cause any non-black area outside the foreground's alpha to be added to the result, thereby lightening the edges. By blending between Additive and Subtractive, you can tweak the edge brightness to be just right for your situation.

Using Composite Modes in the Merge Node

To create a more convincing composite in layer-based systems, you often use Blend modes. Blend modes are located in the Merge node since that is where one layer gets composted over another. Let's take an example where you want to use the Screen mode to make a foreground image look more like a reflection.

The Merge node has a variety of controls built into it for creating just about every compositing effect you need. Items you may be familiar with as Blend modes are located in the Apply Mode pop-up menu. You can use these mathematical compositing modes to combine the foreground and background layers together. A Blend slider allows you to fade the foreground input with the background.



Adjusting the Apply Mode and Blend slider of the Merge node in the Inspector

NOTE: The Subtractive/Additive slider disappears when you choose any other Apply Mode option besides Normal, because the math would be invalid. This isn't unusual; there are a variety of controls in the Inspector that hide themselves when not needed or when a particular input isn't connected.

The Screen node is perfect for simulating reflections, and lowering Blend a bit lets you balance the foreground and background images. It's subtle, but helps sell the shot.

TIP: You may have noticed that the Merge node also has a set of Flip, Center, Size, and Angle controls that you can use to transform the foreground image without needing to add a dedicated Transform node. It's a nice shortcut for simplifying node trees large and small.

Creating and Using Text

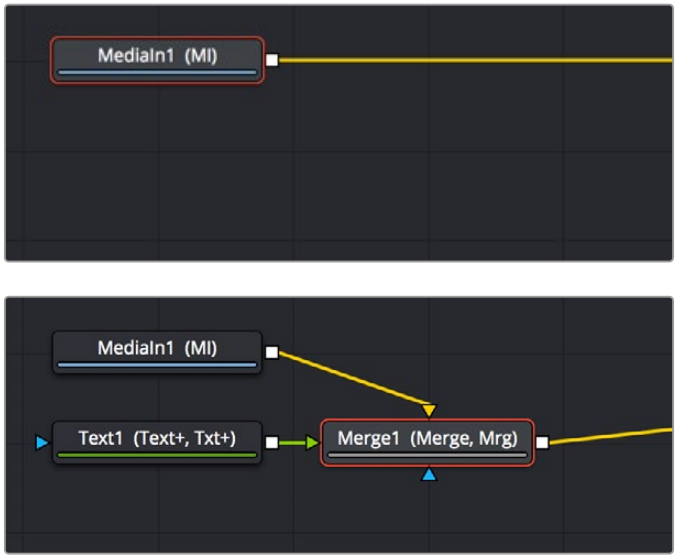
In this next example, we'll look at how to create a simple text object using the Text+ node. Then, we'll see how to use the text generator's alpha channel in another image to create a more complex composite.

Creating Text Using the Text+ Node

The Text+ node is the primary tool for creating 2D text in the Fusion page. If you are using DaVinci Resolve, this is also the same Text+ generator available in the Edit page. It is easily accessible right in the toolbar.

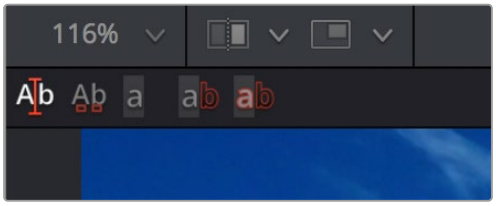
The Text+ node is an incredibly deep tool for creating text effects, with six tabs of controls for adjusting everything from text styling, to different methods of layout, to a variety of shading controls including fills, outlines, shadows, and borders. As sophisticated a tool as this is, we'll only be scratching the surface in this example.

We'll start with a MediaIn node that will serve as our background selected in the Node Editor. Clicking the Text+ button automatically creates a new Text+ node connected as the foreground input of a Merge node. The same behavior occurs if you are using Fusion Studio, with a Loader node.



Selecting a node you want to append another node to (top) Clicking the Text+ button on the toolbar automatically creates a Merge composite with the text as the foreground input connection (bottom)

Selecting a Text node opens the default Text panel parameters in the Inspector, and it also adds a toolbar at the top of the viewer with tools specific to that node. Clicking on the first tool at the left lets you type directly into the viewer, or you can type into the Styled Text field in the Inspector.



The viewer toolbar for the Text node with tools for text entry, kerning, and outline controls

If you're viewing the Merge, the text appears in the viewer superimposed against the background clip. Onscreen controls appear that let you rotate (the circle) and reposition (the red center handle and two arrows) the text, and we can see a faint cursor that lets us edit and kern the text using other tools in the viewer toolbar.



Text that's been typed into the viewer, with onscreen text transform controls

Styling and Adjusting Text

To style the text, you use the controls in the Inspector, modifying text style controls such as Font, Size, and Tracking to change the spacing between the letters.

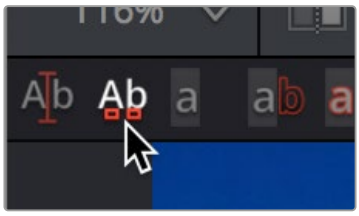


The restyled text

TIP: Holding down the Command key while dragging any control in the Inspector "gears down" the adjustment so that you can make smaller and more gradual adjustments.

Selecting the Manual Kerning tool in the viewer toolbar (second tool from the left) reveals small red dots underneath each letter of text.

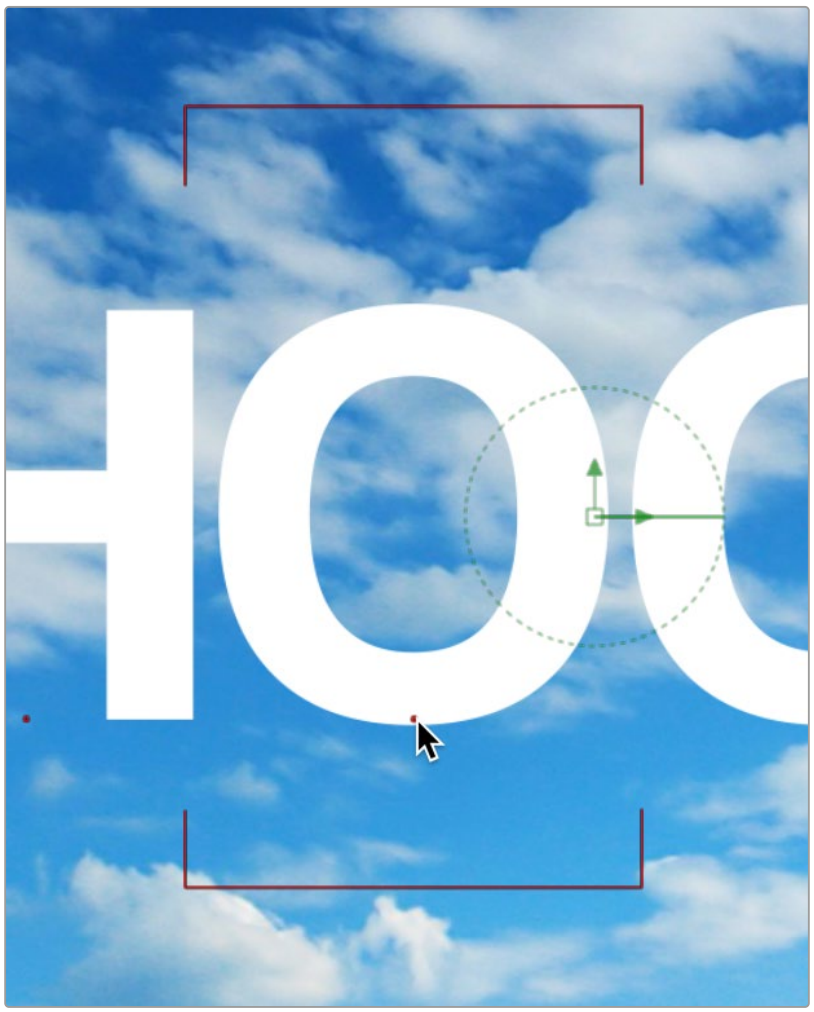
Clicking a red dot under a particular letter puts a kerning highlight over that letter.



The Manual Kerning tool in the viewer toolbar

To make manual kerning adjustments:

1. Option-drag the red dot under any letter of text to adjust that character's kerning while constraining letter movement to the left and right. You can also drag letters up and down for other effects. Depending on your system, the kerning of the letter you're adjusting might not update until you drop the red dot in place.
2. If you don't like what you've done, you can open the Advanced Controls in the Inspector and clear either the kerning of selected letters or all manual kerning before starting over again.



Option-dragging the little red dot revealed by the Manual Kerning tool to manually adjust kerning

Using Text as a Mask

You can fill the text with a color or gradient, or you can use the text as a matte to cut the letters out of another image. First, we'll drag another clip of a chalkboard covered with math from the Media Pool to the Node Editor as a disconnected MediaIn2 node.

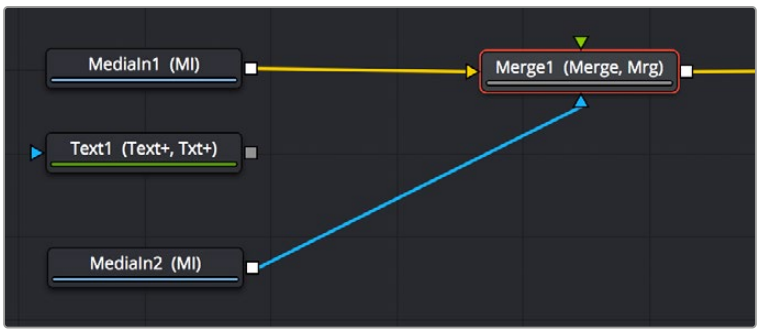
Disconnecting and Reconnecting Nodes

If we start from the previous example that had a MediaIn as the background and the Text+ as the foreground to a Merge, we'll need to do a little rearranging. Clicking the last half of the connection from the Text1 node to the Merge foreground input disconnects it.



Clicking the second half of a connection to disconnect it (previous page), and the result with the text node disconnected (above)

Connecting a MediaIn2 or Loader2 node onto the Merge1 node's foreground input causes the entire viewer to be filled with the MediaIn2 (assuming we're still viewing the Merge node). At this point, we need to insert the Text1 node's image as an alpha channel into the MediaIn2 node's connection, and we can do that using a MatteControl node.

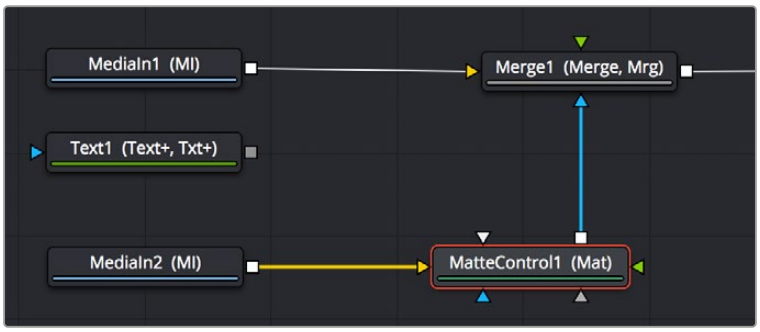


The updated composite, with two video images connected and the text node disconnected

Using Matte Control Nodes

Select the MediaIn2 node, and click the Matte Control button in the toolbar to add it between the MediaIn2 and Merge1 nodes. (To tidy up, I've moved the nodes around a bit in the screenshot.)

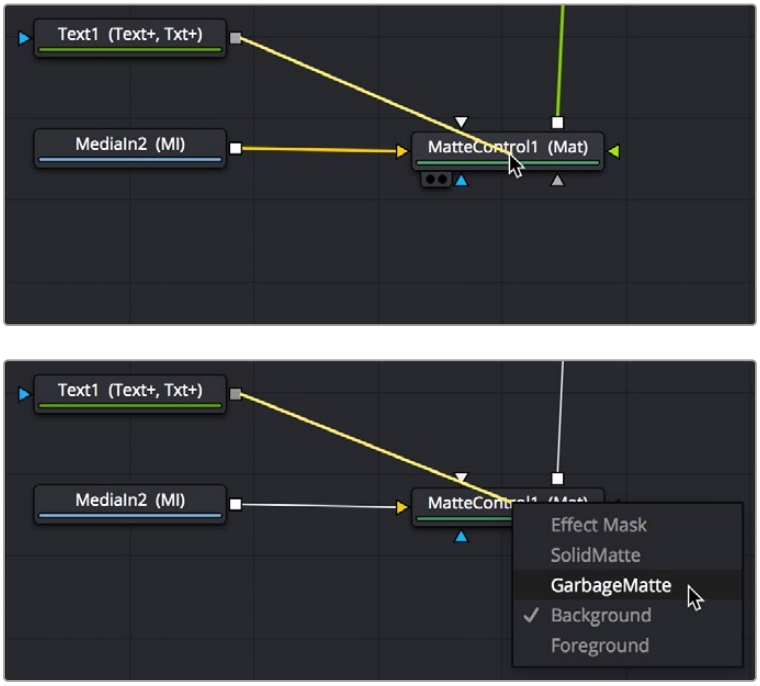
The MatteControl node has numerous uses. Among them is taking one or more alpha channels, mattes, or images that are connected to the Garbage Matte, Solid Matte, and/or foreground inputs, combining them, and using the result as an alpha channel for the image that's connected to the background input. It's critical to make sure that the image you want to add an alpha channel to is connected to the background input of the MatteControl node, or the MatteControl node won't work.



The second image properly connected to the MatteControl node's background input

With this done, connecting the Text+ node's output, which has the alpha channel, to the MatteControl node's Garbage Matte input, is a shortcut we can use to make a mask, matte, or alpha punch out a region of transparency in an image.

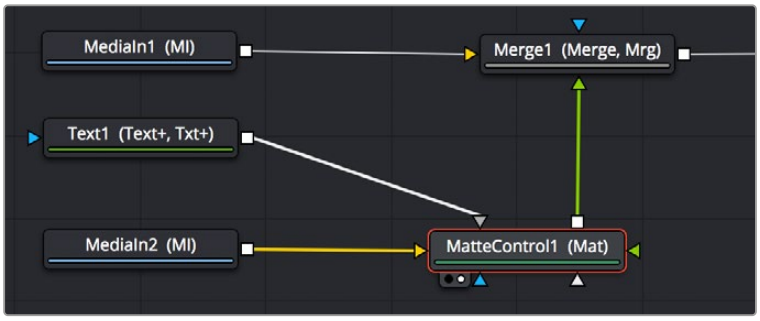
Keep in mind that it's easy to accidentally connect to the wrong input. Because inputs rearrange themselves depending on what's connected and where the node is positioned (and, frankly, the colors can be hard to keep track of when you're first learning), it's key to make sure that you always check the tooltips associated with the input you're dragging a connection over to make sure that you're really connecting to the correct one. If you don't, the effect won't work, and if your effect isn't working, the first thing you should always check is whether you've connected the proper inputs.

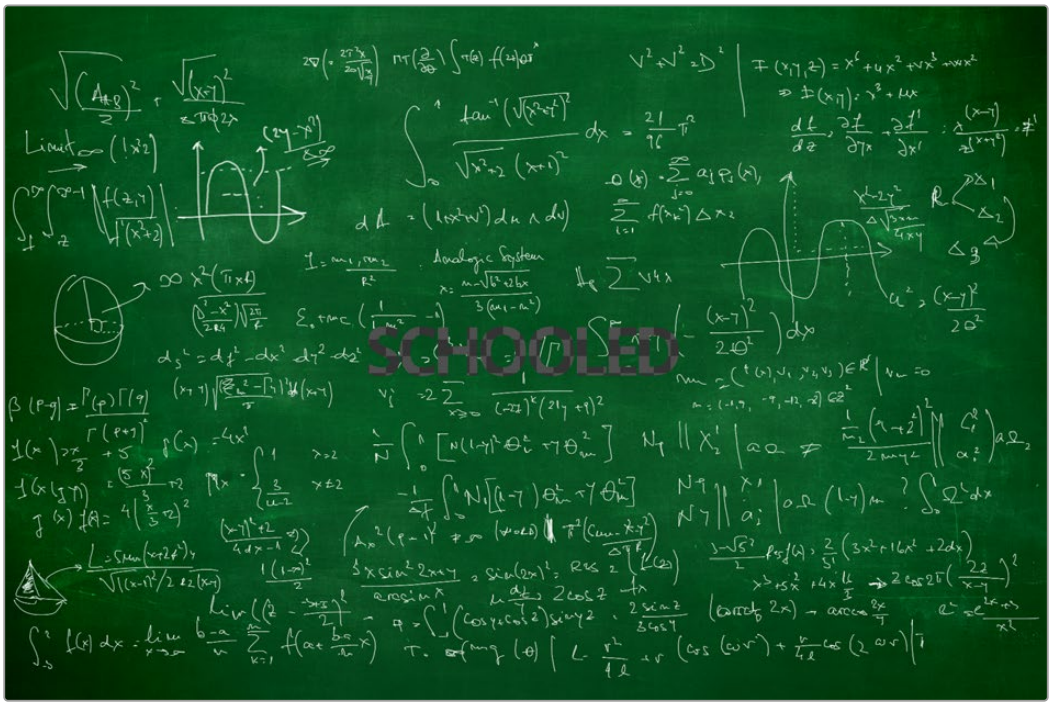


Option-dragging a node connection to drop onto another node exposes a node input menu

One alternate method of connecting nodes together is to hold down the Option key while dragging a connection from one node's output and dropping it onto the body of another node. This opens a pop-up menu from which you can choose the specific input you want to connect to, by name. Note that the menu only appears after you've dropped the connection on the node and released your pointing device's button.

Once the Text1 node is properly connected to the MatteControl node's Garbage Matte input, a text-shaped area of transparency is displayed for the graphic if you load the MatteControl node into the viewer.





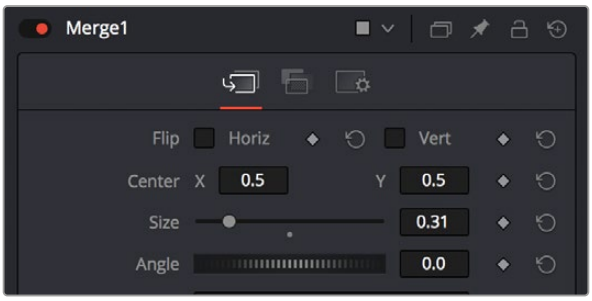
Connecting the Text node to the MatteControl node's Garbage Matte input (previous page), and the resulting hole punched in the image (above)

Customizing Matte Control Nodes

You can use the Inspector to change some parameters to get the result you want. In the Inspector controls for the MatteControl node, revealing the Garbage Matte controls exposes parameters for modifying how the Garbage Matte input is applied to the image. For example, you can choose to have the text mask filled with the image instead of cutting a hole in the image.

Using Transform Controls in the Merge Node

The Merge node includes a set of transform parameters in the Inspector that specifically affect the foreground input's image. This makes it quick and easy to adjust a foreground image to match the background without requiring another node.



The Merge node transform controls that affect the foreground input's image

NOTE: When connecting two images of different sizes to a Merge node, the resolution of the background image defines the output resolution of that node. Keep that in mind when you run into resolution issues.

You can use the Size slider to resize the foreground.



The final composite

Building a Simple Green-Screen Composite

Taking another step forward in compositing, the next example shows how you can equate a multilayered Timeline like the one in DaVinci Resolve's Edit page to nodes in Fusion's node tree. We'll use DaVinci Resolve, but understanding how layers map to nodes can be helpful for anyone new to dealing with a node-based interface. In this example, we'll create a simple composite using a green-screen key and two other layers to create a news story.

Mapping Timeline Layers to Nodes in Fusion

This composite involves three layers in a Timeline. The Timeline consists of a background graphic on video track 1, a green-screen clip on video track 2, and a foreground graphic on video track 3.

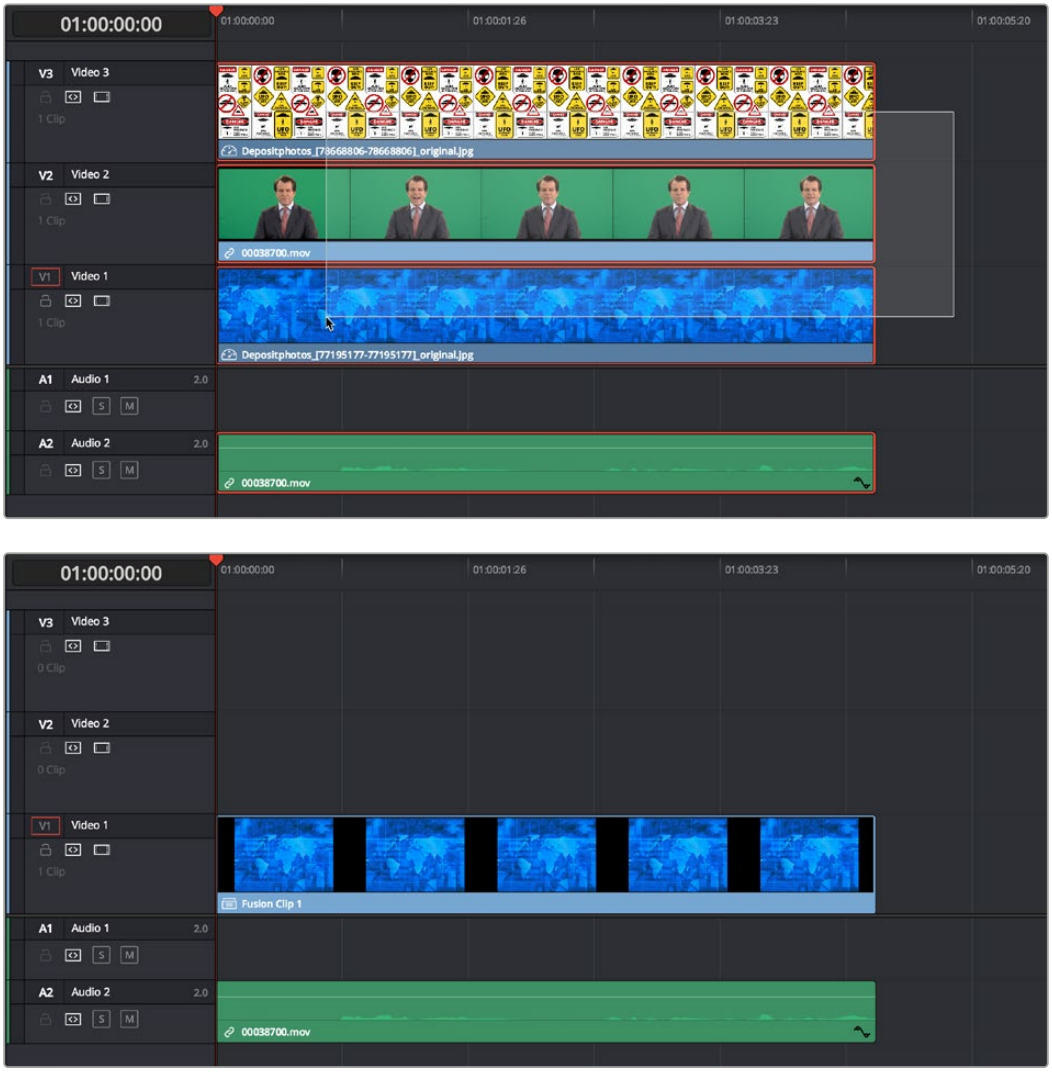


Background on video track 1 (top left), green-screen clip on video track 2 (bottom), and graphic file on video track 3 (top right)

Implied in a timeline-based system is that higher numbered video tracks appear as the more forward, or frontmost, element in the viewer. Video track 1 is the background to all other video tracks. Video track 3 is in the foreground to both video track 1 and video track 2.

TIP: If using DaVinci Resolve, you can bring all three layers from the Edit page into Fusion by creating a Fusion clip.

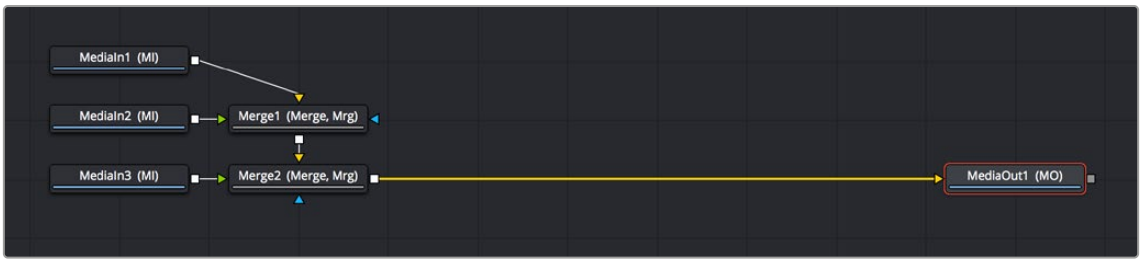
For more information on creating Fusion Clips, *see Chapter 3, "Getting Clips into Fusion."* in the Fusion Reference Manual.



A stack of clips to use in a composite (previous page), and turning that stack into a Fusion clip in DaVinci Resolve's Edit page (above)

In Fusion, each video clip is represented by a MediaIn in the Fusion page or a Loader in Fusion Studio.

In our example below, the MediaIn2 is video track 2, and MediaIn 1 is video track 1. These two elements are composited using a Merge node (foreground over background, respectively). The composite of those two elements becomes the output of the first Merge node, which becomes the background to a second Merge. There is no loss of quality or precomposing when you chain Merges together. MediaIn3 represents video track 3 and is the final foreground in the node tree since it is the topmost layer.

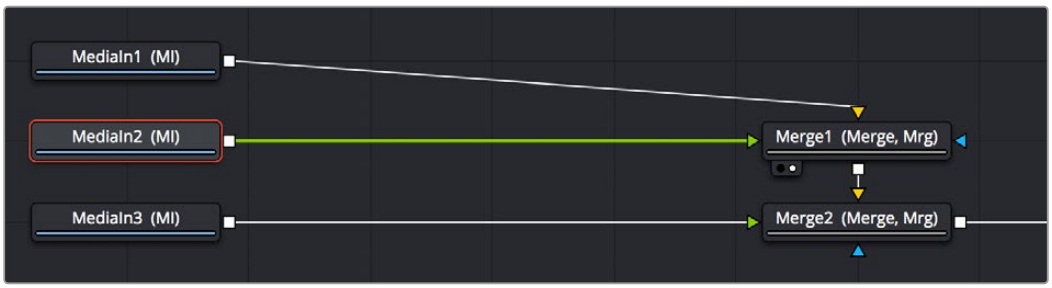


The initial node tree of the three clips we turned into a Fusion clip

With this node tree assembled to mimic the video layers, we can focus the rest of this example on adding the nodes we'll need to each branch of this tree to create the green-screen composite.

Pulling a Green-Screen Key Using the Delta Keyer

To make this composite, you need to create transparency behind the newscaster. When working in a node tree, you must become accustomed to rearranging existing nodes to make room for new ones. You'll often move nodes off to the side or up above to make room for the additional nodes.



Creating space after the MediaIn nodes and selecting the second one in preparation for adding a node

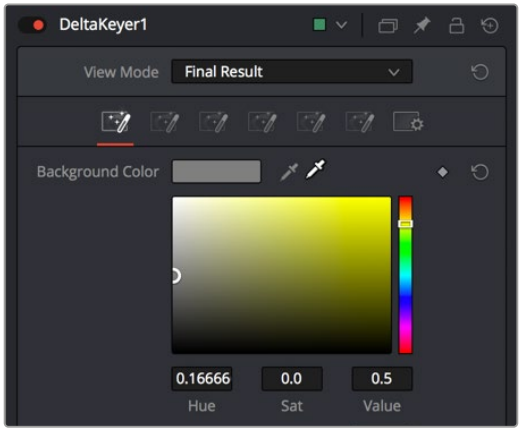
The DeltaKeyer node is the main tool used for green-screen keying. It attaches to the output of the node that represents the green screen—in our example, that is the MediaIn2 node. With the MediaIn2 selected, pressing Shift-Space opens the Select Tool dialog where you can search for and insert any node. Below we have added the DeltaKeyer after the MediaIn2 node but prior to being merged with the background.



Adding a DeltaKeyer node inline after the MediaIn2 node

The DeltaKeyer node is a sophisticated keyer that is capable of impressive results by combining different kinds of mattes and a clean-plate layer, but it can also be used very simply if the background that needs to be keyed is well lit. And once the DeltaKeyer creates a key, it embeds the resulting alpha channel in its output, so in this simple case, it's the only node we need to add. It's also worth noting that, although we're using the DeltaKeyer to key a green screen, it's not limited to keying green or blue only; the DeltaKeyer can create impressive keys on any color in your image.

With the DeltaKeyer selected, we'll use the Inspector controls to pull our key by quickly sampling the shade of green from the background of the image. To sample the green-screen color, drag the Eyedropper from the Inspector over the screen color in the viewer.



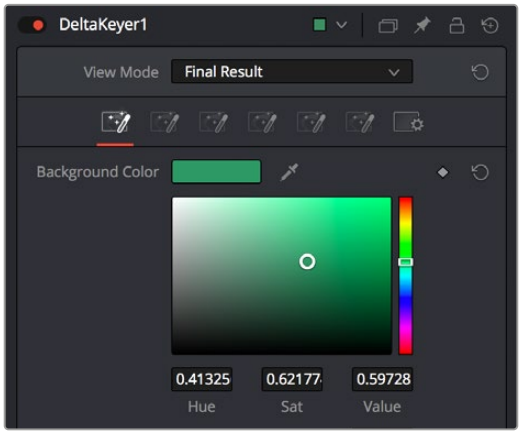
Dragging the Eyedropper to the viewer samples the screen color

As you drag in the viewer, an analysis of the color picked up by the location of the Eyedropper appears within a floating tooltip, giving some guidance as to which color you're really picking. Meanwhile, if viewing the Merge in a second viewer, we get an immediate preview of the transparency and the image we've connected to the background.



The original image (left), and after sampling the green screen using the Eyedropper from the Inspector (right)

When we're happy with the preview, releasing the pointer button samples the color, and the Inspector controls update to display the value we've chosen.



The DeltaKeyer Inspector updates with the sampled color



Loading the DeltaKeyer into the viewer and clicking the Color button to view the alpha channel being produced

No matter how good the composite may look, once you've selected the screen color to pull a key, you need to load the DeltaKeyer node into the viewer itself. This allows you to evaluate the quality or density of the alpha channel created by the key. Above the viewer, click the Color button in the viewer toolbar, or click in the viewer and press C to switch the viewer between the RGB color channels of the image and the alpha channel.

Black in a matte represents the transparent areas, while white represents the opaque areas. Gray areas represent semi-transparency. Unless you are dealing with glass, smoke, or fog, most mattes should be pure white and pure black with no gray gray areas. If a close examination of the alpha channel reveals some fringing in the white foreground of the mask, the DeltaKeyer has integrated controls for post-processing of the key and refining the matte. Following is a quick checklist of the primary adjustments to make.

After making the screen selection with the Eyedropper, try the following adjustments to improve the key.

- Adjust the Gain slider to boost the screen color, making it more transparent. This can adversely affect the foreground transparency, so adjust with care.
- Adjust the Balance slider to tint the foreground between the two non-screen colors. For a green screen, this pushes the foreground more toward red or blue, shifting the transparency in the foreground.

Clicking the third of the seven tabs of controls in the DeltaKeyer Inspector opens up a variety of controls for manipulating the matte.

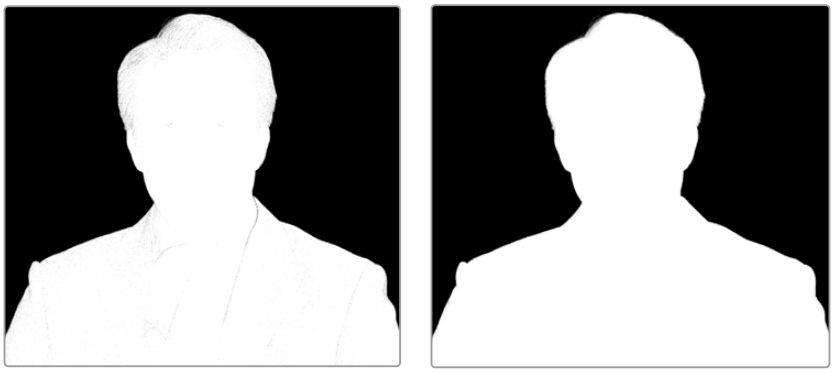
Initial adjustments in the matte tab may include the following parameters:

- Adjust the lower and upper thresholds to increase the density in black and white areas.
- Very subtly adjust the Clean Foreground and Clean Background sliders to fill small holes in the black and white matte. The more you increase these parameters, the more harsh the edges of your matte become.



Adjusting the Clean Foreground slider in the Matte panel of the DeltaKeyer controls

In this case, raising the Clean Foreground slider a bit eliminates the inner fringing we don't want, without noticeably compromising the edges of the key.



The original key (left), and the key after using the Clean Foreground slider (right)

With this accomplished, we're happy with the key, so we load the Merge1 node back into the viewer, and press C to set the Color control of the viewer back to RGB. We can see the graphic in the background, but right now it's too small to cover the whole frame, so we need to make another adjustment.



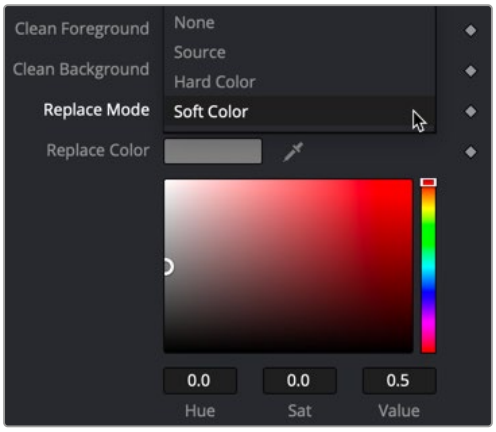
The final key is good, but now we need to work on the background

Dealing with Spill

The DeltaKeyer can handle any screen color that may bounce off the background and land on your subject. This bounce screen light is called spill. However, instead of using one node to handle the matte creation and the spill correction, consider disabling this in the DeltaKeyer and remove the spill using a separate color correction node.

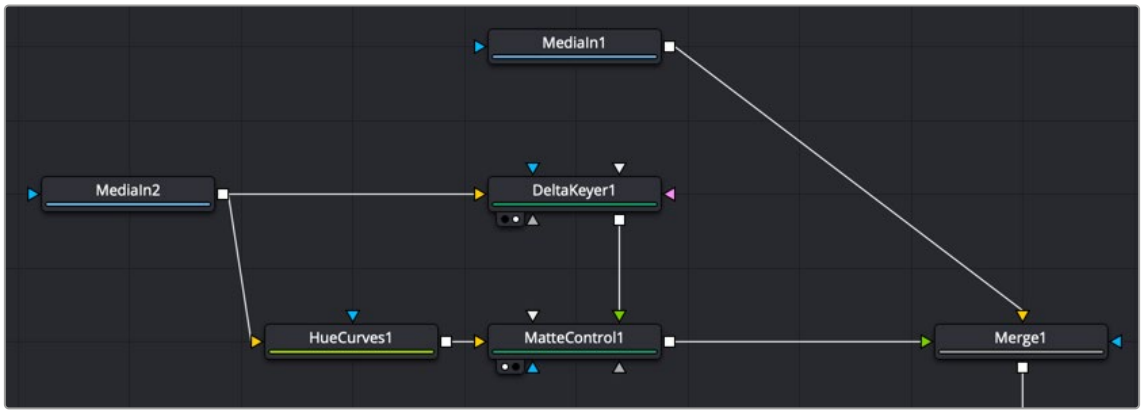
To disable spill suppression in the DeltaKeyer, do the following.

- Select the matte tab in the Inspector.
- From the Replace Mode drop-down menu, choose Source. This uses the color from the original source image, effectively disabling any spill correction in the DeltaKeyer.



The Replace Mode menu determines how spill suppression is handled in the DeltaKeyer

Spill can now be handled using a color correction node placed directly after the DeltaKeyer or branched from the original MediaIn or Loader node and combined with a MatteControl.



Branching the original image with one branch for the DeltaKeyer and a second for color correction

Masking a Graphic

Next, it's time to work on the top video track: the news graphic that will appear to the left of the newscaster. The graphic we will use is actually a sheet of different logos, so we need to cut one out using a mask and position it into place.



A graphic of multiple logos that must be cropped down to isolate just one

The easiest way to crop a MediaIn or Loader node is to add one of the mask shapes from the toolbar directly to it. Selecting the MediaIn or Loader node and clicking the Rectangle mask from the toolbar will crop, or mask off, the graphic.



Masking the logo using a Rectangle mask connected directly to a Merge node

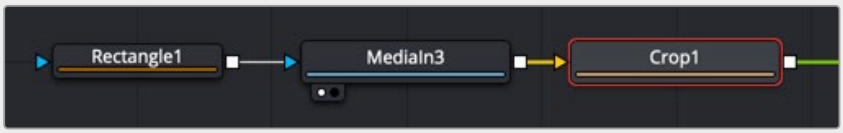
Now, all we need to do is to use the onscreen controls of the Rectangle mask to crop the area we want to use, dragging the position of the mask using the center handle, and resizing it by dragging the top/bottom and left/right handles of the outer border.

As an extra bonus, you can take care of rounded corners when masking a graphic by using the Corner Radius slider in the Inspector controls for the Rectangle mask to add the same kind of rounding.



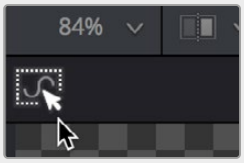
Moving and resizing the mask to fit our logo, and rounding the edges using the Corner Radius Inspector control

For a simple over-the-shoulder graphic, masking the image may be all you need to do, but masking an image does not change the actual dimensions of the graphic. It only changes the area you see. So, accurately positioning the graphic based on the center of the composite becomes more difficult, and any type of match moving would give incorrect results because the graphic has a different resolution than the background. To fix this resolution mismatch, you can place a Crop node after the MediaIn to change the actual dimensions of the graphic layer.



Adding a Crop node after the masked MediaIn to center the cropped logo on the background

With the Crop node selected, the viewer toolbar includes a Crop tool.



Selecting the crop tool in the viewer toolbar

You can crop the image by dragging a bounding box around it. Unlike a mask which creates a small window you view the image through, a crop effectively changes the resolution of the graphic to the crop bounding box size.



Dragging a bounding box using the Crop tool (left), and the cropped logo now centered on the frame (right)

NOTE: The Resize, Letterbox, and Scale nodes also change the resolution of an image.

At this point, we're all set to move the logo into place. Because the logo is the foreground input to a Merge, you can select the Merge2 node, load it into the viewer, and use the built-in Center X and Y controls or the oncreen controls to place the logo where you want it and make it a suitable size.



Placing the logo using the foreground input transform controls of the Merge2 node

Chapter 20

Rotoscoping with Masks

This chapter covers how to use masks to rotoscope, one of the most common tasks in compositing.

Contents

Introduction to Masks and Polylines 505
Mask Nodes 505
Polyline Types 507
Converting Polylines from One Type to Another 508
How to Use Masks with Other Nodes 509
Attaching Masks to an Image for Rotoscoping 510
Combining Multiple Masks 511
Mask Inputs on Other Nodes 512
Creating and Editing Polylines In-Depth 515
The Polyline Toolbar 515
Selecting a Specific Polyline 515
Polyline Creation Modes 515
Protection Modes 517
Closing Polylines 517
Selecting and Adjusting Polylines 518
Polyline Points Selection 518
Moving Polyline Points 518
Smoothing a Polyline Segment 518
Linearizing a Polyline Segment 519
Transforming Individual or Multiple Points 519
Deleting Selected Points 519

Editing Bézier Handles 520
Point Editor 520
Reduce Points 521
Shape Box 521
Showing and Hiding Onscreen Polyline Controls 522
Stop Rendering 523
Roto Assist 523
Creating Softness Using Double Polylines 523
Converting a Single Polyline to a Double Polyline 524
Adding Softness to a Segment 525
Adding Additional Points to the Shape 525
Locking/Unlocking Point Pairs 525
Animating Polyline Masks 526
Removing Animation from a Polyline Mask 526
Adding and Removing Points from an Animated Mask 526
Publishing Specific Control Points 526

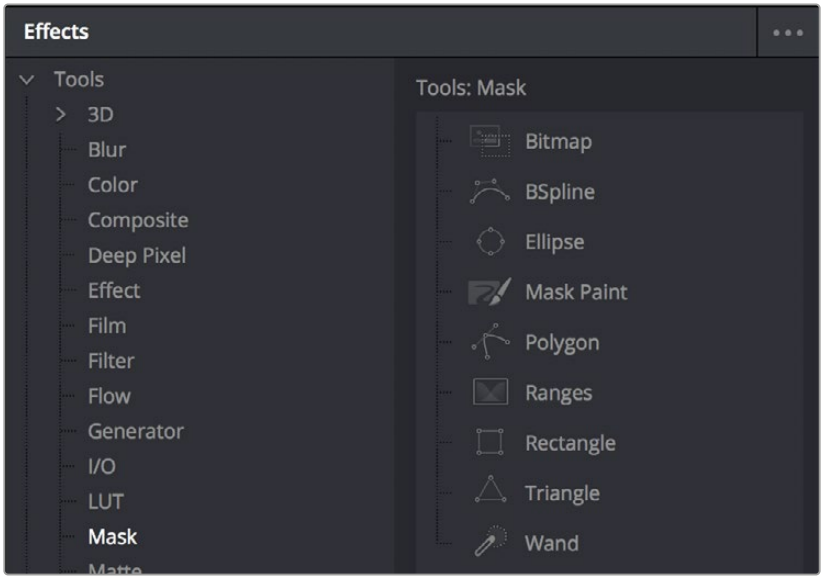
Introduction to Masks and Polylines

Polylines are splines that are used whenever a control is animated with a motion path or when a node's effect is masked with a drawn shape. They are also used in the Paint and Grid Warp nodes. In a more basic form, polylines are used to control the animation in the Spline Editor. Since these splines are used for just about everything, they are extremely flexible, with a considerable amount of controls, modes, and options. This chapter offers an overview of polylines and their operation, with specific information on how to use them for masks.

Mask Nodes

Mask nodes create an image that is used to define transparency in another image. Unlike other image creation nodes in Fusion, mask nodes create a single channel image rather than a full RGBA image. The most used mask tool, the Polygon mask tool, is located in the toolbar.

> For more information on these mask tools, *see Chapter 48, "Mask Nodes,"* in the Fusion Reference Manual.



The available nodes in the Mask bin of the Effects Library

Polygon Mask

Polygon masks are user-created Bézier shapes. This is the most common type of polyline and the basic workhorse of rotoscoping. Polygon mask tools are automatically set to animate as soon as you add them to the Node Editor.

B-Spline Masks

B-Spline masks are user-created shapes made with polylines that are drawn using the B-Splines. They behave identically to polyline shapes when linear, but when smoothed the control points influence the shape through tension and weight. This generally produces smoother shapes while requiring fewer control points. B-Spline mask tools are automatically set to animate as soon as you add them to the Node Editor.

Bitmap Masks

The Bitmap mask allows images from the Node Editor to act as masks for nodes and effects. Bitmap masks can be based on values from any of the color, alpha, hue, saturation, luminance, and the auxiliary coverage channels of the image. The mask can also be created from the Object or Material ID channels contained in certain 3D-rendered image formats.

Mask Paint

Mask Paint allows a mask to be painted using Fusion's built-in vector paint nodes.

Wand Mask

A Wand mask provides a crosshair that can be positioned in the image. The color of the pixel under the crosshair is used to create a mask, where every contiguous pixel of a similar color is also included in the mask. This type of mask is ideal for isolating color adjustments.

Ellipse, Rectangle, and Triangle Masks

These are primitive shape masks.

> For more information on these mask tools, *see Chapter 48, "Mask Nodes,"* in the Fusion Reference Manual.

Ranges Mask

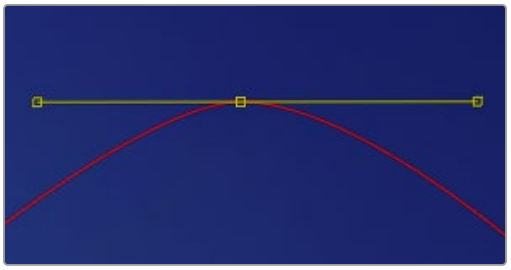
Similar to the Bitmap mask, the Ranges mask allows images from the Node Editor to act as masks for nodes and effects. Instead of creating a simple luminance-based mask from a given channel, Ranges allows spline-based selection of low, mid, and high ranges, similar to to the Color Corrector node.

Polyline Types

You can draw polylines using either B-Spline or Bézier spline types. Which you choose depends on the shape you want to make and your comfort with each spline style.

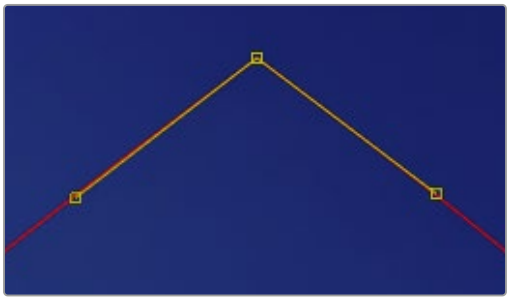
Bézier Polylines

Bézier polylines are shapes composed of control points and handles. Several points together are used to form the overall shape of a polyline.



Bézier control point with direction handles extended to create a smooth curve

Each control point has a pair of handles used to define the exact shape of the polyline segments passing through each control point. Adjusting the angle or length of the direction handles will affect whether that segment of the polyline is smooth or linear.

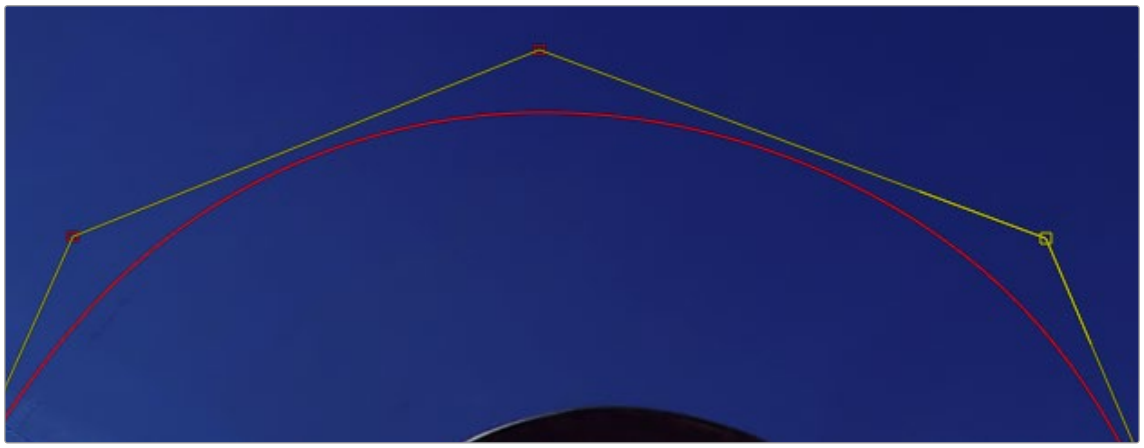


Bézier control point with direction handles aligned to create a linear segment

If you're familiar with applications such as Adobe Photoshop or Illustrator, you'll already be familiar with many of the basic concepts of editing Bézier polylines.

B-Spline Polylines

A B-Spline polyline is similar to a Bézier spline; however, these polylines excel at creating smooth shapes. Instead of using a control point and direction handles for smoothness, the B-Spline polyline uses points without direction handles to define a bounding box for the shape. The smoothness of the polyline is determined by the tension of the point, which can be adjusted as needed.



B-Splines excel at creating smooth curves

Converting Polylines from One Type to Another

Just because you created a shape using a B-Spline or polyline, that doesn't mean you're stuck with the controls you started with. You can convert any shape from B-Spline to Bézier, or Bézier to B-Spline, as needed.

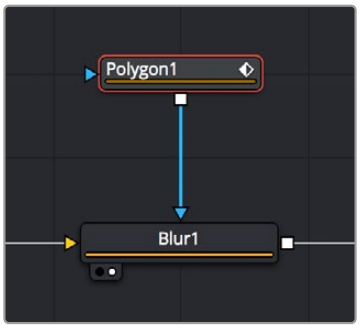
To switch a shape between Polyline and B-Spline controls:

Right-click a shape in the viewer and choose Convert Bézier Spline to B-Spline or Convert B-Spline to Bézier from the spline's contextual menu (only the appropriate option will be displayed).

When converting from one type to another, the original shape is preserved. The new polyline generally has twice as many control points as the original shape to ensure the minimum change to the shape. While animation is also preserved, this conversion process will not always yield perfect results. It's a good idea to review the animation after you convert spline types.

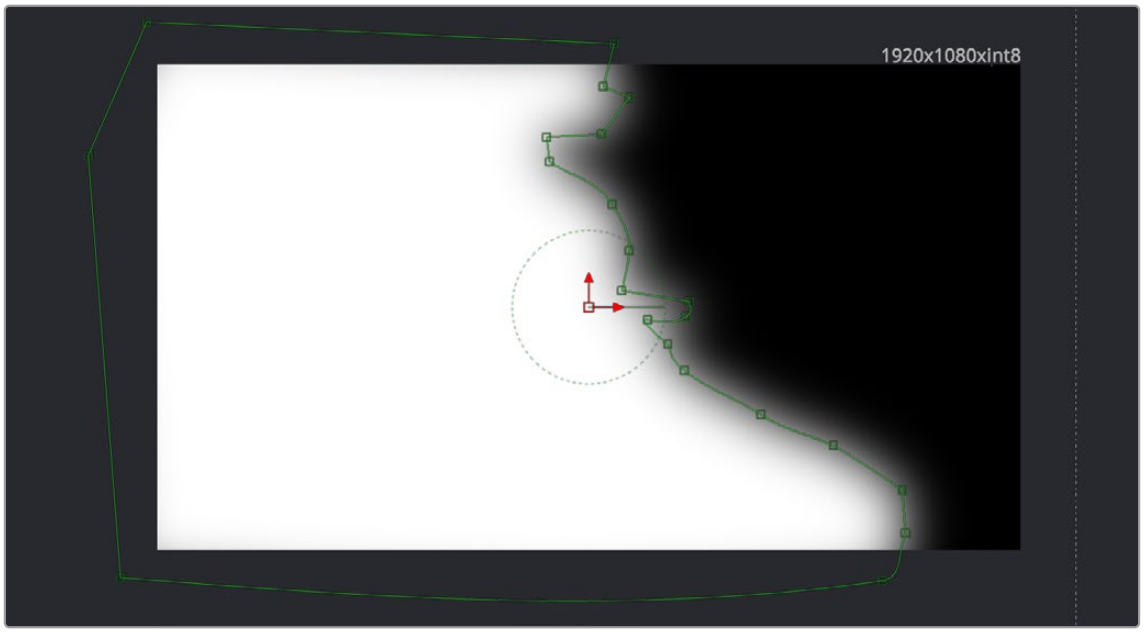
How to Use Masks with Other Nodes

Typically, a node applies its effect to every pixel of an image. However, many nodes have mask inputs that can be used to limit the effect that node has on the image.



A Blur node with a Polygon node masking its effect

Masks are single-channel images that can be used to define which regions of an image you want to affect. Masks can be created using primitive shapes (such as circles and rectangles), complex polyline shapes that are useful for rotoscoping, or by extracting channels from another image.



A Polygon node's mask seen in the viewer

Each mask node is capable of creating a single shape. However, Mask nodes are designed to be added one after the other, so you can combine multiple masks of different kinds to create complex shapes. For example, two masks can be subtracted from a third mask to cut holes into the resulting mask channel.

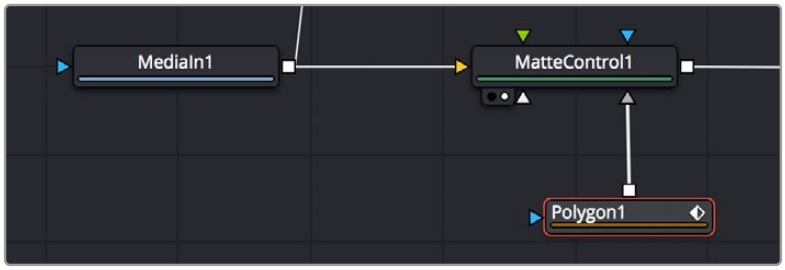
Fusion offers several different ways you can use masks to accomplish different tasks. You can attach Mask nodes after other nodes in which you want to create transparency, or you can attach Mask nodes directly to the specialized inputs of other nodes to limit or create different kinds of effects.

Attaching Masks to an Image for Rotoscoping

There are two ways you'll typically attach a Mask node, such as a Polygon node, so that it adds an alpha channel to an image for compositing later in the node tree.

Using a MatteControl Node

The MatteControl node is the main node used for combining masks in different ways and inserting the result into an image stream. The MatteControl node is attached downstream of the node outputting the image you want to rotoscope. You'll typically attach a Polygon or B-Spline node to the Garbage Matte input of the MatteControl node to use the spline as an alpha channel.



Feeding a Polygon node to a MatteControl node to perform rotoscoping

To use this setup, you'll load the MatteControl node into the viewer and select the Polygon node to expose its controls so you can draw and modify a spline while viewing the image you're rotoscoping. The MatteControl node's Garbage Matte > Invert checkbox lets you choose which part of the image becomes transparent.

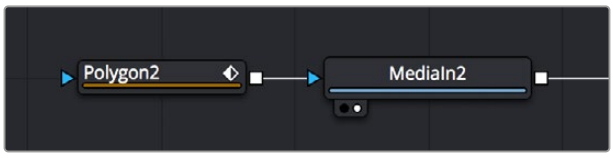
Connecting a Mask to a MediaIn or Loader Node's Input

This method is a bit simpler but requires you to know that you can view one node while adjusting another node, even if that other node is disconnected. If you add an unattached Mask node such as a Polygon or B-Spline node, and then place a MediaIn or Loader node directly into the viewer while selecting the Mask node, you can draw a spline to rotoscope the image.



Rotoscoping a MediaIn node using a disconnected Polygon node

When you're finished rotoscoping, you simply connect the Polygon node's output to the Loader node's input, and an alpha channel is automatically added to that node.

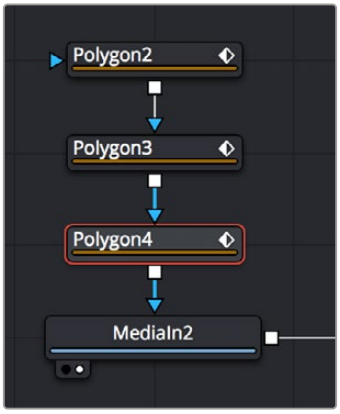


Connecting a Polygon node to a MediaIn node to use a spline as an alpha channel

TIP: If you connect a Mask node to a MediaIn or Loader node's effect input without any shapes drawn, that mask outputs full transparency, so the immediate result is that the image output by the MediaIn or Loader node becomes completely blank. This is why when you want to rotoscope by connecting a mask to the input of a MediaIn or Loader node, you need to work within a disconnected Mask node first. Once the shape you're drawing has been closed, connect the Mask node to the MediaIn or Loader's input, and you're good to go.

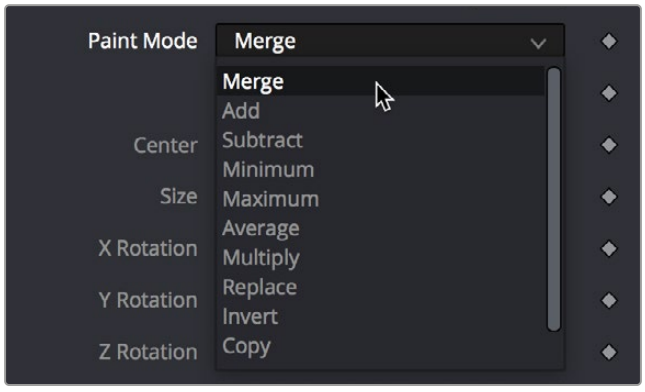
Combining Multiple Masks

Masks are designed to be added one after the other, with each Mask node acting as an additional layer of masking.



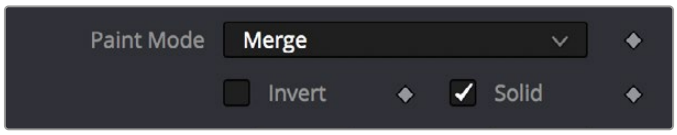
Combining multiple Polygon nodes one after the other in the node tree

When a Mask node's input is attached to another mask, a Paint Mode drop-down menu appears, which allows you to choose how you want to combine the two masks.



The Paint Mode parμameter in the Polygon node Inspector parameters

The default option is Merge, but you can also choose Subtract, Minimum, Maximum, Multiply, or any other operation that will give you the mask boolean interaction you need. Additionally, a pair of Invert and Solid checkboxes let you further customize how to combine the current mask with the one before it.



The Invert and Solid options

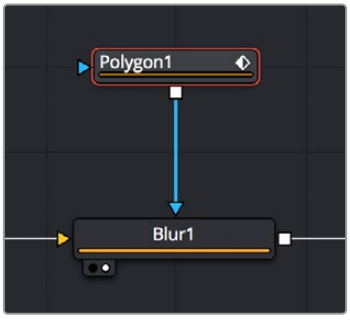
Mask Inputs on Other Nodes

Masks can be used for a variety of reasons, so there are several categories of mask inputs that different nodes include to accommodate these different uses. Incidentally, in most cases you can connect either masks or mattes to a mask input to take advantage of that input's functionality.

TIP: If you select a node with an empty effect mask input, adding a Mask node automatically connects to the open effect mask input.

Effects Mask Inputs

Almost every node in Fusion has an Effect mask input (colored blue), which lets you choose which parts of the image will or will not be affected by that node.

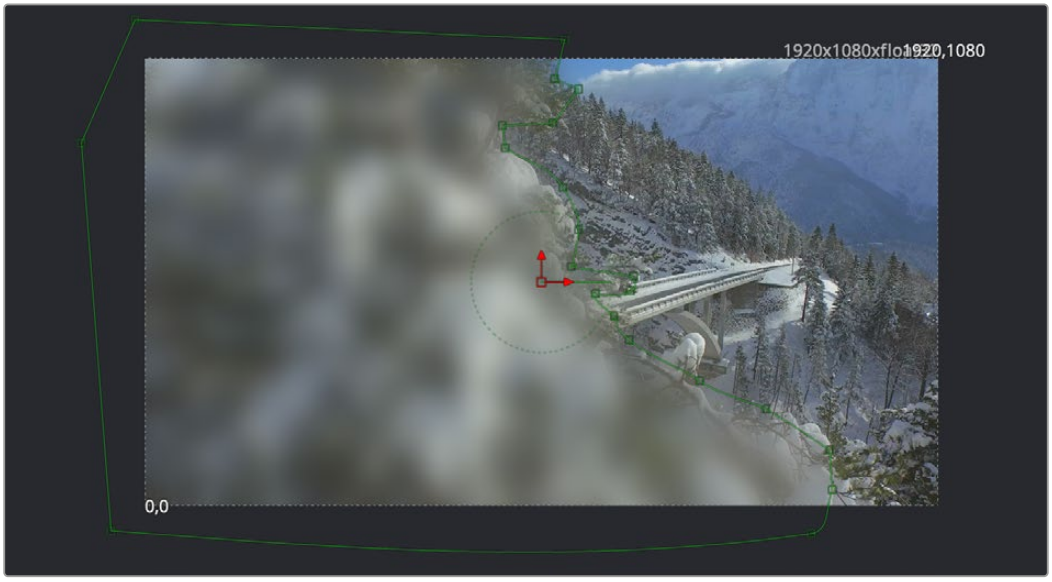


A Blur node with a Polygon node masking its effect

While masks (or mattes) are connected via an input, they are actually applied "post effect," which means the node first applies its effect to the entire image, and then the mask is used to limit the result by copying over unaffected image data from the input.

Although many nodes support effects masking, there are a few where this type of mask does not apply—notably Savers, Time nodes, and Resize, Scale, and Crop nodes.

TIP: Effects masks define the domain of definition (DoD) for that effect, making it more efficient.



A Polygon node used as a mask to limit the Blur node's effect

Pre-Masking Inputs

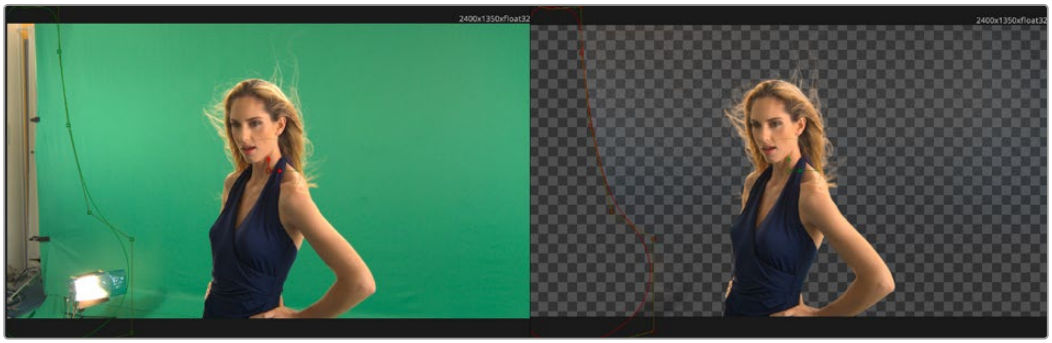
Unlike effect masks, a pre-mask input (the name of which is usually specific to each node using them) is used by the node before the effect is applied. This usually causes the node to render more quickly and to produce a more realistic result. In the case of the Highlight and the Glow nodes, a pre-mask restricts the effect to certain areas of the image but allows the result of that effect to extend beyond the limits of the mask.

The advantage to pre-masking is that the behavior of glows and highlights in the real world can be more closely mimicked. For example, if an actor is filmed in front of a bright light, the light will cause a glow in the camera lens. Because the glow happens in the lens, the luminance of the actor will be affected even though the source of the glow is only from the light.

In the case of the DVE node, a pre-mask is used to apply a transformation to a selected portion of the image, without affecting portions of the image outside of the mask. This is useful for applying transformations to only a region of the image.

Garbage Matte Inputs

Garbage Matte inputs (usually colored gray) are used to exclude light stands, rigging, and boom microphones that intrude upon masks being pulled via blue-screen and green-screen keys. In the following example, a lighting stand to the left is removed from the image via a B-Spline node's mask connected to the Garbage Matte input of the DeltaKeyer node.



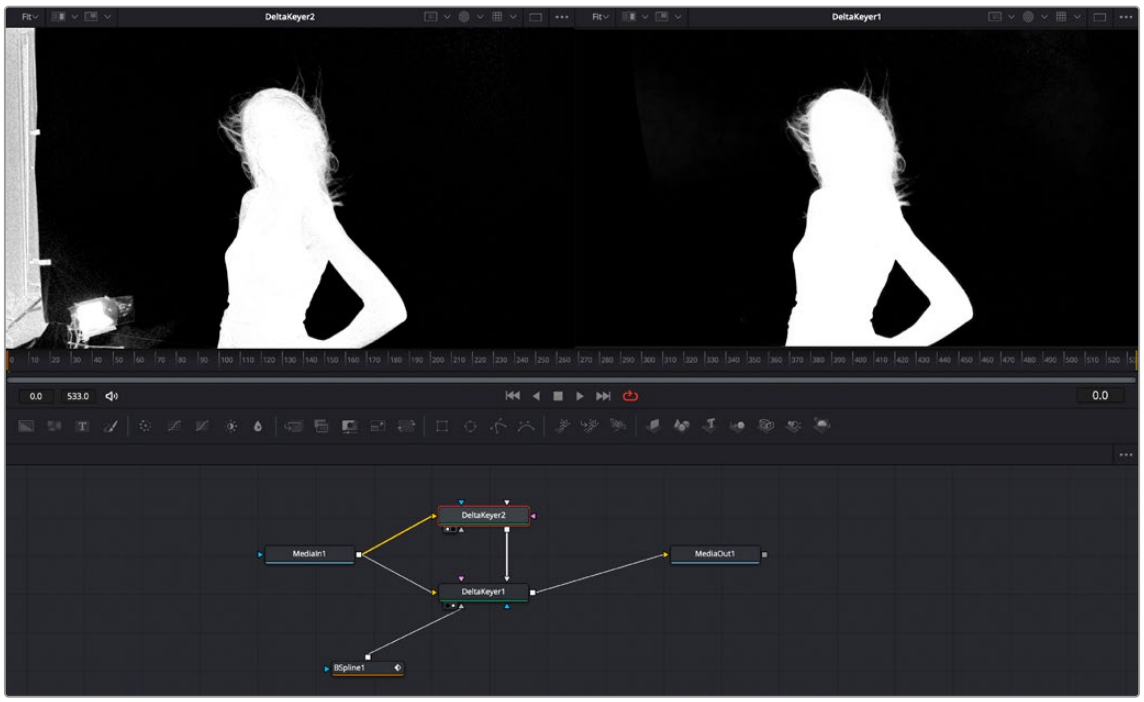
A B-Spline node is connected to the Garbage Matte input of a DeltaKeyer node to eliminate a light stand at the left of frame

TIP: You can quickly add a mask node to the Effect/Solid/Garbage Matte inputs of a keyer node by right-clicking the header bar of that node in the Inspector and choosing whichever mask node you want to use from the Effect Mask, SolidMatte, and GarbageMatte submenus.

You choose whether a garbage matte is applied to a keying node as opaque or transparent in the Inspector for the node to which it's connected.

Solid Matte

Solid Matte inputs (colored white) are intended to fill unwanted holes in a matte, often with a less carefully pulled key producing a dense matte with eroded edges, although you could also use a polygon or mask paint to serve this purpose. In the example below, a gentle key designed to preserve the soft edges of the talent's hair leaves holes in the mask of the woman's face, but using another DeltaKeyer to create a solid matte for the interior of the key that can be eroded to be smaller than the original matte lets you fill the holes while leaving the soft edges alone. This is also sometimes known as a hold-out matte.



Filling in holes in the mask pulled by the DeltaKeyer1 node (left) with another, harder but eroded key in DeltaKeyer2 that's connected to the SolidMatte input of DeltaKeyer1 (right)

Creating and Editing Polylines In-Depth

This section covers the Polygon node's capabilities in depth.

The Polyline Toolbar

Whenever a node that contains one or more polylines is selected, the polyline is shown on all viewers and the Polyline toolbar is displayed along the side of each viewer. The toolbar contains several buttons that make switching polyline modes and options easy to access.



The Polyline toolbar

If you hover the pointer over any of the Polyline toolbar buttons, a tooltip that describes the button's function appears. Clicking on a button will affect the currently active polyline or the selected polyline points, depending on the button.

You can change the size of the toolbar icons, add labels to the buttons, or make other adjustments to the toolbar's appearance in order to make polylines easier to use. All the options can by found by right-clicking on the toolbar and selecting from the options displayed in the contextual menu.

Selecting a Specific Polyline

It is possible to have several polylines in the viewer at once if you select multiple Mask nodes in the Node Editor, so it's important to be able to switch between polylines easily.

To make a polyline active, do one of the following:

- Click one of the polyline's control points or segments.
- Press Tab and Shift-Tab to cycle between available polylines.
- Right-click in the viewer and choose the desired polyline by name from the Controls > Select menu.

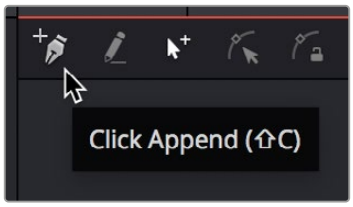
Polyline Creation Modes

There are several different modes available from the toolbar for creating and modifying polylines. The specific mode used when a polyline is first added will depend on whether it is used as a path or a mask.

Each of the modes is described in more detail below.

Click Append

This mode is the default mode for mask creation. It's used to quickly define the rough shape of the mask, after which you switch to Insert and Modify mode to refine the mask further.



The Click Append toolbar button (Shift-C)

To create a mask using the Click Append mode, do the following:

1. Select Click Append from the toolbar or press Shift-C.
2. Click the pointer where you want to start the shape.
3. Move and click the pointer to append a point to the last one.
4. To close the shape, place the mouse pointer over the first point created and click when the pointer changes shape.

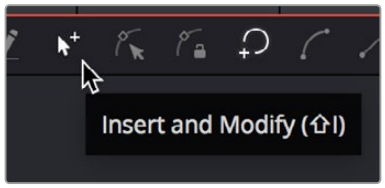
When a shape is closed, the polyline is automatically switched to Insert and Modify mode.

Although the Click Append mode is rarely used with paths, it can be helpful when you know the overall shape of a motion path, but you don't yet know the timing.

> **TIP:** Holding Shift while you draw a mask constrains subsequent points to 45-degree angles relative to the previous point. This can be very helpful when drawing regular geometry.

Insert and Modify

Masks, which are created in Click Append mode, automatically switch to Insert and Modify mode when the mask shape is closed. You can also manually switch to this mode by clicking the Insert and Modify button in the toolbar or using the Shift-I keyboard shortcut. This mode makes it easier to add additional points and refine the shape of the mask. Dragging the control points or direction handles modifies existing points on the polyline.

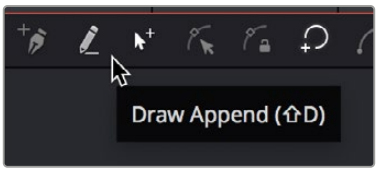


The Insert Modify toolbar button (Shift-I)

Insert and Modify mode is also the default mode for creating motion paths. A new control point is automatically added to the end of the polyline, extending or refining the path, any time a parameter that is animated with a motion path is moved.

Draw Append

The Draw Append mode creates a freehand polyline shape directly on the viewer, like drawing with a pencil or a paintbrush. This mode is ideal to use in conjunction with a tablet and for the creation of garbage mattes and effect masks around complex shapes.



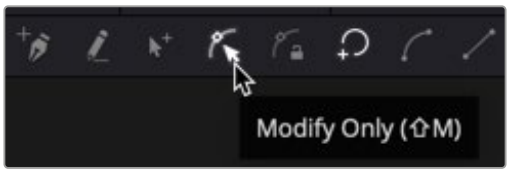
The Draw Append toolbar button (Shift-D)

Protection Modes

In addition to the modes used to create a polyline, two other modes are used to protect the points from further changes after they have been created.

Modify Only

Modify Only mode allows existing points on the polyline to be modified, but new points may not be added to the shape.

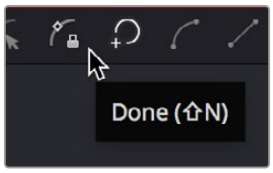


The Modify Only toolbar button (Shift-M)

TIP: Even with Modify Only selected, it is still possible to delete points from a polyline.

Done

The Done mode prohibits the creation of any new points, as well as further modification of any existing points on the polyline.



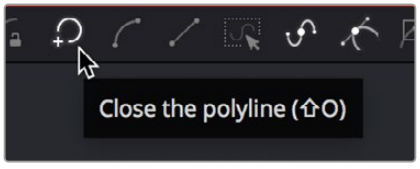
The Done toolbar button (Shift-N)

Closing Polylines

There are several ways to close a polyline, which will connect the last point to the first.

To close a polyline, do one of the following:

- Hover the pointer over the first point created, and then click on the point.
- Press Shift-O on the keyboard.
- Click the Close button on the polyline toolbar.
- Draw a polyline until you are ready to close the shape, and then right-click and choose Polygon:Polyline > Closed.



The Close toolbar button (Shift-O)

All these options are toggles that can also be used to open a closed polygon.

Selecting and Adjusting Polylines

To create the shape you need for a mask or a motion path, you need to know how to manipulate the splines. Fusion provides a number of simple techniques for selecting, moving, and smoothing a spline, but also includes more complex adjustment techniques for scale, skewing, and twisting a spline.

Polyline Points Selection

To select one or more control points on a polyline, do one of the following:

- Click directly on the control points.
- Lasso around the points.

To add or remove points from the current selection, do one of the following:

- Hold the Shift key to select a continuous range of points.
- Hold Command and click each control point you want to add or remove.
- Press Command-A to select all the points on the active polyline.

> **TIP:** Once a control point is selected, you can press Page Down or Page Up on the keyboard to select the next control point in a clockwise or counterclockwise rotation. This can be very helpful when control points are very close to each other.

Moving Polyline Points

The selected polyline points can be moved using either the keyboard or the mouse.

To move selected control points using the pointer, do one of the following:

- Drag on the selected points anywhere in the viewer.
- Hold Shift while dragging to restrict movement to a single axis.
- Hold Option and drag anywhere in the viewer to move the selected control point.

To move selected control points using the keyboard, do one of the following:

- Press the Up or Down Arrow keys on the keyboard to nudge a point up or down in the viewer.
- Hold Command-Up or Down Arrow keys to move in smaller increments.
- Hold Shift-Up or Down Arrow keys to move in larger increments.

Smoothing a Polyline Segment

If you want to shape the polyline and control its slope, you can choose to smooth a spline segment by adjusting the Bézier direction handles.

To smooth the selected points on an active polyline, do one of the following:

- Press Shift-S.
- Click the Smooth button on the Polyline toolbar.
- Choose Smooth from the polyline's contextual menu.



The Smooth button in the toolbar (Shift-S)

Linearizing a Polyline Segment

To make certain that a polyline segment is perfectly straight, that segment must be linearized. A linear segment aligns the Bézier direction handles with the segment and therefore has no curvatures. The segment is always drawn in a straight line between two points on the polyline.

To linearize the selected points on an active polyline, do one of the following:

- Press Shift-L.
- Click the Linear button on the polyline's toolbar.
- Choose Linear from the polyline's contextual menu.



The Linear button in the toolbar (Shift-L)

Transforming Individual or Multiple Points

Select the points to be transformed, and then do one of the following:

- Hold T and drag to twist.
- Hold S and drag to scale.
- Hold X and drag to scale horizontally only.
- Hold Y and drag to scale vertically only.
- Hold O and drag to offset the points perpendicular to the tangent.

The position of the pointer when the transformation begins becomes the center used for the transformation.

Deleting Selected Points

You can delete a selected point or group of points by pressing Delete or Backspace, choosing Delete from the contextual menu, or by clicking the Delete Point button in the toolbar. The shape of the polyline changes to reflect the removal of these points.

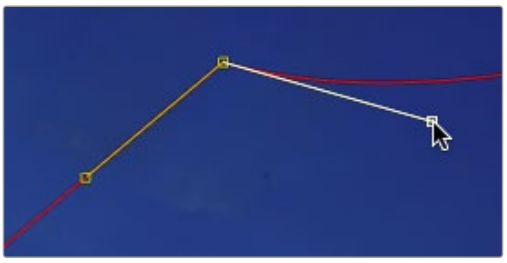
TIP: Deleting all the points in a polyline does not delete the polyline itself. To delete a polyline, you must delete the node or modifier that created the polyline.

Editing Bézier Handles

For Bézier polylines, each control point has two direction handles that adjust the slope of a curve through the control point. These direction handles appear only when the point is selected.

Dragging a direction handle makes adjustments to the curve of the segment that emerges from the control point. The direction handle on the opposing side of the control point will also move to maintain the relationship between these two handles.

To break the relationship between direction handles and adjust one independently, hold Command while dragging a handle. Subsequent changes will maintain the relationship, unless Command is held during each adjustment.

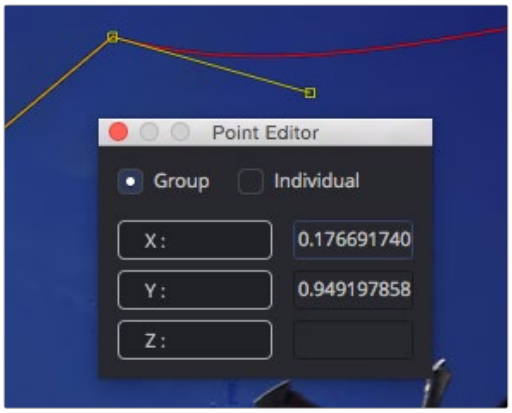


Hold Command to adjust one handle independently

If you want to adjust the length of a handle without changing the angle, hold Shift while moving a direction handle.

Point Editor

The Point Editor dialog can be used to reposition control points using precise X and Y coordinates. Pressing the E key on the keyboard will bring up the Point Editor dialog and allow you to reposition one or more selected control points.

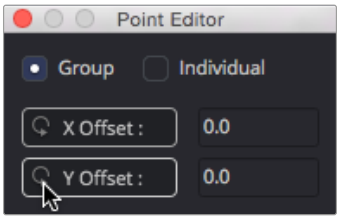


The Point Editor dialog can be used to position control points

The dialog box contains the X- and Y-axis values for that point. Entering new values in those boxes repositions the control point. When multiple control points are selected, all the points move to the same position. This is useful for aligning control points along the X- or Y-axis.

If more than one point is selected, a pair of radio buttons at the top of the dialog box determines whether adjustments are made to all selected points or to just one. If the Individual option is selected, the affected point is displayed in the viewer with a larger box. If the selected point is incorrect, you can use the Next and Previous buttons that appear at the bottom of the dialog to change the selection.

In addition to absolute values for the X- and Y-axis, you can adjust points using relative values from their current position. Clicking once on the label for the axis will change the value to an offset value. The label will change from X to X-offset or from Y to Y-offset.

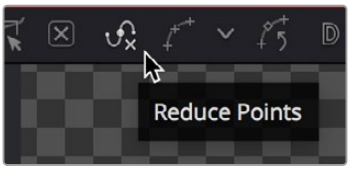


The Point Editor dialog with Offset values

If you are not sure of the exact value, you can also perform mathematical equations in the dialog box. For example, typing 1.0-5 will move the point to 0.5 along the given axis.

Reduce Points

When freehand drawing a polyline or an editable paint stroke, the spline is often created using more control points than you need to efficiently make the shape. If you choose Reduce Points from the polyline's contextual menu or toolbar, a dialog box will open allowing you to decrease the number of points used to create the polyline.



The Reduce Points button in the toolbar

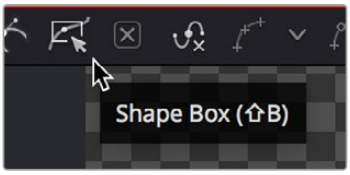
The overall shape will be maintained while eliminating redundant control points from the path. When the value is 100, no points are removed from the spline. As you drag the slider to the left, you reduce the number of points in the path.

Shape Box

If you have a polyline shape or a group of control points you want to scale, stretch, squish, skew, or move, you can use the shape box to easily perform these operations.

To enable the shape box, do one of the following:

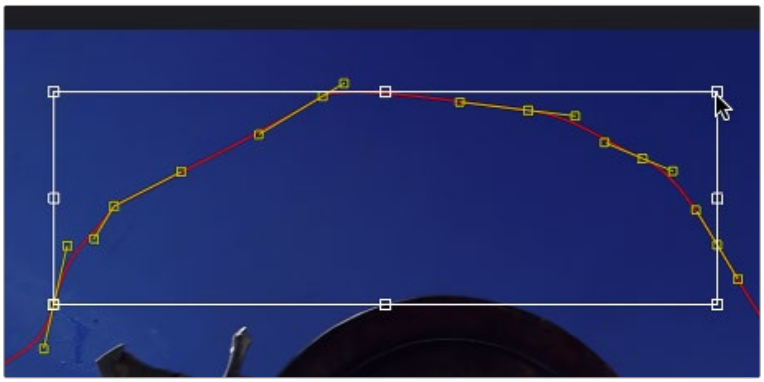
- Click the Shape Box toolbar button.
- Choose Shape Box from the contextual menu.
- Press Shift-B.



The Shape Box button in the Polyline toolbar

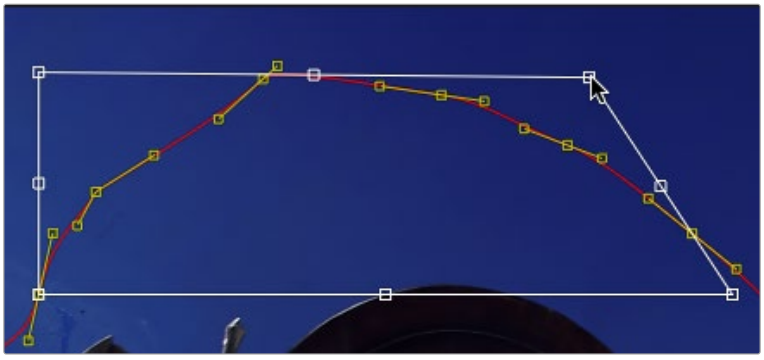
If there are selected points on the polyline when the Shape Box mode is enabled, the shape box is drawn around those points. Otherwise, you can drag the shape box around the area of control points you want to include.

If you want to freely resize the shape box horizontally and vertically, you can drag a corner handle. Dragging a handle on the side of the shape box resizes the polyline along a specific axis.



Dragging a side handle resizes along a specific axis

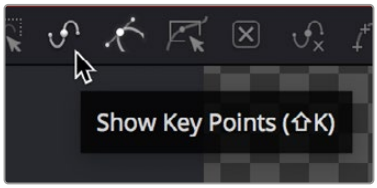
Holding Command while dragging a shape box handle will apply adjustments from the center of the shape box, constraining the transformation to the existing proportions of the shape box. Holding Shift while dragging a corner handle affects only that handle, allowing skewed and non-uniform transformations.



Hold Shift while dragging a corner to perform non-uniform transformations

Showing and Hiding Onscreen Polyline Controls

It is often difficult to identify individual points when they are placed closely together. You can choose to display both points and their direction handles, just points, or just handles. These display mode options are selected using the Show Key Points and Show Handles toolbar buttons, or from the polyline's context menu.



The Show Key Points and Show Handles buttons in the toolbar

You use these options to simplify the screen display when adjusting control points placed closely together and to avoid accidentally modifying controls and handles that are adjacent to the intended target.

Stop Rendering

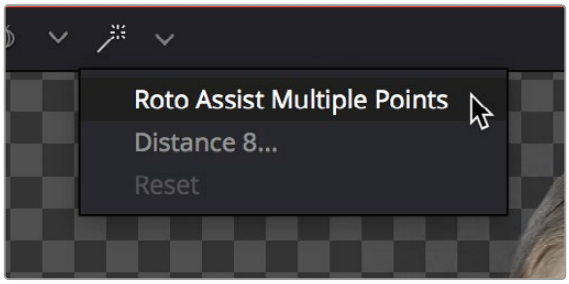
While points along the polyline are being moved, the results are rendered to the viewer to provide constant interactive feedback. Although extremely useful, there are situations where this can be distracting and can slow down performance on a complex effect. To disable this behavior so renders happen only when the points stop moving, you can toggle the Stop Rendering button in the toolbar or select this option from the polyline contextual menu.

Roto Assist

You can enable the Roto Assist button in the toolbar when you begin drawing your shape to have points snap to the closest high-contrast edge as you draw the shape. The points that have snapped to an edge are indicated by a cyan outline.

There are three main Roto Assist options.

- **Multiple Points:** Allows adding multiple points along an entire edge with a single click instead of having to add each point individually.
- **Distance:** Defines the pixel range within which searching for an edge will take place.
- **Reset:** Used for resetting the snap attribute of the snapped points. After resetting, the points will become unavailable for tracking.

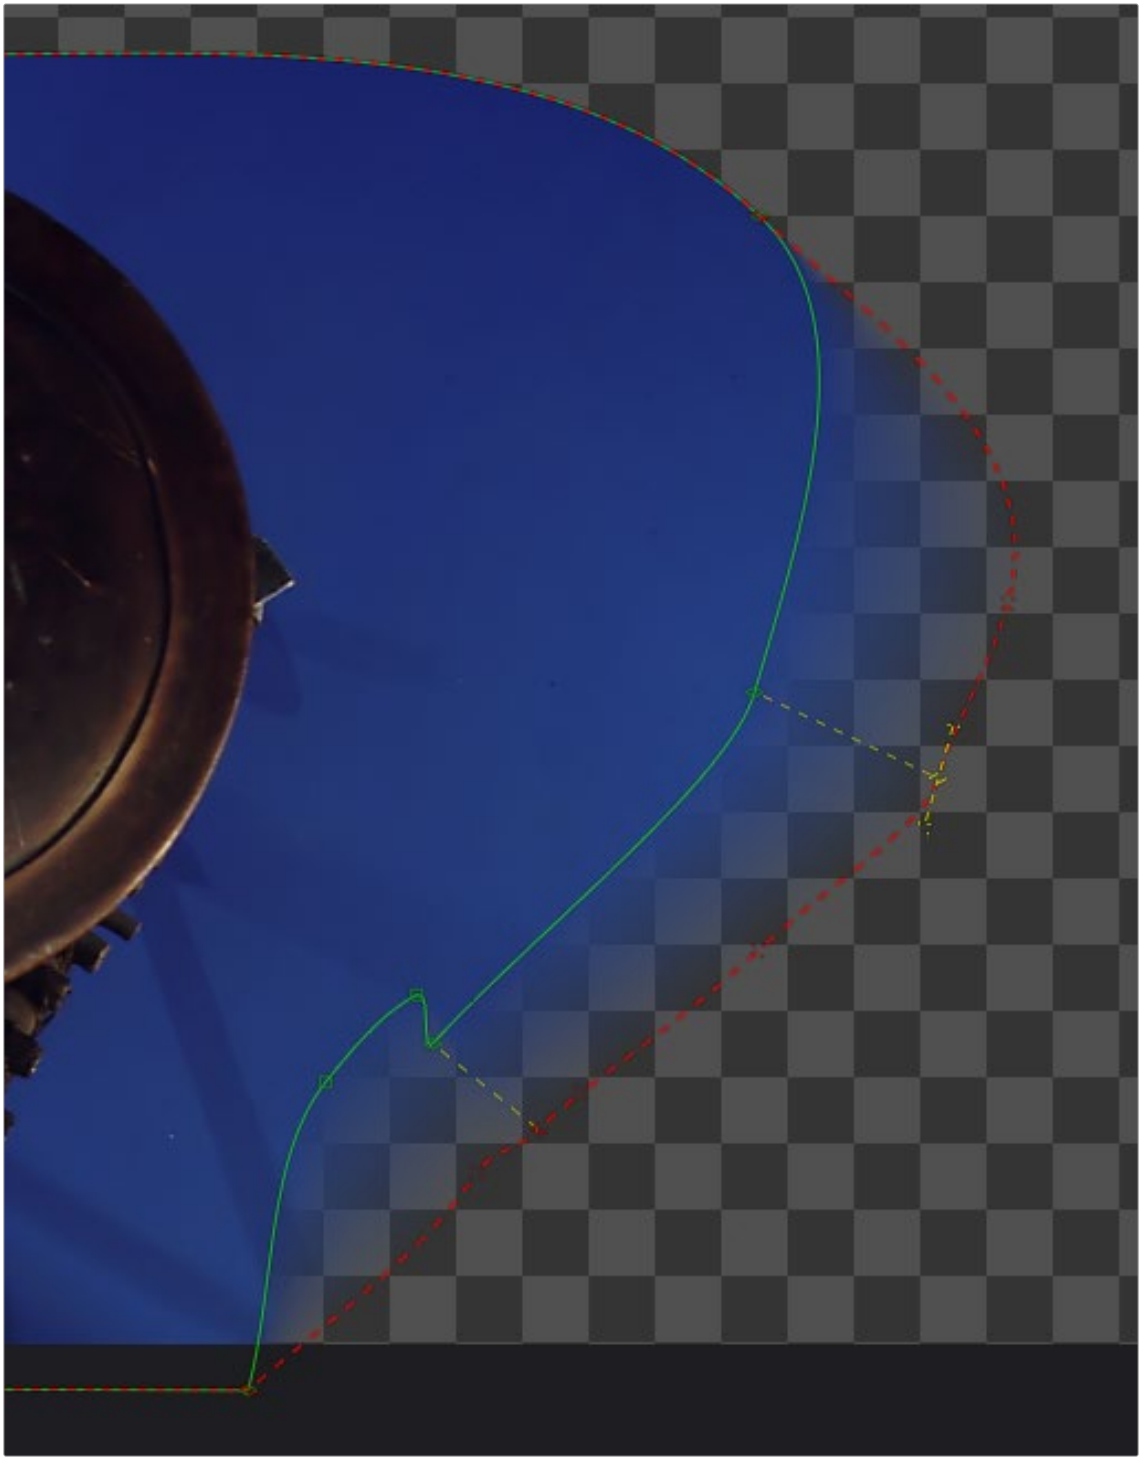


The Roto Assist options in the toolbar

Creating Softness Using Double Polylines

The standard soft edge control available in all Mask nodes softens the entire mask equally. However, there are times, particularly with a lot of motion blur, when softening part of the curve while keeping other portions of the curve sharp is required.

This form of softness is called non-uniform softness, which is accomplished by converting the shape from a single polyline to a double polyline. The double polyline is composed of two shapes: an inner and an outer shape. The inner shape is the original shape from the single polyline, whereas the outer shape is used to determine the spread of the softness. The further the outer shape gets from the inner shape, the softer that segment of the shape becomes.

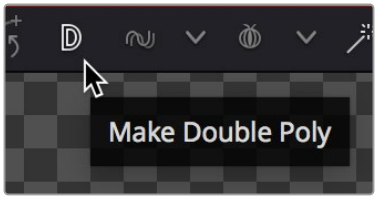


A double polyline uses an inner and outer shape for non-uniform softness

Converting a Single Polyline to a Double Polyline

To convert a mask into a double polyline, click the Double Polyline button in the Polyline toolbar or right-click in the viewer and select Make Outer Polyline from the mask's contextual menu.

The shape will be converted into an inner and an outer polyline spline. Both polylines start with exactly the same shape as the original single polyline. This keeps the mask sharp to start with and allows any animation that may have already been applied to the shape to remain.



Make Double Polyline button

The control points on the outer shape are automatically parented to their matching points on the inner shape. This means that any changes made to the inner shape will also be made to the outer shape. The relationship is one-way; adjustments to the outer shape can be made without affecting the inner shape.

A dashed line drawn between the points indicates the relationship between the points on the inner and outer shapes.

Adding Softness to a Segment

The outer shape is drawn using a green dashed line instead of a solid line to help distinguish it from the inner shape.

To select the outer soft edge shape, do one of the following:

- Use the Tab key to cycle between the onscreen controls until the dashed outline is visible
- Right-click over a spline in the view and choose Controls > Select > Polygon: Outer Polygon.

Once the outer polyline is selected, you can drag any of the points away from the inner polyline to add some softness to the mask.

> **TIP:** Press Shift-A to select all the points on a shape, and then hold O and drag to offset the points from the inner shape. This gives you a starting point to edit the falloff.

The farther the outer shape segment is from the inner shape, the larger the falloff will be in that area.

Adding Additional Points to the Shape

It is not necessary for every point on the inner shape to have a match on the outer shape, or vice versa. You can add additional control points to refine the shape of either shape.

Each polyline stores its animation separately; however, if a point is adjusted on the inner shape that is parented to a point on the outer shape, a keyframe will be set for both splines. Adjusting a parented point on the outer shape only sets a keyframe for the outer shape's spline. If a point that is not parented is adjusted, it will only set a keyframe on the relevant spline. You can disable this behavior entirely for this polyline by selecting Polygon: Outer Polygon > Follow Inner Polyline from the contextual menu.

Locking/Unlocking Point Pairs

If you want to parent additional control points, you can select the points, right-click in the viewer, and choose Lock Point Pairs from the contextual menu for either spline. This will cause the selected point on the outer shape to become parented to the selected point on the inner shape.

Any animation already applied to either point is preserved when the points become parented.

To unlock a point so it is no longer parented, select the point, right-click in the viewer, and deselect Lock Point Pairs from the contextual menu.

Animating Polyline Masks

Animating masks is surprisingly easy. When Polygon or B-Spline masks are added to the Node Editor, the spline's control points are automatically ready to be animated. All you have to do to animate a mask is move the playhead to a new frame and then change the shape of the mask. A new keyframe is added in the Spline Editor and Timeline Editor. This one keyframe controls the position of all control points for that mask at that frame. Once two or more keyframes have been created, the shape of the polygon or B-Spline is automatically interpolated from one keyframe to the next.

> **TIP:** The center point and rotation of a shape are not auto-animated. Only the control points are automatically animated. To animate the center position or rotation, enable keyframes for that parameter in the Inspector.

To adjust the overall timing of the mask animation, you edit the Keyframe horizontal position spline using the Spline Editor or Timeline Editor. Additional points can be added to the mask at any point to refine the shape as areas of the image become more detailed.

Removing Animation from a Polyline Mask

If you want a Polyline mask to remain static, you can remove the automatic animation setting. In the Inspector for the mask, right-click in the bottom of the panel where it says Right Click Here For Shape Animation. From the contextual menu, choose Remove Bézier Spline. If you decide you need to animate the mask at a later time, right-click in the same area again and choose Animate.

Adding and Removing Points from an Animated Mask

When adding points to an animated mask, the new point is fit into the shape at all keyframes. Deleting a point removes that point from all keyframes in the animated mask.

Publishing Specific Control Points

Although you can rapidly animate the entire shape of a polyline using a single keyframe, by default the Spline Editor and Timeline display only one keyframe for the entire shape at any given frame.

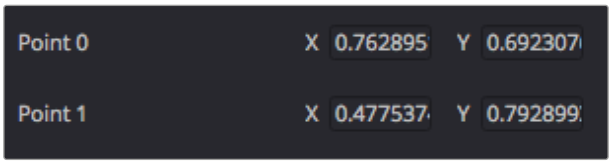
This default keyframing behavior is convenient when quickly animating shapes from one form to another, but it doesn't allow for specific individual control points that need to be keyframed independently of all other control points for a particular shape. If you're working on a complex mask that would benefit from more precise timing or interpolation of individual control points, you can expose one or more specific control points on a polyline by publishing them.

Be aware that publishing a control point on a polyline removes that point from the standard animation spline. From that point forward, that control point can only be animated via its own keyframes on its own animation spline. Once removed, this point will not be connected to paths, modifiers, expressions, or trackers that are connected to the main polyline spline.

To publish a selected point or points, do one of the following:

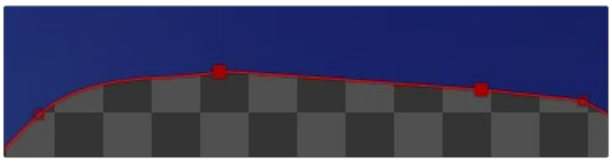
- Click on the Publish Points button in the Polyline toolbar.
- Select Publish Points from the Polyline's contextual menu.

A new coordinate control is added to the Polyline mask controls for each published point, named Point 0, Point 1, and so on.



The Publish Points controls in the Inspector

The onscreen control indicates published points on the polyline by drawing that control point much larger. Once a published point is created, it can be connected to a tracker, path, expression, or modifier by right-clicking on this control and selecting the desired option from the point's contextual menu.



The published point in the viewer

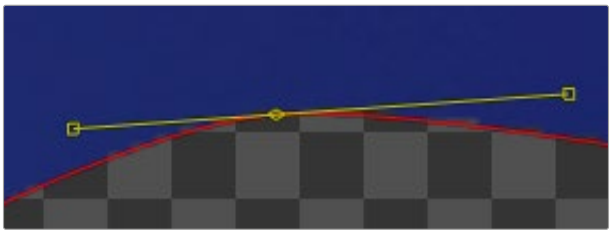
Using "Publish to Path" to Preserve Animation

When a point is published, any animation already applied to that point is removed. However, if you need to keep the animation, you can use the "Publish to Path" option. This Polyline contextual menu option publishes the selected points and converts their existing animation to a path. You can also use the Publish to Path button in the Polyline toolbar.

Using "Follow Published Points" to Add Points

There are times when you will need to have control points that lie between two other published points follow the motion of the published points, while still maintaining their relative offset and shape. For this reason, points in a Polyline mask can be set to "Follow Published Points" using the Polyline's contextual menu.

When a point of an effect mask is set to follow points, the point will be drawn as a diamond shape rather than a small box.



A control point set to Follow Published Points

When this mode is enabled, the new "following" control points will maintain their position relative to the motion of any published points in the mask, while attempting to maintain the shape of that segment of the mask. Unlike published points, the position of the following points can still be animated to allow for morphing of that segment's shape over time.

Chapter 21

Paint

This chapter describes how to use Fusion's non-destructive Paint tool to repair images, remove objects, and add creative elements.

Contents

Paint Overview 529
Types of Paint Nodes 529
Setting Up the Paint Node 530
Setting the Paint Node's Resolution 530
Paint Node Workflow 531
Select the Correct Paint Stroke Type 531
Setting the Brush Size 534
Choosing an Apply Mode 534
Editing Paint Strokes 538
Editing Paint Strokes in the Modifiers Tab 539
Deleting Strokes 539
Animating and Tracking Paint Strokes 540
Animating with Write-On Controls 540
Tracking a Paint Stroke 540
Using the Planar Tracker with the Paint Tool 543
Inverting the Steady Effect to Put the Motion Back In 549
Painting a Clean Plate 551

Paint Overview

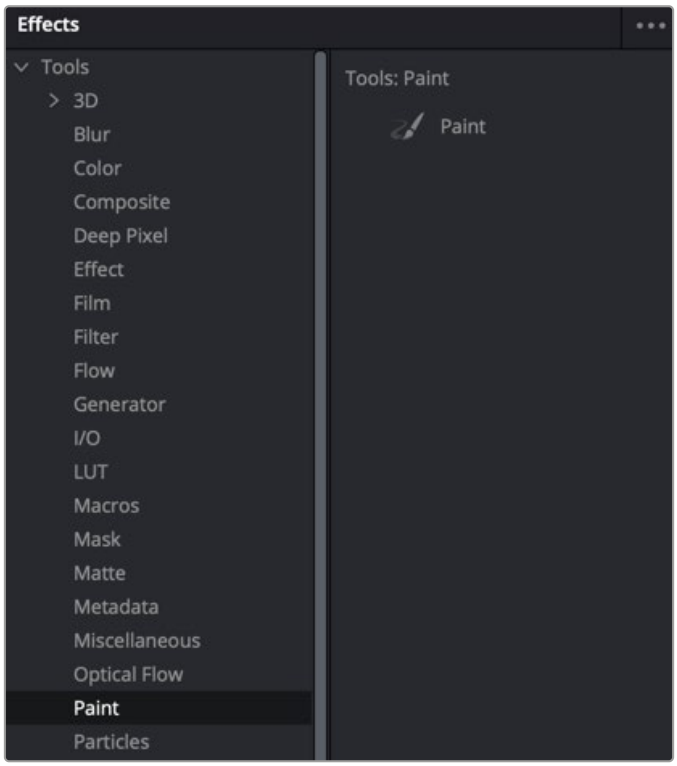
The Paint node is a procedural paint tool, which means that each paint stroke is a live, editable object that's drawn with properties that you can mix and match to address a wide variety of painting tasks. You can use it to paint masks, retouch images, perform beauty work, clone out objects, or even create motion graphics. Each element of a paint stroke can be altered long after you apply it. Since strokes are editable, you can apply, change, ignore, delete, and even reorder them in a node tree.

Types of Paint Nodes

There are two types of paint nodes in Fusion. The Paint node is a full-featured creative and retouch vector-based paint tool that requires an input on which to paint. The Mask Paint node lets you specifically paint an alpha channel to limit the area of an effect. It allows you to create paint strokes on an alpha channel without needing to have an input.

- The Paint node is located in the Paint category of the Effects Library.
- The Mask Paint node is located in the Mask category of the Effects Library.

The main difference between these two Paint tools is that the Mask Paint tool only paints on the Alpha channel, so there are no channel selector buttons. The Paint tool can paint on any or all channels. The majority of this chapter covers the Paint node, since it shares identical parameters and settings with the Mask Paint node.



The Paint node available in the Paint category of the Effects Library

Setting Up the Paint Node

The Paint node has two inputs. Typical of most Fusion nodes, the orange input background is the primary input for connecting the “canvas” or image to paint on, while the second blue input is an Effect Mask. Unlike the Mask Paint node, the Paint node requires a background input to begin painting.

Setting the Paint Node’s Resolution

No matter how you use the Paint node in your node tree, the Paint node assumes the resolution of the background input image as your working resolution for that operation. Although the Paint tool is in fact resolution independent and you can change this assigned resolution at any time, it’s essential to make sure you properly set the resolution of the media you’re working with, because it will affect operations such as motion tracking that you’ll want to use as part of your paint workflow.

Because of this, the Paint node requires a background input to set the resolution of the “canvas” you’ll be painting upon. To do this, you can set up a Paint node in the node tree in one of two ways: painting directly on an image or using Paint as the foreground.

Painting Directly on an Image

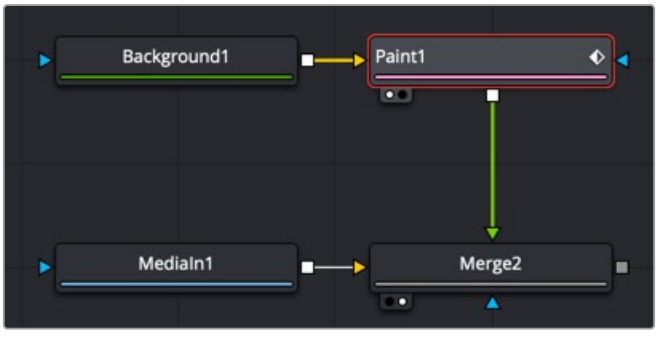
You can connect the image you want to paint on directly to the background input of the Paint node. This is the easiest and cleanest node tree construction, but it doesn’t provide much in terms of compositing flexibility.



The Paint node is inserted directly after the node it is painting on

Using Paint as the Foreground in a Merge Composite

An alternative setup is to use a Background node to set the resolution for the Paint node, compositing the result over the actual background using a Merge node. Working this way lets you use the Merge node’s Apply Mode setting (also referred to as composite modes) to control how your paint strokes are composited against the image, but it does require a bit more setup.



The Paint node is composited over the image you want to paint on using the merge

Setting this up requires some configuration of the nodes. The Background node must be fully transparent and, unless you are doing something simple like using the Stroke tool set to Color to paint over an image, you must drag the image you want to clone or smudge into the Source Tool field in the Paint node's inspector. These steps are described in more detail later in this chapter.

Paint Node Workflow

You begin painting by first selecting the paint stroke type from the Paint toolbar above the viewer. There are ten stroke types to choose from as well as two additional tools for selecting and grouping paint strokes.

The stroke types and tools are described in detail, *see Chapter 53, "Paint Node,"* in the Fusion Reference Manual.

The primary tool for painting and cloning is the Stroke tool. The Stroke tool is a fully animatable and editable vector-based paint stroke and initially uses a duration of the entire global range.



The Stroke tool is most often used for cloning, beauty work, and creative paint

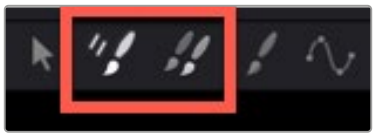
Select the Correct Paint Stroke Type

Instead of having multiple dedicated Paint nodes for different operations, the procedural nature of Fusion's Paint tool means there is just one tool with a variety of different stroke types. Each stroke type serves a different purpose. The Paint toolbar above the viewer allows you to choose between different stroke types and drawing tools. These are grouped into a few categories.

Single-Frame Stroke Types

The Multistroke and Clone Multistroke are explicitly designed for single-frame retouching paint jobs like removing raindrops in a shot that's supposed to be a sunny day or removing scratch marks and dust when restoring vintage content. When handling these types of jobs, these two Multistroke options are faster than the other stroke types, but they're not editable later on. This means you must set up the size and function of the brush as well as the duration for the paint stroke before you paint.

The duration of each of these stoke types is one frame by default, but this can be changed using the Stroke Duration slider in the Inspector. Multistroke and Clone MultiStroke are basically the same tools, except Clone Multistroke automatically configures the tool for cloning. In contrast, the Multistroke requires you to set up the tool for cloning manually.



The Multistroke and Clone Multistroke are non-editable single frame stroke types

Editable Stroke Types

The Stroke and Polyline are similar in that they can be modified and animated at any time. Also, they both begin with a duration lasting the entire comp, but that also can be changed using the Keyframes Editor.

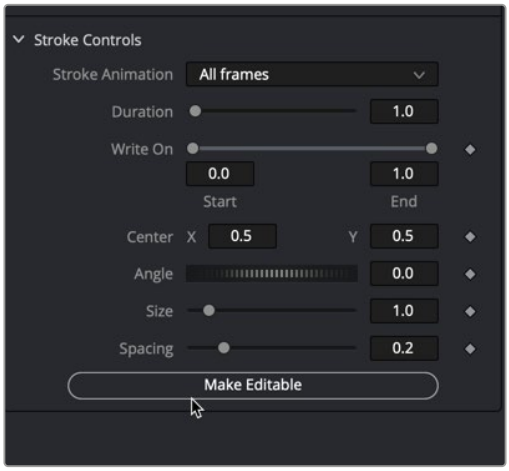


The Stroke and Polyline strokes are editable and last for the entire comp

The Stroke Tool

One of the most flexible editable stroke types you'll use for many tasks is the Stroke, because it is fully animatable and editable. You can animate all elements of the Stroke, and you can use the Write-on/Write-off parameters to control how the stroke appears onscreen. You can also connect to a tracker from the Center point of the Stroke if you want to make the stroke follow specific onscreen motion.

By default, the Stroke type does not expose control points for the shape of the path. You can move and track the center and rotation of the Stroke, but the individual control points that create the spline are hidden. To reveal the control points, you can open the Stroke Controls at the bottom of the Inspector and click the Make Editable button.



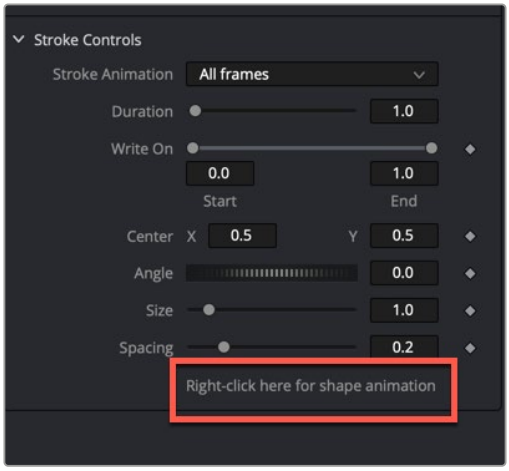
The Stroke's control points can be revealed using the Make Editable button

Although the Stroke type is the most flexible, that flexibility can come at a performance penalty if you're painting hundreds of strokes on a frame. For larger numbers of strokes that do not need to be animated, it's better to use Multistroke or Clone Multistroke, as these are more processor efficient.

The Polyline Stroke Tool

The Polyline Stroke acts more like a drawing tool than a paintbrush. It includes the same functionality as the Stroke tool, except that it is created not by dragging or "painting" like a paintbrush, but by clicking to create a spline path, as you do with masks and motion paths. Without even creating a stroke in the viewer, the Polyline Stroke can connect to existing polylines like a mask or a motion path.

If a motion path is published, right-clicking on the Shape Animation label at the bottom of the Polyline Stroke's Stroke Controls allows you to use the Connect To menu to assume the shape of a motion path or mask. You can also use this method if you import SVG graphics and want to "paint-on" the outlines.



The Polyline Stroke Shape Animation label

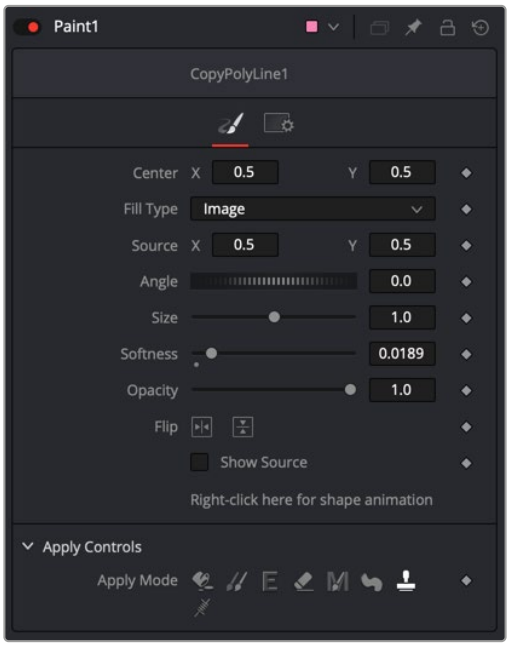
Shape Drawing Tools

Five shape-based drawing tools allow you to draw shapes and either fill them with a color or clone an area from a source image. All these tools act similarly to the Stroke and Polyline stroke type in that they are editable at any time and have a default duration spanning the entire global range of the comp. However, you can edit the duration at any time in the Keyframes Editor.



The Shape Strokes are used to create shapes or clone areas based on shapes

All of the Copy [Shape Name] stroke types require that you connect the source node you are cloning from into the Paint node, and set the Fill Type menu to Image.



The Copy Shapes require a Source to be directly connected to the Paint tool and set to Image

Setting the Brush Size

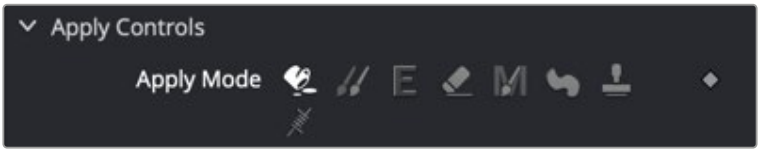
After selecting the Stroke type, the Brush size can be set in the Inspector or more intuitively in the viewer. With the Paint node selected in the Node Editor and the pointer positioned over the viewer, you can see an outline of the current brush size. To change the brush size, hold down the Command key and drag. The circle changes size, so you can set it relative to other objects you may be painting over.



Brush size can be changed interactively in the viewer

Choosing an Apply Mode

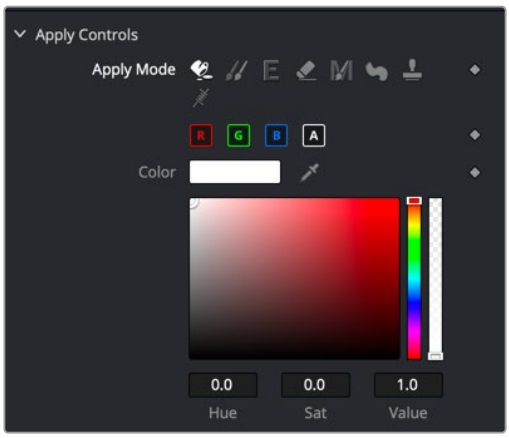
The Apply Mode buttons determine the functionality of the paintbrush. There are eight Apply modes that set the brush to do things like paint a color, clone from a source, smudge an area, or remove thin wires.



Apply Mode buttons determine the paint brush functionality

Picking a Paint Color

There are several ways to pick the paint color and opacity for a colored brush stroke. You use the Fill button in the row of Apply modes when you want to paint with a solid color.



The Color swatch shows the current color
Clicking it opens the OS Color Picker window

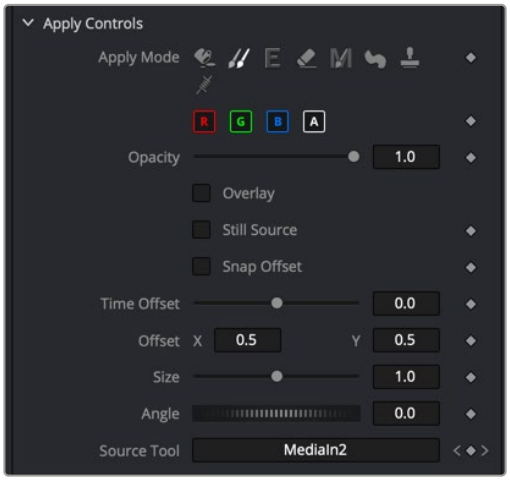
To select a color for the paint brush, do one of the following:

- Click the color swatch to open a standard OS Color Picker window.
- Drag the Eyedropper into the viewer.
- Drag inside the color chooser to select a saturation and luminance. Drag on the sidebars to change the hue and transparency.

When you paint, each stroke is unpremultiplied, so adjusting the Alpha slider in the Inspector does not affect what you apply to the RGB channels. However, changing opacity affects all four channels.

Cloning from the Frame

Choosing the Clone Apply Mode allows you to paint from one area of an image over another area. This is the most common use of the Paint tool. It allows you to remove objects or artifacts from a clip by covering them up with another area of the frame. Depending on the Stroke type chosen, you may clone on either a single frame or for the entire duration of the clip.



The Clone Apply Mode allows you to sample from one area and use it as a source to paint over another area

You can use the Clone Apply Mode to clone from the same image connected to the Paint node's background input or a different source from the node tree.

To clone from a different area of the same frame:

1. Select a Stroke type from the Paint toolbar above the viewer.
2. Using the size slider in the Brush controls section of the Inspector, set the size of the brush.
3. From the Apply Mode buttons, select the Clone mode.
4. Option-click over the area in the viewer you want to use as the source. A dot appears showing you the center of what you're sampling from.



The Clone source starting area identified by the X and the paint brush size represented by the circle

5 Paint over the area you want to cover up using the source pixels.

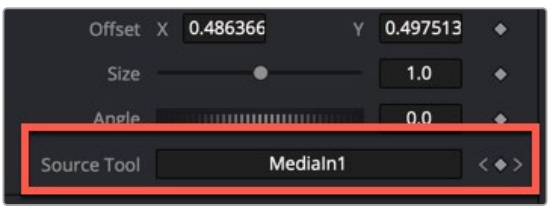


The Clone completed after selecting the source area and painting over the flag pole

When trying to erase objects or artifacts from a clip using the Clone Apply Mode, it can sometimes be easier if you sample from a different frame on the same clip. This works well when the object you are trying to clone out moves during the clip, revealing the area behind the object. Sampling from a different frame can use the revealed background by offsetting the source frame.

To clone from a different frame of the same clip:

1 Select a Stroke type from the Paint toolbar above the viewer.
2 Set the size of the brush.
3 From the Apply Mode buttons, select the Clone mode.
4 Drag the clip (MediaIn or Loader) from the Node Editor into the Source Tool field in the Inspector.



The Paint Inspector with the MediaIn1 dragged from the Node Editor into the Source Tool field

5 Click the Overlay checkbox to see the current frame and the offset frame superimposed.

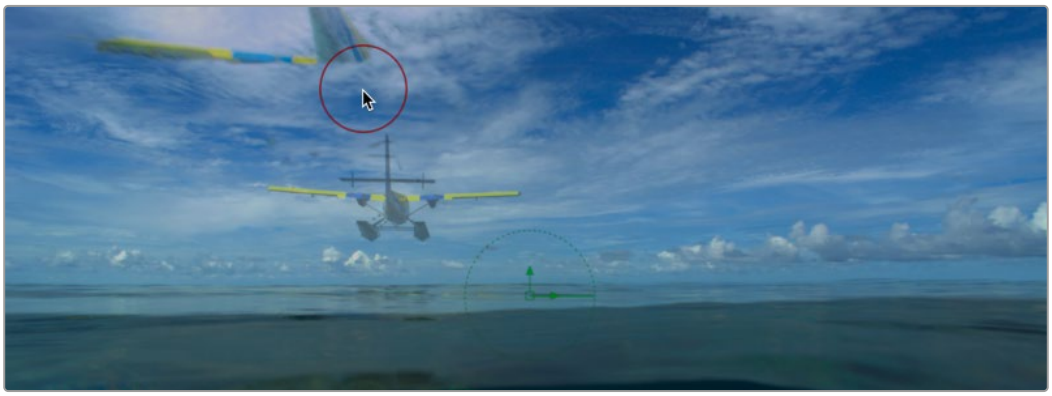
6 Drag the Time Offset slider to select the source frame you want to use.

7 Option-click over the area in the viewer you want to use as the source or to offset the source frame's position.



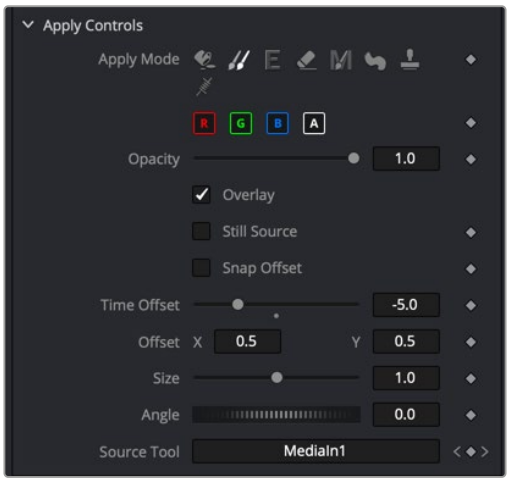
Overlay shows two frames overlapped with Time Offset, allowing you to clone from one frame onto another

8 Paint over the area you want to cover up using the source pixels.



The plane is half painted out using on the Overlay with Time Offset

9 Disable the Overlay checkbox.



The Clone Apply Mode can use a different frame from the same clip

TIP: When using a Clone Apply Mode, you can hold down the O key instead of clicking the Overlay checkbox in Inspector to see the Overlay. Releasing the O key will return to normal viewing without the Overlay.

Editing Paint Strokes

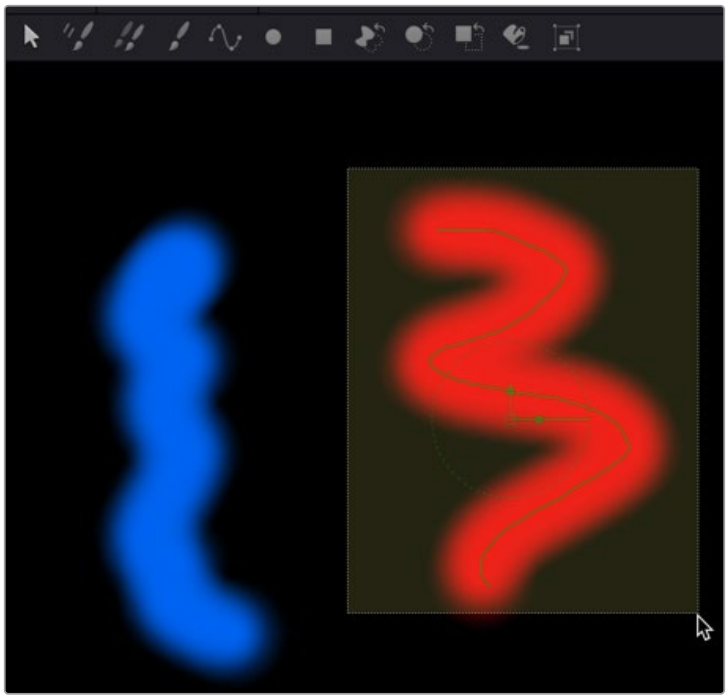
Once you've painted using the Stroke or Polyline stroke type, you can change the look of the stroke by selecting it and updating the parameters in the Inspector. Selecting the stroke requires you to switch to the selection tool in the Paint toolbar above the viewer. Using the Paint node's selection tool, you can either click once on a stroke or drag a bounding box around a stroke to select it for editing.

TIP: To select multiple strokes, you can Shift-click or Command-click to select and deselect multiple specific strokes, or you can drag a selection box around all strokes you want to select.

Although you can make changes in the Tools tab in the Inspector, the Paint node uses both the Tools tab and the Modifiers tab. In the Tools tab, you can create new brush strokes and select a stroke in the viewer to edit. The Modifiers tab presents a list of all the strokes for the selected Paint node, which makes it easy to modify any previously created paint stroke.

NOTE: Multistroke and Clone Multistroke each only appear as one item in the Modifiers tab no matter how many strokes you create using those tools. Those two tools are not editable after creating them.

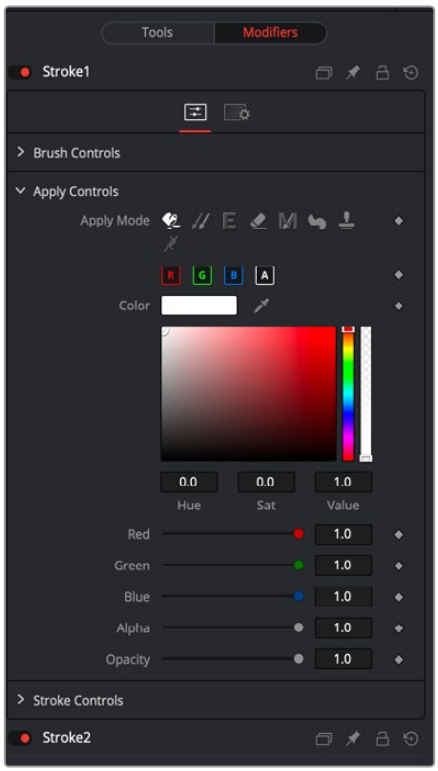
The same controls you used in the Tools tab to create the strokes are located in the Modifier's tab to modify them. You can also animate each individual stroke.



The Stroke or Polyline Stroke type can be edited by selecting the stroke in the viewer

Editing Paint Strokes in the Modifiers Tab

When you paint a stroke, the settings for that stroke appear in the Inspector's Modifiers tab. You can then change the settings in the Tools tab for the next stroke you are about to paint. Each time you click, drag, and release the pointer button, you create a new stroke. Each stroke is numbered in the Modifiers tab, where it can be selected and edited.



Each stroke is listed in the Modifiers tab for editing

Once you stop painting a stroke, it's added to the Modifiers tab along with an additional Stroke modifier that represents the next stroke. For instance, if you paint your first stroke, the Modifiers tab shows your stroke as Stroke1 and then a stroke 2 as well, which represents the next stroke you create. You always have one more stroke in the Modifiers tab than strokes in the viewer.

Deleting Strokes

There are two ways you can delete paint strokes.

To delete any individual stroke, do the following:

1. Select the Paint node.
2. Click the Modifiers tab.
3. Right-click over the Stroke header you want to delete and choose Delete from the menu.

To delete all paint strokes you've made on every frame, do one of the following:

- Click the reset button in the upper-right corner of the Inspector.
- Delete the Paint node in the Node Editor.

Animating and Tracking Paint Strokes

In some ways, animating paint strokes is no different than animating any other effect in the Inspector. Each parameter that can be animated includes a gray diamond Keyframe button along the right side. Clicking the Keyframe button sets a keyframe on the current frame and enables auto-keyframe mode for the parameter. However, more commonly, the paint stroke is tracked using one of Fusion's trackers, or for motion graphics, animated using the Write-On Start and End sliders.

Animating with Write-On Controls

The Stroke and Polyline stroke types include Write-On controls located in the Stroke controls section of the Inspector. These Write-On controls animate the appearance of a stroke along the path. You can animate the Write-On controls using either the Stroke Animation drop-down menu or using the Start and End sliders.

Stroke Animation Drop-Down Menu

The Stroke Animation drop-down menu includes six options for auto-animating a paint stroke. The first two options do not truly animate the stroke as much as set a duration. The Limited Duration option uses the Duration slider to set the number of frames the stroke is onscreen.

To auto-animate the stroke, you can choose one of the three Write options or the Trail option.

Choosing Write On automatically creates a write-on animation. The duration is set by two keyframes that get added when you choose Write On from the menu. The Start keyframe is set on the frame where you first created the stroke. The End keyframe is added on the current frame when you choose Write On from the menu. The remaining options in the menu set their Start and End keyframes similarly but change the direction of the animation based on the menu selection.

Write-On Start and End Parameters

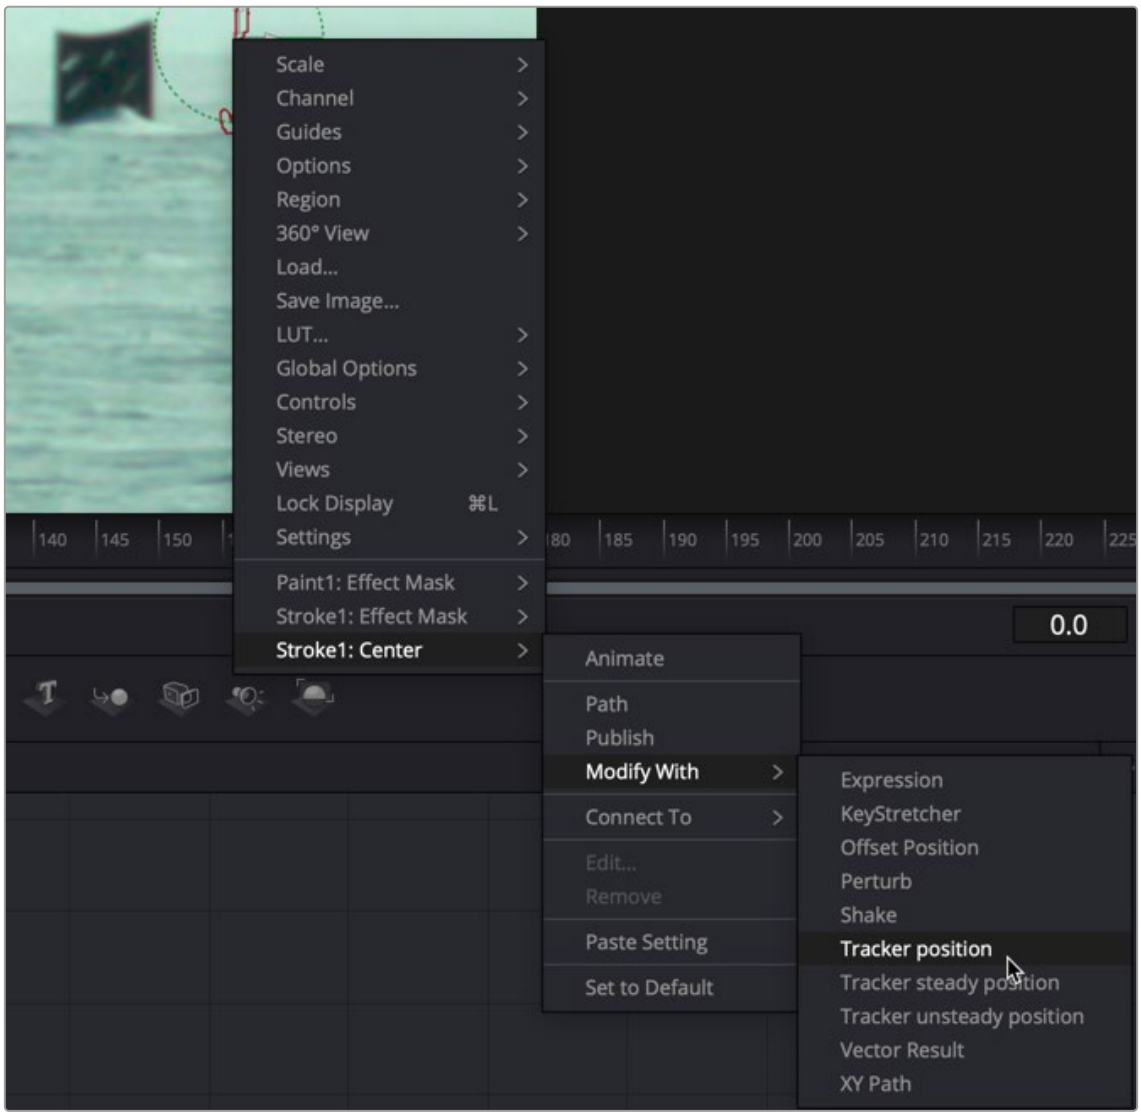
The Write-On Start and End parameters allow you to manually control the start point and end point along any stroke's path, and use keyframes to animate each parameter individually. The Start parameter determines the point at which the stroke begins, measured as a percentage offset from the beginning of the stroke's path. For example, a Start value of 50 moves the starting point of the stroke to the middle of the stroke's path. The End parameter works the same way but from the other end of the stroke. You can animate a stroke onscreen, creating a handwriting effect by setting keyframes for the End parameter from 0 to 100 over several frames.

Tracking a Paint Stroke

You can animate the position of a paint stroke using any of Fusion's trackers. For instance, if you cloned out a flag pole from a clip, but the camera moves, you can track the flag pole and attach the resulting path to the paint stroke.



Trackers can be attached to the Center parameters of a paint stroke



Right-clicking over the Stroke's Center control in the view allows you to apply a tracker modifier

To attach a tracker to a paint stroke:

1. With the Paint node, select the Stroke brush type and clone out an object on a frame.
2. In the Paint toolbar above the viewer, click the Select tool.
3. Drag a selection box around the stroke to select it.
4. Right-click the center control on the stroke, and then choose Stroke1:Center > Modify With > Tracker Position.
5. Click the Modifiers tab to view the Tracker controls.
6. From the Node Editor, drag the MediaIn for the image you painted on, and drag it into the Tracker Source field in the Inspector.



Drag the MediaIn you want to track into the Tracker Source field in the Inspector

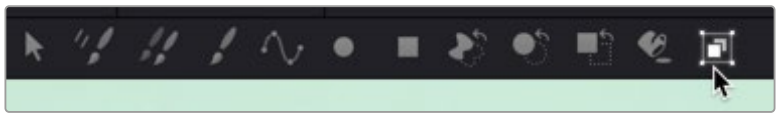
7. Click the Track Forward button.
8. After tracking, at the bottom of the Inspector, use the Tracker 1 X Offset/Y Offset controls to reposition the paint stroke, if necessary.

Tracking a Group of Paint Strokes

You can assign a tracker to multiple strokes by adding the stroke to a group and connecting the tracker to the group. Instead of connecting each individual stroke, the group's center is used for all of the strokes. Assuming the motion of each object is consistent in the same direction, as it would be objects "nailed to the set", then applying the tracker to the group makes cloning multiple objects out with a single paint node very easy.

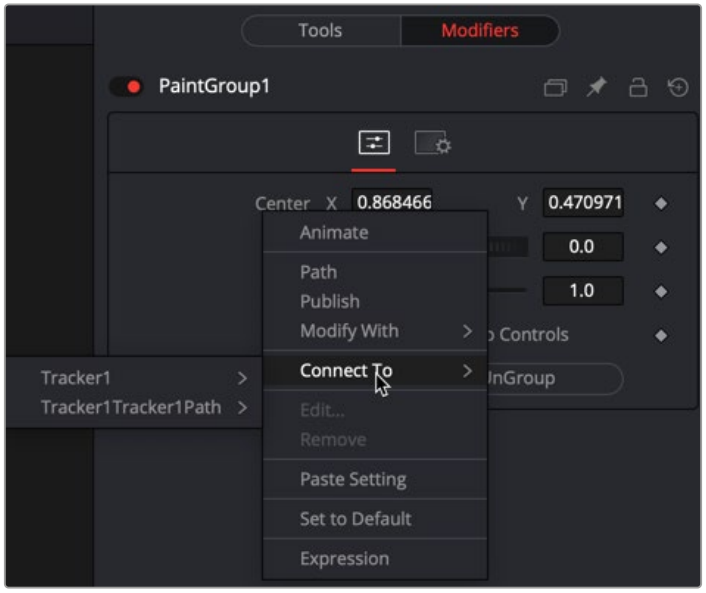
To group paint strokes, do the following:

1. Drag a bounding box, Shift-click, or Command-click to select every stroke that you want to group together.
2. Click the Paint Group button in the Paint toolbar.



Selecting all the strokes and then clicking the Paint Group button collects all the strokes into a single group

The group's onscreen controls replace the controls for each paint stroke, and the Modifiers tab in the Inspector shows the group's parameters. The individual strokes are still editable by selecting Show Subgroup Controls in the Modifiers tab of the Inspector. The group then comes with a Center, Angle, and Size control for connecting to a tracker.



Use the onscreen controls, or in the Modifiers tab, right-click over the Paint Group's Center X label to connect the tracker

Using the Planar Tracker with the Paint Tool

Here's an example that dives deeper into a workflow where we use the Paint tool with the Planar Tracker for retouching a clip. We'll eliminate some facial scars on an actor's forehead in a commercial by combining the Paint node with the PlanarTracker node, illustrating a common way of using these two powerful tools together.



The actor has some scars on his forehead that the director would like painted out

Because this is a clip in motion, we can't just paint out the scars on the man's forehead; we need to deal with the motion so that the paint work we do stays put on his face. In this case, a common workflow is to analyze the motion in the image and use it to apply a "steady" operation, pinning down the area we want to paint in place so we can paint on an unmoving surface.

Setting Up the Planar Tracker for Stabilization

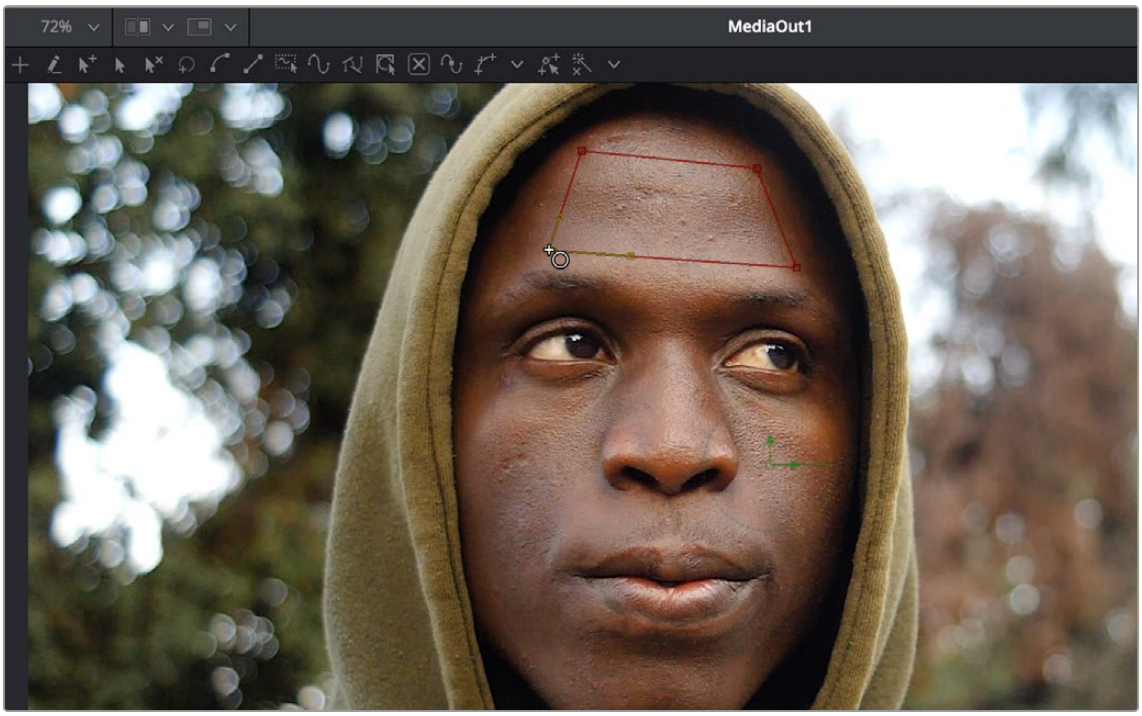
The best way to do this is to use the Planar Tracker, so we'll add the PlanarTracker node after the MediaIn1 node, such that the image connects to the background input of the PlanarTracker node. As always, it's important to be careful about which input you connect for the effect to work properly.



Adding a PlanarTracker node to analyze and steady the part of the image we want to paint on

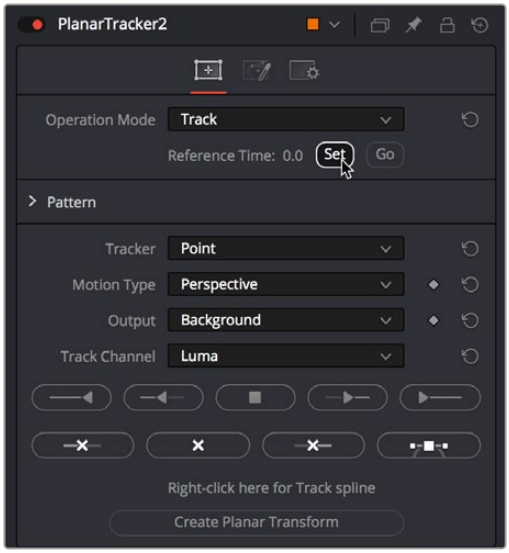
With the PlanarTracker node selected and loaded in the viewer, a viewer toolbar appears with a variety of tools for drawing shapes and manipulating tracking data. The Planar Tracker works by tracking flat surfaces that you define by drawing a shape around the feature you want to track. When you first create a PlanarTracker node, you can immediately begin drawing a shape, so in this case, we draw a simple polygon over the man's forehead since that's the feature we want to steady in preparation for painting.

We draw a simple box by clicking once each on each corner of the man's forehead to create control points, and then clicking the first one we created to close the shape.



Drawing a shape over the man's forehead to prepare for Planar Tracking

In the Inspector, the PlanarTracker node has tracking transport controls that are similar to those of the Tracker. However, there are two buttons, Set and Go, underneath the Operation Mode menu, which defaults to Track, since that's the first thing we need to do. The Set button lets you choose which frame to use as the "reference frame" for tracking, so you click the Set button first before clicking the Track Forward button below.



Setting a reference frame at the beginning of the range of frames we want to track

TIP: The Set button lets you supervise a Planar Track in progress and stop it if you see it slipping, making adjustments as necessary before clicking Set at the new frame to set a new reference before continuing to track forward towards the end of the clip.

The Pattern controls let you set up how you want to handle the analysis. Of these controls, the Motion Type menu is perhaps the most important. In this particular case, Perspective tracking is the analysis we want. Still, in other situations, you may find you get better results with the Translation, Translation/Rotation, and Translation/Rotation/Scale options.

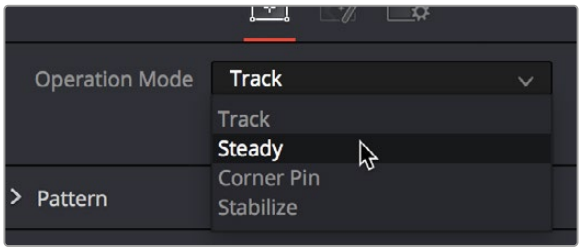
Once you initiate the track, a series of dots appears within the track region shape you created to indicate trackable pixels found. A green progress bar at the bottom of the Timeline ruler lets you see how much of the shot is remaining to track.



Clicking the Track from First Frame button to set the Planar Track in progress; green dots on the image and a green progress bar let you know the track is happening

NOTE: If nothing happens when you track, or it starts to track and then stops, that's your cue that there isn't enough trackable detail within the shape you've drawn for the Planar Tracker to work, and your best bet is to choose a different location of the image to track.

Once the track is complete, you can set the Operation Mode of the PlanarTracker node's controls in the Inspector to Steady.



Setting the PlanarTracker node to Steady

You'll immediately see the image warped as much as is necessary to pin the tracked region in place for whatever operation you want to perform. If you scrub through the clip, you should see that the image dynamically cornerpin-warps as much as is necessary to keep the forehead region within the shape you drew pinned in place. In this case, this sets up the man's head as a canvas for paint.

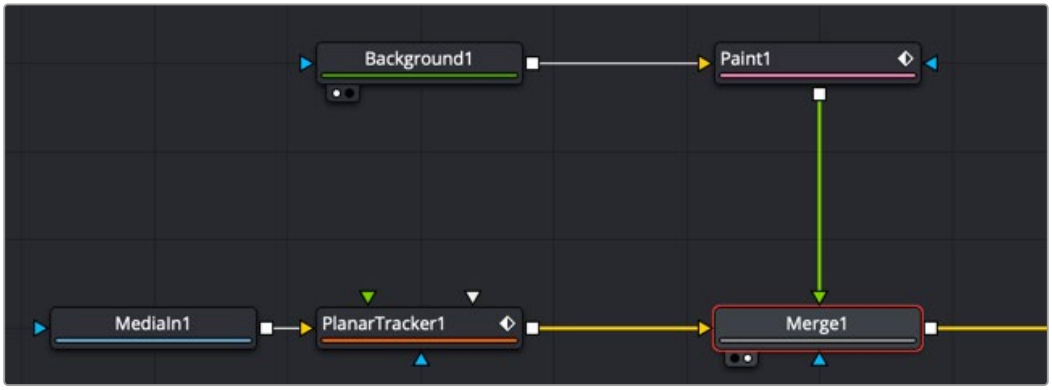


Steadying the image results in warping as the forehead is pinned in place for painting

At this point, you're ready to paint out those scars.

Connecting the Paint Node

Although you could paint directly on the image by connecting the Paint node after the Planar Tracker, it gives you more control over the process if you merge the Paint node over the top of the steadied image. To do that, you'll add a Merge after the Planar Tracker and then connect a Background node into the foreground of the Merge. Insert a Paint node after the background, and you are ready for this process.



Paint node merged over the top of a steadied Planar Tracker

Making the Background Node Transparent

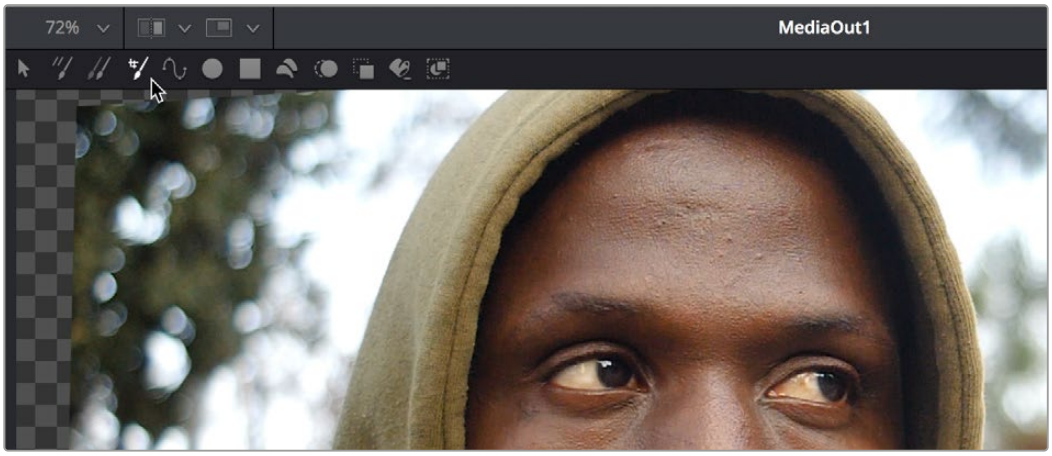
The Background node must be fully transparent; otherwise, you are painting on the solid color instead of the Merge background image. To make the background transparent, drag the Alpha channel sidebar all the way up in the Color Picker.



The Color Picker in the Background node set to be transparent

Selecting the Stroke and Clone Mode

With the Paint node selected and set up, the next thing we want to do is to select the Stroke tool. The Stroke tool is the tool of choice when you want to paint out features or paint in fixes to subjects within the frame that need to remain in place for the entire shot.



Choosing the Stroke tool from the Paint node's tools in the viewer toolbar

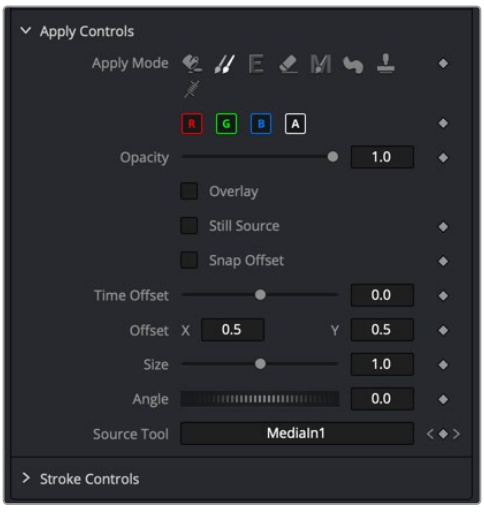
Next, choose the Clone mode from the Apply Controls. In this example, we'll clone part of the man's face over the scars to get rid of them. Choosing the Clone mode switches the controls of the Paint node to those used for cloning.



Choosing the Clone mode in the Inspector

Setting a Clone Source

The Paint node clones from the connected Input image unless you instruct it otherwise. If you are just painting a color over the background image, then you are good to go. However, if you plan on using the merge's steadied background image for cloning with the Paint tool, you must set that image as the Paint node's source. To set the clone source for the Paint node, you drag the PlanarTracker node into the Source Tool field in the Inspector.



Any node can be dragged into the Source Tool field when cloning with the Paint tool

With the Stroke tool selected in the Paint toolbar, the Clone mode selected in the Inspector controls, and the Source for cloning added to the Source Tool field, we're ready to start painting. If we move the pointer over the viewer, a circle shows us the paint tool, ready to go.

To use the clone brush, first hold down the Option key and click somewhere on the image to identify the source area of the clone. In this example, we'll sample from just below the first scar we want to paint. After Option-clicking to sample the image, you can click to begin painting anywhere in the frame.



Setting an offset to sample for cloning (left), and dragging to draw a clone stroke (right)

If you don't like the stroke you've created, you can undo with Command-Z and try again. We repeat the process with the other scar on the man's forehead, possibly adding a few other small strokes to make sure there are no noticeable edges, and in a few seconds, we've taken care of the issue.

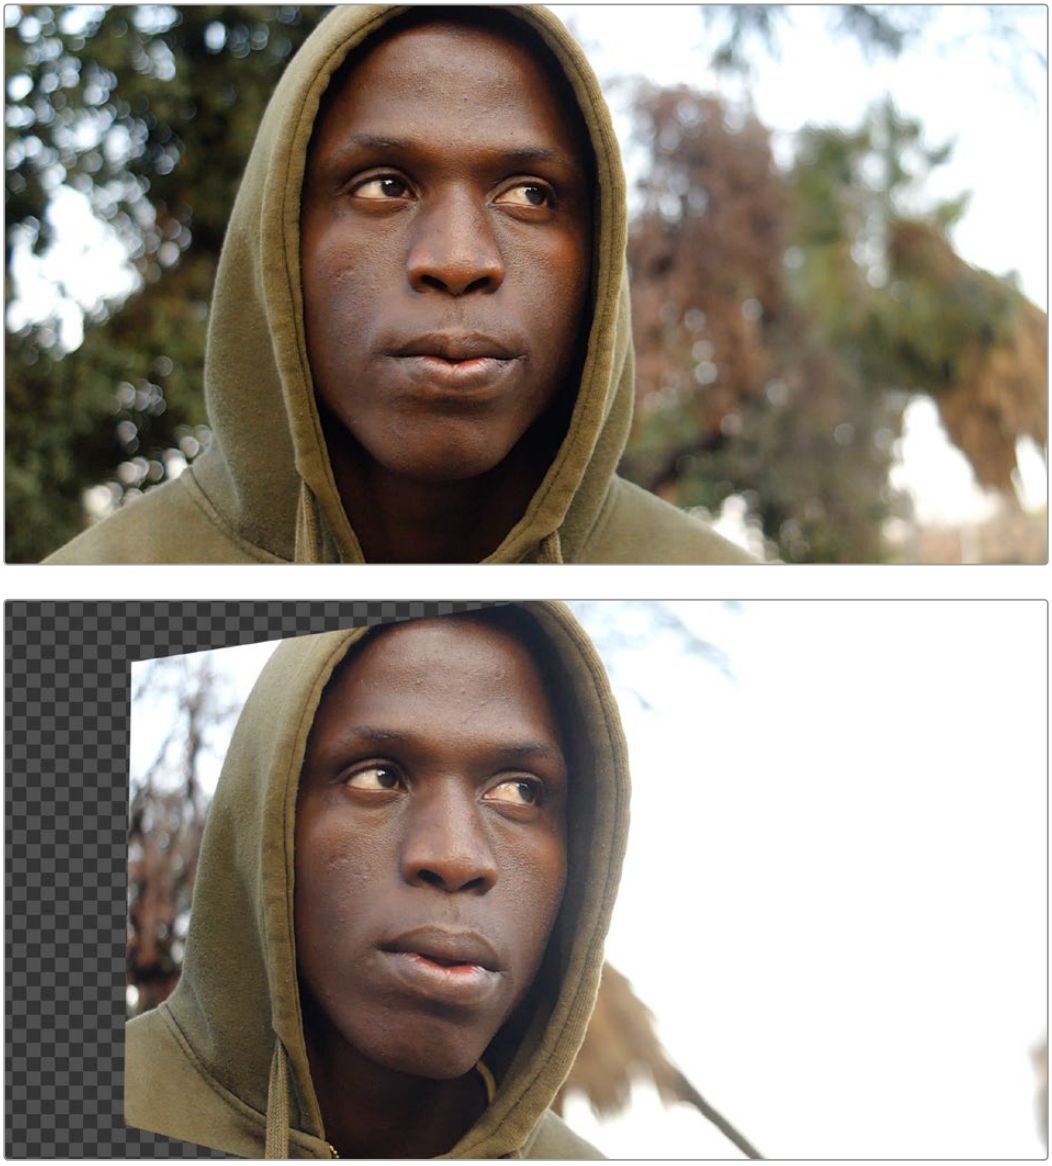


Original image (left), and after painting out two scars on the man's forehead with the Stroke tool set to Clone

TIP: You can adjust the size of the brush right in the viewer, if necessary, by holding down the Command key and dragging the pointer left and right. You'll see the brush outline change size as you do this.

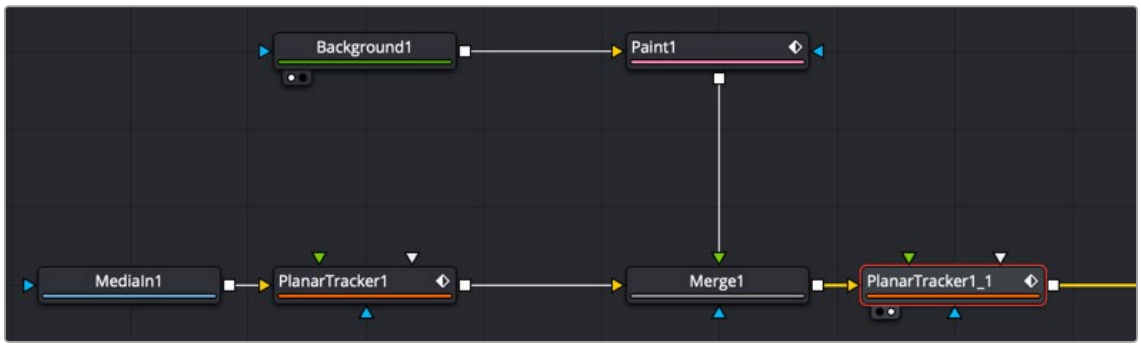
Inverting the Steady Effect to Put the Motion Back In

At this point, scrubbing through the clip shows that the paint strokes we've made are indeed sticking to the man's forehead as we need them to do. Now we just have to invert the transform that the Planar Tracker applied to restore the clip back to the way it was, except now with the painted fix attached in the process. This ends up being a two-part process, but the first part is the simplest.



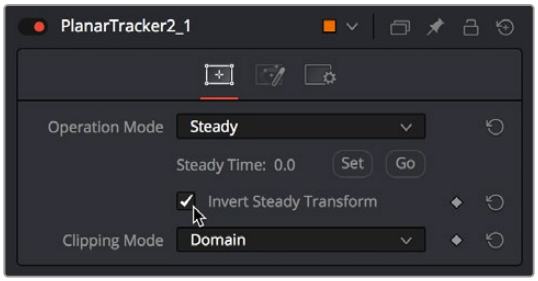
Scrubbing through the steadied clip shows the paint fix is “sticking” to the man’s forehead

We select and copy the PlanarTracker node coming before the Merge node, and paste a copy of it after. This copy has all the analysis and tracking data of the original PlanarTracker node.



Pasting a second copy of the PlanarTracker node after the Paint node

With the second PlanarTracker node selected, we go into the Inspector and turn on the Invert Steady Transform checkbox, which inverts the steady warp transform to restore the image back to the way it was.



Turning on Invert Steady Transform to try to restore the image back to the way it was

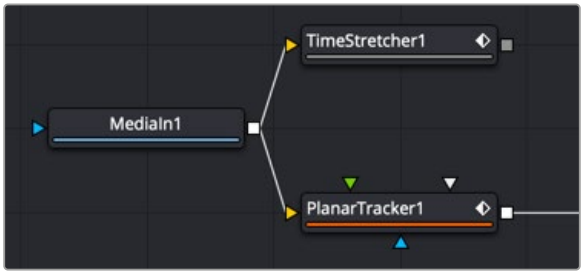
This is just one example of how to set up a Planar Tracker and Paint node. In some instances, you made need to do more work with masks and layering, but the above example gives you a good starting point.

Painting a Clean Plate

On simple clips, planar tracking the clone paint strokes may work fine. In other cases, you may not be able to steady the clip, or the strokes might appear like they are “bubbling.” Paint is just a single stroke repeated over multiple frames. Each painted frame has the potential to appear differently than the one before or after. Often, the more reliable way to use paint is by creating what is referred to as a “clean plate.” Create one good painted frame and freeze it. Then use the Planar Tracker to track the cleaned area over the top of the original. The benefit is that there is only a single paint stroke instead of repeated paint strokes that potentially show unwanted artifacts. For this next example, we continue with the previous clip, but now the technique is to freeze, mask, and composite just a single frame of the fixed forehead over the original clip.

Creating a Freeze Frame and Clean Plate

As in the previous example, you still track the clip using the Planar Tracker. After tracking, you can branch out from the MediaIn (or Loader node if you are using Fusion Studio) into a Time Stretcher to create the freeze frame.



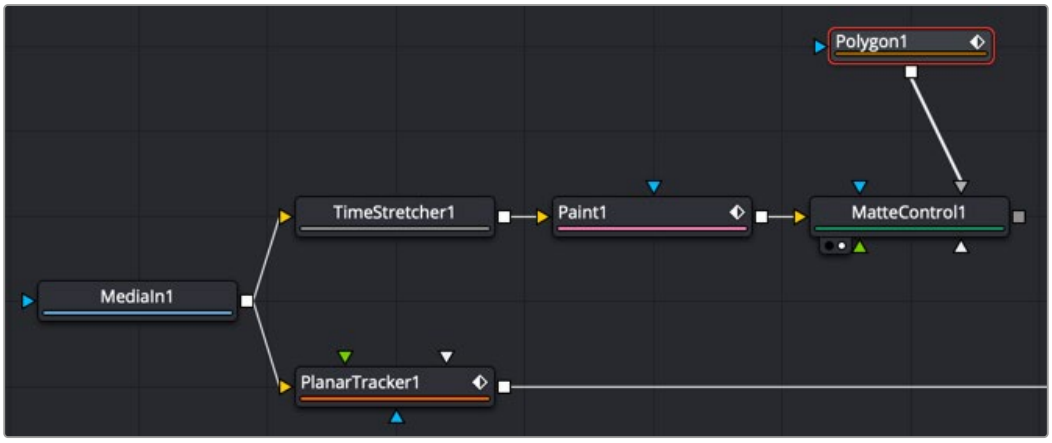
Time Remap is used to retime or freeze a clip

Disable the default Keyframe in the Time Stretcher and enter the frame you want to freeze. If you have already performed a Planar Track, then entering the frame you set as the Reference Frame is usually a good frame to freeze.

To create the clean plate, you connect the paint node to the output of the Time Stretcher. Clone over the areas you want to hide, and you now have a single clean frame. Now you need to composite the clean area over the original.

Isolating the Painted Forehead

First, we need to mask out just the man's painted forehead. We can do this by connecting a MatteControl node to the output of the Paint node and then connect a Polygon node to the garbage matte input. This lets us draw a shape with the Polygon node and use it as a mask to crop out the man's painted forehead.



Add a MatteControl node with a garbage mask to cut out the painted forehead

TIP: When it comes to using masks to create transparency, there are a variety of ways to connect one—for example, (a) attach the image to the background input of a Brightness/Contrast node and attach a Polygon mask node to the effect mask input. On the Brightness/Contrast node, enable the Alpha channel and lower the Gain slider to darken a hole, or (b) using Channel Booleans to copy channel data to the alpha from a Polygon node attached to the foreground input and the image attached to the background input.

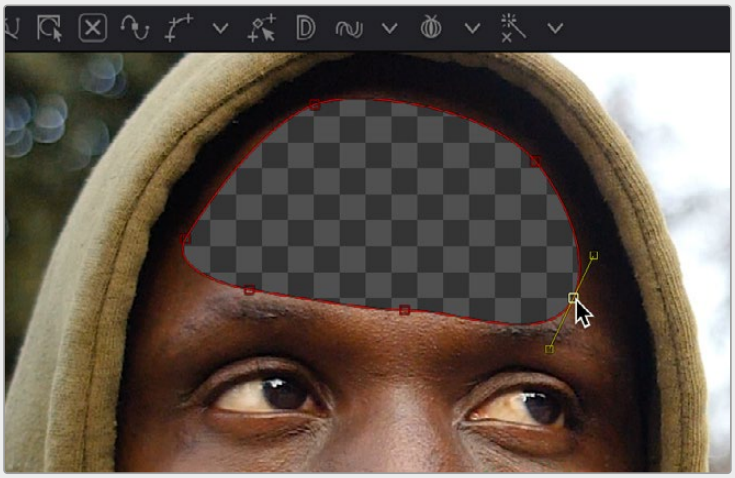
Drawing a Polygon Mask

After moving the playhead to the first frame of the clip, we're ready to draw a mask to isolate the "clean plate" forehead. Loading the MatteControl1 into the viewer and selecting the Polygon1 node so that we see its tools in the viewer toolbar sets us up for drawing a polygon.

Drawing shapes using the Polygon node is similar to shape drawing in other spline-based environments, including the Color page:

- Clicking once draws a corner control point.
- Clicking and dragging creates a Bézier curve.
- Clicking the first control point you created closes a shape.

We click and drag to create a shape that outlines the man's forehead, and when we close the shape, we see exactly the opposite of what we want, a hole in the middle of the image.

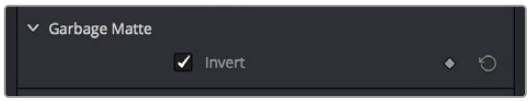


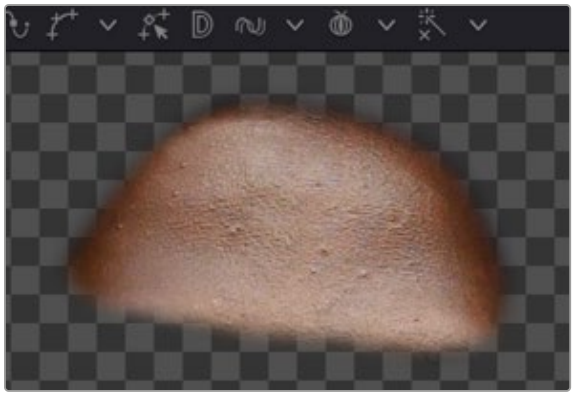
Drawing a shape to isolate the forehead gives an inverted result at first when using the Garbage Matte input of the MatteControl node to attach the Polygon to the MatteControl node

Before fixing this, we drag the Soft Edge slider in the Inspector to the right to blur the edges just a bit.

Inverting the Garbage Input

Selecting the MatteControl1 node, we open the Garbage Matte controls and click the Invert checkbox, which immediately gives us the result we want, of the forehead in isolation, ready for compositing.





Inverting the Garbage Matte input (Left), and the resulting inverted mask inverting the forehead (right)

Compositing the Painted Forehead Against the Original Image

Almost finished. Next, we'll add a Merge node to layer the fixed forehead over the original image being output by the MediaIn node.

We create a Merge node connected to the output of the PlanarTracker node, and then we connect the MatteControl's output to the green foreground input of the Merge node. This puts the cropped and fixed forehead on top of the original image.



The painted forehead composited against the original image

Match Moving the Mask to the Shot

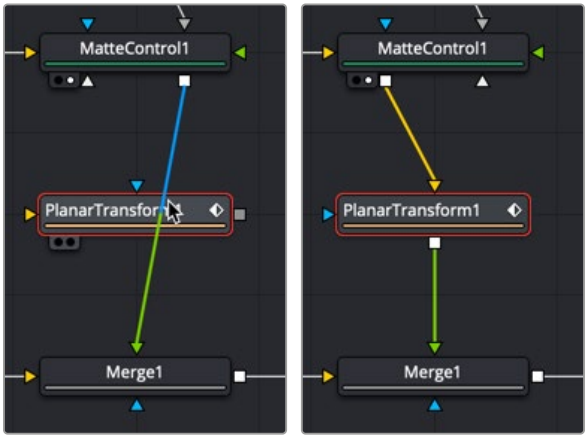
Now we have the best of both worlds: a fixed forehead and the background of the shot looking good. However, if we select the Polygon node and then scrub forward in the clip, the fixed forehead mask drifts out of sync with the motion of the shot, so we have one remaining issue. Fortunately, match moving the mask to move with the shot is really simple.



Because the Polygon isn't animated to match the motion of the shot, it goes out of sync

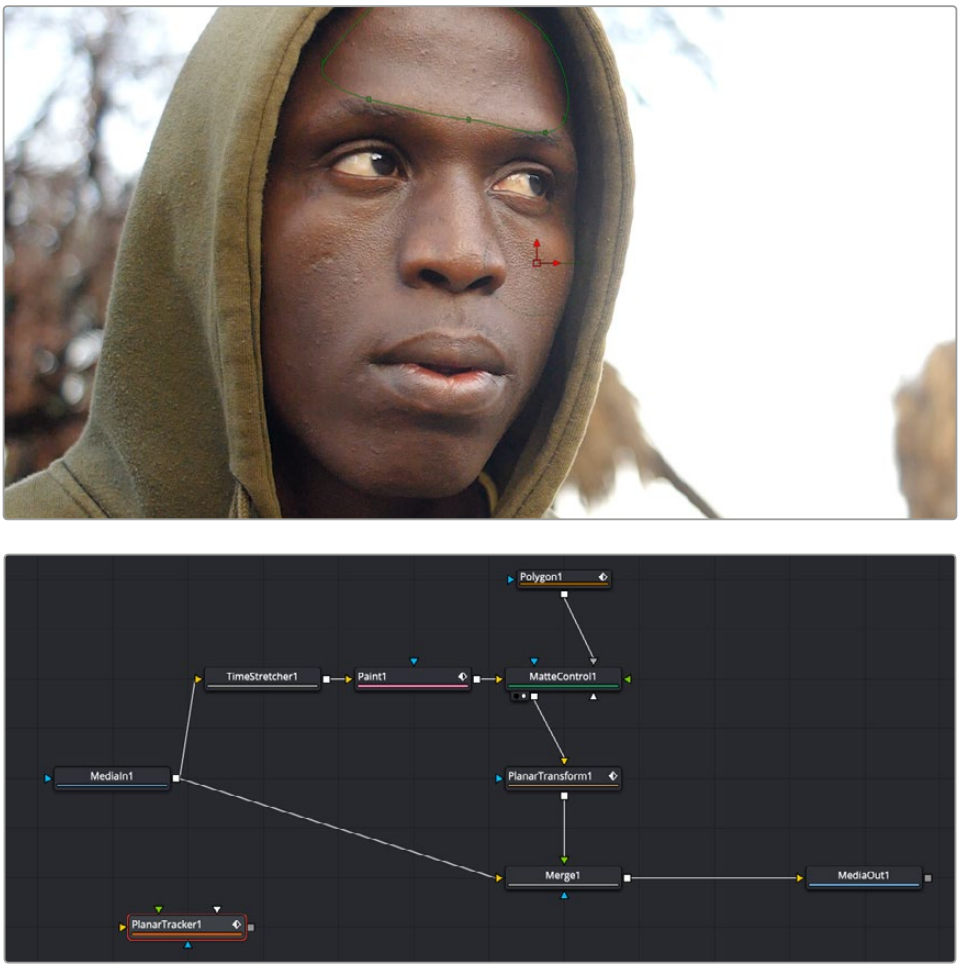
Selecting the first PlanarTracker node that comes right after the MediaIn node, and choosing Track from the Operation Mode menu, reveals a Create Planar Transform button at the bottom of the listed controls. Clicking this button creates a new, disconnected Planar Transform node in the Node Editor, which has the transforms from the Planar Tracker baked in. Unlike the Planar Tracker, the Planar Transform is intended for use with masks. With the Planar Transform created, there is no longer a use for the Planar Tracker, and it can be disconnected or deleted from the Node Editor.

We can insert this new node into the node tree to use it by holding down the Shift key and dragging the node over the connection between the Polygon node and the MatteControl node, dropping it when the connection highlights.



Inserting a PlanarTransform node by holding down the Shift key while dropping over a connection (left), and after inserting the PlanarTransform node (right)

With the new Planar Transform node inserted, the Polygon automatically moves to match the motion of the forehead that was tracked by the original PlanarTracker node, and it animates to follow along with the movement of the shot. At this point, we're finished!



The final painted image, along with the final node tree

Chapter 22

Using the Tracker Node

This chapter shows the many capabilities of the Tracker node in Fusion, starting with how trackers can be connected in your node trees, and finishing with the different tasks that can be performed.

Contents

Introduction to Tracking 557
Tracker Node Overview 558
Modes of the Tracker Node 558
Basic Tracker Node Operation 559
Connect to a Tracker's Background Input 559
Analyze the Image to be Tracked 559
Apply the Tracking Data 560
Viewing Tracking Data in the Spline Editor 563
Tracker Inspector Controls 564
Motion Tracking Workflow In Depth 565
Connect the Image to Track 565
Add Trackers 565
Refine the Search Area 568
Perform the Track Analysis 568
Tips for Choosing a Good Pattern 569
Using the Pattern Flipbooks 570
Using Adaptive Pattern Tracking 571
Dealing with Obscured Patterns 572
Dealing with Patterns That Leave the Frame 572
Setting Up Tracker Offsets 573
Stabilizing with the Tracker Node 574
Stabilization Using the Tracker Match Move Mode 574

Smoothing Motion 575
Using the Tracker Node for Match Moving 576
Simple Match Moving 576
Corner Positioning Operations 577
Perspective Positioning Operations 577
Connecting to Trackers' Operations 577
Steady Position 577
Steady Angle 577
Offset Position 578
Unsteady Position 578
Steady Size 578
Using the Outputs of a Tracker 578
Using the Tracker as a Modifier 580
Match Moving Text Example 583
Adding a Layer to Match Move 583
Setting Up Motion Tracking 583
A Simple Tracking Workflow 585
Connecting Motion Track Data to Match Move 588
Offsetting the Position of a Match Moved Image 590

Introduction to Tracking

Tracking is one of the most useful and essential techniques available to a compositor. It can be roughly defined as the creation of a motion path from analyzing a specific area in a clip over time. Fusion includes a variety of different tracking nodes that let you analyze different kinds of motion. Once you have tracked motion on a clip, you can then use the resulting data for stabilization, motion smoothing, matching the motion of one object to that of another, and a host of other essential tasks.

Types of tracking nodes in Fusion:

Tracker: Follows a relatively small, identifiable feature or pattern in a clip to derive a 2D motion path. This is sometimes referred to as point tracking.

Planar Tracker: Follows a flat, unvarying surface area in a clip to derive a 2 ½D motion path including perspective. A planar tracker is also more tolerant than a point tracker when some tracked pixels move offscreen or become obscured.

Camera Tracker: Tracks multiple points or patterns in a clip and performs a more sophisticated analysis by comparing those moving patterns. The result is a precise recreation of the live-action camera in virtual 3D space.

Each tracker type has its own chapter in this manual. This chapter covers the tracking techniques with the Tracker node.

Tracker Node Overview

The Tracker node is a single node that actually performs tracking, stabilizing, matching moving, and corner-pinning operations. Since the Tracker node can transform the foreground input, it can be used to generate tracks and then operate as a Merge in a match move or corner-pin setup. Or you can use it to produce tracking data only and then publish that data to other nodes in the Node Editor.

Modes of the Tracker Node

The Tracker node is an incredibly flexible tool often used multiple times in a composite to help with dozens of tasks. However, most of those tasks can be boiled down into just a few operations. The Tracker node has four operation modes that cover the majority of tracking situations.

Stabilizing

You can use one or more tracked patterns to remove all the motion from the sequence or to smooth out vibration and shakiness. When you use a single tracker pattern to stabilize, you stabilize only the X and Y position. Using multiple patterns together, you are able to stabilize position, rotation, and scaling.

Match Moving

The reverse of stabilizing is match moving, which detects position, rotation, and scaling in a clip using one or more patterns. Instead of removing that motion, it is applied to another image so that the two images can be composited together.

Corner Positioning

Corner positioning tracks four patterns that are then used to map the four corners of a new foreground into the background. This technique is generally used to replace signs or mobile phone screens. The Planar Tracker node is often a better first choice for these types of tracking tasks.

Perspective Positioning

Perspective positioning again tracks four patterns to identify the four corners of a rectangle. Each corner is then mapped to a corner of the image, rescaling and warping the image to remove all apparent perspective. The Planar Tracker node is often a better first choice for removing perspective from a clip.

Basic Tracker Node Operation

All tracking workflows consist of three fundamental steps.

1. Attach an image you want to track to the yellow background input of the Tracker node.
2. Set the tracking pattern and analyze the clip to create a path.
3. Apply the tracking data to stabilize, match move, corner pin, or remove perspective.

Connect to a Tracker's Background Input

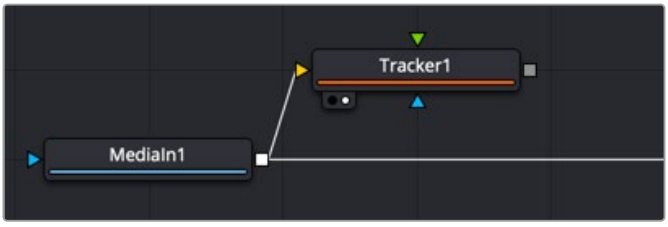
You start by connecting the output of the image you want to track to a Tracker node's background input. The Tracker node analyzes the image that's attached to its background input.

You can insert the Tracker node serially with other nodes if you intend to use the Tracker node itself to do a simple stabilization operation or if you want to use it to perform the function of a Merge node in a match move or corner-pin operation.



Tracker node connected serially so it can both track and transform the input images

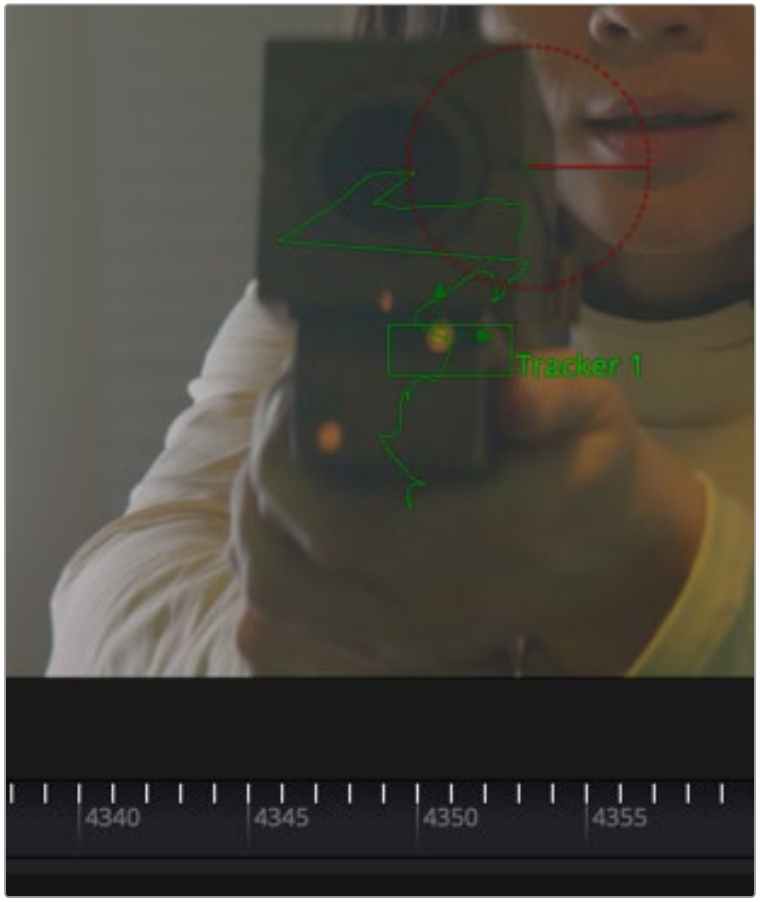
However, if you're just using a Tracker node to analyze data for use with multiple nodes elsewhere in the comp, you could choose to branch it and leave its output disconnected to indicate that Tracker node is a data repository. Please note that this is not necessary; serially connected Tracker nodes can be linked to multiple other nodes as well.



Tracker connected as a branch to indicate it is linked to other nodes and not used directly

Analyze the Image to be Tracked

After constructing the node tree and inserting the Tracker where you want, you can set up the tracker in the viewer. You identify one or more features in the image that you wish to track (referred to as patterns) by adding trackers (there's one by default) and positioning them using the onscreen controls in the viewer. After the Tracker node analyzes the clip, the resulting tracking data is stored within that Tracker node. Keyframes, one per frame, indicate the Tracked Center X and Y data that has been saved, while a motion path shows the path of tracked data in the viewer.



A motion path that indicates the tracked motion, and tic marks that indicate tracking data keyframes

Apply the Tracking Data

The resulting tracking data stored within the Tracker node is used to stabilize, match move, corner pin, or remove perspective in one of two ways.

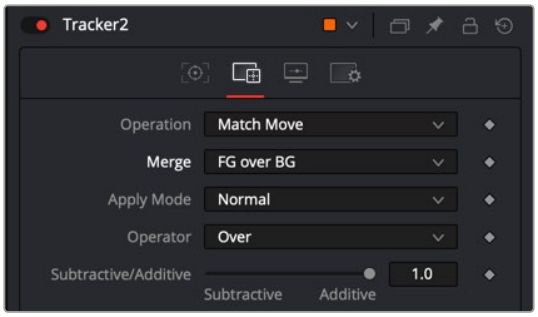
Method 1: Use the Tracker node to match move and merge

You can connect a foreground image to the Tracker node and apply the motion from the analyzed background image.



Using a Tracker node in line for a match move

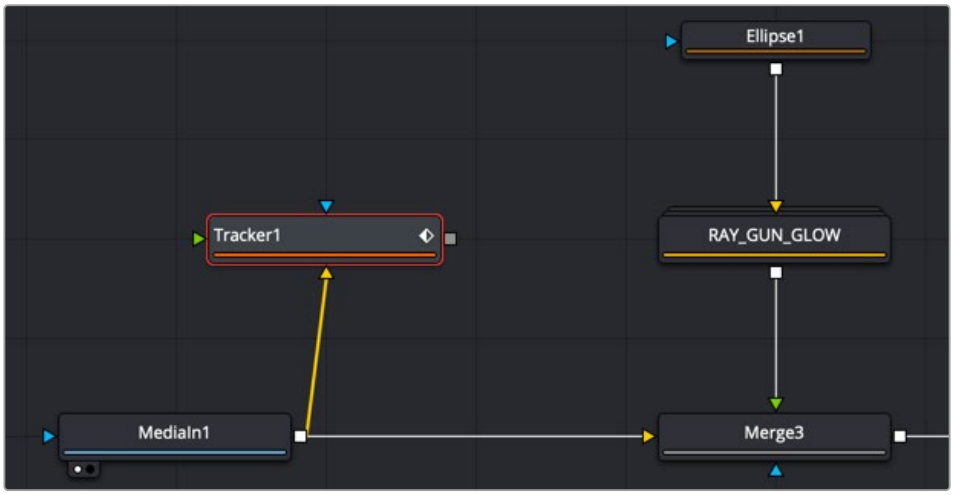
Setting the Operation parameter in the Operation tab in the Inspector to Match Move, Corner Position, or Perspective Position always applies the motion to the foreground input (if one is connected). This is an easy workflow for simple situations. In this scenario, you can use the Tracker node to replace a Merge node since Tracker nodes include all the same functionality as a Merge.



Using a Tracker node to do a match move and merge, all in one

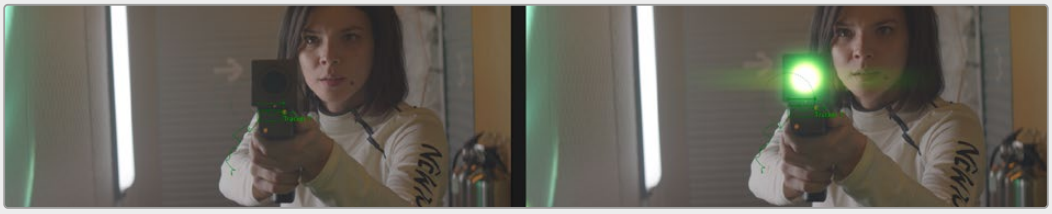
Method 2: Connect specific parameters to the Tracker Node

Alternatively, you can connect the tracking data from the Tracker node to the specific parameters of other nodes that will actually do the work, for instances where setting up a match move isn't just a matter of transforming a foreground image. Each Tracker node and each pattern within the Tracker node publishes its data for other nodes to use without directly linking to them in the node tree. For example, in the following node tree, an Ellipse node is being used to isolate a glow effect for the ray gun prop.



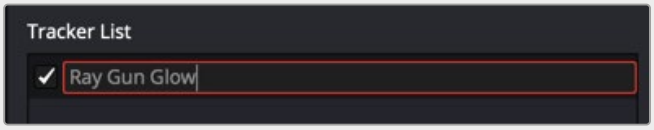
The Tracker set up as a branch and connected using the Connect To menu

The ellipse needs to follow the motion of the ray gun, so a Tracker node is used to analyze the movement of the gun tip so that tracking data can be used to animate the ellipse. The ellipse is not connected to the tracker directly via the foreground input but indirectly through the Connect To contextual menu.



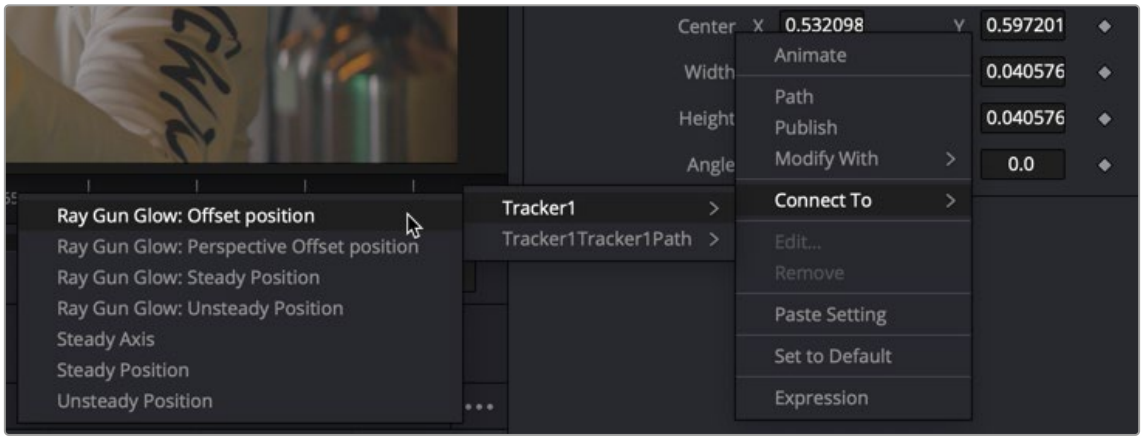
Applying the light of a ray gun by connecting tracking data to the center position of an Ellipse node

This is made easier by renaming the Tracker you created to something descriptive of what's being tracked.



You can rename trackers in the Tracker List by double-clicking them and typing descriptive text

Once the tip of the ray gun has been tracked, this tracking data is then connected to the Center parameter of an Ellipse node that's limiting a Glow effect by right-clicking the label of the Center parameter in the Inspector, and choosing Tracker1 > Ray Gun Glow: Offset position from the Connect to submenu of the contextual menu. All the data from every Tracker node in your node tree and every tracking pattern appears within this submenu, and since we named the Tracker, it's easy to find. Choosing Offset position because it will place the center of the ellipse directly over the path. However, it also gives us the flexibly to offset the ellipse if need be, using the offset controls in the Inspector.



Right-clicking a parameter's label lets you connect tracking data to it to animate it

You can connect the data from a Tracker node to any other node's parameter; however, you'll most typically connect track data to center, pivot, or corner X/Y style parameters. When you use tracking data this way, it's not necessary to connect the output of the Tracker node itself to anything else in your node tree; the data is passed from the Tracker to the Center parameter by linking it with the Connect To submenu.

Viewing Tracking Data in the Spline Editor

Tracking data can be seen in the viewer as a path or as a displacement spline in the Spline Editor. You can manipulate the tracking data in either place.

The Tracker uses a displacement spline by detail that indicates how far the tracking point is based on the original location. It is great for modifying velocity, but it doesn't tell you anything about direction. If you need to nudge a few points in a certain direction, you can convert the displacement spline to an X and Y coordinate spline.



Right-clicking a parameter's label lets you connect tracking data to it to animate it

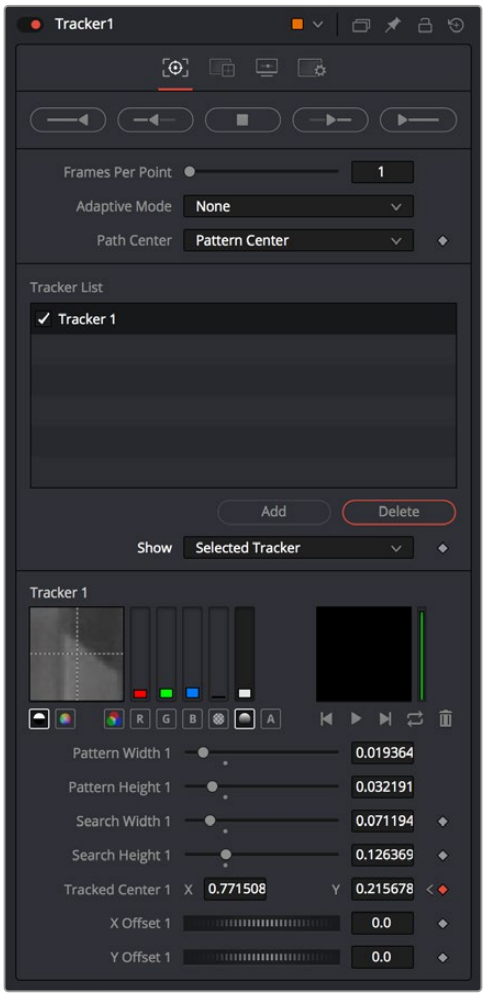
Right-click in the viewer to bring up a contextual menu. At the very bottom is a reference to the path the Tracker created, called Tracker1Tracker1Path:Polyline. Choosing it calls up a longer submenu where you can choose Convert to XY Path.

> For more information on Displacement Splines, *see Chapter 10, "Animating in Fusion's Spline Editor,"* in the Fusion Reference Manual.

Tracker Inspector Controls

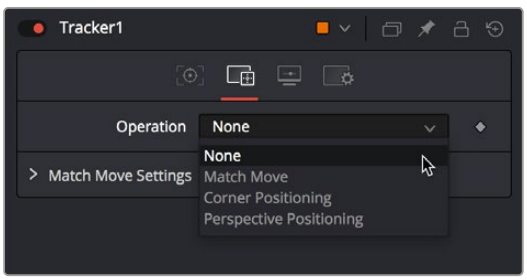
The layout of the Tracker node's tabs in the Inspector reflects this workflow. It's divided into three main Tracker tabs, as well as the common Settings tab.

- **The Tracker Control tab:** This is where you create onscreen trackers with which to target patterns, and where the controls appear that let you perform the required track analysis.

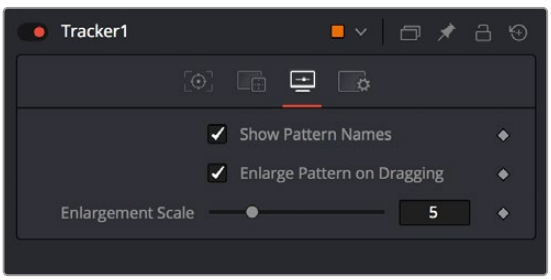


The Tracker Control tab

- **The Operations tab:** This is where you decide how the tracking data is used.



The Tracker Operations tab



The Tracker Display Options tab

- **The Display Options tab:** This is where you can customize how the onscreen controls look in the viewer.

Motion Tracking Workflow In Depth

Tracker nodes serve two purposes. They provide a method to analyze an object you want to follow, and they serve as a container for the resulting track data. This allows you to use one node for analysis and to pass on that analysis to any other node that requires it. Following is a more detailed breakdown of the tracking process.

Connect the Image to Track

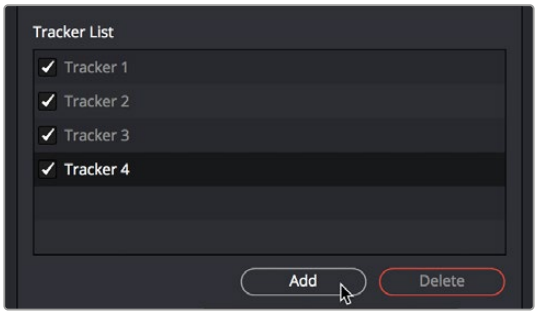
Regardless of whether you actually use the Tracker node itself to do anything with the tracking data, the image you want to track must be connected to the background input (yellow) of a Tracker node for there to be a successful analysis. While the Tracker node has a foreground input, it is initially ignored for purposes of tracking analysis, so even if you connect a foreground, the background is the only input used during the analysis process.



Connecting a MediaIn node's output to a Tracker node's background input

Add Trackers

Although each Tracker node starts with a single tracker pattern, a single node is capable of analyzing multiple tracking patterns that have been added to the Tracker List, enabling you to track multiple features of an image all at once for later use and to enable different kinds of transforms. Additional trackers can be added by clicking the Add button immediately above the Tracker List control.



To add an additional tracker, click the Add button in the Inspector

Multiple patterns are useful when stabilizing, match moving, or removing perspective from a clip. They also help to keep the Node Editor from becoming cluttered by collecting into a single node what would otherwise require several nodes.

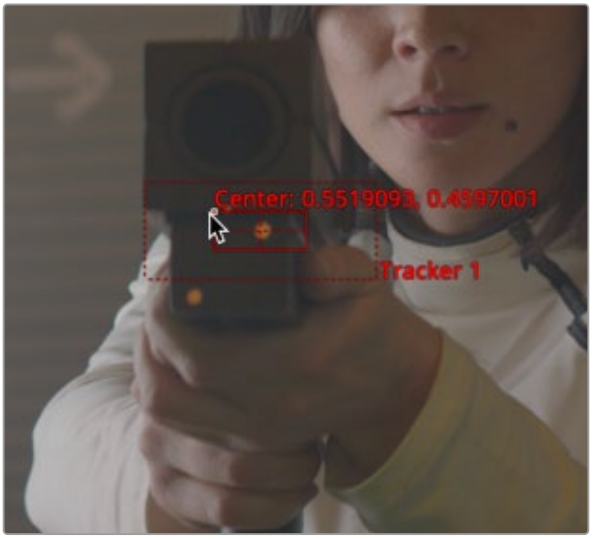
Working in the Tracker list:

- **To select a tracker:** Click the name of the Tracker you want to select.
- **To rename a tracker:** You can rename trackers to make it easier to reference them later. For example, if you're tracking a car door handle, you can name the Tracker "Car Handle" so it's easy to find later. To do so, just double-click the default name of the Tracker in the Tracker list, type a new one, and press Return.
- **To delete a tracker:** Select a tracker and click the Delete button.
- **To disable, suspend, or re-enable trackers:** Click the checkbox to the left of each tracker's name in the Tracker List. It's a three-way toggle that switches between Enabled, Suspended, and Disabled.
- **Enabled:** An Enabled tracker will re-track its pattern every time the track is performed. Its path data is available for use by other nodes, and the data is available for stabilization and corner positioning.
- **Suspended:** When the checkbox appears gray, it is Suspended. In this state, it does not re-track its pattern when the track is performed. The data is locked to prevent additional changes, but the data from the path is still available for other nodes. The data is also available for advanced tracking modes like stabilization and corner positioning.
- **Disabled:** A Disabled tracker does not create a path when tracking is performed. Its data is not available to other nodes or for advanced tracking operations like stabilization and corner positioning.

Position the Pattern Boxes

A pattern is the region of pixels that are targeted for tracking within an image using a pattern box. The pattern box is defined in the viewer by a rectangle when a tracker is active. A single Tracker node can have many pattern boxes, each targeting different patterns. Each tracked pattern will produce its own motion path.

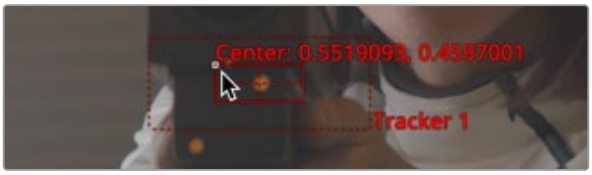
Clicking any part of a tracker's onscreen controls will select it. Selected pattern boxes are red, while deselected pattern boxes are green.



A pattern box positioned over an eye you want to track

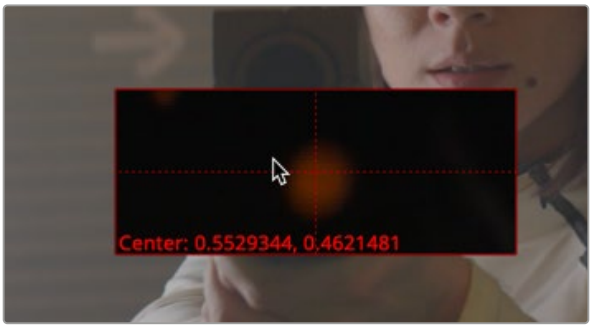
When you add a Tracker node to the Node Editor, you start with one pattern box displayed in the viewer as a small rectangle. When the cursor is placed over the pattern rectangle, the control expands and two rectangles appear. The outer rectangle has a dashed line, and the inner rectangle has a solid line. The outer rectangle is the search area, and the inner rectangle is the pattern.

If you need to select a new pattern, you can move the pattern box by dragging the small (and easily missed) handle at the top left of the inner pattern box.



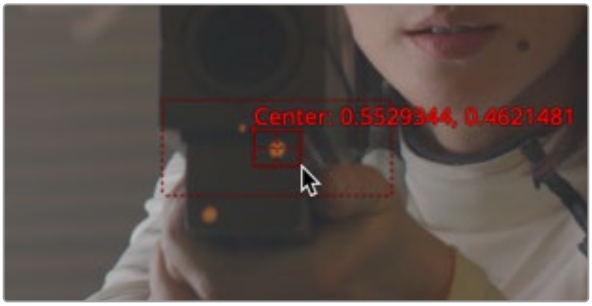
To move a pattern in the viewer, drag it from the upper-left corner

While moving the pattern box, an overlay pop-up appears, showing a zoomed version of the pixels contained within the rectangle to help you precisely position the pattern via the crosshairs within.



A zoomed view appears while you drag a pattern box to help you position it

The pattern rectangle can also be resized by dragging on the edges of the rectangle. You want to size the pattern box so that it fits the detail you want to track, and excludes area that doesn't matter. Ideally, you want to make sure that every pixel of the pattern you're tracking is on the same plane, and that no part of the pattern is actually an occluding edge that's in front of what you're really tracking. When you resize the pattern box, it resizes from the center, so one drag lets you create any rectangle you need.



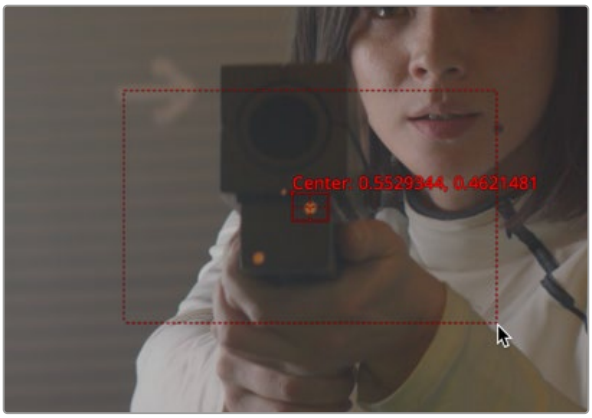
Resizing a pattern box to fit the tracking point on the ray gun

TIP: The magnified pattern box does not take viewer LUTs into account. When using Log content, it may make it easier to position the tracker if you temporarily insert a Brightness Contrast node between the source content and the yellow input of the tracker. You can use the Brightness Contrast node to temporarily increase the visibility of the region you are tracking.

Refine the Search Area

A second rectangle with a dotted border surrounds the pattern box. This is the search area. When progressing from one frame to another while tracking, the Tracker analyzes the region defined by the search area, which surrounds the last known tracker position in an attempt to relocate the pattern. The larger the search area, the better chance you have of successfully tracking fast moving objects, but the longer the track will take. However, there are some ways to optimize tracking for specific content.

For example, tracking a pattern that is moving quickly across the screen from left to right requires a wide search area but does not require a very tall one, since all movement is horizontal. If the search area is smaller than the movement of the pattern from one frame to the next, the Tracker will likely fail and start tracking the wrong pixels, so it's important to take the speed and direction of the motion into consideration when setting the search area.

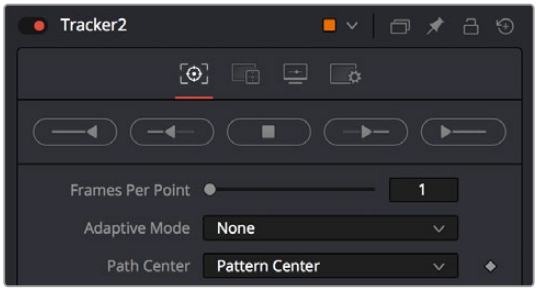


You can resize the search area by dragging the edges of the dotted outline

Perform the Track Analysis

Before you begin analyzing, you'll need to make sure you've set a render range in the Time Ruler that corresponds to the range of frames during which the pattern is visible. This may be an entire clip or only a small portion of that clip. Depending on the type of motion you're tracking, you may want to use the Adaptive Mode option to aid the analysis (see below for more details).

Once your options are set, you can use any of the tracking transport buttons at the top of the Inspector to start tracking. Once tracking has started, you cannot work in the Node Editor until it has completed.



The tracking transport buttons and analysis parameters

To begin tracking, do one of the following:

- Click the Track Reverse button to track from the very end of the render range.
- Click the Track Backward from Current Frame button to track backward from the current playhead position.
- Click the Track Forward button to track from the very start of the render range.
- Click the Track Forward from Current Frame button to track forward from the current playhead position.

Pattern tracking will stop automatically when it reaches the end of the render range (or the start when tracking backward), but you can also interrupt it and stop tracking at any time.

To stop tracking, do one of the following:

- Click the Stop Tracking button in the tracker transports.
- Click Stop Render at the bottom of the Fusion window.
- Press the Escape key.

When tracking is complete, the path will be connected to the pattern. The path from that pattern can now be connected to another node or used for more advanced operations like stabilization and corner positioning.

Once the track is complete, assuming it's good, you can use the various techniques in this chapter to use the track in your composition.

Tips for Choosing a Good Pattern

The Tracker works by searching each frame for the pixels contained in the pattern. In order for a track to be successful, a fairly high contrast and unique region of the image must be located in the footage. This process is known as pattern selection.

The first step in pattern selection is to review the footage to be tracked several times. Watch for candidate patterns that are visible through the entire range of frames, where the contrast is high and the shape of the pattern does not change over time. The more unique the pattern, the more likely the track is to be successful.

In addition to locating high contrast, defined patterns, watch for the frames where the pattern moves the most. Identifying the maximum range of a pattern's motion will help to determine the correct size for the pattern search area.

It is not uncommon to have a scene that requires the use of several different patterns to generate a single path. This most often occurs because the pattern moves out of frame or is temporarily obscured by another scene element. Combining patterns into a single pattern is described later in the chapter.

Selecting the Pattern's Image Channels

When a pattern of pixels is selected, the Tracker automatically selects the color channel used for tracking the pattern based on an analysis of each channel for contrast, clarity, and reliability. The channels selected are highlighted in the bars to the right of the Pattern display window in the node controls.



Highlighted channel bars indicate which channel is selected for tracking

You can override the automatic channel selection by clicking the buttons beneath the bars for each channel to determine the channel used for tracking.

You can choose any one of the color channels, the luminance channels, or the alpha channel to track a pattern.

When choosing a channel, the goal is to choose the cleanest, highest contrast channel for use in the track. Channels that contain large amounts of grain or noise should be avoided. Bright objects against dark backgrounds often track best using the luminance channel.

Selecting Patterns for Stabilization

Selecting patterns for stabilization can be a tricky business. The location of the pattern, when it is selected, is used to determine precisely how the image will be stabilized. At least two patterns are required to correct for rotation; using three patterns will correct for scaling, and more will usually improve the quality of the solution.

Try not to select just any potentially valid pattern in the sequence, as some patterns will make the solution worse rather than better. To help with your selection, use the following guidelines when selecting patterns for stabilization.

- Locate patterns at the same relative depth in the image. Objects further in the background will move in greater amounts compared to objects in the foreground due to perspective distortion. This can confuse the stabilization calculations, which do not compensate for depth.
- Locate patterns that are fixed in position relative to each other. Patterns should not be capable of moving with reference to each other. The four corners of a sign would be excellent candidates, while the faces of two different people in the scene would be extremely poor choices for patterns.

Using the Pattern Flipbooks

Each pattern has a pair of thumbnail windows shown in the Inspector. The left window shows the selected pattern, while the right window is updated during the track to show the actual pattern that has been acquired for each frame.

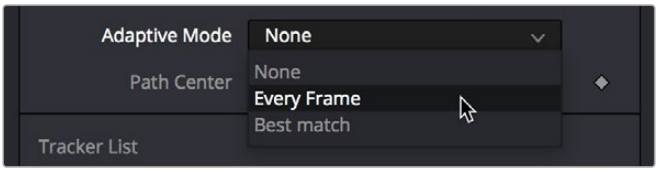


The Tracker Pattern Selection and Flipbook thumbnails

Each pattern that's stored is added to a Flipbook. Once the render is complete, you can play this Pattern Flipbook to help you evaluate the accuracy of the tracked path. If you notice any jumps in the frames, then you know something probably went wrong.

Using Adaptive Pattern Tracking

Even the most ideal pattern will usually undergo shifts in profile, lighting conditions, and other variables. These can adversely affect pattern recognition to the point that a pattern becomes unusable. The Tracker offers three modes of pattern acquisition during tracking that can help to correct these conditions. The modes can be set using the Adaptive Mode options in the Inspector.



The Adaptive Mode options

None

When the Adaptive mode is set to None, the pattern within the rectangle is acquired when the pattern is selected, and that becomes the only pattern used during the track.

Every Frame

When Every Frame is chosen, the pattern within the rectangle is acquired when the pattern is selected, and then reacquired at each frame. The pattern found at frame 1 is used in the search on frame 2, the pattern found on frame 2 is used to search frame 3, and so on. This method helps the Tracker adapt to changing conditions in the pattern.

Every Frame tracking is slower and can be prone to drifting from sub-pixel shifts in the pattern from frame to frame. Its use is therefore not recommended unless other methods fail.

Best Match Tracking

Best Match tracking works in much the same way as Every Frame tracking; however, it will not reacquire the pattern if the difference between the original pattern and the new one is too great. This helps to prevent cases where transient changes in the image cause the Tracker to become confused.

As a comparison between the two Adaptive modes, if a shadow passes over the tracker point, the Every Frame tracking mode may start tracking the shadow instead of the desired pattern. The Best Match mode would detect that the change from the previous frame's pattern was too extreme and would not grab a new pattern from that frame.

The Adaptive mode is applied to all active patterns while tracking. If you only want some patterns to use the Adaptive mode, disable all other patterns in the list before tracking.

Dealing with Obscured Patterns

Often, an otherwise ideal pattern can be temporarily obscured (occluded) or blocked from tracking—for example, when tracking a car that passes behind a telephone pole.

In these situations, you divide the render range up into two ranges, the range before the pattern is obscured and the range after the pattern becomes visible again. After tracking the two ranges individually, the Tracker will automatically interpolate between the end of the first range and the start of the second.

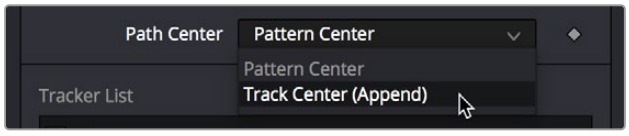
If you need to edit the resulting motion path to account for any non-linear motion that takes place between the two tracked ranges, you can select the track path to expose a Node toolbar with controls for adjusting the control points on this path. For example, you can choose Insert and Modify mode to insert points in the non-tracked range to compensate for any nonlinear motion in the tracked pattern.



Tools for modifying tracker paths in the Node toolbar of the viewer

Dealing with Patterns That Leave the Frame

There are two options when a tracker leaves the frame. If the pattern re-enters the frame, you can treat it like an obscured pattern. If the pattern does not re-enter the frame, or it is undesirable to hand track portions of the movement, you can use the Track Center (Append) mode to select a new pattern.



The Track Center (Append) mode pop-up menu

The Track Center (Append) mode selects a new pattern that will continue to add keyframes to the existing path. The offset between the old pattern and the new pattern is automatically calculated to create one continuous path.

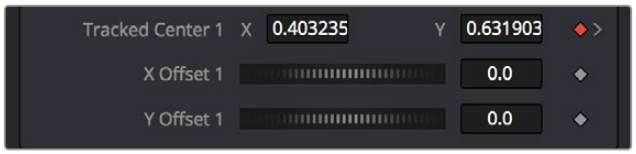
To use the Track Center (Append) mode, do the following:

1. When the pattern has become untrackable for some reason, stop analysis and move the playhead to the last frame that tracked successfully.
2. Choose Track Center (Append) from the Path Center pop-up menu in the Inspector.
3. Now, drag the Pattern selector to a new pattern that can be tracked from that point onward.
4. Restart tracking from the current frame.

When selecting a pattern to use in appending to an existing path, a pattern that is close to the old pattern and at the same apparent depth in the frame generates the best results. The further away the new pattern is, the more likely it is that the difference in perspective and axial rotation will reduce accuracy of the tracked result.

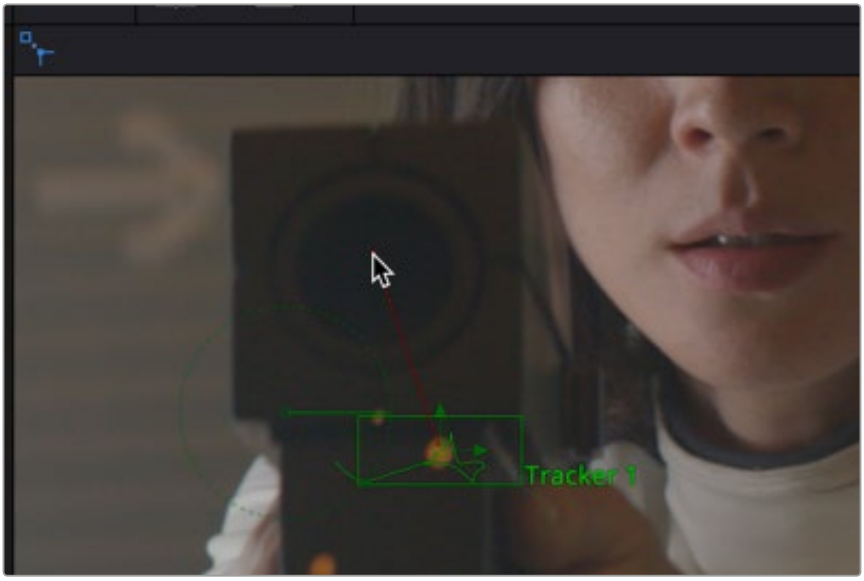
Setting Up Tracker Offsets

Often, it's impossible to track the thing you want to apply an effect to. For example, the only pattern available for an accurate track is a button on an actor's sleeve. However, the effect requires the person's hand to be glowing. To cause the glow's effect mask to be centered on the actor's hand, it's necessary to use the Tracker Offset control.



The Tracker Offset controls in the Inspector

The X and Y Offset controls allow for constant or animated positional offsets to be created relative to the actual Tracker's pattern center. The position of the offset in the viewer will be shown by a dashed line running from the pattern center to the offset position. You can also adjust the offset in the viewer using the Tracker Offset button. Clicking the button enables you to reposition the path while keeping the Tracker pattern in place.



The Tracker Offset tool in the Node toolbar of the viewer; a track of the orange dot is being offset to the center of the ray gun

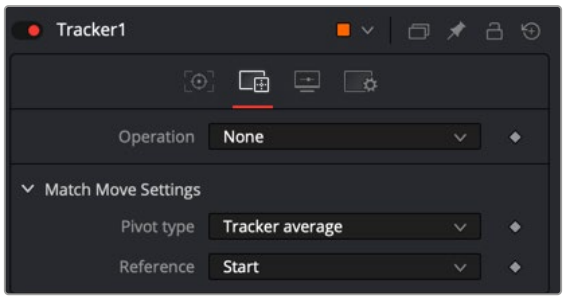
Once an offset for a pattern is set, you can connect other positional controls to the Tracker's Offset menu using the Connect To > Tracker: Offset Position option in the control's contextual menu. The path created during the track remains fixed to the center of the pattern.

Stabilizing with the Tracker Node

When a Tracker node is set to Match Move in the Operations tab, it is capable of a variety of functions. Applying the motion from the background clip to the foreground clip is the obvious functionality. However, the Match Move operation is also used for stabilizing footage to to either completely remove motion from the scene or smooth existing motion.

Here are some common scenarios for stabilization that are handled when the Tracker is set to Match Move.

- A sequence that should be steady has vibrations or undesirable movement.
- A sequence that requires a smooth camera move suffers from jarring.



The Tracker Operation tab Match Move set to BG Only for stabilization

Stabilization Using the Tracker Match Move Mode

Stabilizing motion completely removes the appearance of motion from the image. The motion from frame to frame is calculated, and the contents of the frame are transformed to return the image to a reference position. This position can be either the start or end of the sequence or a manually selected frame from the sequence.

Stabilization can correct for position with as little as one pattern. Two or more patterns are required to correct for rotation or scaling within the image.

When the Operation menu is set to Match Move, choosing BG only from the Merge operation menu stabilizes the background (yellow input) clip. Only the controls that are applicable for stabilization operations will appear in the Operation tab.

Several of the stabilization controls are always available, collected under the Match Move Settings disclosure button. These controls are available at all times because the Steady and Unsteady positions of a tracker are always published. This makes them available for connection by other controls, even when the Tracker's operation is not set to match moving.



The Match Move settings

Merge

The Merge menu determines to which input connection the Tracking data is applied. When stabilizing an image to remove all motion, or smooth the motion, the Merge button must be set to BG Only.

Edges

The Edges menu determines whether the edges of an image that leave the visible frame are cropped, duplicated, or wrapped when the stabilization is applied. Wrapping edges is often desirable for some methods of match moving, although rarely when stabilizing the image for any other purpose.

For more information on the controls, *see Chapter 59, "Tracking Nodes,"* in the Fusion Reference Manual.

Position/Rotation/Scaling

Use the Position, Rotation, and Scaling checkboxes to select what aspects of the motion are corrected.

Match Move Settings

Options for the Match Move settings include Pivot and Reference.

Pivot Type

The Pivot Type for the stabilization is used to calculate the axis of rotation and scaling calculations. This is usually the average of the combined pattern centers but may be changed to the position of a single tracker or a manually selected position.

Reference

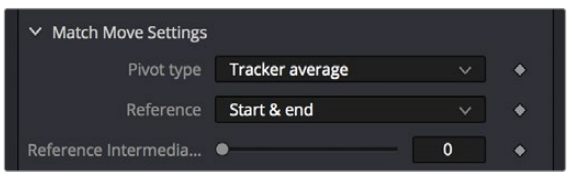
The Reference controls establish whether the image is stabilized to the first frame in the sequence, the last frame, or to a manually selected frame. Any deviation from this reference by the tracked patterns is transformed back to this ideal frame.

As a general rule, when tracking to remove all motion from a clip, set the Merge mode to BG Only, the Pivot Type to Tracker Average or Selected Tracker, and the Reference control to Start, End, or Select Time.

Smoothing Motion

When confronted with an image sequence with erratic or jerky camera motion, instead of trying to remove all movement from the shot, you often need to preserve the original camera movement while losing the erratic motion.

The Start & End reference option is designed for this technique. Instead of stabilizing to a reference frame, the tracked path is simplified. The position of each pattern is evaluated from the start of the path and the end of the path along with intervening points. The result is smooth motion that replaces the existing unsteady move.



The Reference Intermediate Points slider is displayed when Start & End is selected to enable the smoothing of motion

To preserve some of the curvature of the original camera motion, you can increase the value of the Reference Intermediate Points slider that appears when the Start & End reference mode is selected.

When tracking to create smooth camera motion, ensure that the Start & End reference mode is enabled and set the Merge mode to BG Only. It is recommended to leave the Pivot Type control set to Tracker Average.

Using the Tracker Node for Match Moving

A simple match moving example is shown at the beginning of this chapter, but this section presents additional details that you may not have been aware of. Examples of match moving include:

- A static CG element must be believably added to a moving sequence.
- Two sequences with different motions must be composited together.

Some clips may need to be stabilized so that an element from another source can be added to the shot. After the element or effect has been composited, the stabilization should be removed to make the shot look natural again.

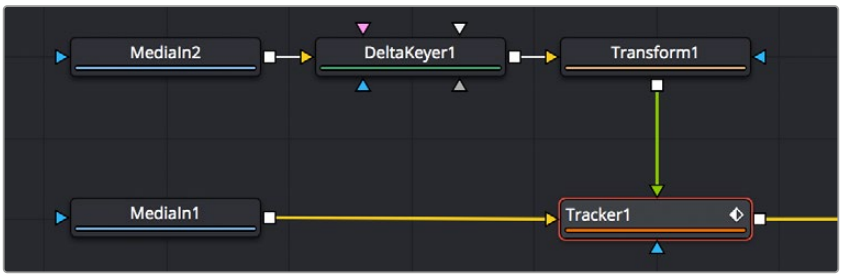
Simple Match Moving

Match moving essentially applies the movement from the tracked clip to another clip. There are two ways to perform match moving. One method involves connecting other nodes, such as Transform or Merge, to a Tracker's outputs. The other method is to stabilize an image by trying to remove all motion, but instead of setting the Merge menu to BG Only, set it to FG Over BG, FG Only, or in rare occasions, BG Over FG.



Set the Merge menu to BG Only, FG Over BG, or BG Over FG

When using this Merge menu, you connect a foreground image to the Tracker node's input connection in the Node Editor.



Connect a foreground image to the Tracker's foreground input

Enabling the FG Only mode will apply the motion from the background to the foreground, and the Tracker will only output the modified FG image. This result can later be merged over the original, allowing further modifications of the foreground to be applied using other nodes before merging the result over the background clip.

Corner Positioning Operations

The Corner Positioning operation maps the four corners of a foreground image to four patterns within the Tracker. This operation, or technique, is most commonly used for sign replacements.

The Corner Positioning operation of the Tracker requires the presence of a minimum of four patterns. If this operation mode is selected and there are not four patterns set up in the Tracker already, additional patterns will automatically be added to bring the total up to four.

When this mode is enabled, a set of drop-down boxes will appear to select which tracker relates to each corner of the rectangle. It has no effect when the Merge control option is set to BG Only.

Perspective Positioning Operations

The Perspective Positioning operation is used to remove perspective from a foreground image or apply the perspective from one sequence to another. This can be useful when you need to paint out an area that is distorted by perspective. Removing the perspective flattens the images for painting, and then another tracker adds the perspective back.

The Perspective Positioning operation of the Tracker requires the presence of a minimum of four patterns. If this operation mode is selected and there are not four patterns set up in the Tracker already, additional patterns will automatically be added to bring the total up to four.

When this mode is enabled, a set of drop-down boxes will appear to select which tracker relates to each corner of the rectangle. It has no effect when the Merge control option is set to BG Only.

Connecting to Trackers' Operations

One of the most common applications for a tracked pattern is using the tracked position or path to drive the position of another node's parameters. For example, tracking an eye in order to color correct the eye to blue using an effect mask. You start off by tracking the eye, and then create a color corrector with the desired settings. You create a mask in the shape of the eye and connect the Tracker's position to the Center of the mask.

In addition to the path (called Offset Position), each pattern in a tracker publishes four other values for use as connections that are available to other nodes in the Node Editor.

You connect a node's position parameters to a tracker by selecting the connection type from the controls contextual menu (for example, Transform 1: Center > Connect To > Tracker 1 > Offset Position).

There are five connection types automatically published by the tracker to connect to a position parameter in another node.

Steady Position

Steady Position can be used to stabilize footage in both X and/or Y to remove camera shake and other unwanted movement. The connection inverts the output of the tracked pattern's motion. When you connect a Center parameter to the Steady Position of the Tracker, it will be placed at 0.5/0.5 (the center of the screen) by default at frame 1. You can change this using the Reference mode in the Tracker's Operation tab.

Steady Angle

The Steady Angle mode can be used to stabilize footage in both X and/or Y to remove camera shake and other unwanted movement. When you connect a control, for example the Angle of a Transform, to the Steady Angle of the Tracker, it will be placed at 0 degrees by default at frame 1. This can be

changed by means of the Reference mode in the Tracker's Operation tab. From there on, the resulting motion of the Steady Angle mode will rotate into the opposite direction of the original motion.

So if the angle at frame 10 is 15 degrees, the result of the Steady Angle will be -15 degrees.

To use Steady Angle, you need at least two tracked patterns in your tracker. With just one point, you can only apply (Un)Steady Position.

Offset Position

An Offset Position is essentially the path generated by the tracker. It is the one you select when you want an object to follow the path. It is available for each single tracker in the Tracker node and refers to that single tracker only. When you connect the Center X and Y parameters to the offset position of the Tracker, the node's center will follow exactly the path of that tracker. Connecting to single trackers is always useful when you want to match elements with object motion in your footage. For example, you could track a hand of your actor and attach a ball to the Tracker's offset position, so that the ball follows the exact motion of the hand. Or you could track an element that needs rotoscoping and attach the mask's center to the Tracker's offset position.

Unsteady Position

After using the Steady Position, the Unsteady Position is used to reintroduce the original movement on an image after an effect or new layer has been added. The resulting motion from Unsteady Position is basically an offset in the same direction as the original motion.

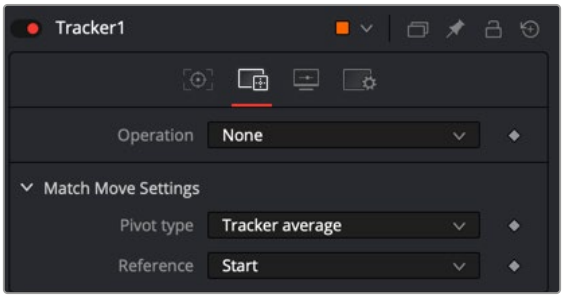
Steady Size

The Steady Size connection outputs the inverse of the tracked pattern's scale. When you connect a parameter, for example the Size of a Transform, to the Steady Size of the Tracker, it will be placed with a Size of 1 (i.e., the original size) by default at frame 1. This can be changed by means of the Reference mode in the Tracker's Operation tab. The resulting size of the Steady Size mode will then counteract the size changes of the original motion. So if the actual size at frame 10 is 1.15, the result of the Steady Size will be 1 - (1.15 - 1) = 0.85.

To use Steady Size, you need at least two tracked patterns in your tracker. With just one point, you can only apply (Un)Steady Position.

Using the Outputs of a Tracker

The tracker outputs described above are published by each tracker pattern created in the Tracker node. Each Tracker node itself also publishes a Steady Position, Angle, Size and an Unsteady Position. The values of these Tracker node outputs are calculated using all the patterns in that tracker, as configured by the Match Move Settings controls in the Tracker's Operation tab.



The Match Move settings determine the reference frame for stabilization and the point around which stabilization pivots

Rather than using the Tracker node to perform the Merge operation, an alternative and common way to use these published outputs is to create a match move by connecting the outputs to multiple nodes. A tracker is used to track a pattern, and then that data can be connected to multiple other nodes using the Connect To submenu.

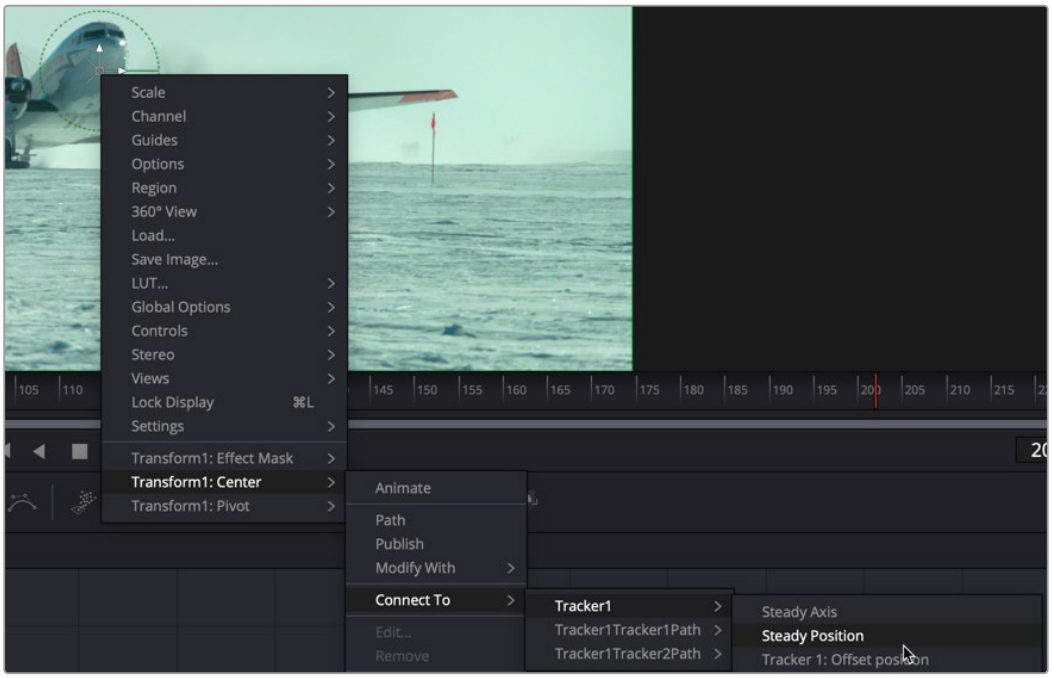
As an example, to use the Connect To menu to perform a match move, do the following:

1. Track the background clip using at least two tracking patterns in the tracker.
2. In a different branch, add a Transform node to the background clip.



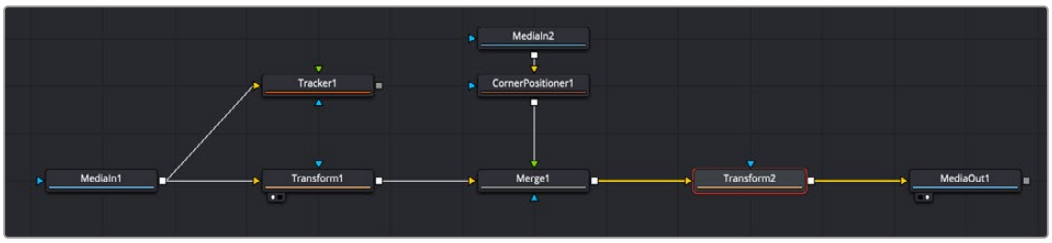
The background branched to go to a Tracker and a Transform

3. Right-click over the Transform's Center and choose Connect to > Tracker1 > Steady Position.



The tracker publishes its output for other nodes to connect to, as done here to stabilize the clip

4. Connect the foreground to a corner-positioned node, so you can position the corners of the foreground appropriately over the background.
5. Add another Transform node to the Node Editor after the Merge.



A second Transform after the Merge is used to add back in the original motion with Unsteady Poisition

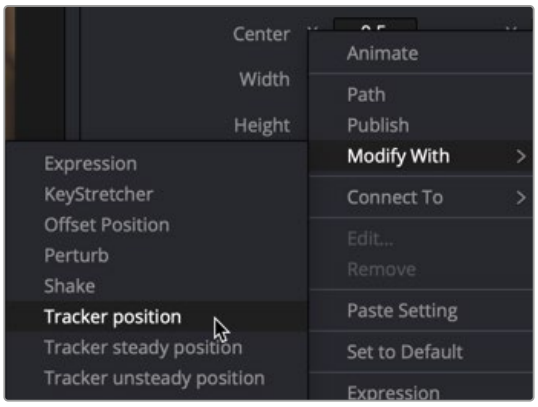
6 Connect the new Transform's Center to the Tracker's Unsteady Position. The image will be restored to its original state with the additional effect included.

To better understand how this works, imagine a pattern that is selected at frame 1, at position 0.5, 0.5. The pattern does not move on frame 2, so its position is still 0.5, 0.5. On the third frame, it moves 10 percent of the image's width to the right. Now its position is 0.6, 0.5.

If a transform center is connected to the Steady Position output of the Tracker, the Transform node's center is 0.5, 0.5 on the first and second frames because there has been no change. On frame 3, the center moves to 0.4, 0.5. This is the inverse of the horizontal motion that was tracked in the pattern, moving the image slightly to the right by 10 percent of the image width to counteract the movement and return the pattern of pixels back to where they were found.

Using the Tracker as a Modifier

Another technique for adding a tracker directly to a control is to add it as a modifier. Choosing the Tracker from the Modify With contextual menu does not use a Tracker node; it adds a modifier in the Inspector with a set of parameters almost identical to those found in the Tracker node itself. The benefit here is that the object that you want to follow the tracked path is automatically connected to the tracker modifier when you apply it.



Applying the tracker as a modifier

The differences between a Tracker modifier and a Tracker node are as follows:

- The Tracker modifier can only track a single pattern.
- A source image must be set for the Tracker modifier.

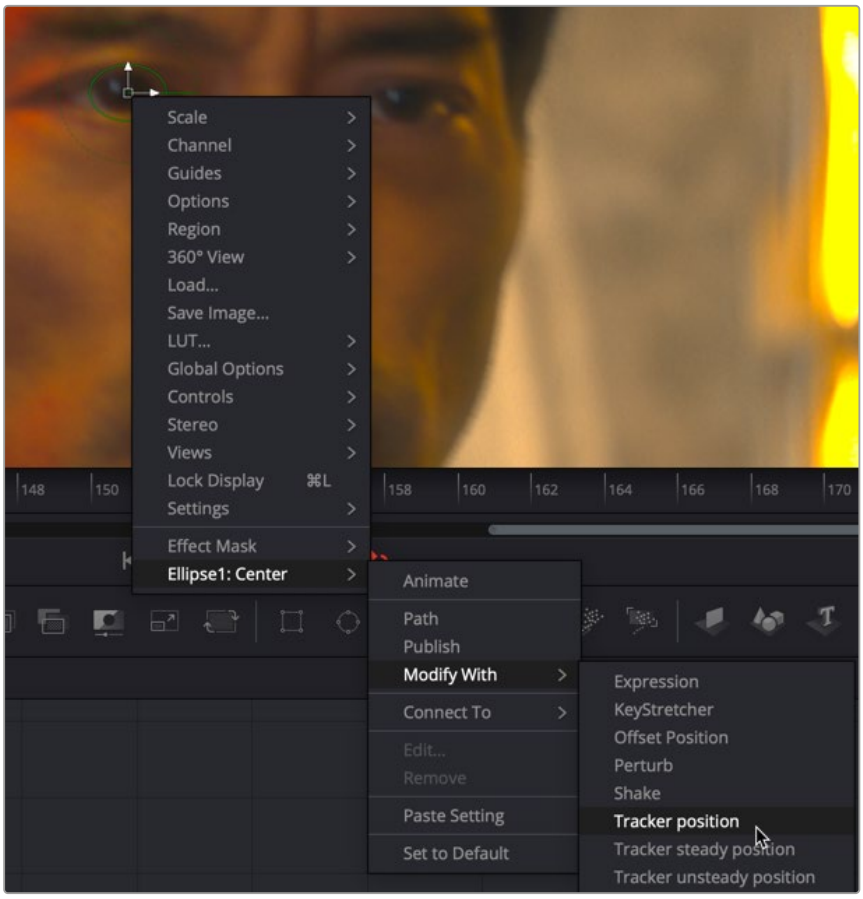
The Tracker modifier can only output a single value and cannot be used for complex stabilization procedures, but it is a nice quick way to apply a tracker to a point that you need to follow.

As an example, to apply the Tracker as a modifier, do the following:

Imagine that you needed to track an actor's eyes so that an unearthly, alien glow could be applied to the eyes.

1 Add an ellipse mask node to cover an actor's eye.

2 In the Inspector, right-click on the mask's Center parameter and from the contextual menu choose Ellipse1 Center > Modify With > Tracker Position.



Right-click over any Center Coordinate Control and choose Modify With > Tracker Position to add a Tracker modifier

3 Click the Modifiers tab in the Inspector and drag the MediaIn1 node that you want to track into the Tracker Source field.



The Modifiers tab includes the tracking controls and a Tracker Source field to indicate which node to use for tracking

4 Click the Track Forward button to begin tracking the person's eye.

5 Insert a Soft Glow node directly after the MediaIn and connect the Ellipse Mask to the white Glow Mask input.



A Tracker modifier applied to the Ellipse to create a green glow on an actor's pupil

You can set a different source image for the Tracker modifier by typing in the name of the node or dragging and dropping the node from the Node Editor into the Tracker Source field control. If you have a node (let's call it node#1) connected to the node that contains the modifier (let's call it node#2), the source image for the Tracker modifier will automatically the node #1

For more information on the Tracking parameters, *see Chapter 59, "Tracking Nodes,"* in the Fusion Reference Manual.

Match Moving Text Example

This example takes you through a complete motion tracking task, and shows how you can create a very simple match-moving effect using the Tracker node, which is the Swiss army knife of trackers in Fusion.

Adding a Layer to Match Move

In the example composition, we have a Text1 node that's creating a "Switzerland" title that's composited over a drone shot flying over and around a mountain bridge. With the Text1 node selected, the onscreen controls that let you position the text it's generating are visible in the viewer, and the text is positioned where we'd like it to start. Note that, with the Text node selected, even the part of the text that's offscreen can still be seen as an outline showing us where it is.

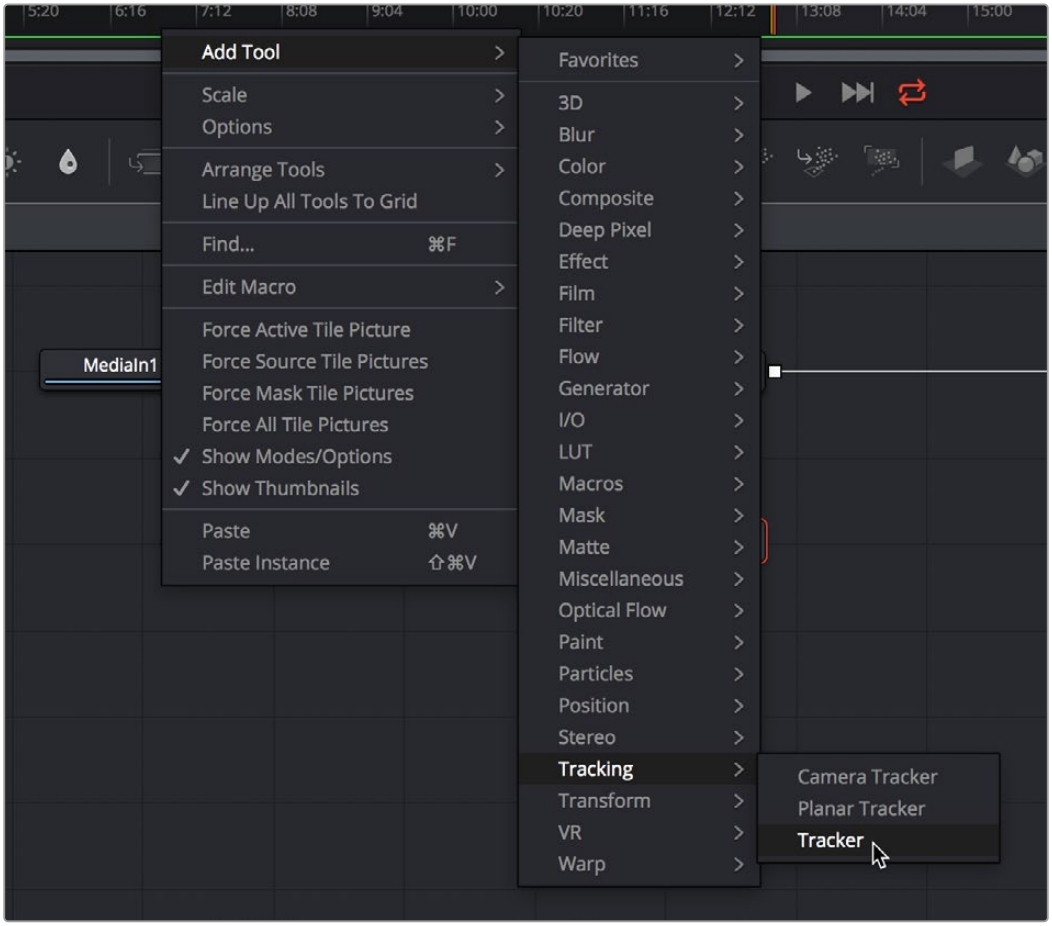


Some text superimposed against a background, ready to track

Our goal for this composition is to motion track the background image so that the text moves along with the scene as the camera flies along.

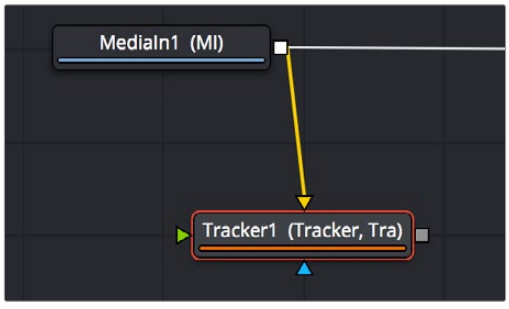
Setting Up Motion Tracking

To set up for the motion track, we'll begin by creating a disconnected Tracker node, using another method other than those seen previously. Right-click anywhere in the background of the Node Editor (preferably where you want the new node to appear), and choose Add Tool > Tracking > Tracker from the contextual menu to create a new Tracker1 node underneath the MediaIn node (or Loader node if you are using Fusion Studio).



Creating a new node using the Node Editor contextual menu

Next, we'll drag a connection from the MediaIn1 node to the Tracker1 node to automatically connect the source clip to the Tracker1 background input. This branches the output from the MediaIn1 node to the Tracker node so that the Tracker1 node processes the image separately from the rest of the node tree. This is not required, but it's a nice organizational way to see that the Tracker node is doing an analysis that must be referred to in a way other than a "physical" connection.



Branching a Tracker node to use to analyze an image

A Simple Tracking Workflow

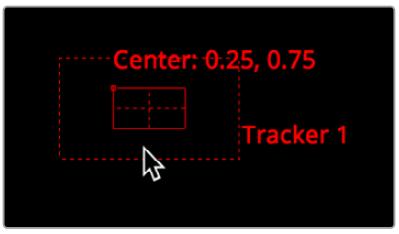
The Tracker node is the simplest tracking operation the Fusion page has, and while there are several ways of using it, an extremely common workflow is to use the Tracker node controls to analyze the motion of a subject in the frame with motion you want to follow, and then use the resulting motion path data by "connecting" it to the Center parameter of another node that's capable of transforming the image you want to match move.

Positioning the Tracker Onscreen Control

When the Tracker node is selected, a single green box appears in the viewer, which is the default onscreen control for the first default tracker that node contains (seen in the Tracker List of the Inspector controls). Keep in mind that you only see onscreen controls for nodes that are selected, so if you don't see the onscreen tracker controls, you know you need to select the tracker you want to work with. Loading the tracker you want to work with into the viewer is also the safest way to make sure you're positioning the controls correctly relative to the actual image that you're tracking.

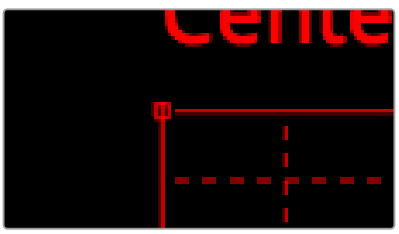
If you position your pointer over this box, the entire onscreen control for that tracker appears, and if you click the onscreen control to select that tracker, it turns red. As with so many other tracker interfaces you've likely used, this consists of two boxes with various handles for moving and resizing them:

- The inner box is the "pattern box," which identifies the "pattern" in the image you're tracking and want to follow the motion of. The pattern box has a tiny handle at its upper-left corner that you use to drag the box to overlap whatever you want to track. You can also resize this box by dragging any corner, or you can squish or stretch the box by dragging any edge to make the box better fit the size of the pattern you're trying to track. The center position of the tracker is indicated via X and Y coordinates.
- The outer box is the "search box," which identifies how much of the image the tracker needs to analyze to follow the motion of the pattern. If you have a slow-moving image, then the default search box size is probably fine. However, if you have a fast-moving image, you may need to resize the search box (using the same kind of corner and side handles) to search a larger area, at the expense of a longer analysis. The name of that tracker is shown at the bottom right of the search box.



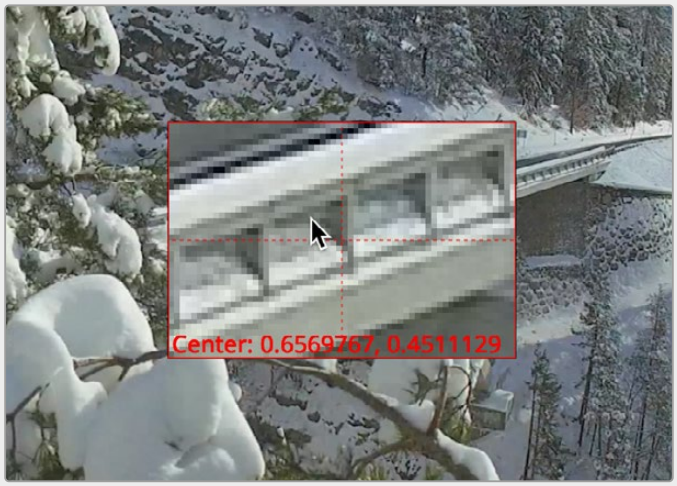
The onscreen controls of a selected tracker seen in isolation

It's worth saying a second time that the handle for moving a tracker's onscreen control is a tiny dot at the upper-left corner of the inner pattern box. It's really easy to miss if you're new to Fusion. You must click on this dot to drag the tracker around.



The handle for dragging the tracker boxes to move them around

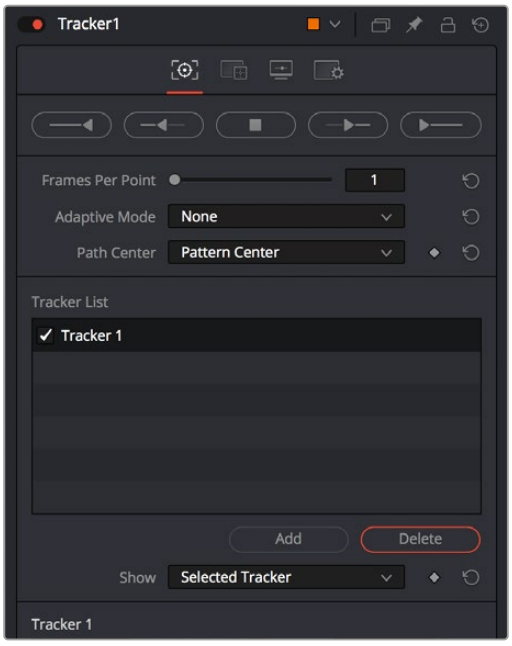
In this example, we'll drag the onscreen control so the pattern box overlaps a section of the bridge right over the leftmost support. As we drag the onscreen control, we see a zoomed-in representation of the part of the image we're dragging over to help us position the tracker with greater precision. For this example, the default sizes of the pattern and search box are fine as is.



The zoomed-in preview that helps you position the pattern box as you drag it

Using the Tracker's Inspector Controls to Perform the Analysis

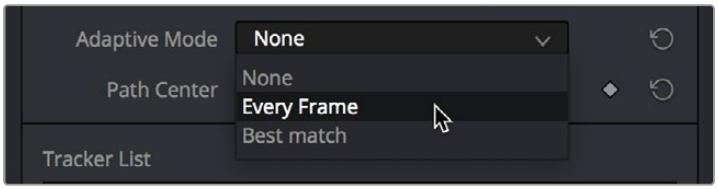
At this point, let's look at the Tracker node's controls in the Inspector. There are a lot of controls, but for this simple example we only care about the main Tracker panel, with the tracking analysis buttons at the top, the tracking options below those, and the Tracker List underneath those. The Tracker List also has buttons for adding and deleting trackers; you have the option of adding multiple trackers that can be analyzed all at once for different workflows, but we don't need that for now.



Tracker Inspector controls, with the tracking analysis buttons at top, the tracker options in the middle, and the Tracker List below

Additional controls over each tracker and the image channels being analyzed appear at the bottom, along with offset controls for each tracker, but we don't need those now (at least not yet).

Again, this track is so simple that we don't need to change the default behaviors that much, but because the drone is flying in a circular pattern, the shape of the pattern area is changing as the clip plays. Fortunately, we can choose Every Frame from the Adaptive Mode menu to instruct the tracker to update the pattern being matched at every frame of the analysis, to account for this.



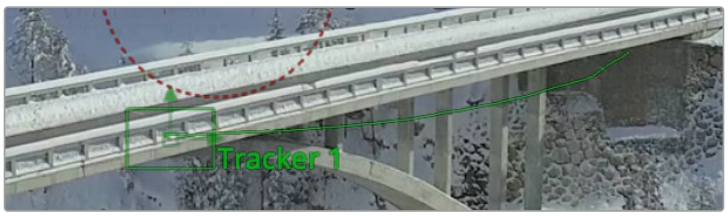
Changing the Adaptive Mode of the Tracker node to Every Frame to account for the camera's shift of perspective

Now, we just need to use the tracker analysis buttons at the top to begin the analysis. These buttons work like transport controls, letting you start and stop analysis as necessary to deal with problem tracks in various ways. Keep in mind that the first and last buttons, Track from Last Frame and Track from First Frame, always begin a track at the last or first frame of the composition, regardless of the playhead's current position, so make sure you've placed your tracker onscreen controls appropriately at the last or first frame.



The analysis buttons, left to right: Track from Last Frame, Track Backward, Stop Tracking, Track Forward, Track from First Frame

For now, clicking the Track from Beginning button will analyze the entire range of this clip, from the first frame to the last. A dialog lets you know when the analysis is completed, and clicking the OK button dismisses it so you can see the nice clean motion path that results.



The analyzed motion path resulting from tracking a section of the bridge as the camera flies past

Viewing Motion Track Data in the Spline Editor

This is not a necessary part of the tracking workflow, but if you have an otherwise nice track with a few bumps in it, you can view the motion tracking data in the Spline Editor by viewing that tracker's Displacement parameter curve. This curve is editable, so you can massage your tracking data in a variety of ways, if necessary.



Viewing motion tracking analysis data in the Spline Editor

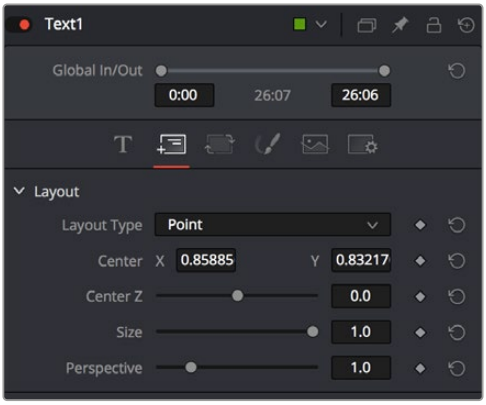
Connecting Motion Track Data to Match Move

Now that we have a successful analysis, it's time to use it to create the Match Move effect. To make this process easier, we'll double-click the tracker's name in the Tracker List of the Inspector, and enter a new name that's easier to keep track of (heh). Adding your own names make that tracker easier to find in subsequent contextual menus and lets you keep track of which trackers are following which subjects as you work on increasingly complex compositions.



Renaming a tracker to make it easier to find

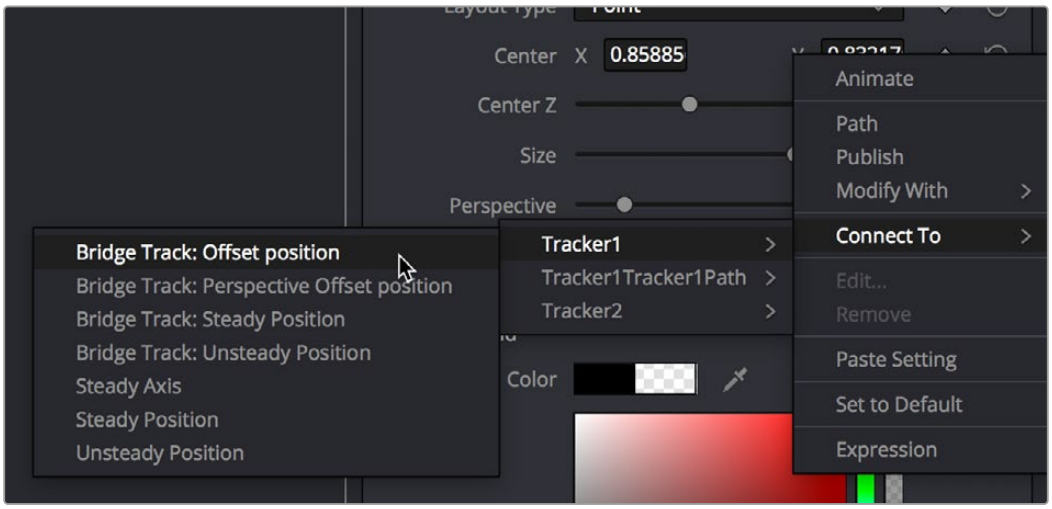
Now it's time to connect the track we've just made to the text in order to start it in motion. After loading the Merge1 node into the viewer to see the text in context with the overall composite we're creating, we'll select the Text1 node to open its parameters in the Inspector, and click the Layout panel icon (second button from the left) to expose the Layout controls, which are the text-specific transform controls used to position the text object in the frame. These are the controls that are manipulated when you use the Text node onscreen controls for repositioning or rotating text.



The Layout controls for a Text node, in the Layout panel

The Center X and Y parameters, while individually adjustable, also function as a single target for purposes of connecting to tracking to quickly set up match moving animation. You set this up via the contextual menu that appears when you right-click any parameter in the Inspector, which contains a variety of commands for adding keyframing, modifiers, expressions, and other automated methods of animation including connecting to motion tracking.

If we right-click anywhere on the line of controls for Center X and Y, we can choose Connect To > Tracker1 > Bridge Track: Offset position from the contextual menu, which connects this parameter to the tracking data we analyzed earlier.



Connecting the Center X and Y parameter to the Bridge Track: Offset position motion path we analyzed

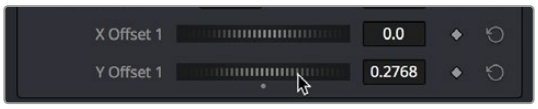
Immediately, the text moves so that the center position coincides with the center of the tracked motion path at that frame. This lets us know the center of the text is being match moved to the motion track path.



The text now aligns with the motion track coordinate

Offsetting the Position of a Match Moved Image

In fact, we want to offset the match-moved text, so it's higher up in the frame. To do this, we select the Tracker1 node again and use the Y Offset 1 dial control to move the text up, since now any changes we make to the Bridge Track dataset now apply to the center of the text that's connected to it.



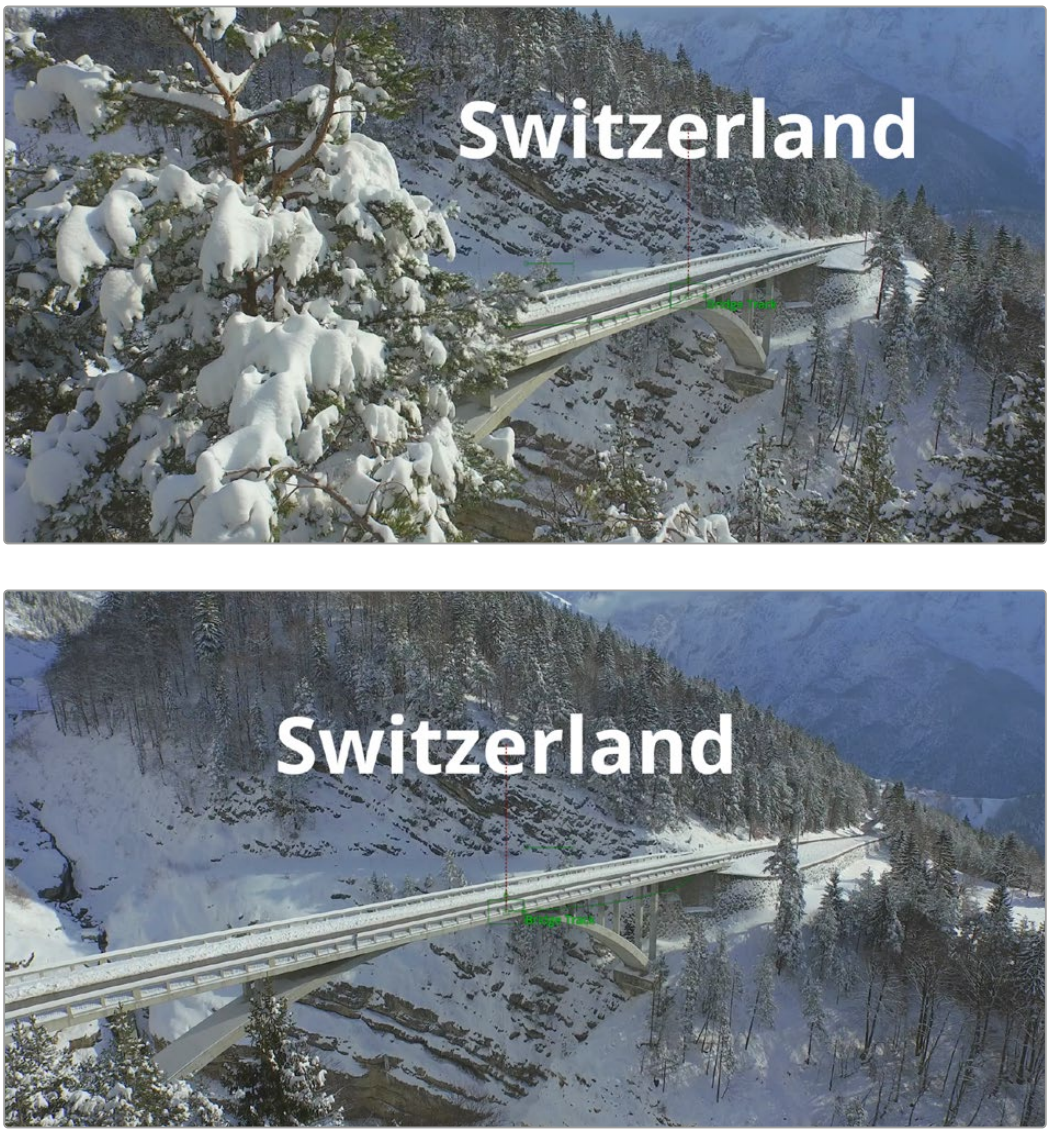
Using the X and Y Offset controls in the Tracker1 node to offset the text layer's position from the tracked motion path

The offset we create is shown as a dotted red line that lets us see the actual offset being created by the X and Y Offset controls. In fact, this is why we connected to the Bridge Track: Offset position option earlier.



The text offset from the tracked motion path; the offset can be seen as a dotted red line in the viewer

Now, if we play through this clip, we can see the text moving along with the bridge.



Two frames of the text being match moved to follow the bridge in the shot

Chapter 23

Planar Tracking

This chapter provides an overview of how to use the Planar Tracker node, and how to use it to make match moves simple.

> For more information about the Planar Tracker node, *see Chapter 59, "Tracking Nodes,"* in the Fusion Reference Manual.

Contents

Introduction to Tracking 593
Using the Planar Tracker 593
Different Ways of Using the Planar Tracker Node 593
Setting Up to Use the Planar Tracker 594
Check for Lens Distortion 594
A Basic Planar Tracker Match Move Workflow 595
Tips for Choosing Good Planes to Track 598

Introduction to Tracking

Fusion includes three different Tracking nodes that let you analyze different kinds of motion. Once you have tracked motion on a clip, you can then use the resulting data for stabilization, motion smoothing, matching the motion of one object to that of another, and a host of other essential tasks. Each tracker type has its own chapter in this manual. This chapter covers the tracking techniques with the Planar Tracker node.

Using the Planar Tracker

The Planar Tracker node is designed to deal with match moving issues on flat, unvarying surfaces that commonly arise during post-production. Examples of flat, unvarying surfaces include clips containing a license plate, a road sign, or a brick wall that often need images merged on top of them, such as replacing the numbers in the license plate, changing the city's name in the road sign, or placing a billboard poster on the empty brick wall.

The Planar Tracker automates this process by analyzing the perspective distortions of a planar surface on a background plate over time, and then re-applying those same perspective distortions to a different foreground.

> **TIP:** Part of using the Planar Tracker is also knowing when to give up and fall back to using Fusion's Tracker node or to manual keyframing. Some shots are simply not trackable, or the resulting track suffers from too much jitter or drift. The Planar Tracker is a time-saving node in the artist's toolbox and, while it can track most shots, no tracker is a 100% solution.

Different Ways of Using the Planar Tracker Node

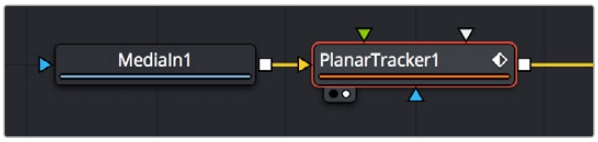
Like the other tracking nodes found in Fusion, the Planar Tracker can both analyze and contain the resulting image tracking data interior to the node, and it can also use that tracking data to transform either another image, paint strokes, or a polygon mask shape.

The Planar Tracker provides four modes of operation:

- **Track:** Used to isolate a planar surface and track its movement over time. Then, you can create a Planar Transform node that uses this data to match move another clip in various ways.
- **Steady:** After analyzing a planar surface, this mode removes all motion and distortions from the planar surface, usually in preparation for some kind of paint or roto task, prior to "unsteadying" the clip to add the motion back.
- **Corner Pin:** After analyzing a planar surface, this mode computes and applies a matching perspective distortion to a foreground image you connect to the foreground input of the Planar Tracker node, and merges it on top of the tracked footage.
- **Stabilize:** After analyzing a planar surface, allows smoothing of a clip's translation, rotation, and scale over time. Good for getting unwanted vibrations out of a clip while retaining the overall camera motion that was intended.

Setting Up to Use the Planar Tracker

Similar to the Tracker node, to do a planar track, you need to connect the output of the image you want to track to the background input of a Planar Tracker node.



Connecting an image to the background input of a PlanarTracker node

Check for Lens Distortion

If the image has barrel distortion, or any other kinds of lens distortion, it can adversely affect your track. The more lens distortion in the footage, the more the resulting track will slide and wobble. If you can see distortion in the image or you're having problems with the track, you'll want to try inserting the Lens Distort node between the image and the Planar Tracker to eliminate this problem.

Fusion's Lens Distort node can be used to remove or add lens distortion in an image. Connecting the MediaIn or Loader node to the Lens Distort node displays controls for manually correcting lens distortion. If you use Synth Eyes, PFTrack or 3D Equalizer software, you can also import lens data from those applications to make the adjustments more automatic.



A Lens Distort node inserted between a MediaIn1 and Planar Tracker to remove lens distortion

For more information about using the Lens Distort node, *see Chapter 63, "Warp Nodes,"* in the Fusion Reference Manual.

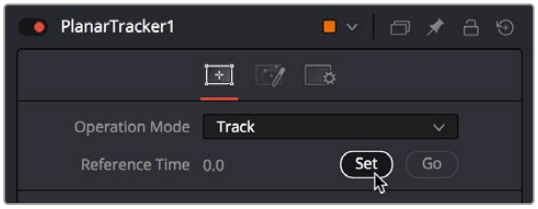
If you are using DaVinci Resolve, you can use the Lens Corrections control in the Cut page or Edit page. This adjustment carries over into the Fusion page. Lens correction in DaVinci Resolve automatically analyzes the frame in the Timeline viewer for edges that are being distorted by a wide angle lens. Clicking the Analyze button moves the Distortion slider to provide an automatic correction. From there, the MediaIn node in the Fusion page will have the correction applied, and you can begin planar tracking.

A Basic Planar Tracker Match Move Workflow

Using the Planar Tracker is a process, but it's straightforward once you've learned how to use it. The following procedure tries to make this process as clear as possible.

To track a surface using the Planar Tracker:

1. Make sure the Operation Mode is set to Track, as you need to analyze an image to track a surface before you do anything else.
2. With the background input of the Planar Tracker connected to an image, and the Planar Tracker open in a viewer, move the playhead to a frame of video where the planar surface you want to track is at its largest, is unoccluded, and is clearly a plane, and then click the Set button in the Track panel of the Inspector to make this the reference frame that will be used to guide the track.



Clicking the Set button to set the reference frame to use for analysis

3. Next, you'll need to identify the specific pattern within the image that you want to track. In most cases, this will probably be a rectangle, but any arbitrary closed polygon can be used. The pixels enclosed by this region will serve as the pattern that will be searched for on other frames. Please note that it is important that the pattern is drawn on the reference frame. In this example, we want to track the wall behind the man, so we draw a polygon around part of the wall that the man won't pass over as he moves during the shot.



Drawing a polygon to identify the part of the image you want to track, which should be a flat, trackable plane

TIP: Do not confuse the pattern you're identifying with the region you're planning to corner pin (which always has four corners and is separately specified in Corner Pin mode.

4 (Optional) If moving objects partially cover up or occlude the planar surface, you may wish to connect a mask that surrounds and identifies these occlusions to the white “occlusion mask” input of the Planar Tracker. This lets the Planar Tracker ignore details that will cause problems.

When using the Hybrid Tracker, providing a mask to deal with occluding objects is nearly mandatory, while with the Point Tracker it is recommended to try tracking without a mask.

5 If necessary, move the playhead back to the reference frame, which in this case was the first frame. Then, click the Track To End button and wait for the track to complete.



The Analyze buttons of the Planar Tracker

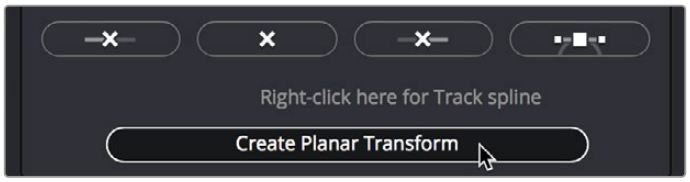
As the clip tracks, you can see track markers and trails (if they’re enabled in the Options tab of the Inspector) that let you see how much detail is contributing to the track, and the direction of motion that’s being analyzed.



During tracking, you can see track markers and trails to follow how well the track is going

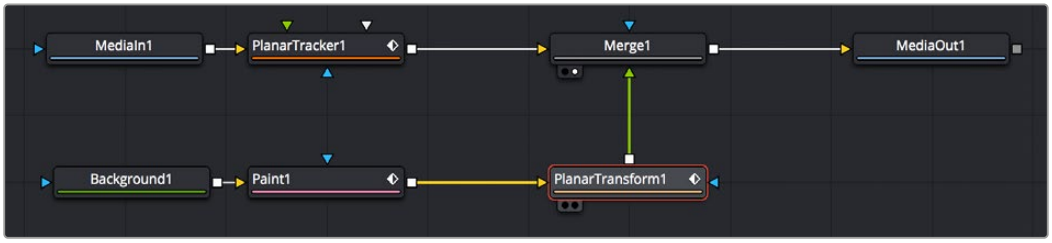
6 Once the track is complete, play through the clip to visually inspect the track so you can evaluate how accurate it is. Does it stick to the surface? Switching to Steady mode can help here, as scrubbing through the clip in Steady mode will help you immediately see unwanted motion in the track.

7 Since we’re doing a match move, click the Create Planar Transform button to export a Planar Transform node that will automatically transform either images or masks to follow the analyzed motion of the plane you tracked.



Clicking Create Planar Transform to create a node to use to transform other images or masks

In this case, the Planar Transform node will be inserted after a pair of Background and Paint nodes that are being used to put some irritatingly trendy tech jargon graffiti on the wall. The Planar Transform will automatically transform the Paint node’s output connected to its background input to match the movement of the wall.



Adding the PlanarTransform node after a Paint node to match move it to the background image, combining it via a Merge node

The result is a seamless match move of the fake graffiti married to the wall in the original clip.



The final result; the paint layer is match moved to the background successfully

TIP: If you want to composite semi-transparent paint strokes on the wall, or use Apply modes with paint stroke, you can attach a Paint node to a Background node set to 100 transparency. The resulting image will be whatever paint strokes you make against transparency and is easy to composite.

Tips for Choosing Good Planes to Track

The region to track is specified by drawing a polygon on the reference frame. Make sure the region selected belongs to a physically planar surface in the shot. Sometimes a region that is only approximately planar can be used. In general, the less planar the surface, the poorer the quality of the resulting track.

As a rule of thumb, the more pixels in the pattern, the better the quality of the track. In particular, this means on the reference frame, the pattern to be tracked should:

- Be as large as possible.
- Be as much in frame as possible.
- Be as unoccluded as possible by any moving foreground objects.
- Be at its maximum size (e.g., when tracking an approaching road sign, it is good to pick a later frame where it is 400 x 200 pixels rather than 80 x 40 pixels).
- Be relatively undistorted (e.g., when the camera orbits around a flat stop sign, it is better to pick a frame where the sign is face on parallel to the camera rather than a frame where it is at a highly oblique angle).

If the pattern contains too few pixels or not enough trackable features, this can cause problems with the resulting track, such as jitter, wobble, and slippage. Sometimes dropping down to a simpler motion type can help in this situation.

Chapter 24

Using Open FX, Resolve FX, and Fuse Plugins

Fusion's capabilities can be extended using different kinds of plugins. All compositions in Fusion Studio and in the Fusion page of DaVinci Resolve support third-party Open FX plugins.

Additionally, the Fusion page of DaVinci Resolve provides access to all of the Resolve FX that come with DaVinci Resolve.

Lastly, you can develop your own plugins without using a computer development environment by scripting Fusion's native Fuse plugins.

Contents

What Are Open FX? 600
What Are Resolve FX? 600
Applying Open FX and Resolve FX Plugins 600
DaVinci Resolve OFX Renderer for Fusion Studio 601
Introduction to Fuse Plugins 603

What Are Open FX?

Fusion is able to use compatible Open FX (OFX) plugins that are installed on your computer. Open FX is an open standard for visual effects plugins. It allows plugins written to the standard to work on both DaVinci Resolve and Fusion Studio as well as other applications that support the standard.

OFX plugins can be purchased and downloaded from third-party suppliers such as BorisFX, Red Giant, and RE:Vision Effects. All OFX appear in the Open FX category of the Effects Library, alongside all other effects that are available in Fusion.

What Are Resolve FX?

The Fusion page is DaVinci Resolve is also able to access Resolve FX. Resolve FX are filter effects within DaVinci Resolve. Although most Resolve FX have the same capabilities on all DaVinci Resolve pages, some are specific to the Color page and require the use of the Color page tracker for full functionality. All Resolve FX appear in the Fusion page Effects Library in the Open FX category alongside third-party OFX plugins. Resolve FX are not available in Fusion Studio.

> For more information about Resolve FX, see Chapter 154, “Resolve FX,” in the *DaVinci Resolve 20 Reference Manual*.

Applying Open FX and Resolve FX Plugins

Resolve FX and OFX plugins are applied to the Node Editor just as you would apply native Fusion nodes. They can help you create fast, interesting looks and effects using the Color page, or imaginative transitions and effects on your clips on the Edit page. Resolve FX are installed with DaVinci Resolve,

After installing a set of OFX plugins, you can access them or Resolve FX plugins in Fusion by opening the Open FX category in the Effects Library.

To add a plugin to the Node Editor, either click the Open FX or Resolve FX plugin name in the Effects Library or drag and drop the plugin onto a connection line to insert it into the node tree. If the plugin has editable settings, you can adjust these in the Inspector.

DaVinci Resolve OFX Renderer for Fusion Studio

You can use the OFX DaVinci Resolve Renderer plugin in Fusion Studio. This plugin allows Fusion Studio to open and apply DaVinci Resolve's color changes using exported DRX stills. Essentially this turns an installation of DaVinci Resolve Studio into one large OFX plugin. This can be very useful when round tripping between DaVinci Resolve Studio and Fusion Studio, where you may want to maintain the exact look created in DaVinci Resolve.

Exporting a DRX still contains much more information about the grade than a simple LUT. DRX files can include data like all the native color and sizing palettes and other items you stored in a Gallery still. This ensures that the grade the colorist created can be exactly replicated in the VFX software.

As of this writing, the following attributes WILL transfer via the DaVinci Resolve OFX Renderer Plugin:

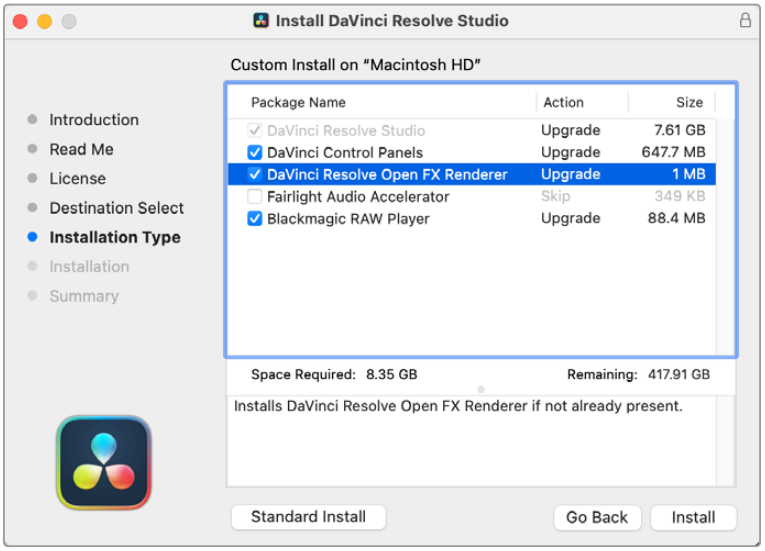
- All native color palettes (Primaries, HDR, Curves, etc.).
- All native sizing palettes (Input, Output, Node, etc.).
- Most Resolve FX.
- Most third-party Open FX.

As of this writing, the following attributes WILL NOT transfer via the DaVinci Resolve OFX Renderer Plugin:

- Resolve FX and Open FX that are temporal based.
- Magic Mask.

To set up the DaVinci Resolve OFX Renderer Plugin:

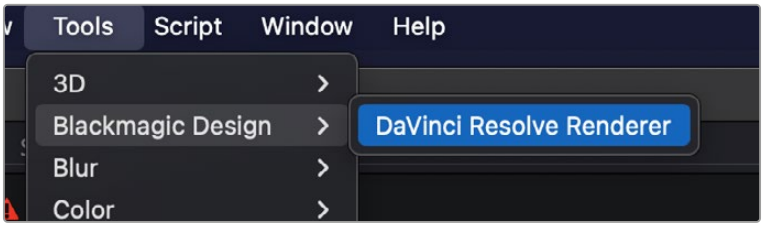
1 When installing DaVinci Resolve Studio, select Custom Install and then check the plugin option in the program installer.



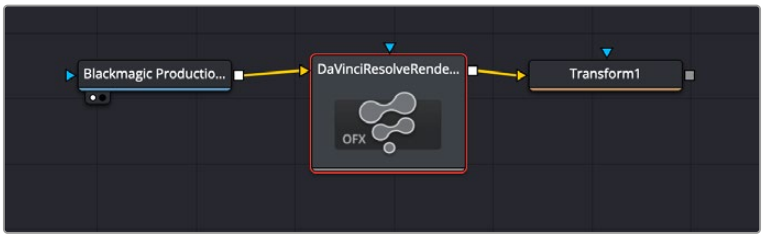
Select Custom Install, then check DaVinci Resolve Open FX Renderer to install the plugin

2 Create your grade for a clip in the Color page, then save a still to the Gallery. From the Gallery, export a .DRX file.

3 In Fusion Studio, apply the DaVinci Resolve Renderer plugin where appropriate for your composition.

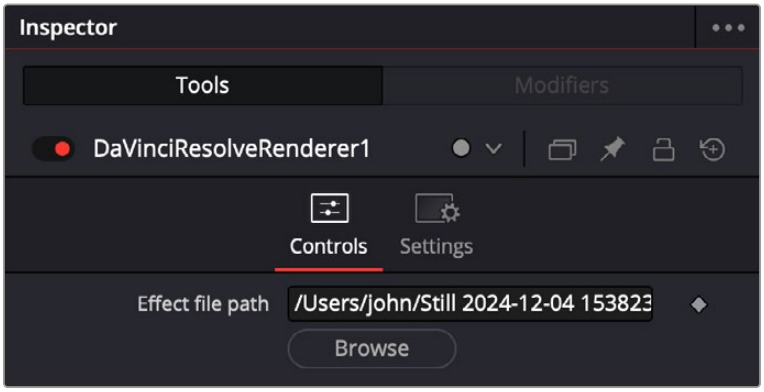


ApplyRendererPluginFusion



Inserting the DaVinci Resolve Renderer OFX into a Fusion Studio node tree

4 In the DaVinci Resolve Renderer Inspector in Fusion Studio, choose the .DRX file you exported from the file browser.



Select the .DRX file you exported from DaVinci Resolve in the plugin

Introduction to Fuse Plugins

Fuses are plugins developed for Fusion using the Lua built-in scripting language. Being script-based, Fuses are compiled on-the-fly in Fusion without the need of a computer programming environment. While a Fuse may be slower than an identical Open FX plugin created using Fusion's C++ SDK, a Fuse will still take advantage of Fusion's existing nodes and GPU acceleration.

To install a Fuse:

1. Use the .fuse extension at the end of the document name.
2. For DaVinci Resolve, save it in one of the following locations:

On macOS: Macintosh HD/Users/username/Library/Application Support/Blackmagic Design/DaVinci Resolve/Fusion/Fuses

On Windows: C:\Users\username\AppData\Roaming\Blackmagic Design\DaVinci Resolve\Support\Fusion\Fuses

On Linux: home/username/.local/share/DaVinciResolve/Fusion/Fuses

For Fusion Studio, save it in one of the following locations:

On macOS: Macintosh HD/Users/username/Library/Application Support/Blackmagic Design/Fusion/Fuses/

On Windows: C:\Users\username\AppData\Roaming\Blackmagic Design\Fusion\Fuses

On Linux: home/username/.fusion/BlackmagicDesign/Fusion/Fuses

You can open and edit Fuses by selecting the Fuse node in the Node Editor and clicking the Edit button at the top of the Inspector. The Fuse opens in the text editor specified in the Global Preferences/Scripting panel.

TIP: Changes made to a Fuse in a text editor do not immediately propagate to other instances of that Fuse in the composition. Reopening a composition updates all Fuses in the composition based on the current saved version. Alternatively, you can click the Reload button in the Inspector to update the selected node without closing and reopening the composition.



3D Compositing

FUSION

CONTENTS

25 3D Compositing Basics 605

26 3D Camera Tracking 660

27 Particle Systems 678

Chapter 25

3D Compositing Basics

This chapter covers many of the nodes used for creating 3D composites, the tasks they perform, and how they can be combined to produce effective 3D scenes.

Contents

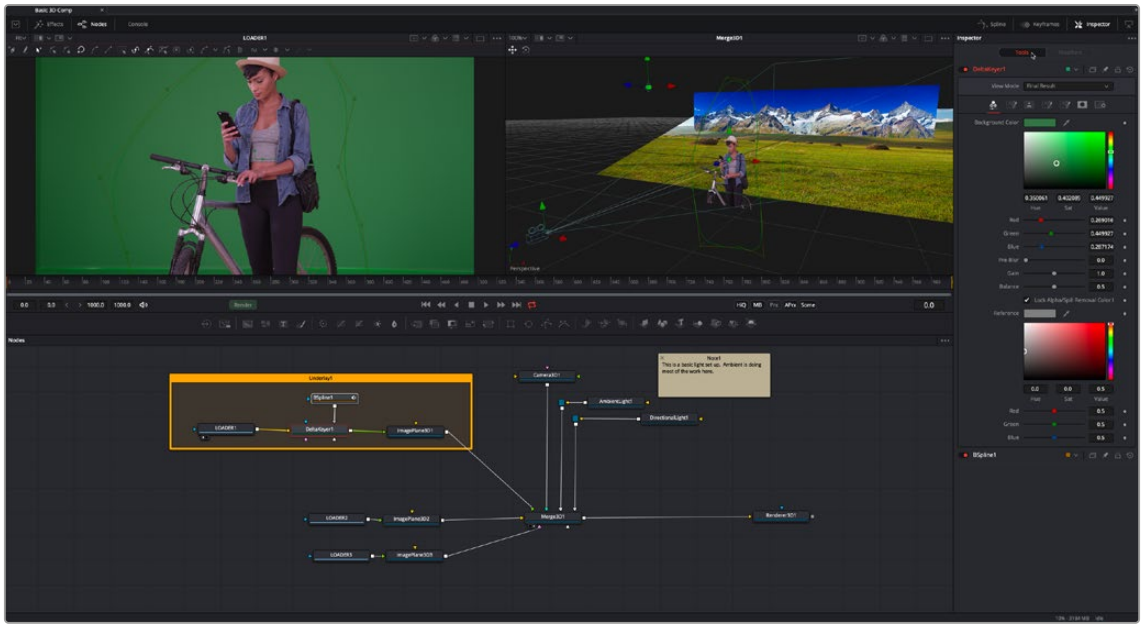
An Overview of 3D Compositing 606
3D Compositing Fundamentals 607
Creating a Minimal 3D Scene 607
The Elements of a 3D Scene 609
Geometry Nodes 609
The Merge3D Node 611
The Renderer3D Node 614
Software vs. GPU Rendering 615
Software Renderer 615
OpenGL Renderer 615
OpenGL UV Renderer 616
Loading 3D Nodes into the Viewer 617
Navigating the 3D View 618
Transforming Cameras and Lights Using the Viewers 619
Transparency Sorting 620
Material Viewer 620
Transformations 621
Onscreen Transform Controls 622
Pivot 623
Target 623
Parenting 624
Cameras 625
Quickly Viewing a Scene Through a Camera 626
Plane of Focus and Depth of Field 626
Importing Cameras 627
Lighting and Shadows 628
Enabling Lighting in the Viewer 628
Enabling Lighting to Be Rendered 628
Cotrolling Lighting within Each 3D Object 628
Lighting Types Explained 629
Materials and Textures 633
Material Components 634
Alpha Detail 636
Illumination Models 638
Textures 639
Reflections and Refractions 640
Bump Maps 642
Projection Mapping 643
Geometry 646
Common Visibility Parameters 647
Adding FBX Models 647
Using Text3D 649
Fog 3D and Soft Clipping 653
Material and Object IDs 655
World Position Pass 655
Point Clouds 657

An Overview of 3D Compositing

Traditional image-based compositing is a two-dimensional process. Image layers have only the amount of depth needed to define one as foreground and another as background. This is at odds with the realities of production, since all images are either captured using a live-action camera with freedom in all three dimensions, in a shot that has real depth, or have been created in a true 3D modeling and rendering application.

Within the Fusion Node Editor, you have a GPU-accelerated 3D compositing environment that includes support for imported geometry, point clouds, and particle systems for taking care of such things as:

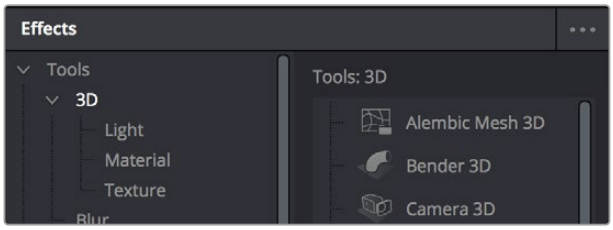
- Converting 2D images into image planes in 3D space
- Creating rough primitive geometry
- Importing mesh geometry from FBX or Alembic scenes
- Creating realistic surfaces using illumination models and shader compositing
- Rendering with realistic depth of field, motion blur, and supersampling
- Creating and using 3D particle systems
- Creating, extruding, and beveling 3D text
- Lighting and casting shadows across geometry
- 3D camera tracking
- Importing cameras, lights, and materials from 3D applications such as Maya, 3ds Max, or LightWave
- Importing matched cameras and point clouds from applications such as SynthEyes or PF Track



An example 3D scene in Fusion

3D Compositing Fundamentals

The 3D category of nodes (which includes the Light, Material, and Texture subcategories) work together to create 3D scenes. Examples are nodes that generate geometry, import geometry, modify geometry, create lights and cameras, and combine all these elements into a scene. Nearly all these nodes are collected within the 3D category of nodes found in the Effects Library.

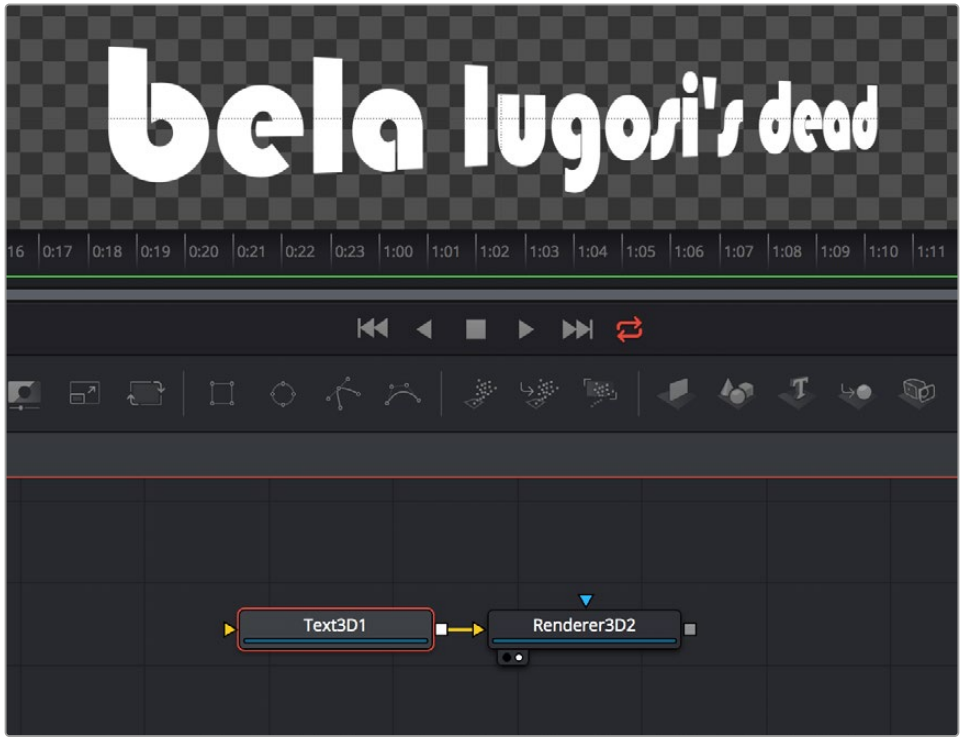


The 3D category of nodes in the Effects Library

Conveniently, at no point are you required to specify whether your overall composition is 2D or 3D, because you can seamlessly combine any number of 2D and 3D "scenes" together to create a single output. However, the nodes that create these scenes must be combined in specific ways for this to work properly.

Creating a Minimal 3D Scene

Creating a 3D scene couldn't be easier, but you need to connect the required nodes in the right way. At minimum, you need only connect a geometry node (such as a Text3D node) to a Renderer3D node to output a 2D image that can be combined with other 2D images in your composition, as seen below. However, you'll only get a simply shaded piece of geometry for your trouble, although you can color and transform it in the Inspector using controls internal to whichever geometry node you're using.

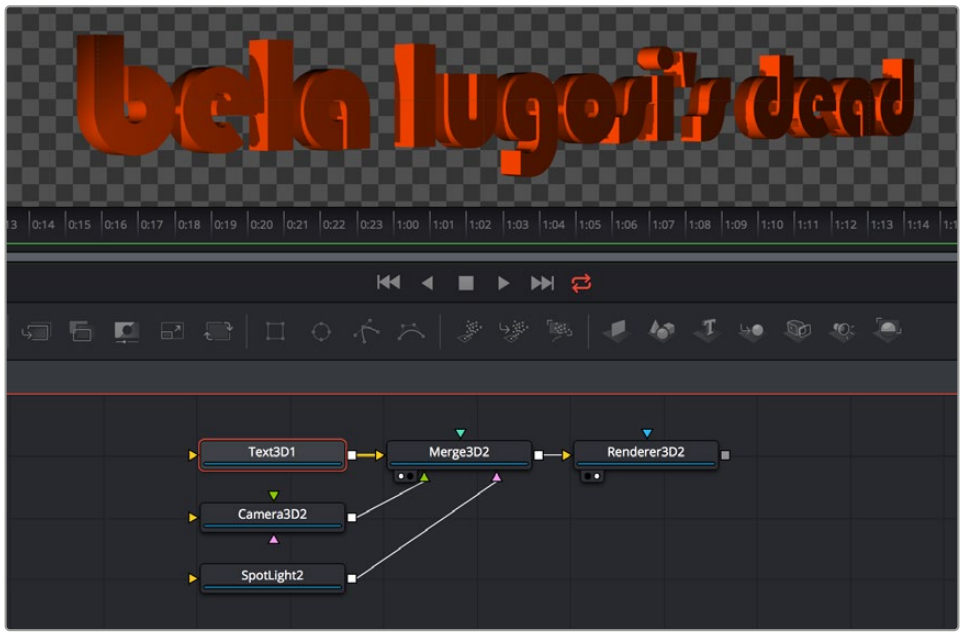


A simple 3D scene with a Text3D node connected directly to a Renderer3D node

More realistically, each 3D scene that you want to create will probably have three to five nodes to give you a better lit and framed result. These include:

- One of the available geometry nodes (such as Text3D or Image Plane 3D)
- A light node (such as DirectionalLight or SpotLight)
- A camera node
- A Merge3D node
- A Renderer3D node

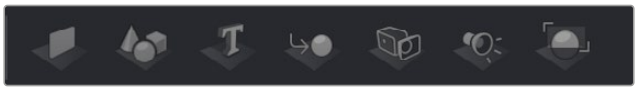
All these should be connected together as seen below, with the resultantly more complex 3D scene shown below.



The same text, this time lit and framed using Text3D, Camera, and SpotLight nodes to a Merge3D node

To briefly explain how this node tree works, the geometry node (in this case Text3D) creates an object for the scene, and then the Merge3D node provides a virtual stage that combines the attached geometry with the light and camera nodes to produce a lit and framed result with highlights and shadows, while the aptly named Renderer3D node renders the resulting 3D scene to produce 2D image output that can then be merged with other 2D images in your composition.

In fact, these nodes are so important that they appear at the right of the toolbar, enabling you to quickly produce 3D scenes whenever you require. You might notice that the order of the 3D buttons on the toolbar, from left to right, corresponds to the order in which these nodes are ordinarily used. So, if you simply click on each one of these buttons from left to right, you cannot fail to create a properly assembled 3D scene, ready to work on, as seen in the previous screenshot.



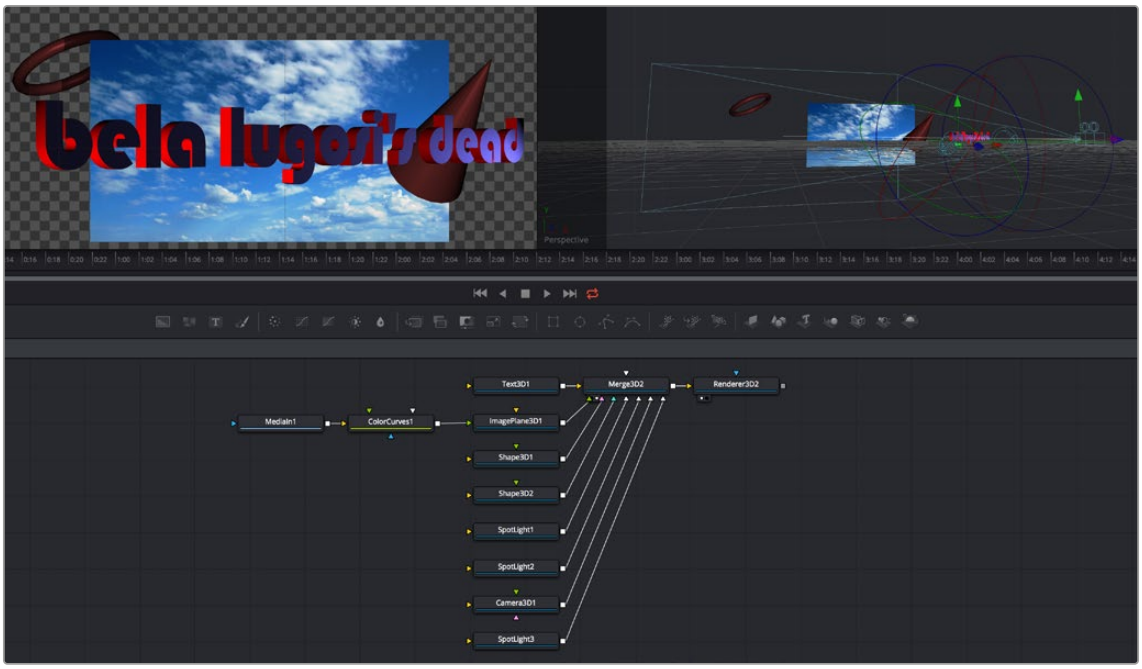
The 3D nodes available from the toolbar include the ImagePlane3D, Shape3D, Text3D, Merge3D, Camera3D, SpotLight3D, and Renderer3D nodes

The Elements of a 3D Scene

All 3D nodes can be divided into a number of categories.

Geometry Nodes

You can add 3D geometry to a composition using the ImagePlane3D node, the Shape3D node, the Cube3D node, the Text3D node, or optionally by importing a model via the FBX Mesh 3D node. Furthermore, you can add particle geometry to scenes from pEmitter nodes. You can connect these to a Merge3D node either singularly or in multiples to create sophisticated results combining multiple elements.

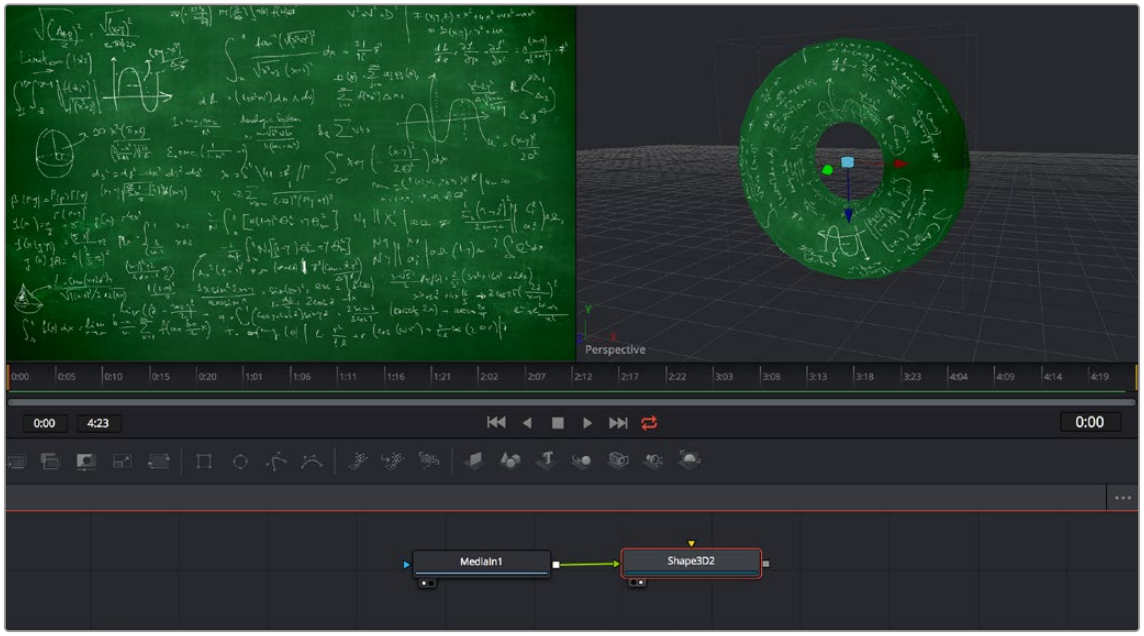


A more complex 3D scene combining several geometry nodes including the Text3D, Shape3D, and ImagePlane3D nodes

Texturing Geometry

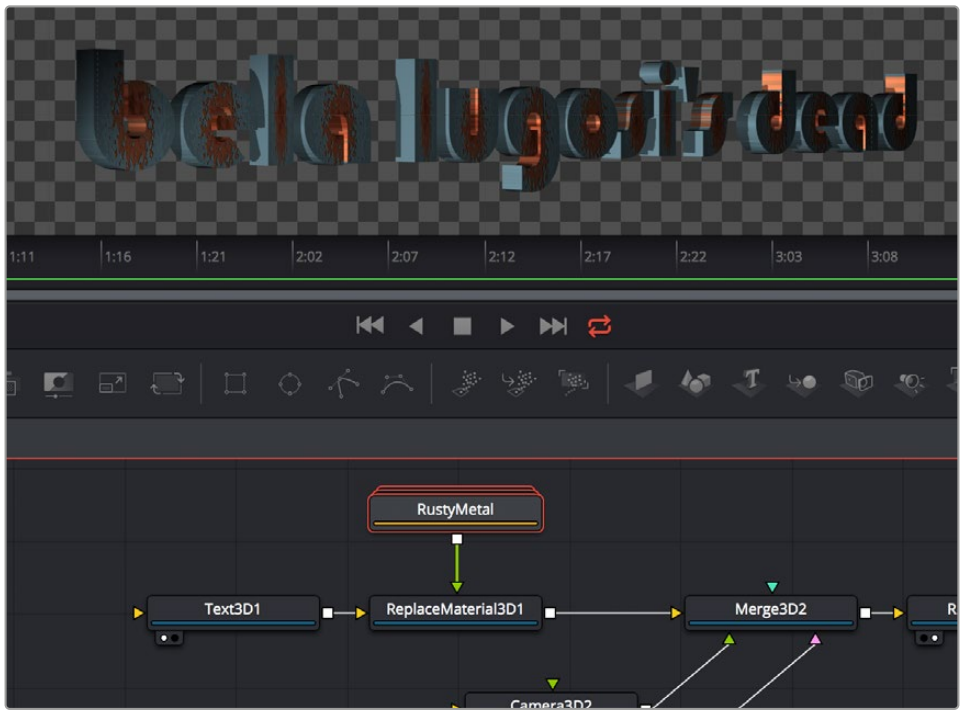
By itself, geometry nodes can only consist of a simple flat color. However, you can alter the look of 3D geometry by texturing it using clips (either still images or movies), using material nodes such as the Blinn and Phong nodes to create more sophisticated textures with combinations of 2D images and environment maps, or you can use a preset shader from the Templates > Shader bin of the Effects Library, which contains materials and texture presets that are ready to use.

If you're working with simple geometric primitives, you can texture them by connecting either an image (a still image or movie) or a shader from the Templates bin of the Effects Library directly to the material input of a Shape3D, Cube3D, or other compatible node, as shown below.



An image connected to the material input of a Shape3D node set to Torus, with the image (left), and the shaded torus (right)

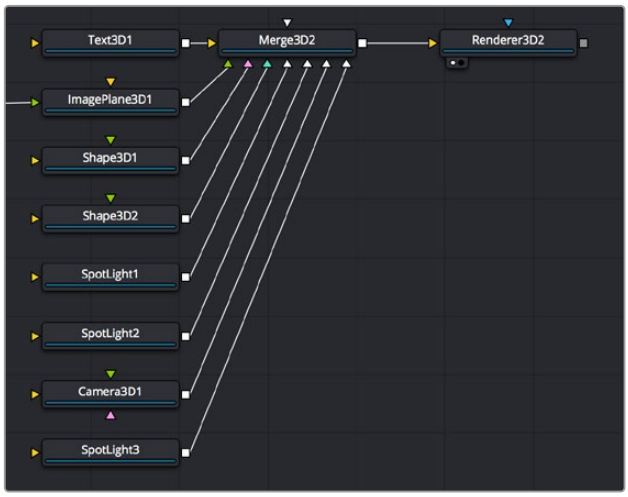
If you're shading or texturing Text3D nodes, you need to add a texture in a specific way since each node is actually a scene with individual 3D objects (the characters) working together. In the following example, the RustyMetal shader preset is applied to a Text3D node using the ReplaceMaterial3D node. The interesting thing about the ReplaceMaterial3D node is that it textures every geometric object within a scene at once, meaning that if you put a ReplaceMaterial3D node after a Text3D node, you texture every character within that node. However, if you place a ReplaceMaterial3D node after a Merge3D node, then you'll end up changing the texture of every single geometric object being combined within that Merge3D node, which is quite powerful.



The geometry created by a Text3D node is textured using a shader connected to a ReplaceMaterial3D node that's connected downstream of the object you want to shade

The Merge3D Node

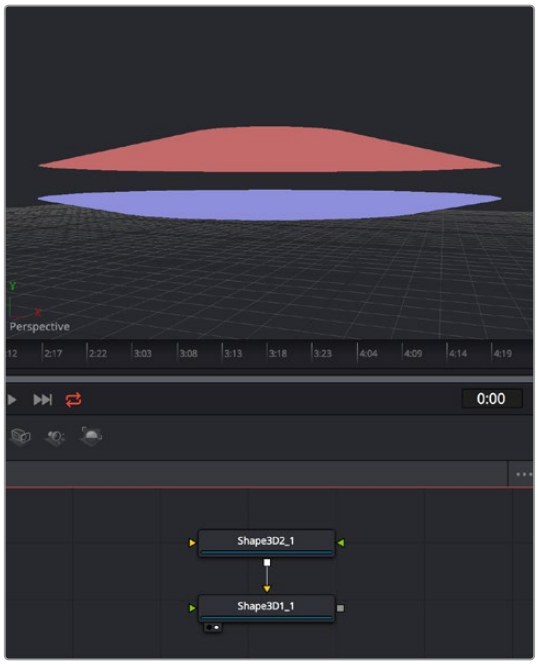
The Merge3D node combines the output of one or more 3D nodes into a single scene. Unlike the Merge2D node, the ordering of elements in the scene is not restricted to only background and foreground inputs. Instead, the Merge3D node lets you connect an unlimited number of inputs, with the resulting output combined according to each object's absolute position in 3D space.



Merging many objects together in a 3D scene using the Merge3D node

Combining Objects Directly

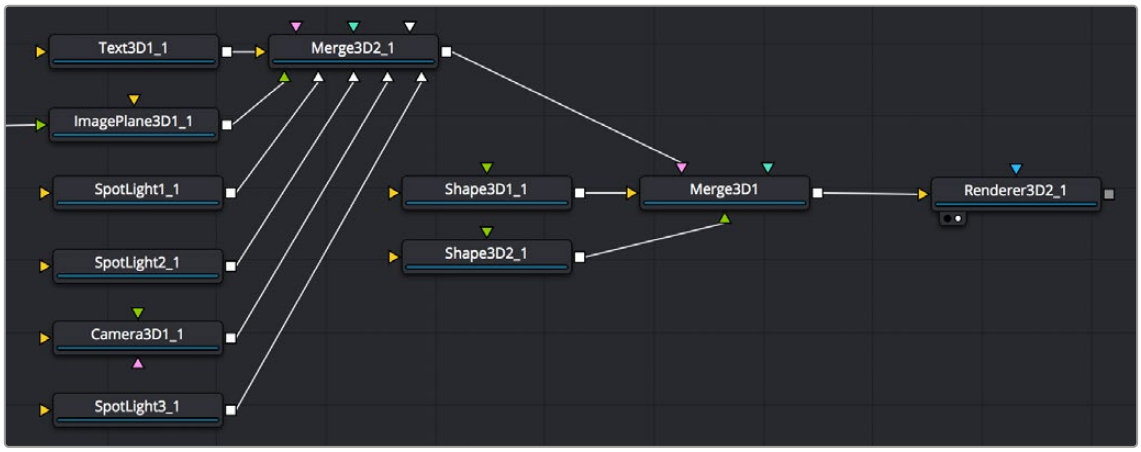
While the Merge3D node provides a structured way of combining objects, you can also combine 3D objects such as Text3D and Shape3D nodes by connecting the output of one 3D object node to the input of another, as seen in the following screenshot. When you do this, you must use each node's internal transform parameters to transform their position, size, and rotation directly, but the transform control of downstream 3D object nodes also transforms all upstream 3D object nodes. This even works for lights and the Camera3D node, giving you a fast way of combining a set of objects that always go together, which you can later connect to a Merge3D node for additional lighting and eventual connection to a Renderer3D node.



Connecting one Shape3D node to another directly to combine them Transforming the last downstream 3D object also transforms all upstream objects; the last Shape3D node is viewed, showing both

Combining Multiple Merge3D Nodes

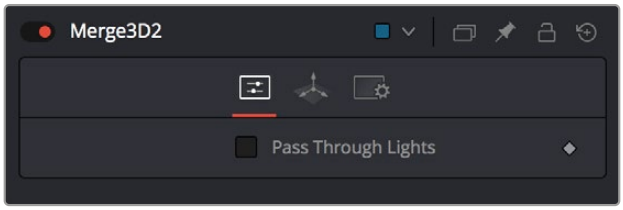
Furthermore, Merge3D nodes can be combined with other Merge3D nodes, allowing you to create composite 3D scenes made up of multiple “sub-scenes,” each put together within individual Merge3D nodes.



You can build elaborate scenes using multiple Merge3D nodes connected together

Lighting Multiple Merge3D Nodes

Once you’ve combined multiple Merge3D nodes, there’s an easy way to control how lights that are connected to upstream Merge3D nodes affect the results of other Merge3D nodes connected downstream. Each Merge3D node’s Controls tab contains a single checkbox, Pass Through Lights, which enables lighting to pass through the output of an upstream Merge3D node in order to shine onto objects connected to downstream Merge3D nodes.



You can light downstream Merge3D scenes with lights connected to upstream Merge3D scenes by turning on Pass Through Lights

This checkbox is disabled by default, which lets you light elements in one Merge3D scene without worrying about how the lighting will affect geometry attached to other Merge3D nodes further downstream. For example, you may want to apply a spotlight to brighten the wall of a building in one Merge3D node without having that spotlight spill over onto the grass or pavement at the foot of the wall modeled in another Merge3D node. In the example shown below, the left image shows how the cone and torus connected to a downstream node remain unlit by the light in an upstream node with Pass Through Lights disabled, while the right image shows how everything becomes lit when turning Pass Through Lights on.



The result of lights on the text in one Merge3D node not affecting the cone and torus added in a downstream Merge3D node (left). Turning on Pass Through Lights in the upstream Merge3D node results in those lights also illuminating the downstream shapes (right).

Transforming Merge3D Scenes

Each Merge3D node includes a Transform tab. These transform parameters adjust the position, scale, and rotation of all objects being combined within that Merge3D node together, including lighting and particles. All transformations take place around a common pivot point. This forms the basis of parenting in the 3D environment.



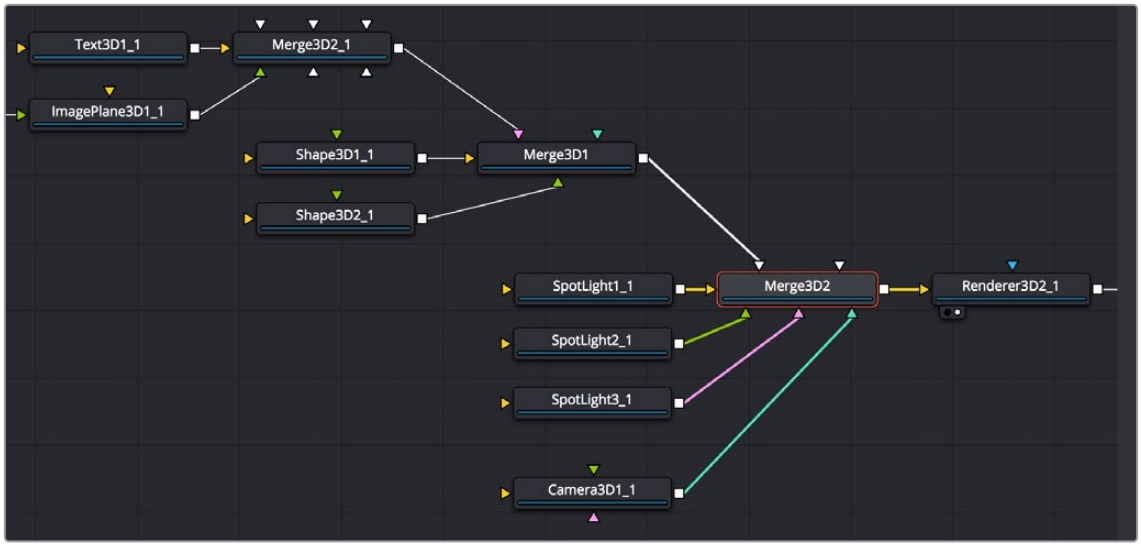
The Transform tab of a Merge3D node

If you transform a Merge3D node that's connected to other Merge3D nodes, what happens depends on which node you're transforming, an upstream node or the downstream node:

- If you transform a downstream Merge3D node, you also transform all upstream nodes connected to it as if they were all a single scene.
- If you transform an upstream Merge3D node, this has no effect on downstream Merge3D nodes, allowing you to make transforms specific to that particular node's scene.

Transforming Upstream, Lighting Downstream

When building complex scenes using multiple Merge3D nodes being combined together, it's common to use one last downstream node to combine light and camera nodes to illuminate the final scene, while leaving the upstream Merge3D nodes free for controlling object transforms and animation. This way, you can transform and animate subsets of your overall scene without worrying about accidentally altering the overall lighting scheme or cameras for that scene, unless you've specifically connected lights or cameras upstream that are meant to be attached to the geometry you're transforming.

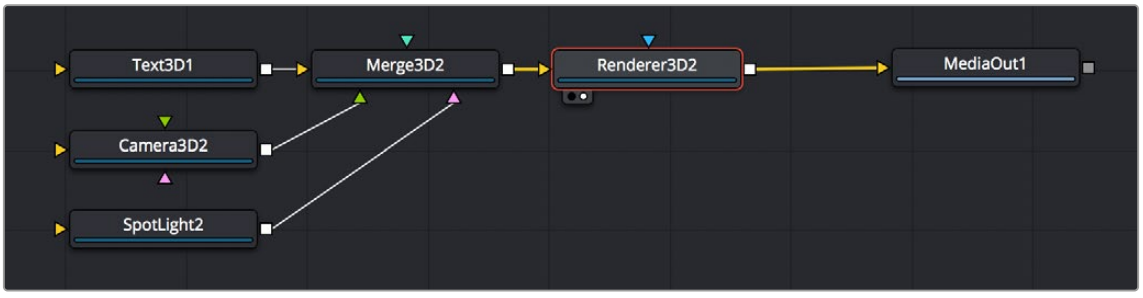


An example of a 3D scene using multiple Merge3D nodes working together; the upstream Merge3D nodes arrange the 3D objects placed within the scene, while the last Merge3D node (orange) lights and frames the scene.

The Renderer3D Node

Every 3D node you add outputs a complete 3D scene. This is unlike most traditional 3D modeling and animation programs, where all objects reside within a global scene environment. This means that the scenes created by a Camera 3D node and an image plane are separate until they're combined into the same scene via a Merge3D node, which itself outputs a complete 3D scene. However, this 3D scene data can neither be composited with other 2D images in your composition nor rendered out without first being rendered within the node tree using a Renderer3D node.

To be more specific, 3D nodes that output 3D scenes cannot be connected directly to inputs that require 2D images. For example, the output of an ImagePlane3D node cannot be connected directly to the input of a Blur node, nor can the output of a Merge3D node be directly connected to a regular Merge node. First, a Renderer3D node must be placed at the end of your 3D scene to render it into 2D images, which may then be composited and adjusted like any other 2D image in your composition.



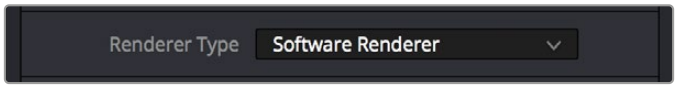
Output of a Merge3D connected to a Renderer3D node to output 2D image data

The Renderer3D uses one of the cameras in the scene (typically connected to a Merge3D node) to produce an image. If no camera is found, a default perspective view is used. Since this default view rarely provides a useful angle, most people build 3D scenes that include at least one camera.

The image produced by the Renderer3D can be any resolution with options for fields processing, color depth, and pixel aspect.

Software vs. GPU Rendering

The Renderer3D node lets you choose between using a software renderer or an OpenGL renderer, trading off certain aspects of rendered image quality for speed, and trading off depth of field rendering for soft shadow rendering, depending on the needs of a particular element of your composition. To choose which method of rendering to use, there's a Renderer Type pop-up menu in the Controls tab of each Renderer3D node's parameters in the Inspector. The default is Software Renderer.

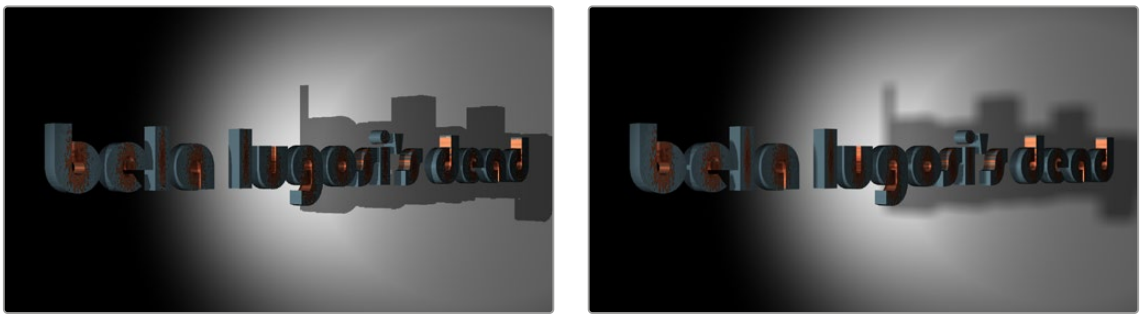


The Renderer Type option in the Controls tab of a Renderer3D node

Software Renderer

The software renderer is generally used to produce the final output. While the software renderer is not the fastest method of rendering, it has twin advantages. First, the software renderer can easily handle textures much larger than one half of your GPU's maximum texture size, so if you're working with texture images larger than 8K you should choose the software renderer to obtain maximum quality.

Second, the software renderer is required to enable the rendering of "constant" and "variable" soft shadows with adjustable Spread, which is not supported by the OpenGL renderer. Soft shadows are more natural, and they're enabled in the Shadows parameters of the Controls tab of light nodes; you can choose Sampling Quality and Softness type, and adjust Spread, Min Softness, and Filter Size sliders. Additionally, the software renderer supports alpha channels in shadow maps, allowing transparency to alter shadow density.



When the Renderer3D node "Renderer Type" drop-down is set to OpenGL Renderer, you cannot render soft shadows or excessively large textures (left). When the Renderer3D node "Renderer Type" drop-down is set to Software Renderer, you can render higher-quality textures and soft shadows (right).

OpenGL Renderer

The OpenGL renderer takes advantage of the GPU in your computer to render the image; the textures and geometry are uploaded to the graphics hardware, and OpenGL shaders are used to produce the result. This can produce high-quality images that can be perfect for final rendering, and can also be potentially orders of magnitude faster than the software renderer, but it does pose some limitations

on some rendering effects, as soft shadows cannot be rendered, and the OpenGL renderer also ignores alpha channels during shadow rendering, resulting in a shadow always being cast from the entire object.

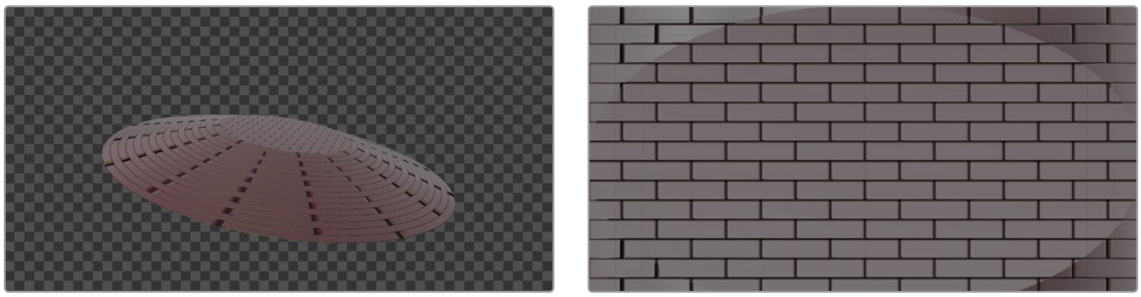
On the other hand, because of its speed, the OpenGL renderer exposes additional controls for Accumulation Effects that let you enable depth of field rendering for creating shallow-focus effects. Unfortunately, you can't have both soft shadow rendering and depth of field rendering, so you'll need to choose which is more important for any given 3D scene you render.

Don't Forget That You Can Combine Rendered Scenes in 2D

While it may seem like an insurmountable limitation that you can't output both soft shadows and depth of field using the same renderer, don't forget that you can create multiple 3D scenes each using different renderers and composite them in 2D later on. Furthermore, you can also render out auxiliary channels that can be used by 2D image processing nodes such as AmbientOcclusion, DepthBlur, and Fog to create pseudo-3D effects using the rendered images.

OpenGL UV Renderer

When you choose the OpenGL UV Renderer option, a Renderer3D node outputs an "unwrapped" version of the textures applied to upstream objects, at the resolution specified within the Image tab of that Renderer3D node.



A normally rendered 3D scene (left), and the same scene rendered using the OpenGL UV Renderer mode of the Renderer3D node (right)

This specially output image is used for baking out texture projections or materials to a texture map for one of two reasons:

- Baking out projections can speed up a render.
- Baking out projections lets you modify the texture using other 2D nodes within your composition, or even using third-party paint applications (if you output this image in isolation as a graphics file) prior to applying it back onto the geometry.

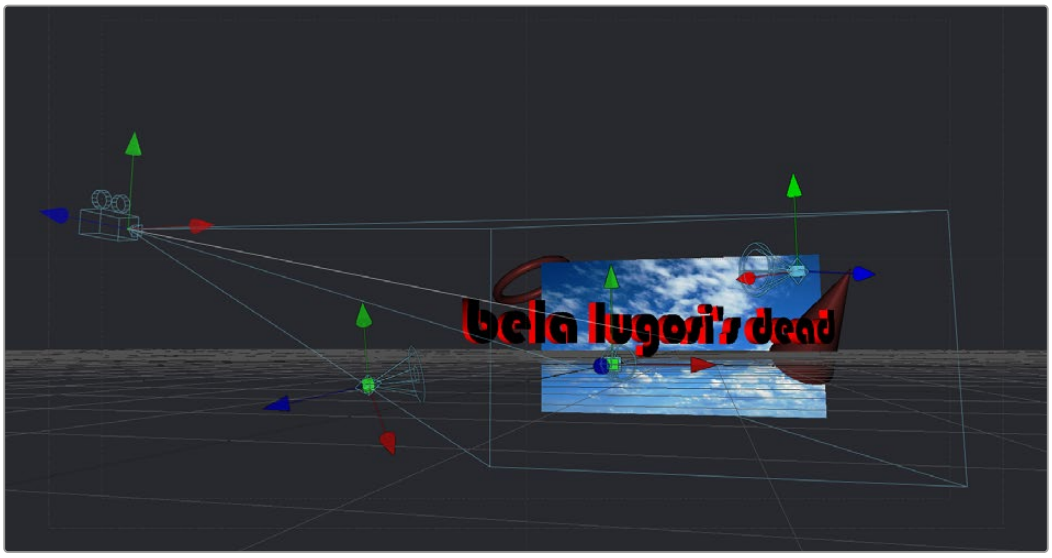
Suppose, for instance, that you have a scene on a street corner, and there's a shop sign with a phone number on it, but you want to change the numbers. If you track the scene and have standing geometry for the sign, you can project the footage onto it, do a UV render, switch the numbers around with a Paint node, and then apply that back to the mesh with a Texture2D.

The UV renderer can also be used for retouching textures. You can combine multiple DSLR still shots of a location, project all those onto the mesh, UV render it out, and then retouch the seams and apply it back to the mesh.

You could project tracked footage of a road with cars on it, UV render out the projection from the geometry, do a temporal median filter on the frames, and then map a "clean" roadway back down.

Loading 3D Nodes into the Viewer

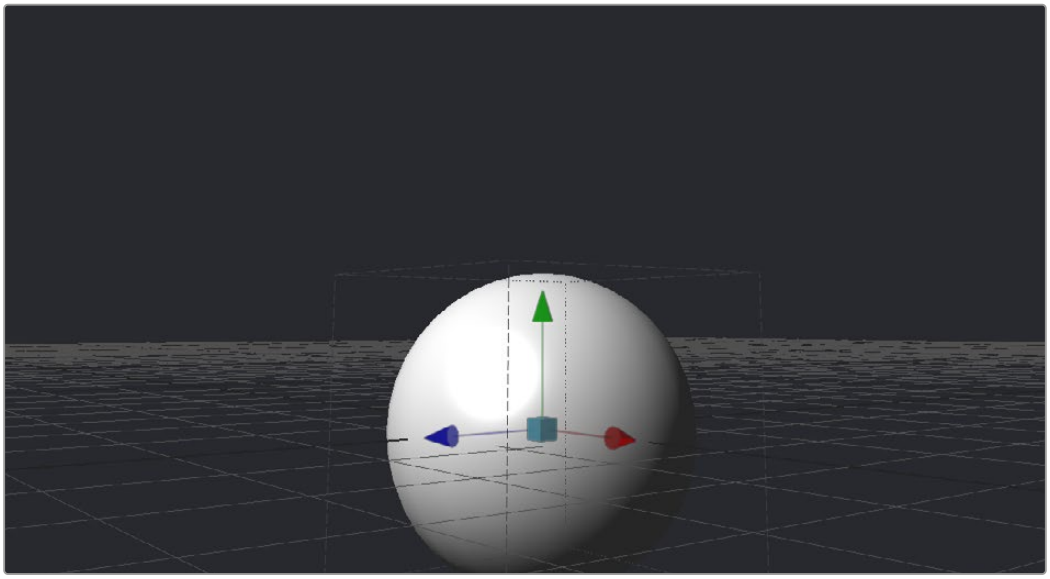
When you load a 3D node into the viewer, it switches to a 3D Viewer, which lets you pan, zoom, and rotate the scene in 3D, making it easy to make adjustments in three dimensions.



The 3D Viewer

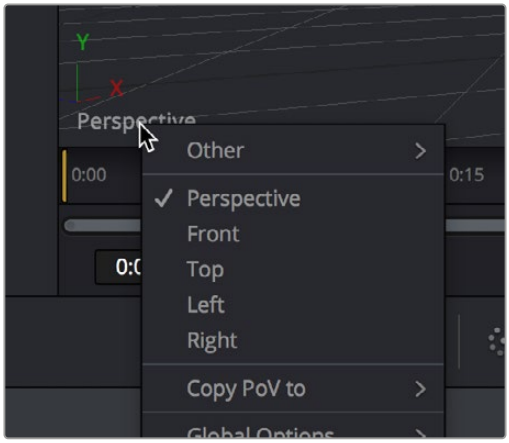
The interactive 3D Viewer is highly dependent on the computer's graphics hardware, relying on support from OpenGL. The amount of onboard memory, as well as the speed and features of your workstation's GPU, make a huge difference in the speed and capabilities of the 3D Viewer.

Displaying a node with a 3D output in any viewer will switch the display type to a 3D Viewer. Initially, the contents of the scene will be displayed through a default perspective view.



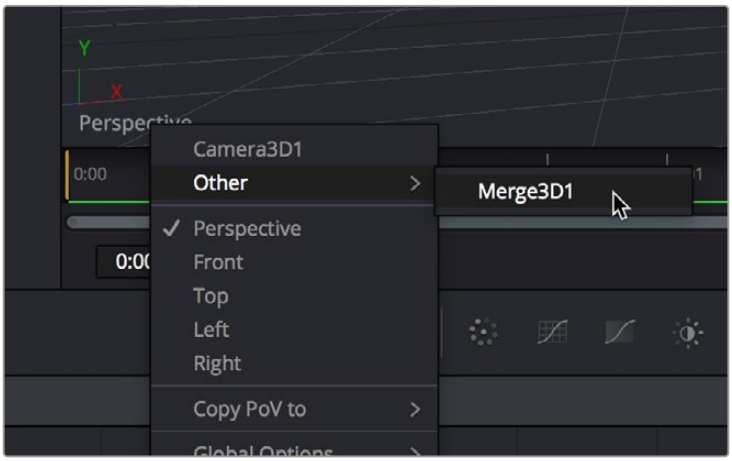
A 3D Viewer's default perspective view

To change the viewpoint, right-click in the viewer and choose the desired viewpoint from the ones listed in the Camera submenu. A shortcut to the Camera submenu is to right-click on the axis label displayed in the bottom corner of the viewer.



Right-click the Axis label of the viewer to change the viewpoint

In addition to the usual Perspective, Front, Top, Left, and Right viewpoints, if there are cameras and lights present in the scene as potential viewpoints, those are shown as well. It's even possible to display the scene from the viewpoint of a Merge3D or Transform3D by selecting it from the contextual menu's Camera > Other submenu. Being able to move around the scene and see it from different viewpoints can help with the positioning, alignment, and lighting, as well as other aspects of your composite.



The Perspective drop-down menu also shows cameras, lights, and Merge3D and Transform3D nodes you can switch to.

Navigating the 3D View

For the most part, panning and scaling of the 3D Viewer uses the same controls as the 2D Viewer.

For more information about the options available in the 3D Viewer, *see Chapter 7, "Using Viewers."* in the Fusion Reference Manual.

To pan in a 3D Viewer, do the following:

- Hold the middle mouse button and drag in the viewer.

To dolly (zoom) in the 3D Viewer, do one of the following:

- Hold down the middle and left mouse buttons and drag left or right in the viewer.
- Hold down the Command key and use your pointing device's scroll control.

To rotate around the 3D Viewer, do the following:

- Hold down the Option key and middle-button-drag left and right in the viewer.

If you want to frame certain objects in the viewer:

1. Select the viewer you want to work in.
2. Do one of the following:
   - Press Shift-F to Fit all objects in the viewer.
   - Press F to Fit to selection (or Fit All if nothing is selected).
   - Press D to Rotate the viewer to look at the center of the currently selected object without moving the viewer's position.

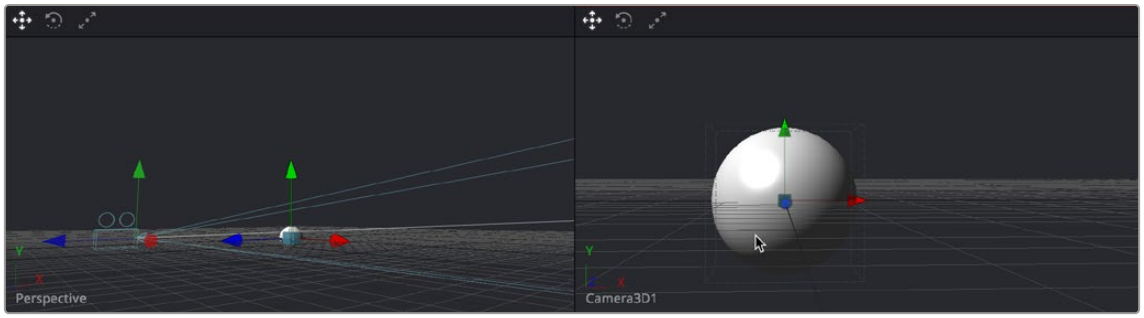
   Furthermore, selecting a 3D node in the Node Editor also selects the associated object in the 3D Viewer.

Transforming Cameras and Lights Using the Viewers

When the viewer is set to look through a 3D object in the scene, such as a camera or spotlight, the usual controls for panning and rotating the viewer will now directly affect the position of the camera or spotlight you're viewing through. Here's an example.

To adjust a camera's position when looking through it in a viewer:

1. Right-click the viewpoint label, and choose a camera from the contextual menu. (Optional) If you're in dual-viewer mode, you can load the camera you've selected in one viewer into the other viewer to see its position as you work.
2. Move the pointer into the viewer that's displaying the camera's viewpoint.
3. Hold the middle and left mouse buttons down and drag to zoom the viewer, or middle-click-drag to pan the viewer, or option-middle-click-drag to rotate the viewer, all while also moving the camera.



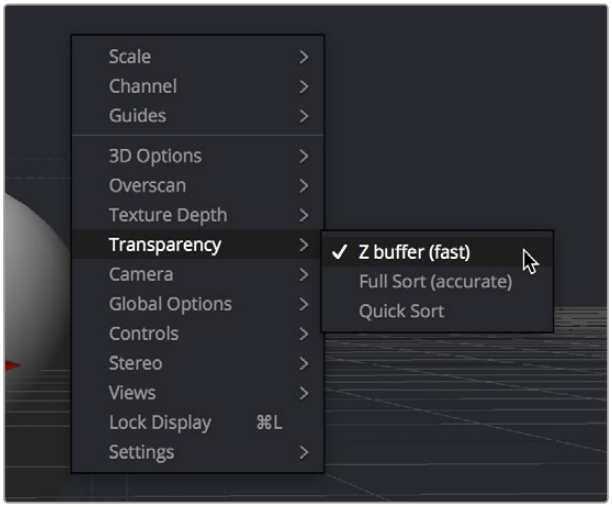
When a viewer is set to display the view of a camera or light, panning, zooming, or rotating the viewer (seen at right) actually transforms the camera or light you're viewing through (seen at left)

It is even possible to view the scene from the perspective of a Merge3D or Transform3D node by selecting the object from the Camera > Others menu. The same transform techniques will then move the position of the object. This can be helpful when you are trying to orient an object in a certain direction.

Transparency Sorting

While generally the order of geometry in a 3D scene is determined by the Z-position of each object, sorting every face of every object in a large scene can take an enormous amount of time. To provide the best possible performance, a Fast Sorting mode is used in the OpenGL renderer and viewers. This is set by right-clicking in the viewer and choosing Transparency > Z-buffer. While this approach is much faster than a full sort, when objects in the scene are partially transparent it can also produce incorrect results.

The Sorted (Accurate) mode can be used to perform a more accurate sort at the expense of performance. This mode is selected from the Transparency menu of the viewer's contextual menu. The Renderer3D also presents a Transparency menu when the Renderer Type is set to OpenGL. Sorted mode does not support shadows in OpenGL. The software renderer always uses the Sorted (Accurate) method.

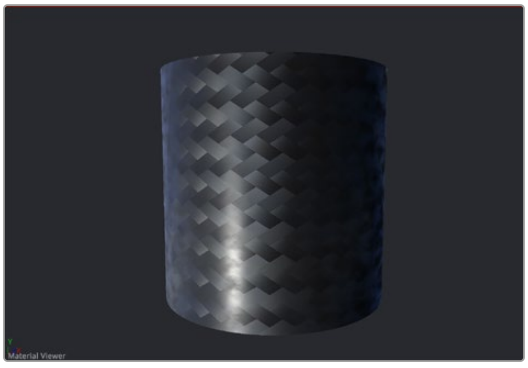


Transparency Sorting in the viewer contextual menu

The basic rule is when a scene contains overlapping transparency, use the Full/Quick Sort modes, and otherwise use the Z-buffer (Fast). If the Full Sort method is too slow, try switching back to Z-buffer (Fast).

Material Viewer

When you view a node that comes from the 3D > Material category of nodes in the Effects Library, the viewer automatically switches to display a Material Viewer. This Material Viewer allows you to preview the material applied to a lit 3D sphere rendered with OpenGL by default.



The Material Viewer mode of the viewer

The type of geometry, the renderer, and the state of the lighting can all be set by right-clicking the viewer and choosing options from the contextual menu. Each viewer supports A and B buffers to assist with comparing multiple materials.

Methods of working with the Material Viewer:

- You can change the shape of the previewed geometry by right-clicking the viewer and choosing an option from the Shape submenu of the contextual menu. The geometry that the material is applied to is locked to the center of the viewer and scaled to fit. It is not possible to pan or scale the Material Viewer.
- The Material Viewer can be rotated to provide a different angle on the material by holding Option while pressing the middle mouse button and dragging to the left and right.
- You can adjust the position of the light used to preview the material by dragging with the middle mouse button. Or, you can right-click the viewer and choose an option from the Lighting > Light Position submenu of the contextual menu.
- You can also toggle lighting off and on by right-clicking the viewer and choosing Lighting > Enable Lighting from the contextual menu.
- You can choose the renderer used to preview the material by right-clicking the viewer and choosing an option from the Renderer submenu of the contextual menu.

Transformations

Merge3D, 3D Objects, and Transform3D all have Transform parameters that are collected together into a Transform tab in the Inspector. The parameters found in this tab affect how the object is positioned, rotated, and scaled within the scene.



The Transform tab of a Merge3D node

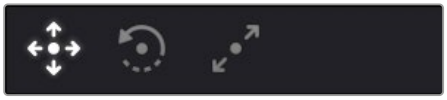
The Translation parameters are used to position the object in local space, the Rotation parameters affect the object's rotation around its own center, and the Scale slider(s) affect its size (depending on whether or not they're locked together). The same adjustments can be made in the viewer using onscreen controls.

Onscreen Transform Controls

When an object is selected, it displays onscreen Transform controls in the viewers that allow you to adjust the object's position, rotation, and scale. Buttons in the Transform toolbar allow you to switch modes, or you can use the keyboard shortcuts.

To switch Transform modes, use the following keyboard shortcuts:

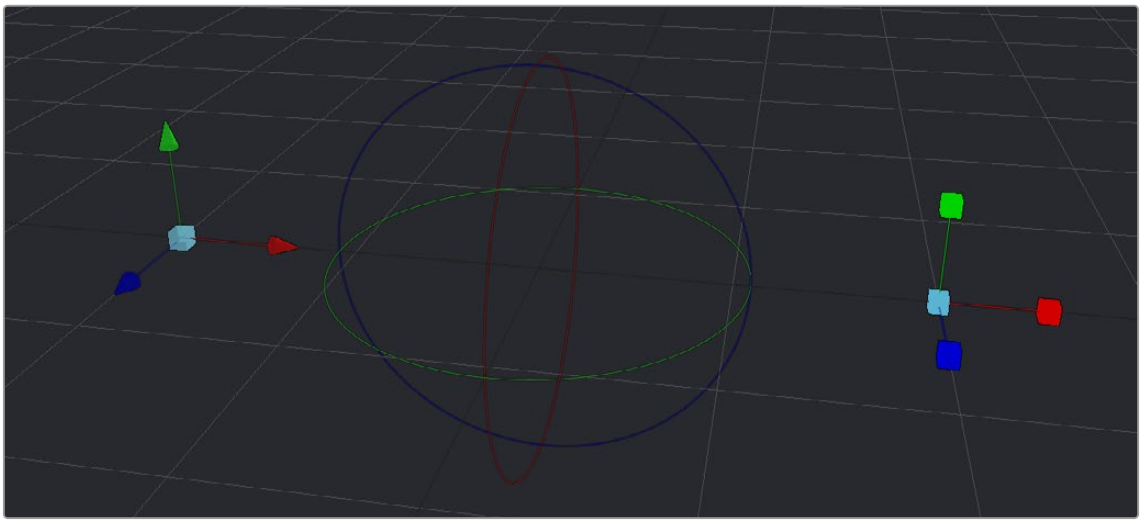
- Press Q for Position
- Press W for Rotation
- Press E for Scaling



The Position, Rotation, and Scale modes in the Transform toolbar

Using Onscreen Transform Controls

In all three modes, red indicates the object's local X-axis, green the Y-axis, and blue the Z-axis, respectively (just remember RGB = XYZ). You can drag directly on the red, green, or blue portion of any onscreen control to constrain the transform to that axis, or if you drag the center of the onscreen control, you can apply a transform without constraints. Holding Option and dragging in the viewer allows you to freely translate in all three axes without clicking on a specific control.



From left to right, the Position, Rotation, and Scale onscreen Transform controls

If the Scale's Lock XYZ checkbox is enabled in the Inspector, only the overall scale of the object is adjusted by dragging the red or center onscreen control, while the green and blue portions of the onscreen control have no effect. If you unlock the parameters, you are able to scale an object along individual axes separately to squish or stretch the object.

Selecting Objects

With the onscreen controls visible in the viewer, you can select any object by clicking on its center control. Alternatively, you can also select any 3D object by clicking its node in the Node Editor.

Pivot

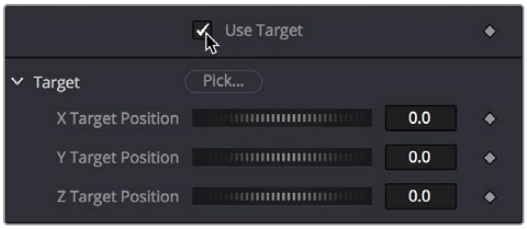
In 3D scenes, objects rotate and scale around an axis called a pivot. By default, this pivot goes through the object's center. If you want to move the pivot so it is offset from the center of the object, you can use the X, Y, and Z Pivot parameters in the Inspector.

Target

Targets are used to help orient a 3D object to a specific point in the scene. No matter where the object moves, it will rotate in the local coordinate system so that it always faces its target, which you can position and animate.

To enable a target for a 3D object:

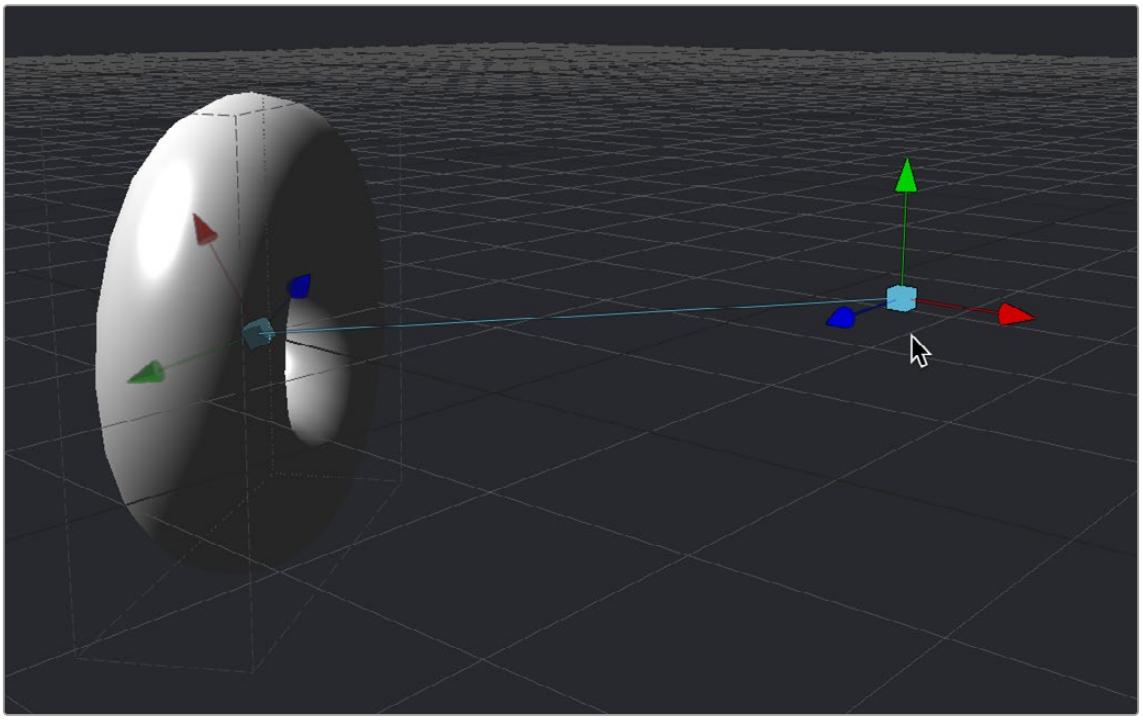
1. Select that object's node.
2. Open the object's Transform panel in the Inspector.
3. Turn on the Use Target checkbox.



Turning on the Use Target checkbox of a 3D object

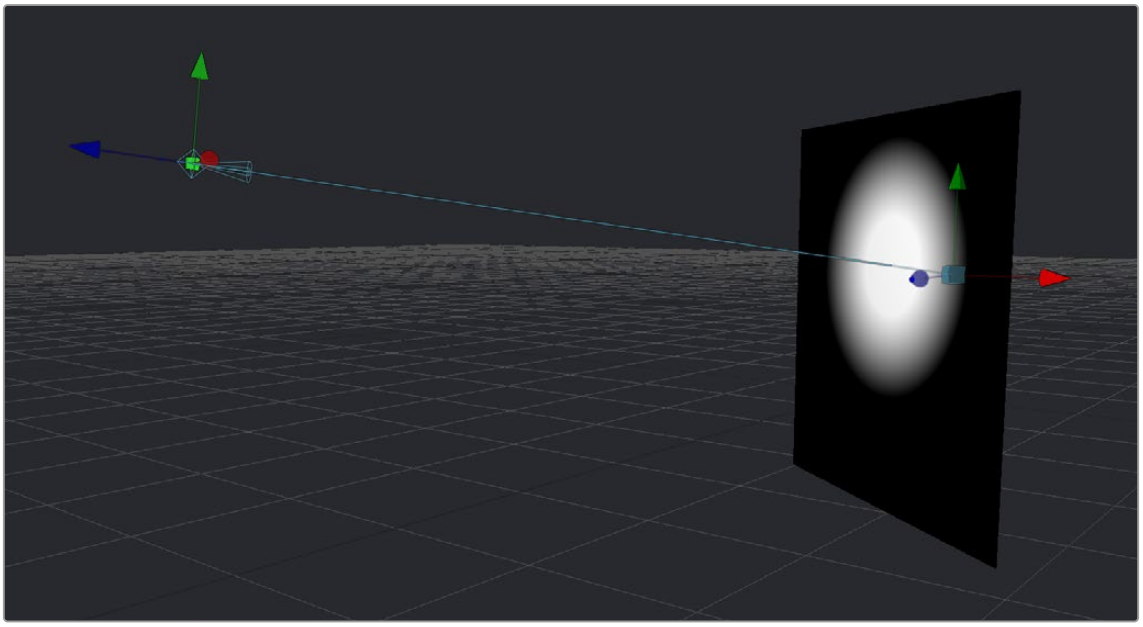
4. Use the X/Y/Z Target Position controls in the Inspector or the Target onscreen control in the viewer to position the target and in turn position the object it's attached to.

In the viewer, a line is drawn between the target and the center of the 3D object it's attached to, to show the relationship between these two sets of controls. Whenever you move the target, the object is automatically transformed to face its new position.



A torus facing its onscreen Target controls

For example, if a spotlight is required in the scene to point at an image plane, enable the spotlight's target in the Transform tab and connect the target's XYZ position to the image plane's XYZ position. Now, no matter where the spotlight is moved, it will rotate to face the image plane.

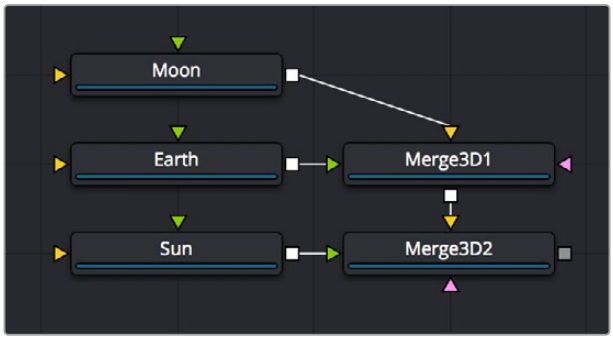


A light made to face the wall using its enabled target control

Parenting

One of the many advantages of the node-based approach to 3D compositing is that parenting between objects becomes implicit in the structure of a 3D node tree. The basis for all parenting is the Merge3D node. If you're careful about how you connect the different 3D objects you create for your scene, you can use multiple Merge3D nodes to control which combinations of objects are transformed and animated together, and which are transformed and animated separately.

For example, picture a scene with two spheres that are both connected to a Merge3D. The Merge3D can be used to rotate one sphere around the other, like the moon around the earth. Then the Merge3D can be connected to another Merge3D to create the earth and the moon orbiting around the sun.



One Merge3D with two spheres parented to another Merge3D and parenting using three connected spheres

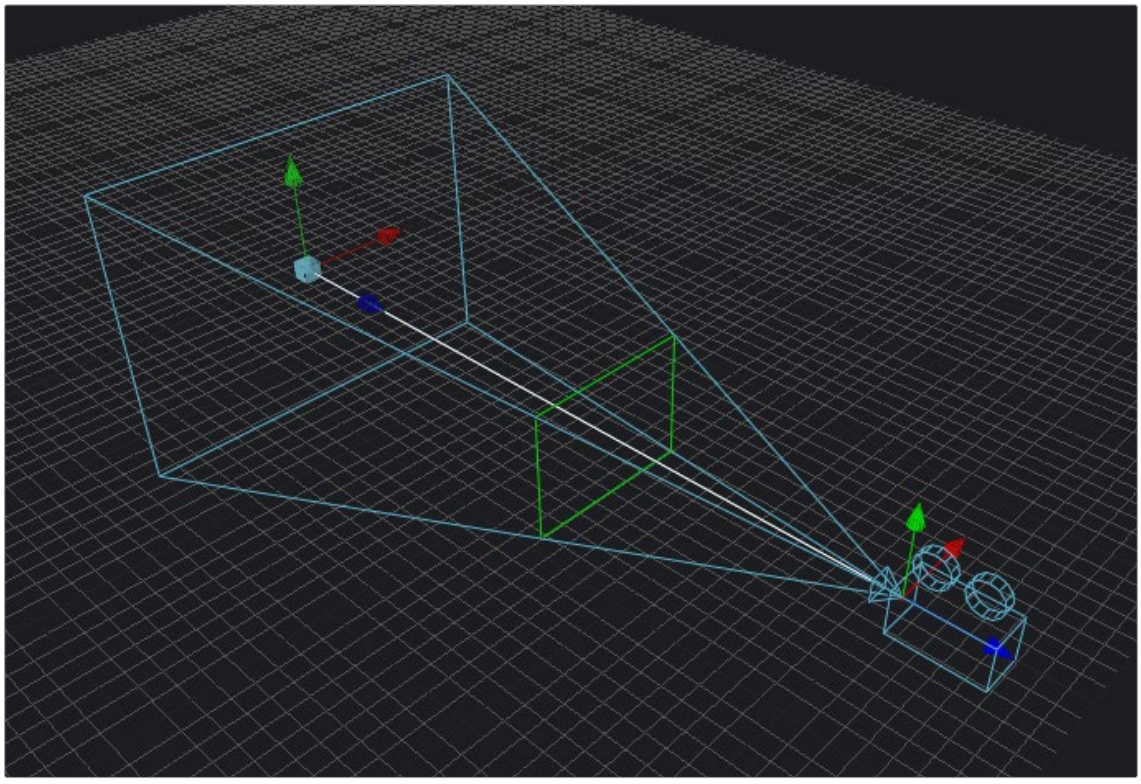
Here are the two simple rules of transforming parented Merge3D nodes:

- Transforms and animation applied to a Merge3D are also applied to all 3D objects connected to that Merge3D node, including cameras, lights, geometry, and other merge nodes connected upstream.
- Transforms and animation applied to upstream merge nodes don't affect downstream merge nodes.

Cameras

When setting up and animating a 3D scene, the metaphor of a camera is one of the most comprehensible ways of framing how you want that scene to be rendered out, as well as animating your way through the scene. Additionally, compositing artists are frequently tasked with matching cameras from live-action clips, or matching cameras from 3D applications.

To accommodate all these tasks, Fusion provides a flexible Camera3D node with common camera controls such as Angle of View, Focal Length, Aperture, and Clipping planes, to either set up your own camera or to import camera data from other applications. The Camera3D node is a virtual camera through which the 3D environment can be viewed.



A camera displayed with onscreen Transform controls in the viewer; the Focal Plane indicator is enabled in green

Cameras are typically connected and viewed via a Merge3D node; however, you can also connect cameras upstream of other 3D objects if you want that camera to transform along with that object when it moves.

Quickly Viewing a Scene Through a Camera

When you've added a camera to a scene, you can quickly view the scene "through the camera" by setting up the following.

To view the scene through the camera:

1. Select the Merge3D node that the camera is connected to, or any node downstream of that Merge3D.
2. Load the selected Merge3D or downstream node into a viewer.
3. Right-click on the axis label in the bottom corner of the viewer and choose the camera name.

The viewer's frame may be different from the camera frame, so it may not match the true boundaries of the image that will be rendered by the Renderer3D node. If there is no Renderer3D node added to your scene yet, you can use Guides that represent the camera's framing.

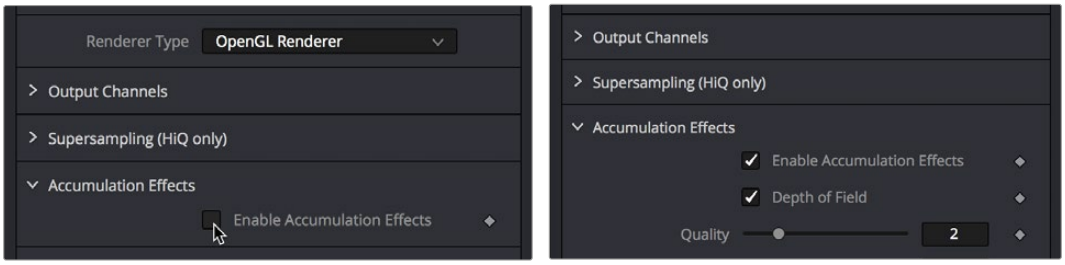
> For more information about Guides, see *Chapter 7, "Using Viewers."* in the Fusion Reference Manual.

Plane of Focus and Depth of Field

Cameras have a plane of focus, for when depth of field rendering is available. Here's the procedure for enabling depth of field rendering in your scenes.

To render depth of field in a 3D scene:

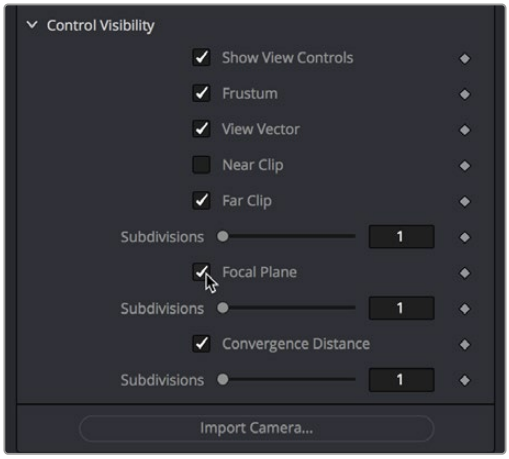
1. You must add a Renderer3D node at the end of your 3D scene.
2. Select the Renderer3D node, and set the Renderer Type to OpenGL Renderer.
3. Open the Accumulation Effects disclosure control that appears, and turn on the Enable Accumulation Effects checkbox in the OpenGL render.



Turning on Enable Accumulation Effects enables additional depth of field controls

Turning on "Enable Accumulation Effects" exposes a Depth of Field checkbox along with Quality and Amount of DoF Blur sliders that let you adjust the depth of field effect. These controls affect only the perceived quality of the depth of field that is rendered. The actual depth of field that's generated depends solely on the setup of the camera and its position relative to the other 3D objects in your scene.

When you select your scene's Camera3D node to view its controls in the Inspector, a new Focal Plane checkbox appears in the Control Visibility group. Turning this on lets you see the green focal plane indicator in the 3D Viewer that lets you visualize the effect of the Focal Plane slider, which is located in the top group of parameters in the Camera3D node's Controls tab.



Turning on the Focal Plane checkbox in the Camera3D node

For more information about these specific camera controls, *see Chapter 29, "3D Nodes,"* in the Fusion Reference Manual.

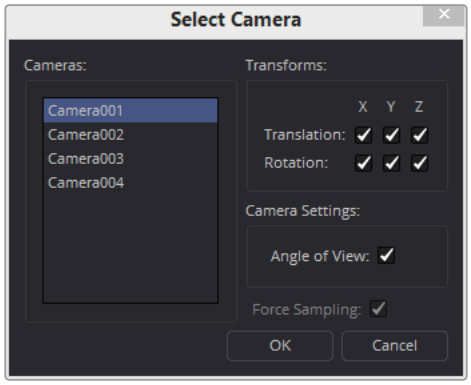
Importing Cameras

If you want to match cameras between applications, you can import camera paths and positions from a variety of popular 3D applications. Fusion is able to import animation splines from Maya and XSI directly with their own native spline formats. Animation applied to cameras from 3ds Max and LightWave are sampled and keyframed on each frame.

To import a camera from another application, do the following:

1. Select the camera in the Node Editor.
2. At the bottom of the Inspector, click the Import Camera button.
3. In the file browser, navigate to and select the scene that contains the camera you want to import.

   A dialog box with several options will appear. When the Force Sampling checkbox is enabled, Fusion will sample each frame of the motion, regardless of the format.



The Select Camera dialog

TIP: When importing parented or rigged cameras, baking the camera animation in the 3D application before importing it into Fusion often produces more reliable results.

Lighting and Shadows

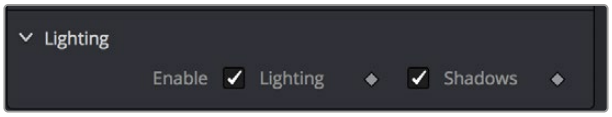
You can add light sources to a scene to create very detailed lighting environments and atmosphere. There are four different types of lights you can use in 3D scenes: ambient, directional, point, and spotlights.

Enabling Lighting in the Viewer

A scene without lights uses a default directional light, but this automatically disappears once you add a 3D light object. However, even when you add light objects to your scene, lighting and shadows won't be visible in the viewer unless you first enable lighting in the viewer contextual menu by right-clicking anywhere within a viewer and choosing 3D Options > Lighting or Shadows to turn on one or both.

Enabling Lighting to Be Rendered

Lighting effects won't be rendered in the Renderer3D node until the Enable Lighting and/or Shadows checkboxes are checked in the Inspector.



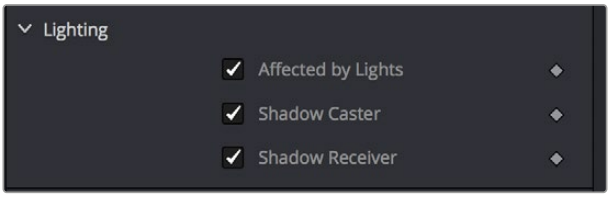
The Lighting button under the viewer

NOTE: When lighting is disabled in either the viewer or final renders, the image will appear to be lit by a 100% ambient light.

Cotrolling Lighting within Each 3D Object

All nodes that create or merge geometry also include lighting options that are used to choose how each object is affected by light:

- Merge3D nodes have a Pass Through Lights checkbox that determines whether lights attached to an upstream Merge3D node also illuminate objects attached to downstream Merge3D nodes.
- ImagePlane3D, Cube3D, Shape3D, Text3D, and FBXMesh3D nodes have a set of Lighting controls that let you turn three controls on and off: Affected by Lights, Shadow Caster, and Shadow Receiver.



3D objects have individual lighting controls that let you control how each object interacts with light and shadows

Lighting Types Explained

Here's a more detailed explanation of each type of light in Fusion.

Ambient Light

You use ambient light to set a base light level for the scene, since it produces a general uniform illumination of the scene. Ambient light exists everywhere without appearing to come from any particular source; it cannot cast shadows and will tend to fill in shadowed areas of a scene.

Directional Light

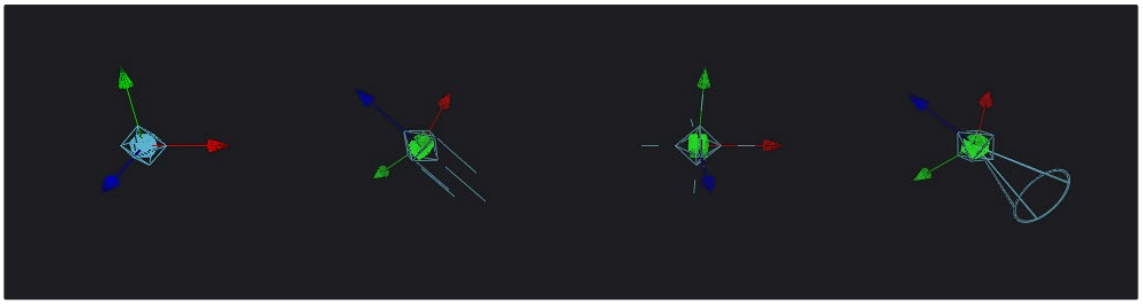
A directional light is composed of parallel rays that light up the entire scene from one direction, creating a wall of light. The sun is an excellent example of a directional light source.

Point Light

A point light is a well defined light that has a small clear source, like a light bulb, and shines from that point in all directions.

Spotlight

A spotlight is an advanced point light that produces a well defined cone of light with falloff. This is the only light that produces shadows.

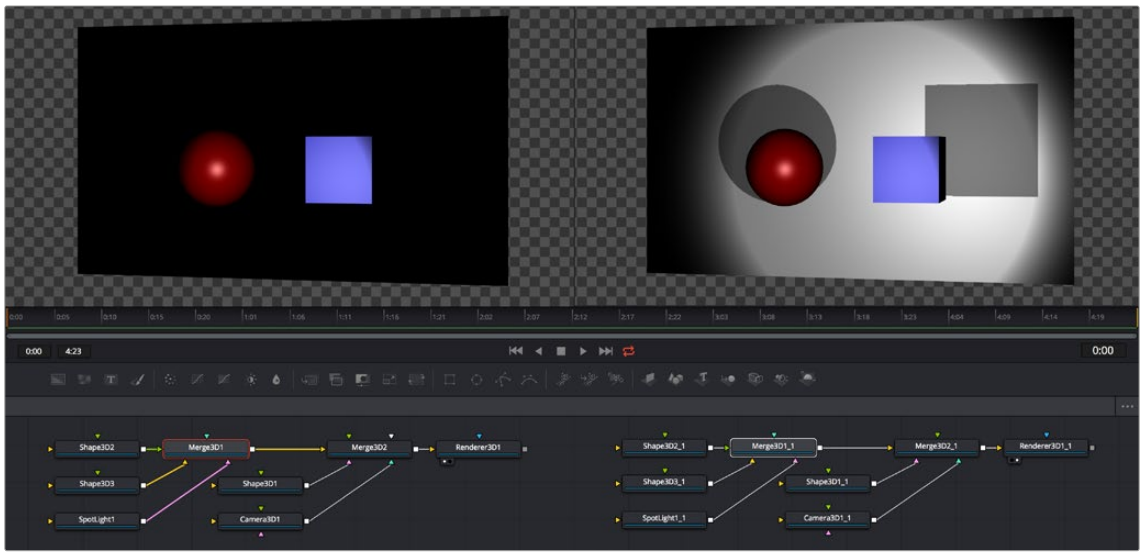


From left to right: Directional light, point light, and spotlight

All of the Light nodes display onscreen controls in the viewer, although not all controls affect every light type. In the case of the ambient light, the position has no effect on the results. The directional light can be rotated, but position and scale will be ignored. The point light ignores rotation. Both position and rotation apply to the spotlight.

Lighting Hierarchies

Lights normally do not pass through a Merge, since the Pass Through Lights checkbox is off by default. This provides a mechanism for controlling which objects are lit by which lights. For example, in the following two node trees, two shapes and an ambient light are combined with a Merge3D node, which is then connected to another Merge3D node that's also connected to a plane and a spotlight. At the left, the first Merge3D node of this tree has Pass Through Lights disabled, so you can only see the two shapes lit. At the right, Pass Through Lights has been enabled, so both the foreground shapes and the background image plane receive lighting.



Pass Through Lights is disabled, so only the front two shapes are illuminated (left) Pass Through Lights is enabled, so all shapes connected to both Merge3D nodes are illuminated (right)

Lighting Options

Most nodes that generate geometry have additional options for lighting. These options are used to determine how each individual object reacts to lights and shadows in the scene.

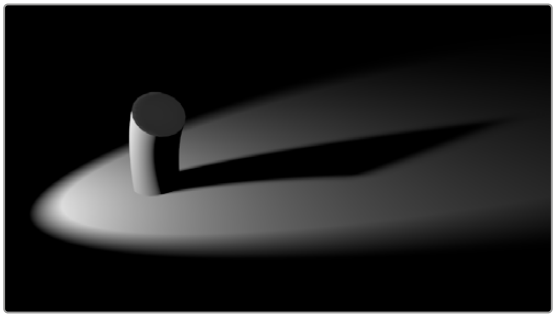
- **Affected By Lights:** If the Affected By Lights checkbox is enabled, lights in the scene will affect the geometry.
- **Shadow Caster:** When enabled, the object will cast shadows on other objects in the scene.
- **Shadow Receiver:** If this checkbox is enabled, the object will receive shadows.



3D objects have individual lighting controls that let you control how each object interacts with light and shadows

Shadows

Spotlight, Point Light, and Directional light can all cast shadows in the 3D viewer, however Point and Directional lights need to have Hardware Renderer set as the Rendering Type in the Renderer3D node. Spotlight can work with either the Hardware or Software renderers. These lights cast shadows by default, although these shadows will not be visible in the viewer until shadows are enabled using the viewer toolbar button. Shadows will not appear in the output of the Renderer3D unless the Shadows option is enabled for that renderer. If you want to prevent a spotlight from casting shadows, you can disable the Enable Shadows checkbox in the node's Inspector.



An image with spotlight casting a variable soft shadow

For more information on shadow controls, see the "Spotlight" section of *Chapter 30, "3D Light Nodes,"* in the Fusion Reference Manual.

Shadow Maps

A shadow map is an internal depth map that specifies each pixel's depth in the scene. This information is used to assemble the shadow layer created from a spotlight. All the controls for the shadow map are found in the Spotlight Inspector.

The quality of the shadow produced depends greatly on the size of the shadow map. Larger maps generate better-looking shadows but will take longer to render. The wider the cone of the spotlight, or the more falloff in the cone, the larger the shadow map will need to be to produce useful quality results. Setting the value of the Shadow Map Size control sets the size of the depth map in pixels.

Generally, through trial and error, you'll find a point of diminishing returns where increasing the size of the shadow map no longer improves the quality of the shadow. It is not recommended to set the size of the shadow maps any larger than they need to be.

The Shadow Map Proxy control is used to set a percentage by which the shadow map is scaled for fast interactive previews, such as Autoproxy and LoQ renders. A value of 4, for example, represents a 40% proxy.

Shadow Softness

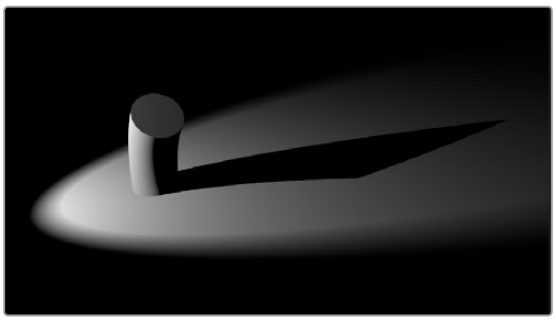
By default, the spotlight generates shadows without soft edges, but there are options for constant and variable soft shadows. Hard-edged shadows will render significantly faster than either of the Soft Shadow options. Shadows without softness will generally appear aliased, unless the shadow map size is large enough. In many cases, softness is used to hide the aliasing rather than increasing the shadow map to preserve memory and avoid exceeding the graphics hardware capabilities.



Soft Shadow controls in the Control panel

Setting the spotlight's shadow softness to None will render crisp and well-defined shadows. The Constant option will generate shadows where the softness is uniform across the shadow,

regardless of the shadow's distance from the casting geometry. The Variable option generates shadows that become softer as they get farther from the geometry that is casting the shadow. This is a more realistic effect, but the shadows are somewhat harder to control. When this option is selected, additional controls for adjusting the falloff of the shadow will appear, as well as sliders for the minimum and maximum softness.



Hard shadow cast by a spotlight

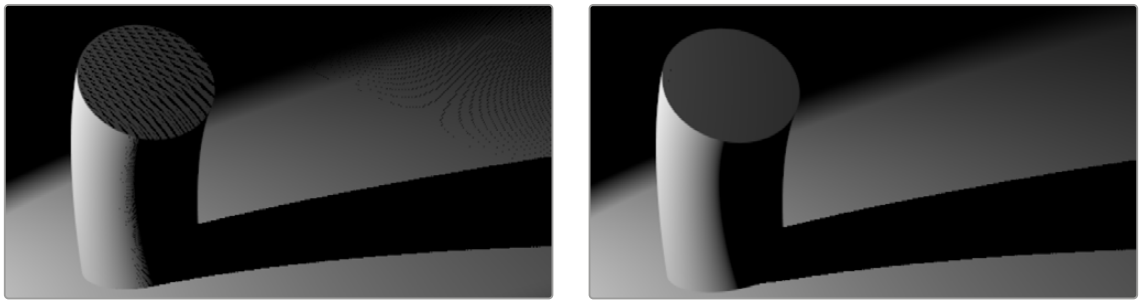
Selecting the Variable option reveals the Spread, Min Softness, and Filter Size sliders. A side effect of the method used to produce variable softness shadows is that the size of the blur applied to the shadow map can become effectively infinite as the shadow's distance from the geometry increases. These controls are used to limit the shadow map by clipping the softness calculation to a reasonable limit.

The filter size determines where this limit is applied. Increasing the filter size increases the maximum possible softness of the shadow. Making this smaller can reduce render times but may also limit the softness of the shadow or potentially even clip it. The value is a percentage of the shadow map size.

For more information, see "Spotlight" in *Chapter 30, "3D Light Nodes,"* in the Fusion Reference Manual.

Multiplicative and Additive Bias

Shadows are essentially textures applied to objects in the scene that occasionally result in "fighting." Z-fighting results when portions of an object that should be receiving shadows instead render over the top of the shadow because they effectively exist in the same exact location in 3D space.



Results of shadow map Z-fighting (Left), and the corrected shadow shown using Biasing (Righht)

Two Biasing sliders in the Shadows group of Spotlight parameters work by adding a small depth offset to move the shadow away from the surface it is shadowing, eliminating the Z-fighting. When too little bias is added, the objects can self shadow themselves. When too much is added, the shadow can become separated from the surface.



The Multiplicative and Additive Bias sliders, and the Non-Transmissive Materials checkbox, all in the Spotlight Inspector controls

The goal is to adjust the Multiplicative Bias slider until the majority of the Z-fighting is resolved, and then adjust the Additive Bias slider to eliminate the rest. The softer the shadow, the higher the bias will probably have to be. You may even need to animate the bias to get a proper result for some particularly troublesome frames.

Force All Materials Non-Transmissive

How light passes through a semi-transparent material plays an important role in determining the appearance of the shadow an object casts. Normally, this transmittance behavior is defined in each object's Materials tab. However, selecting Force All Materials Non-Transmissive in the Spotlight Inspector overrides this, causing the shadow map produced by the node to ignore transmittance entirely.

Materials and Textures

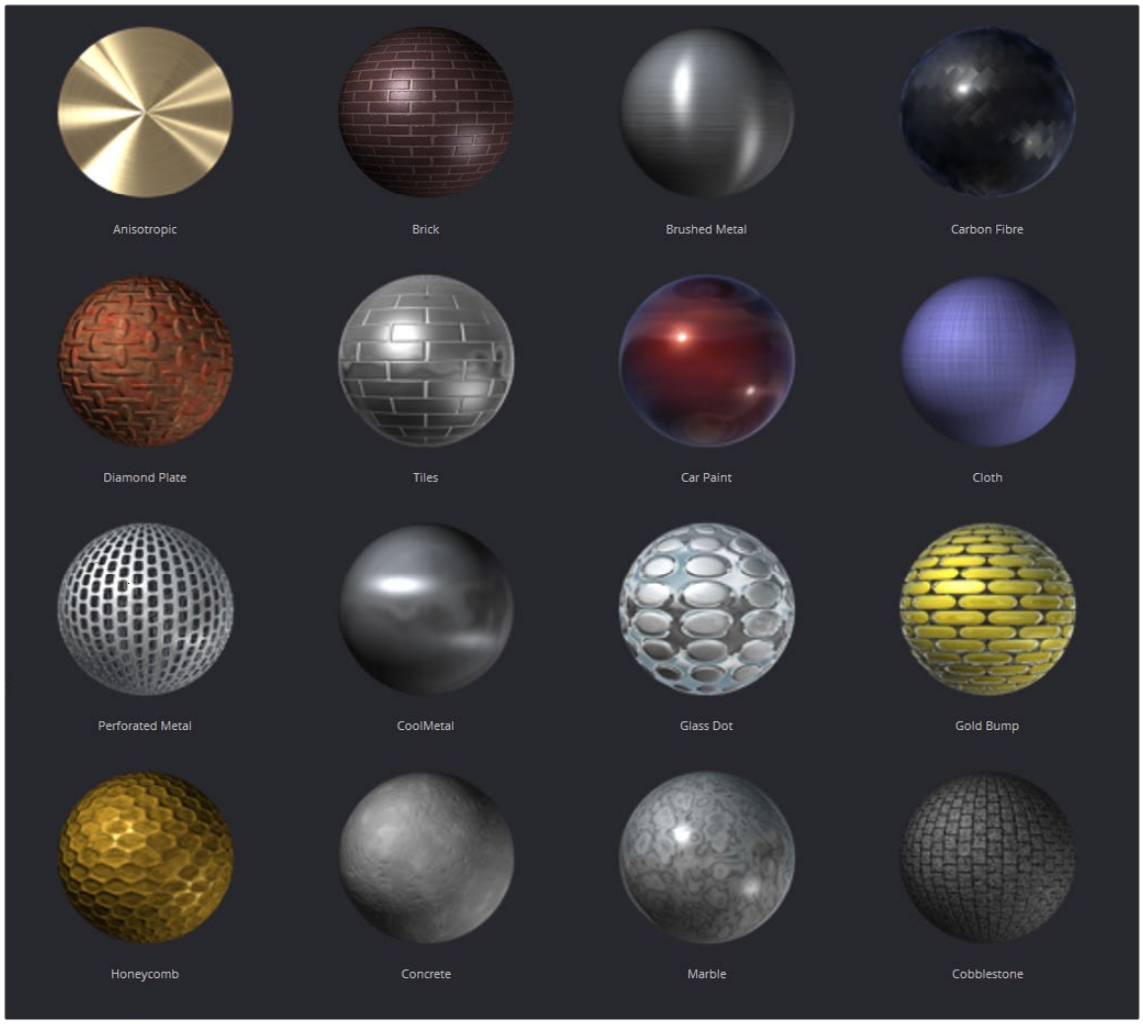
To render a 3D scene, the renderer must take into account the shape of the object as well as its appearance. The geometry of an object defines the shape of the object, while the material applied to the object defines its appearance. Fusion provides a range of options for applying materials and textures to geometry, so you can give your 3D objects the surface qualities you want.

Nodes that describe the geometry's response to light are called illumination models. Blinn, Cook-Torrance, Ward, and Phong are the included illumination models. These nodes are found in the 3D > Material category of nodes in the Effects Library.

Most materials also accept textures, which are typically 2D images. Textures are used to refine the look of an object further, by adding photorealistic details, transparency, or special effects. More complex textures like bump maps, 3D textures, and reflection maps are also available in the 3D > Texture category.

Materials can also be combined to produce elaborate and highly detailed composite materials.

Each node that creates or loads geometry into a 3D scene also assigns a default material. The default material is the Blinn illumination model, but you can override this material using one of several nodes that output a 3D material. Some of these materials provide a greater degree of control over how the geometry reacts to light, providing inputs for diffuse and specular texture maps, bump mapping, and environmental maps, which mimic reflection and refraction.



Material examples from the bin

Material Components

All the standard illumination models share certain characteristics that must be understood.

Diffuse

The Diffuse parameters of a material control the appearance of an object where light is absorbed or scattered. This diffuse color and texture are the base appearance of an object, before taking into account reflections. The opacity of an object is generally set in the diffuse component of the material.

Alpha

The Alpha parameter defines how much the object is transparent to diffuse light. It does not affect specular levels or color. However, if the value of alpha, either from the slider or a Material input from the diffuse color, is very close to or at zero, those pixels, including the specular highlights, will be skipped and disappear.

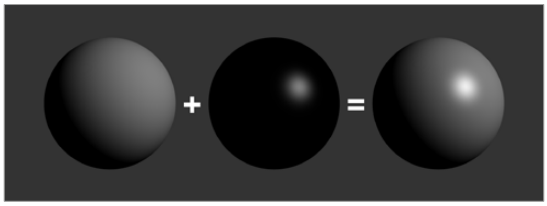
Opacity

The Opacity parameter fades out the entire material, including the specular highlights. This value cannot be mapped; it is applied to the entire material.

Specular

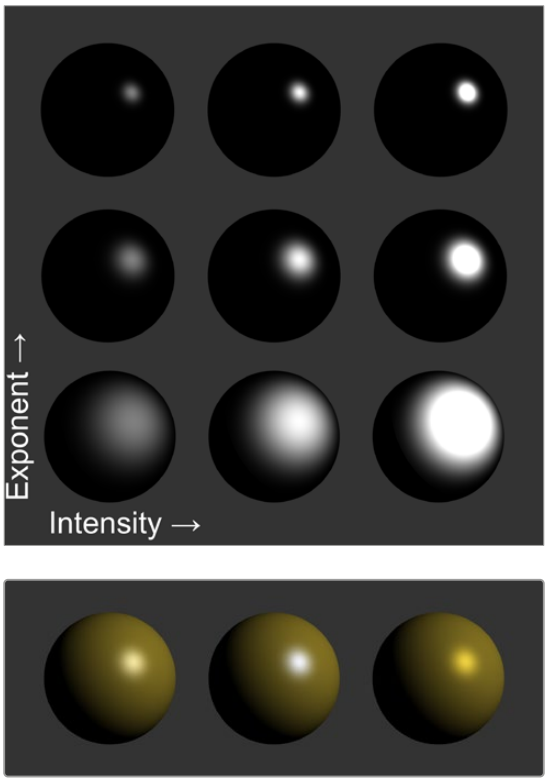
The Specular parameters of a material control the highlight of an object where the light is reflected to the current viewpoint. This causes a highlight that is added to the diffuse component. The more specular a material is, the glossier it appears. Surfaces like plastics and glass tend to have white specular highlights, whereas metallic surfaces like gold have specular highlights that tend to inherit their color from the material color.

Specularity is made up of color, intensity, and exponent. The specular color determines the color of light that reflects from a shiny surface. Specular intensity is how bright the highlight will be.



Three spheres, left to right: diffuse only, specular only, and combined

The specular exponent controls the falloff of the specular highlight. The larger the value, the sharper the falloff and the smaller the specular component will be.



Left to right: white, complimentary, and matching specular colors

Transmittance

When using the software renderer, the Transmittance parameters control how light passes through a semi-transparent material. For example, a solid blue pitcher will cast a black shadow, but one made of translucent blue plastic would cast a much lower density blue shadow. The transmittance parameters are essential to creating the appearance of stained glass.

TIP: You can adjust the opacity and transmittance of a material separately. It is possible to have a surface that is fully opaque yet transmits 100% of the light arriving upon it, so in a sense it is actually a luminous/emissive surface.

Transmissive surfaces can be further limited using the Alpha and Color Detail control.

Attenuation

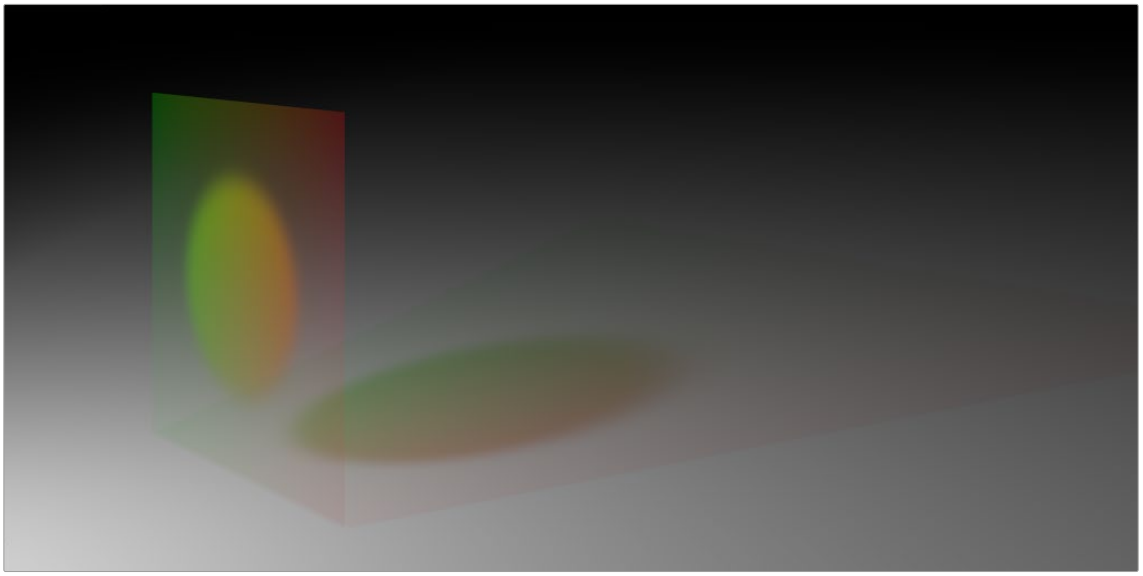
The transmittance color determines how much color is passed through the object. For an object to have fully transmissive shadows, the transmittance color must be set to to RGB = (1, 1, 1), which means 100% of green, blue, and red light passes through the object. Setting this color to RGB = (1, 0, 0) means that the material will transmit 100% of the red arriving at the surface but none of the green or blue light.

Alpha Detail

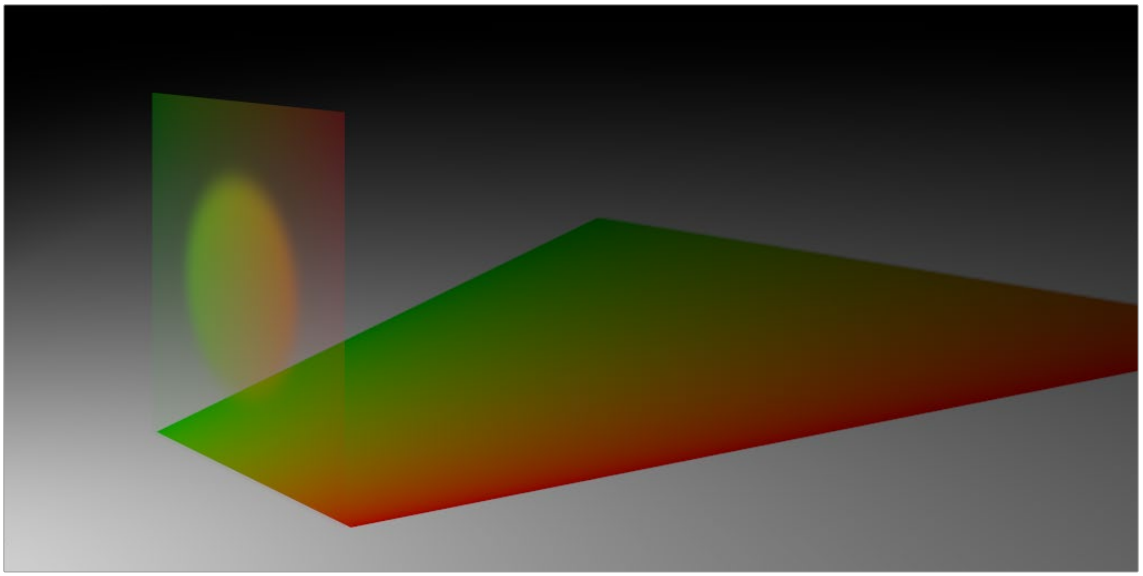
When the Alpha Detail slider is set to 0, the non-zero portions of the alpha channel of the diffuse color are ignored and the opaque portions of the object casts a shadow. If it is set to 1, the alpha channel determines how dense the object casts a shadow.

NOTE: The OpenGL renderer ignores alpha channels for shadow rendering, resulting in a shadow always being cast from the entire object. Only the software renderer supports alpha in the shadow maps.

The following examples for Alpha Detail and Color Detail cast a shadow using this image. It is a green-red gradient from left to right. The outside edges are transparent, and inside is a small semi-transparent circle.



Alpha Detail set to 1; the alpha channel determines the density of the shadow

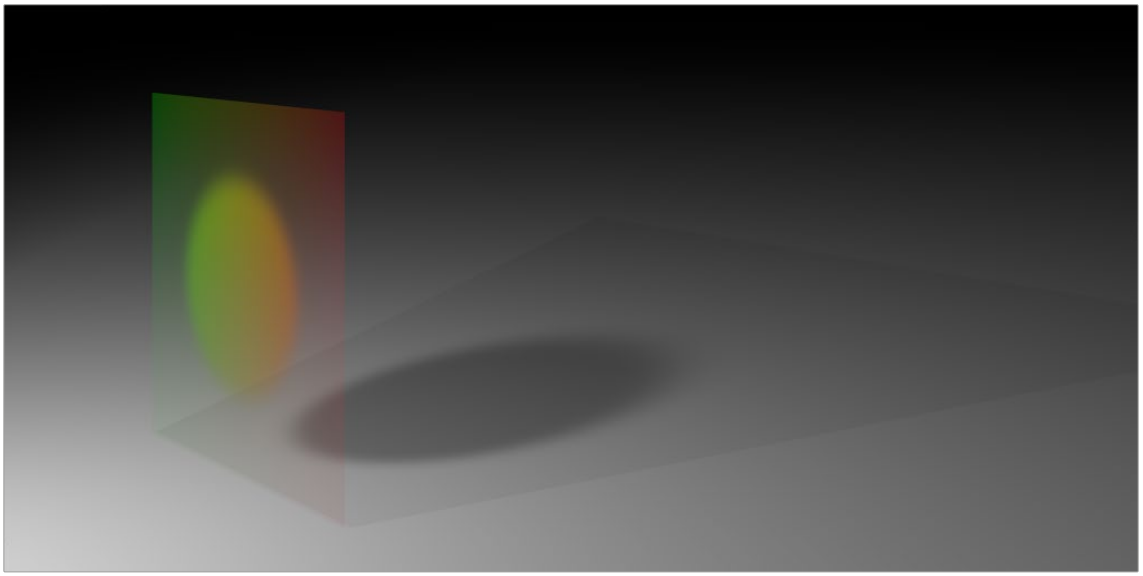


Alpha Detail set to 0; a dense-colored shadow results

Color Detail

Color Detail is used to color the shadow with the object's diffuse color. Increasing the Color Detail slider from 0 to 1 brings in more diffuse color and texture into the shadow.

> **TIP:** The OpenGL renderer will always cast a black shadow from the entire object, ignoring the color. Only the software renderer supports color in the shadow maps.



Color Detail set to 0; no color is visible in the shadow.

Saturation

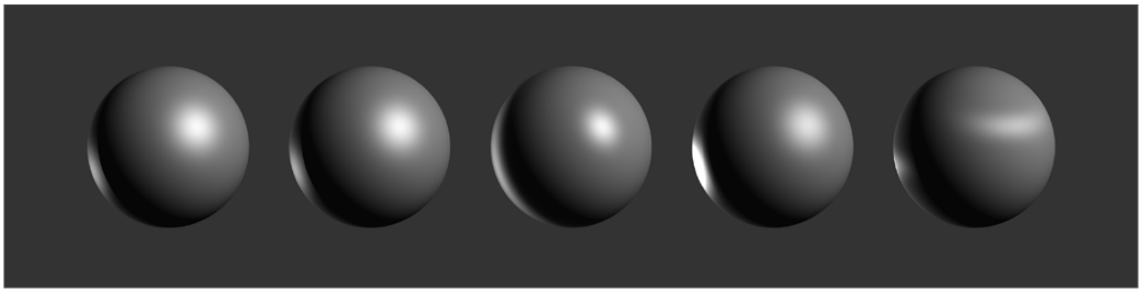
Saturation will allow the diffuse color texture to be used to define the density of the shadow without affecting the color. This slider lets you blend between the full color and luminance only.

Transmittance and Shadows

The transmittance of an object's material plays an important role in determining the appearance of the shadow it casts. Normally, the transmittance behavior is defined in each object's Materials tab as explained above. However, selecting Force All Materials Non-Transmissive in the Spotlight Inspector overrides this, causing the shadow map produced by the spotlight to ignore transmittance entirely.

Illumination Models

Now that you understand the different components that make up a material or shader, we'll look at them more specifically. Illumination models are advanced materials for creating realistic surfaces like plastic, wood, or metal. Each illumination model has advantages and disadvantages, which make it appropriate for particular looks. An illumination model determines how a surface reacts to light, so these nodes require at least one light source to affect the appearance of the object. Four different illumination models can be found in the Nodes > 3D > Material menu.



Illumination models left to right: Standard, Blinn, Phong, Cook-Torrance, and Ward

Standard

The Standard material provides a default Blinn material with basic control over the diffuse, specular, and transmittance components. It only accepts a single texture map for the diffuse component with the alpha used for opacity. The Standard Material controls are found in the Material tab of all nodes that load or create geometry. Connecting any node that outputs a material to that node's Material Input will override the Standard material, and the controls in the Material tab will be hidden.

Blinn

The Blinn material is a general purpose material that is flexible enough to represent both metallic and dielectric surfaces. It uses the same illumination model as the Standard material, but the Blinn material allows for a greater degree of control by providing additional texture inputs for the specular color, intensity, and exponent (falloff), as well as bump map textures.

Phong

The Phong material produces the same diffuse result as Blinn, but with wider specular highlights at grazing incidence. Phong is also able to make sharper specular highlights at high exponent levels.

Cook-Torrance

The Cook-Torrance material combines the diffuse illumination model of the Blinn material with a combined microfacet and Fresnel specular model. The microfacets need not be present in the mesh or bump map; they are represented by a statistical function, Roughness, which can be mapped. The Fresnel factor attenuates the specular highlight according to the Refractive Index, which can be mapped.

Ward

The Ward material shares the same diffuse model as the others but adds anisotropic highlights, ideal for simulating brushed metal or woven surfaces, as the highlight can be elongated in the U or V directions of the mapping coordinates. Both the U and V spread functions are mappable.

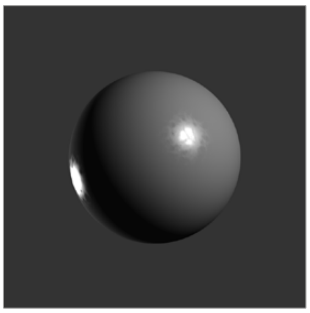
This material does require properly structured UV coordinates on the meshes it is applied to.

Textures

Texture maps modify the appearance of a material on a per-pixel basis. This is done by connecting an image or other material to the inputs on the Material nodes in the Node Editor. When a 2D image is used, the UV mapping coordinates of the geometry are used to fit the image to the geometry, and when each pixel of the 3D scene is rendered, the material will modify the material input according to the value of the corresponding pixel in the map.

TIP: UV Mapping is the method used to wrap a 2D image texture onto 3D geometry. Similar to X and Y coordinates in a frame, U and V are the coordinates for textures on 3D objects.

Texture maps are used to modify various material inputs, such as diffuse color, specular color, specular exponent, specular intensity, bump map, and others. The most common uses of texture maps is the diffuse color/opacity component.



The Fast Noise texture used to control the roughness of a Cook-Torrance material

A node that outputs a material is frequently used, instead of an image, to provide other shading options. Materials passed between nodes are RGBA samples; they contain no other information about the shading or textures that produced them.

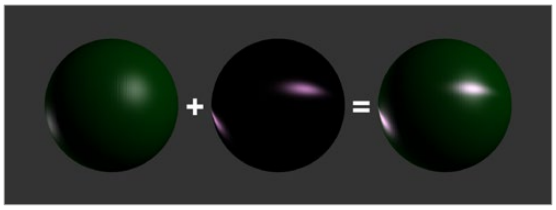


The Texture2D node is used to translate a texture in the UV space of the object, as well as set the filtering and wrap mode

Composite Materials

Building complex materials is as easy as connecting the output of a Material node to one of the Material inputs of another Material or Texture node. When a Material input is supplied just as with a 2D image, its RGBA values are used per pixel as a texture. This allows for very direct compositing of shaders.

For instance, if you want to combine an anisotropic highlight with a Blinn material, you can take the output of the Blinn, including its specular, and use it as the diffuse color of the Ward material. Or, if you do not want the output of the Blinn to be relit by the Ward material, you can use the Channel Boolean material to add the Ward material's anisotropic specular component to the Blinn material with a greater degree of control.



Combining an anisotropic highlight with a Blinn material using the Channel Boolean material

Reflections and Refractions

Environment maps can be applied with the Reflect material in the 3D > Material category. This node can be used to simulate reflections and refractions on an object. Reflections are direct-bounce light that hits an object, while refractions simulate the distortion of light seen through semi-translucent surfaces.

The reflections and refractions use an environment mapping technique to produce an approximation that balances realistic results with greater rendering performance. Environment maps assume an object's environment is infinitely distant from the object and rendered into a cubic or spherical texture surrounding the object.

The Nodes > 3D > Texture > Cube Map and Sphere Map nodes can be used to help create environment maps, applying special processing and transforms to create the cubic or spherical coordinates needed.



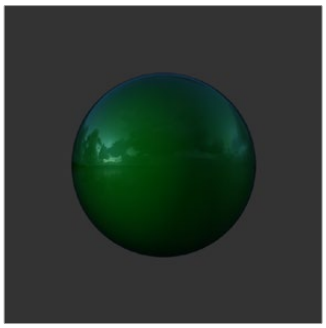
Sphere map example

To produce reflections with real-time interactive feedback at a quality level appropriate for production environment maps, you make some trade-offs on functionality when compared with slower but physically accurate raytraced rendering. Environment-mapped reflections and refractions do not provide self-reflection or any other kind of interaction between different objects. In particular, this infinite distance assumption means that objects cannot interact with themselves (e.g., the reflections on the handle of a teapot will not show the body of the teapot). It also means that objects using the same cube map will not inter-reflect with each other. For example, two neighboring objects would not reflect each other. A separate cube map must be rendered for each object.

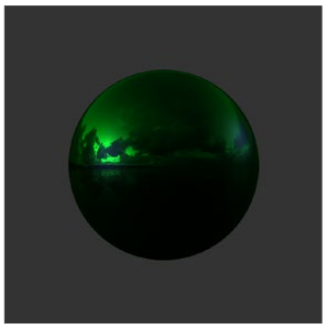
The Reflect node outputs a material that can be applied to an object directly, but the material does not contain an illumination model. As a result, objects textured directly by the Reflect node will not respond to lights in the scene. For this reason, the Reflect node is usually combined with the Blinn, Cook-Torrance, Phong, or Ward nodes.

Reflection

Reflection outputs a material making it possible to apply the reflection or refraction to other materials either before or after the lighting model with different effects.



A Blinn material connected to a background material input of the Reflect This causes the reflection to be added to the Blinn output



A Reflect is connected to the Diffuse Color component of the Blinn, causing the reflection to be multiplied by the diffuse color and modulated by the lighting

Refraction

Refraction occurs only where there is transparency in the background material, which is generally controlled through the Opacity slider and/or the alpha channel of any material or texture used for the Background Material Texture input. The Reflect node provides the following material inputs:

- **Background Material:** Defines both the opacity for refraction and the base color for reflection.
- **Reflection Color Material:** The environment reflection.
- **Reflection Intensity Material:** A multiplier for the reflection.
- **Refraction Tint Material:** The environment refraction.
- **Bump Map Texture:** Normal perturbing map for environment reflection/refraction vectors.

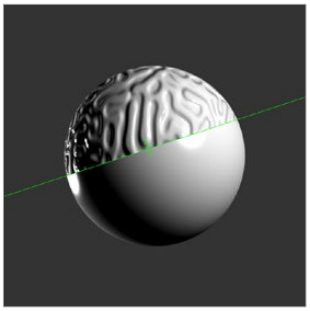
Working with reflection and refraction can be tricky. Here are some techniques to make it easier:

- Typically, use a small amount of reflection, between 0.1 and 0.3 strength. Higher values are used for surfaces like chrome.
- Bump maps can add detail to the reflections/refractions. Use the same bump map in the Illumination model shader that you combine with Reflect.

- When detailed reflections are not required, use a relatively small cube map, such as 128 x 128 pixels, and blur out the image.
- The alpha of refracted pixels is set to 1 even though the pixels are technically transparent. Refracted pixels increase their alpha by the reflection intensity.
- If the refraction is not visible even when a texture is connected to the Refraction Tint Material input, double-check the alpha/opacity values of the background material.

Bump Maps

Bump mapping helps add details and small irregularities to the surface appearance of an object. Bump mapping modifies the geometry of the object or changes its silhouette.



Split screen of a sphere—half with bump map, half without

To apply a bump map, you typically connect an image containing the bump information to the BumpMap node. The bump map is then connected to the Bump input of a Material node. There are two ways to create a bump map for a 3D material: a height map and a bump map.



Image connected to a BumpMap connected to a CookTorrance material node

Using a Height Map

A height map is an image where the value of a pixel represents the height. It is possible to select which color channel is used for bump creation. White means high and black means low; however, it is not the value of a pixel in the height map that determines the bumpiness, but rather how the value changes in the neighborhood of a pixel.

Using a Bump Map

A bump map is an image containing normals stored in the RGB channels.

TIP: Normals are generated by 3D modeling and animation software as a way to trick the eye into seeing smooth surfaces, even though the geometry used to create the models uses only triangles to build the objects.

Normals are 3 float values (nx, ny, nz) whose components are in the range [–1, +1]. Because you can store only positive values in Fusion's integer images, the normals are packed from the range [–1, +1]

to the range [0, 1] by multiplying by 0.5 and adding 0.5. You can use Brightness Contrast or a Custom node to do the unpacking.

If you were to connect a bump map directly to the bump map input of a material, it will result in incorrect lighting. Fusion prevents you from doing this, however, because Fusion uses a different coordinate system for doing the lighting calculation. You first must use a BumpMap that expects a packed bump map or height map and will do the conversion of the bump map to work correctly.

If your bump mapping doesn't appear correct, here are a few things to look for:

- Make sure you have the nodes connected correctly. The height/bump map should connect into a BumpMap and then, in turn, should connect into the bump map input on a material.
- Change the precision of the height map to get less banding in the normals. For low frequency images, float32 may be needed.
- Adjust the Height scale on the BumpMap. This scales the overall effect of the bump map.
- Make sure you set the type to HeightMap or BumpMap to match the image input. Fusion cannot detect which type of image you have.
- Check to ensure High Quality is on (right-click in the transport controls bar and choose High Quality from the contextual menu). Some nodes like Text+ produce an anti-aliased version in High Quality mode that will substantially improve bump map quality.
- If you are using an imported normal map image, make sure it is packed [0–1] in RGB and that it is in tangent space. The packing can be done in Fusion, but the conversion to tangent space cannot.

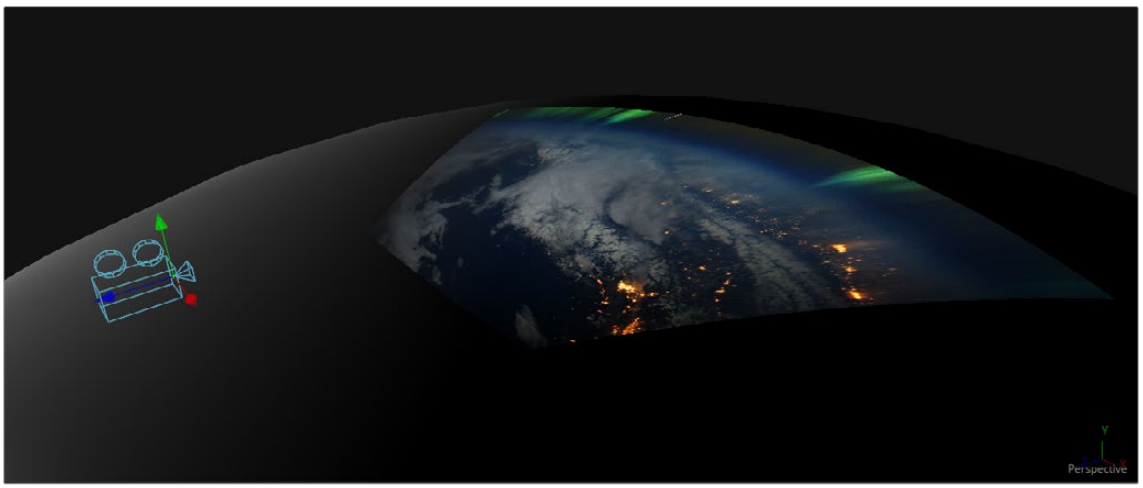
Projection Mapping

Projection is a technique for texturing objects using a camera or projector node. This can be useful for texturing objects with multiple layers, applying a texture across multiple separate objects, projecting background shots from the camera's viewpoint, image-based rendering techniques, and much more.

There are three ways to do projection mapping in Fusion.

Using the Projector/Camera Tool to Project Light

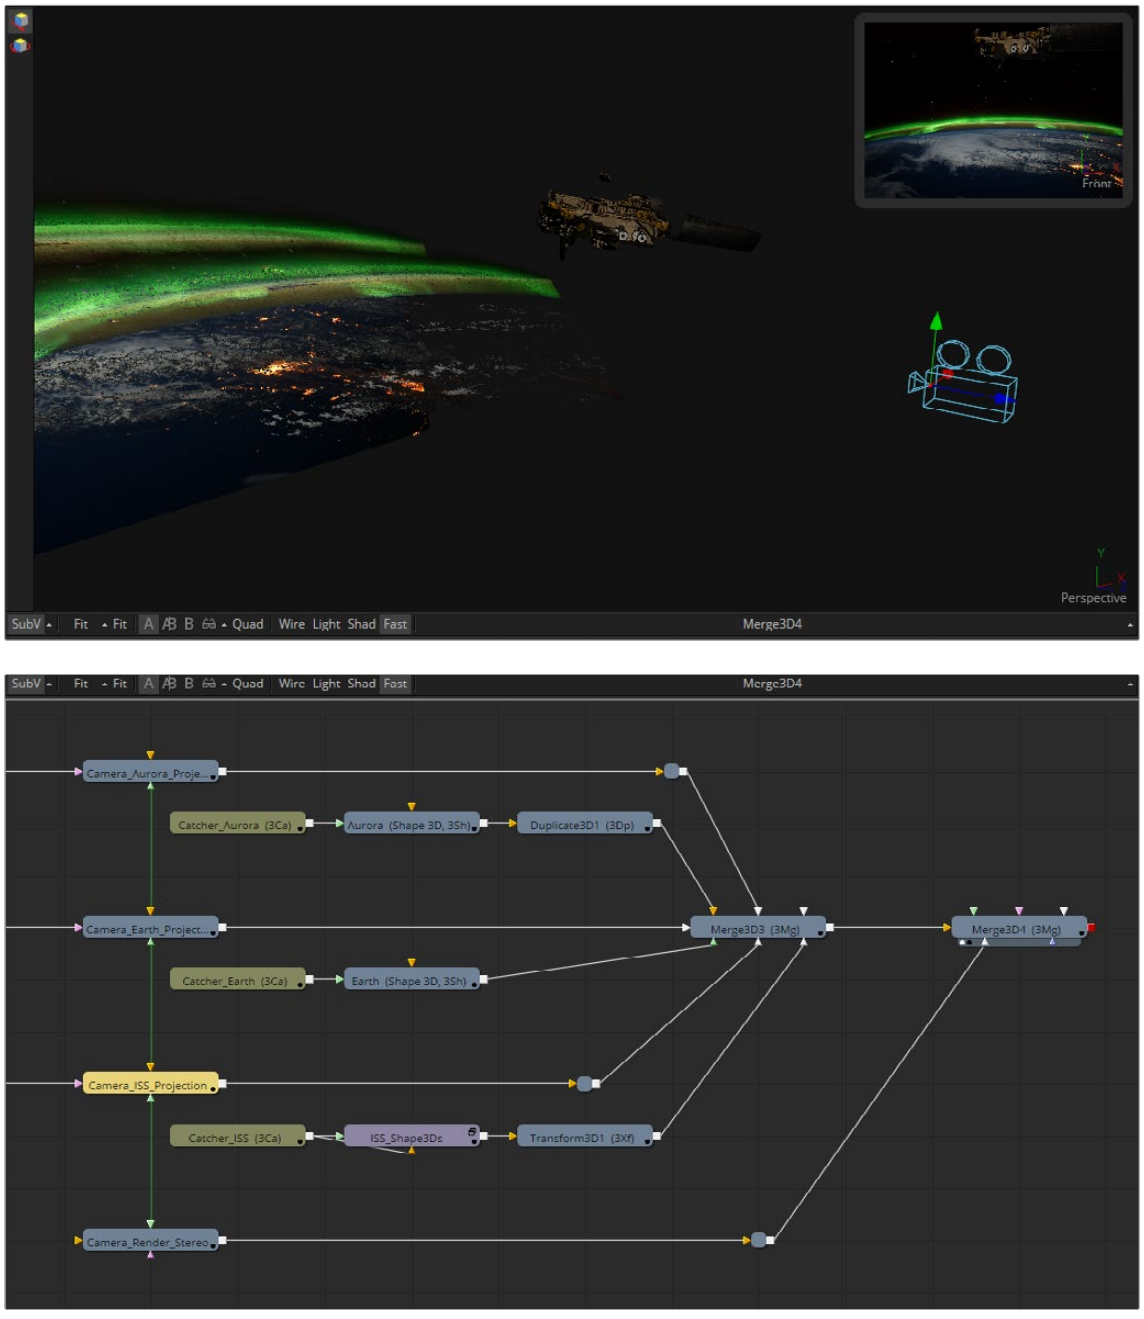
When lighting is enabled, a Camera 3D or Projector 3D can act as a light with all the lighting features. When Camera Projection is enabled or you use a projector, you can choose whether the projection behaves like a spotlight or an ambient light; however, alpha channels cannot be projected. Overlapping projections add together like any other light node. An internal clipping plane (at around 0.01 distance from camera) limits how close the projector or camera can get to the receivers of the projection.



Camera node used for a projection map

Project a Texture onto a Catcher Material

If you do not want to work with light sources, you can use the projector or camera as a texture projector. To work without lighting, a catcher is required in order to receive the texture and apply it to a material. Only objects using this material will receive the projection. This offers some advantages, like the projection of alpha channels, and texturing other channels like specular color or roughness. If the software renderer is used, overlapping projections can be combined in various ways (mean, median, blend, and so on) via the Catcher node. When using the OpenGL renderer, one catcher per projector is used, and the results can be combined using another material. Similar to the Light Projection technique, an internal clipping plane (at around 0.01 distance from camera) limits how close the projector/camera can get to the projection receivers.



Camera projection used with a Catcher node (example from an older version of Fusion)

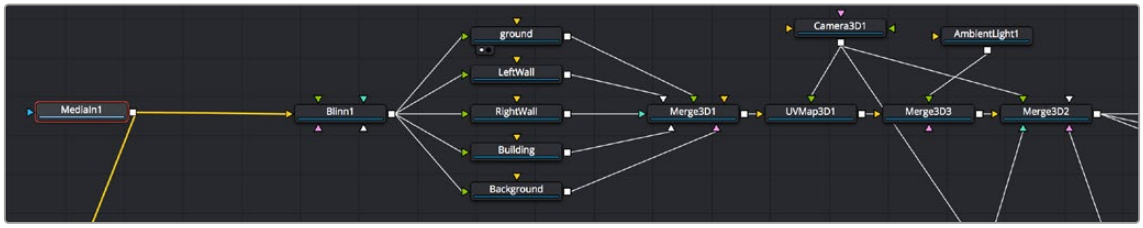
Project Using the UVMap Node

This mode requires a camera and a UVMap3D node downstream of the objects to which the texture is being projected. In the Inspector, when the UVMap Map mode is set to Camera, it gathers the information from the camera and creates new UVs for the input objects, which are used for texturing. Because the UVs are stored in the vertices of the mesh, the object must be tessellated sufficiently.

Textures are assigned to the object like any other texturing technique. The UVs can be locked to the vertices at a chosen frame using the Ref Time slider. This locking only works as long as vertices are not created, destroyed, or reordered (e.g., projection locking will not work on particles because they get created/destroyed, nor will they work on a Cube3D when its subdivision level slider is animated).

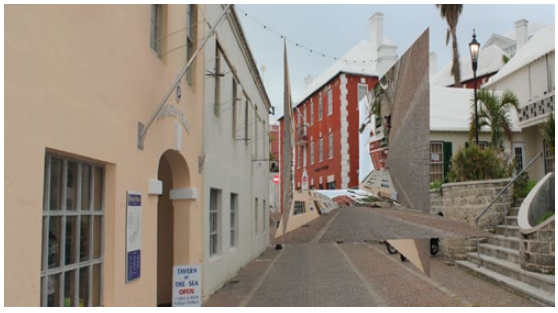
> **TIP:** Projected textures can be allowed to slide across an object. If the object moves relative to the Projector 3D, or alternatively, by grouping the two together with a Merge3D, they can be moved as one and the texture will remain locked to the object.

In the following section of a much larger composition, an image (the Loader1 node) is projected into 3D space by mapping it onto five planes (Shape3D nodes renamed ground, LeftWall, RightWall, Building, and Background), which are positioned as necessary within a Merge3D node to apply reflections onto a 3D car to be composited into that scene.



Excerpt of a composition that's projecting an image of a street scene into 3D space

The output of the Merge3D node used to assemble those planes into a scene is then fed to a UV Map node, which in conjunction with a Camera3D node correctly projects all of these planes into 3D space so they appear as they would through that camera in the scene. Prior to this UVMap projection, you can see the planes arranged in space at left, where each plane has the scene texture mapped to it. At right is the image after the UVMap projection, where you can see that the scene once again looks "normal," with the exception of a car-shaped hole introduced to the scene.





Five planes positioning a street scene in 3D space in preparation for UV Projection (left), and the UV Map node being used to project these planes so they appear as through a camera in the scene (right)

However, this is now a 3D scene, ready for a digital car to be placed within it, receiving reflections and lighting and casting shadows into the scene as if it were there.



The new 3D scene casting reflections and lighting onto a 3D car, and receiving shadows caused by the car

Geometry

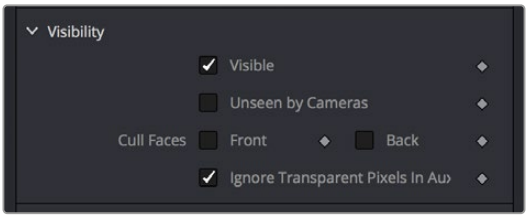
There are five nodes used for creating geometry in Fusion. These nodes can be used for a variety of purposes. For instance, the Image Plane 3D is primarily used to place image clips into a 3D scene, while the Shapes node can add additional building elements to a 3D set, and Text 3D can add three-dimensional motion graphics for title sequences and commercials. Although each node is covered in more detail in the "3D Nodes" chapter, a summary of the 3D creation nodes is provided below.

- **Cube 3D**
  The Cube 3D creates a cube with six inputs that allow mapping of different textures to each of the cube's faces.
- **Image Plane 3D**
  The Image Plane 3D is the basic node used to place a 2D image into a 3D scene with an automatically scaled plane.
- **Shape 3D**
  This node includes several basic primitive shapes for assembling a 3D scene. It can create planes, cubes, spheres, cylinders, cones, and toruses.
- **Text 3D**
  The Text 3D is a 3D version of the Text+ node. This version supports beveling and extrusion but does not have support for the multi-layered shading model available from Text+.
- **Particles**
  When a pRender node is connected to a 3D view, it will export its particles into the 3D environment. The particles are then rendered using the Renderer3D instead of the Particle renderer.

For more information, *see Chapter 54, "Particle Nodes,"* in the Fusion Reference Manual.

Common Visibility Parameters

Visibility parameters are found in the Controls tab of most 3D geometry-producing nodes, exposed via a disclosure control. These parameters let you control object visibility in the viewers and in the final render.



A 3D geometry node's visibility parameters

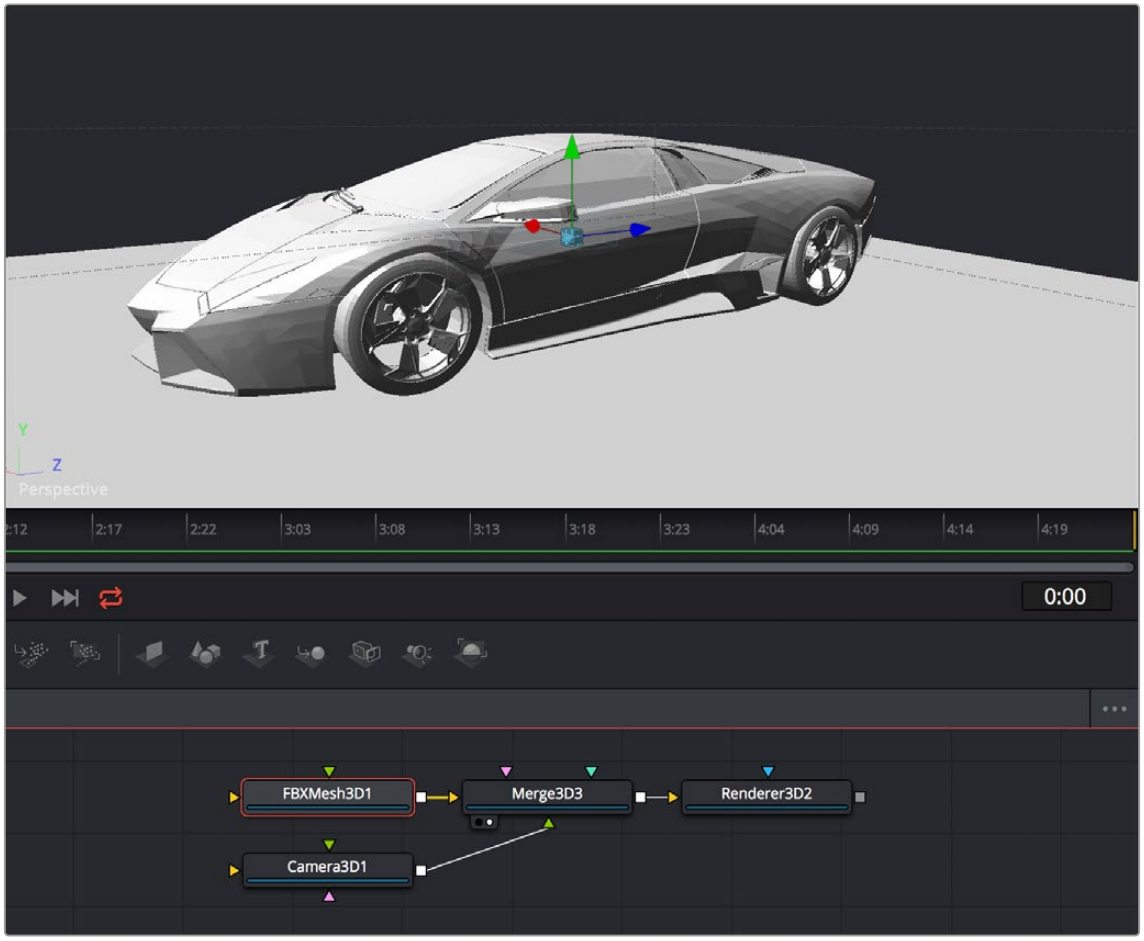
- **Visible**
  If the Visibility checkbox is not selected, the object will not be visible in a viewer, nor will it be rendered into the output image by a Renderer3D. A non-visible object does not cast shadows. This is usually enabled by default, so objects that you create are visible in both the viewers and final renders.
- **Unseen by Cameras**
  If the Unseen by Cameras checkbox is selected, the object will be visible in the viewers but invisible when viewing the scene through a camera, so the object will not be rendered into the output image by a Renderer3D. Shadows cast by an Unseen object will still be visible.
- **Cull Front Face/Back Face**
  Use these options to cull (exclude) rendering of certain polygons in the geometry. If Cull Back Face is selected, all polygons with normals pointing away from the view will not be rendered and will not cast shadows. If Cull Front Face is selected, all polygons with normals pointing away from the view will likewise be excluded. Selecting both checkboxes has the same effect as deselecting the Visible checkbox.
- **Ignore Transparent Pixels in Aux Channels**
  For any piece of geometry, the Renderer3D rejects transparent pixels in the auxiliary image channels. The reason this is the default is to prevent aux channels (e.g., normals, Z-channel, UVs) from filling in where there should be blank space or full transparency. For example, suppose in post you want to add some fog to the rendered image. If you had fully transparent geometry in the foreground affecting the Z-channel, you would get incorrect fog rendering. By deselecting this checkbox, the transparency will not be considered and all the aux channels will be filled for all the pixels. This could be useful if you wanted to replace texture on a 3D element that is fully transparent in certain areas with a texture that is transparent in different areas; it would be useful to have the whole object set aux channels (in particular UVs).

Adding FBX Models

The Filmbox FBX format is a scene interchange format that facilitates moving 3D scene information from one application to another. Fusion's FBX format support extends model import support to other 3D files such as Collada and OBJ.

Importing an FBX Scene

To import an entire FBX scene, you add an FBXMesh3D node to your node tree. After being prompted to choose a scene or object file, Fusion imports it to create a composition with the same lights, cameras, materials, and geometry found in an FBX file.



An imported model, via the FBXMesh3D node

FBX Scene Import Dialog

The FBX Mesh node is used to import mesh geometry from an FBX file. The first texture applied to a mesh will also be imported, if available.

Since different 3D applications use different units to measure their 3D scenes, the imported geometry may be enormous compared to the rest of the scene, because Fusion treats its scale of measurement as equal to its own system. For example, if your 3D application defaults to using millimeters as its scale, an object that was 100 millimeters in size will import as a massive 100 units.

You can use the Size slider in the FBX Mesh Inspector parameters to reduce the scale of such files to something that matches Fusion's 3D scene.

FBX Exporter

You can export a 3D scene from Fusion to other 3D packages using the FBX Exporter node. On render, it saves geometry, cameras lights, and animation into different file formats such as .dae or .fbx. The animation data can be included in one file, or it can be baked into sequential frames. Textures and materials cannot be exported.

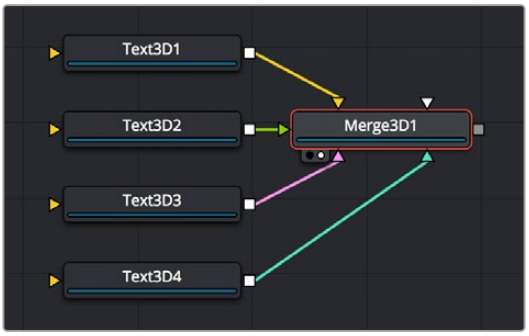
Using Text3D

The Text3D node is probably the most ubiquitous node employed by motion graphics artists looking to create titles and graphics from Fusion. It's a powerful node filled with enough controls to create nearly any text effect you might need, all in three dimensions. This section seeks to get you started quickly with what the Text3D node is capable of.

For more information, *see Chapter 29, "3D Nodes,"* in the Fusion Reference Manual.

Assembling Text Objects

Each Text3D node is a self-contained scene within which each character of text is an individual object. Because of this, the ideal way to combine numerous text objects that you might want to animate or style independently from one another is to connect as many Text3D objects as you want to be able to independently animate or style to one or more Merge3D nodes.

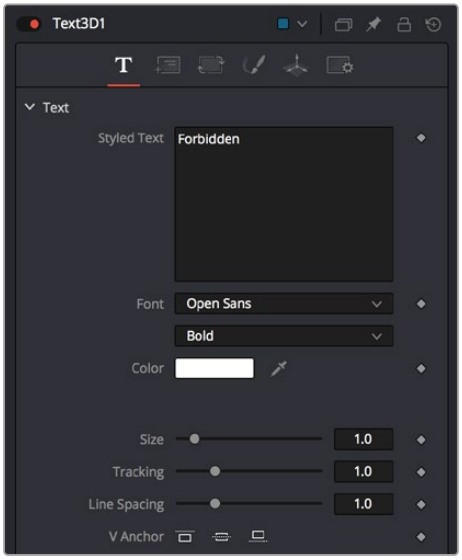


Merging multiple text objects to create an intricately styled scene

TIP: If you click the Text icon in the toolbar to create a Text3D node, and then you click it again while the Text3D node you just created is selected, a Merge3D node is automatically created and selected to connect the two. If you keep clicking the Text icon, more Text3D nodes will be added to the same selected Merge3D node.

Entering Text

When you select a Text3D node and open the Inspector, the Text tab shows a "Styled Text" text entry field at the very top into which you can type the text you want to appear onscreen. Below, a set of overall styling parameters are available to set the Font, Color, Size, Tracking, and so on. All styling you do in this tab affects the entire set of text at once, which is why you need multiple text objects if you want differently styled words in the same scene.



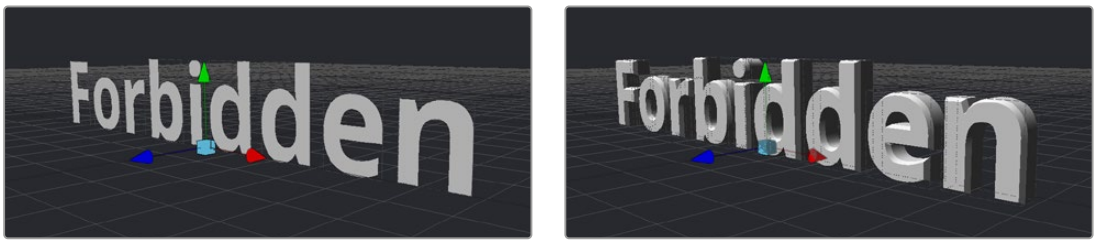
The text entry and styling parameters in the Text tab

Near the bottom of the Text tab are the Extrusion parameters, available within a disclosure control.



The Extrusion parameters near the bottom of the Text tab

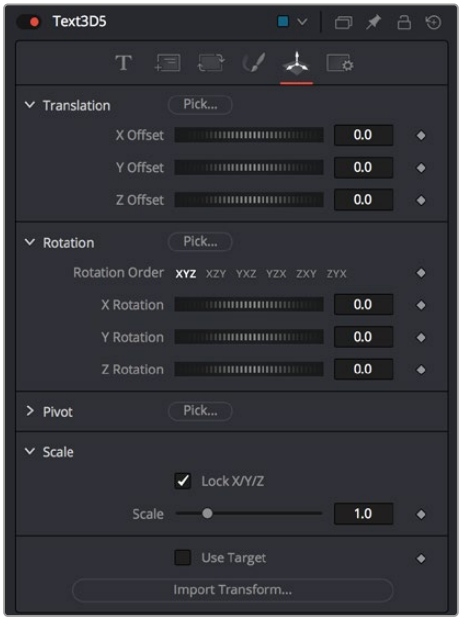
By default, all text created with the Text3D node is flat, but you can use the Extrusion Style, Extrusion Depth, and various Bevel parameters to give your text objects thickness.



Unextruded text (left), and Extruded text (right)

Positioning and Transforming Text

By default, every new Text3D node is positioned at 0, 0, 0, so when you add multiple Text3D nodes, they're all in the same place. Fortunately, every Text3D node has built-in transform controls in the Transform tab.



Text3D nodes also have Transform parameters built-in

Additionally, selecting a Text3D node exposes all the onscreen transform controls discussed elsewhere in this chapter. Using these controls, you can position and animate each text object independently.



Repositioned text objects to create a title sequence

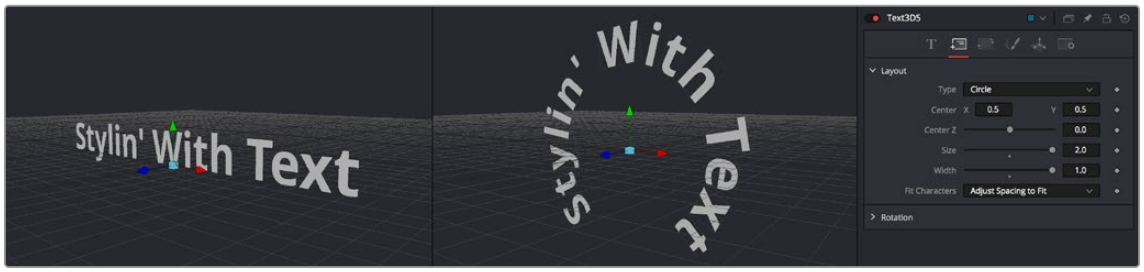
Combining Text3D nodes using Merge3D nodes doesn't just create a scene; it also enables you to transform your text objects either singly or in groups:

Selecting an individual Text3D node or piece of text in the viewer lets you move that one text object around by itself, independently of other objects in the scene.

Selecting a Merge3D node exposes a transform control that affects all objects connected to that Merge3D node at once, letting you transform the entire scene.

Layout Parameters

The Layout tab presents parameters you can use to choose how text is drawn: on a straight line, a frame, a circle, or a custom spline path, along with contextual parameters that change depending on which layout you've selected (all of which can be animated).



Text using two different layouts

"Sub" Transforms

Another Transform tab (which the documentation has dubbed the "Sub" Transform tab) lets you apply a separate level of transform to either characters, words, or lines of text, which lets you create even more layout variations. For example, choosing to Transform by Words lets you change the spacing between words, rotate each word, and so on. You can apply simultaneous transforms to characters, words, and lines, so you can use all these capabilities at once if you really need to go for it. And, of course, all these parameters are animatable.



Transforming individual words in two different ways

Shading

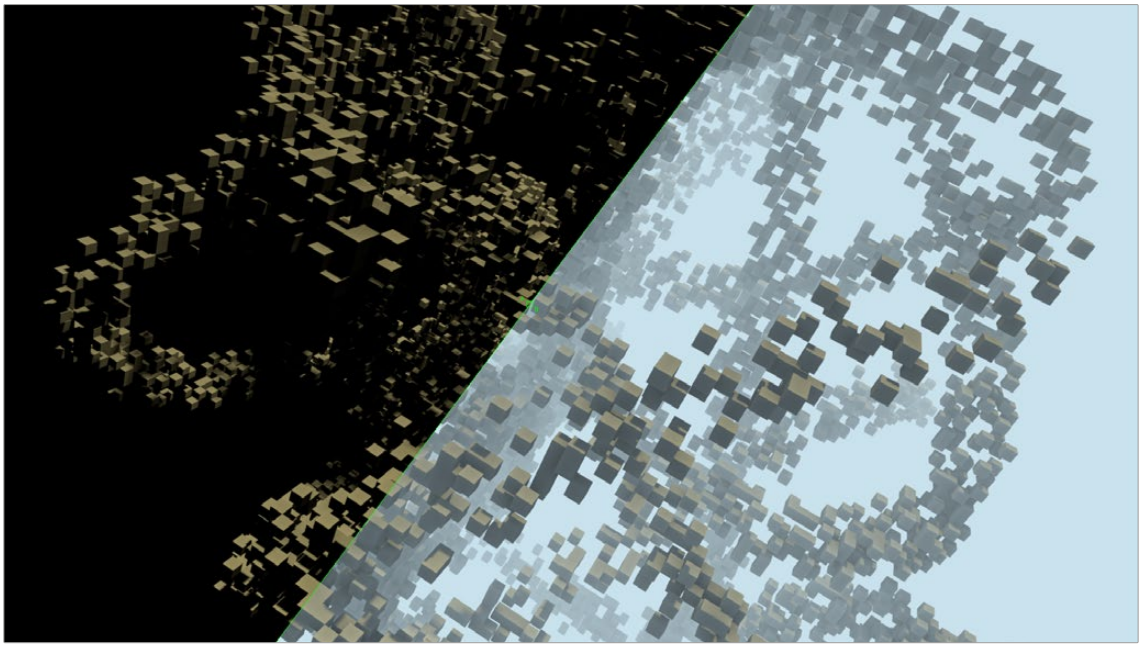
The Shading tab lets you shade or texture a text object using standard Material controls.



Shading controls for text objects

Fog 3D and Soft Clipping

The Fog3D node helps to create atmospheric depth cues.



Split screen with and without fog

The Fog3D node works well with depth of field and antialiasing supported by the OpenGL renderer. Since it is not a post-processing node (like the VolumeFog node found in the Nodes > Position menu or Fog node in Nodes > Deep Pixel), it does not need additional channels like Position or Z-channel color. Furthermore, it supports transparent objects.

The SoftClip node uses the distance of a pixel from the viewpoint to affect opacity, allowing objects to gradually fade away when too close to the camera. This prevents objects from “popping off” should the camera pass through them. This is especially useful with particles that the camera may be passing through.

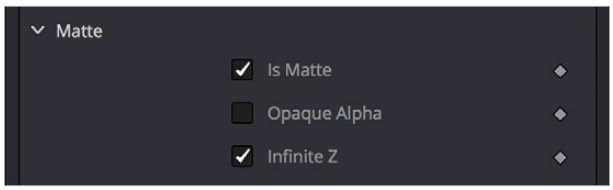
Geometry nodes such as the Shape3D node use a Matte Objects checkbox to enable masking out parts of the 3D scene. Effectively, everything that falls behind a matte object doesn’t get rendered. However, matte objects can contribute information into the Z-channel and the Object ID channel, leaving all other channels at their default values. They do not remove or change any geometry; they can be thought of as a 3D garbage matte for the renderer.



Circle shape used as a Matte object to see the floor

Matte Object Parameters

Opening the Matte disclosure control reveals the Is Matte option, which when turned on enables two more options.



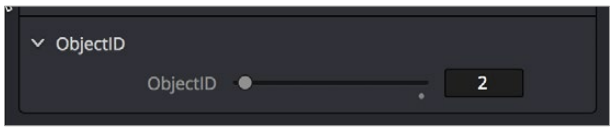
Matte parameters in the Shape3D node; enabling Is Matte reveals additional options

- **Is Matte**
  Located in the Controls tab for the geometry, this is the main checkbox for matte objects. When enabled, objects whose pixels fall behind the matte object's pixels in Z do not get rendered.
- **Opaque Alpha**
  When the Is Matte checkbox is enabled, the Opaque Alpha checkbox is displayed. Enabling this checkbox sets the alpha value of the matte object to 1. Otherwise the alpha, like the RGB, will be 0.
- **Infinite Z**
  When the Is Matte checkbox is enabled, the Infinite Z checkbox is displayed. Enabling this checkbox sets the value in the Z-channel to infinite. Otherwise, the mesh will contribute normally to the Z-channel.

  Matte objects cannot be selected in the viewer unless you right-click in the viewer and choose 3D Options > Show Matte Objects in the contextual menu. However, it's always possible to select the matte object by selecting its node in the node tree.

Material and Object IDs

Most nodes in Fusion that support effect masking can use Object ID and Material ID auxiliary channels to generate a mask. The parameters used to accomplish this are found in the Common Controls tab of each node.



Material ID parameters in a Shape3D node's Inspector controls

The Material ID is a value assigned to identify what material is used on an object. The Object ID is roughly comparable to the Material ID, except it identifies objects and not materials.

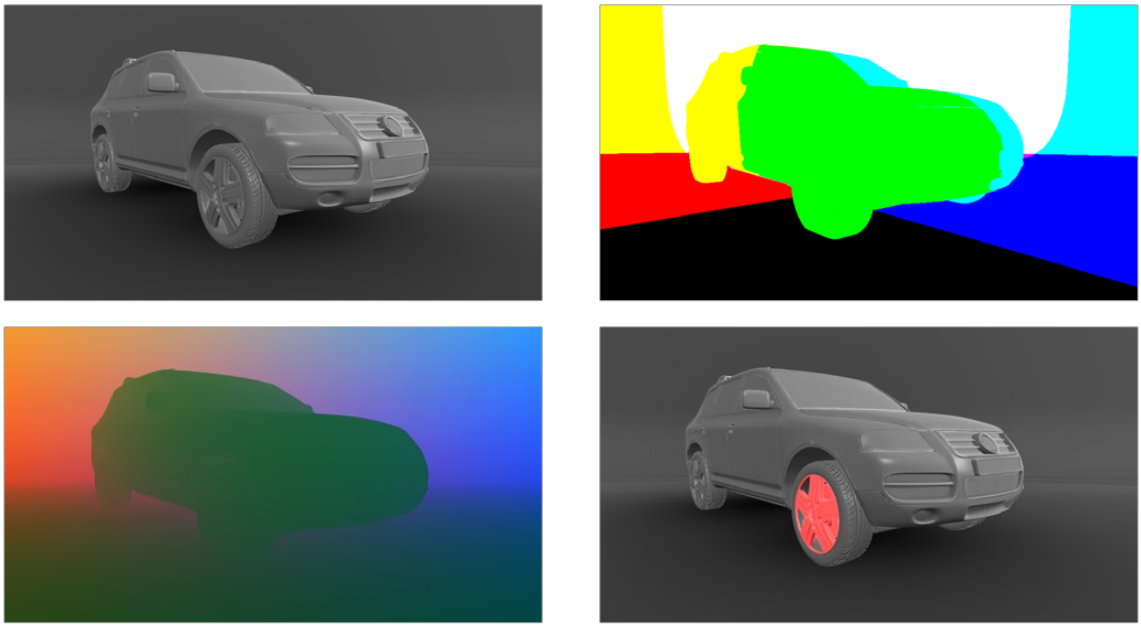
Both the Object ID and Material ID are assigned automatically in numerical order, beginning with 1. It is possible to set the IDs to the same value for multiple objects or materials even if they are different. Override 3D offers an easy way to change the IDs for several objects. The Renderer will write the assigned values into the frame buffers during rendering, when the output channel options for these buffers are enabled. It is possible to use a value range from 0 to 65534. Empty pixels have an ID of 0, so although it is possible to assign a value of 0 manually to an object or material, it is not advisable because a value of 0 tells Fusion to set an unused ID when it renders.



Object ID for ground plane and object set to the same numeric value

World Position Pass

The World Position Pass, or WPP, is a render pass generated from 3D applications. Each pixel is assigned the XYZ position where the pixel was generated in the world coordinates. So if the face from which the pixel was derived in the scene sits at (0,0,0), the resulting pixel will have a Position value of (0,0,0). If we visualize this as RGB, the pixel will be black. If a face sits at (1,0,0) in the original scene, the resulting RGB pixel will be red. Due to the huge range of possible positions in a typical 3D scene, and 7⁄8 of those possible positions containing negative coordinates, the Position channel is always rendered in 32-bit float.



A World Position Pass rendering of a scene with its center at (0,0,0) The actual image is on the left

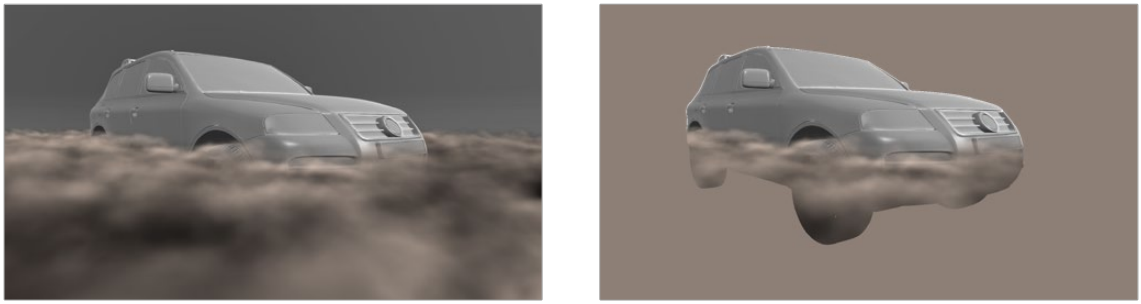
3D Scene Input

Nodes that utilize the World Position channel are located under the Position category. VolumeFog and Z to WorldPos require a camera input matching the camera that rendered the Position channels, which can either be a Camera3D or a 3D scene containing a camera. Just as in the Renderer3D, you can choose which camera to use if more than one are in the scene. The VolumeFog can render without a camera input from the Node Editor if the world space Camera Position inputs are set to the correct value. VolumeMask does not use a camera input. Nodes that support the World Position Pass, located under the Position category, offer a Scene input, which can be either a 3D Camera or a 3D scene containing a camera.

There are three Position nodes that can take advantage of World Position Pass data.

- Nodes > Position > Volume Fog
- Nodes > Position > Volume Mask
- Nodes > Position > Z to World
- The “Dark Box”

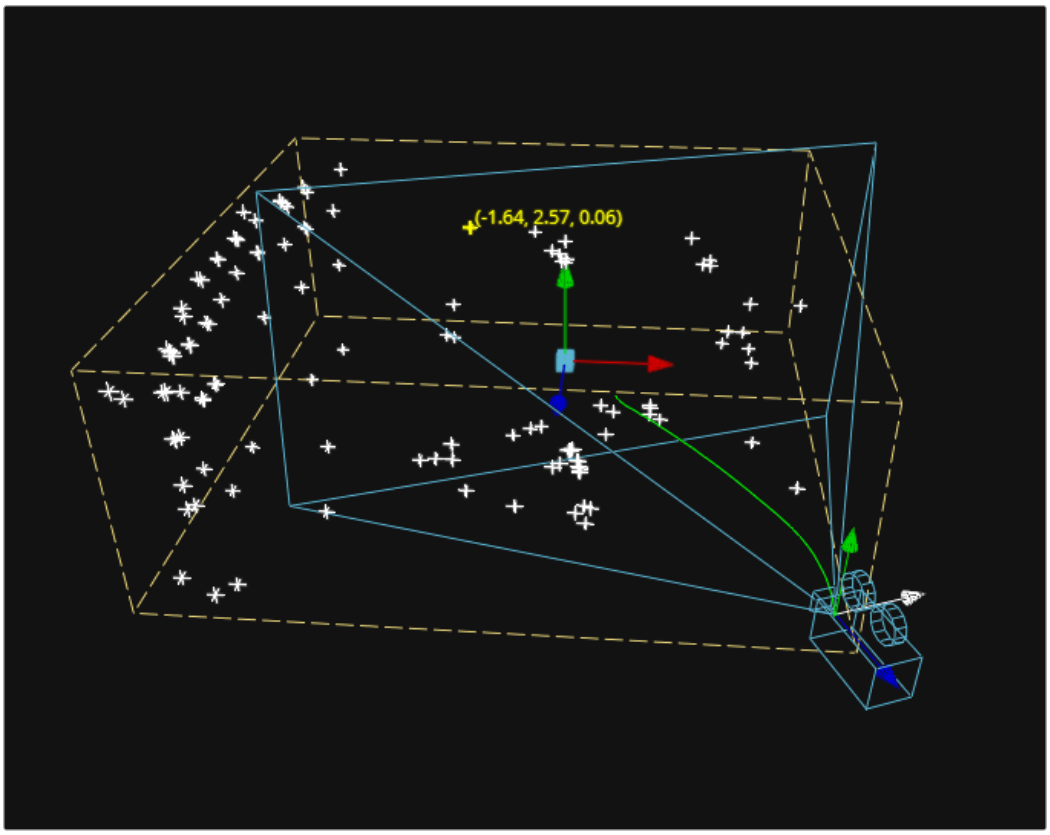
Empty regions of the render will have the Position channel incorrectly initialized to (0,0,0). To get the correct Position data, add a bounding sphere or box to your scene to create distant values and allow the Position nodes to render correctly.



Without a bounding mesh to generate Position values, the fog fills in the background incorrectly

Point Clouds

The Point Cloud node is designed to work with locator clouds generated from 3D tracking software. 3D camera tracking software, such as SynthEyes and PF Track, will often generate hundreds or even thousands of tracking points. Seeing these points in the scene and referencing their position in 3D and screen space is important to assist with lining up live action and CG, but bringing each point in as an individual Locator3D would impact performance dramatically and clutter the node tree.



Point cloud in the viewer

The Point Cloud node can import point clouds written into scene files from match moving or 3D scanning software.

To import a point cloud, do the following:

1. Add the PointCloud3D node to your composition.
2. Click the Import Point Cloud button in the Control panel.
3. Browse to the scene file and select a cloud to import from the scene.

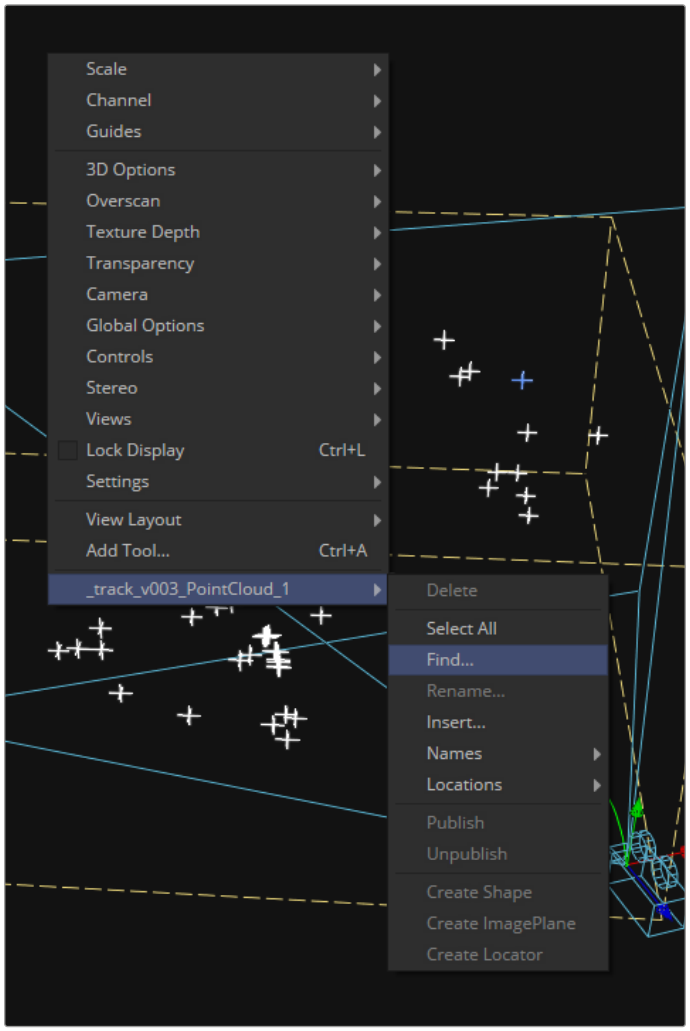
The entire point cloud is imported as one object, which is a significantly faster approach.

Finding, Naming, and Publishing Points

Many 3D trackers allow for the naming of individual tracking points, as well as setting tracking points on points of interest. The Point Cloud 3D will quickly find these points and publish them. A published point in the cloud can be used to drive the animation of other parameters.

To find a point in the point cloud, do the following:

1. Right-click anywhere within a viewer.
2. Choose Find from the Point Cloud's submenu in the contextual menu.
3. Type the name of the point and click OK.



Finding a point cloud using the viewer contextual menu

If a point that matches the name you entered is found, it will be selected in the point cloud and highlighted yellow.

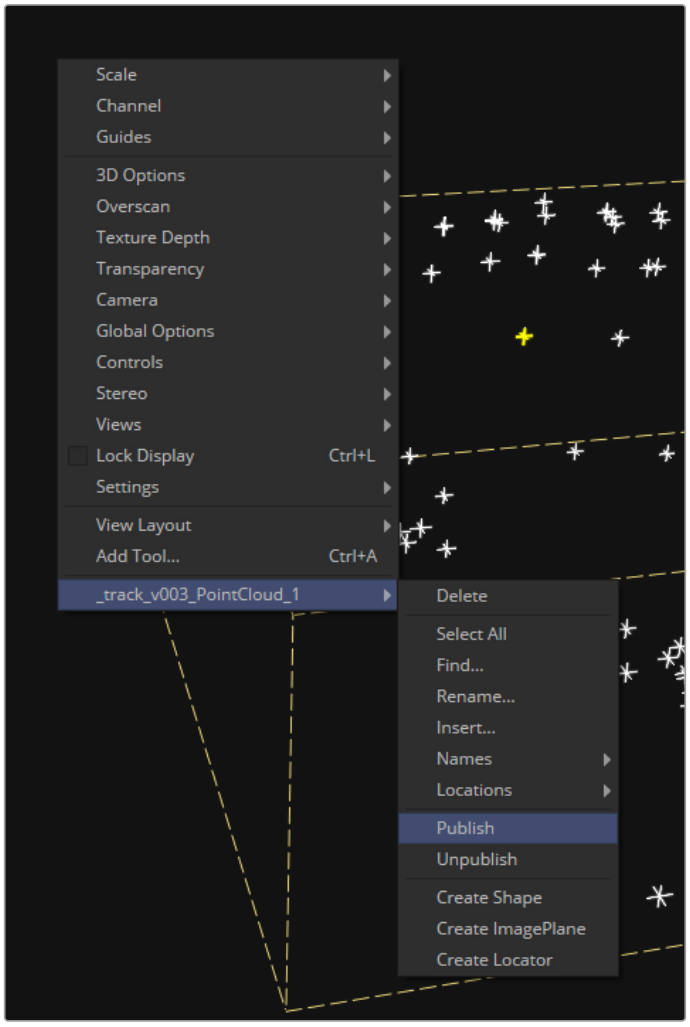
TIP: The Point Cloud Find function is a case-sensitive search. A point named "tracker15" will not be found if the search is for "Tracker15".

Renaming a Point in the Cloud

You can use the Point Cloud contextual menu to rename a selected point. This works only for a single point. A group of points cannot be renamed.

Publishing a Point

If you want to use a point's XYZ positions for connections to other controls in the scene, you can publish the point. This is useful for connecting objects to the motion of an individual tracker. To publish a point, right-click it and choose Publish from the contextual menu.



Publishing a point using the viewer contextual menu

Chapter 26

3D Camera Tracking

This chapter presents an overview of using the Camera Tracker node and the workflow it involves. Camera tracking is used to create a virtual camera in Fusion's 3D environment based on the movement or a live-action camera in a clip. You can then use the virtual camera to composite 3D models, text, or 2D images into a live-action clip that has a moving camera.

For more information on other types of tracking in Fusion, *see Chapter 22, "Using the Tracker Node,"* in the Fusion Reference Manual.

Contents

Introduction to Tracking 661
What Is 3D Camera Tracking? 661
How Camera Tracking Works 662
The Camera Tracking Workflow 662
Clips That Don't Work Well for Camera Tracking 663
Outputting from the Camera Tracker 664
2D View 664
3D View 665
Auto-Tracking in the Camera Tracker 666
Increasing Auto-Generated Tracking Points 666
Masking Out Objects 667
Matching the Live-Action Camera 668
Running the Solver 669
How Do You Know When to Stop? 670
Using Seed Frames 670
Cleaning Up Camera Solves 671
Exporting a 3D Scene for Efficiency 673
Unalign the 3D Scene Transforms 674
Setting the Ground Plane 674
Setting the Origin 675
Setting the Scale 675
Realign the Scene 675
Viewing the Exported Results 675

Introduction to Tracking

Tracking is one of the most useful and essential techniques available to a compositor. It can roughly be defined as the creation of a motion path from analyzing a specific area in a clip over time. Fusion provides a variety of different tracking nodes that let you analyze different kinds of motion.

Each tracker type has its own chapter in this manual. This chapter covers the tracking techniques with the Camera Tracker node.

What Is 3D Camera Tracking?

Camera tracking is used for match moving, and it's a vital link between 2D scenes and 3D scenes, allowing compositors to integrate 3D CGI elements into live-action clips. The Camera Tracker node calculates the path of a live-action camera and generates a virtual camera in 3D space. This virtual camera is intended to be identical to the actual camera that shot the scene, not only in terms of motion but in matching the lens focal length as well. The calculated position and movement of the virtual camera is central to realistically compositing 3D elements with live action.



An example of 3D elements integrated in a live-action scene

How Camera Tracking Works

Camera tracking begins by tracking the movement of fixed features from one frame to the next. To put it another way, camera tracking algorithms follow features that are “nailed to the set.” Objects in the scene that move independently of the camera movement in the shot, such as cars driving or people walking, cause poor tracks, so masks can be used to restrict the features that are tracked in order to improve the results. Additionally, it is helpful to provide specific camera metadata, such as the sensor size and the focal length of the lens. This information guides the scene reconstruction calculation, called a solver, toward generating a more accurate virtual camera.

The Camera Tracker’s purpose is to create a 3D animated camera and point cloud of the scene. A point cloud is a large group of points generated by the solver that roughly recreates the 3D positions of the tracked features in a scene. The point cloud can then be used as a guide when integrating other 2D or 3D elements alongside live-action features.

The Camera Tracking Workflow

Camera tracking has two main phases:

1. *Tracking*, which is the analysis of a scene.
2. *Solving*, which calculates the virtual 3D scene.

Once you complete these steps, an animated camera and point cloud are exported from the Inspector into a 3D composite. The Camera Tracker encompasses this complete workflow within one tool. Five tabs at the top of the Inspector are roughly laid out in the order in which you’ll use them. These five tabs are:

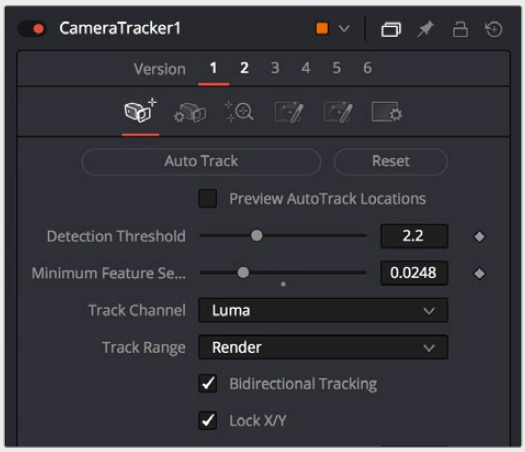
Track: Used to track a clip.

Camera: Configures the basic Camera parameters.

Solve: Calculates the 3D placement of the 2D tracking points and reconstructs the camera.

Export: Generates a Camera 3D node, a Point Cloud node, and a 3D scene in the node tree.

Options: Used to customize the look of the onscreen overlays.



The Camera Tracker tab

Clips That Don't Work Well for Camera Tracking

Even though the Camera Tracker is somewhat automatic, it sometimes needs your help. If you can identify potential issues before you even track or solve the shot, you can save yourself much time. Certain types of clips will cause more significant problems for camera tracking than others. Some are fixable, while for others you just have to admit defeat and figure out another solution. Here is a list of the types of shots to look out for, as they can be big headaches for camera tracking:

- **Lack of depth:** Camera tracking requires parallax in a clip in order to work. You must be able to identify objects further away and objects that are nearer as the camera moves. If everything is at the same distance from the camera, there is no way to calculate depth. In this case, it's better to skip the Camera Tracker node and find another solution.
- **Locked-off shots:** If the camera does not move, there is no way to calculate which objects are closer and which are nearer. Again, don't spend too much time in this situation; it is better to skip the Camera Tracker node and find another solution.
- **Tripod pans:** Similar to a locked-off shot, there is no way to calculate which objects are closer and which are nearer from a pan that remains centered on a locked-off tripod. Skip the Camera Tracker node and find another solution.
- **No detail:** Clips like green screens without tracking markers lack enough detail to track. If you are lucky enough to be involved in the shooting of these types of shots, including tracker markers makes it much easier to get a good track. Without detail, camera tracking will fail and you will need to find a more manual solution.
- **Motion blur:** Fast camera motion or slow shutter speeds can introduce motion blur, which will make it difficult to find patterns to track. It's worth trying shots like these to see if there are enough details to get a good solve, but know when give up and turn to another solution.
- **Rolling shutter:** CMOS-based cameras sometimes introduce distortion due to the shutter capturing different lines at slightly different times. This distortion can create significant problems for camera tracking. Sometimes it is possible to create motion vectors with the Optical Flow node to create new in-between frames without the wobble distortion of the rolling shutter. Then you can use the corrected image to connect to the Camera Tracker.
- **Parallax issues:** When objects at different distances in a shot overlap in the frame, the overlapping area can be misinterpreted as a corner. Having a tracker assigned to an overlapping angle like this will cause errors as the parallax starts to shift and the overlapping area slides. This can be solved in Fusion by removing that tracker before running the solver.
- **Moving objects:** It's difficult to capture a shot where objects in the clip do not move. People, cars, animals, or other object may move in and out of a shot. These objects move independent of the camera movement and must be eliminated or they will cause solving errors. You can fix these issues by masking out objects that are not "nailed to the set." The masks are then connected to the Track Mask input on the Camera Tracker node.

TIP: Some shots that cannot be tracked using Fusion's Camera Tracker can be performed in dedicated 3D camera-tracking software like 3D Equalizer and PF Track. Camera tracking data from these applications can then be imported in the Camera3D node in Fusion.

Outputting from the Camera Tracker

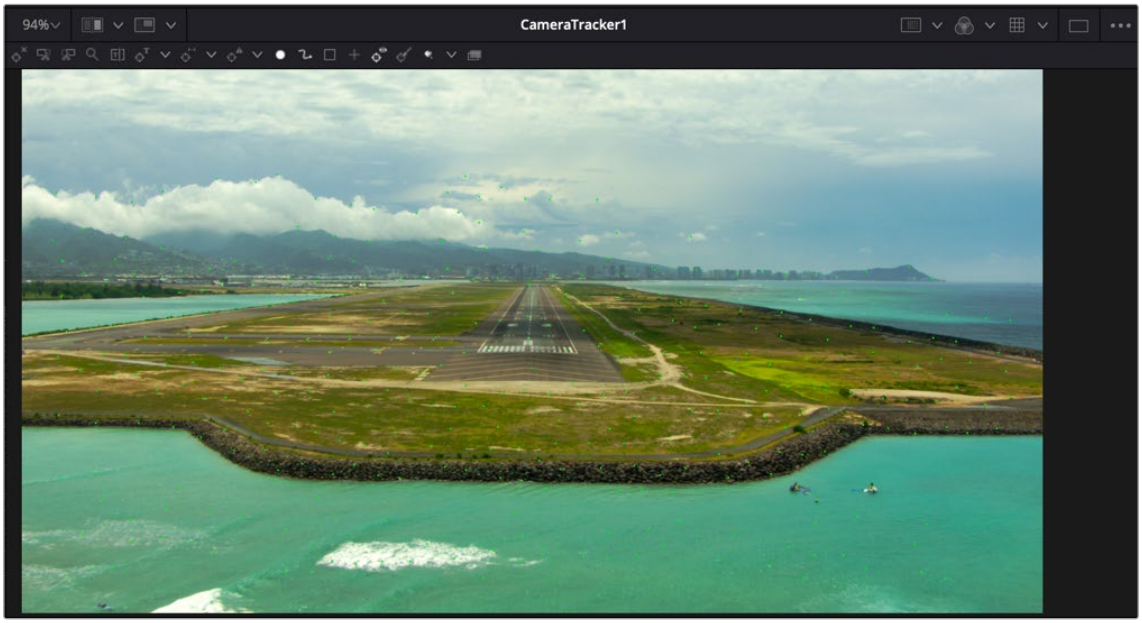
Unlike most Fusion nodes, the Camera Tracker node has two outputs:

- The primary output is a 2D view used when you are setting up the Track, refining the camera, and performing your initial solve.
- There is also a 3D output used after your initial solve for viewing the camera path and point cloud in 3D space. This view can be helpful when you are refining tracks to increase the accuracy of the solve and aligning your ground plane. It can be used simultaneously with the 2D output in side-by-side views.

Note that the selection of tracks in the 2D view and their corresponding locators (in the point cloud) in the 3D view are synchronized. There are also viewer menus available in both the 2D and 3D views to give quick control of the functionality of this tool.

2D View

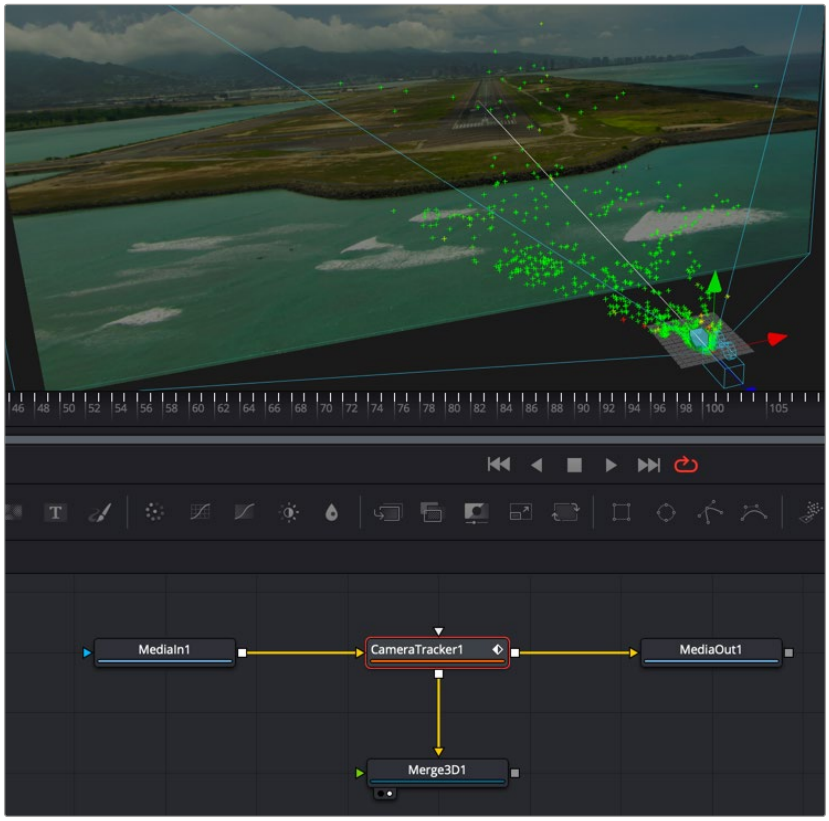
The 2D view is the primary display for the node. Viewing the node displays the image being tracked as well as overlay tracker markers and their motion paths. A dedicated toolbar gives you access to the common features used to track and solve a clip.



The Camera Tracker 2D output with toolbar and auto-track points

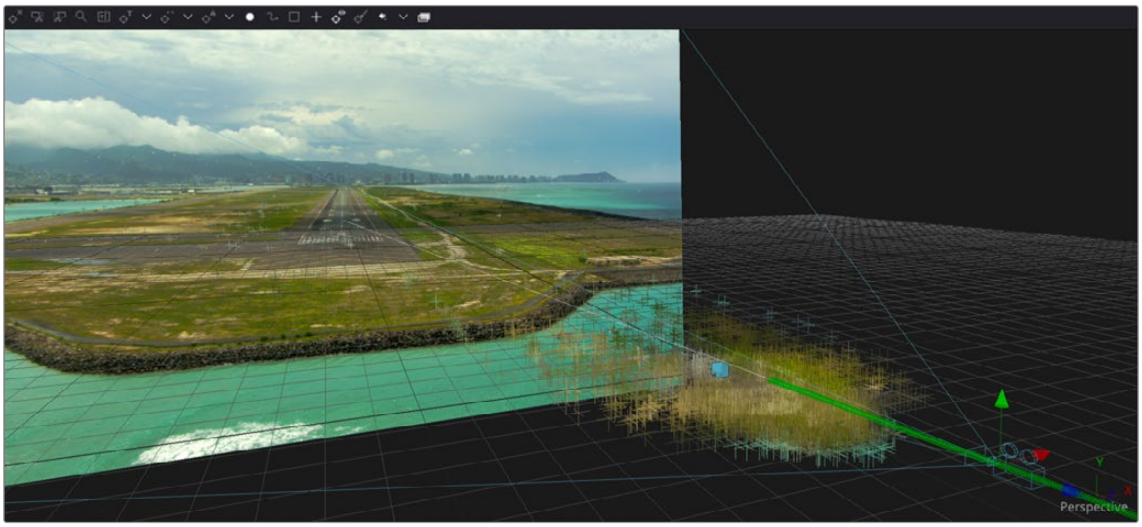
3D View

The second output of the Camera Tracker node displays a 3D scene. To view this, connect this 3D output to a 3D Transform or Merge 3D node and view that tool.



After solving, connecting the second output of the Camera Tracker node to a Merge 3D displays the point cloud in 3D

After an initial solve, the 3D output displays the point cloud and the camera, along with the image connected to it. Selecting points displays the Camera Tracker toolbar above the viewer, which gives control of various functions, such as renaming, deleting, and changing the colors of points in the point cloud.



3D output of a point cloud and a solved camera path

Auto-Tracking in the Camera Tracker

Tracking is the term used to describe the task of observing or analyzing a sequence of frames (or clip). The Camera Tracker node must take into account the movement of the source footage before it can determine the location and movement of the virtual camera. The Camera Tracker tool automatically searches for features that are high-contrast patterns within the clip and assigns trackers to those features. Having a wide distribution of tracking points across the frame and having points with long durations results in the best track.

Increasing Auto-Generated Tracking Points

Unlike the Tracker node, setting tracking points is entirely automatic in the Camera Tracker, but the Detection Threshold and Minimum Feature Separation sliders let you adjust the criteria by which tracking points are found. Lowering these parameters lets you increase the number of tracking points that will be found. This can be useful if the scene has too few points, which will prevent the solver from generating an accurate camera and point cloud. However, make these adjustments with care, since adding too many points may generate redundant trackers that slow down the entire process with minimal benefit.

Previewing Tracking Points

You can view the tracking points currently generated on the clip by viewing the Camera Tracker node and turning on the Preview AutoTrack Locations checkbox. This causes green tracking points to be displayed in the viewer as you play through the clip. Using this preview, you can decide whether you need to make adjustments to the Detection Threshold or Minimum Feature Separation to increase or decrease the number of tracking points that are found automatically.



Green tracker marks are added automatically to the features in an image

Bi-Directional Tracking

When performing a track, you can enable the Bidirectional Tracking checkbox, which first tracks forward from the start of the clip, and then tracks a second pass in reverse. This two-pass approach can potentially extend the duration of any given point by re-analyzing points initially identified in the forward pass. There is very little reason not to have this enabled unless you are very short on time. Bidirectional tracking takes longer, but it's usually worth it, and the process is reasonably quick considering the benefit.

Choosing a Tracking Algorithm

There are three available choices for the algorithm to use when tracking. The three options in the New Track Defaults section of the Inspector include:

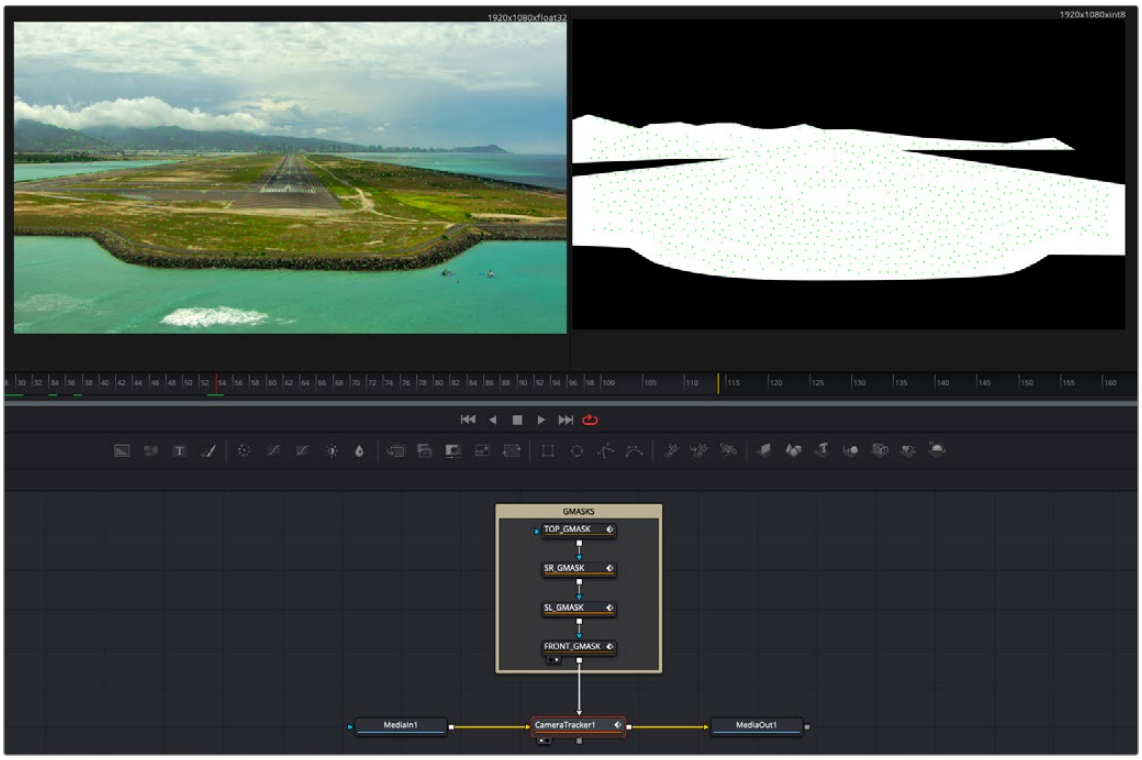
- **Optical Flow:** Usually your best choice, unless you have a great deal of criss-crossing objects in a clip.
- **Tracker:** A good second choice when Optical Flow can't be used due to motion estimation errors like criss-crossing objects.
- **Planar:** Mostly used in simpler clips, where the majority of the image consists of planar surfaces such as the facades of buildings.

Masking Out Objects

When tracking a clip, the Camera Tracker automatically generates trackers on feature details. However, not all features that stand out in a clip are appropriate for camera tracking. You only want to track features that are "nailed to the set." In other words, objects that move independently of the camera motion, like moving cars and people, cause inaccuracies when camera tracking. You must eliminate these types of objects from the analysis.

The primary way of avoiding these problem areas is by masking. You connect a mask to the Camera Tracker node's Track Mask input to identify areas of a scene that the Camera Tracker can analyze. For example, if you have a clip of an airport runway along a shoreline, the waves of the water and moving clouds in the sky must be masked since they move independently of the camera.

When creating a mask, the fixed areas of the image to be analyzed for tracking should be encompassed in the white portion of the mask. All moving objects that need to be ignored should be encompassed in the black portion. The mask should then be attached to the Camera Tracker Track Mask input.



Masks used to omit the moving clouds and waves from being tracked by the Camera Tracker

By doing this, the tracker ignores the waves of the water and moving clouds. Unlike drawing a mask for an effect, the mask in this case does not have to be perfect. You are just trying to identify the rough area to occlude from the tracking analysis.



The original image to be tracked (left), and the occlusion mask of the clouds and water (right)

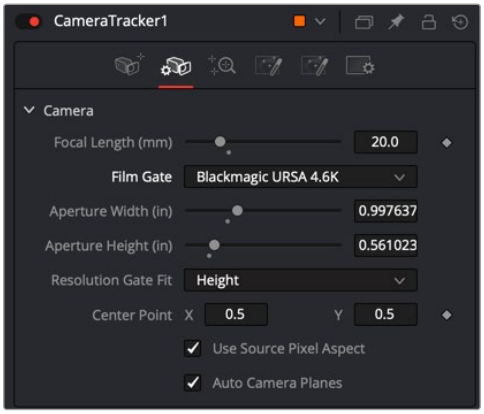
TIP: If there's a lot of motion in a shot, you can use the Tracker or Planar Tracker nodes to make your occlusion mask follow the area you want to track. Just remember that, after using the PlanarTracker or PlanarTransform node to transform your mask, you need to use a Bitmap node to turn it back into a mask that can be connected to the Camera Tracker node's Track Mask input.

Matching the Live-Action Camera

Once you have completed tracking, the next stage of this workflow requires the controls in the Camera tab. This is where you define the actual camera used on set, primarily the film gate size and focal length. This information should have been logged on the set to make available for post-production. When using camera-original media, you can sometimes locate this information in the file metadata.

To locate camera metadata, do the following:

- If you are using DaVinci Resolve, select the MediaIn node with the camera clip, open the Metadata Editor, and view the Camera metadata preset.
- If you are using Fusion Studio, display the metadata subview from the viewer toolbar.



The Camera tab in the Camera Tracker tool

If the actual values are not known, try a best guess. The solver attempts to find a camera near these parameters, and it helps the solver by giving parameters as close to the live action as possible. The more accurate the information you provide, the more accurate the solver calculation. At a minimum, try to at least choose the correct camera model from the Film Gate menu. If the film gate is incorrect, the chances that the Camera Tracker correctly calculates the lens focal length become very low.

Unlike the Track and Solve tabs, the Camera tab does not include a button at the top of the Inspector that executes the process. There is no process to perform on the Camera tab once you configure the camera settings. After you set the camera settings to match the live-action camera, you move to the Solve tab.

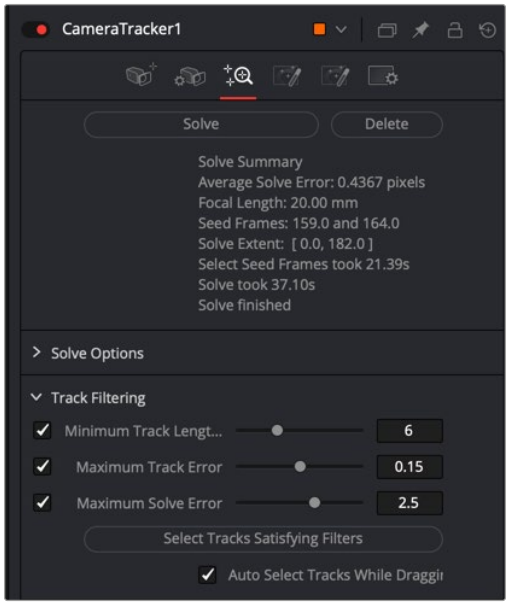
Running the Solver

The next step in this workflow involves the controls found in the Solve tab. Solving is a compute-intensive process in which the Camera Tracker analyzes the currently existing tracks to create a 3D scene. It generates a virtual camera that matches the live action and a point cloud consisting of 3D locators that recreate the tracked features in 3D space. The analysis is based on parallax in the frame, which is the perception that features closer to the camera move quicker than features further away. This is much like when you look out the side window of a car and can see objects in the distance move more slowly than items near the roadside.

The trackers found in the Track phase of this workflow have a great deal to do with the success or failure of the solver, making it critical to deliver the best set of tracking points from the very start. Although the masking you create to occlude objects from being tracked helps to omit problematic tracking points, you almost always need to further filter and delete poor quality tracks in the Solver tab. That's why, from a user's point of view, solving should be thought of as an iterative process.

To solve a camera's motion:

1. Click the Solve button to run the solver.
2. Filter out and delete poor tracks.
3. Rerun the solver.



The Solver tab after it has run and produced an average solve error of 04367 pixels

How Do You Know When to Stop?

At the end of the solve process, an Average Solve Error (sometimes called a reprojection error) appears at the top of the Inspector. This is the crucial value that tells you how well the calculation has gone. A good Average Solve Error for HD content is below 1.0.

You can interpret a value of 1.0 as a pixel offset; at any given time, the track could be offset by 1 pixel. The higher the resolution, the lower the solve error should be. If you are working with 4K material, your goal should be to achieve a solve error below 0.5.

Tips for Solving Camera Motion

When solving camera movement, it's important to provide accurate live-action camera information, such as focal length and film gate size, which can significantly improve the accuracy of the camera solve. For example, if the provided focal length is too far away from the correct physical value, the solver can fail to converge, resulting in a useless solution.

Additionally, for the solver to accurately triangulate and reconstruct the camera and point cloud, it is important to have:

- A good balance of tracks across objects at different depths, with not too many tracks in the distant background or sky (these do not provide any additional perspective information to the solver).
- Tracks distributed evenly over the image and not highly clustered on a few objects or one side of the image.
- The track starts and ends staggered over time, with not too many tracks ending on the same frame.

Using Seed Frames

The solver works by first constructing a partial solution between two seed frames. These seed frames are selected automatically. However, automatic selection adds time to the process. The time is reported in the solve summary at the top of the Inspector once you run the solver. You can select your own seed frames to speed the process and potentially get a better solve on trickier clips. The solver uses these seed frames to create an intermediate solution and then extends this forward and backward for the duration of the clip.

Selecting appropriate seed frames is not necessarily recommended unless you have some experience with camera tracking. Keeping the default Auto Select Seed Frames checkbox enabled in the Solve Options section of the Solver tab selects the best frames in most cases. However, you can disable the checkbox and use the Seed Frame 1 and Seed Frame 2 slider to select frames you believe achieve better results.

When choosing seed frames, it is important to satisfy two main characteristics:

- Seed frames should have lots of tracks in common.
- Seed frames should be far apart in perspective.

Sometimes There's Nothing You Can Do

Some shots that do not have enough camera motion to triangulate feature locations cannot be reconstructed with any useful accuracy. Ensuring that a shot is camera-trackable begins on set, with proper usage of track markers, and by ensuring that camera moves have enough perspective shifts for the solver to glean useful data from.

Cleaning Up Camera Solves

Sometimes the first solve will be good enough. Other times, it may take a few hours of cleaning up tracks to get a good solve, and sometimes it is impossible. With experience, one gets a feel for which tracks should be deleted and which should be kept, and which shots will be easy, difficult, or impossible to solve.

Be aware that deleting too many tracks can cause the Average Solve Error to increase, as the solver has too little information to work with. In particular, if there are fewer than eight tracks on any frame, mathematically there is not enough information to solve the clip. However, it is strongly recommended to use a lot more than eight tracks to get a robust and accurate solve.

IMPORTANT: If you are not familiar with camera tracking, it may be tempting to try to directly edit the resulting 3D splines in the Spline Editor in order to improve a solved camera's motion path. This option should be used as an absolute last resort. It's preferable, instead, to modify the 2D tracks being fed into the solver.

How to Judge Track Accuracy

The automatic color coding of tracking markers makes deleting false or poor tracks easier. After the solver runs, each tracker is assigned a solve error color that indicates which 3D locators match their 2D tracking points well, and which match up poorly.

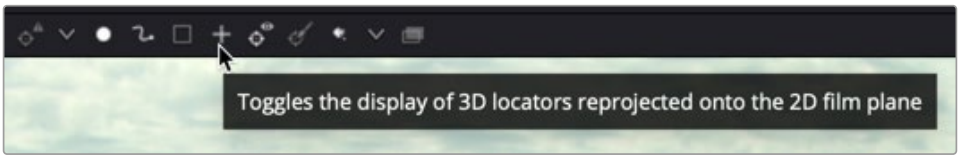
Green: Good. Tracked very well.

Yellow: Moderate confidence. Usually an acceptable track.

Orange: Low Confidence. May be acceptable in some situations.

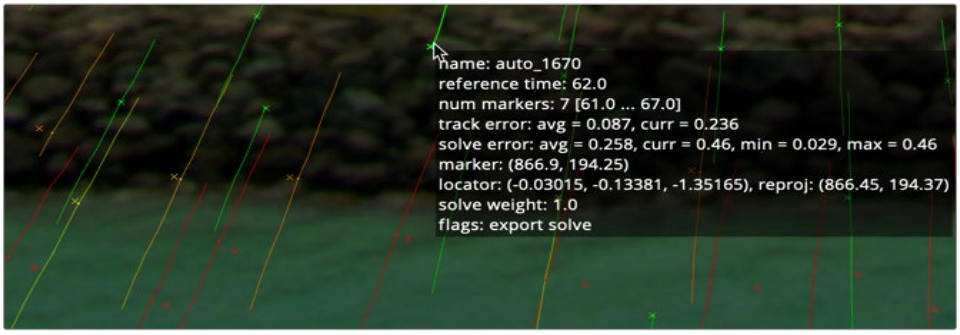
Red: No Confidence. The tracks have not solved well.

Hovering the pointer over any tracking point displays a large metadata tooltip that includes the solve error for the point. For a more visual representation of the accuracy, you can enable the display of 3D locators in the viewer by clicking the Reprojection Locators button in the viewer toolbar.



After a solve, the Camera Tracker toolbar can display Reprojection locators

When the tracking points are converted into a point cloud by the solver, it creates 3D reprojection locators for each tracking point. These Reprojection locators appear as small X marks near the corresponding tracking point. The more the two objects overlap, the lower the solve error.



Reprojection locators displayed with tracking points, and tooltip

The goal when filtering the trackers is to remove all red tracker marks and keep all the green marks. Whether you decide to keep both the yellow and orange or just the yellow is more a question of how many marks you have in the clip. You produce a better solve if you retain only the yellow marks; however, if you do not have enough marks to calculate the 3D scene, you will have to keep some of the better orange marks as well.

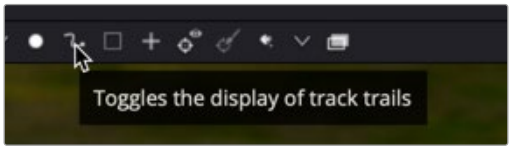
Tips for What to Keep and What to Delete

Understanding what false tracks look like, and then manually cleaning the track data to reduce it to a precise set of clear tracks, will result in a more accurate solve. When cleaning up any track—particularly yellow and orange color coded tracks—keep the following in mind:

- Keep all tracks with motion that's completely determined by the motion of the live-action camera.
- Delete tracks on moving objects or people and tracks that have parallax issues.
- Delete tracks that are reflected in windows or water.
- Delete tracks of highlights that move over a surface.
- Delete tracks that do not do a good job of following a feature.
- Delete tracks that follow false corners created by the superposition of foreground and background layers.
- Consider deleting tracks that correspond to locators that the solver has reconstructed at an incorrect Z-depth.

Deleting Tracks

You can manually delete tracks in the viewer or use filters to select groups of tracks. When deleting tracks in the viewer, it is best to modify the viewer just a bit to see the tracks more clearly. From the Camera Tracker toolbar above the viewer, clicking the Track Trails button hides the trails of the tracking points. This cleans up the viewer to show points only, making it easier to make selections. At the right end of the toolbar, clicking the Darken Image button slightly darkens the image, again making the points stand out a bit more in the viewer.



Hiding trails and darkening the viewer can make it easier to see and select poor tracks

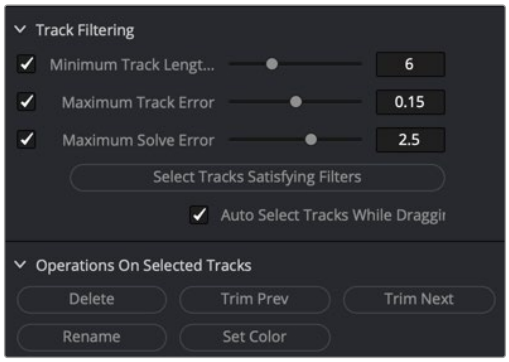
To begin deleting poor-quality tracks, you can drag a selection box around a group of tracks you want to remove and then either click the Delete Tracks button in the Camera Tracker toolbar or press Command-Delete.

You can hold down the Command key to select discontiguous tracking marks that are not near each other. If you accidentally select tracks you want to keep, continue holding the Command key and drag over the selected tracks to deselect them.

When deleting tracks, take note of the current Average Solve Error at the top of the Inspector and then rerun the solver. It is better to delete small groups of tracks and then rerun the solver than to delete one or two large sections. As mentioned previously, deleting too many tracks can have adverse effects and increase the Average Solve Error.

Using Filters to Delete Problem Tracks

The Solve tab includes filters that can be used to select groups of similar tracks by track length, track error, and solve error. These can be used to quickly select and delete poorly performing tracks that may be misleading to the resulting camera, leaving a concise list of accurate tracks.



Tracks can be selected using filters and deleted using the Operations On Selected Tracks buttons

For instance, it is generally best to run the solver using tracks with longer durations. Since shorter tracks tend to be less accurate when calculating the camera, you can remove them using the Filter section in the Inspector.

Increasing the Minimum Track Length parameter sets a threshold that each tracker must meet. Tracks falling below the threshold appear red. You can then click the Select Tracks Satisfying Filters button to select the shorter tracks and click Delete from the Options section in the Inspector.

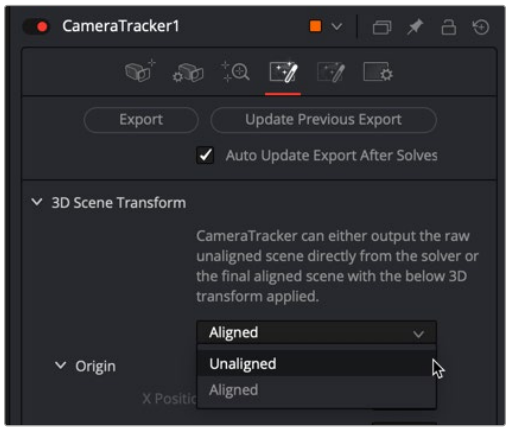
Exporting a 3D Scene for Efficiency

The Camera Tracker saves all its 2D tracks into the composition, sometimes resulting in a rather large file on disk. If you are dealing with a large clip with many 2D tracks over a long duration, the saved composition can reach over a gigabyte in size. Using a Camera Tracker node in a composition can make it cumbersome to load and operate. While it is possible to use the Camera Tracker node directly to composite via the 3D output, you'll achieve better performance by exporting. Once the quality of the solve is satisfactory, the Export tab can generate a "low memory" alternative by producing individual Camera 3D, Point Cloud, Ground Plane, and 3D Renderer nodes.

Before you can export the 3D scene, you must provide a bit more information about it. You'll do this using controls found in the Export tab. Cameras do not include tiltmeters, so clips do not contain metadata that indicates how the camera is tilted or oriented. This is critical information when recreating the virtual camera. It is also useful to determine the location for the center of this 3D scene. The Export tab provides various translation, rotation, and scale controls to set these options.

Unalign the 3D Scene Transforms

By default, the Export tab is set to Aligned in the 3D Scene Transform section. The Aligned setting locks the orientation and scale of the 3D scene to prevent accidentally changing it. So, before you can set the ground plane and origin location, you must change the Camera Tracker to be unaligned using the 3D Scene Transform menu in the Export tab. After you have gone through the Export settings and configured them how you want, you must set the menu back to Aligned before exporting.



Set the 3D Scene Transform menu to Unaligned before setting the ground plane

Setting the Ground Plane

The Camera Tracker has no idea if the camera is on its side or tilted in some way. So, it is up to you to indicate where the ground plane is in a clip. After choosing Unaligned from the 3D Scene Transform menu, you can begin identifying the ground plane.



Drag a selection box around marks that represent the ground

To set the ground plane, do the following:

1. Move to a frame with lots of green 3D locators where you can see a large part of the ground.
2. In the viewer, drag a selection box around the marks located on the ground in the clip.
3. In the Inspector Orientation section, click the Set from Selection button.

> **TIP:** In some cases, the clip you are tracking may not have the ground in the frame. If necessary, you can set the Selection menu to XY, which indicates you are selecting points on a wall.

Setting the Origin

When it comes time to add and position new objects into your 3D scene, you can make it easier by setting the origin or center location.

To set the origin of the 3D scene, do the following:

1. Move to a frame that clearly shows the area you want to select as the center of the scene.
2. In the viewer, either select a single point or drag a selection box around a few marks located where you want to position the center of the 3D scene.
3. In the inspector Origin section, click the Set from Selection button.

Setting the Scale

The Camera Tracker has no idea of the size of the 3D scene, so the scale parameter is used to scale the scene output. This makes it possible to match the scale of two or more clips.

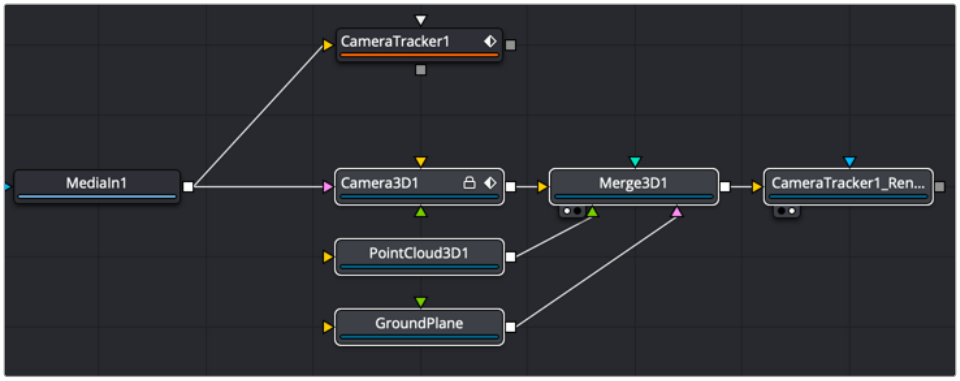
Realign the Scene

Before exporting the scene from the Camera Tracker, you must set the 3D scene Transform menu back to Aligned. Now you're ready to export.

Viewing the Exported Results

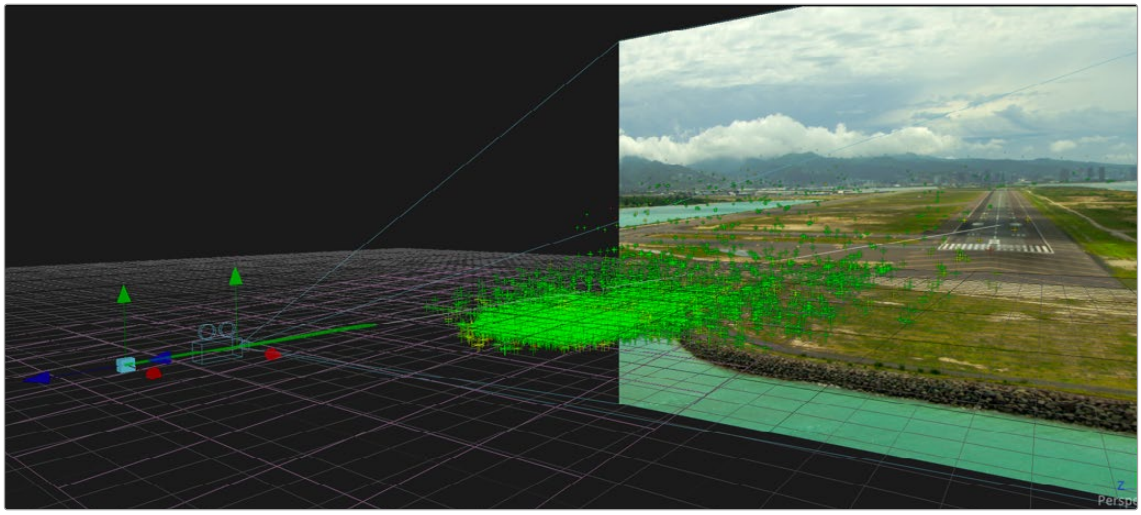
Clicking the Export button at the top of the Inspector creates a functional 3D scene with five new nodes automatically added to the node tree.

- Camera 3D
- Point Cloud
- Ground Plane
- Merge 3D
- Camera Tracker Renderer (3D Renderer)



Five nodes created as a result of exporting from the Camera Tracker

To work with the 3D scene, you can select the Merge 3D and load it into one of the viewers, and then select the Camera Tracker Renderer and load that into a second viewer.





Viewing the Merge 3D shows the point cloud, ground plane, and camera

When the Merge 3D is selected, a toolbar above the viewer can add 3D test geometry like an image plane or cube to verify the precision of the 3D scene and camera. You can then connect actual 3D elements into the Merge 3D as you would any manually created 3D scene. The point cloud can help align and guide the placement of objects, and the CameraTracker Renderer is a Renderer 3D node with all the same controls.



Use the point cloud to accurately place different elements into a 3D scene

At this point, there is no need for the Camera Tracker node unless you find that you need to rerun the solver. Otherwise, you can save some memory by deleting the Camera Tracker node.

Chapter 27

Particle Systems

This chapter is designed to give you a brief introduction to the creation of fully 3D particle systems, one of Fusion's most powerful features.

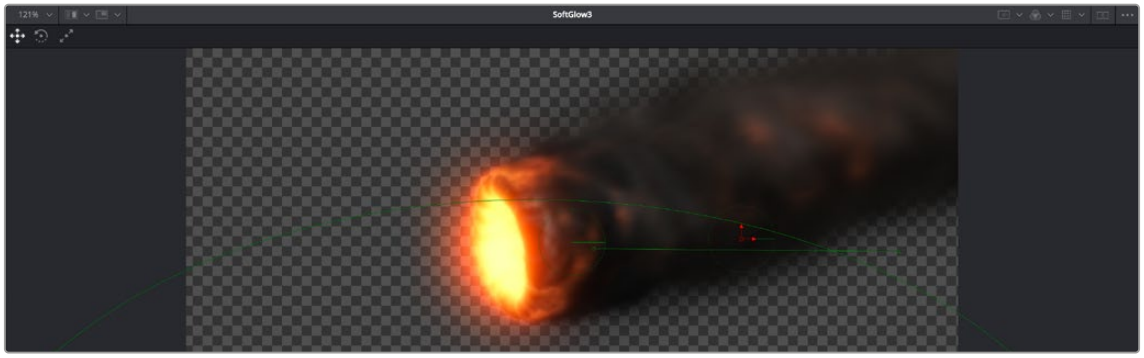
Once you understand these basics, for more Information on each particle system node that's available, *see Chapter 54, "Particle Nodes,"* in the Fusion Reference Manual.

Contents

Introduction to Particle Systems 679
Anatomy of a Simple Particle System 680
Particle System Distribution 682
Particle Nodes Explained by Type 684
Emitters 684
Forces 684
Compositing 685
Rendering 685
Example Particle Systems 685

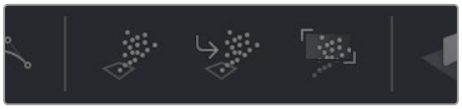
Introduction to Particle Systems

Particle systems are computer simulations that use customizable rules to automatically generate and animate large numbers of elements to simulate smoke, dust, fire, leaves, sparks, or any other animated system of shapes. As Fusion is a full-featured 3D compositing environment, particle systems can be created in 2D or 3D, which makes them incredibly flexible and capable of producing all kinds of visual effects or abstract animated content for use in motion graphics.



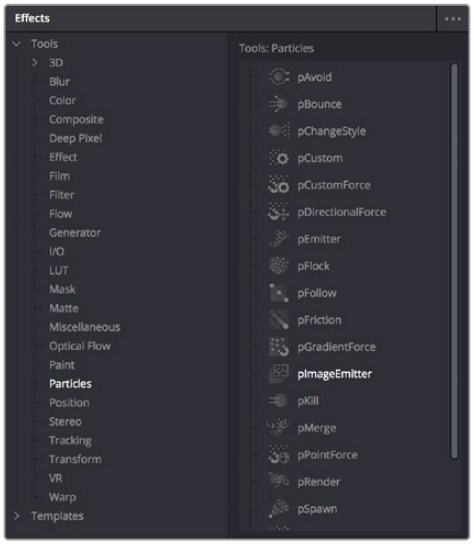
A 3D particle system, also created entirely within Fusion

The three most fundamental nodes required for creating particle systems are found on the toolbar. As with the 3D nodes to the right, these are arranged, from left to right, in the order in which they must be connected to work, so even if you can't remember how to hook up a simple particle system, all you need to do is click the three particle system nodes from left to right to create a functional particle system.



The pEmitter, pMerge, and pRender particle system nodes available from the toolbar

However, these three nodes are only the tip of the iceberg. Opening the Particle category in the Effects Library reveals many, many particle nodes designed to work together to create increasingly complex particle interactions.



A sample of the nodes available in the Particles bin of the Effects Library

All particle nodes begin with the letter “p,” and they’re designed to work together to produce sophisticated effects from relatively simple operations and settings. The next section shows different ways particle nodes can be connected to produce different effects.

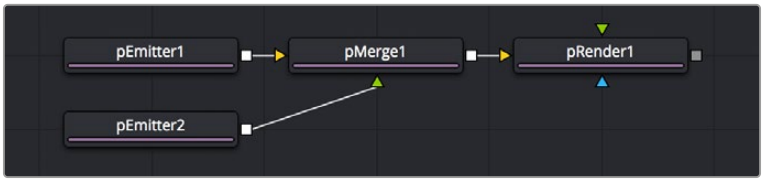
Anatomy of a Simple Particle System

The simplest particle system you can create is a pEmitter node connected to a pRender node. The pEmitter node includes the core controls for creating various kinds of particles in different ways, while the pRender node is required to render a 2D or 3D result that can be composited with other scenes within your composition.



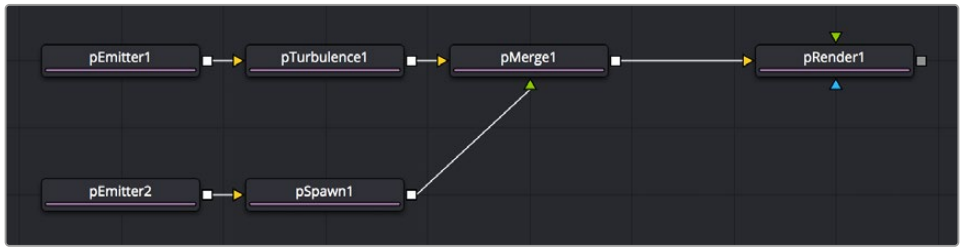
The minimum node tree required to create a simple particle system

If your needs are more complicated, you can combine two or more pEmitter nodes using a pMerge node (the particle system version of a Merge node), to create compound particle systems where multiple types of particles combine with one another to create a result.



Compositing two pEmitter nodes to create a compound particle system, combining two kinds of particles together

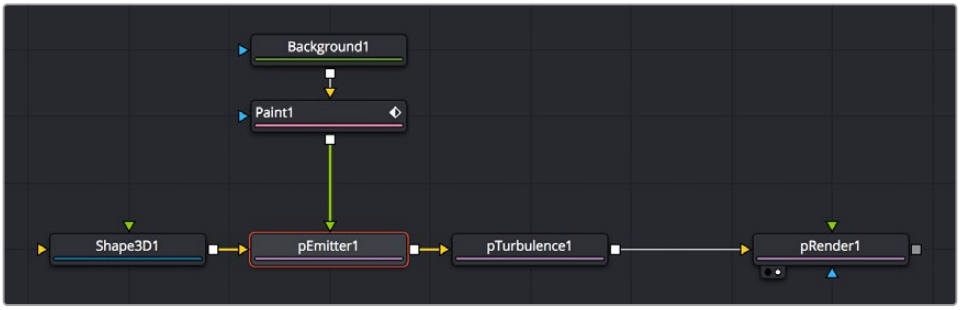
If you’re trying to create particle systems with more natural effects, you can add “forces” to each emitter. These forces are essentially physics or behavioral simulations that automatically cause the particles affected by them to be animated with different kinds of motion, or to be otherwise affected by different objects within scenes.



Customizing the effect of pEmitter nodes using different forces to add complexity to the particle animation

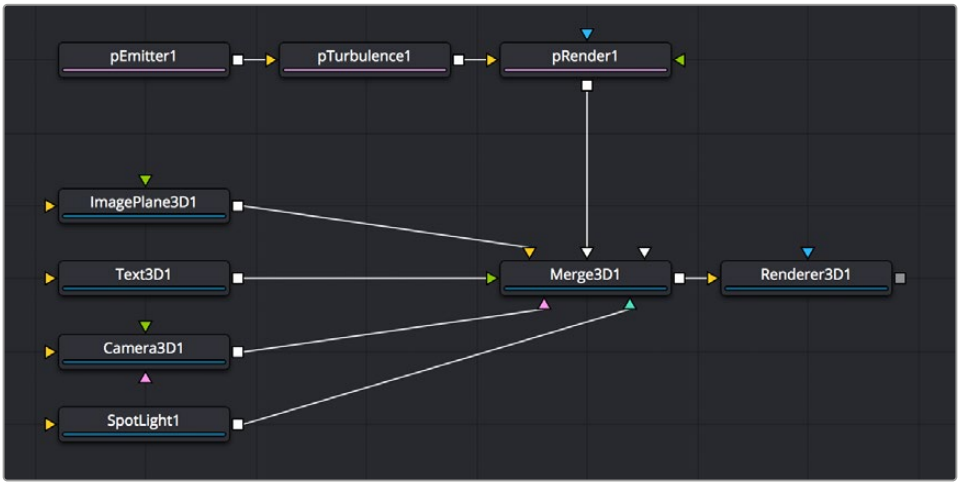
You can also attach the following types of nodes to a pEmitter node to deeply customize a particle system:

- Attach a 2D image to a pEmitter node to create highly customized particle shapes. Make sure your image has an appropriate alpha channel.
- Attach a Shape3D or other 3D geometry node to a pEmitter node to create a more specific region of emission (by setting Region to Mesh in the Region tab).



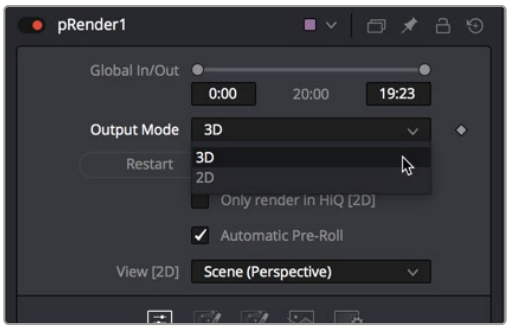
Customizing pEmitter nodes using mesh geometry to define regions and 2D images to define particle shape

The above examples assume that you'll output 2D renders to combine into the rest of a 2D composition. However, because particle systems are fully 3D, you also have the option of outputting your particle system in such a way as to be used from within other 3D scenes in your composition.



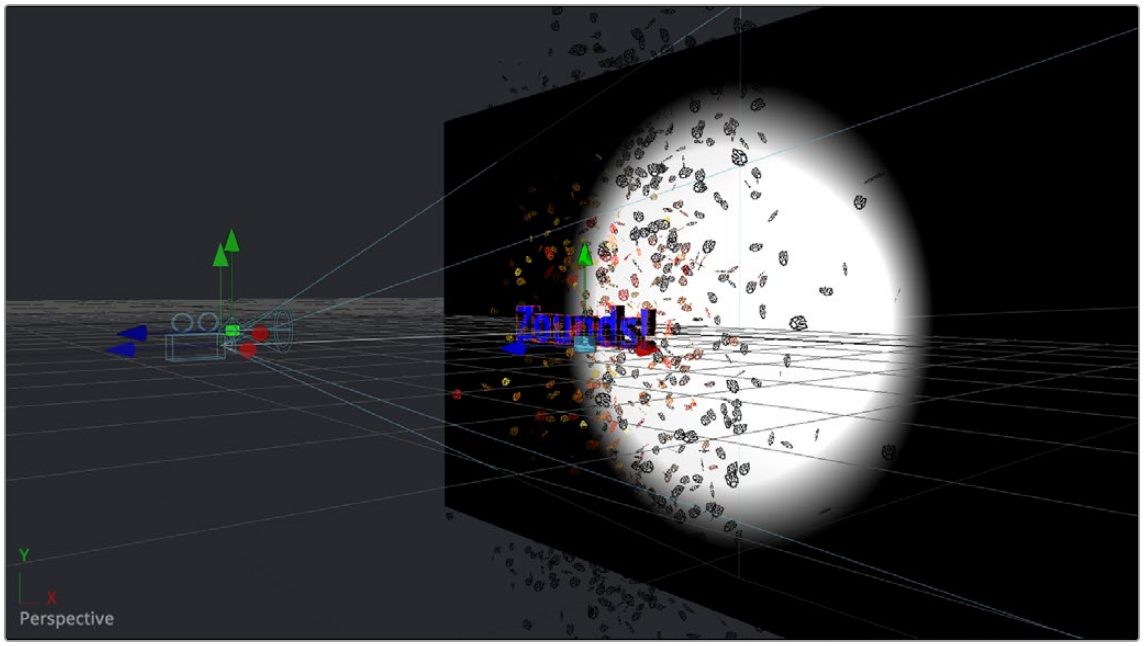
Connecting a particle system to a Merge3D node so the particles are subject to lighting and shadows within a 3D scene

The Output Mode of the pRender node, at the very top of the controls exposed in the Inspector, can be set to either 2D or 3D, depending on whether you want to combine the result of the particle system with 2D layers or with objects in a 3D scene.



Choosing whether a particle system's output is 2D or 3D in the pRender node's Inspector controls

If you connect a pRender node to a Merge3D node, the Output Mode is locked to 3D, meaning that 3D geometry is output by the pRender node for use within the Merge3D node's scene. This means that the particles can be lit, they can cast shadows, and they can interact with 3D objects within that scene.

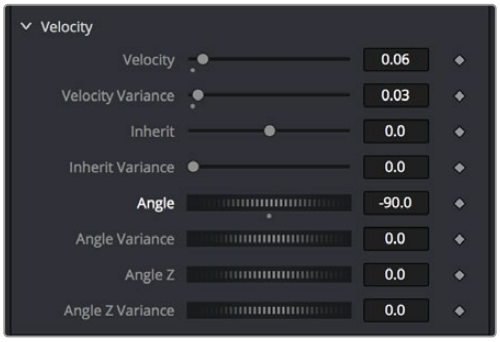


The result of using a particle system within a 3D scene

NOTE: Once you set the pRender node to either 2D or 3D and make any change to the nodes in the Inspector, you cannot change the output mode.

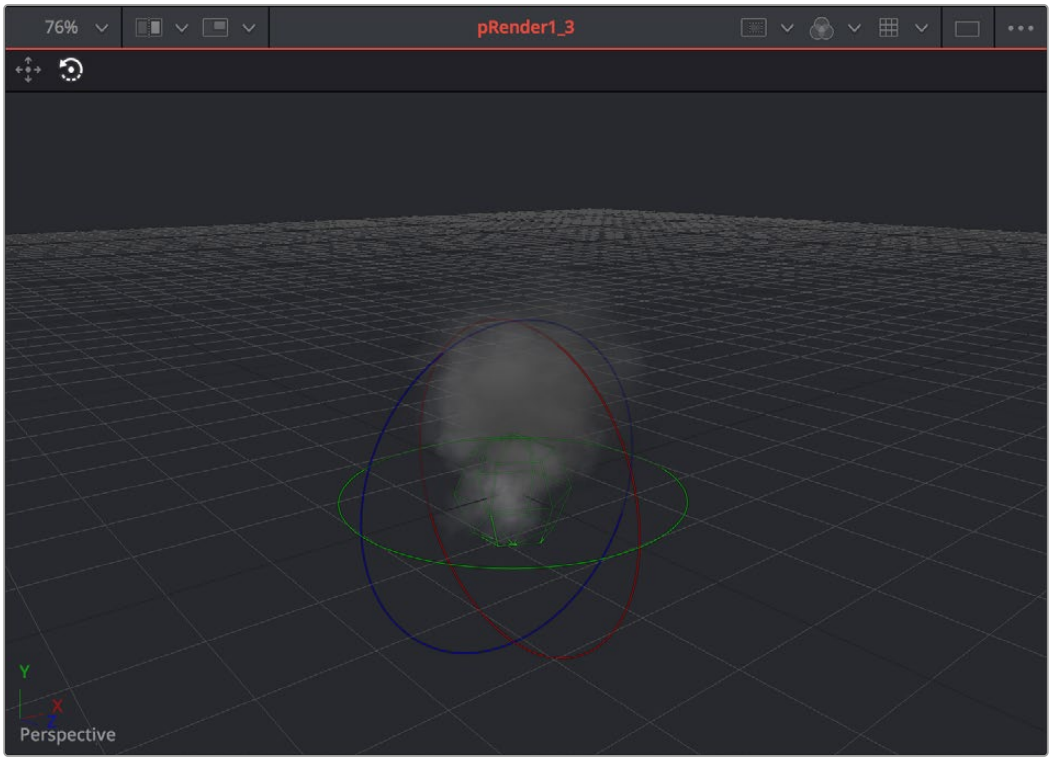
Particle System Distribution

To adjust the distribution of particles being emitted, select the pEmitter node to expose its controls in the Inspector, then open the Velocity controls in the Controls tab, and use the Angle, Angle Variance, Angle Z, and Angle Z Variance controls to adjust the direction and width over which particles are emitted. All these controls can be animated.



A pEmitter node's Velocity Angle and Angle Variance controls let you adjust the direction and width of particle distribution

Particle systems can be positioned and rotated by loading the pEmitter nodes that generate particles into a viewer and using the onscreen 3D position and Rotation controls provided to move the particle system around.



A pEmitter node loaded into the viewer with the rotation onscreen controls enabled

Alternatively, you can use the controls of the pEmitter's Region tab in the Inspector to adjust Translation, Rotation, and Pivot. All these controls can be animated.



A pEmitter node's Region controls open in the Inspector

Particle Nodes Explained by Type

This section introduces the four types of particle system nodes available in the Effects Library.

Emitters

pEmitter nodes are the source of all particles. Each pEmitter node can be set up to generate a single type of particle with enough customization so that you'll never create the same type of particle twice. Along with the pRender node, this is the only other node that's absolutely required to create a particle system.

pEmitter nodes have four parameters tabs:

Controls: The primary controls governing how many particles are generated (Number), how long they live (Lifespan), how fast they move (Velocity) and how widely distributed they are (Angle and Angle Variance), their rotation (Rotation Mode with X, Y, and Z controls), and whether there's spin (Spin X, Y, and Z controls). For each parameter of particle generation, there's an accompanying Variance control that lets you make that parameter less uniform and more natural by introducing random variation.

Sets: This tab contains settings that affect the physics of the particles emitted by the node. These settings do not directly affect the appearance of the particles. Instead, they modify behaviors such as velocity, spin, quantity, and lifespan.

Style: While the Controls tab has a simple control for choosing a color for particles, the Style tab has more comprehensive controls including color variance and Color Over Life controls. Additionally, size controls including Size Over Life, fade controls, and blur controls let you create sophisticated particle animations with a minimum of adjustments, while Merge controls give you an additional level of control over how overlapping particles combine visually. A set of controls at the bottom lets you choose how animated effects are timed.

Region: The Region tab lets you choose what kind of geometric region is used to disperse particles into space and whether you're emitting particles from the region's volume or surface. The Winding Rule and Winding Ray Direction controls determine how the mesh region will handle particle creation with geometric meshes that are not completely closed, as is common in many meshes imported from external applications. Tweaking these last parameters is common when using imported mesh geometry as a region for emitting particles, since even geometry that appears closed will frequently appear to "leak" particles thanks to improperly welded vertices.

Forces

Many of the particle nodes found in the Particles bin of the Effects Library are "forces" that enhance a particle simulation by simulating the effect of various forces acting upon the particles generated by an emitter.

Some forces, including pDirectionalForce, pFlock, pFriction, pTurbulence, and pVortex, are rules that act upon particles without the need for any other input. These are simply "acts of nature" that cause particles to behave in different ways.

Other forces, such as pAvoid, pBounce, pFollow, and pKill, work in conjunction with 3D geometry in a scene such as shapes or planes to cause things to happen when a particle interacts or comes near that geometry. Note that some of the particles described previously can also use geometry to direct their actions, so these two categories of forces are not always that clear-cut.

Compositing

The pMerge node is a simple way to combine multiple emitters so that different types of particles work together to create a sophisticated result. The pMerge node has no parameters; you simply connect emitters to it, and they're automatically combined.

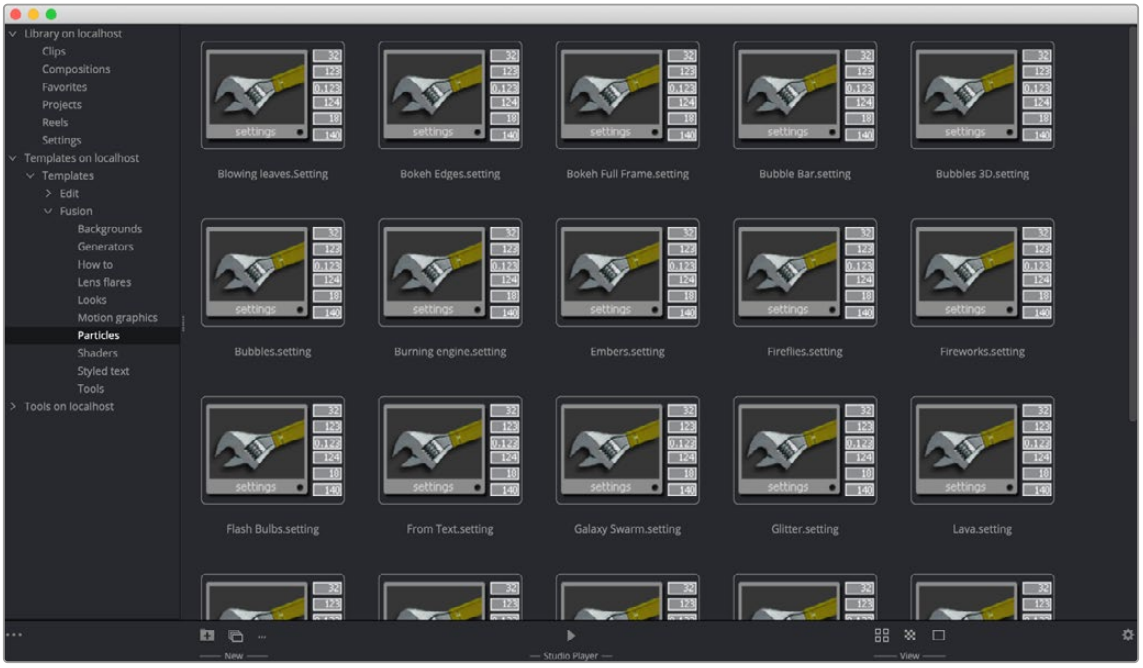
Rendering

The pRender node is required whether you're connecting a particle system's output to a 2D Merge node or to a Merge3D node for integration into a 3D scene. Along with the pEmitter node, this is the only other node that's absolutely required to create a particle system.

- **Controls:** The main controls that let you choose whether to output 2D or 3D image data, and whether to add blur or glow effects to the particle systems, along with a host of other details controlling how particles will be rendered.
- **Scene:** These controls let you transform the overall particle scene all at once.
- **Grid:** The grid is a helpful, non-rendering guide used to orient 2D particles in 3D space. The grid is never output in renders. The width, depth, number of lines, and grid color can be set using the controls found in this tab.
- **Image:** Controls the output of the pRender node, with controls over the process mode, resolution, and color space settings of the output.

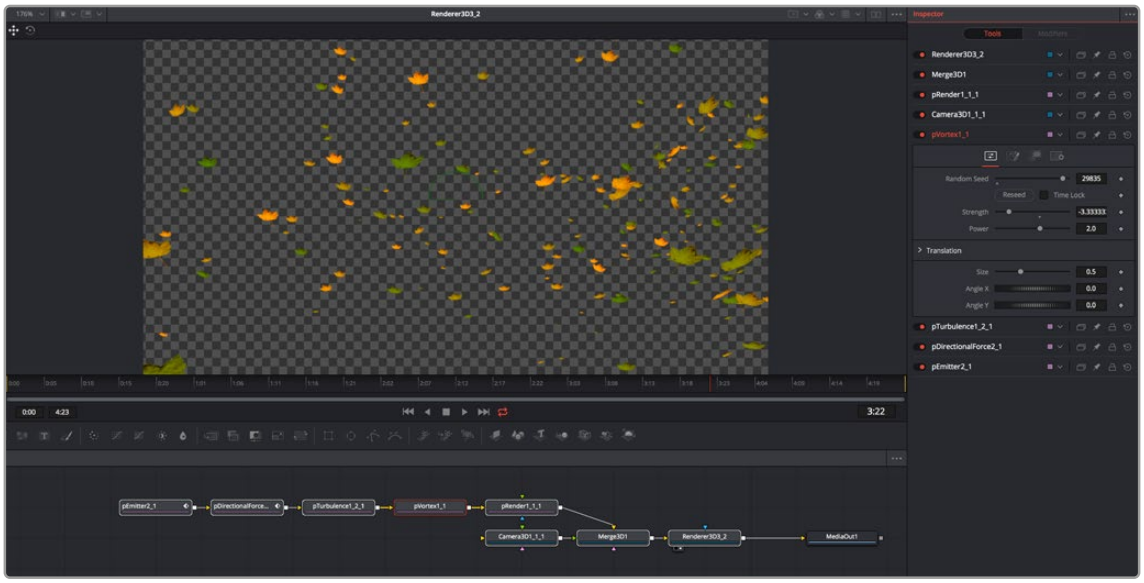
Example Particle Systems

The Templates category in the Inspector in the Fusion page of DaVinci Resolve or in the Bins window in Fusion Studio includes over 20 different examples of particle systems creating a variety of effects. One of the best ways of learning how to create and customize particle systems is to open these and investigate how they're made.



Different particle system presets in the Templates category of the Bins window in Fusion Studio

Simply drag and drop any of the particle presets into the Node Editor, load the last node into the viewer, and you'll see how things are put together.



The Blowing Leaves preset from the Templates category



Advanced Compositing Techniques

FUSION

CONTENTS

28 Optical Flow and Stereoscopic Nodes 688

Chapter 28

Optical Flow and Stereoscopic Nodes

This chapter covers the numerous stereoscopic and optical flow-based nodes available in Fusion and their related workflows.

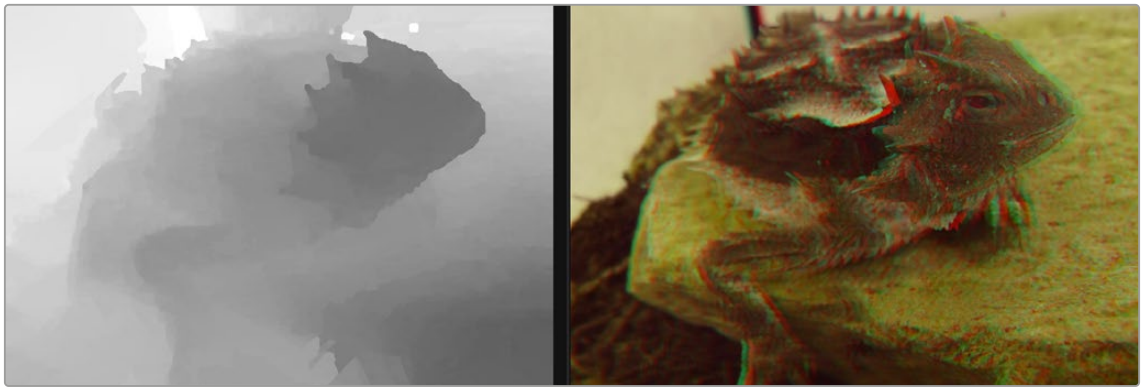
Contents

Overview 689
Stereoscopic Overview 689
Optical Flow Overview 689
Toolset Overview 690
Working with Aux Deep Channels 690
Optical Flow Workflows 691
OpticalFlow 691
TimeSpeed, TimeStretcher 692
SmoothMotion 692
Repair Frame, Tween 692
Advanced Optical Flow Processing 692
Stereoscopic Workflows 693
Stereo Camera 693
Stereo Materials 694
Disparity 694
NewEye, StereoAlign 694
DisparityToZ, ZToDisparity 695
Separate vs. Stack 695
Setting Up Stereo in the Node Editor 695
About the Disparity Channel 696
Viewing of Disparity and Vector Channels 696

Stereo and Optical Flow Best Practices 697
Semi-Transparent Objects 697
Motion Blur 697
Depth of Field 697
Where to Calculate Disparity and Optical Flow? 697
Cropping the Source 698
Nodes with Multiple Outputs 698
Picking from Aux Channels 698
Vector and Disparity Channels 698

Overview

Fusion includes 3D stereoscopic and optical flow-based nodes, which can work together or independently of each other to create, repair, and enhance 3D stereoscopic shots.



Stereoscopic comp displayed in the viewers.

Stereoscopic Overview

All stereoscopic features are fully integrated into Fusion's 3D environment. Stereoscopic images can be created using a single camera, which supports eye separation and convergence distance, and a Renderer 3D for the virtual left and right eye. It is also possible to combine two different cameras for a stereo camera rig.

Stereoscopic nodes can be used to solve 3D stereoscopic shooting issues, like 3D rig misalignment, image mirror polarization differences, camera timing sync issues, color alignment, convergence, and eye separation issues. The stereo nodes can also be used for creating depth maps.

Optical Flow Overview

Optical Flow analyzes the motion in a clip and generates motion vectors between neighboring frames. It generates X and Y vectors from the previous frame to the current frame (Back Vectors) and to the next frame in sequence (Forward Vectors). Once calculated, optical flow data can be used by other nodes to create smooth slow motion and variable retiming of clips, repair missing frames, and even correct disparity in stereo 3D clips.

NOTE: The stereoscopic nodes in the Fusion page work independently of the stereoscopic tools in the other DaVinci Resolve pages.

Toolset Overview

Here is an overview of the available nodes.

Optical Flow Nodes

- **Optical Flow > OpticalFlow:** Analyzes motion between neighboring frames in a sequence to generate motion vectors, which can then be used by other nodes for retiming, motion blur, and other effects.
- **Miscellaneous > TimeSpeed:** Retimes a clip at a constant speed using Flow Interpolation mode.
- **Miscellaneous > TimeStretcher:** Retimes a clip at variable speeds using Flow Interpolation mode.
- **Optical Flow > RepairFrame:** Generates a new frame using the motion vectors between two neighboring frames.
- **Optical Flow > SmoothMotion:** Smoothes the color or aux channels using motion vectors.
- **Optical Flow > Tween:** Interpolates between two non-sequential images to generate a new frame.
- **Color > CopyAux:** Copies aux channels, including motion vectors, into RGBA more efficiently than Channel Booleans.

Stereoscopic Nodes

- **Stereo > Anaglyph:** Combines stereo images to create a single anaglyph image for viewing.
- **Stereo > Combiner:** sStacks a separate stereo images into a single stacked pair, so they can be processed together.
- **Stereo > Disparity:** Generates disparity between left/right images.
- **Stereo > DisparityToZ:** Converts disparity to Z-depth.
- **Stereo > Global Align:** Shifts each stereo eye manually to do basic alignment of stereo images.
- **Stereo > NewEye:** Replaces left and/or right eye with interpolated eyes.
- **Stereo > Splitter:** Separates a stacked stereo image into to left and right images.
- **Stereo > StereoAlign:** Adjusts vertical alignment, convergence, and eye separation.
- **Stereo > ZToDisparity:** Converts Z-depth to disparity.

Working with Aux Deep Channels

Certain image formats can contain channels other than RGBA color, called aux deep channels. Stereo Disparity and OpticalFlow deal directly with auxiliary deep channels.

Aux channels supported in Fusion include:

- **RGBA:** These are the standard colors.
- **Z:** The eyespace Z coordinate is almost always negative because in eyespace, Fusion's camera sits at (0, 0, 0) looking down the Z-axis. Z values start at Z = 0 at the camera focal point and progressively become more negative for objects deeper in the scene.
- **Coverage:** The percentage of the pixel covered by the frontmost pixel, used for antialiased Z-compositing.

- **Object ID:** These are user-assigned integers to meshes.
- **Material ID:** These are user-assigned integers to materials.
- **Texture Coords:** Normalized texture coordinates stored as (u, v) pairs.
- **Normal Vector:** Normal vector (nx, ny, nz) where the components are typically in the range [–1, +1].
- **Background Color:** The color of the pixel if the frontmost layer were removed, used for antialiased Z-compositing.
- **Vector:** The forward motion vector is an offset (vx, vy) that compares every pixel's position in one frame to the same pixel's position in the next frame.
- **Back Vector:** The backward motion vector is an offset (vx, vy) that compares every pixel's position in one frame to the same pixel's position in the previous frame.
- **World Position:** The position (wx, wy, wz) of the pixel in world coordinates.
- **Disparity:** An offset (dx, dy) that maps a pixel in the Left > Right or Right > Left frames.

Some extra channels are used by specific Fusion nodes.

For example:

- Merge can use the Z channel to perform a depth merge. If the Coverage and BackgroundColor channels are present, it can do a better job on antialiased edges during the Z merge.
- Most image-processing nodes (e.g., BrightnessContrast) have options on their common controls tab to limit their processing by MaterialID and ObjectID.
- The Fog and DepthBlur nodes make use of the Z channel.
- The Texture node makes use of the TexCoord channel.
- The Shader node makes use of the Normal channel.

There are a couple of ways to retrieve or generate those extra channels within Fusion.

For example:

- The Renderer3D node is capable of generating most of these channels.
- The OpticalFlow node generates the Vector and BackVector channels, and then TimeStretcher and TimeSpeed can make use of these channels.
- The Disparity node generates the Disparity channels, and then DisparityToZ, NewEye, and StereoAlign nodes can make use of the Disparity channels.
- The OpenEXR format can be used to import or export aux channels into Fusion by specifying a mapping from EXR attributes to Fusion Aux channels using CopyAux.

Optical Flow Workflows

The Optical Flow analysis is a non real-time process, and depending on your computer, the clip's resolution, and the duration of the clip, it can take some time. Because of this, the general idea is that you pre-generate the motion vectors, either by performing the analysis overnight or using a render farm, and save results into an OpenEXR sequence. The Optical Flow toolset is designed around four types of nodes that either generate, destroy, pass through, or construct the motion vectors.

OpticalFlow

The Optical Flow node generates the Vector and BackVector data. Typically, for optimal performance, you connect the Optical Flow output to a Saver to save the image as OpenEXR files with the motion vectors stored in an aux channel.

TimeSpeed, TimeStretcher

You can create smooth constant or variable slow-motion effects using the TimeSpeed or TimeStretcher nodes. When Optical Flow motion vectors are available in the aux channel of an image, enabling Flow mode in the TimeSpeed or TimeStretcher Interpolation settings will take advantage of the Vector and BackVector channels. For the Flow mode to work, there must be either an upstream OpticalFlow node generating the hidden channels or an OpenEXR Loader bringing these channels in. These nodes use the Vector/BackVector data to do interpolation on the motion channel and then destroy the data on output since the input Vector/BackVector channels are invalid.

For more detail on TimeSpeed or TimeStretcher, *see Chapter 51, "Miscellaneous Nodes,"* in the Fusion Reference Manual.

SmoothMotion

SmoothMotion can be used to smooth the Vector and BackVector channels or smooth the disparity in a stereo 3D clip. This node passes through, modifies, or generates new aux channels, but does not destroy them.

Repair Frame, Tween

The Tween and Repair Frame nodes are different from standard optical flow nodes because they have the OpticalFlow analysis and motion vector generation built in. Tween will compare two frames and create an in-between frame, which is good for recreating a missing or flawed frame. Repair Frame will look at frames on either side of the current frame and repair scratches, dust marks, and so on. Because these nodes work with flow values between non-sequential frames, they cannot use the optical flow stored in the input image's Vector/BackVector channels, but rather must regenerate the flow of each frame, do their processing, and then destroy the flow channels. This being the case, these nodes are computationally expensive.

For more detail on Tween or Repair Frame, *see Chapter 52, "Optical Flow,"* in the Fusion Reference Manual.

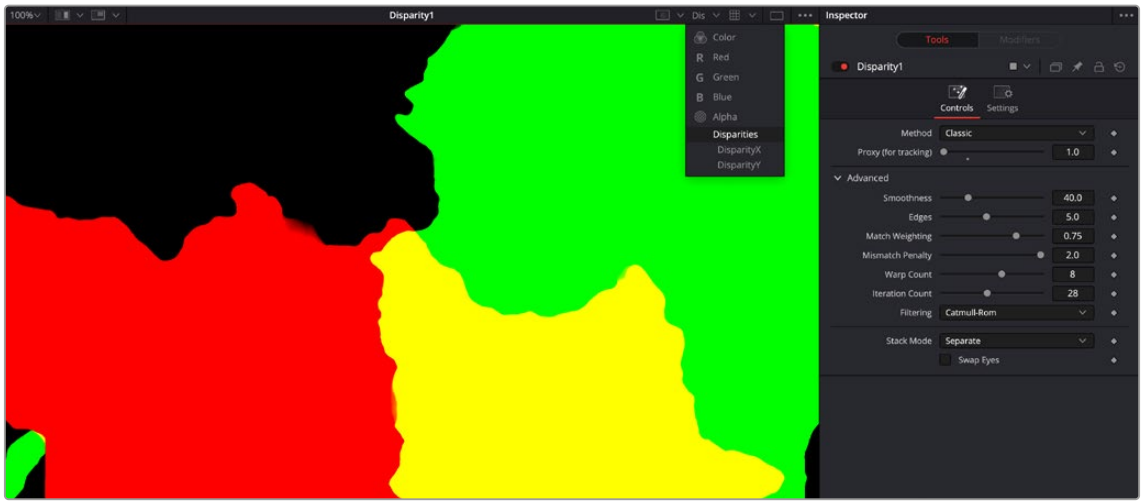
Advanced Optical Flow Processing

The Optical Flow, Repair Frame, and Tween nodes include a faster GPU-based Optical Flow algorithm. When you add the Optical Flow, Repair Frame, or Tween node to a comp, the Inspector includes a Method drop-down menu where you can choose Advanced to enable the GPU-based algorithm. This Advanced method is the same Optical Flow algorithm used in other DaVinci Resolve pages.

By choosing Classic from the Method drop-down menu in the Inspector, you can use the older CPU-based algorithm to maintain compatibility with comps created in previous versions. This method may also be better suited for some Stereo3D processing.

Stereoscopic Workflows

Disparity is the difference between the left and right image. The Disparity map is used by nodes to align and massage the stereo pair of images.



The Disparity node analyzes a stereo pair of images and generates an X&Y disparity map.

The workflow is to load a left and right stereo image pair and process those in the Disparity node. Once the Disparity map is generated, other nodes can process the images.

> **TIP:** When connectng stereo pairs in the node tree, make sure that the left and right images are connected to the left and right inputs of the Disparity node.

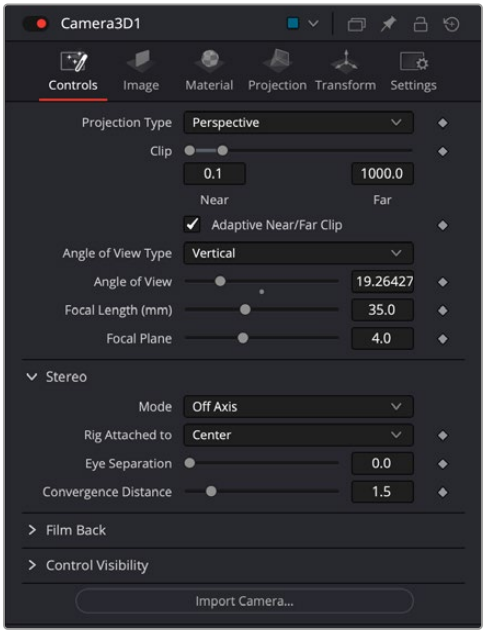
Disparity generation, like Optical Flow, is computationally expensive, so the general idea is that you can pre-generate these channels, either overnight or on a render farm, and save them into an EXR sequence.

The toolset is designed around this philosophy.

Stereo Camera

There are two ways to set up a stereoscopic camera. The common way is to simply add a Camera 3D and adjust the eye separation and convergence distance parameters.

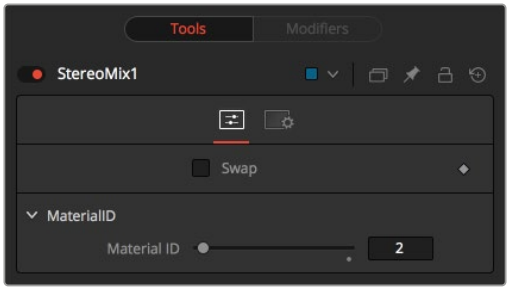
The other way is to connect another camera to the RightStereoCamera input port of the Camera 3D. When viewing the scene through the original camera or rendering, the connected camera is used for creating the right-eye content.



Stereoscopic cameras can be done with a single camera or two connected cameras.

Stereo Materials

Using the Stereo Mix material node, it is possible to assign different textures per eye.



Material viewer showing stereoscopic material.

Disparity

The Disparity node does the heavy lifting of generating disparity maps. This generates the Disparity channel and stores it in the hidden aux channels of their output image.

NewEye, StereoAlign

NewEye and StereoAlign use and destroy the Disparity channel to do interpolation on the color channel.

The hidden channels are destroyed in the process because, after the nodes have been applied, the original Disparity channels would be invalid.

For these nodes to work, there must be either an upstream Disparity node generating the hidden channels or an OpenEXR Loader bringing these channels in.

DisparityToZ, ZToDisparity

These nodes pass through, modify, or generate new aux channels, but do not destroy any.

TIP: If the colors between shots are different, use Color Corrector or Color Curves to do a global alignment first before calculating the Disparity map. Feed the image you will change into the orange input and the reference into the green input. In the Histogram section of the Color Corrector, select Match, and also select Snapshot Match Time. In the Color Curves' Reference section, select Match Reference.

Separate vs. Stack

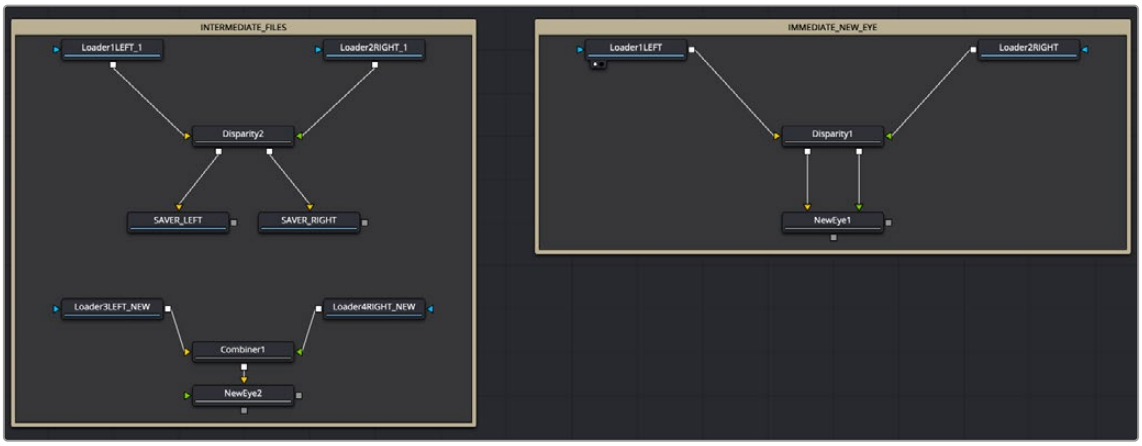
Stereo nodes can work in Separate or Stack modes. When in Stack mode, the left/right eyes are stacked horizontally or vertically, forming one image with double width or height, respectively.

The advantage to using Stack mode is that you do not have to have duplicate branches of the Node Editor for the left and right eyes. As a consequence, you will see Stereo nodes with two inputs and two outputs labeled as "Left" and "Right."

When in Stack mode, the stack should be connected to the left eye input and the Left output should be used for connecting further nodes. In Stack mode, the respective Right eye inputs and outputs are hidden.

Setting Up Stereo in the Node Editor

The disparity generation is the first operation. This can be configured in the Node Editor in two different ways.



Two stereoscopic workflows.

In the above example, the workflow on the right takes the left and right eye, generates the disparity, and then NewEye is used to generate a new eye for the image right away.

The example on the left renders the frames with disparity to intermediate EXR images. These images are then loaded back into Stereo nodes and used to create the NewEye images.

By using Render nodes to compute the disparity first, the later processing of the creative operations can be a much faster and interactive experience.

Although not shown in the above diagram, it is usually a good idea to color correct the right eye to be similar to the left eye before disparity generation, as this helps with the disparity-tracking algorithm. The color matching does not need to be perfect—for example, it can be accomplished using the "Match" option in a Color Corrector's histogram options.

About the Disparity Channel

The Disparity channel stores the displacement vectors that match pixels in one eye to the other eye. The left image's Disparity channel will contain vectors that map left>right and the right image's Disparity channel will contain vectors that map right>left. For example:

```
(xleft, yleft) + (Dleft. x, Dleft. y) -> (xright, yright) (xright, yright) +
(Dright. x, Dright. y) -> (xleft, yleft)
```

You would expect for non-occluded pixels that Dleft = -Dright, although, due to the disparity generation algorithm, this is only an approximate equality.

NOTE: Disparity stores both X and Y values because rarely are left/right images perfectly registered in Y, even when taken through a carefully set up camera rig.

Both Disparity and Optical Flow values are stored as un-normalized pixel shifts. In particular, note that this breaks from Fusion's resolution-independent convention. After much consideration, this convention was chosen so the user wouldn't have to worry about rescaling the Disparity/Flow values when cropping an image or working out scale factors when importing/exporting these channels to other applications. Because the Flow and Disparity channels store things in pixel shifts, this can cause problems with Proxy and AutoProxy. Fusion follows the convention that, for proxied images, these channels store unscaled pixel shifts valid for the full-sized image. So if you wish to access the Disparity values in a script or via a probe, you need to remember to always scale them by (image. Width/image. OriginalWidth, image. Height/ image. OriginalHeight).

Viewing of Disparity and Vector Channels

Aux channels can be displayed directly in the viewers through the Channel viewer button's menu.

The CopyAux node is used to copy those channels directly into the RGB channels for viewing or further processing. The advantage of using the CopyAux node is that it does static normalization, which reduces a lot of flicker that the viewer's time-variant normalization causes. When viewing long sequences of aux channels, the CopyAux node has the option to kill off aux channels and keep only the current RGB channels, freeing up valuable memory so you can cache more frames.

TIP: Although you can use the Channel Booleans to copy any aux channel into RGBA, it involves a few additional clicks when compared to CopyAux.

One thing to be aware of is that aux channels tend to consume a lot of memory. A float-32 1080p image containing just RGBA uses about 32 MB of memory, but with all the aux channels enabled it consumes around 200 MB of memory.

Stereo and Optical Flow Best Practices

How you create your composition, the images you are using, and the type of shot you are working on can all have an impact on the success of the Disparity generation and Optical Flow analysis. Below, we'll look at some of the situations to be aware of and how you can avoid some pitfalls when dealing with optical flow.

Semi-Transparent Objects

The Optical Flow and Disparity generation algorithms Fusion uses assume there is only one layer per pixel when tracking pixels from frame to frame. In particular, transparent objects and motion blur will cause problems. For example, a shot flying through the clouds with the semi-transparent clouds in the foreground and a distant landscape background will confuse the Optical Flow/Stereo algorithms, as they do not recognize overlapping objects with different motions. Usually the optical flow will end up tracking regions of one object or the other. If the transparent object and the background are near the same depth and consequently have the same disparity, then it is not a problem.

Motion Blur

Motion blur is also a serious problem for the reason explained in the previous point. The Disparity and Optical Flow algorithms are unsure whether to assign a pixel in the motion blur to the moving object or the background pixel. Because the algorithms used are global in nature, not only the vectors on the motion blur will be wrong, but it will confuse the algorithm on regions close to the motion blur.

Depth of Field

Depth of field is also another problem related to the above two problems. The problem occurs when you have a defocused foreground object over a background object that is moving (Optical Flow case) or shifts between L/R (Stereo Disparity case). The blurred edges will confuse the tracking because they can't figure out that the edges are actually two separate objects.

Where to Calculate Disparity and Optical Flow?

Where you choose to generate optical flow or disparity in your composition can drastically affect the results.

For example, if you have composited a lens flare in, it is better to compute OpticalFlow/Disparity before that, since the semi-transparent lens flare will confuse the tracking algorithms.

If you are color correcting the left/right eyes to match or for deflickering, it is better to apply the OpticalFlow/Disparity afterward, since it will be easier for the tracking algorithm to find matches if the color matches between frames.

If you are removing lens distortion, think carefully about whether you want to do it before or after Disparity computation. If you do it after, your Disparity map will also act as a lens distortion map, combining the two effects as one.

As a general rule of thumb, it is best to use OpticalFlow/Disparity before any compositing operations except an initial color matching correction and a lens distortion removal.

Cropping the Source

As a general tip, if you are cropping down your input images for any reason, it is probably better to compute the optical flow or disparity before the crop and then afterward crop the flow/disparity along with the color.

The reason is that flow/disparity matching works well when there is common pixel data to match in both frames, but when there are pixels that show up in just one frame (or one eye), then the Disparity/OpticalFlow nodes must make a guess and fill in the data. The biggest occlusions going from L <–> R are usually pixels along the L/R edges of the images that get moved outside. This is similar for optical flow when you have a moving camera.

Another thing to be aware of are black borders around the edges of your frames, which you should crop away.

Nodes with Multiple Outputs

Many of the stereo nodes in the Fusion toolset have multiple outputs. This can cause some confusion to new users. One particularly confusing thing is that when you drag a Stereo node to the view, it will always display the left output. There is no way to view the right output without connecting another node like BC (BrightnessContrast) to the right output and viewing that.

Picking from Aux Channels

Some nodes, like StereoAlign, allow one to drag pick from the Z or Disparity auxiliary channels. You must pick from a node upstream of the StereoAlign, not from the output of the StereoAlign. If you try to pick a disparity from the output of a StereoAlign node, you will get nothing because StereoAlign consumes/destroys the Disparity aux channel (and even if it did not destroy the Disparity channel, you would still be picking the wrong value since you would be picking from the aligned result).

The typical workflow for picking is:

1. View StereoAlign in the left view.
2. View the node upstream of StereoAlign in the right view.
3. Pick the Disparity value from the left eye in the right view.

Although this picking functionality does not operate any differently from normal picking of color channels, this issue may cause some confusion. If it helps, the analogous workflow mistake with color nodes would be a user trying to pick a gradient color for a Background node from a view showing the Background node itself (you are trying to pick a color for a node from its own output).

Another issue that you need to be aware of is which eye you are picking. To avoid problems, it's a good idea to always pick from the left eye. The reason is that the Disparity channels for the left and right eyes are different, and when you pick from a horizontal/vertical stereo stack, Fusion has no way of knowing whether you picked the Disparity value from the left or right eye.

The above are not hard and fast rules; rather, they are guidelines to prevent foot shootings. If you understood the above reasoning fully, you'll realize there are exceptions, like picking disparity from the left output of DisparityToZ and Z from the left/right output of ZToDisparity, where everything is okay.

Vector and Disparity Channels

The Vector and BackVector channels store the forward and reverse optical flow.

The Vector channel might be better named "forward vector" or "forward flow," since the name "Vector" to describe a channel is "not technically correct," as the more mathematically-inclined user might recognize that all the channels except the scalar channels Z/ID are technically "vector" channels.

A frames Vector aux channel will store the flow forward from the current frame to the next frame in the sequence, and the BackVector aux channel will store the flow backward from the current frame to the previous frame. If either the previous or next frames do not exist (either not on disk or the global range of a Loader does not allow OpticalFlow to access them), Fusion will fill the corresponding channels with zeros (transparent black).

The Disparity channel stores the displacement vectors that match pixels in one eye to the other eye. The left image's Disparity channel will contain vectors that map left > right and the right image's Disparity channel will contain vectors that map right > left.

For example:

```
(xleft, yleft) + (Dleft. x, Dleft. y) -> (xright, yright) (xright, yright) +
(Dright. x, Dright. y) -> (xleft, yleft)
```

You would expect for non-occluded pixels that Dleft = -Dright, although due to the disparity generation algorithm, this is only an approximate equality. Note that Disparity stores both X and Y values because rarely are left/right images perfectly registered in Y, even when taken through a carefully set up camera rig.

Disparity and Optical Flow values are stored as un-normalized pixel shifts. In particular, note that this breaks from Fusion's resolution-independent convention. After much consideration, this convention was chosen so the user wouldn't have to worry about rescaling the Disparity/Flow values when cropping an image or working out scale factors when importing/exporting these channels to other applications. Because the Flow and Disparity channels store things in pixel shifts, this can cause problems with Proxy and AutoProxy. The convention that Fusion follows is that, for proxied images, these channels store unscaled pixel shifts valid for the full-sized image. So if you wish to access the disparity values in a script or via a probe, you need to remember to always scale them by (image. Width/image. OriginalWidth, image. Height/image. OriginalHeight).

When using Vector and BackVector aux channels, remember that all nodes expect these aux channels to be filled with the flow between sequential frames.

More precisely, if you have sequence of three frames A, B, C, then:

B	Vector will contain the flow B>C
B	BackVector will contain the flow B>A
A	Vector will contain the flow A>B
A	BackVector is written with zeros as there is no frame before A
C	Vector is written with zeros as there is no frame D to flow C>D
C	BackVector will contain the flow C>B

When working with these channels, it is the user's responsibility to follow these rules (or for clever users to abandon them). Nodes like TimeStretcher will not function correctly since they still expect the channels to contain flow forward/back by 1 frame.

NOTE: Currently DoD/RoI is not supported for all Fusion nodes.



Fusion Page Effects

FUSION

CONTENTS

29 3D Nodes 701
30 3D Light Nodes 819
31 3D Material Nodes 838
32 3D Texture Nodes 867
33 Blur Nodes 892
34 Color Nodes 917
35 Composite Nodes 992
36 Deep Image Nodes 1012
37 Deep Pixel Nodes 1031
38 Effect Nodes 1047
39 Film Nodes 1084
40 Filter Nodes 1103
41 Flow Nodes 1120
42 Flow Organizational Nodes 1123
43 Fuses 1128
44 Generator Nodes 1130
45 I/O Nodes 1170
46 Layer Nodes 1197
47 LUT Nodes 1211
48 Mask Nodes 1220
49 Matte Nodes 1261
50 Metadata Nodes 1332
51 Miscellaneous Nodes 1340
52 Optical Flow 1379
53 Paint Node 1399
54 Particle Nodes 1409
55 Position Nodes 1469
56 Resolve Connect 1488
57 Shape Nodes 1494
58 Stereo Nodes 1534
59 Tracking Nodes 1566
60 Transform Nodes 1622
61 USD Nodes 1650
62 VR Nodes 1703
63 Warp Nodes 1714
64 Modifiers 1744

Chapter 29

3D Nodes

This chapter covers, in great detail, the nodes used for creating 3D composites.

The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Alembic Mesh 3D [Abc] 702
Bender 3D [3Bn] 705
Camera 3D [3Cm] 707
Cube 3D [3Cb] 717
Custom Vertex 3D [3CV] 719
Displace 3D [3Di] 725
Duplicate 3D [3Dp] 727
Extrude 3D [3Ex] 733
FBX Exporter 3D [FBX] 735
FBX Mesh 3D [FBX] 738
Fog 3D [3Fo] 740
Image Plane 3D [3Im] 743
Locator 3D [3Lo] 745
Merge 3D [3Mg] 748
Override 3D [3Ov] 750
Point Cloud 3D [3PC] 752
Projector 3D [3Pj] 756
Renderer 3D [3Rn] 761
Replace Material 3D [3RpI] 769
Replace Normals 3D [3RpN] 771
Replicate 3D [3Rep] 774
Ribbon 3D [3Ri] 779
Shape 3D [3Sh] 782
Soft Clip [3SC] 785
Spherical Camera [3SC] 787
Text 3D [3Txt] 790
Transform 3D [3Xf] 800
Triangulate 3D [3Tri] 803
UV Map 3D [3UV] 804
Weld 3D [3We] 807
Modifier 810
The Common Controls 811

Alembic Mesh 3D [Abc]



The Alembic Mesh 3D node

Alembic Mesh Node Introduction

At times, you may need to import 3D geometry from applications like Blender, Cinema4D, or Maya. One of the formats you can use for importing 3D geometry is the Alembic file format. This file type is a 3D scene interchange format that contains baked animation with its geometry. In other words, it eliminates the animation calculation times by embedding fixed, uneditable animation with 3D geometry. The animation is typically embedded using a point cache, which saves the dynamic data such as velocity after it has been calculated. Alembic objects can contain mesh geometry, cameras, points, UVs, normals, and baked animation.

You can import Alembic files (.abc) into Fusion in two ways:

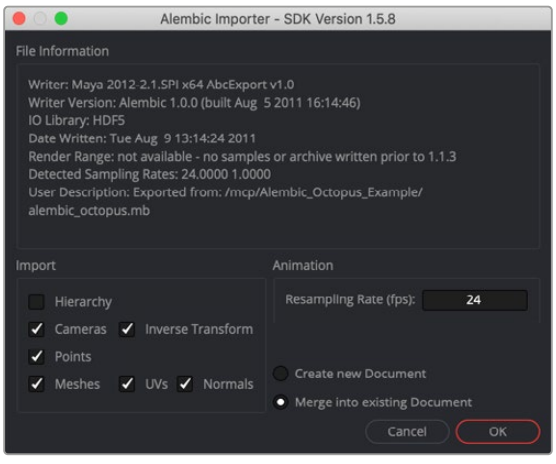
- Choose File > Import > Alembic Scene in Fusion or Fusion > Import > Alembic Scene in DaVinci Resolve's Fusion page.
- Add an AlembicMesh3D node to the Node Editor.

The first method is the preferred method; both Alembic and FBX nodes by themselves import the entire model as one object. However, the Import menu breaks down the model, lights, camera, and animation into a string of individual nodes. This makes it easy to edit and modify and use subsections of the imported Alembic mesh. Also, transforms in the file are read into Fusion splines and into the Transform 3D nodes, which get saved with the comp. Later, when reloading the comp, the transforms are loaded from the comp and not the Alembic file. Fusion handles the meshes differently, always reloading them from the Alembic file.

Arbitrary user data varies depending on the software creating the Alembic file, and therefore this type of metadata is mostly ignored.

Alembic Import Dialog

An Alembic Import dialog is displayed once you select the file to import



Alembic Import options

The top half of the Import dialog displays information about the selected file including the name of the plugin/application that created the Alembic file, the version of the Alembic software developer kit used during the export, the duration of the animation in seconds, if available, and the frame rate(s) in the file.

Various objects and attributes can be imported by selecting the checkboxes in the Import section.

- **Hierarchy:** When enabled, the full parenting hierarchy is recreated in Fusion using multiple Transform 3D nodes. When disabled, the transforms in the Alembic file are flattened down into the cameras and meshes. The flattening results in several meshes/cameras connected to a single Merge node in Fusion. It is best to have this disabled when the file includes animation. If enabled, the many rigs used to move objects in a scene will result in an equally large number of nodes in Fusion, so flattening will reduce the number of nodes in your node tree.
- **Orphaned Transforms:** When the hierarchy option is enabled, an Orphaned Transforms setting is displayed. Activating this Orphan Transforms setting imports transforms that parent a mesh or camera. For example, if you have a skeleton and associated mesh model, the model is imported as an Alembic mesh, and the skeleton as a node tree of Merge3Ds. If this is disabled, the Merge3Ds are not created.
- **Cameras:** When enabled, importing a file includes cameras along with Aperture, Angles of View, Plane of Focus, as well as Near and Far clipping plane settings. The resolution Gate Fit may be imported depending on whether the application used to export the file correctly tagged the resolution Gate Fit metadata. If your camera does not import successfully, check the setting for the Camera3D Resolution Gate Fit. Note that 3D Stereoscopic information is not imported.
- **InverseTransform:** Imports the Inverse Transform (World to Model) for cameras.
- **Points:** Alembic files support a Points type. This is a collection of 3D points with position information. Some 3D software exports particles as points. However, keep in mind that while position is included, the direction and orientation of the particles are lost.
- **Meshes:** This setting determines whether importing includes 3D models from the Alembic file. If it is enabled, options to include UVs and normals are displayed.

Animation

This section includes one option for the Resampling rate. When exporting an Alembic animation, it is saved to disk using frames per second (fps). When importing Alembic data into Fusion, the fps are detected and entered into the Resample Rate field unless you have changed it previously in the current comp. Ideally, you should maintain the exported frame rate as the resample rate, so your samples match up with the original. The Detected Sampling Rates information at the top of the dialog can give an idea of what to pick if you are unsure. However, using this field, you can change the frame rate to create effects like slow motion.

Not all objects and properties in a 3D scene have an agreed upon universal convention in the Alembic file format. That being the case, Lights, Materials, Curves, Multiple UVs, and Velocities are not currently supported when you import Alembic files.

Since the FBX file format does support materials and lights, we recommend the use of FBX for lights, cameras, and materials. Use Alembic for meshes only.

Inputs

The AlembicMesh3D node has two inputs in the Node Editor. Both are optional since the node is designed to use the imported mesh.

SceneInput: The orange input can be used to connect an additional 3D scene or model. The imported Alembic objects combine with the other 3D geometry.

MaterialInput: The optional green input is used to apply a material to the geometry by connecting a 2D bitmap image. It applies the connected image to the surface of the geometry in the scene.

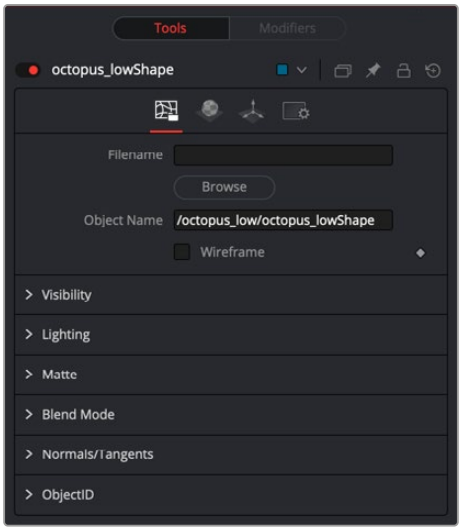
Basic Node Setup

The AlembicMesh3D node is designed to be part of a larger 3D scene. Typically, when imported, a 3D geometry model is represented by one node, and any transforms are in another node. The nodes imported as part of the Alembic file connect into a Merge 3D node along with a camera, lights, and other elements that may be required for the scene.



Alembic node structure

Inspector



Alembic mesh 3D controls

Controls Tab

The first tab in the Inspector is the Controls tab. It includes a series of unique controls specific to the Alembic Mesh 3D node as well as six groupings of controls that are common to most 3D nodes. The Common Controls section at the end of this chapter includes detailed descriptions of the common controls.

Below are descriptions of the Alembic Mesh 3D specific controls.

Filename

The complete file path of the imported Alembic file is displayed here. This field allows you to change or update the file linked to this node.

Object Name

This text field shows the name of the imported Alembic mesh, which is also used to rename the Alembic Mesh 3D node in the Node Editor.

When importing with the Alembic Mesh 3D node, if this text field is blank, the entire contents of the Alembic geometry are imported as a single mesh. When importing geometry using File > Import > Alembic Scene, this field is set by Fusion.

Wireframe

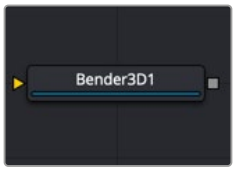
Enabling this option causes the mesh to display only the wireframe for the object in the viewer. When enabled, there is a second option for wireframe anti-aliasing. You can also render these wireframes out to a file if the Renderer 3D node has the OpenGL render type selected.

Common Controls

Controls, Materials, Transform, and Settings Tabs

The controls for Visibility, Lighting, Matte, Blend Mode, Normals/Tangents, and Object ID in the Controls tab are common in many 3D nodes. The Materials tab, Transforms tab and Settings tab in the Inspector are also duplicated in other 3D nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Bender 3D [3Bn]



The Bender 3D node

Bender 3D Introduction

The Bender 3D node is used to bend, taper, twist, or shear 3D geometry based on the geometry’s bounding box. It works by connecting any 3D scene or object to the orange input on the Bender 3D node, and then adjusting the controls in the Inspector. Only the geometry in the scene is modified. Any lights, cameras, or materials are passed through unaffected.

The Bender node does not produce new vertices in the geometry; it only alters existing vertices in the geometry. So, when applying the Bender 3D node to primitives, like the Shape 3D, or Text 3D nodes, increase the Subdivision setting in the primitive’s node to get a higher-quality result.

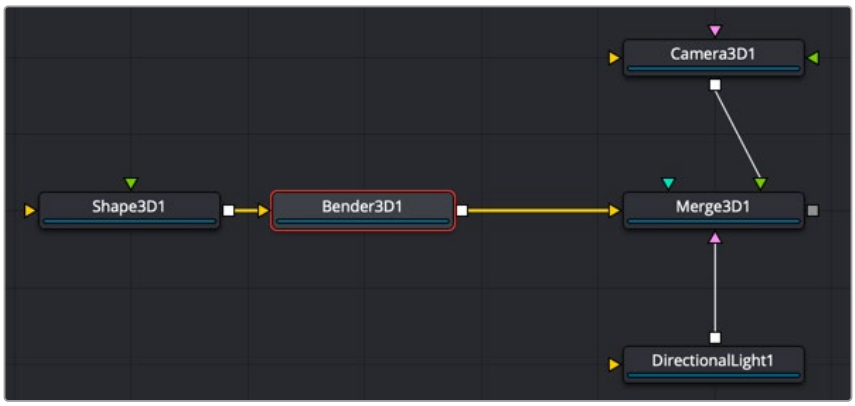
Inputs

The following inputs appear on the Bender 3D node in the Node Editor.

SceneInput: The orange scene input is the required input for the Bender 3D node. You use this input to connect another node that creates or contains a 3D scene or object.

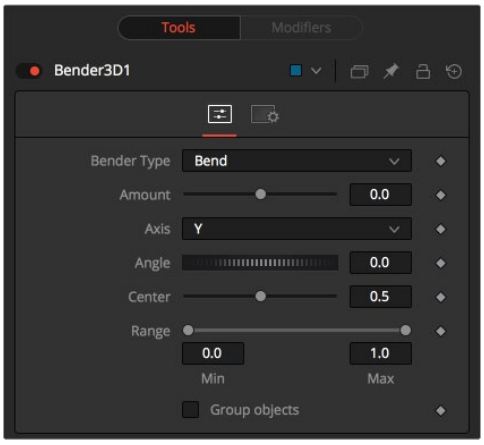
Basic Node Setup

The Bender 3D node works by connecting a 3D node that contains geometry, like an image plane 3D, Shape 3D or Text 3D. The element you connect to the Bender 3D node will be distorted based on the controls in the Inspector. The Bender 3D node is designed to be part of a larger 3D scene, with the output typically connecting into a Merge 3D.



Bender 3D node structure

Inspector



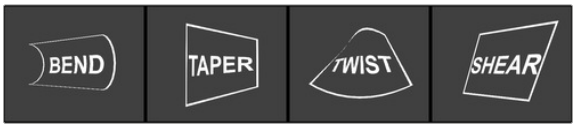
Bender 3D controls

Controls Tab

The first tab in the Inspector is the Controls tab. It includes all the controls for the Bender 3D node.

Bender Type

The Bender Type menu is used to select the type of deformation to apply to the geometry. There are four modes available: Bend, Taper, Twist, and Shear.



Amount

Adjusting the Amount slider changes the strength of the deformation.

Axis

The Axis control determines the axis along which the deformation is applied. It has a different meaning depending on the type of deformation. For example, when bending, this selects the elbow in conjunction with the Angle control. In other cases, the deform is applied around the specified axis.

Angle

The Angle thumbwheel control determines what direction about the axis a bend or shear is applied. It is not visible for taper or twist deformations.

Range

The Range control can be used to limit the effect of a deformation to a small portion of the geometry. The Range control is not available when the Bender Type is set to Shear.

Group Objects

If the input of the Bender 3D node contains multiple 3D objects, either through a Merge 3D or strung together, the Group Objects checkbox treats all the objects in the input scene as a single object, and the common center is used to deform the objects, instead of deforming each component object individually.

Common Controls

Settings

The Settings tab in the Inspector is common to all 3D nodes. This common tab is described in detail at the end of this chapter in “The Common Controls” section.

Camera 3D [3Cm]



The Camera 3D node

Camera 3D Node Introduction

The Camera 3D node generates a virtual camera for viewing the 3D environment. It closely emulates the settings used in real cameras to make matching live-action or 3D-rendered elements as seamless as possible. Adding any cameras to a 3D composite allows you to frame the elements in a composite how you want and animate the camera during a scene to create moving camera shots.

Camera Projection

The Camera 3D node can also be used to perform Camera Projection by projecting a 2D image through the camera into 3D space. Projecting a 2D image can be done as a simple Image Plane aligned with the camera, or as an actual projection, similar to the behavior of the Projector 3D node, with the added advantage of being aligned precisely with the camera. The Image Plane, Projection, and Materials tabs do not appear until you connect a 2D image to the magenta image input on the Camera 3D node in the Node Editor.

Stereoscopic

The Camera node has built-in stereoscopic features. They offer control over eye separation and convergence distance. The camera for the right eye can be replaced using a separate camera node connected to the green left/right stereo camera input. Additionally, the plane of focus control for depth of field rendering is also available here.

If you add a camera by dragging the camera icon from the toolbar onto the 3D view, it automatically connects to the Merge 3D you are viewing. Also, the current viewer is set to look through the new camera.

Alternatively, it is possible to copy the current viewer to a camera (or spotlight or any other object) by selecting the Copy PoV To option in the viewer's contextual menu, under the Camera submenu.

Inputs

There are three optional inputs on the Camera 3D node in the Node Editor.

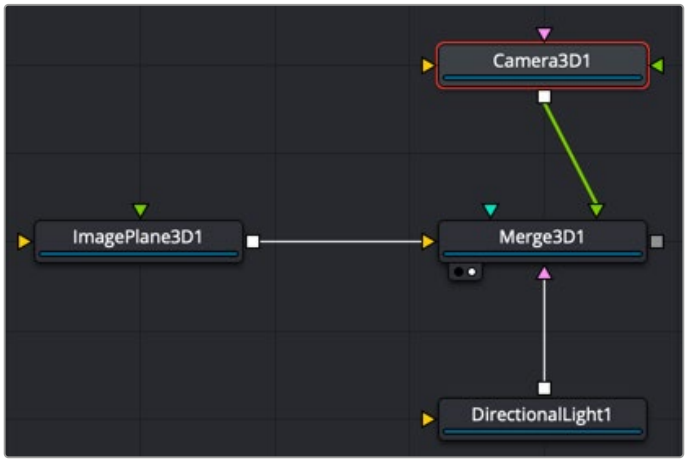
SceneInput: The orange input is used to connect a 3D scene or object. When connected, the geometry links to the camera's field of view. It acts similarly to an image attached to the Image Plane input. If the camera's Projection tab has projection enabled, the image attached to the orange image input projects on to the geometry.

ImageInput: The optional magenta input is used to connect a 2D image. When camera projection is enabled, the image can be used as a texture. Alternatively, when the camera's image plane controls are used, the parented planar geometry is linked to the camera's field of view.

RightStereoCamera: The green input should be connected to another Camera 3D node when creating 3D stereoscopic effects. It is used to override the internal camera used for the right eye in stereoscopic renders and viewers.

Basic Node Setup

The output of a camera 3D node should be connected to a Merge 3D node. You then view the Merge 3D node and select the camera from the viewer's right-click menu or by right-clicking over the axis label in the viewer.

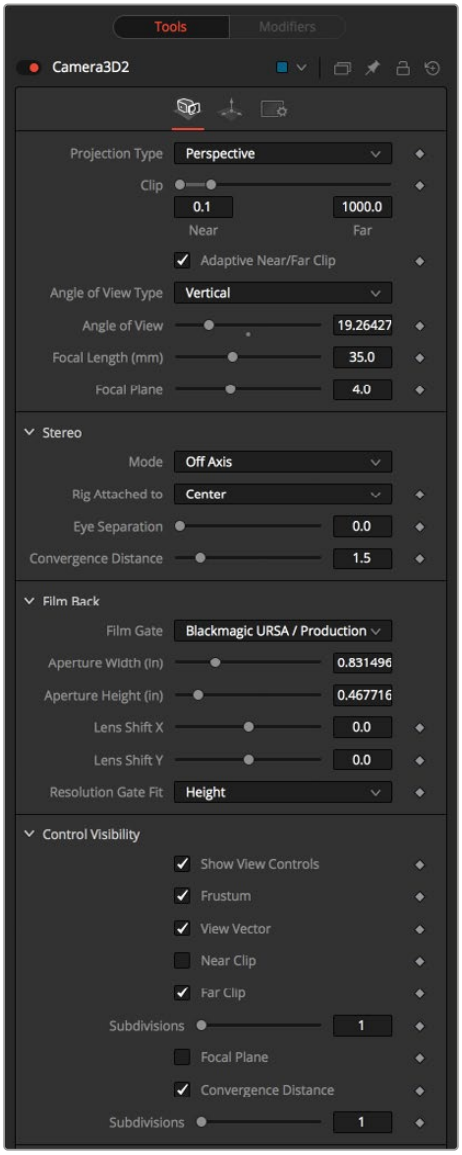


Camera node connected to and viewed through the Merge 3D

Displaying a camera node directly in the viewer shows only an empty scene; there is nothing for the camera to see. To view the scene through the camera, view the Merge 3D node where the camera is connected, or any node downstream of that Merge 3D. Then right-click on the viewer and select Camera > [Camera name] from the contextual menu. Right-clicking on the axis label found in the lower corner of each 3D viewer also displays the Camera submenu.

The aspect of the viewer may be different from the aspect of the camera, so the camera view may not match the actual boundaries of the image rendered by the Renderer 3D node. Guides can be enabled to represent the portion of the view that the camera sees and assist you in framing the shot. Right-click on the viewer and select an option from the Guides > Frame Aspect submenu. The default option uses the format enabled in the Composition > Frame Format preferences. To toggle the guides on or off, select Guides > Show Guides from the viewers' contextual menu, or use the Command-G (macOS) or Ctrl-G (Windows) keyboard shortcut when the viewer is active.

Inspector



Camera 3D controls

Controls Tab

The Camera3D Inspector includes six tabs along the top. The first tab, called the Controls tab, contains some of the most fundamental camera settings, including the camera's clipping plains, field of view, focal length, and stereoscopic properties. Some tabs are not displayed until a required connection is made to the Camera 3D node.

Projection Type

The Projection Type menu is used to select between Perspective and Orthographic cameras. Generally, real-world cameras are perspective cameras. An orthographic camera uses parallel orthographic projection, a technique where the view plane is perpendicular to the viewing direction. This produces a parallel camera output that is undistorted by perspective.

Orthographic cameras present controls only for the near and far clipping planes, and a control to set the viewing scale.

Near/Far Clip

The clipping planes are used to limit what geometry in a scene is rendered based on an object's distance from the camera's focal point. Clipping planes ensure objects that are extremely close to the camera, as well as objects that are too far away to be useful, are excluded from the final rendering.

The default perspective camera ignores this setting unless the Adaptive Near/Far Clip checkbox located under the Near/Far Clip control is disabled.

The clip values use units, so a far clipping plane of 20 means that any object more than 20 units from the camera is invisible to the camera. A near clipping plane of 0.1 means that any object closer than 0.1 units is also invisible.

> **NOTE:** A smaller range between the near and far clipping planes allows greater accuracy in all depth calculations. If a scene begins to render strange artifacts on distant objects, try increasing the distance for the Near Clip plane.

Adaptive Near/Far Clip

When selected, the renderer automatically adjusts the camera's near/far clipping plane to match the extents of the scene. This setting overrides the values of the Near and Far clip range controls described above. This option is not available for orthographic cameras.

Viewing Volume Size

When the Projection Type is set to Orthographic, the viewing volume size adjustment appears. It determines the size of the box that makes up the camera's field of view.

The Z-distance of an orthographic camera from the objects it sees does not affect the scale of those objects, only the viewing size does.

Angle of View Type

Use the Angle of View Type buttons to choose how the camera's angle of view is measured. Some applications use vertical measurements, some use horizontal, and others use diagonal measurements. Changing the Angle of View type causes the Angle of View control below to recalculate.

Angle of View

Angle of View defines the area of the scene that can be viewed through the camera. Generally, the human eye can see more of a scene than a camera, and various lenses record different degrees of the total image. A large value produces a wider angle of view, and a smaller value produces a narrower, or more tightly focused, angle of view.

Just as in a real-world camera, the angle of view and focal length controls are directly related. Smaller focal lengths produce a wider angle of view, so changing one control automatically changes the other to match.

Focal Length

In the real world, a lens' Focal Length is the distance from the center of the lens to the film plane. The shorter the focal length, the closer the focal plane is to the back of the lens. The focal length is measured in millimeters. The angle of view and focal length controls are directly related. Smaller focal lengths produce a wider angle of view, so changing one control automatically changes the other to match.

The relationship between focal length and angle of view is angle = 2 * arctan[aperture / 2 / focal_length].

Use the vertical aperture size to get the vertical angle of view and the horizontal aperture size to get the horizontal angle of view.

Plane of Focus (For Depth of Field)

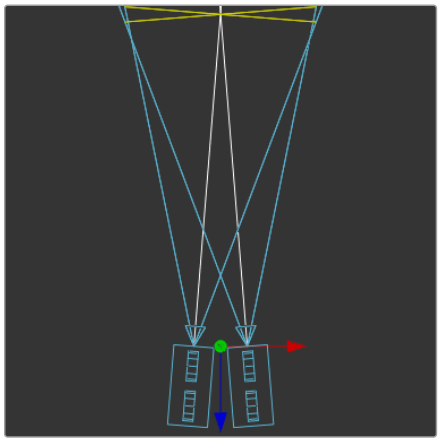
Like a focal point on a real-world camera, this setting defines the distance from the camera to an object. It is used by the OpenGL renderer in the Renderer 3D node to calculate depth of field.

Stereo

The Stereo section includes options for setting up 3D stereoscopic cameras. 3D stereoscopic composites work by capturing two slightly different views, displayed separately to the left and right eyes. The mode menu determines if the current camera is a stereoscopic setup or a mono camera. When set to the default mono setting, the camera views the scene as a traditional 2D film camera. Three other options in the mode menu determine the method used for 3D stereoscopic cameras.

Toe In

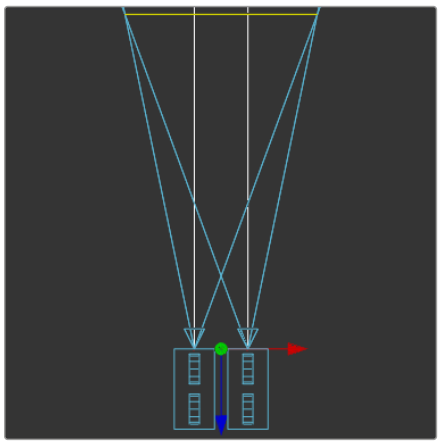
In a toe-in setup, both cameras are rotating in on a single focal point. Though the result is stereoscopic, the vertical parallax introduced by this method can cause discomfort by the audience. Toe-in stereoscopic works for convergence around the center of the images but exhibits keystoning, or image separation, to the left and right edges. This setup is can be used when the focus point and the convergence point need to be the same. It is also used in cases where it is the only way to match a live-action camera rig.



Toe In 3D camera setup

Off Axis

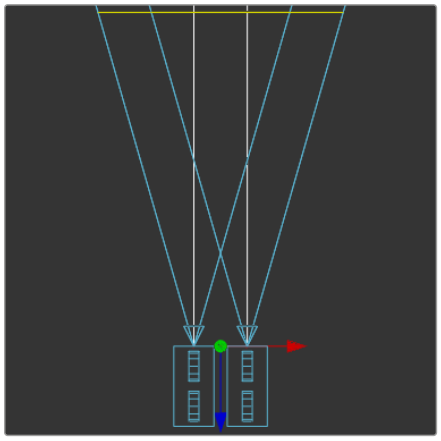
Regarded as the correct way to create stereo pairs, this is the default method in Fusion. Off Axis introduces no vertical parallax, thus creating stereo images with less eye strain. Sometimes called a skewed-frustum setup, this is akin to a lens shift in the real world. Instead of rotating the two cameras inward as in a toe-in setup, Off Axis shifts the lenses inward.



Off axis 3D camera setup

Parallel

The cameras are shifted parallel to each other. Since this is a purely parallel shift, there is no Convergence Distance control that limits your control over placing objects in front of or behind the screen. However, Parallel introduces no vertical parallax, thus creating less strain on the eyes.



Parallel 3D camera setup

Rig Attached To

This drop-down menu allows you to control which camera is used to transform the stereoscopic setup. Based on this menu, transform controls appear in the viewer either on the right camera, left camera, or between the two cameras. The ability to switch the transform controls through rigging can assist in matching the animation path to a camera crane or other live-action camera motion. The Center option places the transform controls between the two cameras and moves each evenly as the separation and convergence are adjusted. Left puts the transform controls on the left camera, and the right camera moves as the separation and convergence are adjusted. Right puts the transform controls on the right camera, and the left camera moves as adjustments are made to separation and convergence.

Eye Separation

Eye Separation defines the distance between both stereo cameras. Setting Eye Separation to a value larger than 0 shows controls for each camera in the viewer when this node is selected. Note that there is no Convergence Distance control in Parallel mode.

Convergence Distance

This control sets the stereoscopic convergence distance, defined as a point located along the Z-axis of the camera that determines where both left- and right-eye cameras converge.

The Convergence Distance controls are only available when setting the Mode menu to Toe-In or Off Axis.

Film Back

Film Gate

The size of the film gate represents the dimensions of the aperture. Instead of setting the aperture's width and height, you can choose it using the list of preset camera types in the Film Gate menu. Selecting one of the options automatically sets the aperture width and aperture height to match.

Aperture Width/Height

The Aperture Width and Height sliders control the dimensions of the camera's aperture or the portion of the camera that lets light in on a real-world camera. In video and film cameras, the aperture is the mask opening that defines the area of each frame exposed. The Aperture control uses inches as its unit of measurement.

Resolution Gate Fit

Determines how the film gate is fitted within the resolution gate. This only has an effect when the aspect of the film gate is not the same aspect as the output image.

> **NOTE:** This setting corresponds to Maya's Resolution Gate. The modes Overscan, Horizontal, Vertical, and Fill correspond to Inside, Width, Height, and Outside.

- **Inside:** The image source defined by the film gate is scaled uniformly until one of its dimensions (X or Y) fits the inside dimensions of the resolution gate mask. Depending on the relative dimensions of image source and mask background, either the image source's width or height may be cropped to fit the dimension of the mask.
- **Width:** The image source defined by the film gate is scaled uniformly until its width (X) fits the width of the resolution gate mask. Depending on the relative dimensions of image source and mask, the image source's Y-dimension might not fit the mask's Y-dimension, resulting in either cropping of the image source in Y or the image source not covering the mask's height entirely.
- **Height:** The image source defined by the film gate is scaled uniformly until its height (Y) fits the height of the resolution gate mask. Depending on the relative dimensions of image source and mask, the image source's X-dimension might not fit the mask's X-dimension, resulting in either cropping of the image source in X or the image source not covering the mask's width entirely.
- **Outside:** The image source defined by the film gate is scaled uniformly until one of its dimensions (X or Y) fits the outside dimensions of the resolution gate mask. Depending on the relative dimensions of image source and mask, either the image source's width or height may be cropped or not fit the dimension of the mask.
- **Stretch:** The image source defined by the film gate is stretched in X and Y to accommodate the full dimensions of the generated resolution gate mask. This might lead to visible distortions of the image source.

Control Visibility

This section allows you to selectively activate the onscreen controls that are displayed along with the camera.

- **Show View Controls:** Displays or hides all camera onscreen controls in the viewers.
- **Frustum:** Displays the actual viewing cone of the camera.
- **View Vector:** Displays a white line inside the viewing cone, which can be used to determine the shift when in Parallel mode.
- **Near Clip:** The Near clipping plane. This plane can be subdivided for better visibility.
- **Far Clip:** The Far clipping plane. This plane can be subdivided for better visibility.
- **Focal Plane:** The plane based on the Plane of Focus slider explained in the Controls tab above. This plane can be subdivided for better visibility.
- **Convergence Distance:** The point of convergence when using Stereo mode. This plane can be subdivided for better visibility.

Import Camera

The Import Camera button displays a dialog to import a camera from another application.

It supports the following file types:

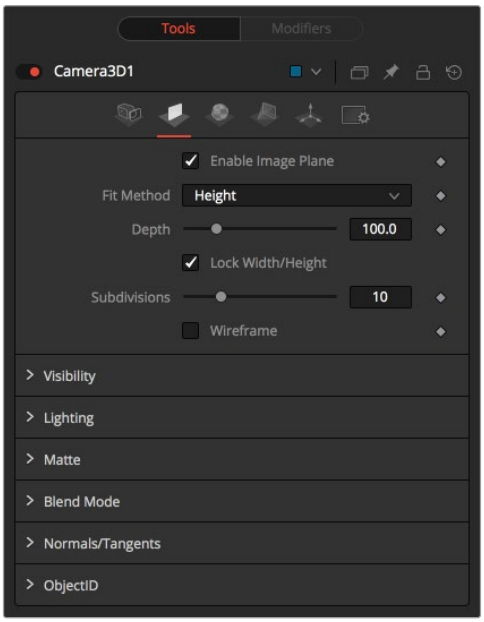
*LightWave Scene	`.lws`
*Max Scene	`.ase`
*Maya Ascii Scene	`.ma`
*dotXSI	`.xsi`

NOTE: FBX cameras can be imported using DaVinci Resolve's Fusion > Import > FBX Scene menu or File > Import > FBX Scene in Fusion Studio.

Image Tab

When a 2D image is connected to the magenta image input on the Camera3D node, an Image tab is created at the top of the inspector. The connected image is always oriented so it fills the camera's field of view.

Except for the controls listed below, the options in this tab are identical to those commonly found in other 3D nodes. For more detail on visibility, lighting, matte, blend mode, normals/tangents, and Object ID, see "The Common Controls" section at the end of this chapter.



Camera 3D image plane tab

Enable Image Plane

Use this checkbox to enable or disable the usage of the Image Plane.

Fill Method

This menu configures how to scale the image plane if the camera has a different aspect ratio.

- **Inside:** The image plane is scaled uniformly until one of its dimensions (X or Y) fits the inside dimensions of the resolution gate mask. Depending on the relative dimensions of image source and mask background, either the image source's width or height may be cropped to fit the dimensions of the mask.
- **Width:** The image plane is scaled uniformly until its width (X) fits the width of the mask. Depending on the relative dimensions of image source and the resolution gate mask, the image source's Y-dimension might not fit the mask's Y-dimension, resulting in either cropping of the image source in Y or the image source not covering the mask's height entirely.
- **Height:** The image plane is scaled uniformly until its height (Y) fits the height of the mask. Depending on the relative dimensions of image source and the resolution gate mask, the image source's X-dimension might not fit the mask's X-dimension, resulting in either cropping of the image source in X or the image source not covering the mask's width entirely.
- **Outside:** The image plane is scaled uniformly until one of its dimensions (X or Y) fits the outside dimensions of the resolution gate mask. Depending on the relative dimensions of image source and mask, either the image source's width or height may be cropped or not fit the respective dimension of the mask.
- **Depth:** The Depth slider controls the image plane's distance from the camera.

NOTE: The Camera Z position has no effect on the image plane's distance from the camera.

Materials Tab

The options presented in the Materials tab are identical to those commonly found in other 3D nodes. For more detail on Diffuse, Specular, Transmittance, and Martial ID controls, see "The Common Controls" section at the end of this chapter.

Projection Tab

When a 2D image is connected to the camera node, a fourth projection tab is displayed at the top of the Inspector. Using this Projection tab, it is possible to project the image into the scene. A projection is different from an image plane in that the projection falls onto the geometry in the scene exactly as if there were a physical projector present in the scene. The image is projected as light, which means the Renderer 3D node must be set to enable lighting for the projection to be visible.

Enable Camera Projection

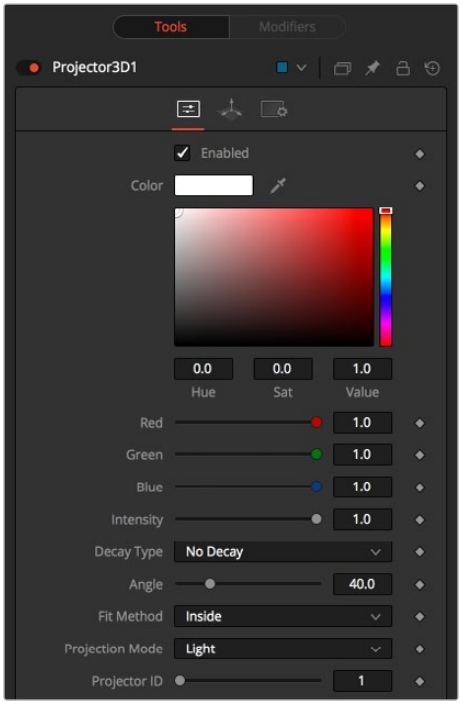
Select this checkbox to enable projection of the 2D image connected to the magenta input on the Camera node.

Projection Fit Method

This menu can be used to select the method used to match the aspect of the projected image to the camera's field of view.

Projection Mode

- **Light:** Defines the projection as a spotlight.
- **Ambient Light:** Defines the projection as an ambient light.
- **Texture:** Allows a projection that can be relighted using other lights. Using this setting requires a Catcher node connected to the applicable inputs of the specific material.



Camera 3D projection tab

Common Controls

Transform and Settings Tabs

The options presented in the Transform and Settings tabs are commonly found in other 3D nodes. For more detail on the controls found in these tabs, see “The Common Controls” section at the end of this chapter.

Tips for Camera 3D

Camera Projection: When importing a camera from a 3D application that is also used as a projector, make sure that the Fit Resolution Gate options on the Controls tab as well as the Projection tab are in sync. Only the first one automatically sets to what the 3D app was using. The latter might have to be adjusted manually.

Image Plane: The camera’s image plane isn‘t just a virtual guide for you in the viewers. It‘s actual geometry that you can also project on to. To use a different image on the image plane, you need to insert a Replace Material node after your Camera node.

Parallel Stereo: There are three ways you can achieve real Parallel Stereo mode:

- Connect an additional external (right) camera to the green Right Stereo Camera input of your camera.
- Create separate left and right cameras.
- When using Toe-In or Off Axis, set the Convergence Distance slider to a very large value of 999999999.

Rendering Overscan: If you want to render an image with overscan, you also must modify your scene‘s Camera3D. Since overscan settings aren’t exported along with camera data from 3D applications, this is also necessary for cameras you’ve imported via .fbx or .ma files. The solution is to increase the film back’s width and height by the factor necessary to account for extra pixels on each side.

Cube 3D [3Cb]



The Cube 3D node

Cube 3D Node Introduction

The Cube 3D node is a basic primitive geometry type capable of generating a simple cube.

The node also provides six additional image inputs that can be used to map a texture onto the six faces of the cube. Cubes are often used as shadow casting objects and for environment maps. For other basic primitives, see the Shape 3D node in this chapter.

Inputs

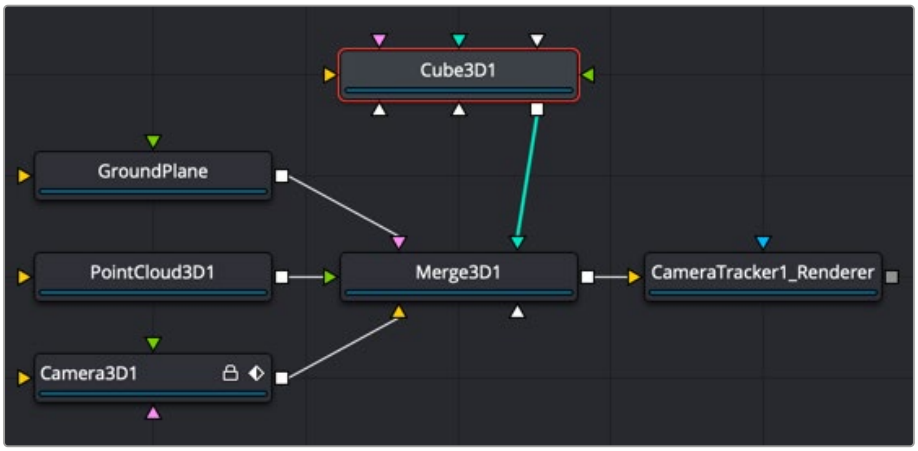
The following are optional inputs that appear on the Cube3D node in the Node Editor:

SceneInput: The orange scene input is used to connect another node that creates or contains a 3D scene or object. The additional geometry gets added to the Cube3D.

NameMaterialInput: These six inputs are used to define the materials applied to the six faces of the cube. You can connect either a 2D image or a 3D material to these inputs. Textures or materials added to the Cube3D do not get added to any 3D objects connected to the Cube's SceneInput.

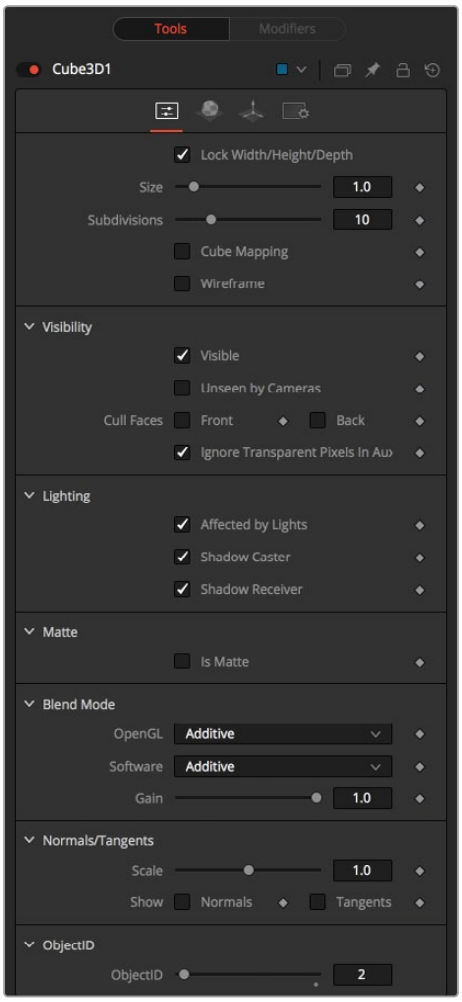
Basic Node Setup

The output of a Cube 3D node typically connects to a Merge 3D node, integrating it into a larger scene. When 3D tracking, the Cube 3D is often used as a placeholder for proper geometry that is not available at the current time.



Cube 3D node connected to a 3D scene exported from the Camera Tracker node

Inspector



Cube 3D controls

Controls Tab

The first tab in the Inspector is the Controls tab. It includes the primary controls for determining the overall size and shape of the Cube 3D node.

Lock Width/Height/Depth

This checkbox locks the Width, Height, and Depth dimensions of the cube together. When selected, only a Size control is displayed; otherwise, separate Width, Height, and Depth sliders are shown.

Size or Width/Height/Depth

If the Lock checkbox is selected, then only the Size slider is shown; otherwise, separate sliders are displayed for Width, Height, and Depth. The Size and Width sliders are the same control renamed, so any animation applied to Size is also applied to Width when the controls are unlocked.

Subdivision Level

Use the Subdivision Level slider to set the number of subdivisions used when creating the image plane.

The 3D viewers and renderer use vertex lighting, meaning all lighting is calculated at the vertices on the 3D geometry and then interpolated from there. Therefore, the more subdivisions in the mesh, the more vertices are available to represent the lighting. For example, make a sphere and set the subdivisions to be small so it looks chunky. With lighting on, the object looks like a sphere but has some amount of fracturing resulting from the large distance between vertices. When the subdivisions are high, the vertices are closer and the lighting becomes more even. So, increasing subdivisions can be useful when working interactively with lights.

Cube Mapping

Enabling the Cube Mapping checkbox causes the cube to wrap its first texture across all six faces using a standard cubic mapping technique. This approach expects a texture laid out in the shape of a cross.

Wireframe

Enabling this checkbox causes the mesh to render only the wireframe for the object when rendering with the OpenGL renderer in the Renderer 3D node.

Common Controls

Controls, Materials, Transform, and Settings Tabs

The remaining controls for Visibility, Lighting, Matte, Blend Mode, Normals/Tangents, and Object ID are common to many 3D nodes. The same is true of the Materials, Transform, and Settings tabs. Their descriptions can be found in “The Common Controls” section at the end of this chapter.

Custom Vertex 3D [3CV]



The Custom Vertex 3D node

Custom Vertex 3D Node Introduction

The Custom Vertex 3D node is an advanced custom node for 3D geometry that performs per vertex manipulations. If you have moderate experience with scripting or C++ programming, you should find the structure and terminology used by the Custom node familiar.

Using scripting math functions and lookup tables from images, you can move vertex positions on 3D geometry. Vertices can be more than just positions in 3D space. You can manipulate normals, texture coordinates, vectors, and velocity.

For example, Custom Vertex 3D can be used to make a flat plane wave like a flag, or create spiral models.

Besides providing a 3D scene input and three image inputs, the Inspector includes up to eight number fields and as many as eight XYZ position values from other controls and parameters in the node tree.

NOTE: Modifying the X, Y, and Z positions of a 3D object does not modify the normals/tangents. You can use a ReplaceNormals node afterward to recompute the normals/tangents.

TIP: Not all geometry has every attribute. For example, most Fusion geometry does not have vertex colors, with the exception of particles and some imported FBX/Alembic meshes. No geometry currently has environment coordinates, and only particles have velocities. If an attribute is not present on the input geometry, it is assumed to have a default value.

Inputs

The Custom Vertex 3D node includes four inputs. The orange scene input is the only one of the four that is required.

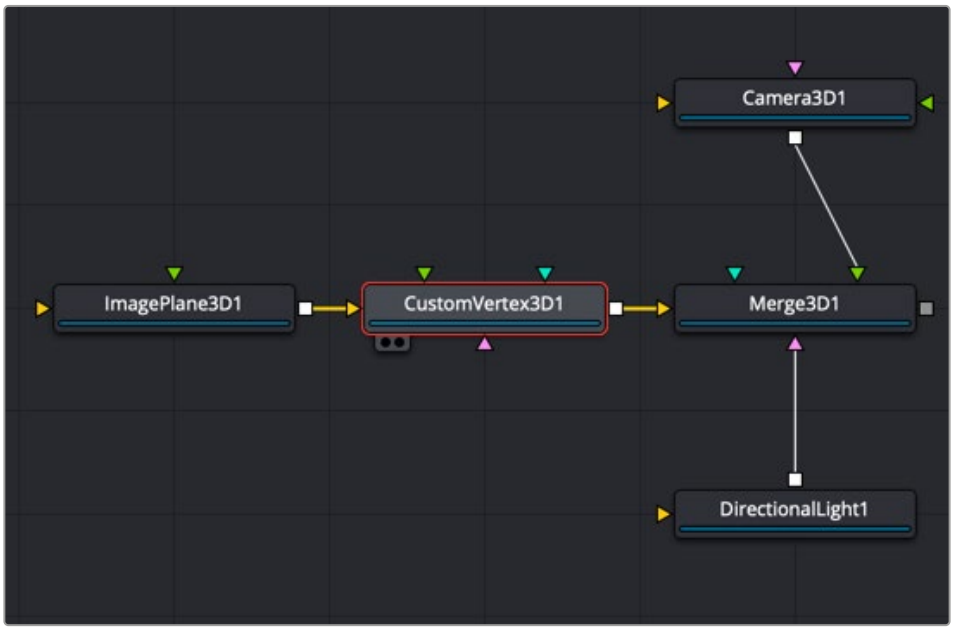
SceneInput: The orange scene input takes 3D geometry or a 3D scene from a 3D node output. This is the 3D scene or geometry that is manipulated by the calculations in the Custom Vertex 3D node.

ImageInput1, ImageInput2, ImageInput3: The three image inputs using green, magenta, and teal colors are optional inputs that can be used for compositing.

NOTE: Missing attributes on the input geometry are created if the expression for an attribute is nontrivial. The values for the attributes are given as in the above point. For example, if the input geometry does not have normals, then the values of (nx, ny, nz) is always (0,0,1). To change this, you could use a ReplaceNormals node beforehand to generate them.

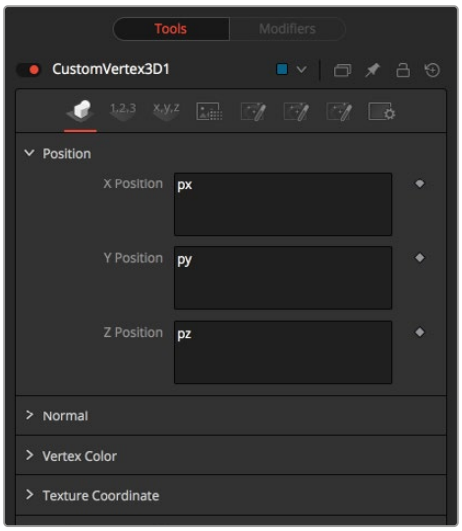
Basic Node Setup

The object you want to manipulate connects to the orange scene input of the Custom Vertex 3D node. The output typically connects to a Merge 3D node, integrating it into a larger scene.



Custom Vertex 3D node manipulating an Image Plane 3D node

Inspector



Custom Vertex 3D node Vertex tab

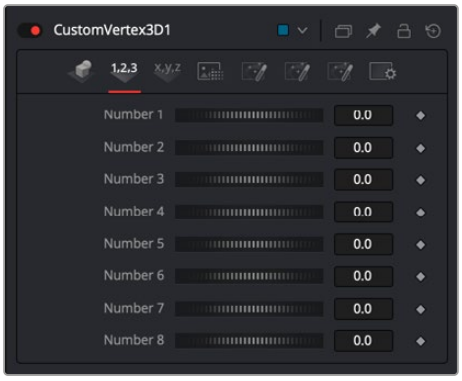
Vertex Tab

Using the fields in the Vertex tab, vertex calculations can be performed on the Position, Normals, Vertex Color, Texture Coordinates, Environment Coordinates, UV Tangents, and Velocity attributes.

The vertices are defined by three XYZ Position values in world space as px, py, pz. Normals, which define as a vector the direction the vertex is pointing as nx, ny, nz.

Vertex color is the Red, Green, Blue, and Alpha color of the point as vcr, vcg, vcb, vca.

Numbers Tab



Custom Vertex 3D node Numbers tab

Numbers 1-8

Numbers are variables with a dial control that can be animated or connected to modifiers exactly as any other control might. The numbers can be used in equations on vertices at current time: n1, n2, n3, n4,... or at any time: n1_at(float t), n2_at(float t), n3_at(float t), n4_at(float t), where t is the time you want. The values of these controls are available to expressions in the Setup and Intermediate tabs. They can be renamed and hidden from the viewer using the Config tab.

Points Tab

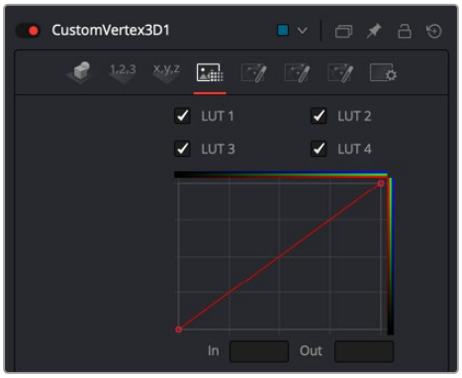


Custom Vertex 3D node Points tab

Points 1-8

The point controls represent points in the Custom Vertex 3D tool, not the vertices. These eight point controls include 3D X,Y,Z position controls for positioning points at the current time: (p1x, p1y, p1z, p2x, p2y, p2z) or at any time: p1x_at(float t), p1y_at(float t), p1z_at(float t), p2x_at(float t), p2y_at(float t), p2z_at(float t), where t is the time you want. For example, you can use a point to define a position in 3D space to rotate the vertices around. They can be renamed and hidden from the viewer using the Config tab. They are normal positional controls and can be animated or connected to modifiers as any other node might.

LUT Tab



Custom Vertex 3D node LUT tab

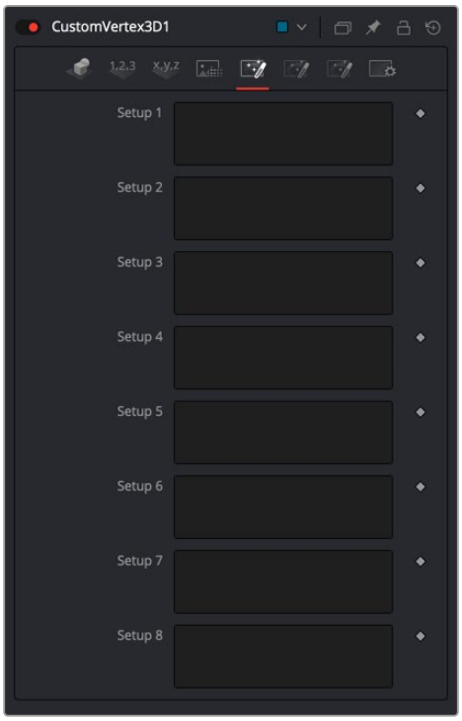
LUTs 1-4

The Custom Vertex 3D node provides four LUT splines. A LUT is a lookup table that will return a value from the height of the LUT spline. For example, getlut1(float x), getlut2(float x),... where x = 0 ... 1 accesses the LUT values.

The values of these controls are available to expressions in the Setup and Intermediate tabs using the getlut# function. For example, setting the R, G, B, and A expressions to getlut1(r1), getlut2(g1), getlut3(b1), and getlut4(a1) respectively, would cause the Custom Vertex 3D node to mimic the Color Curves node.

These controls can be renamed using the options in the Config tab to make their meanings more apparent, but expressions still see the values as lut1, lut2,...lut8.

Setup Tab



Custom Vertex 3D node Setup tab

Setups 1-8

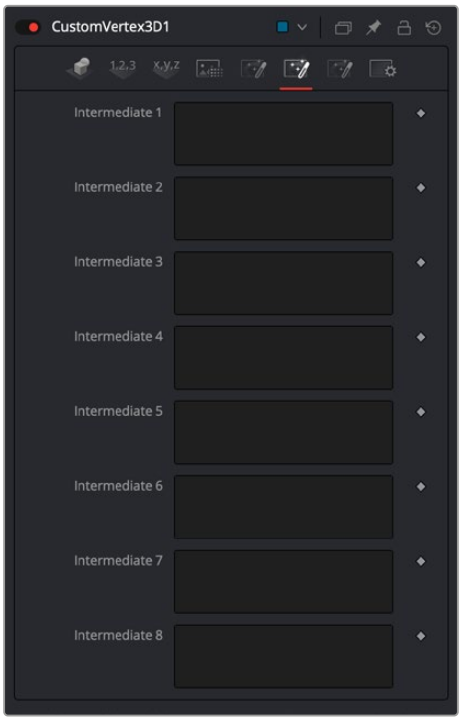
Up to eight separate expressions can be calculated in the Setup tab of the Custom Vertex 3D node. The Setup expressions are evaluated once per frame, before any other calculations are performed. The results are then made available to the other expressions in the node as variables s1, s2, s3, and s4.

Think of them as global setup scripts that can be referenced by the intermediate and channel scripts for each vertex.

For example, Setup scripts can be used to transform vertex from model to world space.

NOTE: Because these expressions are evaluated once per frame only and not for each pixel, it makes no sense to use per-pixel variables like X and Y or channel variables like r1, g1, b1, and so on. Allowable values include constants, variables like n1...n8, time, W and H, and so on, and functions like sin() or getr1d().

Intermediate Tab



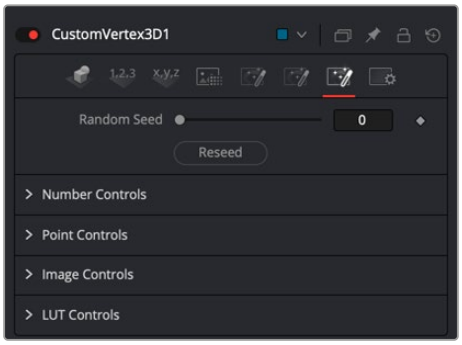
Custom Vertex 3D Node Intermediate tab

Intermediates 1-8

An additional eight expressions can be calculated in the Intermediate tab. The Intermediate expressions are evaluated once per vertex, after the Setup expressions are evaluated. Results are available as variables i1, i2, i3, i4, i5, i6, i7, i8, which can be referenced by channel scripts. Think of them as “per vertex setup” scripts.

For example, you can run the script to produce the new vertex (i.e., new position, normal, tangent, UVs, etc.) or transform from world space back to model space.

Config Tab



Custom Vertex 3D node Config tab

Random Seed

Use this to set the seed for the rand() and rands() functions. Click the Reseed button to set the seed to a random value. This control may be needed if multiple Custom Vertex 3D nodes are required with different random results for each.

Number Controls

There are eight sets of Number controls, corresponding to the eight sliders in the Numbers tab. Disable the Show Number checkbox to hide the corresponding Number slider, or edit the Name for Number text field to change its name.

Point Controls

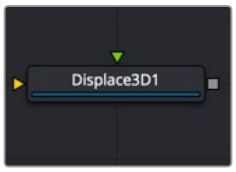
There are eight sets of Point controls, corresponding to the eight controls in the Points tab. Disable the Show Point checkbox to hide the corresponding Point control and its crosshair in the viewer. Similarly, edit the Name for Point text field to change the control's name.

Common Controls

Settings Tab

The Settings tab controls are common to many 3D nodes, and their descriptions can be found in "The Common Controls" section at the end of this chapter.

Displace 3D [3Di]



The Displace 3D node

Displace 3D Node Introduction

The Displace 3D node is used to displace the vertices of an object along their normals based on a reference image. The texture coordinates on the geometry are used to determine where to sample the image.

When using Displace 3D, keep in mind that it only displaces existing vertices and does not subdivide surfaces to increase detail. To obtain a more detailed displacement, increase the subdivision amount for the geometry that is being displaced. Note that the pixels in the displacement image may contain negative values.

TIP: Passing a particle system through a Displace 3D node disables the Always Face Camera option set in the pEmitter. Particles are not treated as point-like objects; each of the four particle vertices are individually displaced, which may or may not be the preferred outcome.

Inputs

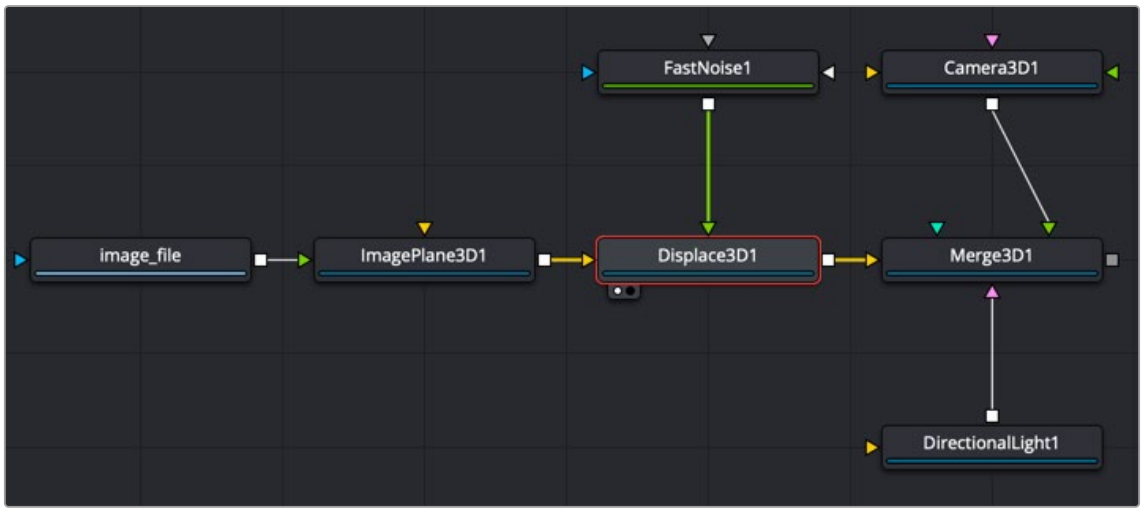
The following two inputs appear on the Displace 3D node in the Node Editor:

SceneInput: The orange scene input is the required input for the Displace 3D node. You use this input to connect another node that creates or contains a 3D scene or object.

Input: This green input is used to connect a 2D image that is used to displace the object connected to the Scene input. If no image is provided, this node effectively passes the scene straight through to its output. So, although not technically a required input, there isn't much use for adding this node unless you connect this input correctly.

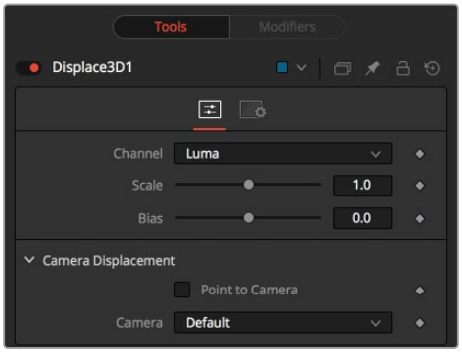
Basic Node Setup

The output of a Displace 3D node typically connects to a Merge 3D node, integrating it into a larger scene. The 3D geometry you want to displace is connected to the orange input, and in this example, a Fast Noise node is used to displace the geometry.



An image on an Image Plane 3D is displaced by a Fast Noise node

Inspector



Displace 3D controls

Controls Tab

The Displace 3D Inspector includes two tabs along the top. The primary tab, called the Controls tab, contains the dedicated Displace 3D controls.

Channel

Determines which channel of the connected input image is used to displace the geometry.

Scale and Bias

Use these sliders to scale (magnify) and bias (offset) the displacement. The bias is applied first and the scale afterward.

Camera Displacement

- **Point to Camera:** When the Point to Camera checkbox is enabled, each vertex is displaced toward the camera instead of along its normal. One possible use of this option is for displacing a camera's image plane. The displaced camera image plane would appear unchanged when viewed through the camera but is deformed in 3D space, allowing one to comp-in other 3D layers that correctly interact in Z.
- **Camera:** This menu is used to select which camera in the scene is used to determine the camera displacement when the Point to Camera option is selected.

Common Controls

Settings Tab

The Settings tab controls are common to many 3D nodes, and their descriptions can be found in "The Common Controls" section at the end of this chapter.

Duplicate 3D [3Dp]



The Duplicate 3D node

Duplicate 3D Node Introduction

Similar to the 2D version called the Duplicate node, the Duplicate 3D node can be used to duplicate any geometry in a scene, applying a successive transformation to each, and creating repeating patterns and complex arrays of objects. The options in the Jitter tab allow non-uniform transformations, such as random positioning or sizes.

Inputs

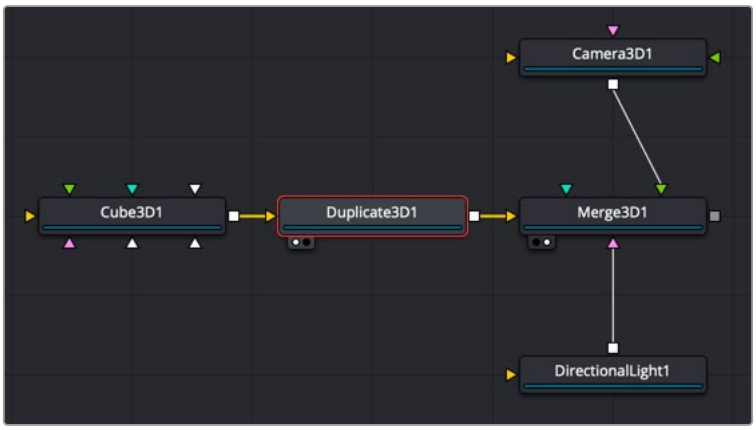
The Duplicate 3D node has a single input by default where you connect a 3D scene. An optional Mesh input appears based on the settings of the node.

SceneInput: The orange Scene Input is a required input. The scene or object you connect to this input is duplicated based on the settings in the Control tab of the Inspector.

MeshInput: A green optional mesh input appears when the Region's tab Region menu is set to mesh. The mesh can be any 3D model, either generated in Fusion or imported.

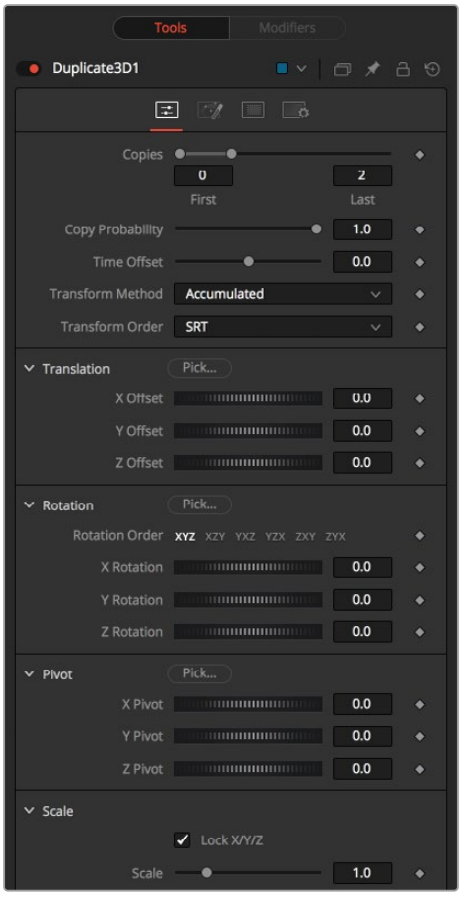
Basic Node Setup

The output of a Duplicate 3D node typically connects to a Merge 3D node, integrating it into a larger scene. The 3D geometry you want duplicated, in this case a Cube 3D, is connected to the orange input.



A Cube 3D is duplicated

Inspector



Duplicate 3D controls

Controls Tab

The Controls tab includes all the parameters you can use to create, offset, and scale copies of the object connected to the scene input on the node.

Copies

Use this range control to set the number of copies made. Each copy is a copy of the last copy, so if this control is set to [0,3], the parent is copied, then the copy is copied, then the copy of the copy is copied, and so on. This allows some interesting effects when transformations are applied to each copy using the controls below.

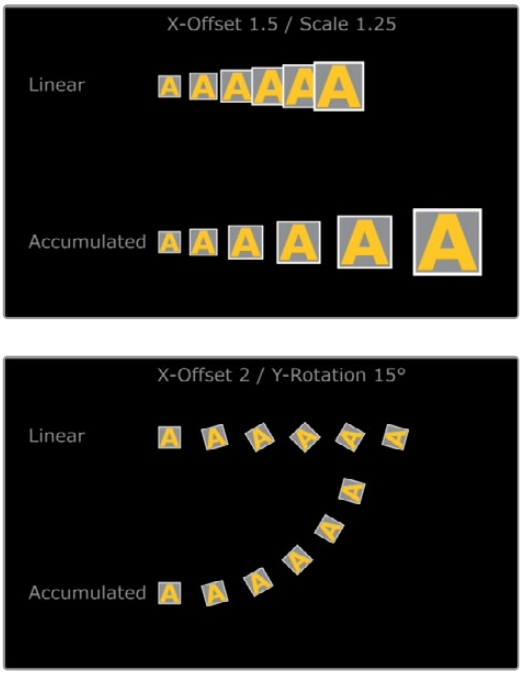
Setting the First Copy to a value greater than 0 excludes the original object and shows only the copies.

Time Offset

Use the Time Offset slider to offset any animations that are applied to the source geometry by a set amount per copy. For example, set the value to -1.0 and use a cube set to rotate on the Y-axis as the source. The first copy shows the animation from a frame earlier; the second copy shows animation from a frame before that, etc. This can be used with great effect on textured planes—for example, where successive frames of a clip can be shown.

Transform Method

- **Linear:** When set to Linear, transforms are multiplied by the number of the copy, and the total scale, rotation, and translation are applied in turn, independent of the other copies.
- **Accumulated:** When set to Accumulated, each object copy starts at the position of the previous object and is transformed from there. The result is transformed again for the next copy.



Transform Order

With this menu, the order in which the transforms are calculated can be set. It defaults to Scale-Rotation-Transform (SRT).

Using different orders results in different positions of your final objects.

Translation

The X, Y, and Z Offset sliders set the offset position applied to each copy. An X offset of 1 would offset each copy 1 unit along the X-axis from the last copy.

Rotation

The buttons along the top of this group of rotation controls set the order in which rotations are applied to the geometry. Setting the rotation order to XYZ would apply the rotation on the X-axis first, followed by the Y-axis rotation, then the Z-axis rotation.

The three Rotation sliders set the amount of rotation applied to each copy.

Pivot

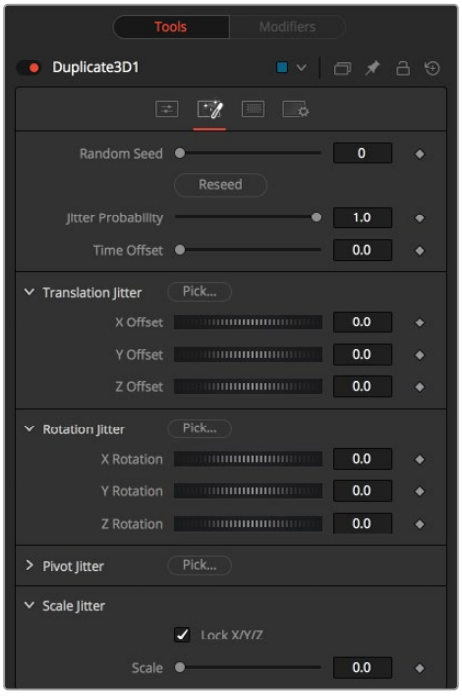
The pivot controls determine the position of the pivot point used when rotating each copy.

Scale

- **Lock:** When the Lock XYZ checkbox is selected, any adjustment to the duplicate scale is applied to all three axes simultaneously. If this checkbox is disabled, the Scale slider is replaced with individual sliders for the X, Y, and Z scales.
- **Scale:** The Scale controls tell Duplicate how much scaling to apply to each copy.

Jitter Tab

The options in the Jitter tab allow you to randomize the position, rotation, and size of all the copies created in the Controls tab.



Duplicate 3D Jitter tab

Random Seed

The Random Seed slider is used to generate a random starting point for the amount of jitter applied to the duplicated objects. Two Duplicate nodes with identical settings but different random seeds produce two completely different results.

Randomize

Click the Randomize button to auto generate a random seed value.

Jitter Probability

Adjusting this slider determines the percentage of copies that are affected by the jitter. A value of 1.0 means 100% of the copies are affected, while a value of 0.5 means 50% are affected.

Time Offset

Use the Time Offset slider to offset any animations that are applied to the source geometry by a set amount per copy. For example, set the value to –1.0 and use a cube set to rotate on the Y-axis as the source. The first copy shows the animation from a frame earlier; the second copy shows animation from a frame before that, etc. This can be used with great effect on textured planes—for example, where successive frames of a clip can be shown.

Translation Jitter

Use these three controls to adjust the amount of variation in the X, Y, and Z translation of the duplicated objects.

Rotation Jitter

Use these three controls to adjust the amount of variation in the X, Y, and Z rotation of the duplicated objects.

Pivot Jitter

Use these three controls to adjust the amount of variation in the rotational pivot center of the duplicated objects. This affects only the additional jitter rotation, not the rotation produced by the Rotation settings in the Controls tab.

Scale Jitter

Use this control to adjust the amount of variation in the scale of the duplicated objects. Disable the Lock XYZ checkbox to adjust the scale variation independently on all three axes.

Region Tab

The options in the Region tab allow you to define an area in the viewer where the copies can appear or are prevented from appearing. Like most parameters in Fusion, this area can be animated to cause the copied object to pop on and off the screen based on the region's shape and setting.

Region Tab

The Region section includes two settings for controlling the shape of the region and the affect the region has on the duplicate objects.

- **Region Mode:** There are three options in the Region Mode menu. The default, labeled "Ignore region" bypasses the node entirely and causes no change to the copies of objects from how they are set in the Controls and Jitter tabs. The menu option labeled "When inside region" causes the copied objects to appear only when their position falls inside the region defined in this tab. The last menu option, "When not Inside region" causes the copied objects to appear only when their position falls outside the region defined in this tab.
- **Region:** The Region menu determines the shape of the region. The five options include cube, sphere, and rectangle primitive shapes. The mesh option allows you to connect a 3D model into the green mesh input on the node. The green input appears only after the Region menu is set to Mesh.

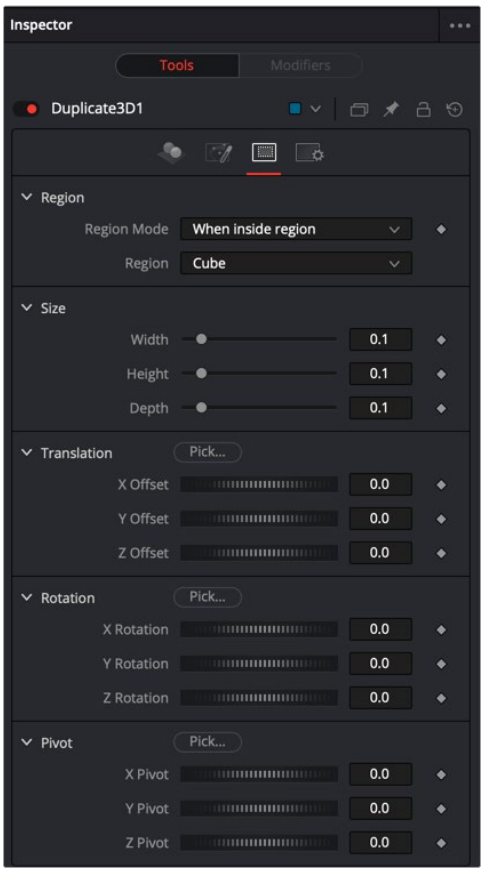
The All setting refers to the entire scene. This allows the copies to pop on and off if the Region mode is animated. When the Region menu is set to Mesh, four other options are displayed. These are described below.

Winding Rule: Using four common techniques, the Winding Rule menu determines how the mesh of polygons is determined as an area of volume and consequently how copies locate the vertices in the mesh. Complex overlapping regions of a mesh can cause an irregular fit. Trying a different technique from this menu can sometimes create a better match between the mesh and how the copies interpret the mesh shape.

Winding Ray Direction: A 3D model is a mesh of vertices made up of flat polygons. When making this a volume for a region, the Winding Ray Direction is used to determine in which direction the volume of each polygon (like depth extrude) is aligned.

Limit by Object ID: When a scene with multiple meshes is connected to the green Mesh input on the node, all the meshes are used as the region. Enabling this checkbox allows you to use the Object ID slider to select the ID for the mesh you want to use as the Region.

Object ID: When the Limit by Object ID checkbox is enabled, this slider selects the number ID for the mesh object you want to use for the Region.



Duplicate 3D Region tab

Common Controls

Settings Tab

The Settings tab controls are common to many 3D nodes, and their descriptions can be found in “The Common Controls” section at the end of this chapter.

Extrude 3D [3Ex]



The Extrude 3D node

Extrude 3D Node Introduction

The Extrude3D tool operates on Fusion's shape tools, taking flat shapes and transforming them into dynamic, three-dimensional objects through controlled extrusion, pushing the shape's edges into the Z-axis. Additional beveling can be applied to create polished edges and chamfers.

Inputs

There are three inputs to this node, one for inputing the shape itself, and two for materials.

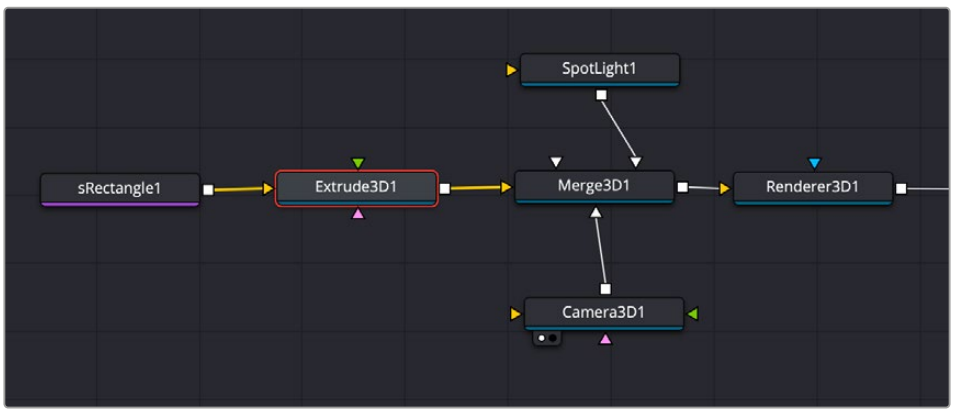
BevelMaterialInput: The pink-colored material input accepts either a 2D image or a 3D material. It provides the texture for the bevel based on the connected source, such as a Loader node in Fusion Studio or a MediaIn node in DaVinci Resolve. The 2D image is used as a diffuse texture map for the Basic Material tab in the Inspector. If a 3D material is connected, then the Basic Material tab is disabled.

ShapeInput: This yellow input expects a Shape node. This is the 2D shape that you want to bevel and extrude into 3D space.

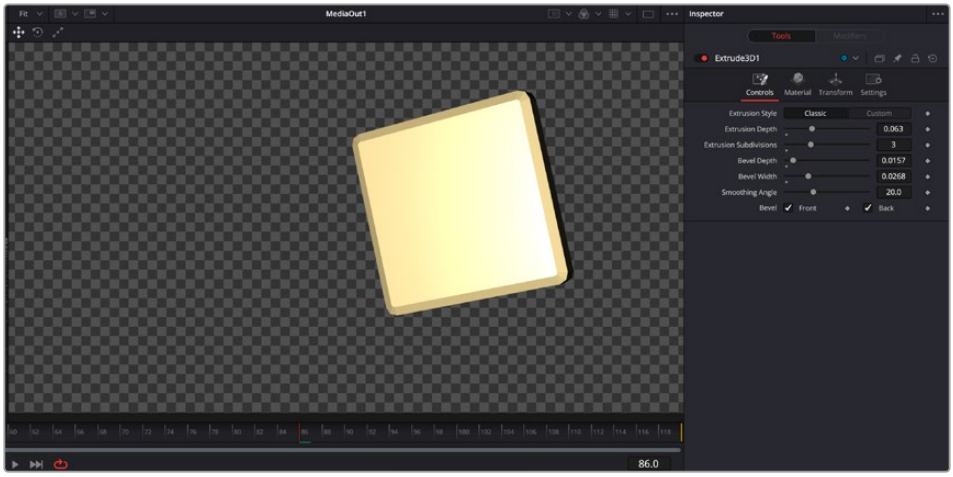
MaterialInput: The green-colored material input accepts either a 2D image or a 3D material. It provides the texture for the shape based on the connected source, such as a Loader node in Fusion Studio or a MediaIn node in DaVinci Resolve. The 2D image is used as a diffuse texture map for the Basic Material tab in the Inspector. If a 3D material is connected, then the Basic Material tab is disabled.

Basic Node Setup

The Extrude 3D node is primarily used to create a 3D shape from a 2D shape. The Shape Node node is connected to the Extrude 3D node, and the Extrude 3D node is then connected to a Merge 3D node. Viewing the Merge 3D node will show the Extrude 3D node and other elements connected to it.

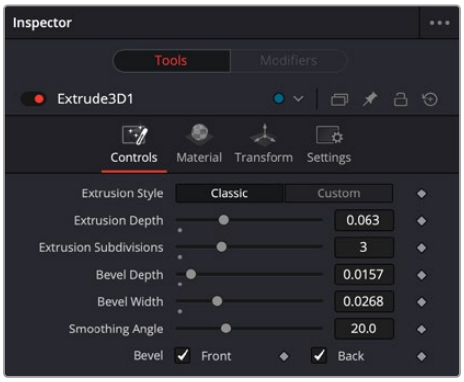


A simple sRectangle Shape node connected to an Extrude 3D node brings the rectangle into 3D space.



A simple sRectangle extruded into 3D space and beveled

Inspector



The Extrude3D Controls tab

Controls

The Controls tab is used to define the polygon characteristics, including extrusion depth and beveling.

Extrusion Style

There are two choices: Classic and Custom. Classic gives a standard uniform extrusion, while Custom exposes an Extrusion Profile graph that lets you add points and manipulate them to create unique custom extrusions, like picture frames and knurled buttons.

Extrusion Depth

This slider determines the extruded width of the polygon.

Extrusion Subdivisions

This slider determines the number of subdivisions within the smoothed portions of the extrusion profile.

Bevel Depth

Increase the value of the Bevel Depth slider above zero to add a bevel to the polygon.

Bevel Width

This slider determines the width of the polygon's bevel.

Smoothing Angle

Use this control to adjust the smoothing angle applied to the edges of the bevel.

Bevel Front/Back

Use these checkboxes to enable beveling for the front and back faces of the polygon separately.

Common Controls

Materials, Transform, and Settings Tabs

The Materials tab, Transforms tab, and Settings tab in the Inspector are also duplicated in other 3D nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

FBX Exporter 3D [FBX]



The FBX Exporter node

FBX Exporter Node Introduction

The FBX Exporter node provides a method of exporting a Fusion 3D scene to the FBX scene interchange format. Each node in Fusion is a single object in the exported file. Objects, lights, and cameras use the name of the node that created them. The node can be set to export a single file for the entire scene, or to output one frame per file.

Setting the Preferences > Global > General > Auto Clip Browse option in the Fusion Studio application, or the Fusion > Fusion Settings > General > Auto Clip Browse option in DaVinci Resolve to Enabled (default), and then adding this node to a composition automatically displays a file browser allowing you to choose where to save the file.

Once you have set up the node, the FBX Exporter is used similarly to a Saver node: clicking the Render button in the toolbar renders out the file.

Besides the FBX format, this node can also export to the 3D Studio's .3ds, Collada's .dae, Autocad's .dxf, and the Alias .obj formats.

Inputs

The FBX Exporter node has a single orange input.

Input: The output of the 3D scene that you want to export connects to the orange input on the FBX Exporter node.

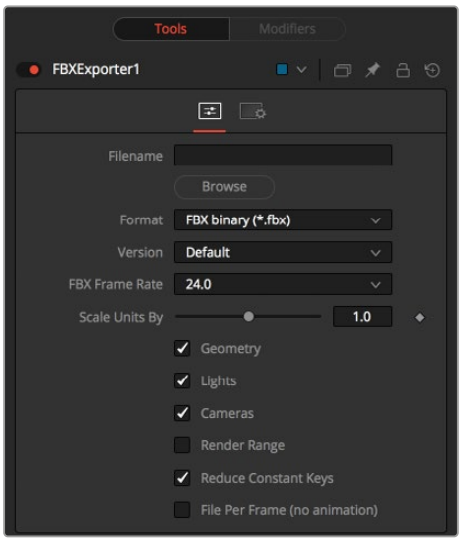
Basic Node Setup

The input to the FBX Exporter 3D node is any 3D scene you want to export. Below, the node is placed as a separate branch off of the Duplicate 3D node. Only the objects generated by the Duplicate 3D node are exported.



An FBX Exporter 3D branched off from the Duplicate 3D node

Inspector



FBX Exporter controls

Controls Tab

The Controls tab includes all the parameters you used to decide how the FBX file is created and what elements in the scene get exported.

Filename

This Filename field is used to display the location and file that is output by the node. You can click the Browse button to open a file browser dialog and change the location where the file is saved.

Format

This menu is used to set the format of the output file.

Not all features of this node are supported in all file formats. For example, the .obj format does not handle animation.

Version

The Version menu is used to select the available versions for the chosen format. The menu's contents change dynamically to reflect the available versions for that format. If the selected format provides only a single option, this menu is hidden.

Choosing Default for the FBX formats uses FBX2011.

Frame Rate

This menu sets the frame rate that is in the FBX scene.

Scale Units By

This slider changes the working units in the exported FBX file. Changing this can simplify workflows where the destination 3D software that you have uses a different scale.

Geometry/Lights/Cameras

These three checkboxes determine whether the node attempts to export the named scene element. For example, deselecting Geometry and Lights but leaving Cameras selected would output only the cameras currently in the scene.

Render Range

Enabling this checkbox saves the Render Range information in the export file, so other applications know the time range of the FBX scene.

Reduce Constant Keys

Enabling this option automatically removes keyframes if the adjacent keyframes have the same value.

File Per Frame (No Animation)

Enabling this option forces the node to export one file per frame, resulting in a sequence of numbered files. This disables the export of animation. Enable this checkbox to reveal the Sequence Start Frame control where you can set the first frame in the sequence to a custom value.

Sequence Start Frame

Enabling this checkbox displays a thumbwheel control to set a specific start frame for the number sequence applied to the rendered filenames. For example, if Global Start is set to 1 and frames 1–30 are rendered, files are normally numbered 0001–0030. If the Sequence Start frame is set to 100, the rendered output is numbered from 100–131.

Common Controls

Settings Tab

The Settings tab controls are common to many 3D nodes, and their descriptions can be found in "The Common Controls" section at the end of this chapter.

FBX Mesh 3D [FBX]



The FBX Mesh node

FBX Mesh 3D Node Introduction

The FBX Mesh 3D node is used to import polygonal geometry from scene files that are saved in the FilmBox (FBX) format. It is also able to import geometry from OBJ, 3DS, DAE, and DXF scene files. This provides a method for working with more complex geometry than is available using Fusion's built-in primitives.

When importing geometry with this node, all the geometry in the FBX file is combined into one mesh with a single pivot and transformation. The FBX Mesh node ignores any animation applied to the geometry.

Alternatively, in Fusion Studio, the File > Import > FBX Scene or in DaVinci Resolve, the Fusion > Import > FBX Scene menu can be used to import an FBX scene. This option creates individual nodes for each camera, light, and mesh in the file. This menu option can also be used to preserve the animation of the objects.

Setting the Preferences > Global > General > Auto Clip Browse option in Fusion Studio, or the Fusion > Fusion Settings > General > Auto Clip Browse option in DaVinci Resolve to Enabled (default), and then adding this node to a composition automatically displays a file browser allowing you to choose the file to import.

Inputs

SceneInput: The orange scene input is an optional connection if you wish to combine other 3D geometry nodes with the imported FBX file.

Material Input: The green input is the material input that accepts either a 2D image or a 3D material. If a 2D image is provided, it is used as a diffuse texture map for the basic material tab in the node. If a 3D material is connected, then the basic material tab is disabled.

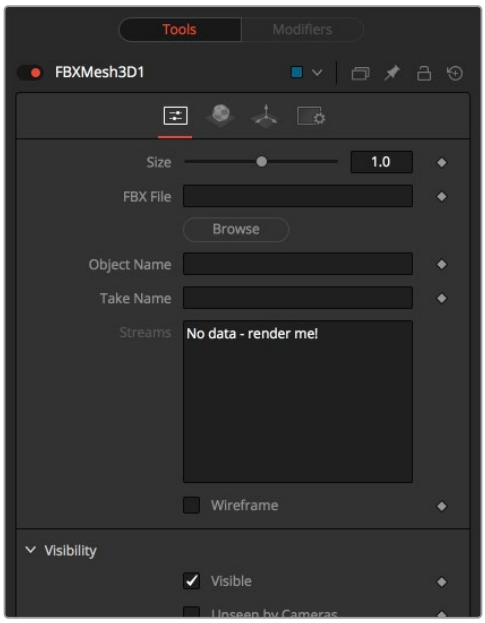
Basic Node Setup

The FBX Mesh 3D node can be used as a stand-alone node without any other nodes connected to it. The output is connected to a Merge 3D, integrating the FBX model into a larger scene. Below, the FBX Mesh 3D node also has a chrome material connected to its material input.



An FBX Mesh 3D node with chrome material applied

Inspector



FBX Mesh 3D controls

Controls Tab

Most of the Controls tab is taken up by common controls. The FBX-specific controls included on this tab are primarily information and not adjustments.

Size

The Size slider controls the size of the FBX geometry that is imported. FBX meshes have a tendency to be much larger than Fusion's default unit scale, so this control is useful for scaling the imported geometry to match the Fusion environment.

FBX File

This field displays the filename and file path of the currently loaded FBX mesh. Click the Browse button to open a file browser that can be used to locate a new FBX file. Despite the node's name, this node is also able to load a variety of other formats.

FBX ascii	(*.fbx)
FBX 5.0 binary	(*.fbx)
Autocad DXF	(*.dxf)
3D Studio 3Ds	(*.3ds)
Alias OBJ	(*.obj)
Collada DAE	(*.dae)

Object Name

This input shows the name of the mesh from the FBX file that is being imported. If this field is blank, then the contents of the FBX geometry are imported as a single mesh. You cannot edit this field; it is set by Fusion when using the File > Import > FBX Scene menu.

Take Name

FBX files can contain multiple instances of an animation, called Takes. This field shows the name of the animation take to use from the FBX file. If this field is blank, then no animation is imported. You cannot edit this field; it is set by Fusion when using the File > Import > FBX Scene menu.

Wireframe

Enabling this checkbox causes the mesh to render only the wireframe for the object. Only the OpenGL renderer in the Renderer 3D node supports wireframe rendering.

Common Controls

Controls, Materials, Transform, and Settings Tabs

The remaining controls for Visibility, Lighting, Matte, Blend Mode, Normals/Tangents, and Object ID are common to many 3D nodes. The same is true of the Materials, Transform, and Settings tabs. Their descriptions can be found in “The Common Controls” section at the end of this chapter.

Fog 3D [3Fo]



The Fog 3D node

Fog 3D Node Introduction

The Fog 3D node applies fog to the scene based on a depth cue. It is the 3D version of the Fog node in the Deep Pixel category. It is designed to work completely in 3D space and takes full advantage of anti-aliasing and depth of field effects during rendering.

The Fog 3D node essentially retextures the geometry in the scene by applying a color correction based on the object’s distance from the camera. An optional density texture image can be used to apply variation to the correction.

Inputs

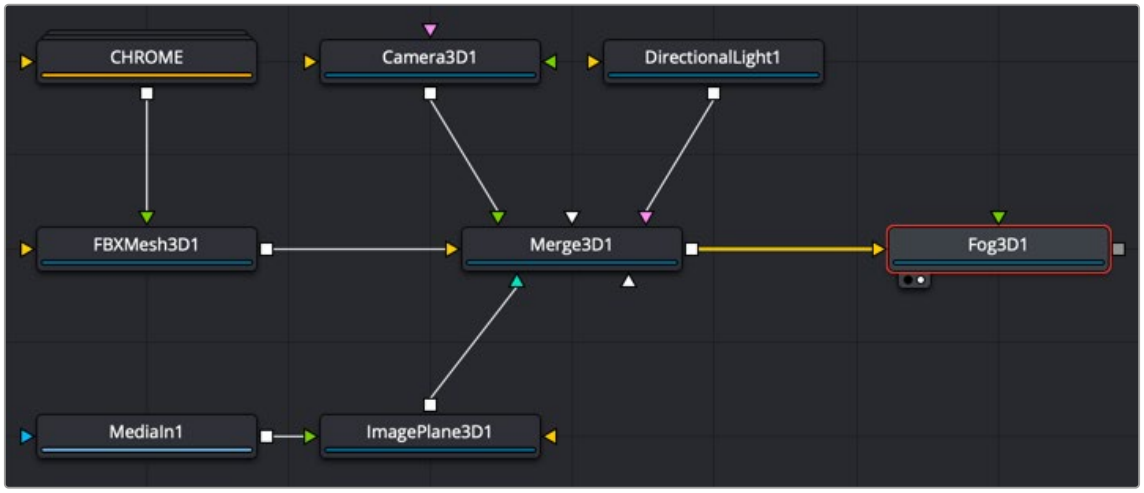
The Fog 3D node has two inputs in the Node Editor, only one of which is required for the Fog 3D to project onto a 3D scene.

SceneInput: The required orange-colored input accepts the output of a 3D scene on which the fog is “projected.”

DensityTexture: This optional green-colored input accepts a 2D image. The color of the fog created by this node is multiplied by the pixels in this image. When creating the image for the density texture, keep in mind that the texture is effectively projected onto the scene from the camera.

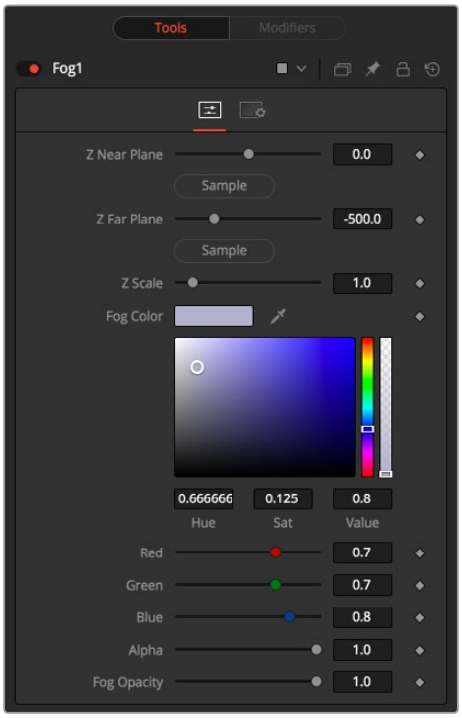
Basic Node Setup

The Fog 3D node is placed after the Merge 3D node that contains the scene. Viewing the Fog node will show the fog applied to the objects in the 3D scene based on their Z position.



A Fog 3D node placed after the Merge 3D scene

Inspector



Fog node controls

Controls Tab

The Controls tab includes all the parameters you use to decide how the Fog looks and projects onto the geometry in the scene.

Enable

Use this checkbox to enable or disable parts of the node from processing. This is not the same as the red switch in the upper-left corner of the inspector. The red switch disables the tool altogether and passes the image on without any modification. The Enable checkbox is limited to the effect part of the tool. Other parts like scripts in the Settings tab still processes as normal.

Show Fog in View

By default, the fog created by this node is visible only when the scene is viewed using a camera node. When this checkbox is enabled, the fog becomes visible in the scene from all points of view.

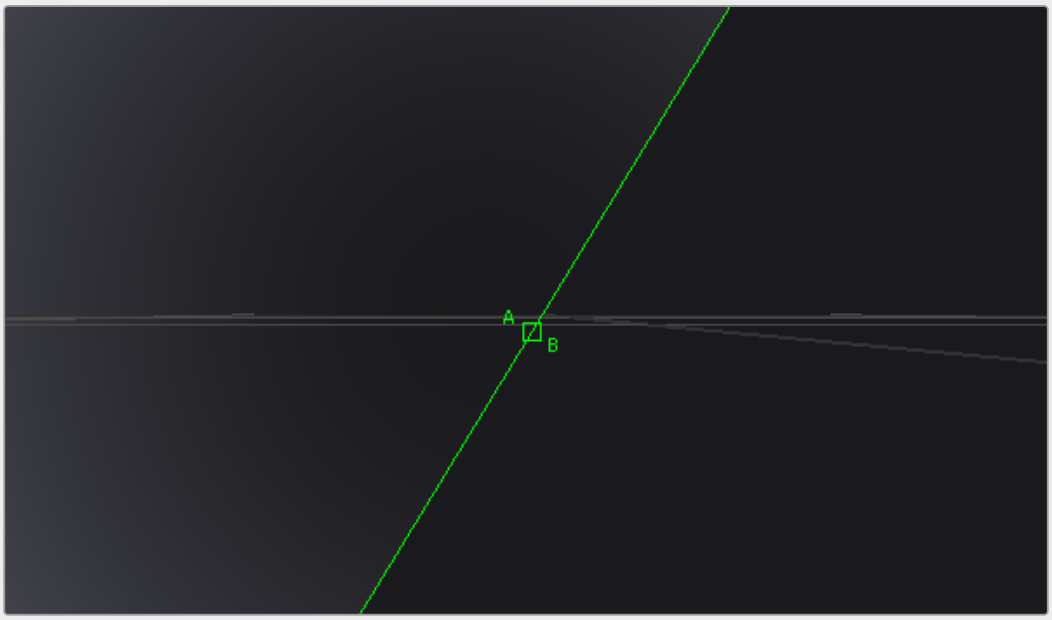
Color

This control can be used to set the color of the fog. The color is also multiplied by the density texture image, if one is connected to the green input on the node.

Radial

By default, the fog is created based on the perpendicular distance to a plane (parallel with the near plane) passing through the eye point. When the Radial option is checked, the radial distance to the eye point is used instead of the perpendicular distance. The problem with perpendicular distance fog is that when you move the camera about, as objects on the left or right side of the frustum move into the center, they become less fogged although they remain the same distance from the eye. Radial fog fixes this. Radial fog is not always desirable, however.

For example, if you are fogging an object close to the camera, like an image plane, the center of the image plane could be unfogged while the edges could be fully fogged.



Radial vs. Perpendicular fog effect

Type

This control is used to determine the type of falloff applied to the fog.

- **Linear:** Defines a linear falloff for the fog.
- **Exp:** Creates an exponential nonlinear falloff.
- **Exp2:** Creates a stronger exponential falloff.

Near/Far Fog Distance

This control expresses the range of the fog in the scene as units of distance from the camera. The Near Distance determines where the fog starts, while the Far Distance sets the point where the fog has its maximum effect. Fog is cumulative, so the farther an object is from the camera, the thicker the fog should appear.

Common Controls

Settings Tab

The Settings tab controls are common to many 3D nodes, and their descriptions can be found in “The Common Controls” section at the end of this chapter.

Image Plane 3D [3Im]



The Image Plane 3D node

Image Plane 3D Node Introduction

The Image Plane node produces 2D planar geometry in 3D space. The node is most commonly used to represent 2D images on “cards” in the 3D space. The aspect of the image plane is determined by the aspect of the image connected to the material input. If you do not want the aspect ratio of the image to modify the “card” geometry, then use a Shape 3D node instead.

Inputs

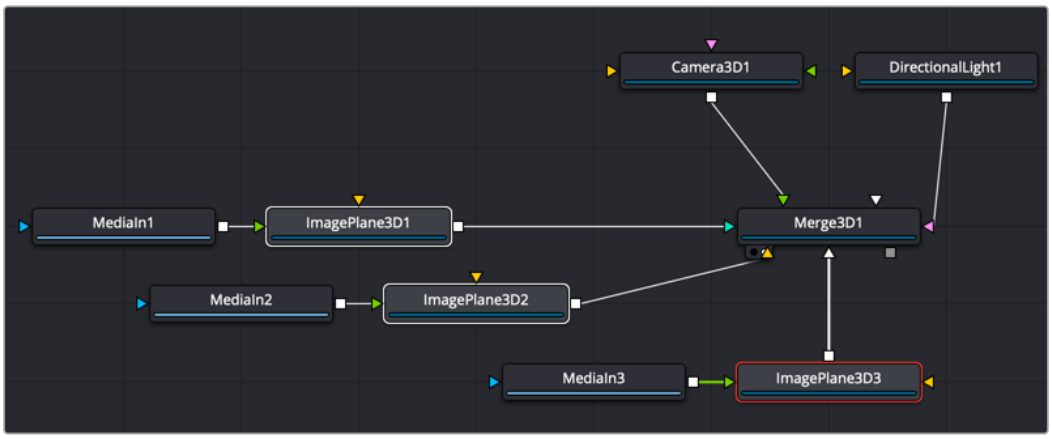
Of the two inputs on this node, the material input is the primary connection you use to add an image to the planar geometry created in this node.

SceneInput: This orange input expects a 3D scene. As this node creates flat, planar geometry, this input is not required.

MaterialInput: The green-colored material input accepts either a 2D image or a 3D material. It provides the texture and aspect ratio for the rectangle based on the connected source such as a Loader node in Fusion Studio or a MediaIn node in DaVinci Resolve. The 2D image is used as a diffuse texture map for the basic material tab in the Inspector. If a 3D material is connected, then the basic material tab is disabled.

Basic Node Setup

The Image Plane 3D node is primarily used to bring a video clip into a 3D composite. The MediaIn or Loader node is connected to the Image Plane 3D node, and the Image Plane 3D is then connected to a Merge 3D node. Viewing the Merge 3D node will show all the Image Plane 3D nodes and other elements connected to it.



Multiple Image Plane 3D nodes connected to a Merge 3D

Inspector

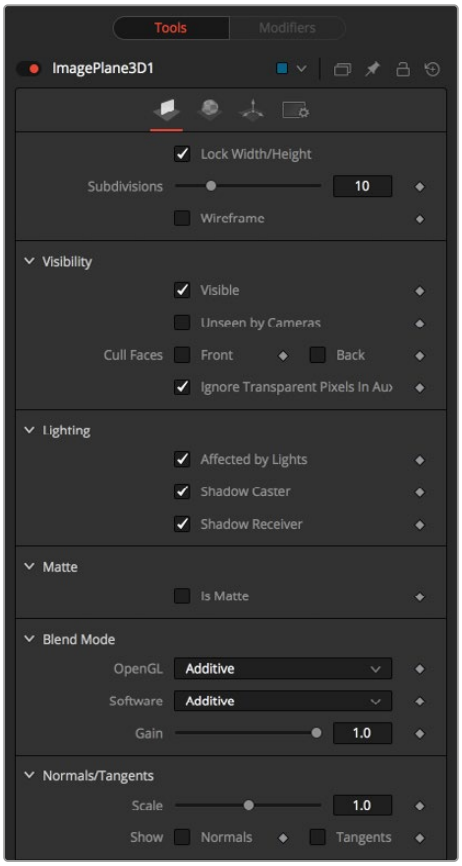


Image Plane 3D controls

Controls Tab

Most of the Controls tab is taken up by common controls. The Image Plane specific controls at the top of the Inspector allow minor adjustments.

Lock Width/Height

When checked, the subdivision of the plane is applied evenly in X and Y. When unchecked, there are two sliders for individual control of the subdivisions in X and Y. This defaults to on.

Subdivision Level

Use the Subdivision Level slider to set the number of subdivisions used when creating the image plane. If the Open GL viewer and renderer are set to Vertex lighting, the more subdivisions in the mesh, the more vertices are available to represent the lighting. So, high subdivisions can be useful when working interactively with lights.

Wireframe

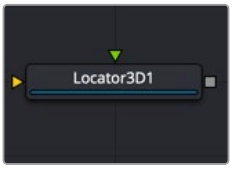
Enabling this checkbox causes the mesh to render only the wireframe for the object when using the OpenGL renderer.

Common Controls

Controls, Materials, Transform, and Settings Tabs

The remaining controls for Visibility, Lighting, Matte, Blend Mode, Normals/Tangents, and Object ID are common to many 3D nodes. The same is true of the Materials, Transform, and Settings tabs. Their descriptions can be found in “The Common Controls” section at the end of this chapter.

Locator 3D [3Lo]



The Locator 3D node

Locator 3D Node Introduction

The Locator 3D node’s purpose is to transform a point in 3D space to 2D coordinates that other nodes can use as part of expressions or modifiers.

When the Locator is provided with a camera and the dimensions of the output image, it transforms the coordinates of a 3D control into 2D screen space. The 2D position is exposed as a numeric output that can be connected to/from other nodes. For example, to connect the center of an ellipse to the 2D position of the Locator, right-click on the Mask center control and select Connect To > Locator 3D > Position.

Inputs

Two inputs accept 3D scenes as sources. The orange scene input is required, while the green Target input is optional.

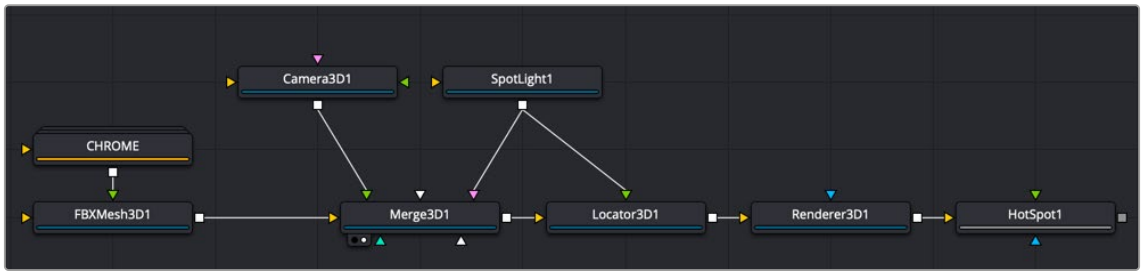
SceneInput: The required orange scene input accepts the output of a 3D scene. This scene should contain the object or point in 3D space that you want to covert to 2D coordinates.

Target: The optional green target input accepts the output of a 3D scene. When provided, the transform center of the scene is used to set the position of the Locator. The transformation controls for the Locator become offsets from this position.

Basic Node Setup

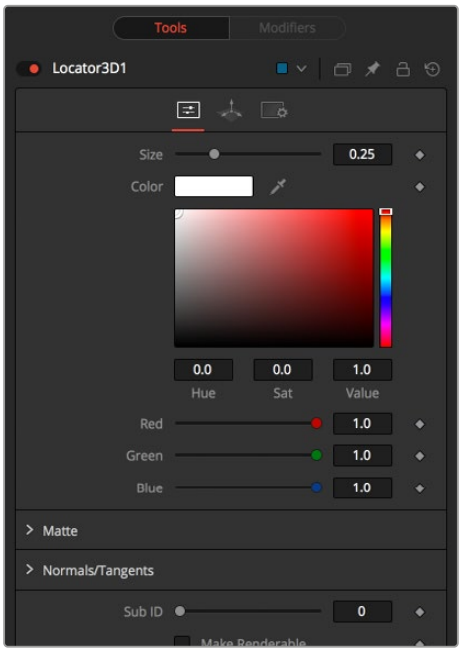
The scene provided to the Locator's input must contain the camera through which the coordinates are projected. So, the best practice is to place the Locator after the Merge that introduces the camera to the scene.

If an object is connected to the Locator node's target input, the Locator is positioned at the object's center, and the Transformation tab's translation XYZ sliders function in the object's local coordinate space instead of global scene space. This is useful for tracking an object's position despite any additional transformations applied further downstream.



Locator 3D connected after a Merge 3D with the SpotLight as the target

Inspector



Locator 3D controls

Controls Tab

Most of the controls for the locator 3D are cosmetic, dealing with how the locator appears and whether it is rendered in the final output. However, the Camera Settings are critical to getting the results you're looking for.

Size

The Size slider is used to set the size of the Locator's onscreen crosshair.

Color

A basic Color control is used to set the color of the Locator's onscreen crosshair.

Matte

Enabling the Is Matte option applies a special texture to this object, causing this object to not only become invisible to the camera, but also making everything that appears directly behind the camera invisible as well. This option overrides all textures.

- **Is Matte:** When activated, objects whose pixels fall behind the matte object's pixels in Z do not get rendered.
- **Opaque Alpha:** Sets the Alpha value of the matte object to 1. This checkbox is visible only when the Is Matte option is enabled.
- **Infinite Z:** Sets the value in the Z-channel to infinity. This checkbox is visible only when the Is Matte option is enabled.

For more information, *see Chapter 25, "3D Compositing Basics,"* in the Fusion Reference Manual.

Sub ID

The Sub ID slider can be used to select an individual subelement of certain geometry, such as an individual character produced by a Text 3D node or a specific copy created by a Duplicate 3D node.

Make Renderable

Defines whether the Locator is rendered as a visible object by the OpenGL renderer. The software renderer is not currently capable of rendering lines and hence ignores this option.

Unseen by Camera

This checkbox control appears when the Make Renderable option is selected. If the Unseen by Camera checkbox is selected, the Locator is visible in the viewers but not rendered into the output image by the Renderer 3D node.

Camera

This drop-down control is used to select the Camera in the scene that defines the screen space used for 3D to 2D coordinate transformation.

Use Frame Format Settings

Select this checkbox to override the width, height, and pixel aspect controls, and force them to use the values defined in the composition's Frame Format preferences instead.

Width, Height, and Pixel Aspect

In order for the Locator to generate a correct 2D transformation, it must know the dimensions and aspect of the image. These controls should be set to the same dimensions as the image produced by a renderer associated with the camera specified above. Right-clicking on these controls displays a contextual menu containing the frame formats configured in the composition's preferences.

Common Controls

Transform and Settings tabs

The remaining Transform and Settings tabs are common to many 3D nodes. Their descriptions can be found in "The Common Controls" section at the end of this chapter.

Merge 3D [3Mg]



The Merge 3D node

Merge 3D Introduction

The Merge 3D node is the primary node in Fusion that you use to combine separate 3D elements into the same 3D environment.

For example, in a scene created with an image plane, a camera, and a light, the camera would not be able to see the image plane and the light would not affect the image plane until all three objects are introduced into the same environment using the Merge 3D node.

The Merge provides the standard transformation controls found on most nodes in Fusion's 3D suite. Unlike those nodes, changes made to the translation, rotation, or scale of the Merge affect all the objects connected to the Merge. This behavior forms the basis for all parenting in Fusion's 3D environment.

Inputs

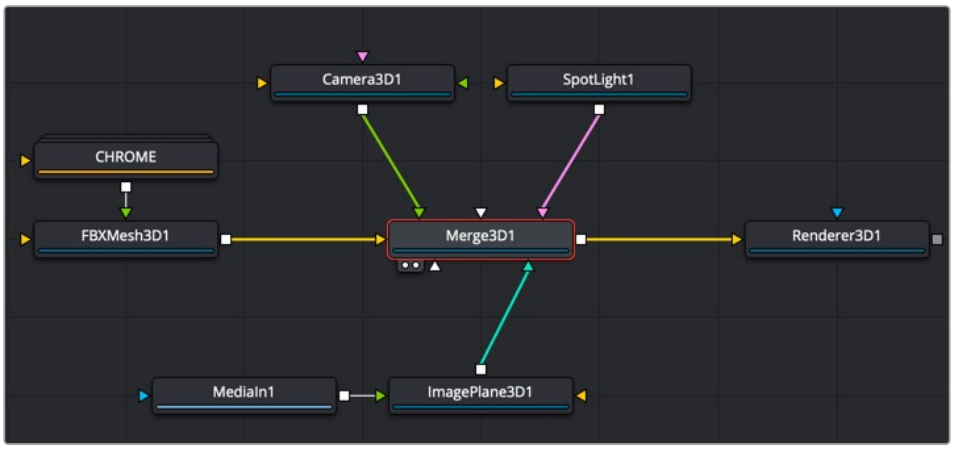
The Merge node displays only two inputs initially, but as each input is connected a new input appears on the node, assuring there is always one free to add a new element into the scene.

SceneInput[#]: These multicolored inputs are used to connect image planes, 3D cameras, lights, entire 3D scenes, as well as other Merge 3D nodes. There is no limit to the number of inputs this node can accept. The node dynamically adds more inputs as needed, ensuring that there is always at least one input available for connection.

Basic Node Setup

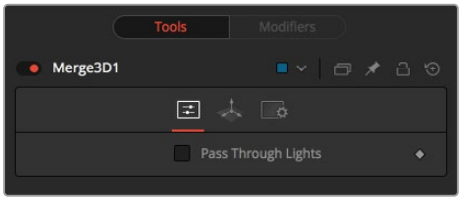
The Merge 3D is the hub of a 3D composite. All elements in a 3D scene connect into a Merge 3D.

Multiple Merge 3D nodes can be strung together to control lighting or for neater organization. The last Merge 3D in a string must connect to a Renderer 3D to be output as a 2D image.



Merge 3D with a connected Image Plane, FBX Mesh object, SpotLight, and camera

Inspector



Merge 3D controls

Controls Tab

The Controls tab is used only to pass through any lights connected to the Merge 3D node.

Pass Through Lights

When the Pass Through Lights checkbox is selected, lights are passed through the Merge into its output to affect downstream elements. Normally, the lights are not passed downstream to affect the rest of the scene. This is frequently used to ensure projections are not applied to geometry introduced later in the scene.

Common Controls

Transform and Settings Tabs

The remaining controls for the Transform and Settings tabs are common to most 3D nodes. Their descriptions can be found in “The Common Controls” section at the end of this chapter.

Override 3D [3Ov]



The Override 3D node

Override 3D Node Introduction

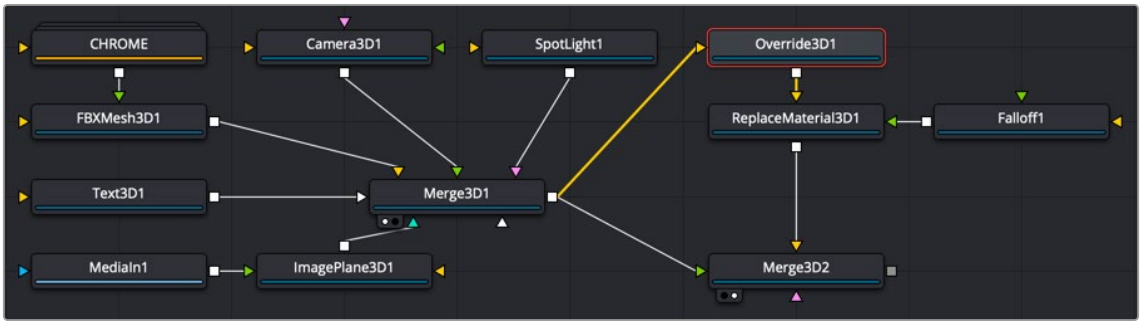
The Override node lets you change object-specific options for every object in a 3D scene simultaneously. This is useful, for example, when you wish to set every object in the input scene to render as a wireframe. Additionally, this node is the only way to set the wireframe, visibility, lighting, matte, and ID options for 3D particle systems and the Text 3D node.

Inputs

SceneInput: The orange Scene input accepts the output of a Merge 3D node or any node creating a 3D scene.

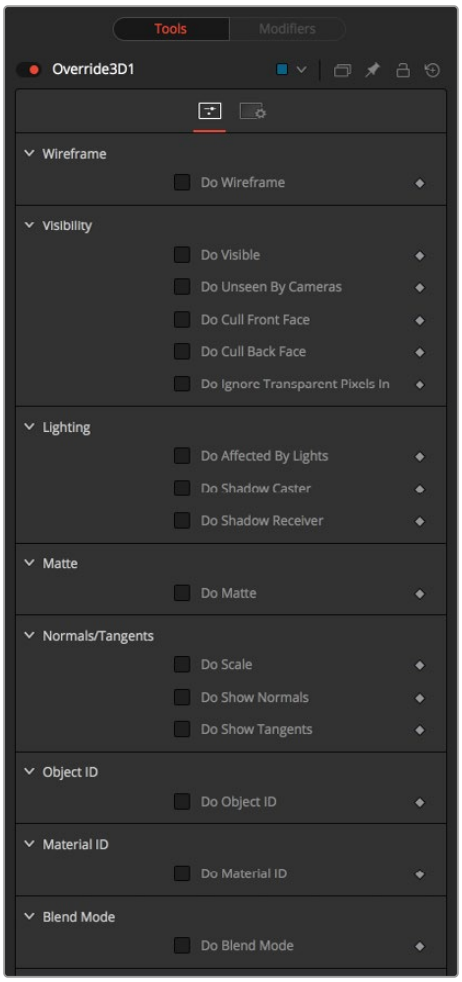
Basic Node Setup

The Override 3D node is frequently used in conjunction with the Replace Material node to produce isolated passes. For example, in the node tree below, a scene branches out to an Override node that turns off the Affected by Lights property of each node, then connects to a Replace Material node that applies a Falloff shader to produce a falloff pass of the scene.



Override 3D connected to a Replace Material node

Inspector



Override 3D controls

Controls Tab

The function of the controls found in the Controls tab is straightforward. First, you select the option to override using the Do [Option] checkbox. That reveals a control that can be used to set the value of the option itself. The individual options are not documented here; a full description of each can be found in any geometry creation node in this chapter, such as the Image Plane, Cube, or Shape nodes.

Do [Option]

Enables the override for this option.

[Option]

If the Do [Option] checkbox is enabled, then the control for the property itself becomes visible. The control values of the properties for all upstream objects are overridden by the new value.

Common Controls

Settings Tabs

The Settings tab includes controls common to most 3D nodes. Their descriptions can be found in “The Common Controls” section at the end of this chapter.

Point Cloud 3D [3PC]



The PointCloud 3D node

Point Cloud 3D Node Introduction

A Point Cloud is generally many null objects created by 3D tracking or modeling software.

When produced by 3D tracking software, the points typically represent each of the patterns tracked to create the 3D camera path. These point clouds can be used to identify a ground plane and to orient other 3D elements with the tracked image. The Point Cloud 3D node creates a point cloud either by importing a file from a 3D tracking application or generating it when you use the Camera Tracker node.

NOTE: A null object is an invisible 3D object that has all the same transform properties of a visible 3D object.

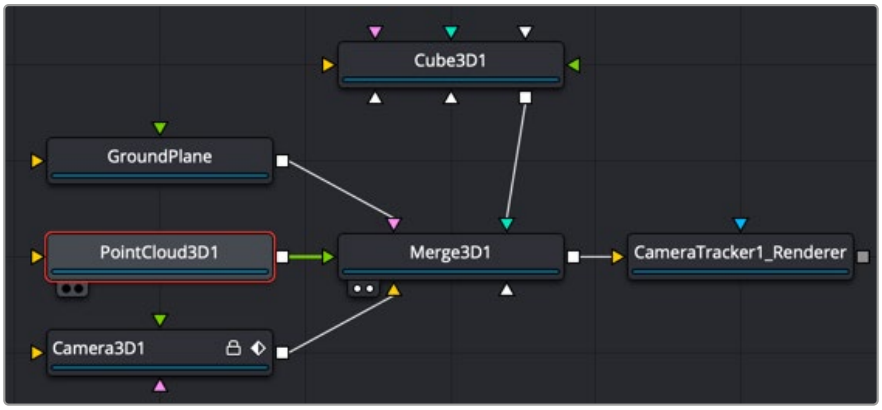
Inputs

The Point Cloud has only a single input for a 3D scene.

SceneInput: This orange input accepts a 3D scene.

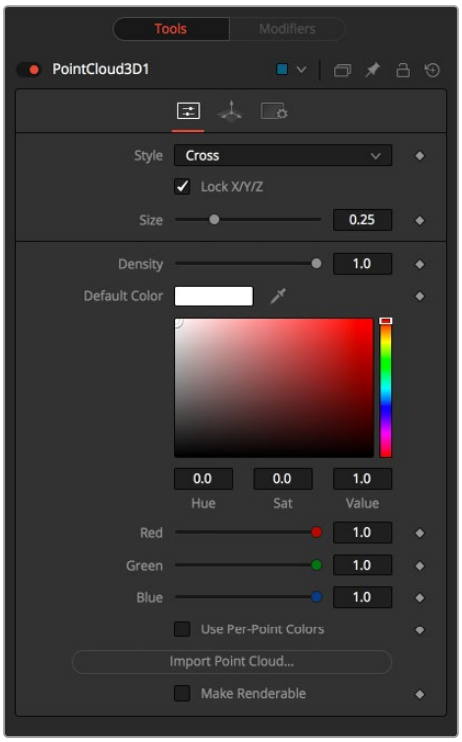
Basic Node Setup

The Point Cloud 3D node is viewed and connected through a Merge 3D node, integrating it into the larger 3D scene.



Point Cloud 3D connected and viewed through a Merge 3D

Inspector



Point Cloud 3D controls

Controls Tab

The Controls tab is where you can import the point cloud from a file and controls its appearance in the viewer.

Style

The Style menu allows you to display the point cloud as cross hairs or points in the viewer.

Lock X/Y/Z

Deselect this checkbox to provide individual control over the size of the X, Y, and Z arms of the points in the cloud.

Size X/Y/Z

These sliders can be used to increase the size of the onscreen crosshairs used to represent each point.

Density

This slider defines the probability of displaying a specific point. If the value is 1, then all points are displayed. A value of 0.2 shows only every fifth point.

Color

Use the standard Color control to set the color of onscreen crosshair controls.

Import Point Cloud

The Import Point Cloud button displays a dialog to import a point cloud from another application.

Supported file types are:

Alias's Maya	`.ma`
3DS Max ASCII Scene Export	`.ase`
NewTek's LightWave	`.lws`
Softimage XSI's	`.xsi.`

Make Renderable

Determines whether the point cloud is visible in the OpenGL viewer and in final renderings made by the OpenGL renderer. The software renderer does not currently support rendering of visible crosshairs for this node.

Unseen by Camera

This checkbox control appears when the Make Renderable option is selected. If the Unseen by Cameras checkbox is selected, the point cloud is visible in the viewers but not rendered into the output image by the Renderer 3D node.

Common Controls

Transform and Settings Tabs

The remaining Transform and Settings tabs are common to many 3D nodes. Their descriptions can be found in “The Common Controls” section at the end of this chapter.

Onscreen Contextual Menu

Frequently, one or more of the points in an imported point cloud is manually assigned to track the position of a specific feature. These points usually have names that distinguish them from the rest of the points in the cloud. To see the current name for a point, hover the mouse pointer directly over a point, and after a moment a small tooltip appears with the name of the point.

When the Point Cloud 3D node is selected, a submenu is added to the viewer’s contextual menu with several options that make it simple to locate, rename, and separate these points from the rest of the point cloud.

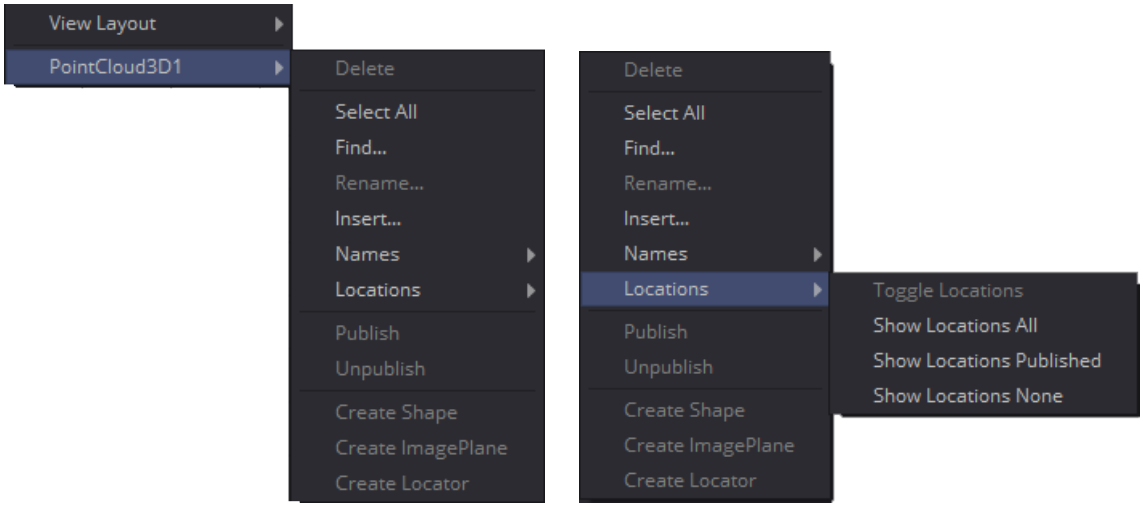
The contextual menu contains the following options:

Find: Selecting this option from the viewer contextual menu opens a dialog to search for and select a point by name. Each point that matches the pattern is selected.

Rename: Rename any point by selecting Rename from the contextual menu. Type the new name into the dialog that appears and press Return. The point now has that name, with a four-digit number added to the end. For example, the Name window is window0000, and multiple points would be window0000, window0001, and so on. Names must be valid Fusion identifiers (i.e., no spaces allowed, and the name cannot start with a number).

Delete: Selecting this option deletes the currently selected points.

Publish: Normally, the exact position of a point in the cloud is not exposed. To expose the position, select the points, and then select the Publish option from this contextual menu. This adds a coordinate control to the control panel for each published point that displays the point’s current location.



The Point Cloud 3D contextual menu options

Additional Toolbar and Shortcuts

Action	Shortcuts
Delete Selected Points	Del
Select All	Shift+A
Find Points	Shift+F
Rename Selected Points	F2
Create New Point	Shift+C
Toggle Names on None/Selected/Published/All Points	Shift+N
Toggle Locations on None/Selected/Published/All Points	Shift+L
Publish Selected Points	Shift+P
Unpublish Selected Points	Shift+U
Create a Shape at Selected Points	Shift+S
Create and Fit an ImagePlane to Selected Points	Shift+I
Create a Locator at Selected Points	Shift+O

Projector 3D [3Pj]



The Projector 3D node

Projector 3D Node Introduction

The Projector 3D node is used to project an image upon 3D geometry. This can be useful in many ways: texturing objects with multiple layers, applying a texture across multiple separate objects, projecting background shots from the camera's viewpoint, image-based rendering techniques, and more. The Projector node is just one of several nodes capable of projecting images and textures. Each method has advantages and disadvantages.

> For more information, *see Chapter 25, "3D Compositing Basics,"* in the Fusion Reference Manual.

Projected textures can be allowed to "slide" across the object if the object moves relative to the Projector 3D, or, alternatively, by grouping the two with a Merge 3D so they can be moved as one and the texture remains locked to the object.

The Projector 3D node's capabilities and restrictions are best understood if the Projector is considered to be a variant on the SpotLight node. The fact that the Projector 3D node is actually a light has several important consequences when used in Light or Ambient Light projection mode:

- Lighting must be turned on for the results of the projection to be visible.
- The light emitted from the projector is treated as diffuse/specular light. This means that it is affected by the surface normals and can cause specular highlights. If this is undesirable, set the Projector 3D to project into the Ambient Light channel.
- Enabling Shadows causes Projector 3D to cast shadows.
- Just as with other lights, the light emitted by a Projector 3D only affects objects that feed into the first Merge 3D that is downstream of the Projector 3D node in the node tree.
- Enabling Merge 3D's Pass Through Lights checkbox allows the projection to light objects further downstream.
- The light emitted by a Projector 3D is controlled by the Lighting options settings on objects and the Receives Lighting options on materials.
- Alpha values in the projected image do not clip geometry in Light or Ambient Light mode. Use Texture mode instead.
- If two projections overlap, their light contributions are added.

To project re-lightable textures or textures for non-diffuse color channels (like Specular Intensity or Bump), use the Texture projection mode instead:

- Projections in Texture mode only strike objects that use the output of the Catcher node for all or part of the material applied to that object.
- Texture mode projections clip the geometry according to the Alpha channel of the projected image.

See the section for the Catcher node for additional details.

Camera Projection vs. Projection 3D Node

The Camera 3D node also provides a projection feature, and should be used when the projection is meant to match a camera, as this node has more control over aperture, film back, and clip planes. The Projector 3D node was designed to be used as a custom light in 3D scenes for layering and texturing. The projector provides better control over light intensity, color, decay, and shadows.

Inputs

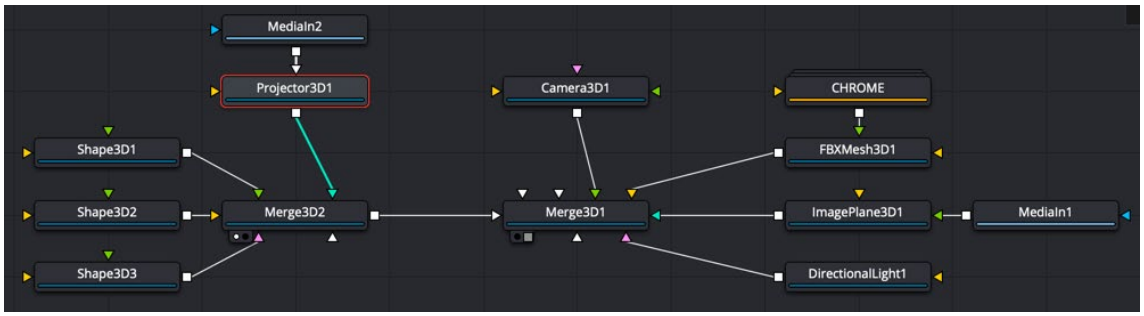
The Projector 3D has two inputs: one for the scene you are projecting on to and another for the projected image.

SceneInput: The orange scene input accepts a 3D scene. If a scene is connected to this input, then transformations applied to the spotlight also affect the rest of the scene.

ProjectiveImage: The white input expects a 2D image to be used for the projection. This connection is required.

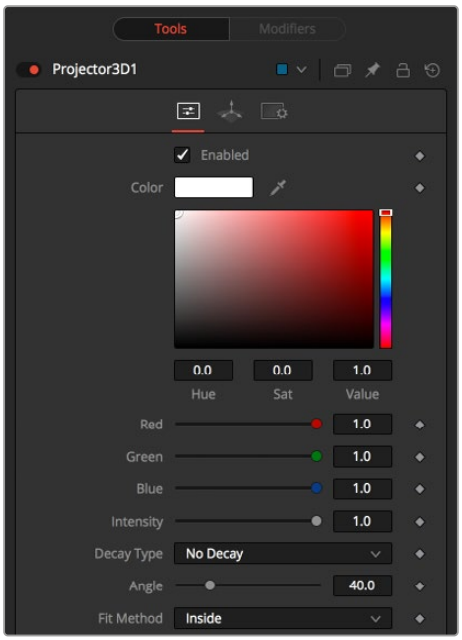
Basic Node Setup

As an example, the Projector 3D node below is used to project a texture (MediaIn2) onto 3D primitives as a way to create a simple 3D set. All the set elements are connected into a Merge 3D, which outputs the projected set into a larger scene with camera, lights, and other elements. As an alternative, the Projector 3D node could be inserted between the two Merge 3D nodes; however, the transform controls in the Projector 3D node would then affect the entire scene.



Projector 3D texturing groups of shapes to construct a set

Inspector



Projector 3D controls

Controls Tab

Enabled

When this checkbox is enabled, the projector affects the scene. Disable the checkbox to turn off the projector. This is not the same as the red switch in the upper-left corner of the Inspector. The red switch disables the tool altogether and passes the image on without any modification. The Enabled checkbox is limited to the effect part of the tool. Other parts, like scripts in the Settings tab still process as normal.

Color

The input image is multiplied by this color before being projected into the scene.

Intensity

Use this slider to set the Intensity of the projection when the Light and Ambient Light projection modes are used. In Texture mode, this option scales the Color values of the texture after multiplication by the color.

Decay Type

A projector defaults to No Falloff, meaning that its light has equal intensity on geometry, despite the distance from the projector to the geometry. To cause the intensity to fall off with distance, set the Decay type to either Linear or Quadratic modes.

Angle

The Cone Angle of the node refers to the width of the cone where the projector emits its full intensity. The larger the angle, the wider the cone angle, up to a limit of 90 degrees.

Fit Method

The Fit Method determines how the projection is fitted within the projection cone.

The first thing to know is that although this documentation may call it a “cone,” the Projector 3D and Camera 3D nodes do not project an actual cone; it’s more of a pyramid of light with its apex at the camera/projector. The Projector 3D node always projects a square pyramid of light—i.e., its X and Y angles of view are the same. The pyramid of light projected by the Camera 3D node can be non-square depending on what the Film Back is set to in the camera. The aspect of the image connected into the Projector 3D/Camera 3D does not affect the X/Y angles of the pyramid, but rather the image is scaled to fit into the pyramid based upon the fit options.

When both the aspect of the pyramid (AovY/AovX) and the aspect of the image (height * pixelAspectY)/(width * pixelAspectX) are the same, there is no need for the fit options, and in this case the fit options all do the same thing. However, when the aspect of the image and the pyramid (as determined by the Film Back settings in Camera 3D) are different, the fit options become important.

For example, Fit by Width fits the width of the image across the width of the Camera 3D pyramid. In this case, if the image has a greater aspect ratio than the aspect of the pyramid, some of the projection extends vertically outside of the pyramid.

There are five options:

- **Inside:** The image is uniformly scaled so that its largest dimension fits inside the cone. Another way to think about this is that it scales the image as big as possible subject to the restriction that the image is fully contained within the pyramid of the light. This means, for example, that nothing outside the pyramid of light ever receives any projected light.
- **Width:** The image is uniformly scaled so that its width fits inside the cone. Note that the image could still extend outside the cone in its height direction.
- **Height:** The image is uniformly scaled so that its height fits inside the cone. Note that the image could still extend outside the cone in its width direction.
- **Outside:** The image is uniformly scaled so that its smallest dimension fits inside the cone. Another way to think about this is that it scales the image as small as possible subject to the restriction that the image covers the entire pyramid (i.e., the pyramid is fully contained within the image). This means that any pixel of any object inside the pyramid of light always gets illuminated.
- **Stretch:** The image is non-uniformly scaled, so it exactly covers the cone of the projector.

Projection Mode

- **Light:** Projects the texture as a diffuse/specular light.
- **Ambient Light:** Uses an ambient light for the projection.
- **Texture:** When used in conjunction with the Catcher node, this mode allows re-lightable texture projections. The projection strikes only objects that use the catcher material as part of their material shaders.

  One useful trick is to connect a Catcher node to the Specular Texture input on a 3D Material node (such as a Blinn). This causes any object using the Blinn material to receive the projection as part of the specular highlight. This technique can be used in any material input that uses texture maps, such as the Specular and Reflection maps.

Shadows

Since the projector is based on a spotlight, it is also capable of casting shadows using shadow maps. The controls under this reveal are used to define the size and behavior of the shadow map.

- **Enable Shadows:** The Enable Shadows checkbox should be selected if the light is to produce shadows. This defaults to selected.
- **Shadow Color:** Use this standard Color control to set the color of the shadow. This defaults to black (0, 0, 0).
- **Density:** The Shadow Density determines the transparency of the shadow. A density of 1.0 produces a completely transparent shadow, whereas lower values make the shadow transparent.

- **Shadow Map Size:** The Shadow Map Size control determines the size of the bitmap used to create the shadow map. Larger values produce more detailed shadow maps at the expense of memory and performance.
- **Shadow Map Proxy:** The Shadow Map Proxy determines the size of the shadow map used for proxy and auto proxy calculations. A value of 0.5 would use a 50% shadow map.
- **Multiplicative/Additive Bias:** Shadows are essentially textures applied to objects in the scene, so there is occasionally Z-fighting, where the portions of the object that should be receiving the shadows render over the top of the shadow instead.
- **Multiplicative and Additive Bias:** Bias works by adding a small depth offset to move the shadow away from the surface it is shadowing, eliminating the Z-fighting. Too little bias and the objects can self-shadow themselves. Too much bias and the shadow can become separated from the surface. Adjust the multiplicative bias first, then fine tune the result using the additive bias control.
- **Force All Materials Non-Transmissive:** Normally, an RGBAZ shadow map is used when rendering shadows. By enabling this option, you are forcing the renderer to use a Z-only shadow map. This can lead to significantly faster shadow rendering while using a fifth as much memory. The disadvantage is that you can no longer cast “stained-glass”-like shadows.
- **Shadow Map Sampling:** Sets the quality for sampling of the shadow map.
- **Softness:** Soft edges in shadows are produced by filtering the shadow map when it is sampled. Fusion provides three separate filtering methods that produce different effects when rendering shadows.

None: Shadows have a hard edge. No filtering of the shadow map is done at all. The advantage of this method is that you only have to sample one pixel in the shadow map, so it is fast.
Constant: Shadow edges have a constant softness. A filter with a constant width is used when sampling the shadow map. Adjusting the Constant Softness slider controls the size of the filter. Note that the larger you make the filter, the longer it takes to render the shadows. If the Softness is set to constant, then a Constant slider appears. It can be used to set the overall softness of the shadow.
Variable: The softness of shadow edges grows the farther away the shadow receiver is from the shadow caster. The variable softness is achieved by changing the size of the filter based on the distance between the receiver and caster. When this option is selected, the Softness Falloff, Min Softness and Max Softness sliders appear. **Softness Falloff:** The Softness Falloff slider appears when the Softness is set to variable. This slider controls how fast the softness of shadow edges grows with distance. More precisely, it controls how fast the shadow map filter size grows based on the distance between shadow caster and receiver. Its effect is mediated by the values of the Min and Max Softness sliders. **Min Softness:** The Min Softness slider appears when the Softness is set to variable. This slider controls the Minimum Softness of the shadow. The closer the shadow is to the object casting the shadow, the sharper it is up to the limit set by this slider. **Max Softness:** The Max Softness slider appears when the Softness is set to variable. This slider controls the Maximum Softness of the shadow. The farther the shadow is from the object casting the shadow, the softer it is up to the limit set by this slider.

Common Controls

Transform and Settings Tabs

The remaining Transform and Settings tabs are common to many 3D nodes. Their descriptions can be found in “The Common Controls” section at the end of this chapter.

Renderer 3D [3Rn]



The Renderer 3D node

Renderer 3D Node Introduction

The Renderer 3D node converts the 3D environment into a 2D image using either a default perspective camera or one of the cameras found in the scene. Every 3D scene in a composition terminates with at least one Renderer 3D node. The Renderer node includes a software and OpenGL render engine to produce the resulting image. Additional render engines may also be available via third-party plugins.

The software render engine uses the system's CPU only to produce the rendered images. It is usually much slower than the OpenGL render engine, but produces consistent results on all machines, making it essential for renders that involve network rendering. The Software mode is required to produce soft shadows, and generally supports all available illumination, texture, and material features.

The OpenGL render engine employs the GPU processor on the graphics card to accelerate the rendering of the 2D images. The output may vary slightly from system to system, depending on the exact graphics card installed. The graphics card driver can also affect the results from the OpenGL renderer. The OpenGL render engines speed makes it possible to provide customized supersampling and realistic 3D depth of field options. The OpenGL renderer cannot generate soft shadows. For soft shadows, the software renderer is recommended.

Like most nodes, the Renderer's motion blur settings can be found under the Common Controls tab. Be aware that scenes containing particle systems require that the Motion Blur settings on the pRender nodes exactly match the settings on the Renderer 3D node.

Otherwise, the subframe renders conflict producing unexpected (and incorrect) results.

NOTE: The Open GL renderer respects the Color Depth option in the Image tab of the Renderer 3D node. This can cause slowdowns on certain graphics cards when rendering to int16 or float32.

Inputs

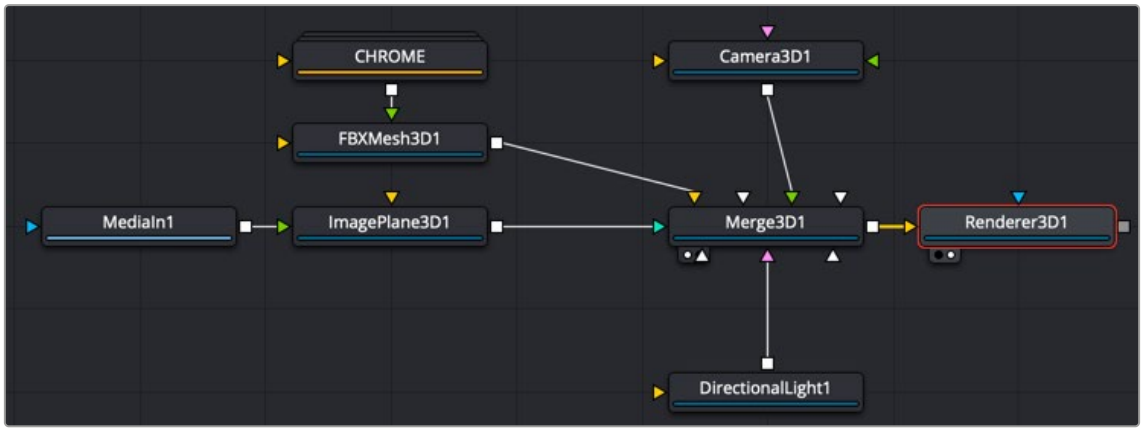
The Renderer 3D node has two inputs on the node. The main scene input takes in the Merge 3D or other 3D nodes that need to be converted to 2D. The effect mask limits the Renderer 3D output.

SceneInput: The orange scene input is a required input that accepts a 3D scene that you want to convert to 2D.

EffectMask: The blue effects mask input uses a 2D image to mask the output of the node.

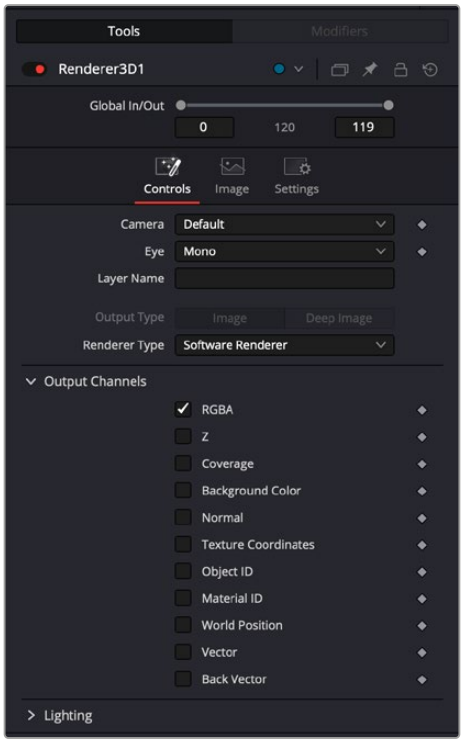
Basic Node Setup

All 3D scenes must end with a Renderer 3D node. The Renderer 3D node is used to convert a 3D scene into a 2D image. Below, the Renderer 3D node takes the output of a Merge 3D node, and renders the 3D scene into a 2D image.



Renderer 3D connected directly after a Merge 3D, rendering the 3D scene to a 2D image

Inspector



Render 3D controls

Controls Tab

Camera

The Camera menu is used to select which camera from the scene is used when rendering. The Default setting uses the first camera found in the scene. If no camera is located, the default perspective view is used instead.

Eye

The Eye menu is used to configure rendering of stereoscopic projects. The Mono option ignores the stereoscopic settings in the camera. The Left and Right options translate the camera using the stereo Separation and Convergence options defined in the camera to produce either left- or right-eye outputs. The Stacked option places the two images one on top of the other instead of side by side. Layers allows the left and right eyes to be output as individual layers.

Reporting

The first two checkboxes in this section can be used to determine whether the node prints warnings and errors produced while rendering to the console. The second set of checkboxes tells the node whether it should abort rendering when a warning or error is encountered. The default for this node enables all four checkboxes.

Renderer Type

This menu lists the available render engines. Fusion provides three: the software renderer, OpenGL renderer, and the OpenGL UV render engine. Additional renderers can be added via third-party plugins.

All the controls found below this drop-down menu are added by the render engine. They may change depending on the options available to each renderer. So, each renderer is described in its own section below.

Software Controls

Output Channels

Besides the usual Red, Green, Blue, and Alpha channels, the software renderer can also embed the following channels into the image. Enabling additional channels consumes additional memory and processing time, so these should be used only when required.

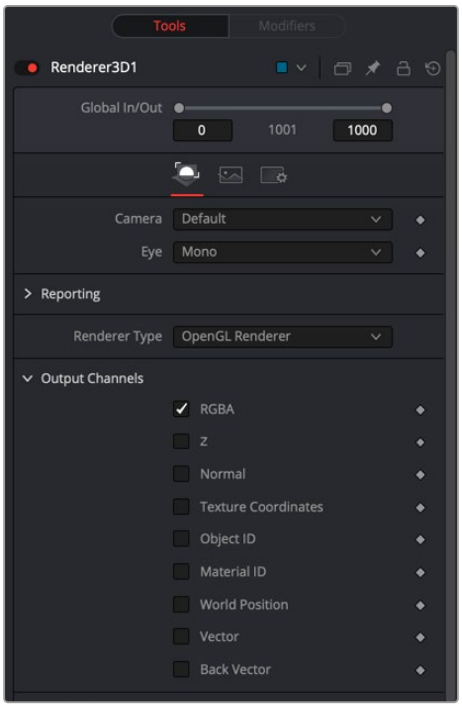
- **RGBA:** This option tells the renderer to produce the Red, Green, Blue, and Alpha color channels of the image. These channels are required, and they cannot be disabled.
- **Z:** This option enables rendering of the Z-channel. The pixels in the Z-channel contain a value that represents the distance of each pixel from the camera. Note that the Z-channel values cannot include anti-aliasing. In pixels where multiple depths overlap, the frontmost depth value is used for this pixel.
- **Coverage:** This option enables rendering of the Coverage channel. The Coverage channel contains information about which pixels in the Z-buffer provide coverage (are overlapping with other objects). This helps nodes that use the Z-buffer to provide a small degree of anti-aliasing. The value of the pixels in this channel indicates, as a percentage, how much of the pixel is composed of the foreground object.
- **BgColor:** This option enables rendering of the BgColor channel. This channel contains the color values from objects behind the pixels described in the Coverage channel.
- **Normal:** This option enables rendering of the X, Y, and Z Normals channels. These three channels contain pixel values that indicate the orientation (direction) of each pixel in the 3D space. A color channel containing values in a range from [–1,1] represents each axis.
- **TexCoord:** This option enables rendering of the U and V mapping coordinate channels. The pixels in these channels contain the texture coordinates of the pixel. Although texture coordinates are processed internally within the 3D system as three-component UVW, Fusion images store only UV components. These components are mapped into the Red and Green color channel.
- **ObjectID:** This option enables rendering of the ObjectID channel. Each object in the 3D environment can be assigned a numeric identifier when it is created. The pixels in this floating-point image channel contain the values assigned to the objects that produced the pixel. Empty pixels have an ID of 0, and the channel supports values as high as 65534. Multiple objects can share a single Object ID. This buffer is useful for extracting mattes based on the shapes of objects in the scene.

- **MaterialID:** This option enables rendering of the Material ID channel. Each material in the 3D environment can be assigned a numeric identifier when it is created. The pixels in this floating-point image channel contain the values assigned to the materials that produced the pixel. Empty pixels have an ID of 0, and the channel supports values as high as 65534. Multiple materials can share a single Material ID. This buffer is useful for extracting mattes based on a texture; for example, a mask containing all the pixels that comprise a brick texture.

Lighting

- **Enable Lighting:** When the Enable Lighting checkbox is selected, objects are lit by any lights in the scene. If no lights are present, all objects are black.
- **Enable Shadows:** When the Enable Shadows checkbox is selected, the renderer produces shadows, at the cost of some speed.

OpenGL Controls



Render 3D Open GL controls

Output Channels

In addition to the usual Red, Green, Blue, and Alpha channels, the OpenGL render engine can also embed the following channels into the image. Enabling additional channels consumes additional memory and processing time, so these should be used only when required.

- **RGBA:** This option tells the renderer to produce the Red, Green, Blue, and Alpha color channels of the image. These channels are required, and they cannot be disabled.
- **Z:** This option enables rendering of the Z-channel. The pixels in the Z-channel contain a value that represents the distance of each pixel from the camera. Note that the Z-channel values cannot include anti-aliasing. In pixels where multiple depths overlap, the frontmost depth value is used for this pixel.
- **Normal:** This option enables rendering of the X, Y, and Z Normals channels. These three channels contain pixel values that indicate the orientation (direction) of each pixel in the 3D space. A color channel containing values in a range from [–1,1] is represented by each axis.

- **TexCoord:** This option enables rendering of the U and V mapping coordinate channels. The pixels in these channels contain the texture coordinates of the pixel. Although texture coordinates are processed internally within the 3D system as three-component UVW, Fusion images store only UV components. These components are mapped into the Red and Green color channels.
- **ObjectID:** This option enables rendering of the ObjectID channel. Each object in the 3D environment can be assigned a numeric identifier when it is created. The pixels in this floating-point image channel contain the values assigned to the objects that produced the pixel. Empty pixels have an ID of 0, and the channel supports values as high as 65534. Multiple objects can share a single Object ID. This buffer is useful for extracting mattes based on the shapes of objects in the scene.
- **MaterialID:** This option enables rendering of the Material ID channel. Each material in the 3D environment can be assigned a numeric identifier when it is created. The pixels in this floating-point image channel contain the values assigned to the materials that produced the pixel. Empty pixels have an ID of 0, and the channel supports values as high as 65534. Multiple materials can share a single Material ID. This buffer is useful for extracting mattes based on a texture—for example, a mask containing all the pixels that comprise a brick texture.

Anti-Aliasing

Anti-aliasing can be enabled for each channel through the Channel menu. It produces an output image with higher quality anti-aliasing by brute force, rendering a much larger image, and then rescaling it down to the target resolution. Rendering a larger image in the first place, and then using a Resize node to bring the image to the desired resolution can achieve the exact same results. Using the supersampling built in to the renderer offers two distinct advantages over this method.

The rendering is not restricted by memory or image size limitations. For example, consider the steps to create a float-16 1920 x 1080 image with 16x supersampling. Using the traditional Resize node would require first rendering the image with a resolution of 30720 x 17280, and then using a Resize to scale this image back down to 1920 x 1080. Simply producing the image would require nearly 4 GB of memory. When anti-aliasing is performed on the GPU, the OpenGL renderer can use tile rendering to significantly reduce memory usage.

The GL renderer can perform the rescaling of the image directly on the GPU more quickly than the CPU can manage it. Generally, the more GPU memory the graphics card has, the faster the operation is performed.

Interactively, Fusion skips the anti-aliasing stage unless the HiQ button is selected in the Time Ruler. Final quality renders always include supersampling, if it is enabled.

Because of hardware limitations, point geometry (particles) and lines (locators) are always rendered at their original size, independent of supersampling. This means that these elements are scaled down from their original sizes, and likely appear much thinner than expected.

Anti-Aliasing of Aux Channels in the OpenGL Renderer

The reason Fusion supplies separate anti-aliasing options for color and aux channels in the Anti-Aliasing preset is that supersampling of color channels is quite a bit slower than aux channels. You may find that 1 x 3 LowQ/HiQ Rate is sufficient for color, but for world position or Z, you may require 4 x 12 to get adequate results. The reasons color anti-aliasing is slower are that the shaders for RGBA can be 10x to even 100x or 1000x more complex, and color is rendered with sorting enabled, while aux channels get rendered using the much faster Z-buffer method.

TIP: For some things, sometimes using an SS Z-buffer improves quality, but for other things like using the merge's PerformDepthMerge option, it may make things worse.

Do not mistake anti-aliasing with improved quality. Anti-aliasing an aux channel does not mean it's better quality. In fact, anti-aliasing an aux channel in many cases can make the results much worse. The only aux channels we recommend you enable anti-aliasing on are WorldCoord and Z.

TIP: We strongly recommend disabling Anti-Aliasing on Material ID and Object ID channels, TexCoord, Normal, BackVector, and Vector channels. The issue arises when you have multiple 3D surfaces with radically different TexCoord values in one pixel. The anti-aliasing does not restrict itself to sampling the main surface but samples both surfaces. For example, if one surface has TexCoords that are approximately (u,v) = (0, 0) within that pixel, and the other surface has (0.5, 0.5), you get a blending of these two. The blended area of the texture could have colors like (0, 0) or (0.5, 0.5), resulting in an oddly colored pixel artifact being output from the 2D Texture node. The same problem can happen for normals.

Enable (LowQ/HiQ)

These two check boxes are used to enable anti aliasing of the rendered image.

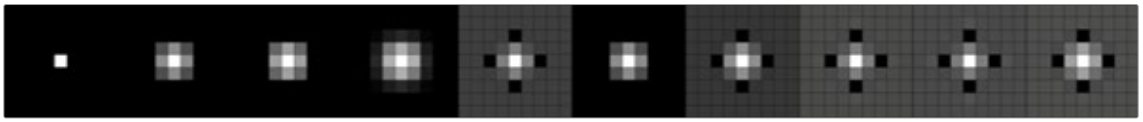
Supersampling LowQ/HiQ Rate

The LowQ and HiQ rate tells the OpenGL render how large to scale the image. For example, if the rate is set to 4 and the OpenGL renderer is set to output a 1920 x 1080 image, internally a 7680 x 4320 image is rendered and then scaled back to produce the target image. Set the multiplier higher to get better edge anti-aliasing at the expense of render time. Typically 8 x 8 supersampling (64 samples per pixel) is sufficient to reduce most aliasing artifacts.

The rate doesn't exactly define the number of samples done per destination pixel; the width of the reconstruction filter used may also have an impact.

Filter Type

When downsampling the supersized image, the surrounding pixels around a given pixel are often used to give a more realistic result. There are various filters available for combining these pixels. More complex filters can give better results but are usually slower to calculate. The best filter for the job often depends on the amount of scaling and on the contents of the image itself.



The functions of these filters are shown in the image above. From left to right these are:

Box	This is a simple interpolation scale of the image.
Bi-Linear (triangle)	This uses a simplistic filter, which produces relatively clean and fast results.
Bi-Cubic (quadratic)	This filter produces a nominal result. It offers a good compromise between speed and quality.
Bi-Spline (cubic)	This produces better results with continuous tone images but is slower than Quadratic. If the images have fine detail in them, the results may be blurrier than desired.
Catmul-Rom	This produces good results with continuous tone images which are scaled down, producing sharp results with finely detailed images.
Gaussian	This is very similar in speed and quality to Quadratic.

Mitchell	This is similar to Catmull-Rom but produces better results with finely detailed images. It is slower than Catmull-Rom.
Lanczos	This is very similar to Mitchell and Catmull-Rom but is a little cleaner and also slower.
Sinc	This is an advanced filter that produces very sharp, detailed results, however, it may produce visible `ringing' in some situations.
Bessel	This is similar to the Sinc filter but may be slightly faster.

Window Method

The Window Method menu appears only when the reconstruction filter is set to Sinc or Bessel.

- **Hanning:** This is a simple tapered window.
- **Hamming:** Hamming is a slightly tweaked version of Hanning.
- **Blackman:** A window with a more sharply tapered falloff.

Accumulation Effects

Accumulation effects are used for creating depth of field effects. Enable both the Enable Accumulation Effects and Depth of Field checkboxes, and then adjust the quality and Amount sliders.

The blurrier you want the out-of-focus areas to be, the higher the quality setting you need. A low amount setting causes more of the scene to be in focus.

The accumulation effects work in conjunction with the Focal plane setting located in the Camera 3D node. Set the Focal Plane to the same distance from the camera as the subject you want to be in focus. Animating the Focal Plane setting creates rack of focus effects.

Lighting

- **Enable Lighting:** When the Enable Lighting checkbox is selected, any lights in the scene light objects. If no lights are present, all objects are black.
- **Enable Shadows:** When the Enable Shadows checkbox is selected, the renderer produces shadows, at the cost of some speed.

Texturing

- **Texture Depth:** Lets you specify the bit depth of texture maps.
- **Warn about unsupported texture depths:** Enables a warning if texture maps are in an unsupported bit depth that Fusion can't process.

Lighting Mode

The Per-vertex lighting model calculates lighting at each vertex of the scene's geometry. This produces a fast approximation of the scene's lighting but tends to produce blocky lighting on poorly tessellated objects. The Per-pixel method uses a different approach that does not rely on the detail in the scene's geometry for lighting, so it generally produces superior results.

Although the per-pixel lighting with the OpenGL renderer produces results closer to that produced by the more accurate software renderer, it still has some disadvantages. The OpenGL renderer is less capable of dealing correctly with semi-transparency, soft shadows, and colored shadows, even with per-pixel lighting. The color depth of the rendering is limited by the capabilities of the graphics card in the system.

Transparency

The OpenGL renderer reveals this control for selecting which ordering method to use when calculating transparency.

- **Z Buffer (fast):** This mode is extremely fast and is adequate for scenes containing only opaque objects. The speed of this mode comes at the cost of accurate sorting; only the objects closest to the camera are certain to be in the correct sort order. So, semi-transparent objects may not be shown correctly, depending on their ordering within the scene.
- **Sorted (accurate):** This mode sorts all objects in the scene (at the expense of speed) before rendering, giving correct transparency.
- **Quick Mode:** This experimental mode is best suited to scenes that almost exclusively contain particles.

Shading Model

Use this menu to select a shading model to use for materials in the scene. Smooth is the shading model employed in the viewers, and Flat produces a simpler and faster shading model.

Wireframe

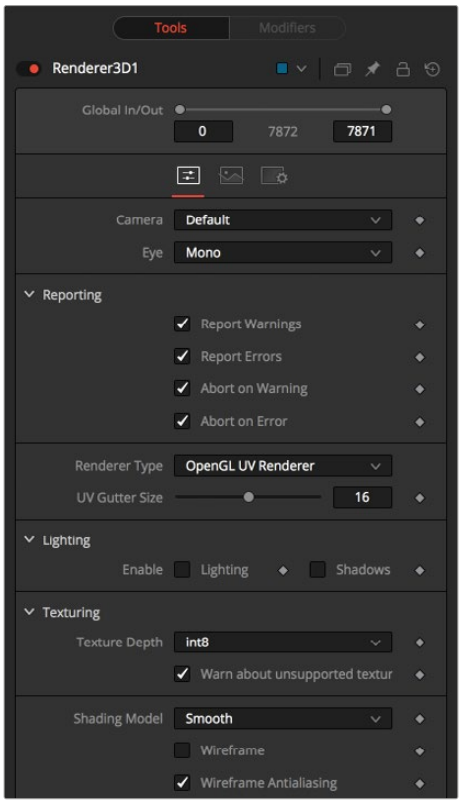
Renders the whole scene as wireframe. This shows the edges and polygons of the objects. The edges are still shaded by the material of the objects.

Wireframe Anti-Aliasing

Enables anti-aliasing for the Wireframe render.

OpenGL UV Renderer

The OpenGL UV renderer is a special case render engine. It is used to take a model with existing textures and render it out to produce an unwound flattened 2D version of the model. Optionally, lighting can be baked in. This is typically done so you can then paint on the texture and reapply it.



Render 3D Open GL UV controls

Below are some issues to be aware of when using the OpenGL UV renderer.

Baked-in lighting: After you have baked lighting into a model's texture, you need to be careful to turn lighting off on the object later when you render it with the baked-in lighting texture.

Single textures/multiple destinations: Beware of cases where a single area of the texture map is used on multiple areas of the model. This is often done to save texture memory and decrease modeling time. An example is the texture for a person where the artist mirrored the left side mesh/uvs/texture to produce the right side. Trying to bake in lighting in this case won't work.

Unwrapped more the one mesh: Unwrapping more than one mesh at once can cause problems. The reason is that most models are authored so they make maximum usage of (u,v) in [0,1] x [0,1], so that in general models overlap each other in UV space.

Seams: When the UV gutter size is left at 0, this produces seams when the model is retextured with the unwrapped texture.

UV Gutter Size: Increase this value to hide seams between faces.

Common Controls

Image and Settings Tabs

The remaining controls for the Image and Settings tabs are common to many 3D nodes. Their descriptions can be found in "The Common Controls" section at the end of this chapter.

Replace Material 3D [3Rpl]



The Replace Material 3D node

Replace Material 3D Node Introduction

The Replace Material 3D node replaces the material applied to all the geometry in the input scene with its own material input. Any lights or cameras in the input scene are passed through unaffected.

The scope of the replacement can be limited using Object and Material identifiers in the Inspector. The scope can also be limited to individual channels, making it possible to use a completely different material on the Red channel, for example.

Since the Text 3D node does not include a material input, you can use the Replace Material to add material shaders to the text.

Inputs

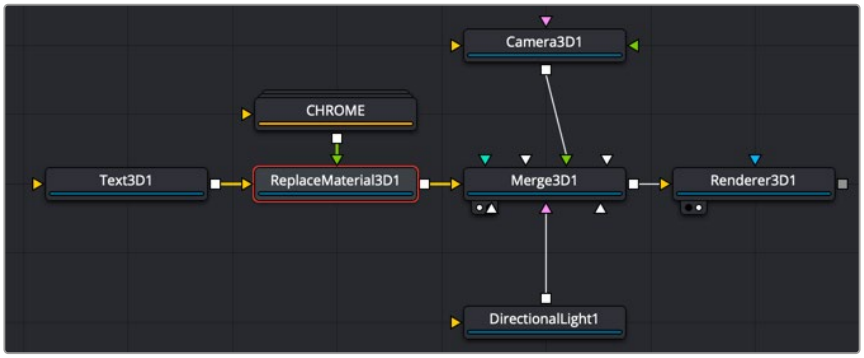
The Replace Material node has two inputs: one for the 3D scene, object, or 3D text that contains the original material, and a material input for the new replacement material.

SceneInput: The orange scene input accepts a 3D scene or 3D text that you want to replace the material.

MaterialInput: The green material input accepts either a 2D image or a 3D material. If a 2D image is provided, it is used as a diffuse texture map for the basic material built into the node. If a 3D material is connected, then the basic material is disabled.

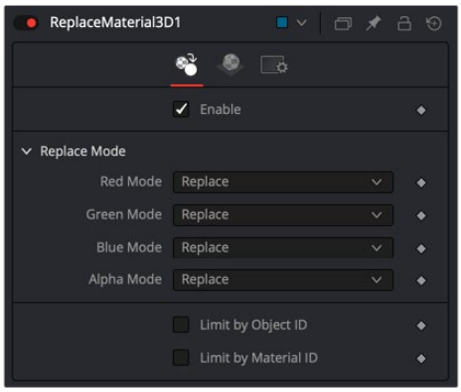
Basic Node Setup

The Replace Material 3D node is inserted directly after the 3D object or scene whose material you want to replace. Below, it is used to replace the default material on a Text 3D node with a chrome shader.



Replace Material 3D used to replace the default material of a Text 3D node

Inspector



Replace Material 3D controls

Controls Tab

Enable

This checkbox enables the material replacement. This is not the same as the red switch in the upper-left corner of the Inspector. The red switch disables the tool altogether and passes the image on without any modification. The enable checkbox is limited to the effect part of the tool. Other parts, like scripts in the Settings tab, still process as normal.

Replace Mode

The Replace Mode section offers four methods of replacing each RGBA channel:

- **Keep:** Prevents the channel from being replaced by the input material.
- **Replace:** Replaces the material for the corresponding color channel.
- **Blend:** Blends the materials together.
- **Multiply:** Multiplies the channels of both inputs.

Limit by Object ID/Material ID

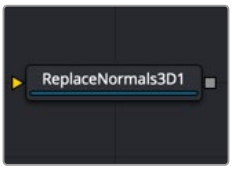
When enabled, a slider appears where the desired IDs can be set. Other objects keep their materials. If both options are enabled, an object must satisfy both conditions.

Common Controls

Material and Settings Tabs

The remaining controls for the Material and Settings tabs are common to many 3D nodes. Their descriptions can be found in “The Common Controls” section at the end of this chapter.

Replace Normals 3D [3RpN]



The Replace Normals 3D node

Replace Normals Node Introduction

In 3D modeling, normals are vectors used to determine the direction light reflects off surfaces. The Replace Normals node is used to replace the normals/tangents on incoming geometry, effectively adjusting the surface of an object between smooth and flat. All geometry connected to the scene input on the node is affected. Lights/Cameras/PointClouds/Locators/Materials, and other non-mesh nodes are passed through unaffected. The normals/tangents affected by this node are Per-vertex normals/tangents, not Per-face normals/tangents. The input geometry must have texture coordinates in order for tangents to be computed. Sometimes geometry does not have texture coordinates, or the texture coordinates were set to All by the FBX import because they were not present on the FBX.

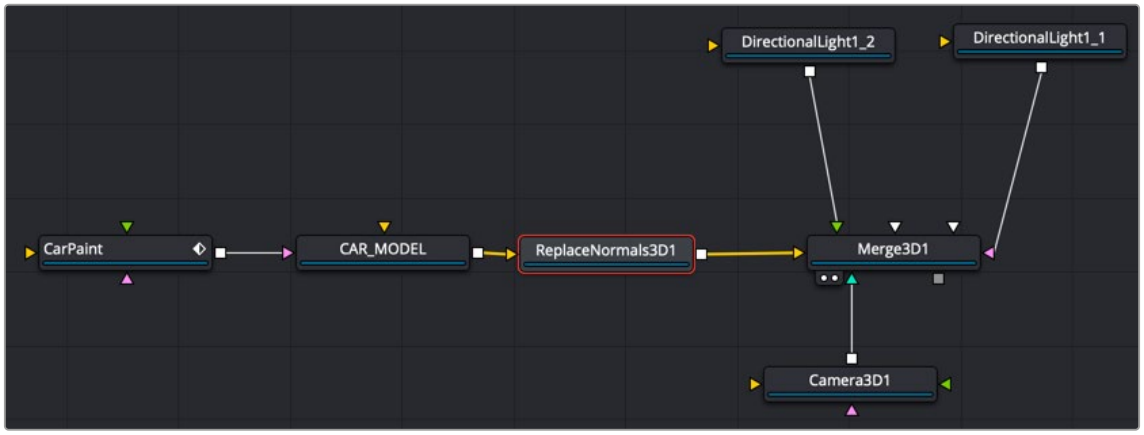
Inputs

The Replace Normals node has a single input for the 3D scene or incoming geometry.

SceneInput: The orange scene input accepts a 3D scene or 3D geometry that contains the normal coordinates you want to modify.

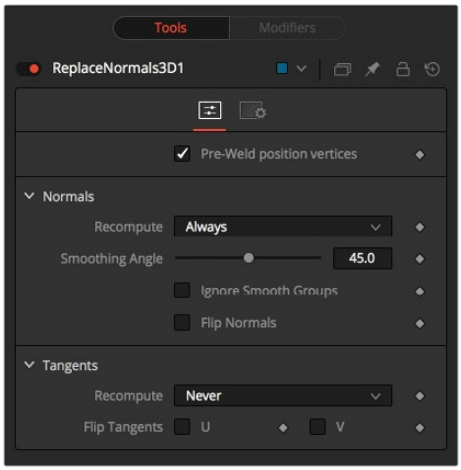
Basic Node Setup

The Replace Normals 3D node is inserted directly after the 3D object or scene whose normals you want to modify. Below, it is used to smooth the material on an imported 3D model.



Replace Normals 3D used to smooth normals on 3D geometry

Inspector



Replace Normals 3D controls

Control Tab

The options in the Control tab deal with repairing 3D geometry and then recomputing normals/tangents.

Pre-Weld Position Vertices

Sometimes position vertices are duplicated in a geometry, even though they have the same position, causing normals/tangents to be miscomputed. The results of pre-welding are thrown away; they do not affect the output geometry's position vertices.

Recompute

Controls when normals/tangents are recomputed.

- **Always:** The normals on the mesh are always recomputed.
- **If Not Present:** The normals on the mesh are recomputed only if they are not present.
- **Never:** The normals are never computed. This option is useful when animating.

Smoothing Angle

Adjacent faces with angles in degrees smaller than this value have their adjoining edges smoothed across. A typical value one might choose for the Smoothing Angle is between 20 and 60 degrees. There is special case code for 0.0f and 360.0f (f stands for floating-point value). When set to 0.0f, faceted normals are produced; this is useful for artistic effect.

Ignore Smooth Groups

If set to False, two faces that have different Smooth Groups are not smoothed across (e.g., the faces of a cube or the top surfaces of a cylinder have different Smooth Groups). If you check this On and set the smoothing angle large enough, the faces of a cube are smoothed across. There is currently no way to visualize Smooth Groups within Fusion.

Flip Normals

Flipping of tangents can sometimes be confusing. Flip has an effect if the mesh has tangent vectors. Most meshes in Fusion don't have tangent vectors until they reach a Renderer 3D, though. Also, when viewing tangent vectors in the viewers, the tangent vectors are created if they don't exist. The confusing thing is if you view a Cube 3D that has no tangent vectors and press the FlipU/FlipV button, nothing happens. This is a result of there being no tangent vectors to create, but later the GL renderer can create some (unflipped) tangent vectors.

There are five items you should be aware of when dealing with normals.

#1 The FBX importer recomputes the normals if they don't exist, but you can get a higher-quality result from the Replace Normals node.

#2 Bump maps can sometimes depend on the model's normals. Specifically, when you simplify a complex high polygon model to a low polygon model + bump map, the normals and bump map can become "linked." Recomputing the normals in this case can make the model look funny. The bump map was intended to be used with the original normals.

#3 Most primitives in Fusion are not generated with tangents; when needed, they are generated on the fly by a Renderer 3D and cached.

#4 Tangents currently are only needed for bump mapping. If a material needs bump mapping, then tangents are created. These tangents are created with some default settings (e.g., Smoothing Angle, and so on). If you don't want Fusion automatically creating tangents, you can use the Replace Normals node to create them manually.

#5 All computations are done in the local coordinates of the geometries instead of in the coordinate system of the Replace Normals 3D node. This can cause problems when there is a non-uniform scale applied to the geometry before Replace Normals 3D is applied.

Common Controls

Settings Tab

The Settings tab is common to many 3D nodes. The description of these controls can be found in "The Common Controls" section at the end of this chapter.

Replicate 3D [3Rep]



The Replicate 3D node

Replicate 3D Node Introduction

The Replicate 3D node replicates input geometry at positions of destination vertices. The vertices can be mesh vertices as well as particle positions. For each copy of the replicated input geometry, various transformations can be applied. The options in the Jitter tab allow non-uniform transformations, such as random positioning or sizes.

Inputs

There are two inputs on the Replicate 3D node: one for the destination geometry that contains the vertices, and one for the 3D geometry you want to replicate.

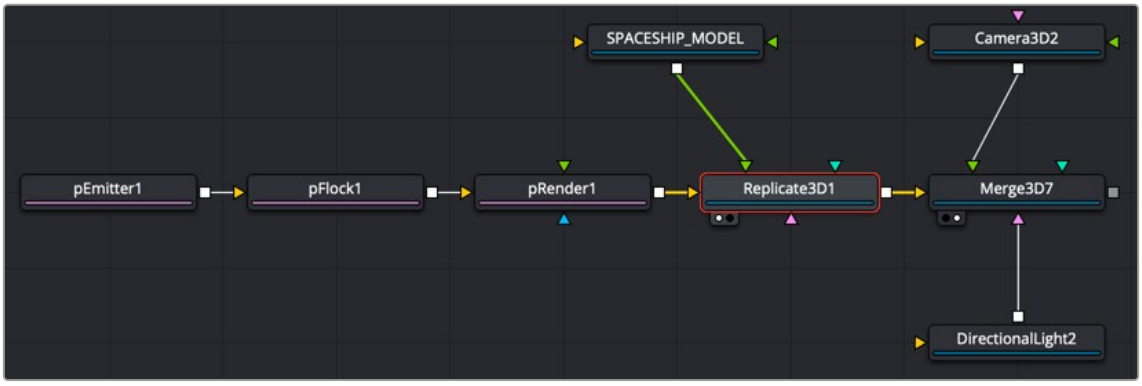
Destination: The orange destination input accepts a 3D scene or geometry with vertex positions, either from the mesh or 3D particle animations.

Input[#]: The input accepts the 3D scene or geometry for replicating. Once this input is connected, a new input for alternating 3D geometry is created.

At least one connected input is required.

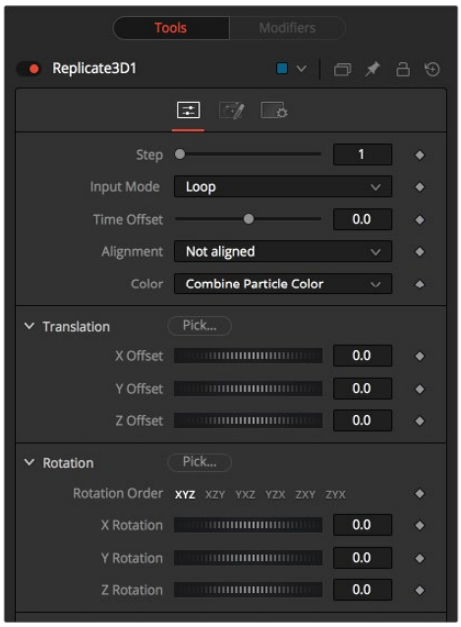
Basic Node Setup

In the example below, a Replicate 3D node is inserted directly after the pRender node. A spaceship FBX node is connected to the green input representing the object that will be replicated based on the particles. Each particle cell takes on the shape of the 3D geometry connected to the input.



Replicate 3D used to create an armada of swarming spaceships

Inspector



Replicated 3D Jitter controls

Controls Tab

Step

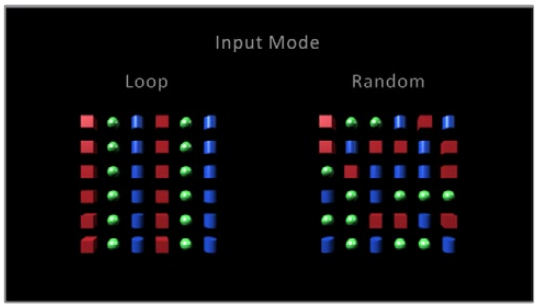
Defines how many positions are skipped. For example, a step of 3 means that only every third vertice of the destination mesh is used, while a step of 1 means that all positions are used.

The Step setting helps to keep reasonable performance for big destination meshes. On parametric geometry like a torus, it can be used to isolate certain parts of the mesh.

Point clouds are internally represented by six points once the Make Renderable option has been set. To get a single point, use a step of 6 and set an X offset of –0.5 to get to the center of the point cloud. Use –0.125 for Locator 3Ds. Once these have been scaled, the offset may differ.

Input Mode

This menu defines in which order multiple input scenes are replicated at the destination. No matter which setting you choose, if only one input scene is supplied this setting has no effect.



The Input Mode Loop vs. Random order

When set to Loop, the inputs are used successively. The first input is at the first position, the second input at the second position, and so on. If there are more positions in the destination present than inputs, the sequence is looped.

When set to Random, a definite but random input for each position is used based on the seed in the Jitter tab. This input mode can be used to simulate variety with few input scenes.

The Death of Particles setting causes the input geometries' IDs to change; therefore, their copy order may change.

Time Offset

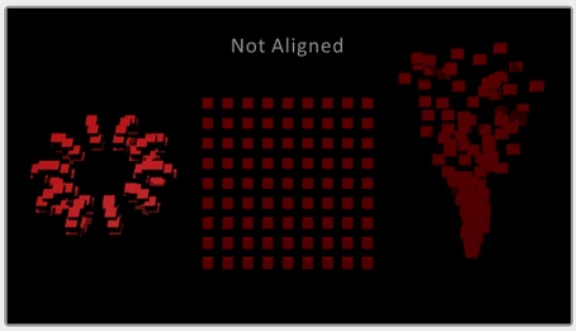
Use the Time Offset slider to offset any animations that are applied to the input geometry by a set amount per copy. For example, set the value to –1.0 and use a cube set to rotate on the Y-axis as the source. The first copy shows the animation from a frame earlier; the second copy shows animation from a frame before that, etc.

This can be used with great effect on textured planes—for example, where successive frames of a video clip can be shown.

Alignment

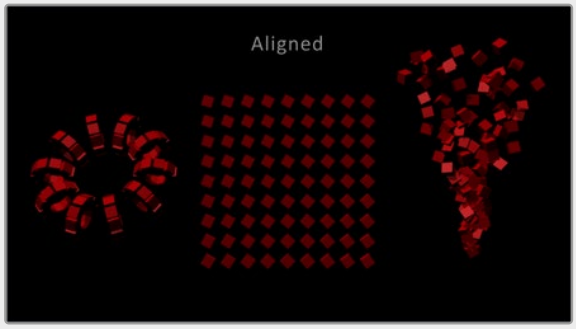
Alignment specifies how to align the copies in respect of the destination mesh normal or particle rotation.

Not Aligned: Does not align the copy. It stays rotated in the same direction as its input mesh.



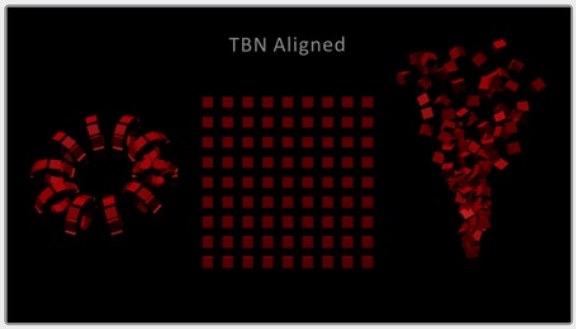
Replicate 3D Not Aligned layout

Aligned: This mode uses the point's normal and tries to reconstruct an upvector. It works best with organic meshes that have unwelded vertices, like imported FBX meshes, since it has the same rotations for vertices at the same positions. On plane geometric meshes, a gradual shift in rotation is noticeable. For best results, it is recommended to use this method at the origin before any transformations.



Replicate 3D Aligned layout

Aligned TBN: This mode results in a more accurate and stable alignment based on the tangent, binormal, and normal of the destination point. This works best for particles and geometric shapes. On unwelded meshes, two copies of multiple unwelded points at the same position may lead to different alignments because of their individual normals.



Replicate 3D Aligned TBN layout

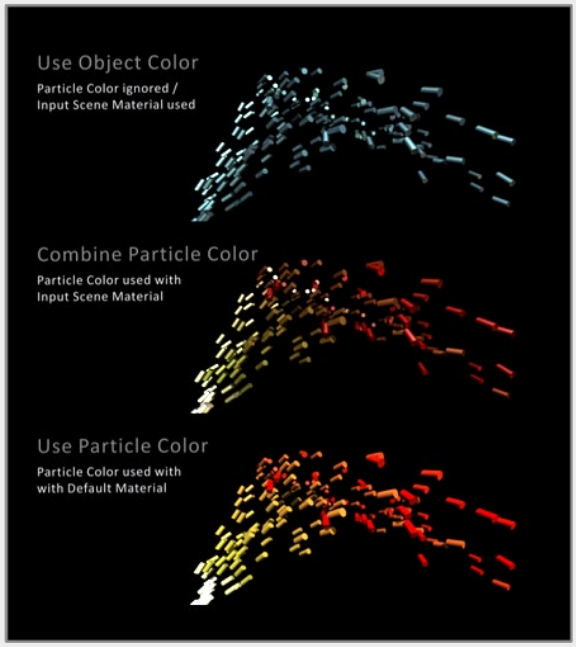
Color

Affects the diffuse color or shader of each copy based on the input's particle color.

Use Object Color: Does not use the color of the destination particle.

Combine Particle Color: Uses the shader of any input mesh and modifies the diffuse color to match the color from the destination particle.

Use Particle Color: Replaces the complete shader of any input mesh with a default shader. Its diffuse color is taken from the destination particle.



Replicate 3D Color options

Translation

These three sliders tell the node how much offset to apply to each copy. An X Offset of 1 would offset each copy one unit, one unit along the X-axis from the last copy.

Rotation Order

These buttons can be used to set the order in which rotations are applied to the geometry. Setting the rotation order to XYZ would apply the rotation on the X-axis first, followed by the Y-axis rotation, and then the Z-axis rotation.

XYZ Rotation

These three rotation sliders tell the node how much rotation to apply to each copy.

XYZ Pivot

The pivot controls determine the position of the pivot point used when rotating each copy.

Lock XYZ

When the Lock XYZ checkbox is selected, any adjustment to the scale is applied to all three axes simultaneously.

If this checkbox is disabled, the Scale slider is replaced with individual sliders for the X, Y, and Z scales.

Scale

The Scale control sets how much scaling to apply to each copy.

Jitter Tab

The Jitter tab can be used to introduce randomness to various parameters.



Replicated 3D Jitter controls

Random Seed/Randomize

The Random Seed is used to generate the jitter applied to the replicated objects. Two Replicate nodes with identical settings but different random seeds will produce two completely different results. Click the Randomize button to assign a Random Seed value.

Time Offset

Use the Time Offset slider to offset any animations that are applied to the source geometry. Unlike Time Offset on the Controls tab, Jitter Time Offset is random, based on the Random Seed setting.

Translation XYZ Jitter

Use these three controls to adjust the variation in the translation of the replicated objects.

Rotation XYZ Jitter

Use these three controls to adjust the variation in the rotation of the replicated objects.

Pivot XYZ Jitter

Use these three controls to adjust the variation in the rotational pivot center of the replicated objects. This affects only the additional jitter rotation, not the rotation produced by the rotation settings in the Controls tab.

Scale XYZ Jitter

Use this control to adjust the variation in the scale of the replicated objects. Uncheck the Lock XYZ checkbox to adjust the scale variation independently on all three axes.

Common Controls

Settings Tab

The Settings tab is common to many 3D nodes. The description of these controls can be found in “The Common Controls” section at the end of this chapter.

Ribbon 3D [3Ri]



The Ribbon 3D node

Ribbon 3D Node Introduction

Ribbon 3D generates an array of subdivided line segments or a single line between two points. It is quite useful for motion graphics, especially in connection with Replicate 3D to attach other geometry to the lines, and with Displace3D for creating lightning bolt-like structures. The array of lines is, by default, assigned with texture coordinates, so they can be used with a 2D texture. As usual, UVMap 3D can be used to alter the texture coordinates. This node relies heavily on certain OpenGL features and does not produce any visible result when the Renderer 3D node is set to use the software renderer.

Furthermore, the way lines are drawn is completely up to the graphics card capabilities, so the ribbon appearance may vary based on your computer's graphics card.

Inputs

There are two inputs on the Ribbon 3D node: one for the destination geometry that contains the vertices, and one for the 3D geometry you want to replicate.

3D Scene: The orange input accepts a 3D scene or geometry.

Material: The input accepts the 2D texture for the ribbon.

Neither connected input is required.

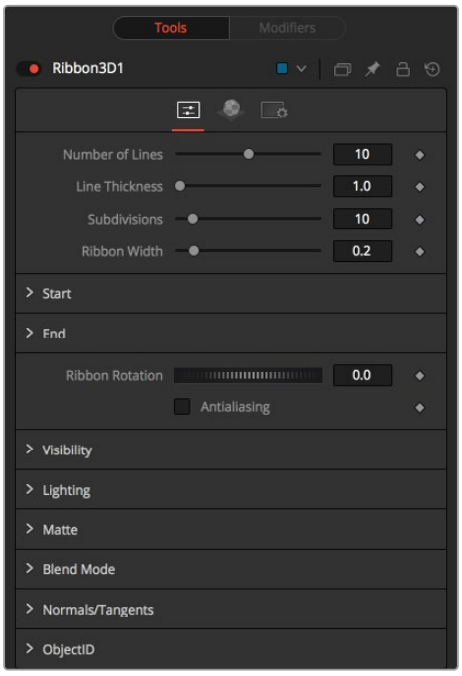
Basic Node Setup

In the example below, a Ribbon 3D node is used to generate lines. A gradient background is connected to "colorize" the lines. Additional nodes are then used after the Ribbon 3D to bend and distort the lines.



Ribbon 3D generates lines distorted by additional nodes

Inspector



Ribbon 3D controls

Controls Tab

The Controls tab determines the number of ribbon strands, their size, length, and spacing.

Number of Lines

The number of parallel lines drawn between the start point and end point.

Line Thickness

Line thickness is allowed in the user interface to take on a floating-point value, but some graphics cards allow only integer values. Some cards may only allow lines equal to or thicker than one, or max out at a certain value.

Subdivision Level

The number of vertices on each line between start point and end points. The higher the number, the more precise and smoother 3D displacement appears.

Ribbon Width

Determines how far the lines are apart from each other.

Start

XYZ control to set the start point of the ribbon.

End

XYZ control to set the end point of the ribbon.

Ribbon Rotation

Allows rotation of the ribbon around the virtual axis defined by start point and end points.

Anti-Aliasing

Allows you to apply anti-aliasing to the rendered lines. Using anti-aliasing isn't necessarily recommended. When activated, there may be be gaps between the line segments. This is especially noticeable with high values of line thickness. Again, the way lines are drawn is completely up to the graphics card, which means that these artifacts can vary from card to card.

Common Controls

Controls, Materials, and Settings Tabs

The controls for Visibility, Lighting, Matte, Blend Mode, Normals/Tangents, and Object ID in the Controls tab are common in many 3D nodes. The Materials tab and Settings tab in the Inspector are also duplicated in other 3D nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Shape 3D [3Sh]



The Shape 3D node

Shape 3D Node Introduction

The Shape 3D node is used to produce several basic primitive 3D shapes, including planes, cubes, spheres, and cylinders.

Inputs

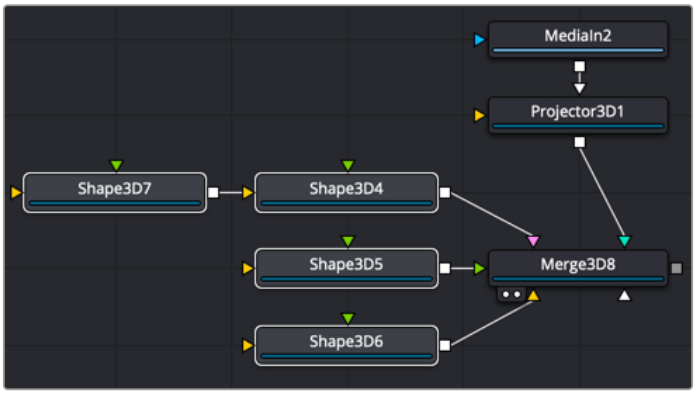
There are two optional inputs on the Shape 3D. The scene input can be used to combine additional geometry with the Shape 3D, while the material input can be used to texture map the Shape 3D object.

SceneInput: Although the Shape 3D creates its own 3D geometry, you can use the orange scene input to combine an additional 3D scene or geometry.

MaterialInput: The green input accepts either a 2D image or a 3D material. If a 2D image is provided, it is used as a diffuse texture map for the basic material built into the node. If a 3D material is connected, then the basic material is disabled.

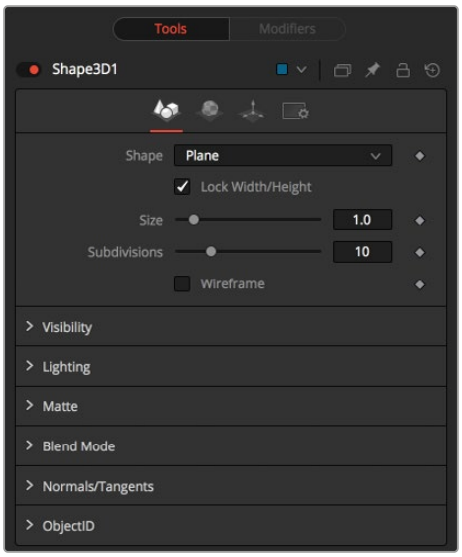
Basic Node Setup

In the example below, four Shape 3D nodes are used to create the primitives of a 3D set. Two of the Shape 3D nodes are connected creating a more complex primitive shape. Those shapes can then be used with a Projector 3D to texture them with a realistic material.



Shape 3D nodes combined with Projector 3D to create a realistic 3D set

Inspector



Shape 3D controls

Controls Tab

The Controls tab allows you to select a shape and modify its geometry. Different controls appear based on the specific shape that you choose to create.

Shape

This menu allows you to select the primitive geometry produced by the Shape 3D node. The remaining controls in the Inspector change to match the selected shape.

- **Lock Width/Height/Depth:** [plane, cube] If this checkbox is selected, the width, height, and depth controls are locked together as a single size slider. Otherwise, individual controls over the size of the shape along each axis are provided.
- **Size Width/Height/Depth:** [plane, cube] Used to control the size of the shape.

Cube Mapping

When Cube is selected in the shape menu, the Cube uses cube mapping to apply the Shape node's texture (a 2D image connected to the material input on the node).

Radius

When a Sphere, Cylinder, Cone, or Torus is selected in the shape menu, this control sets the radius of the selected shape.

Top Radius

When a cone is selected in the Shape menu, this control is used to define a radius for the top of a cone, making it possible to create truncated cones.

Start/End Angle

When the Sphere, Cylinder, Cone, or Torus shape is selected in the Shape menu, this range control determines how much of the shape is drawn. A start angle of 180° and end angle of 360° would only draw half of the shape.

Start/End Latitude

When a Sphere or Torus is selected in the Shape menu, this range control is used to crop or slice the object by defining a latitudinal subsection of the object.

Bottom/Top Cap

When Cylinder or Cone is selected in the Shape menu, the Bottom Cap and Top Cap checkboxes are used to determine if the end caps of these shapes are created or if the shape is left open.

Section

When the Torus is selected in the Shape menu, Section controls the thickness of the tube making up the torus.

Subdivision Level/Base/Height

The Subdivision controls are used to determine the tessellation of the mesh on all shapes. The higher the subdivision, the more vertices each shape has.

Wireframe

Enabling this checkbox causes the mesh to render only the wireframe for the object.

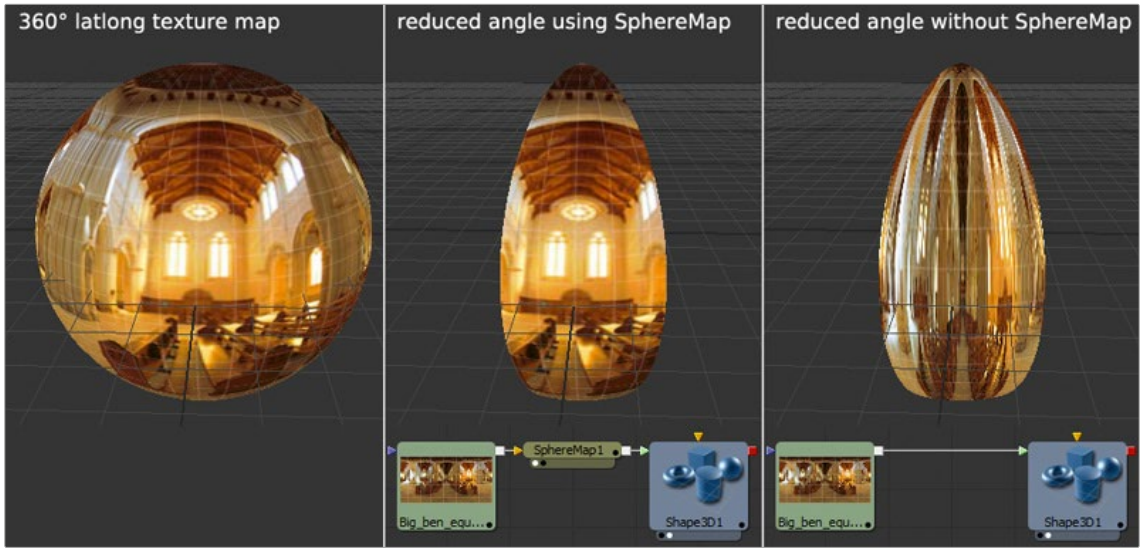
Common Controls

Controls, Materials, Transform and Settings Tabs

The controls for Visibility, Lighting, Matte, Blend Mode, Normals/Tangents, and Object ID in the Controls tab are common in many 3D nodes. The Materials tab, Transforms tab, and Settings tab in the Inspector are also duplicated in other 3D nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Sphere Map vs. Connecting the Texture to a Sphere Directly

You can connect a LatLong (equirectangular) texture map directly to a sphere instead of piping it through the Sphere Map node first. This results in a different rendering if you set the start/end angle and latitude to less than 360°/180°. In the first case, the texture is squashed. When using the Sphere Map node, the texture is cropped. Compare:



Spherical mapping differences

NOTE: If you pipe the texture directly into the sphere, it is also mirrored horizontally. You can change this by using a Transform node first.

Soft Clip [3SC]



The Soft Clip node

Soft Clip Node Introduction

The Soft Clip node is used to fade out geometry and particles that get close to the camera. This helps avoid the visible “popping off” that affects many particle systems and 3D flythroughs.

This node is very similar to the Fog 3D node, in that it is dependent on the geometry’s distance from the camera.

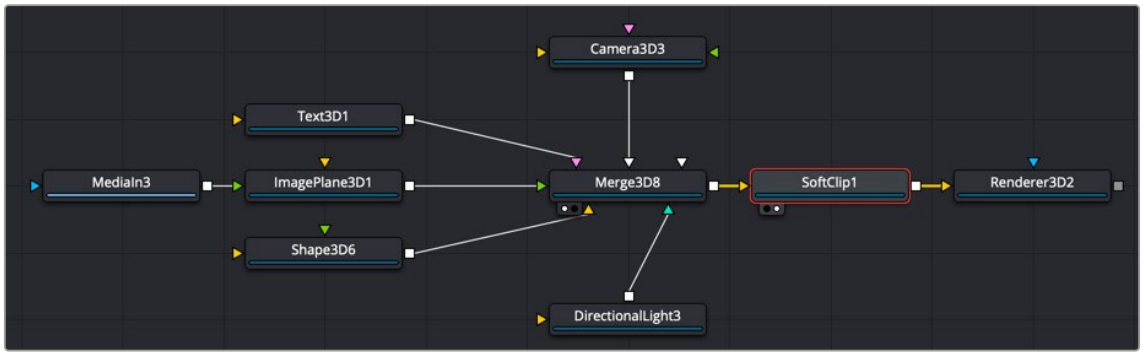
Inputs

The Soft Clip includes only a single input for a 3D scene that includes a camera connected to it.

SceneInput: The orange scene input is a required connection. It accepts a 3D scene input that includes a Camera 3D node.

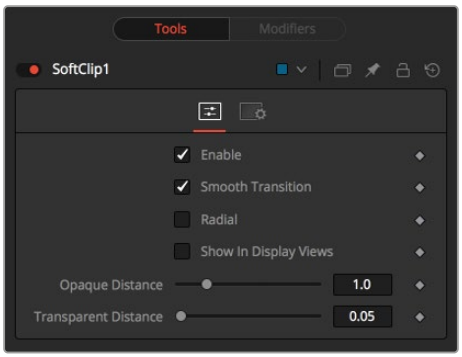
Basic Node Setup

The Soft Clip node is usually placed just before the Renderer 3D node to ensure that downstream adjustments to lighting and textures do not affect the result. It can be placed in any part of the 3D portion of the node tree if the soft clipping effect is only required for a portion of the scene.



Soft Clip placed between a Merge 3D and a Renderer 3D node

Inspector



Soft Clip controls

Controls Tab

The Controls tab determines how an object transitions between opaque and transparent as it moves closer to the camera.

Enable

This checkbox can be used to enable or disable the node. This is not the same as the red switch in the upper-left corner of the Inspector. The red switch disables the tool altogether and passes the image on without any modification. The Enable checkbox is limited to the effect of the tool. Other parts, like scripts in the Settings tab, still process as normal.

Smooth Transition

By default, an object coming closer and closer to the camera slowly fades out with a linear progression. With the Smooth Transition checkbox enabled, the transition changes to a nonlinear curve, arguably a more natural-looking transition.

Radial

By default, the soft clipping is done based on the perpendicular distance to a plane (parallel with the near plane) passing through the eye point. When the Radial option is checked, the Radial distance to the eye point is used instead of the Perpendicular distance. The problem with Perpendicular distance soft clipping is that when you move the camera about, as objects on the left or right side of the frustum move into the center, they become less clipped, although they remain the same distance from the eye. Radial soft clip fixes this. Sometimes Radial soft clipping is not desirable.

For example, if you apply soft clip to an object that is close to the camera, like an image plane, the center of the image plane could be unclipped while the edges could be fully clipped because they are farther from the eye point.

Show In Display Views

Normally, the effect is only visible when the scene is viewed using a Camera node. When enabled, the soft clip becomes visible in the scene from all points of view.

Transparent/Opaque Distance

Defines the range of the soft clip. The objects begin to fade in from an opacity of 0 at the Transparent distance and are fully visible at the Opaque distance. All units are expressed as distance from the camera along the Z-axis.

Common Controls

Settings Tab

The Settings tab in the Inspector is duplicated in other 3D nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Spherical Camera [3SC]



The Spherical Camera node

Spherical Camera Node Introduction

The Spherical Camera allows the 3D Renderer node to output an image covering all viewing angles, laid out in several different formats. This image may be used, for example, as a skybox texture or reflection map or viewed in a VR headset. The Image Width setting in the 3D Renderer sets the size of each square cube face, so the resulting image may be a multiple of this size horizontally and vertically.

Inputs

The Spherical camera node has two inputs.

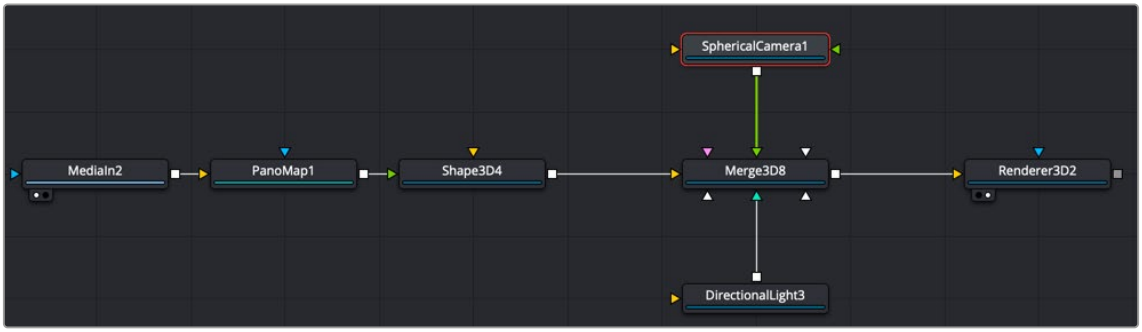
Image: This orange image input requires an image in a spherical layout, which can be any of LatLong (2:1 equirectangular), VR180, Horizontal/Vertical Cross, or Horizontal/Vertical Strip.

Stereo Input: The green input for a right stereo camera if you are working in stereo VR.

Neither input is required.

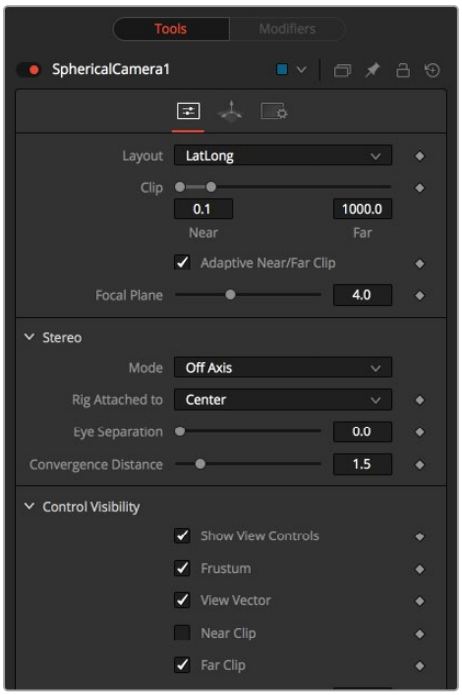
Basic Node Setup

In many ways, the Spherical Camera is set up identically to the regular Camera 3D node. The output of the camera connects into a Merge 3D. Typically, the Merge 3D has an image from a LatLong or H Cross/V Cross formatted image either directly or through a Panomap node. The image is wrapped around a sphere, and the camera is placed inside the sphere.



Spherical Camera placed inside a sphere

Inspector



Spherical Camera controls

Controls Tab

Layout

- **VCross and HCross:** VCross and HCross are the six square faces of a cube laid out in a cross, vertical or horizontal, with the forward view in the center of the cross, in a 3:4 or 4:3 image.
- **VStrip and HStrip:** VStrip and HStrip are the six square faces of a cube laid vertically or horizontally in a line, ordered as Left, Right, Up, Down, Back, Front (+X, -X, +Y, -Y, +Z, -Z), in a 1:6 or 6:1 image.
- **LatLong:** LatLong is a single 2:1 image in equirectangular mapping.
- **VR 180** is a 1:1 180 degree stereoscopic image.

Near/Far Clip

The clipping plane is used to limit what geometry in a scene is rendered based on the object's distance from the camera's focal point. This is useful for ensuring that objects that are extremely close to the camera are not rendered and for optimizing a render to exclude objects that are too far away to be useful in the final rendering.

The default perspective camera ignores this setting unless the Adaptively Adjust Near/Far Clip checkbox control below is disabled.

The values are expressed in units, so a far clipping plane of 20 means that any objects more than 20 units from the camera are invisible to the camera. A near clipping plane of 0.1 means that any objects closer than 0.1 units are also invisible.

Adaptively Adjust Near/Far Clip

When selected, the Renderer automatically adjusts the camera's near/far clipping plane to match the extents of the scene. This setting overrides the values of the Near and Far clip range control described above. This option is not available for orthographic cameras.

Viewing Volume Size

The Viewing Volume Size control appears only when the Projection Type is set to Orthographic. It determines the size of the box that makes up the camera's field of view. The Z distance of an orthographic camera from the objects it sees does not affect the scale of those objects, only the viewing size does.

NOTE: A smaller range between the near and far clipping plane allows greater accuracy in all depth calculations. If a scene begins to render strange artifacts on distant objects, try increasing the distance for the near clip plane. Use the vertical aperture size to get the vertical angle of view and the horizontal aperture size to get the horizontal angle of view.

Plane of Focus (for Depth of Field)

This value is used by the OpenGL renderer to calculate depth of field. It defines the distance to a virtual target in front of the camera.

Stereo Method

This control allows you to adjust your stereoscopic method to your preferred working model.

Toe In

Both cameras point at a single focal point. Though the result is stereoscopic, the vertical parallax introduced by this method can cause discomfort by the audience.

Off Axis

Often regarded as the correct way to create stereo pairs, this is the default method in Fusion. Off Axis introduces no vertical parallax, thus creating less stressful stereo images.

Parallel

The cameras are shifted parallel to each other. Since this is a purely parallel shift, there is no Convergence Distance control. Parallel introduces no vertical parallax, thus creating less stressful stereo images.

Eye Separation

Defines the distance between both stereo cameras. If the Eye Separation is set to a value larger than 0, controls for each camera are shown in the viewer when this node is selected. There is no Convergence Distance control in Parallel mode.

Convergence Distance

This control sets the stereoscopic convergence distance, defined as a point located along the Z-axis of the camera that determines where both left and right eye cameras converge.

Control Visibility

Allows you to selectively activate the onscreen controls that are displayed along with the camera.

- **Frustum:** Displays the actual viewing cone of the camera.
- **View Vector:** Displays a white line inside the viewing cone, which can be used to determine the shift when in Parallel mode.
- **Near Clip:** The Near clipping plane. This plane can be subdivided for better visibility.
- **Far Clip:** The Far clipping plane. This plane can be subdivided for better visibility.

- **Plane of Focus:** The camera focal point according to the Plane of Focus slider explained above. This plane can be subdivided for better visibility.
- **Convergence Distance:** The point of convergence when using Stereo mode. This plane can be subdivided for better visibility.

Common Controls

Settings Tab

The Settings tab in the Inspector is duplicated in other 3D nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Text 3D [3Txt]



The Text 3D node

Text 3D Node Introduction

The Text3D node is a 3D version of the 2D Text+ node. The controls for this node are mostly identical to the controls for the 2D version in almost all respects, except Text 3D supports only one shading element.

The Text 3D node was based on a tool that predates the Fusion 3D environment. So, some of the controls found in the basic primitive shapes and geometry loaders, such as many of the material, lighting, and matte options, are not found in this node’s controls. The Text 3D node has a built-in material, but unlike the other 3D nodes it does not have a material input. The Shading tab contains controls to adjust the diffuse and specular components. To replace this default material with a more advanced material, follow the Text Plus node with a Replace Material 3D node. The Override 3D node can be used to control the lighting, visibility, and matte options for this node.

When network rendering a comp that contains Text 3D nodes, each render machine is required to have the necessary fonts installed or the network rendering fails. Fusion does not share or copy fonts to render slaves.

Inputs

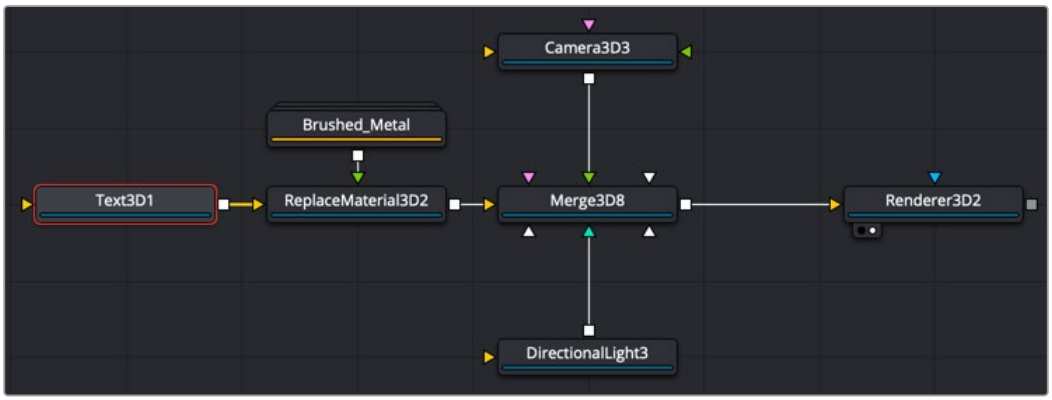
SceneInput: The orange scene input accepts a 3D scene that can be combined with the 3D text created in the node.

ColorImage: The green color image input accepts a 2D image and wraps it around the text as a texture. This input is visible only when Image is selected in the Material Type menu located in the Shading tab.

BevelTexture: The magenta bevel texture input accepts a 2D image and wraps it around the bevel as a texture. This input is visible only when one Material is disabled in the Shader tab and Image is selected in the Bevel Type menu.

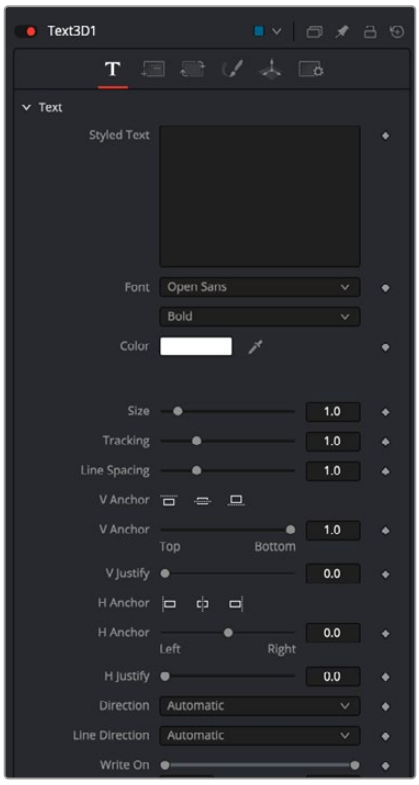
Basic Node Setup

The Text 3D node generates text, so most often this node starts a branch of your node tree. However, to apply more realistic materials, a replace Material node is often added after the Text 3D, prior to connecting into a Merge 3D node.



Text 3D with Replace Material node added to apply alternative material

Inspector



Text 3D controls

Text Tab

The Text 3D text tab in the Inspector is divided into three sections: Text, Extrusion, and Advanced Controls. The Text section includes parameters that are familiar to anyone who has used a word processor. It includes commonly used text formatting options. The Extrusion section includes controls to extrude the text and create beveled edges for the text. The Advanced controls are used for kerning options.

Styled Text

The Edit box in this tab is where the text to be created is entered. Any common character can be typed into this box. The common OS clipboard shortcuts (Command-C or Ctrl-C to copy, Command-X or Ctrl-X to cut, Command-V or Ctrl-V to paste) also work; however, right-clicking on the Edit box displays a custom contextual menu with several modifiers you can add for more animation and formatting options.

Font

Two Font menus are used to select the font family and typeface such as Regular, Bold, and Italic.

Color

This control sets the basic tint color of the text. This is the same Color control displayed in the Material type section of the Shader tab.

Size

This control is used to increase or decrease the size of the text. This is not like selecting a point size in a word processor. The size is relative to the width of the image.

Tracking

The Tracking parameter adjusts the uniform spacing between each character of text.

Line Spacing

Line Spacing adjusts the distance between each line of text. This is sometimes called leading in word-processing applications.

V Anchor

The Vertical Anchor controls consist of three buttons and a slider. The three buttons are used to align the text vertically to the top, middle, or bottom baseline of the text. The slider can be used to customize the alignment. Setting the Vertical Anchor affects how the text is rotated but also the location for line spacing adjustments. This control is most often used when the Layout type is set to Frame in the Layout tab.

V Justify

The Vertical Justify slider allows you to customize the vertical alignment of the text from the V Anchor setting to full justification so it is aligned evenly along the top and bottom edges. This control is most often used when the Layout type is set to Frame in the Layout tab.

H Anchor

The Horizontal Anchor controls consist of three buttons and a slider. The three buttons justify the text alignment to the left edge, middle, or right edge of the text. The slider can be used to customize the justification. Setting the Horizontal Anchor affects how the text is rotated but also the location for tracking (leading) spacing adjustments. This control is most often used when the Layout type is set to Frame in the Layout tab.

H Justify

The Horizontal Justify slider allows you to customize the justification of the text from the H Anchor setting to full justification so it is aligned evenly along the left and right edges. This control is most often used when the Layout type is set to Frame in the Layout tab.

Direction

This menu provides options for determining the direction in which the text is to be written.

Line Direction

These menu options are used to determine the text flow from top to bottom, bottom to top, left to right, or right to left.

Write On

This range control is used to quickly apply simple Write On and Write Off animation to the text. To create a Write On effect, animate the End portion of the control from 1 to 0 over the length of time required. To create a Write Off effect, animate the Start portion of the range control from 0 to 1.

Extrusion Depth

An extrusion of 0 produces completely 2D text. Any value greater than 0 extrudes the text to generate text with depth.



Text 3D extrusion and bevel controls

Bevel Depth

Increase the value of the Bevel Depth slider to bevel the text. The text must have extrusion before this control has any effect.

Bevel Width

Use the Bevel Width control to increase the width of the bevel.

Smoothing Angle

Use this control to adjust the smoothing angle applied to the edges of the bevel.

Front/Back Bevel

Use these checkboxes to enable beveling for the front and back faces of the text separately

Custom Extrusion

In Custom mode, the Smoothing Angle controls the smoothing of normals around the edges of a text character. The spline itself controls the smoothing along the extrusion profile. If a spline segment is smoothed, for example by using the shortcut Shift-S, the normals are smoothed as well. If the control point is linear, the shading edge is sharp. The first and last control points on the spline define the extent of the text.

Custom Extrusion Subdivisions

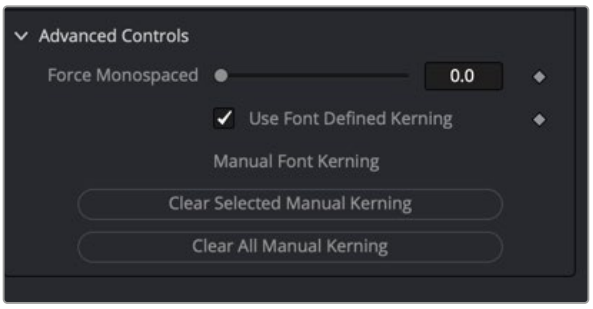
Controls the number of subdivisions within the smoothed portions of the extrusion profile.

TIP: Splines can also be edited from within the Spline Editor panel. It provides a larger working space for working with any spline including the Custom Extrusion.

Extrusion profile spline control: Do not try to go to zero size at the Front/Back face. This results in Z-fighting resulting from self-intersecting faces. To avoid this problem, make sure the first and last point have their profiles set to 0.

Force Monospaced

This slider control can be used to override the kerning (spacing between characters) that is defined in the font. Setting this slider to zero (the default value) causes Fusion to rely entirely on the kerning defined with each character. A value of one causes the spacing between characters to be completely even, or monospaced.



Text 3D Advanced Controls can be used to manually kern letters.

Use Font Defined Kerning

This enables kerning as specified in the True Type font and is on by default.

Manual Font Kerning

Manual Font Kerning is only performed using the Text+ node. To perform manual kerning on Text3D, create the text using the Text+ node and kern it in that tool. Then, right-click over the tool's name in the Inspector and choose Copy. Once the settings are copied, select the Text 3D node and choose Paste Settings from the Inspector's contextual menu. Once the manual kerning is pasted in the Text 3D node, the two buttons in the Inspector clear either the selected character's kerning or all the kerning adjustment in the current text.

Layout Tab

The Layout Tab is used to position the text in one of four different layout types.



Text 3D Layout tab for changing the layout of the text block

Layout Type

This menu selects the layout type for the text.

- **Point:** Point layout is the simplest of the layout modes. Text is arranged around an adjustable center point.
- **Frame:** Frame layout allows you to define a rectangular frame used to align the text. The alignment controls are used to justify the text vertically and horizontally within the boundaries of the frame.
- **Circle:** Circle layout places the text around the curve of a circle or oval. Control is offered over the diameter and width of the circular shape. When the layout is set to this mode, the Alignment controls determine whether the text is positioned along the inside or outside of the circle's edge, and how multiple lines of text are justified.
- **Path:** Path layout allows you to shape your text along the edges of a path. The path can be used simply to add style to the text, or it can be animated using the Position on Path control that appears when this mode is selected.

Center X, Y, and Z

These controls are used to position the center of the layout. For instance, moving the center X, Y, and Z parameters when the layout is set to Frame moves the position of the frame the text is within.

Size

This slider is used to control the scale of the layout element. For instance, increasing size when the layout is set to Frame increases the frame size the text is within.

Width and Height

The Width and Height controls are visible when the Layout mode is set to Circle or Frame. The Width and Height controls are visible only when the Layout mode is set to Frame. They are used to adjust the dimensions and aspect of the Layout element.

Rotation Order

These buttons allow you to select the order in which 3D rotations are applied to the text.

X, Y, and Z

These angle controls can be used to adjust the angle of the Layout element along any axis.

Fit Characters

This menu control is visible only when the Layout type is set to Circle. This menu is used to select how the characters are spaced to fit along the circumference.

Position on Path

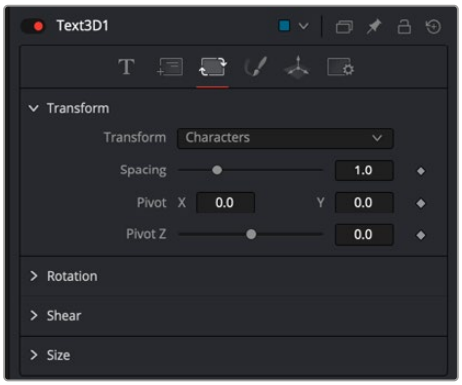
The Position on Path control is used to control the position of the text along the path. Values less than zero or greater than one cause the text to move beyond, continuing in the same direction set by the last two points on the path.

Right-Click Here for Shape Animation

This label appears only when the Layout type is set to Path. It is used to provide access to a contextual menu that provides options for connecting the path to other paths in the node tree, and animating the spline points on the path over time.

Transform Tab

There are actually two Transform tabs in the Text 3D Inspector. The first Transform tab is unique to the Text 3D tool, while the second is the common Transform tab found on many 3D nodes. The Text 3D-specific Transform tab is described below since it contains some unique controls for this node.



Text 3D Transform tab

Transform

This menu determines the portion of the text affected by the transformations applied in this tab. Transformations can be applied to line, word, and character levels simultaneously. This menu is only used to keep the number of visible controls to a reasonable number.

Characters: Each character of text is transformed along its own center axis.

Words: Each word is transformed separately on the word's center axis.

Lines: Each line of the text is transformed separately on that line's center axis.

Spacing

The Spacing slider is used to adjust the amount of space between each line, word, or character. Values less than one usually cause the characters to begin overlapping.

Pivot X, Y, and Z

This provides control over the exact position of the axis. By default, the axis is positioned at the calculated center of the line, word, or character. The pivot control works as an offset, such that a value of 0.1, 0.1 in this control would cause the axis to be shifted downward and to the right for each of the text elements. Positive values in the Z-axis slider move the axis further along the axis (away from the viewer). Negative values bring the axis of rotation closer.

Rotation Order

These buttons are used to determine the order in which transforms are applied. X, Y, and Z would mean that the rotation is applied to X, then Y, and then Z.

X, Y, and Z

These controls can be used to adjust the angle of the text elements in any of the three dimensions.

Shear X and Y

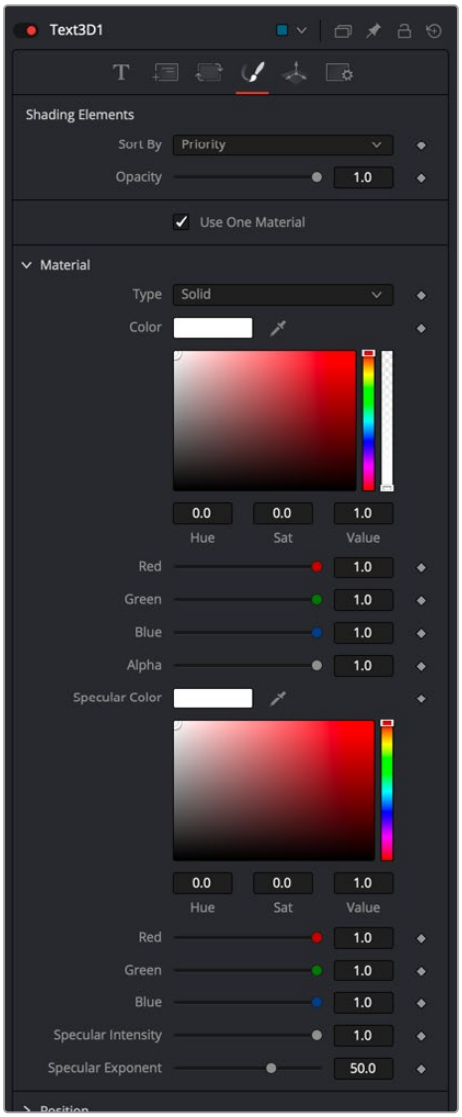
Adjust these sliders to modify the slanting of the text elements along the X- and Y-axis.

Size X and Y

Adjust these sliders to modify the size of the text elements along the X- and Y-axis.

Shading

The Shading tab for the Text 3D node controls the overall appearance of the text and how lights affect its surface.



Text 3D Shading tab

Opacity

Reducing the material's opacity decreases the color and Alpha values of the specular and diffuse colors equally, making the material transparent and allowing hidden objects to be seen through the material.

Use One Material

Deselecting this option reveals a second set of Material controls for the beveled edge of the text.

Type

To use a solid color texture, select the Solid mode. Selecting the Image mode reveals a new external input on the node that can be connected to another 2D image.

Specular Color

Specular Color determines the color of light that reflects from a shiny surface. The more specular a material is, the glossier it appears. Surfaces like plastics and glass tend to have white specular highlights, whereas metallic surfaces like gold have specular highlights that tend to inherit their color from the material color. The basic shader material does not provide an input for textures to control the specularity of the object. Use nodes from the 3D Material category when more precise control is required over the specular appearance.

Specular Intensity

Specular Intensity controls the strength of the specular highlight. If the specular intensity texture port has a valid input, then this value is multiplied by the Alpha value of the input.

Specular Exponent

Specular Exponent controls the falloff of the specular highlight. The greater the value, the sharper the falloff, and the smoother and glossier the material appears. The basic shader material does not provide an input for textures to control the specular exponent of the object. Use nodes from the 3D Material category when more precise control is required over the specular exponent.

Image Source

This control determines the source of the texture applied to the material. If the option is set to Tool, then an input appears on the node that can be used to apply the output of a 2D node as the texture. Selecting Clip opens a file browser that can be used to select an image or image sequence from disk. The Brush option provides a list of clips found in the Fusion\brushes folder.

Bevel Material

This option appears only when the Use One Material checkbox control is selected. The controls under this option are an exact copy of the Material controls above but are applied only to the beveled edge of the text.

Position, Rotation, Shear, and Size

These transform controls act similarly to the transform controls in the Transform tab when a single shading element is enabled from the top of the Shading tab. However, when two or more shading elements are enabled, these transform controls are applied to the currently selected shading element. This allows you to independently control the position, rotation, shearing, and size of borders, fill colors, and shadows.

Uncapped 3D Text To hide the front face of extruded text, uncheck Use One Material on the Shading tab and reduce the first material's color to black, including its Alpha value.

Common Controls

Transform and Settings Tabs

The Transform and Settings tabs in the Inspector are duplicated in other 3D nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Text 3D Modifiers

Right-clicking within the Styled Text box displays a menu with the following text modifiers. Only one modifier can be applied to a Text 3D Styled Text box. Below is a brief list of the text specific modifiers.

> For more information *see Chapter 64, "Modifiers,"* in the Fusion Reference Manual.

Animate

Use the Animate command to set to a keyframe on the entered text and animate the content over time.

Character Level Styling

The Text 3D node doesn't support Character Level Styling directly. However, you can create a Text+ node first and modify its text field with a Character Level Styling modifier. Then either connect the Text 3D's text field to the modifier that is now available or copy the Text+ node and paste its settings to the Text 3D node (right-click > Paste Settings)

Comp Name

Comp Name puts the name of the composition in the Styled Text box and is generally used as a quick way to create slates.

Follower

Follower is a text modifier that can be used to ripple animation applied to the text across each character in the text. See "Text Modifiers" at the end of this chapter.

Publish

Publish the text for connection to other text nodes.

Text Scramble

A text modifier ID is used to randomize the characters in the text. See "Text Modifiers" at the end of this chapter.

Text Timer

A text modifier is used to count down from a specified time or to output the current date and time. See "Text Modifiers" at the end of this chapter.

Time Code

A text modifier is used to output Time Code for the current frame. See "Text Modifiers" at the end of this chapter.

Connect To

Use this option to connect the text generated by this Text node to the published output of another node.

Transform 3D [3Xf]



The Transform 3D node

Transform 3D Node Introduction

The Transform 3D node can be used to translate, rotate, or scale all the elements within a scene without requiring a Merge 3D node. This can be useful for hierarchical transformations or for offsetting objects that are merged into a scene multiple times. Its controls are identical to those found in other 3D nodes' Transformation tabs.

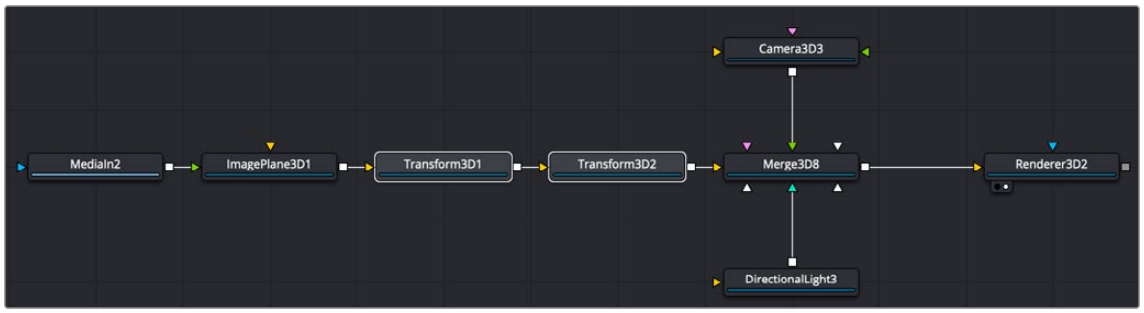
Inputs

The Transform node has a single required input for a 3D scene or 3D object.

Scene Input: The orange scene input is connected to a 3D scene or 3D object to apply a second set of transformation controls.

Basic Node Setup

The Transform 3D node adds 3D position, rotation, and pivot control onto any existing transforms in the 3D node prior to it. You can combine multiple Transform 3D nodes together to build parenting or hierarchical movement.



Transform 3D nodes strung together to create a parenting hierarchy

Inspector



Transform 3D controls

Controls Tab

The Controls tab is the primary tab for the Transform 3D node. It includes controls to translate, rotate, or scale all elements within a scene without requiring a Merge 3D node.

Translation

- **X, Y, Z Offset:** Controls are used to position the 3D element in 3D space.

Rotation

- **Rotation Order:** Use these buttons to select the order used to apply the rotation along each axis of the object. For example, XYZ would apply the rotation to the X-axis first, followed by the Y-axis, and then the Z-axis.
- **X, Y, Z Rotation:** Use these controls to rotate the object around its pivot point. If the Use Target checkbox is selected, then the rotation is relative to the position of the target; otherwise, the global axis is used.

Pivot Controls

- **X, Y, Z Pivot:** A pivot point is the point around which an object rotates. Normally, an object rotates around its own center, which is considered to be a pivot of 0,0,0. These controls can be used to offset the pivot from the center.

Scale

- **X, Y, Z Scale:** If the Lock X/Y/Z checkbox is checked, a single scale slider is shown. This adjusts the overall size of the object. If the Lock checkbox is unchecked, individual X, Y, and Z sliders are displayed to allow scaling in any dimension.

NOTE: If the Lock checkbox is checked, scaling of individual dimensions is not possible, even when dragging specific axes of the Transformation widget in Scale mode.

Use Target

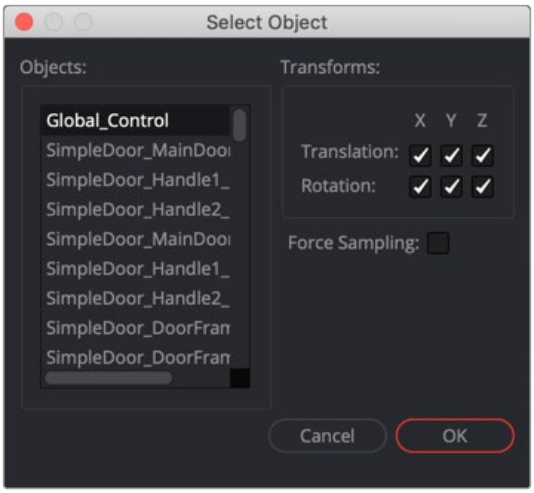
Selecting the Use Target checkbox enables a set of controls for positioning an XYZ target. When Use Target is enabled, the object always rotates to face the target. The rotation of the object becomes relative to the target.

Import Transform

Opens a file browser where you can select a scene file saved or exported by your 3D application. It supports the following file types:

LightWave Scene	`.lws`
Max Scene	`.ase`
Maya Ascii Scene	`.ma`
dotXSI	`.xsi`

The Import Transform button imports only transformation data. For 3D geometry, lights and cameras, consider using the File > FBX Import option from the menus.



Import Transform Browser

Onscreen Transformation Controls

Onscreen Transformation controls provide an alternative way of using the controls in the Inspector. The viewer includes modes for transformation, rotation, and scaling. To change the mode of the onscreen controls, select one of the three buttons in the toolbar along the side of the viewer. The modes can also be toggled using the keyboard shortcut Q for translation, W for rotation, and E for scaling. In all three modes, an individual axis of the control may be dragged to affect just that axis, or the center of the control may be dragged to affect all three axes.

The scale sliders for most 3D nodes default to locked, which causes uniform scaling of all three axes. Unlock the Lock X/Y/Z Scale checkbox to scale an object on a single axis only.



Transform 3D onscreen transformation controls

Triangulate 3D [3Tri]



The Triangulate 3D node

Triangulate 3D Node Introduction

The Triangulate 3D node is a unique node in that it has no controls. This node turns polygon shapes into triangles. For instance, a quad that is four points becomes two triangles. It is used to convert complex polygon shapes into a mesh for easier processing.

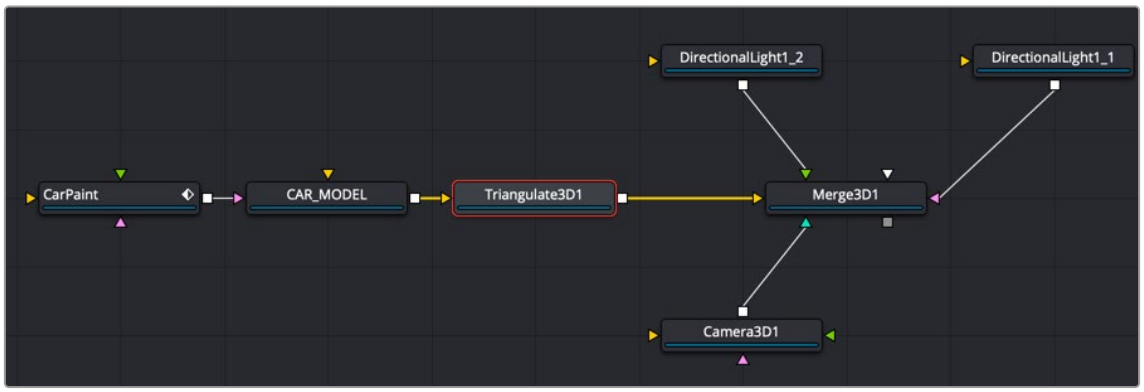
Inputs

The Triangulate 3D node has a single required input for a 3D scene or 3D object.

Scene Input: The orange scene input is connected to the 3D scene or 3D object you want to triangulate.

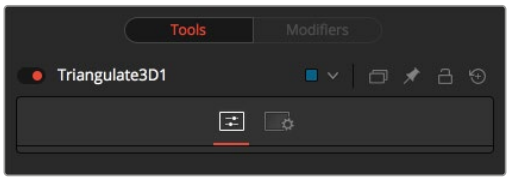
Basic Node Setup

The Triangulate 3D node is placed after the geometry you want to triangulate.



Triangulate 3D nodes connected after an imported model

Inspector



Triangulate 3D controls

Controls Tab

There are no controls for this node.

Common Controls

Settings Tab

The Settings tab in the Inspector is duplicated in other 3D nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

UV Map 3D [3UV]



The UVMap 3D node

UV Map 3D Node Introduction

The UV Map 3D node replaces the UV texture coordinates on the geometry in the scene. These coordinates tell Fusion how to apply a texture to an object. While it is possible to adjust the global properties of the selected mapping mode, it is not possible to manipulate the UV coordinates of individual vertices directly from within Fusion. The onscreen controls drawn in the viewers are for reference only and cannot be manipulated.

Camera Projections with UV Map 3D

The Camera Mapping mode makes it possible to project texture coordinates onto geometry through a camera. Once you select Camera from the Mapping mode menu, then connect the Camera 3D node that you want to use to create the UV coordinates.

NOTE: That this does not directly project an image through the camera. The image to be projected should be connected to the diffuse texture input of whatever material is assigned to the objects. When the texture is applied, it uses the UV coordinates created by the camera. Because this is a texture projection and not light, the Alpha channel of the texture correctly sets the opacity of the geometry.

See the Camera 3D and Projector 3D nodes for alternate approaches to projection.

The projection can optionally be locked to the vertices as it appears on a selected frame.

This fails if the number of vertices in the mesh changes over time, as Fusion must be able to match up the mesh at the reference time and the current time. To be more specific, vertices may not be created or destroyed or reordered. So, projection locking does not work for many particle systems, or for primitives with animated subdivisions, or with duplicate nodes using non-zero time offsets.

NOTE: The UV Map 3D node does not put a texture or material on the mesh; it only modifies the texture coordinates that the materials use. This may be confusing because the material usually sits upstream, as seen in the Basic Node Setup example below.

Inputs

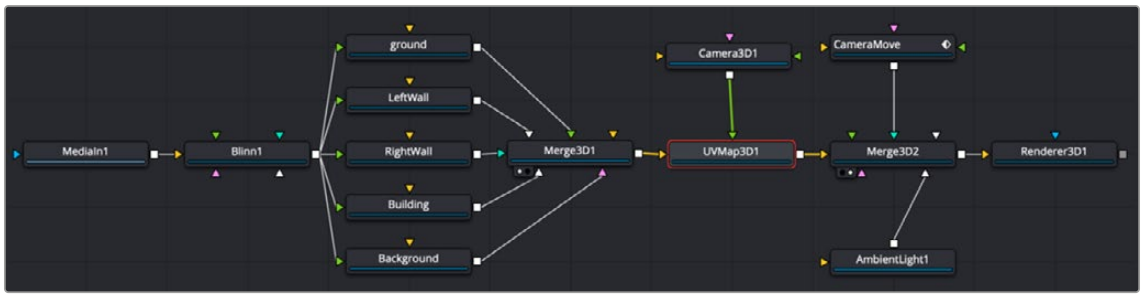
The UV Map 3D node has two inputs: one for a 3D scene or 3D object and another optional input for a Camera 3D node.

Scene Input: The orange scene input is connected to the 3D scene or 3D object you want to triangulate.

CameraInput: This input expects the output of the Camera 3D node. It is only visible when the Camera Map mode menu is set to Camera.

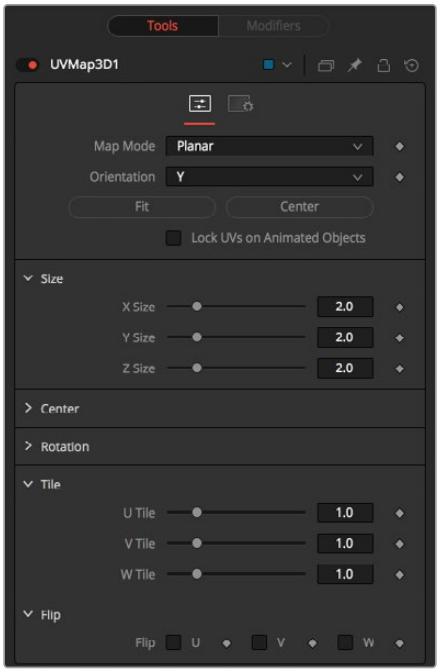
Basic Node Setup

The UV Map 3D node is placed after all the geometry and set to Camera Map. Connecting a camera to the UV map allows you to line up the texture based on a centered camera position and 3D geometry.



UV Map 3D is placed after the Merge 3D, with a camera connected to line up the texture

Inspector



UV Map 3D controls

Controls Tab

The UV Map 3D Controls tab allows you to select Planar, Cylindrical, Spherical, XYZ, and Cubic mapping modes, which can be applied to basic Fusion primitives as well as imported geometry. The position, rotation, and scale of the texture coordinates can be adjusted to allow fine control over the

texture's appearance. An option is also provided to lock the UV produced by this node to animated geometry according to a reference frame. This can be used to ensure that textures applied to animated geometry do not slide.

Map Mode

The Map mode menu is used to define how the texture coordinates are created. You can think of this menu as a way to select the virtual geometry that projects the UV space on the object.

- **Planar:** Creates the UV coordinates using a plane.
- **Cylindrical:** Creates the UV coordinates using a cylindrical-shaped object.
- **Spherical:** The UVs are created using a sphere.
- **XYZ to UVW:** The position coordinates of the vertices are converted to UVW coordinates directly. This is used for working with procedural textures.
- **CubeMap:** The UVs are created using a cube.
- **Camera:** Enables the Camera input on the node. After connecting a camera to the node, the texture coordinates are created based on camera projection.

Orientation X/Y/Z

Defines the reference axis for aligning the Map mode.

Fit

Clicking this button fits the Map mode to the bounding box of the input scene.

Center

Clicking this button moves the center of the Map mode to the bounding box center of the input scene.

Lock UVs on Animated Objects

If the object is animated, the UVs can be locked to it by enabling this option. The option also reveals the Ref Time slider, where it is possible to choose a reference frame for the UV mapping. Using this feature, it is not required to animate the UV map parameters. It is enough to set up the UV map at the reference time.

Size X/Y/Z

Defines the size of the projection object.

Center X/Y/Z

Defines the position of the projection object.

Rotation/Rotation Order

Use these buttons to select which order is used to apply the rotation along each axis of the object. For example, XYZ would apply the rotation to the X-axis first, followed by the Y-axis, and then the Z-axis.

Rotation X/Y/Z

Sets the orientation of the projection object for each axis, independent from the rotation order.

Tile U/V/W

Defines how often a texture fits into the projected UV space on the applicable axis. Note that the UVW coordinates are transformed, not a texture. This works best when used in conjunction with the Create Texture node.

Flip U/V/W

Mirrors the texture coordinates around the applicable axis.

Flip Faces (Cube Map Mode Only)

Mirrors the texture coordinates on the individual faces of the cube.

NOTE: To utilize the full capabilities of the UV Map 3D node, it helps to have a basic understanding of how 2D images are mapped onto 3D geometry. When a 2D image is applied to a 3D surface, it is converted into a texture map that uses UV coordinates to determine how the image translates to the object. Each vertex on a mesh has a (U, V) texture coordinate pair that describes the appearance the object takes when it is unwrapped and flattened. Different mapping modes use different methods for working out how the vertices transform into a flat 2D texture. When using the UV Map 3D node to modify the texture coordinates on a mesh, it's best to do so using the default coordinate system of the mesh or primitive. So the typical workflow would look like Shape 3D > UV Map 3D > Transform 3D. The Transformation tab on the Shape node would be left to its default values, and the Transform 3D node following the UV Map 3D does any adjustments needed to place the node in the scene. Modifying/animating the transform of the Shape node causes the texture to slide across the shape, which is generally undesirable. The UV Map 3D node modifies texture coordinates per vertex and not per pixel. If the geometry the UV map is applied to is poorly tessellated, then undesirable artifacts may appear.

Common Controls

Settings Tab

The Settings tab in the Inspector is duplicated in other 3D nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Weld 3D [3We]



The Weld 3D node

Weld 3D Node Introduction

Sometimes 3D geometry has vertices that should have been joined when the geometry was created, but for one reason or another they are not joined. This can cause artifacts, especially when the two vertices have different normals.

For example, you may find:

- The different normals produce hard shading/lighting edges where none were intended.
- If you try to Displace 3D the vertices along their normals, they crack.
- Missing pixels or doubled-up pixels in the rendered image.
- Particles pass through the tiny invisible cracks.

Instead of round tripping back to your 3D modeling application to fix the “duplicated” vertices, the Weld 3D node allows you to do this in Fusion. Weld 3D welds together vertices with the same or nearly the same positions. This can be used to fix cracking issues when vertices are displaced by welding the geometry before the Displace. There are no user controls to pick vertices. Currently, this node welds together just position vertices; it does not weld normals, texcoords, or any other vertex stream. So, although the positions of two vertices have been made the same, their normals still have their old values. This can lead to hard edges in certain situations.

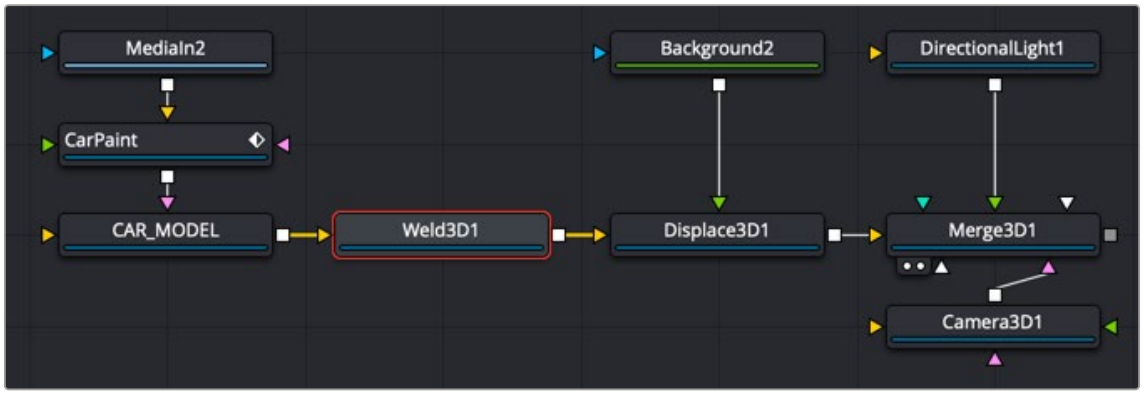
Inputs

The Weld 3D node has a single input for a 3D scene or 3D object you want to repair.

Scene Input: The orange scene input is connected to the 3D scene or 3D object you want to fix.

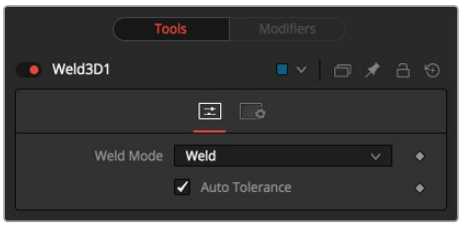
Basic Node Setup

The Weld 3D node is placed after the geometry that has duplicate vertices problems. Sometimes problems are exposed when displacing the geometry. In that case, placing the weld after the geometry but before the Displace 3D can repair the issues.



Weld 3D is placed after the 3D geometry that needs repair

Inspector



Weld 3D controls

Controls Tab

The Controls tab for the Weld 3D node includes a simple Weld Mode menu. You can choose between welding vertices or fracturing them.

Fracture

Fracturing is the opposite of welding, so all vertices are unwelded. This means that all polygon adjacency information is lost. For example, an Image Plane 3D normally consists of connected quads that share vertices. Fracturing the image plane causes it to become a bunch of unconnected quads.

Tolerance

In auto mode, the Tolerance value is automatically detected. This should work in most cases. It can also be adjusted manually if needed.

USAGE Use Weld 3D when issues occur with the geometry. Don't use it everywhere just because it's there, as it influences render time.

Weld 3D is intended to be used as a mesh robustness tool and not as a mesh editing tool to merge vertices. If you can see the gap between the vertices you want to weld in the 3D view, you are probably misusing Weld 3D. Unexpected things may happen when you do this; do so at your own peril.

LIMITATIONS Setting the tolerance too large can cause edges/faces to collapse to points.

If your model has detail distributed over several orders of scale, picking a tolerance value can be hard or impossible.

For example, suppose you have a model of the International Space Station and there are lots of big polygons and lots of really tiny polygons. If you set the tolerance too large, small polygons that shouldn't merge do; if you set the tolerance too small, some large polygons won't be merged.

Vertices that are far from the origin can fail to merge correctly. This is because bignumber + epsilon can exactly equal bignumber in float math. This is one reason it may be best to merge in local coordinates and not in world coordinates.

Sometimes Weld 3D-ing a mesh can make things worse. Take Fusion's cone, for instance. The top vertex of the cone is currently duplicated for each adjoining face, and they all have different normals. If you weld the cone, the top vertices merge and only have one normal, making the lighting look weird.

Weld 3D is not multithreaded.

WARNING Do not misuse Weld 3D to simplify (reduce the polygon count of) meshes. It is designed to efficiently weld vertices that differ by only very small values, like a 0.001 distance.

Common Controls

Settings Tab

The Settings tab in the Inspector is duplicated in other 3D nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Modifier

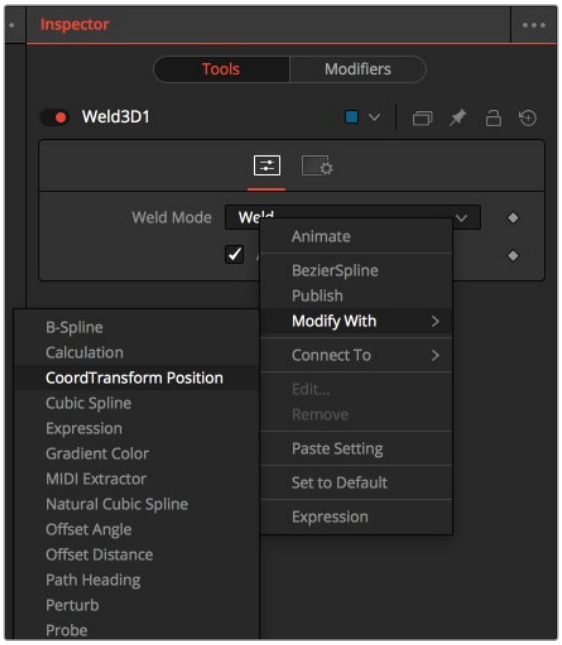
Coordinate Transform 3D

Because of the hierarchical nature of the Fusion 3D node tree, the original position of an object in the 3D scene often fails to indicate the current position of the object. For example, an image plane might initially have a position at 1, 2, 1, but then be scaled, offset, and rotated by other nodes further downstream in the 3D scene, ending up with an absolute location of 10, 20, 5.

This can complicate connecting an object further downstream in the composition directly to the position of an upstream object. The Coordinate Transform modifier can be added to any set of XYZ coordinate controls and calculate the current position of a given object at any point in the scene hierarchy.

To add a Coordinate Transform modifier, simply right-click a number field on any node and select Modify With/CoordTransform Position from the Controls' contextual menu.

Inspector



Weld 3D modifier tools

Target Object

This control should be connected to the 3D node that produces the original coordinates to be transformed. To connect a node, drag and drop a node from the node tree into the Text Edit control, or right-click the control and select the node from the contextual menu. It is also possible to type the node's name directly into the control.

Sub ID

The Sub ID slider can be used to target an individual sub-element of certain types of geometry, such as an individual character produced by a Text 3D node or a specific copy created by a Duplicate 3D node.

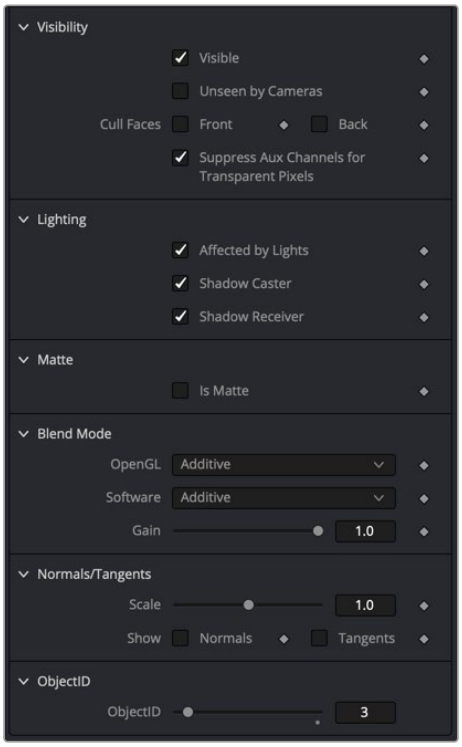
Scene Input

This control should be connected to the 3D node that outputs the scene containing the object at the new location. To connect a node, drag and drop a node from the node tree into the Text Edit control, or right-click the control and select an object from the Connect To submenu.

The Common Controls

Nodes that handle 3D geometry share several identical controls in the Inspector. This section describes controls that are common among 3D nodes.

Common Controls Tab



Common Controls 3D tab

These controls are often displayed in the lower half of the Controls tab. They appear in nodes that create or contain 3D geometry.

Visibility

- **Visible:** If this option is enabled, the object is visible in the viewers and in final renders. When disabled, the object is not visible in the viewers nor is it rendered into the output image by the Renderer 3D node. Also, a non-visible object does not cast shadows.
- **Unseen by Cameras:** When the Unseen by Cameras checkbox is enabled, the object is visible in the viewers (unless the Visible checkbox is disabled), except when viewed through a camera. Also, the object is not rendered into the output image by the Renderer 3D node. However, shadows cast by an unseen object are still visible when rendered by the software renderer in the Renderer 3D node, though not by the OpenGL renderer.
- **Cull Front Face/Back Face:** Use these options to eliminate rendering and display of certain polygons in the geometry. If Cull Back Face is selected, polygons facing away from the camera are not rendered and do not cast shadows. If Cull Front Face is selected, polygons facing toward the camera are not rendered and do not cast shadows. Enabling both options has the same effect as disabling the Visible checkbox.
- **Suppress Aux Channels for Transparent Pixels:** In previous versions of Fusion, transparent pixels were excluded by the software and Open GL render options in the Renderer 3D node. To be more specific, the software renderer excluded pixels with R,G,B,A set to 0, and the GL renderer

excluded pixels with A set to 0. This is now optional. The reason you might want to do this is to get aux channels (e.g., Normals, Z, UVs) for the transparent areas. For example, suppose you want to replace the texture on a 3D element that is transparent in certain areas with a texture that is transparent in different areas. It would then be useful to have transparent areas set aux channels (particularly UVs). As another example, suppose you are adding depth of field. You probably do not want the Z-channel to be set on transparent areas, as this gives you a false depth. Also, keep in mind that the exclusion is based on the final pixel color including lighting, if it is on. So, if you have a specular highlight on a clear glass material, this checkbox does not affect it.

Lighting

- **Affected by Lights:** Disabling this checkbox causes lights in the scene to not affect the object. The object does not receive nor cast shadows, and it is shown at the full brightness of its color, texture, or material.
- **Shadow Caster:** Disabling this checkbox causes the object not to cast shadows on other objects in the scene.
- **Shadow Receiver:** Disabling this checkbox causes the object not to receive shadows cast by other objects in the scene.

Matte

Enabling the Is Matte option applies a special texture, causing the object to not only become invisible to the camera, but also making everything that appears directly behind the camera invisible as well. This option overrides all textures.

> For more information on Fog 3D and Soft Clipping, *see Chapter 25, "3D Compositing Basics,"* in the Fusion Reference Manual.

- **Is Matte:** When activated, objects whose pixels fall behind the matte object's pixels in Z do not get rendered. Two additional options are displayed when the Is Matte checkbox is activated.
- **Opaque Alpha:** When the Is Matte checkbox is enabled, the Opaque Alpha checkbox sets the Alpha value of the matte object to 1.
- **Infinite Z:** This option sets the value in the Z-channel to infinite. This checkbox is visible only when the Is Matte option is enabled.

Blend Mode

A Blend mode specifies which method is used by the renderer when combining this object with the rest of the scene. The blend modes are essentially identical to those listed in the section for the 2D Merge node. For a detailed explanation of each mode, see the section for that node.

The blending modes were originally designed for use with 2D images. Using them in a lit 3D environment can produce undesirable results. For best results, use the Apply modes in unlit 3D scenes using the software option in the Renderer 3D node.

- **OpenGL Blend Mode:** Use this menu to select the blending mode that is used when the geometry is processed by the OpenGL renderer in the Renderer 3D node. This is also the mode used when viewing the object in the viewers. Currently the OpenGL renderer supports a limited number of blending modes.
- **Software Blend Mode:** Use this menu to select the blending mode that is used when the geometry is processed by the software renderer. Currently, the software renderer supports all the modes described in the Merge node documentation, except for the Dissolve mode.

Normal/Tangents

Normals are imaginary lines perpendicular to each point on the surface of an object. They are used to illustrate the exact direction and orientation of every polygon on 3D geometry. Knowing the direction and orientation determines how the object gets shaded. Tangents are lines that exists along the surface's plane. These lines are tangent to a point on the surface. The tangent lines are used to describe the direction of textures you apply to the surface of 3D geometry.

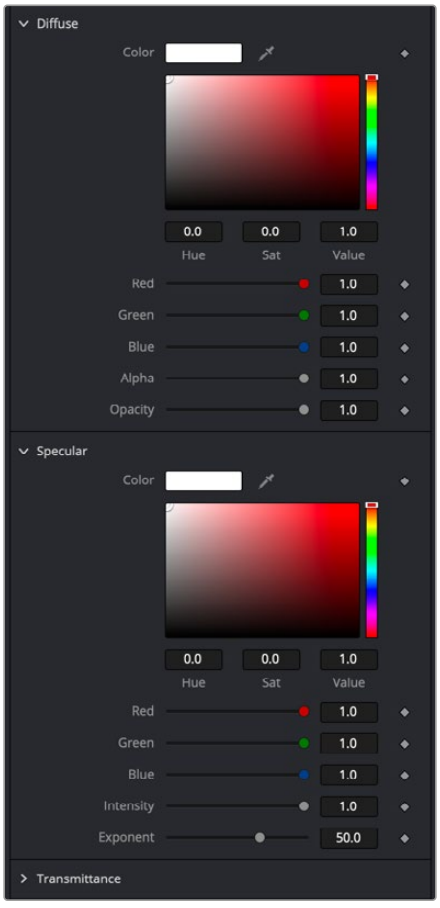
- **Scale:** This slider increases or decreases the length of the vectors for both normals and tangents.
- **Show Normals:** Displays blue vectors typically extending outside the surface of the geometry. These normal vectors help indicate how different areas of the surface are illuminated based on the angle at which the light hits it.
- **Show Tangents:** Displays green vectors for Y and red vectors of X. The X and Y vectors represent the direction of the image or texture you are applying to the geometry.

Object ID

Use this slider to select which ID is used to create a mask from the object of an image. Use the Sample button in the same way as the Color Picker to grab IDs from the image displayed in the viewer. The image or sequence must have been rendered from a 3D software package with those channels included.

Common Materials Tab

The controls in the Materials tab are used to determine the appearance of the 3D object when lit. Most of these controls directly affect how the object interacts with light using a basic shader. For more advanced control over the objects appearance, you can use tools from the 3D Materials category of the Effects Library. These tools can be used to assemble a more finely detailed and precise shader.



Common Materials 3D tab

When a shader is constructed using the 3D Material tools and connected to the 3D Object's material input, the controls in this tab are replaced by a label that indicates that an external material is currently in use.

Diffuse

Diffuse describes the base surface characteristics without any additional effects like reflections or specular highlights.

Diffuse Color

The Diffuse Color determines the basic color of an object when the surface of that object is either lit indirectly or lit by an ambient light. If a valid image is provided to the tools diffuse texture input, then the RGB values provided here are also multiplied by the color values of the pixels in the diffuse texture. The Alpha channel of the diffuse material can be used to control the transparency of the surface.

Alpha

This slider sets the material's Alpha channel value. This affects diffuse and specular colors equally, and affects the Alpha value of the material in the rendered output. If the tools diffuse texture input is used, then the Alpha value provided here is multiplied by the Alpha channel of the pixels in the image.

Opacity

Reducing the material's Opacity decreases the color and Alpha values of the specular and diffuse colors equally, making the material transparent and allowing hidden objects to be seen through the material.

Specular

The Specular section provides controls for determining the characteristics of light that reflects toward the viewer. These controls affect the appearance of the specular highlight that appears on the surface of the object.

Specular Color

Specular Color determines the color of light that reflects from a shiny surface. The more specular a material is, the glossier it appears. Surfaces like plastics and glass tend to have white specular highlights, whereas metallic surfaces like gold have specular highlights that tend to inherit their color from the material color. The basic shader material does not provide an input for textures to control the specularity of the object. Use tools from the 3D Material category when more precise control is required over the specular appearance.

Specular Intensity

Specular Intensity controls how strong the specular highlight is. If the specular intensity texture input has a valid connection, then this value is multiplied by the Alpha value of the input.

Specular Exponent

Specular Exponent controls the falloff of the specular highlight. The greater the value, the sharper the falloff, and the smoother and glossier the material appears. The basic shader material does not provide an input for textures to control the specular exponent of the object. Use tools from the 3D Material category when more precise control is required over the specular exponent.

Transmittance

Transmittance controls the way light passes through a material. For example, a solid blue sphere casts a black shadow, but one made of translucent blue plastic would cast a much lower density blue shadow.

There is a separate opacity option. Opacity determines how transparent the actual surface is when it is rendered. Fusion allows adjusting both opacity and transmittance separately. This might be a bit counter-intuitive to artists who are unfamiliar with 3D software at first. It is possible to have a surface that is fully opaque but transmits 100% of the light arriving upon it, effectively making it a luminous/emissive surface.

Attenuation

Attenuation determines how much color is transmitted through the object. For an object to have transmissive shadows, set the attenuation to (1, 1, 1), which means 100% of green, blue, red light passes through the object. Setting this color to RGB (1, 0, 0) means that the material transmits 100% of the red arriving at the surface but none of the green or blue light. This allows “stained glass” shadows.

Alpha Detail

When the Alpha Detail slider is set to 0, the Alpha channel of the object is ignored and the entire object casts a shadow. If it is set to 1, the Alpha channel determines what portions of the object cast a shadow.

Color Detail

The Color Detail slider modulates light passing through the surface by the diffuse color + texture colors. Use this to throw a shadow that contains color details of the texture applied to the object. Increasing the slider from 0 to 1 brings in more of diffuse color + texture color into the shadow. Note that the Alpha and opacity of the object are ignored when transmitting color, allowing an object with a solid Alpha to still transmit its color to the shadow.

Saturation

The Saturation slider controls the saturation of the color component transmitted to the shadow. Setting this to 0.0 results in monochrome shadows.

Receives Lighting/Shadows

These checkboxes control whether the material is affected by lighting and shadows in the scene. If turned off, the object is always fully lit and/or unshadowed.

Two-Sided Lighting

This makes the surface effectively two-sided by adding a second set of normals facing the opposite direction on the back side of the surface. This is normally off, to increase rendering speed, but can be turned on for 2D surfaces or for objects that are not fully enclosed, to allow the reverse or interior surfaces to be visible as well.

Normally, in a 3D application, only the front face of a surface is visible and the back face is culled, so that if a camera were to revolve around a plane in a 3D application, when it reached the backside, the plane would become invisible. Making a plane two sided in a 3D application is equivalent to adding another plane on top of the first but rotated by 180 degrees so the normals are facing the opposite direction on the backside. Thus, when you revolve around the back, you see the second image plane that has its normals facing the opposite way.

Fusion does exactly the same thing as 3D applications when you make a surface two sided. The confusion about what two-sided lighting does arises because Fusion does not cull backfacing polygons by default. If you revolve around a one-sided plane in Fusion, you still see it from the backside (but you are seeing the frontside bits duplicated through to the backside as if it were transparent). Making the plane two sided effectively adds a second set of normals to the backside of the plane.

Note that this can become rather confusing once you make the surface transparent, as the same rules still apply and produce a result that is counterintuitive. If you view from the frontside a transparent two-sided surface illuminated from the backside, it looks unlit.

Material ID

This control is used to set the numeric identifier assigned to this material. The Material ID is an integer number that is rendered into the MatID auxiliary channel of the rendered image when the Material ID option is enabled in the Renderer 3D tool.

> For more information, *see Chapter 25, "3D Compositing Basics,"* in the Fusion Reference Manual.

Common Transform Tab



Common Transform 3D tab

Many tools in the 3D category include a Transform tab used to position, rotate, and scale the object in 3D space.

Translation

X, Y, Z Offset

These controls can be used to position the 3D element.

Rotation

Rotation Order

Use these buttons to select which order is used to apply rotation along each axis of the object. For example, XYZ would apply the rotation to the X axis first, followed by the Y axis and then finally the Z axis.

X, Y, Z Rotation

Use these controls to rotate the object around its pivot point. If the Use Target checkbox is selected, then the rotation is relative to the position of the target; otherwise, the global axis is used.

Pivot

X, Y, Z Pivot

A Pivot point is the point around which an object rotates. Normally, an object rotates around its own center, which is considered to be a pivot of 0,0,0. These controls can be used to offset the pivot from the center.

Scale

X, Y, Z Scale

If the Lock X/Y/Z checkbox is checked, a single Scale slider is shown. This adjusts the overall size of the object. If the Lock checkbox is unchecked, individual X, Y, and Z sliders are displayed to allow individual scaling in each dimension. Note: If the Lock checkbox is checked, scaling of individual dimensions is not possible, even when dragging specific axes of the Transformation Widget in scale mode.

Use Target

Selecting the Use Target checkbox enables a set of controls for positioning an XYZ target. When target is enabled, the object always rotates to face the target. The rotation of the object becomes relative to the target.

Import Transform

Opens a file browser where you can select a scene file saved or exported by your 3D application. It supports the following file types:

LightWave Scene	`.lws`
Max Scene	`.ase`
Maya Ascii Scene	`.ma`
dotXSI	`.xsi`

The Import Transform button imports only transformation data. For 3D geometry, lights, and cameras, consider using the File > FBX Import option.

Onscreen Transformation Controls

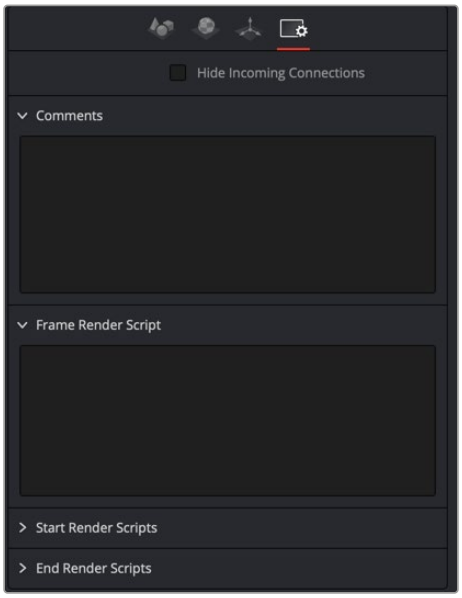


Viewer Transform controls

Most of the controls in the Transform tab are represented in the viewer with onscreen controls for transformation, rotation, and scaling. To change the mode of the onscreen controls, select one of the three buttons in the toolbar in the upper left of the viewer. The modes can also be toggled using the keyboard shortcut Q for translation, W for rotation, and E for scaling. In all three modes, individual axes of the control may be dragged to affect just that axis, or the center of the control may be dragged to affect all three axes.

The scale sliders for most 3D tools default to locked, which causes uniform scaling of all three axes. Unlock the Lock X/Y/Z Scale checkbox to scale an object on a single axis only.

Settings Tab



Common Settings 3D controls

The Common Settings tab can be found on most tools in Fusion. The following controls are specific settings for 3D nodes.

Hide Incoming Connections

Enabling this checkbox can hide connection lines from incoming nodes, making a node tree appear cleaner and easier to read. When enabled, empty fields for each input on a node are displayed in the Inspector. Dragging a connected node from the node tree into the field hides that incoming connection line as long as the node is not selected in the node tree. When the node is selected in the node tree, the line reappears.

Comment Tab

The Comment tab contains a single text control that is used to add comments and notes to the tool. When a note is added to a tool, a small red dot icon appears next to the setting's tab icon and a text bubble appears on the node. To see the note in the Node Editor, hold the mouse pointer over the node for a moment. The contents of the Comments tab can be animated over time, if required.

Scripting Tab

The Scripting tab is present on every tool in Fusion. It contains several edit boxes used to add scripts that process when the tool is rendering. For more details on the contents of this tab, please consult the scripting documentation.

Chapter 30

3D Light Nodes

This chapter details the 3D Light nodes available when creating 3D composites in Fusion. The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Ambient Light [3AL] 820
Directional Light [3DL] 822
DomeLight [3Do] 825
Point Light [3PL] 827
Spot Light [3SL] 830
The Common Controls 834

Ambient Light [3AL]



The Ambient Light node

Ambient Light Node Introduction

An Ambient Light is a directionless light that globally illuminates a scene. It has no real position or rotation, although an onscreen control appears in the viewer to indicate that a light is present in the scene. Position controls for the viewer are provided to move the widget out of the way of other geometry, if necessary.

Similar to a Camera 3D, you connect lights into a Merge 3D and view them in the scene by viewing the Merge 3D node. Selecting a light node and loading it into the viewer does not show anything.

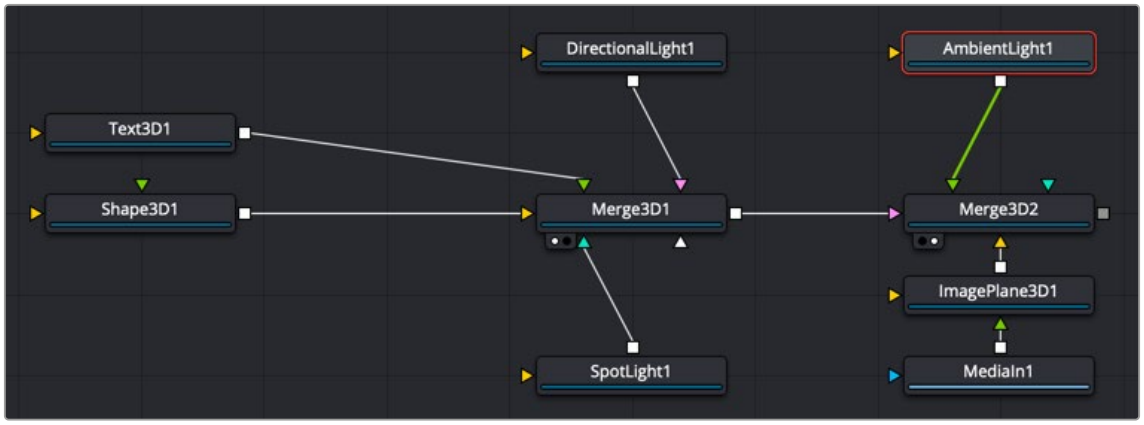
Inputs

The Ambient Light node includes a single optional orange input for a 3D scene or 3D geometry.

SceneInput: The orange input is an optional input that accepts a 3D scene. If a scene is provided, the Transform controls in this node apply to the entire scene provided.

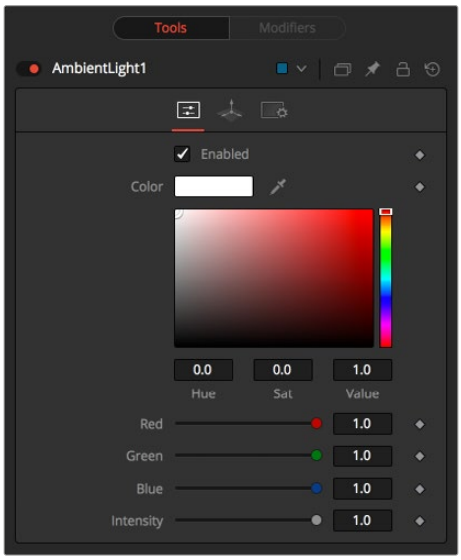
Basic Node Setup

The Ambient Light node is designed to be part of a larger 3D scene. You connect the light directly into a Merge 3D. Separating lights into different Merge 3D nodes allows you to control which lights affect which objects.



Ambient Light node structure

Inspector



Ambient Light controls

Controls Tab

The Controls tab is used to set the color and brightness of the ambient light.

Enabled

When the Enabled checkbox is turned on, the ambient light affects the scene. When the checkbox to turned off, the light is turned off. This checkbox performs the same function as the red switch to the left of the node's name in the Inspector.

Color

Use this standard Color control to set the color of the light.

Intensity

Use this slider to set the Intensity of the ambient light. A value of 0.2 indicates 20% percent light. A perfectly white texture lit only with a 0.2 ambient light would render at 20% gray (.2, .2, .2).

Common Controls

Transform and Settings Tabs

The options presented in the Transform and Settings tabs are commonly found in other lighting nodes. For more detail on the controls found in these tabs, see "The Common Controls" section at the end of this chapter.

Directional Light [3DL]



The Directional Light node

Directional Light Node Introduction

A directional light is a light with a clear direction but without a clear source or distance, similar to sunlight. This light shows an onscreen control, but the position of the control has no meaning. The rotation of the control is used to determine from where in the scene the light appears to be coming.

Similar to a Camera 3D, you connect lights into a Merge 3D and view them in the scene by viewing the Merge 3D node. Selecting a light node and loading it into the viewer does not show anything.

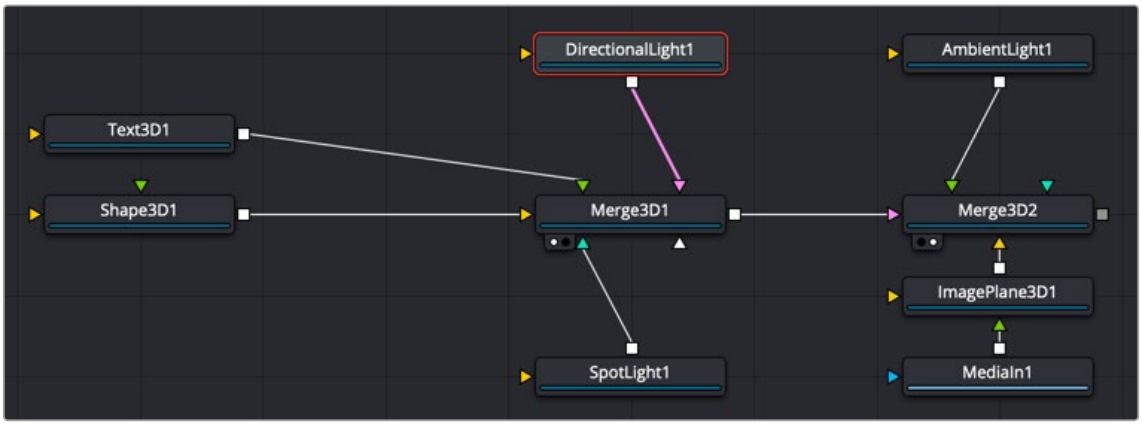
Inputs

The Directional Light node includes a single optional orange input for a 3D scene or 3D geometry.

SceneInput: The orange input is an optional input that accepts a 3D scene. If a scene is provided, the Transform controls in this node apply to the entire scene provided.

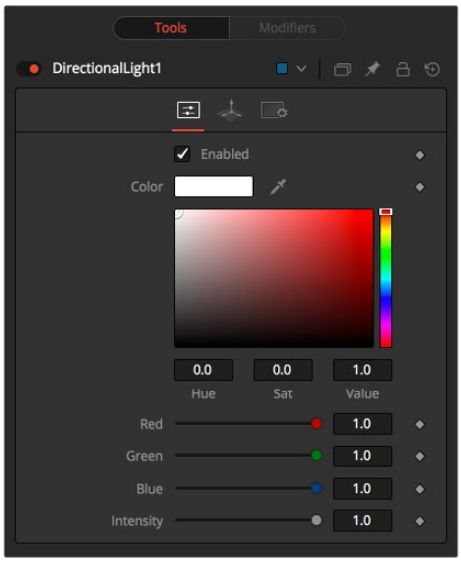
Basic Node Setup

The Directional Light node is designed to be part of a larger 3D scene. You connect the light directly into a Merge 3D. Separating lights into different Merge 3D nodes allows you to control which lights affect which objects.



Directional Light node structure

Inspector



Directional Light controls

Controls Tab

The Controls tab is used to set the color and brightness of the directional light. The direction of the light source is controlled by the rotation controls in the Transform tab.

Enabled

When the Enabled checkbox is turned on, the directional light affects the scene. When the checkbox is turned off, the light is turned off. This checkbox performs the same function as the red switch to the left of the node's name in the Inspector.

Color

Use this standard Color control to set the color of the light.

Intensity

Use this slider to set the Intensity of the directional light. A value of 0.2 indicates 20% percent light.

Shadows

This section provides several controls used to define the shadow map used when this spotlight creates shadows.

For more information, *see Chapter 25, "3D Compositing Basics,"* in the Fusion Reference Manual.

NOTE: The Directional light requires Hardware Renderer set as the Rendering Type in the Renderer3D node, for shadows to be seen in the 3D viewer.

Enable Shadows

The Enable Shadows checkbox should be selected if the light is to produce shadows. This defaults to selected.

Shadow Color

Use this standard Color control to set the color of the shadow. This defaults to black (0, 0, 0).

Density

The shadow density determines the transparency of the shadow. A density of 1.0 produces a completely opaque shadow, whereas lower values make the shadow more transparent.

Shadow Map Size

The Shadow Map Size control determines the size of the bitmap used to create the shadow map. Larger values produce more detailed shadow maps at the expense of memory and performance.

Shadow Map Proxy

Shadow Map Proxy determines the size of the shadow map used when the Proxy or Auto Proxy modes are enabled. A value of 0.5 would produce a shadow map at half the resolution defined in the Shadow Map Size.

Multiplicative/Additive Bias

Shadows are essentially textures applied to objects in the scene, so there is occasionally Z-fighting, where the portions of the object that should be receiving the shadows render over the top of the shadow. Biasing works by adding a small depth offset to move the shadow away from the surface it is shadowing, eliminating the Z-fighting. Too little bias and the objects can self-shadow themselves. Too much bias and the shadow can become separated from the surface. Adjust the Multiplicative Bias first, and then fine tune the result using the Additive Bias control.

> For more information, see the *Multiplicative and Additive Bias* section of Chapter 25, "3D Compositing Basics," in the Fusion Reference Manual.

Force All Materials Non-Transmissive

Normally, an RGBAZ shadow map is used when rendering shadows. By enabling this option, you are forcing the renderer to use a Z-only shadow map. This can lead to significantly faster shadow rendering while using a fifth as much memory. The disadvantage is that you can no longer cast "stained glass"-like shadows.

Shadow Map Sampling

Sets the quality for sampling of the shadow map.

Softness

Soft edges in shadows are produced by filtering the shadow map when it is sampled. Fusion provides two separate filtering methods for rendering shadows, which produce different effects.

> **NOTE:** Shadows have a hard edge. No filtering of the shadow map is done at all. The advantage of this method is that you only have to sample one pixel in the shadow map, so it is fast.

- **Constant:** Shadows edges have a constant softness. A filter with a constant width is used when sampling the shadow map. Adjusting the Constant Softness slider controls the size of the filter. Note that the larger you make the filter, the longer it takes to render the shadows.
- **Variable:** The shadow edge softness grows the farther the shadow receiver is positioned from the shadow caster. The variable softness is achieved by changing the size of the filter based on the distance between the receiver and caster. When this option is selected, the Softness Falloff, Min Softness, and Max Softness sliders appear.

Constant Softness

If the Softness is set to Constant, then this slider appears. It can be used to set the overall softness of the shadow.

Softness Falloff

The Softness Falloff slider appears when the Softness is set to variable. This slider controls how fast the softness of shadow edges grows with distance. More precisely, it controls how fast the shadow map filter size grows based upon the distance between the shadow caster and receiver. Its effect is mediated by the values of the Min and Max Softness sliders.

Min Softness

The Min Softness slider appears when the Softness is set to Variable. This slider controls the Minimum Softness of the shadow. The closer the shadow is to the object casting the shadow, the sharper it is, up to the limit set by this slider.

Max Softness

The Max Softness slider appears when the Softness is set to Variable. This slider controls the Maximum Softness of the shadow. The farther the shadow is from the object casting the shadow, the softer it is, up to the limit set by this slider.

Common Controls

Transform and Settings Tabs

The options presented in the Transform and Settings tabs are commonly found in other lighting nodes. For more detail on the controls found in these tabs, see “The Common Controls” section at the end of this chapter.

DomeLight [3Do]



The Dome light

Dome Light Node Introduction

A Dome Light uses SDR and HDR images and surrounds the entire scene in a sphere light mapped with the image, similar to environment light. This light shows an onscreen control.

Inputs

This node includes an optional orange input for a 3D scene or 3D geometry, and a white Dome Texture input for connecting an image.

SceneInput: The orange input is an optional input that accepts a 3D scene. If a scene is provided, the Transform controls in this node apply to the entire scene provided.

DomeTexture: The white input connects to an SDR or HDR image.

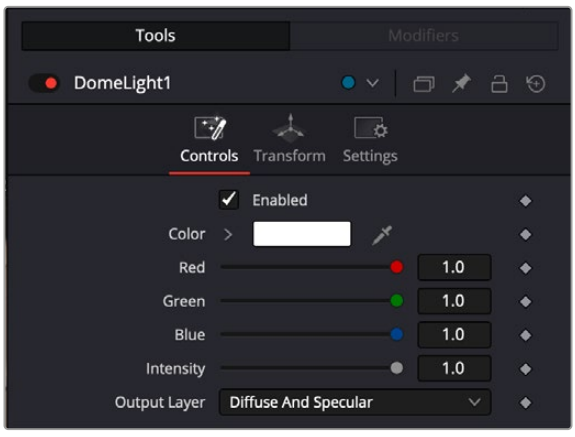
Basic Node Setup

The Dome Light node is designed to be part of a larger 3D scene. You connect the light directly into a Merge 3D. Separating lights into different Merge 3D nodes allows you to control which lights affect which objects.



Dome Light node structure

Inspector



The Dome Light controls

Controls Tab

The Controls tab is used to set the color and brightness of the uDome light. The position, direction, and scale of the light source is controlled in the Transform tab.

Color

Use this standard Color control to set the color of the light.

Intensity

Use this slider to set the intensity of the light. A value of 0.2 indicates 20% percent light.

Point Light [3PL]



The Point Light node

Point Light Node Introduction

A point light is a light source with a clear position in space that emits light in all directions. A light bulb is a good example of a point light.

This light shows an onscreen control, although only the position and distance of the control affect the light. Since the light is a 360-degree source, rotation has no meaning. Additionally, a point light may fall off with distance, unlike an ambient or directional light.

Similar to a Camera 3D, you connect lights into a Merge 3D and view them in the scene by viewing the Merge 3D node. Selecting a light node and loading it into the viewer does not show anything.

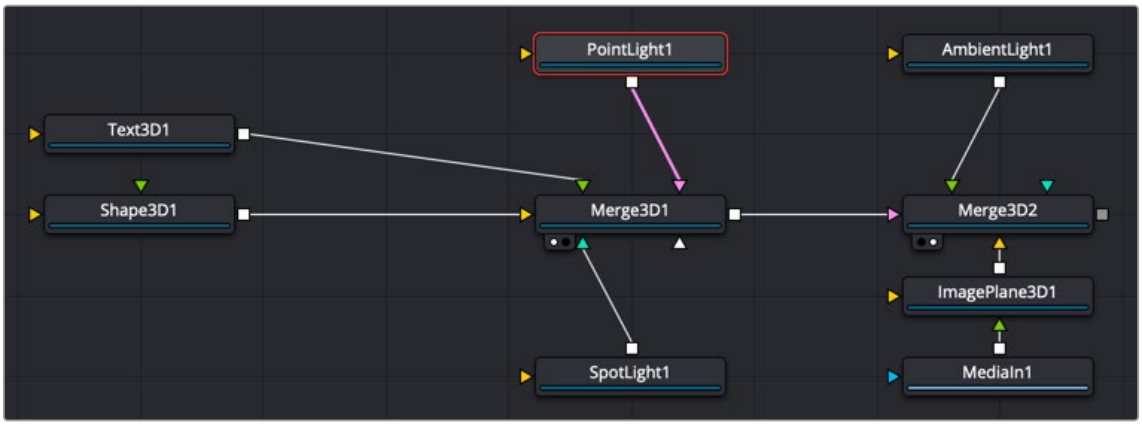
Inputs

The Point Light node includes a single optional orange input for a 3D scene or 3D geometry.

SceneInput: The orange input is an optional input that accepts a 3D scene. If a scene is provided, the Transform controls in this node apply to the entire scene provided.

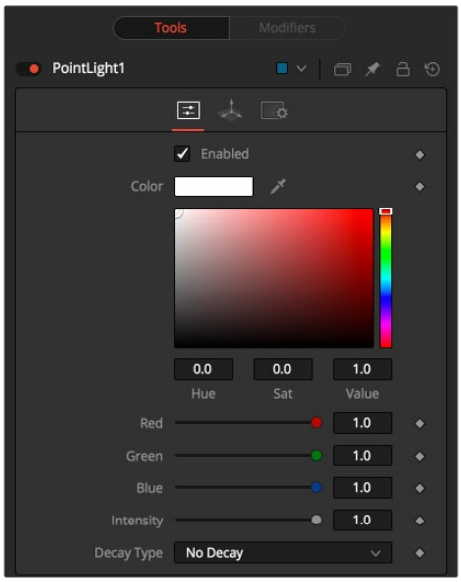
Basic Node Setup

The Point Light node is designed to be part of a larger 3D scene. You connect the light directly into a Merge 3D. Separating lights into different Merge 3D nodes allows you to control which lights affect which objects.



Point Light node structure

Inspector



Point Light controls

Controls Tab

The Controls tab is used to set the color and brightness of the point light. The position and distance of the light source are controlled in the Transform tab.

Enabled

When the Enabled checkbox is turned on, the point light affects the scene. When the checkbox is turned off the light is turned off. This checkbox performs the same function as the red switch to the left of the node's name in the Inspector.

Color

Use this standard Color control to set the color of the light.

Intensity

Use this slider to set the Intensity of the point light. A value of 0.2 indicates 20% percent light.

Decay Type

A point light defaults to No Decay, meaning that its light has equal intensity at all points in the scene. To cause the intensity to fall off with distance, set the Decay Type to either Linear or Quadratic modes.

Shadows

This section provides several controls used to define the shadow map used when this spotlight creates shadows.

For more information, *see Chapter 25, "3D Compositing Basics,"* in the Fusion Reference Manual.

NOTE: The Point light requires Hardware Renderer set as the Rendering Type in the Renderer3D node for shadows to be seen in the 3D viewer.

Enable Shadows

The Enable Shadows checkbox should be selected if the light is to produce shadows. This defaults to selected.

Shadow Color

Use this standard Color control to set the color of the shadow. This defaults to black (0, 0, 0).

Density

The shadow density determines the transparency of the shadow. A density of 1.0 produces a completely opaque shadow, whereas lower values make the shadow more transparent.

Shadow Map Size

The Shadow Map Size control determines the size of the bitmap used to create the shadow map. Larger values produce more detailed shadow maps at the expense of memory and performance.

Shadow Map Proxy

Shadow Map Proxy determines the size of the shadow map used when the Proxy or Auto Proxy modes are enabled. A value of 0.5 would produce a shadow map at half the resolution defined in the Shadow Map Size.

Multiplicative/Additive Bias

Shadows are essentially textures applied to objects in the scene, so there is occasionally Z-fighting, where the portions of the object that should be receiving the shadows render over the top of the shadow. Biasing works by adding a small depth offset to move the shadow away from the surface it is shadowing, eliminating the Z-fighting. Too little bias and the objects can self-shadow themselves. Too much bias and the shadow can become separated from the surface. Adjust the Multiplicative Bias first, and then fine tune the result using the Additive Bias control.

> For more information, see the *Multiplicative and Additive Bias* section of Chapter 25, "3D Compositing Basics," in the Fusion Reference Manual.

Force All Materials Non-Transmissive

Normally, an RGBAZ shadow map is used when rendering shadows. By enabling this option, you are forcing the renderer to use a Z-only shadow map. This can lead to significantly faster shadow rendering while using a fifth as much memory. The disadvantage is that you can no longer cast "stained glass"-like shadows.

Shadow Map Sampling

Sets the quality for sampling of the shadow map.

Softness

Soft edges in shadows are produced by filtering the shadow map when it is sampled. Fusion provides two separate filtering methods for rendering shadows, which produce different effects.

- **Constant:** Shadows edges have a constant softness. A filter with a constant width is used when sampling the shadow map. Adjusting the Constant Softness slider controls the size of the filter. Note that the larger you make the filter, the longer it takes to render the shadows.
- **Variable:** The shadow edge softness grows the farther the shadow receiver is positioned from the shadow caster. The variable softness is achieved by changing the size of the filter based on the distance between the receiver and caster. When this option is selected, the Softness Falloff, Min Softness, and Max Softness sliders appear.

Constant Softness

If the Softness is set to Constant, then this slider appears. It can be used to set the overall softness of the shadow.

Softness Falloff

The Softness Falloff slider appears when the Softness is set to variable. This slider controls how fast the softness of shadow edges grows with distance. More precisely, it controls how fast the shadow map filter size grows based upon the distance between the shadow caster and receiver. Its effect is mediated by the values of the Min and Max Softness sliders.

Min Softness

The Min Softness slider appears when the Softness is set to Variable. This slider controls the Minimum Softness of the shadow. The closer the shadow is to the object casting the shadow, the sharper it is, up to the limit set by this slider.

Max Softness

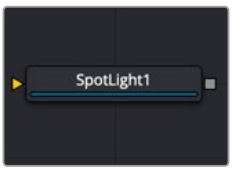
The Max Softness slider appears when the Softness is set to Variable. This slider controls the Maximum Softness of the shadow. The farther the shadow is from the object casting the shadow, the softer it is, up to the limit set by this slider.

Common Controls

Transform and Settings Tabs

The options presented in the Transform and Settings tabs are commonly found in other lighting nodes. For more detail on the controls found in these tabs, see “The Common Controls” section at the end of this chapter.

Spot Light [3SL]



The Spot Light node

Spot Light Node Introduction

A spotlight is a light that comes from a specific point and that has a clearly defined cone, with falloff of the light to the edges. Experienced stage and theatre lighting technicians may recognize the spotlight as being very similar to practical lights used in live productions. This is the only type of light capable of casting shadows.

Similar to a Camera 3D, you connect lights into a Merge 3D and view them in the scene by viewing the Merge 3D node. Selecting a light node and loading it into the viewer does not show anything.

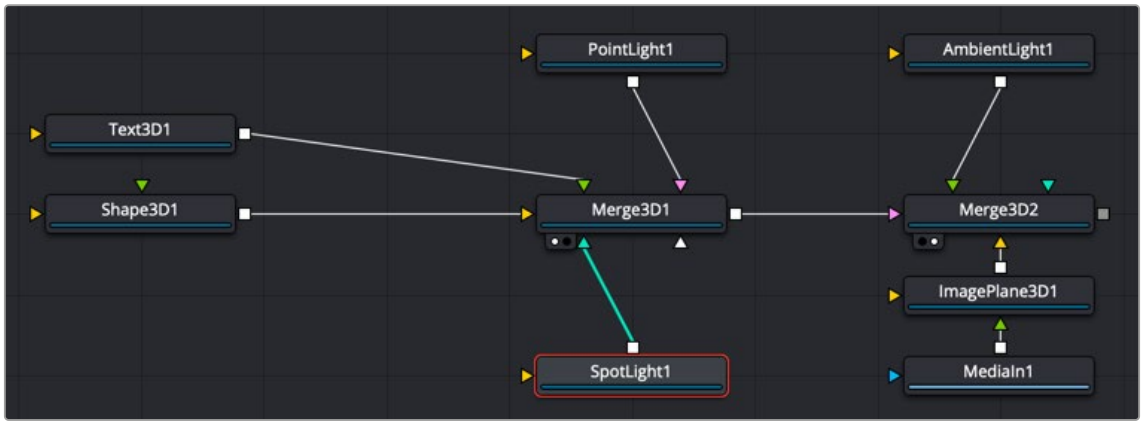
Inputs

The Spot Light node includes a single optional orange input for a 3D scene or 3D geometry.

SceneInput: The orange input is an optional input that accepts a 3D scene. If a scene is provided, the Transform controls in this node apply to the entire scene provided.

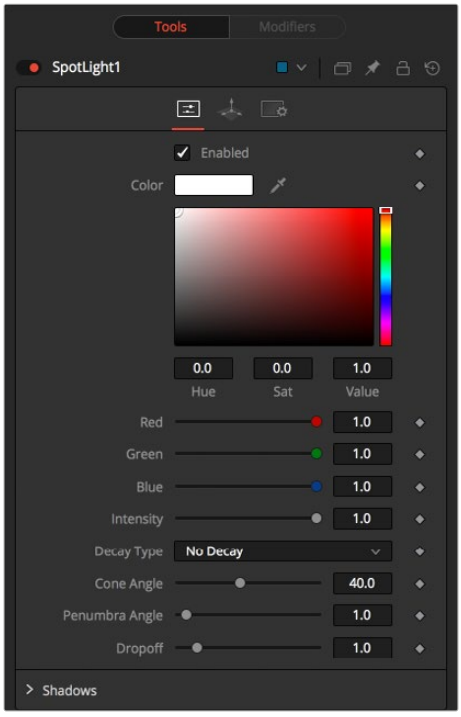
Basic Node Setup

The Spot Light node is designed to be part of a larger 3D scene. You connect the light directly into a Merge 3D. Separating lights into different Merge 3D nodes allows you to control which lights affect which objects.



Spot Light node structure

Inspector



Spot Light control

Controls Tab

The Controls tab is used to set the color and brightness of the spotlight. The position, rotation, and distance of the light source are controlled in the Transform tab.

Enabled

When the Enabled checkbox is turned on, the spotlight affects the scene. When the checkbox is turned off the light is turned off. This checkbox performs the same function as the red switch to the left of the node's name in the Inspector.

Color

Use this standard Color control to set the color of the light.

Intensity

Use this slider to set the Intensity of the spot light. A value of 0.2 indicates 20% percent light.

Decay Type

A spotlight defaults to No Falloff, meaning that its light has equal intensity on geometry despite the distance from the light to the geometry. To cause the intensity to fall off with distance, set the Decay type to either Linear or Quadratic modes.

Cone Angle

The Cone Angle of the light refers to the width of the cone where the light emits its full intensity. The larger the angle, the wider the cone angle, up to a limit of 90 degrees.

Penumbra Angle

The Penumbra Angle determines the area beyond the cone angle where the light's intensity falls off toward 0. A larger penumbra angle defines a larger falloff, while a value of 0 generates a hard-edged light.

Dropoff

The Dropoff controls how quickly the penumbra angle falls off from full intensity to 0.

Shadows

This section provides several controls used to define the shadow map used when this spotlight creates shadows.

> For more information, *see Chapter 25, "3D Compositing Basics,"* in the Fusion Reference Manual.

Enable Shadows

The Enable Shadows checkbox should be selected if the light is to produce shadows. This defaults to selected.

Shadow Color

Use this standard Color control to set the color of the shadow. This defaults to black (0, 0, 0).

Density

The shadow density determines the transparency of the shadow. A density of 1.0 produces a completely opaque shadow, whereas lower values make the shadow more transparent.

Shadow Map Size

The Shadow Map Size control determines the size of the bitmap used to create the shadow map. Larger values produce more detailed shadow maps at the expense of memory and performance.

Shadow Map Proxy

Shadow Map Proxy determines the size of the shadow map used when the Proxy or Auto Proxy modes are enabled. A value of 0.5 would produce a shadow map at half the resolution defined in the Shadow Map Size.

Multiplicative/Additive Bias

Shadows are essentially textures applied to objects in the scene, so there is occasionally Z-fighting, where the portions of the object that should be receiving the shadows render over the top of the shadow. Biasing works by adding a small depth offset to move the shadow away from the surface it is shadowing, eliminating the Z-fighting. Too little bias and the objects can self-shadow themselves. Too much bias and the shadow can become separated from the surface. Adjust the Multiplicative Bias first, and then fine tune the result using the Additive Bias control.

For more information, see the Multiplicative and Additive Bias section of *see Chapter 25, "3D Compositing Basics,"* in the Fusion Reference Manual.

Force All Materials Non-Transmissive

Normally, an RGBAZ shadow map is used when rendering shadows. By enabling this option, you are forcing the renderer to use a Z-only shadow map. This can lead to significantly faster shadow rendering while using a fifth as much memory. The disadvantage is that you can no longer cast "stained glass"-like shadows.

Shadow Map Sampling

Sets the quality for sampling of the shadow map.

Softness

Soft edges in shadows are produced by filtering the shadow map when it is sampled. Fusion provides two separate filtering methods for rendering shadows, which produce different effects.

NOTE: Shadows have a hard edge. No filtering of the shadow map is done at all. The advantage of this method is that you only have to sample one pixel in the shadow map, so it is fast.

- **Constant:** Shadows edges have a constant softness. A filter with a constant width is used when sampling the shadow map. Adjusting the Constant Softness slider controls the size of the filter. Note that the larger you make the filter, the longer it takes to render the shadows.
- **Variable:** The shadow edge softness grows the further the shadow receiver is positioned from the shadow caster. The variable softness is achieved by changing the size of the filter based on the distance between the receiver and caster. When this option is selected, the Softness Falloff, Min Softness, and Max Softness sliders appear.

Constant Softness

If the Softness is set to Constant, then this slider appears. It can be used to set the overall softness of the shadow.

Softness Falloff

The Softness Falloff slider appears when the Softness is set to variable. This slider controls how fast the softness of shadow edges grows with distance. More precisely, it controls how fast the shadow map filter size grows based upon the distance between the shadow caster and receiver. Its effect is mediated by the values of the Min and Max Softness sliders.

Min Softness

The Min Softness slider appears when the Softness is set to Variable. This slider controls the Minimum Softness of the shadow. The closer the shadow is to the object casting the shadow, the sharper it is, up to the limit set by this slider.

Max Softness

The Max Softness slider appears when the Softness is set to Variable. This slider controls the Maximum Softness of the shadow. The further the shadow is from the object casting the shadow, the softer it is, up to the limit set by this slider.

Common Controls

Transform and Settings Tabs

The options presented in the Transform and Settings tabs are commonly found in other lighting nodes. For more detailed information on the controls found in these tabs, see “The Common Controls” section at the end of this chapter.

The Common Controls

Nodes that handle 3D lighting share several identical controls in the Inspector. This section describes controls that are common among 3D lighting nodes.

Common Transform Tab



Common Transform 3D tab

Many tools in the 3D category include a Transform tab used to position, rotate, and scale the object in 3D space.

Translation

X, Y, Z Offset

These controls can be used to position the 3D element.

Rotation

Rotation Order

Use these buttons to select which order is used to apply Rotation along each axis of the object. For example, XYZ would apply the rotation to the X axis first, followed by the Y axis, and finally the Z axis.

X, Y, Z Rotation

Use these control to rotate the object around its pivot point. If the Use Target checkbox is selected, then the rotation is relative to the position of the target; otherwise, the global axis is used.

Pivot

X, Y, Z Pivot

A pivot point is the point around which an object rotates. Normally, an object rotates around its own center, which is considered to be a pivot of 0,0,0. These controls can be used to offset the pivot from the center.

Scale

X, Y, Z Scale

If the Lock X/Y/Z checkbox is checked, a single Scale slider is shown. This adjusts the overall size of the object. If the Lock checkbox is unchecked, individual X, Y, and Z sliders are displayed to allow individual scaling in each dimension. Note: If the Lock checkbox is checked, scaling of individual dimensions is not possible, even when dragging specific axes of the Transformation Widget in scale mode.

Use Target

Selecting the Use Target checkbox enables a set of controls for positioning an XYZ target. When Target is enabled, the object always rotates to face the target. The rotation of the object becomes relative to the target.

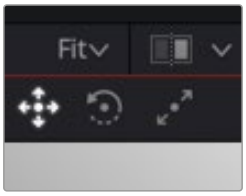
Import Transform

Opens a file browser where you can select a scene file saved or exported by your 3D application. It supports the following file types:

LightWave Scene	`.lws`
Max Scene	`.ase`
Maya Ascii Scene	`.ma`
dotXSI	`.xsi`

The Import Transform button imports only transformation data. For 3D geometry, lights, and cameras, consider using the File > FBX Import option.

Onscreen Transformation Controls



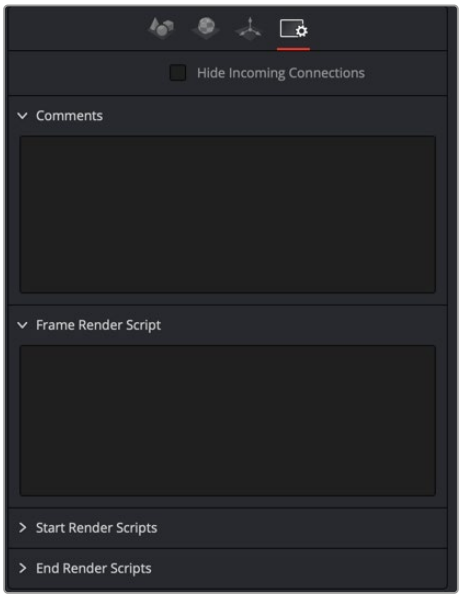
Viewer Transform buttons

Viewer Transform buttons

Most of the controls in the Transform tab are represented in the viewer with onscreen controls for transformation, rotation, and scaling. To change the mode of the onscreen controls, select one of the three buttons in the toolbar in the upper left of the viewer. The modes can also be toggled using the keyboard shortcut Q for translation, W for rotation, and E for scaling. In all three modes, individual axes of the control may be dragged to affect just that axis, or the center of the control may be dragged to affect all three axes.

The Scale sliders for most 3D tools default to locked, which causes uniform scaling of all three axes. Unlock the Lock X/Y/Z Scale checkbox to scale an object on a single axis only.

Settings Tab



Common Settings 3D controls

The Common Settings tab can be found on almost every tool found in Fusion. The following controls are specific settings for 3D nodes.

Hide Incoming Connections

Enabling this checkbox can hide connection lines from incoming nodes, making a node tree appear cleaner and easier to read. When enabled, fields for each input on a node are displayed. Dragging a connected node from the node tree into the field hides that incoming connection line as long as the node is not selected in the node tree. When the node is selected in the node tree, the line reappears.

Comment Tab

The Comment tab contains a single text control that is used to add comments and notes to the tool. When a note is added to a tool, a small red dot icon appears next to the setting's tab icon, and a text bubble appears on the node. To see the note in the Node Editor, hold the mouse pointer over the node for a moment. The contents of the Comments tab can be animated over time, if required.

Scripting Tab

The Scripting tab is present on every tool in Fusion. It contains several edit boxes used to add scripts that process when the tool is rendering. For more details on the contents of this tab, please consult the scripting documentation.

Chapter 31

3D Material Nodes

This chapter details the 3D Material nodes available when creating 3D composites in Fusion. The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Blinn [3Bl] 839
Channel Boolean [3Bol] 843
Cook Torrance [3CT] 846
Material Merge 3D [3MM] 851
Phong [3Ph] 852
Reflect [3Rr] 857
Stereo Mix [3SMM] 860
Ward [3Wd] 861
The Common Controls 866

Blinn [3BI]



The Blinn node

Blinn Node Introduction

The Blinn node is a basic illumination material that can be applied to geometry in the 3D scene. It describes how the object responds to light and provides multiple texture map inputs to allow fine control over the diffuse, specular, and bump map components of the material.

The standard basic material provided in the Materials tab of most geometry nodes is a simplified version of the Blinn node. The primary difference is that the Blinn node provides additional texture map inputs beyond just diffuse.

The Blinn node outputs a 3D Material that can be connected to the material inputs on any 3D geometry node.

The Blinn model in Fusion calculates the highlight as the dot product of the surface normal and the half angle vector between light source and viewer (dot(N, H)). This may not always match the Blinn model illumination model used by other 3D applications.

Inputs

There are five inputs on the Blinn node that accept 2D images or 3D materials. These inputs control the overall color and image used for the 3D object as well as the color and texture used in the specular highlight. Each of these inputs multiplies the pixels in the texture map by the equivalently named parameters in the node itself. This provides an effective method for scaling parts of the material.

Diffuse Texture: The orange Diffuse Texture input accepts a 2D image or a 3D material to be used as a main object texture map.

Specular Color Material: The green Specular Color material input accepts a 2D image or a 3D material to be used as the color texture map for specula highlight areas.

Specular Intensity Materials: The magenta Specular Intensity material input accepts a 2D image or a 3D material to be used to alter the intensity of specular highlights. When the input is a 2D image, the Alpha channel is used to create the map, while the color channels are discarded.

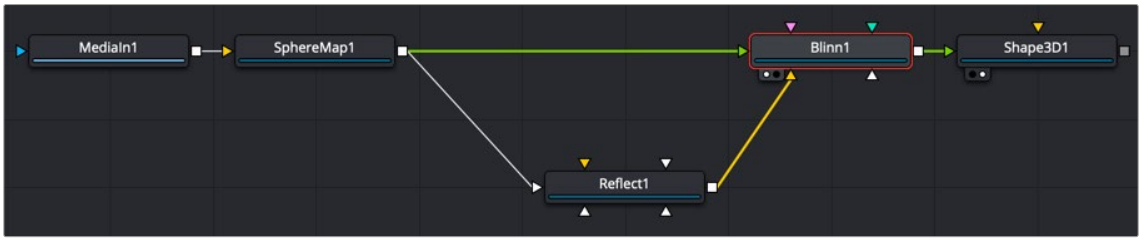
Specular Exponent Material: The teal Specular Exponent material input accepts a 2D image or a 3D material that is used as a falloff map for the material's specular highlights. When the input is a 2D image, the Alpha channel is used to create the map, while the color channels are discarded.

Bump Map Material: The white Bump Map material input accepts only a 3D material. Typically, you connect the texture into a Bump Map node, and then connect the Bump Map node to this input. This input uses the RGB information as texture-space normals.

When nodes have as many inputs as this one does, it is often difficult to make connections with any precision. Hold down the Option (macOS) or Alt (Windows) key while dragging the output from another node over the node tile, and keep holding Option or Alt when releasing the left mouse button. A small drop-down menu listing all the inputs provided by the node appears. Click on the desired input to complete the connection. Alternatively, you can drag the output from a node with the right mouse button to activate the same menu.

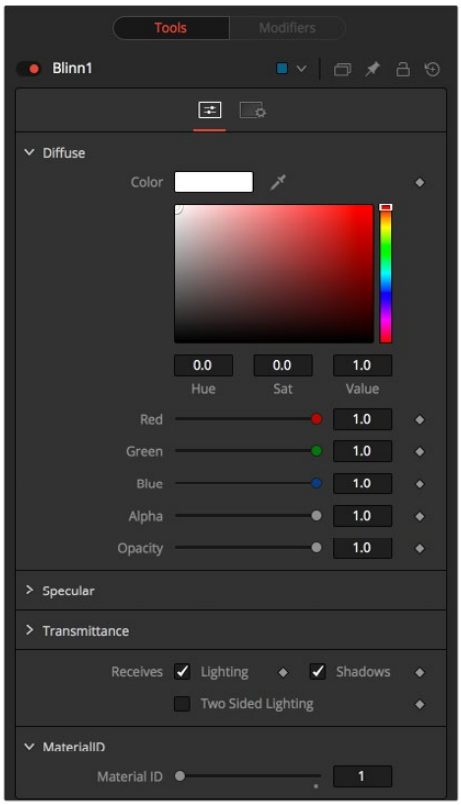
Basic Node Setup

The output of a Blinn node output is connected to the material input on a 3D scene or 3D geometry node to which you want the shader applied. The Blinn inputs can use images as the diffuse color material (orange) and specular color material (green). This can lead to a smooth, shiny material.



A Blinn shader with diffuse and specular color materials connected

Inspector



Blinn controls

Controls Tab

The Controls tab is the primary tab for the Blinn node. It controls the color and shininess applied to the surface of the 3D geometry.

Diffuse

Diffuse describes the base surface characteristics without any additional effects like reflections or specular highlights. Besides defining the base color of an object, the diffuse color also defines the transparency of the object. The Alpha in a diffuse texture map can be used to make portions of the surface transparent.

Diffuse Color

A material's Diffuse Color describes the base color presented by the material when it is lit indirectly or by ambient light. If a diffuse texture map is provided, then the color value provided here is multiplied by the color values in the texture.

Alpha

This slider sets the material's Alpha channel value. This affects diffuse and specular colors equally and affects the Alpha value of the material in the rendered output. If a diffuse texture map is provided, then the Alpha value set here is multiplied by the Alpha values in the texture map.

Opacity

Reducing the material's opacity decreases the color and Alpha values of the specular and diffuse colors equally, making the material transparent.

Specular

The parameters in the Specular section describe the look of the specular highlight of the surface. These values are evaluated in a different way for each illumination model.

Specular Color

Specular Color determines the color of light that reflects from a shiny surface. The more specular a material is, the glossier it appears. Surfaces like plastics and glass tend to have white specular highlights, whereas metallic surfaces like gold have specular highlights that inherit their color from the material color. If a specular texture map is provided, then the value provided here is multiplied by the color values from the texture.

Specular Intensity

Specular Intensity controls how strong the specular highlight is. If the specular intensity texture is provided, then this value is multiplied by the Alpha value of the texture.

Specular Exponent

Specular Exponent controls the falloff of the specular highlight. The greater the value, the sharper the falloff, and the smoother and glossier the material appears. If the specular exponent texture is provided, then this value is multiplied by the Alpha value of the texture map.

Transmittance

Transmittance controls the way light passes through a material. For example, a solid blue sphere casts a black shadow, but one made of translucent blue plastic would cast a much lower density blue shadow.

There is a separate Opacity option. Opacity determines how transparent the actual surface is when it is rendered. Fusion allows adjusting both opacity and transmittance separately. At first, this might be a bit counterintuitive to those who are unfamiliar with 3D software. It is possible to have a surface that is fully opaque but transmits 100% of the light arriving upon it, effectively making it a luminous/emissive surface.

Attenuation

Attenuation determines how much color is passed through the object. For an object to have transmissive shadows, set the attenuation to (1, 1, 1), which means 100% of green, blue, and red light passes through the object. Setting this color to RGB (1, 0, 0) means that the material transmits 100% of the red arriving at the surface but none of the green or blue light. This can be used for “stained glass”-styled shadows.

Alpha Detail

When the Alpha Detail slider is set to 0, the Alpha channel of the object is ignored and the entire object casts a shadow. If it is set to 1, the Alpha channel determines what portions of the object cast a shadow.

Color Detail

The Color Detail slider modulates light passing through the surface by the diffuse color + texture colors. Use this to throw a shadow that contains color details of the texture applied to the object. Increasing the slider from 0 to 1 brings in more diffuse color + texture color into the shadow. Note that the Alpha and opacity of the object are ignored when transmitting color, allowing an object with a solid Alpha to still transmit its color to the shadow.

Saturation

The Saturation slider controls the saturation of the color component transmitted to the shadow. Setting this to 0.0 results in monochrome shadows.

Receives Lighting/Shadows

These checkboxes control whether the material is affected by lighting and shadows in the scene. If turned off, the object is always fully lit and/or unshadowed.

Two-Sided Lighting

This effectively makes the surface two sided by adding a second set of normals facing the opposite direction on the backside of the surface. This is normally off to increase rendering speed, but it can be turned on for 2D surfaces or for objects that are not fully enclosed, to allow the reverse or interior surfaces to be visible as well.

Normally, in a 3D application, only the front face of a surface is visible and the back face is culled, so that if a camera were to revolve around a plane in a 3D application, when it reached the backside, the plane would become invisible. Making a plane two sided in a 3D application is equivalent to adding another plane on top of the first but rotated by 180 degrees so the normals are facing the opposite direction on the backside. Thus, when you revolve around the back, you see the second image plane, which has its normals facing the opposite way.

Fusion does exactly the same thing as 3D applications when you make a surface two sided. The confusion about what two-sided lighting does arises because Fusion does not cull back-facing polygons by default. If you revolve around a one-sided plane in Fusion, you still see it from the backside (but you are seeing the frontside duplicated through to the backside as if it were transparent). Making the plane two sided effectively adds a second set of normals to the backside of the plane.

NOTE: This can become rather confusing once you make the surface transparent, as the same rules still apply and produce a result that is counterintuitive. If you view from the frontside a transparent two-sided surface illuminated from the backside, it looks unlit.

Material ID

This slider sets the numeric identifier assigned to this material. This value is rendered into the MatID auxiliary channel if the corresponding option is enabled in the renderer.

Common Controls

Settings Tab

The Settings tab in the Inspector is duplicated in other 3D nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Channel Boolean [3Bol]



The Channel Boolean node

Channel Boolean Node Introduction

The Channel Boolean (not to be confused with the 2D Channel Booleans) can be used to remap and modify channels of 3D materials using mathematical operations. For example, if you want to use the red channel of a material to control a scalar input of an illumination model that uses the Alpha channel (e.g., Blinn. SpecularExponent), you can remap the channels here. Furthermore, it allows the use of geometry-specific information like texture space coordinates and normals.

Inputs

There are two inputs on the Channel Boolean Node: one for the foreground material, and one for the background material. Both inputs accept either a 2D image or a 3D material like Blinn, Cook-Torrence, or Phong node.

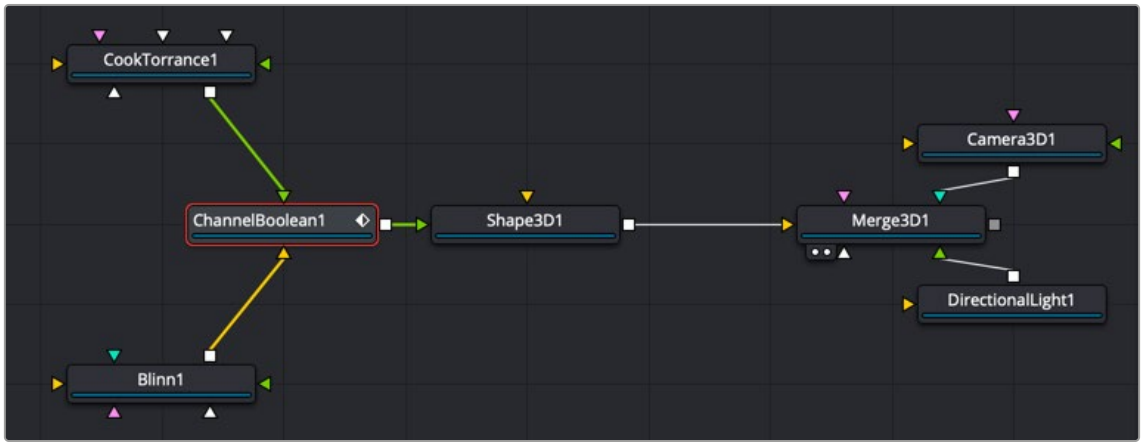
BackgroundMaterial: The orange background material input accepts a 2D image or a 3D material.

ForegroundMaterial: The green foreground input also accepts a 2D image or a 3D material.

Basic Node Setup

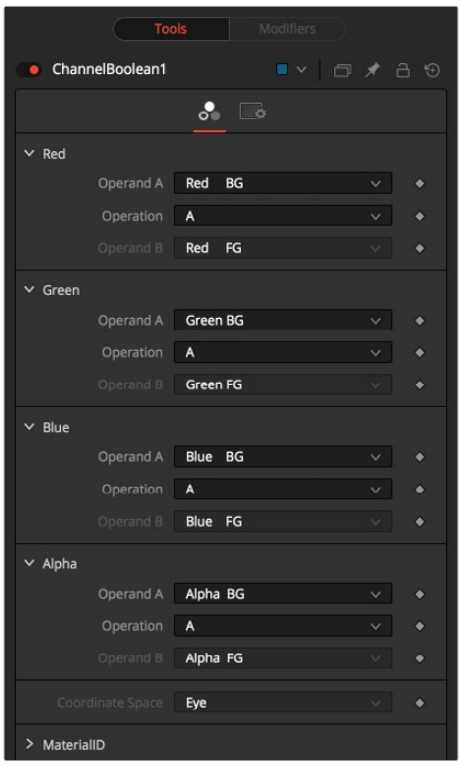
There are many uses for the material 3D Channel Boolean. Most often it is used to combine material looks or manipulate UV texture coordinates

In the below example, the Channel Boolean node combines the Cook Torrance and Blinn materials. It uses the math operands in the Channel Boolean to switch, invert, and mix the two inputs, creating a neon flickering effect.



A Channel Boolean used to combine and operate on Cook Torrance and Blinn nodes

Inspector



Channel Boolean controls

Controls Tab

The Controls tab includes a section for each RGBA channel. Within each channel are two input menus called Operand A and Operand B. The function performed on these two inputs is selected in the Operation menu.

Operand A/B

The Operand menus, one for each output RGBA channel, allow you to set the desired input information for the corresponding channel.

Red/Green/Blue/Alpha FG	Reads the color information of the foreground material.
Red/Green/Blue/Alpha BG	Reads the color information of the background material.
Black/White/Mid Gray	Sets the value of the channel to 0, 0.5, or 1.
Hue/Lightness/Saturation FG	Reads the color information of the foreground material, converts it into the HLS color space, and puts the selected information into the corresponding channel.
Hue/Lightness/Saturation BG	Reads the color information of the background material, converts it into the HLS color space, and puts the selected information into the corresponding channel.
Luminance FG	Reads the color information of the foreground material and calculates the luminance value for the channel.
Luminance BG	Reads the color information of the background material and calculates the luminance value for the channel.
X/Y/Z Position FG	Sets the value of the channel to the position of the pixel in 3D space. The vector information is returned in eye space.
U/V/W Texture FG	Applies the texture space coordinates of the foreground material to the channels.
U/V/W EnvCoords FG	Applies the environment texture space coordinates to the channels. Use it upstream of nodes modifying the environment texture coordinates like the Reflect 3D node.
X/Y/Z Normal	Sets the value of the channel to the selected axis of the normal vector. The vector is returned in eye space.

Operation

Determines the Operation of how the operands are combined.

A	Uses Operand A only for the output channel.
B	Uses Operand B only for the output channel.
1-A	Subtracts the value of Operand A from 1.
1-B	Subtracts the value of Operand B from 1.
A+B	Adds the value of Operand A and B.
A-B	Subtracts the value of Operand B from A.
A*B	Multiplies the value of both Operands.
A/B	Divides the value of Operand B from A.
min(A,B)	Compares the values of Operands A and B and returns the smaller one.
max(A,B)	Compares the values of Operands A and B and returns the bigger one.
avg(A,B)	Returns the average value of both Operands.

Material ID

This slider sets the numeric identifier assigned to this material. This value is rendered into the MatID auxiliary channel if the corresponding option is enabled in the renderer.

Common Controls

Settings Tab

The Settings tab in the Inspector is duplicated in other 3D nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Cook Torrance [3CT]



The Cook Torrance node

Cook Torrance Node Introduction

The Cook Torrance node is a basic illumination material that can be applied to geometry in the 3D scene. The diffuse calculation for this node is similar to that used in the basic material and the Blinn node, but the specular highlights are evaluated using an optimized Fresnel/Beckmann equation. This illumination model is primarily used for shading metal or other shiny and highly reflective surfaces.

The Cook Torrance node outputs a 3D Material that can be connected to the material inputs on any 3D geometry node.

Inputs

There are six inputs on the Cook Torrance node that accept 2D images or 3D materials. These inputs control the overall color and image used for the 3D object as well as controlling the color and texture used in the specular highlight. Each of these inputs multiplies the pixels in the texture map by the equivalently named parameters in the node itself. This provides an effective method for scaling parts of the material.

Diffuse Color Material: The orange Diffuse Color material input accepts a 2D image or a 3D material to be used as overall color and texture of the object.

Specular Color Material: The green Specular Color material input accepts a 2D image or a 3D material to be used as the color and texture of the specular highlight.

Specular Intensity Material: The magenta Specular Intensity material input accepts a 2D image or a 3D material to alter the intensity of the specular highlight. When the input is a 2D image, the Alpha channel is used to create the map, while the color channels are discarded.

Specular Roughness Material: The white Specular Roughness material input accepts a 2D image or a 3D material to be used as a map for modifying the roughness of the specular highlight. The Alpha of the texture map is multiplied by the value of the roughness control.

Specular Refractive Index Material: The white Specular Refractive Index material input accepts a 2D image or a 3D material, using the RGB channels as the refraction texture.

Bump Map Material: The white Bump Map material input accepts only a 3D material. Typically, you connect the texture into a Bump Map node, and then connect the Bump Map node to this input. This input uses the RGB information as texture-space normals.

Each of these inputs multiplies the pixels in the texture map by the equivalently named parameters in the node itself. This provides an effective method for scaling parts of the material.

When nodes have as many inputs as this one does, it is often difficult to make connections with any precision. Hold down the Option (macOS) or Alt (Windows) key while dragging the output from another node over the node tile, and keep holding Option or Alt when releasing the left mouse button. A small drop-down menu listing all the inputs provided by the node appears. Click on the desired input to complete the connection.

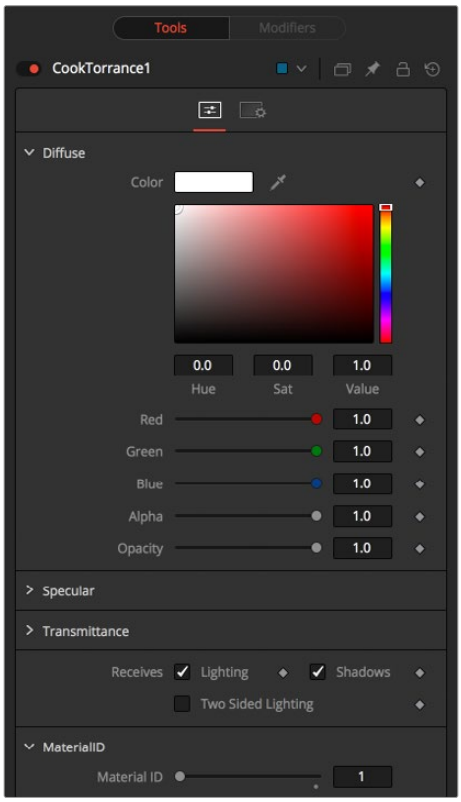
Basic Node Setup

The output of a Cook Torrance node output is connected to the material input on a 3D scene or 3D geometry node to which you want the shader applied. The Cook Torrance inputs can use images as the diffuse color material (yellow) and specular color material (green). This can result in a smooth, shiny material.



A Cook Torrance shader with diffuse and specular color materials connected

Inspector



Cook Torrance controls

Controls Tab

The Controls tab contains parameters for adjusting the main color, highlight, and lighting properties of the Cook Torrance shader node.

Diffuse

Diffuse describes the base surface characteristics without any additional effects like reflections or specular highlights. Besides defining the base color of an object, the diffuse color also defines the transparency of the object. The Alpha in a diffuse texture map can be used to make portions of the surface transparent.

Diffuse Color

A material's Diffuse Color describes the base color presented by the material when it is lit indirectly or by ambient light. If a diffuse texture map is provided, then the color value provided here is multiplied by the color values in the texture.

Alpha

This slider sets the material's Alpha channel value. This affects diffuse and specular colors equally, and affects the Alpha value of the material in the rendered output. If a diffuse texture map is provided, then the Alpha value set here is multiplied by the Alpha values in the texture map.

Opacity

Reducing the material's Opacity decreases the color and Alpha values of the specular and diffuse colors equally, making the material transparent.

Specular

The parameters in the Specular section describe the look of the specular highlight of the surface. These values are evaluated in a different way for each illumination model.

Specular Color

Specular Color determines the color of light that reflects from a shiny surface. The more specular a material is, the glossier it appears. Surfaces like plastics and glass tend to have white specular highlights, whereas metallic surfaces like gold have specular highlights that inherit their color from the material color. If a specular texture map is provided, then the value provided here is multiplied by the color values from the texture.

Specular Intensity

Specular Intensity controls how strong the specular highlight is. If the specular intensity texture is provided, then this value is multiplied by the Alpha value of the texture.

Roughness

The Roughness of the specular highlight describes diffusion of the specular highlight over the surface. The greater the value, the wider the falloff, and the more brushed and metallic the surface appears. If the roughness texture map is provided, then this value is multiplied by the Alpha value from the texture.

Do Fresnel

Selecting this checkbox adds Fresnel calculations to the materials illumination model. This provides more realistic-looking metal surfaces by taking into account the refractiveness of the material.

Refractive Index

This slider appears when the Do Fresnel checkbox is selected. The Refractive Index applies only to the calculations for the highlight; it does not perform actual refraction of light through transparent surfaces. If the refractive index texture map is provided, then this value is multiplied by the Alpha value of the input.

Transmittance

Transmittance controls the way light passes through a material. For example, a solid blue sphere casts a black shadow, but one made of translucent blue plastic would cast a much lower density blue shadow.

There is a separate Opacity option. Opacity determines how transparent the actual surface is when it is rendered. Fusion allows adjusting both opacity and transmittance separately. At first, this might be a bit counterintuitive to those who are unfamiliar with 3D software. It is possible to have a surface that is fully opaque but transmits 100% of the light arriving upon it, effectively making it a luminous/emissive surface.

Attenuation

Attenuation determines how much color is passed through the object. For an object to have transmissive shadows, set the attenuation to (1, 1, 1), which means 100% of green, blue, and red light passes through the object. Setting this color to RGB (1, 0, 0) means that the material transmits 100% of the red arriving at the surface but none of the green or blue light. This can be used to create “stained glass”-styled shadows.

Alpha Detail

When the Alpha Detail slider is set to 0, the Alpha channel of the object is ignored and the entire object casts a shadow. If it is set to 1, the Alpha channel determines what portions of the object cast a shadow.

Color Detail

The Color Detail slider modulates light passing through the surface by the diffuse color + texture colors. Use this to throw a shadow that contains color details of the texture applied to the object. Increasing the slider from 0 to 1 brings in more diffuse color + texture color into the shadow. Note that the Alpha and opacity of the object are ignored when transmitting color, allowing an object with a solid Alpha to still transmit its color to the shadow.

Saturation

The Saturation slider controls the saturation of the color component transmitted to the shadow. Setting this to 0.0 results in monochrome shadows.

Receives Lighting/Shadows

These checkboxes control whether the material is affected by lighting and shadows in the scene. If turned off, the object is always fully lit and/or unshadowed.

Two-Sided Lighting

This effectively makes the surface two sided by adding a second set of normals facing the opposite direction on the backside of the surface. This is normally off to increase rendering speed, but it can be turned on for 2D surfaces or for objects that are not fully enclosed, to allow the reverse or interior surfaces to be visible as well.

Normally, in a 3D application, only the front face of a surface is visible and the back face is culled, so that if a camera were to revolve around a plane in a 3D application, when it reached the backside, the plane would become invisible. Making a plane two sided in a 3D application is equivalent to adding another plane on top of the first but rotated by 180 degrees so the normals are facing the opposite direction on the backside. Thus, when you revolve around the back, you see the second image plane, which has its normals facing the opposite way.

> **NOTE:** This can become rather confusing once you make the surface transparent, as the same rules still apply and produce a result that is counterintuitive. If you view from the frontside a transparent two-sided surface illuminated from the backside, it looks unlit.

Material ID

This slider sets the numeric identifier assigned to this material. This value is rendered into the MatID auxiliary channel if the corresponding option is enabled in the renderer.

Common Controls

Settings Tab

The Settings tab in the Inspector is duplicated in other 3D nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Material Merge 3D [3MM]



The Material Merge node

Material Merge Node Introduction

The Material Merge node can be used to combine two separate materials together. This node can be used to composite Material nodes, combining multiple illumination materials (Blinn, Cook Torrance) with texture nodes (Bump Map, Reflection) to create complex shader networks.

The node also provides a mechanism for assigning a new material identifier to the combined material.

Inputs

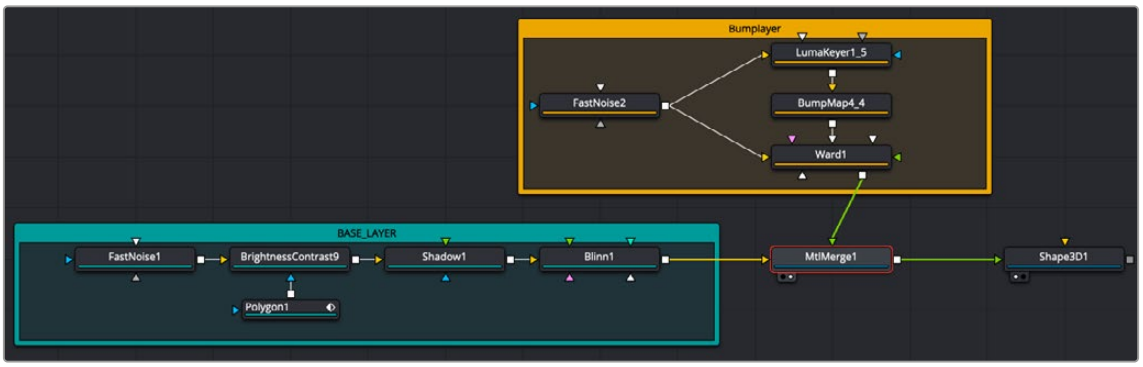
The Material Merge node includes two inputs for the two materials you want to combine.

Background Material: The orange Background material input accepts a 2D image or a 3D material to be used as the background material.

Foreground Material: The green Foreground material input accepts a 2D image or a 3D material to be used as the foreground material. A 2D image is treated as a diffuse texture map in the basic shading model.

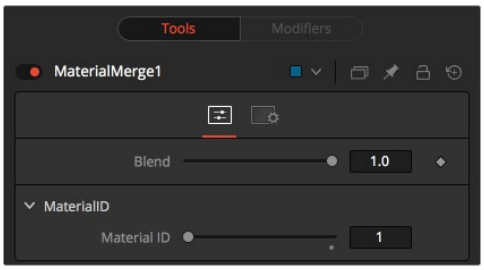
Basic Node Setup

The output of a Material Merge node is connected to the material input on a 3D scene or 3D geometry node. The Material Merge node below is taking in a background base layer from the Blinn shader and combining it with a more textured bump map layer.



A Material Merge node combining a Blinn-based shader (teal underlay) and a Ward-based shader (orange underlay)

Inspector



Material Merge controls

Controls Tab

The Controls tab includes a single slider for blending the two materials together.

Blend

The Blend behavior of the Material Merge is similar to the Dissolve (DX) node for images. The two materials/textures are mixed using the value of the slider to determine the percentage each input contributes. While the background and foreground inputs can be a 2D image instead of a material, the output of this node is always a material.

Unlike the 2D Dissolve node, both foreground and background inputs are required.

Material ID

This slider sets the numeric identifier assigned to the resulting material. This value is rendered into the MatID auxiliary channel if the corresponding option is enabled in the renderer.

Common Controls

Settings Tab

The Settings tab in the Inspector is duplicated in other 3D nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Phong [3Ph]



The Phong node

Phong Node Introduction

The Phong node is a basic illumination material that can be applied to geometry in the 3D scene. It describes how the object responds to light and provides multiple texture map inputs to allow fine control over the diffuse, specular, and bump map components of the material.

While producing a highlight similar to that produced by the Blinn model, it is more commonly used for shiny/polished plastic surfaces.

Inputs

There are five inputs on the Phong node that accept 2D images or 3D materials. These inputs control the overall color and image used for the 3D object as well as controlling the color and texture used in the specular highlight. Each of these inputs multiplies the pixels in the texture map by the equivalently named parameters in the node itself. This provides an effective method for scaling parts of the material.

Diffuse Material: The orange Diffuse material input accepts a 2D image or a 3D material to be used as a main color and texture of the object.

Specular Color Material: The green Specular Color material input accepts a 2D image or a 3D material to be used as a highlight color and texture of the object.

Specular Intensity Material: The magenta Specular Intensity material input accepts a 2D image or a 3D material to be used as an intensity map for the material's highlights. When the input is a 2D image, the Alpha channel is used to create the map, while the color channels are discarded.

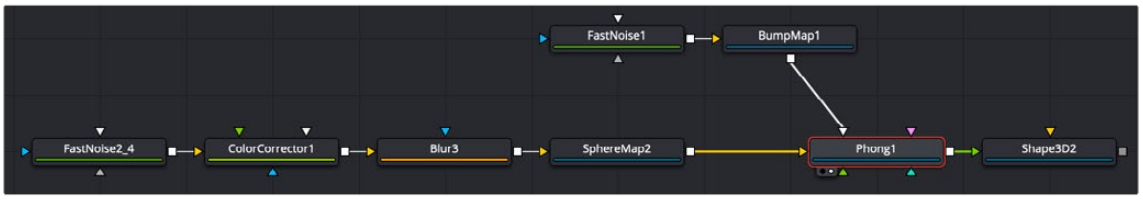
Specular Exponent Material: The teal Specular Exponent material input accepts a 2D image or a 3D material to be used as a falloff map for the material's specular highlights. When the input is a 2D image, the Alpha channel is used to create the map, while the color channels are discarded.

Bump Map Material: The white Bump Map texture input accepts only a 3D material. Typically, you connect the texture into a Bump Map node, and then connect the Bump Map node to this input. This input uses the RGB information as texture-space normals.

When nodes have as many inputs as this one does, it is often difficult to make connections with any precision. Hold down the Option or Alt key while dragging the output from another node over the node tile, and keep holding Option or Alt when releasing the left mouse button. A small drop-down menu listing all the inputs provided by the node appears. Click on the desired input to complete the connection.

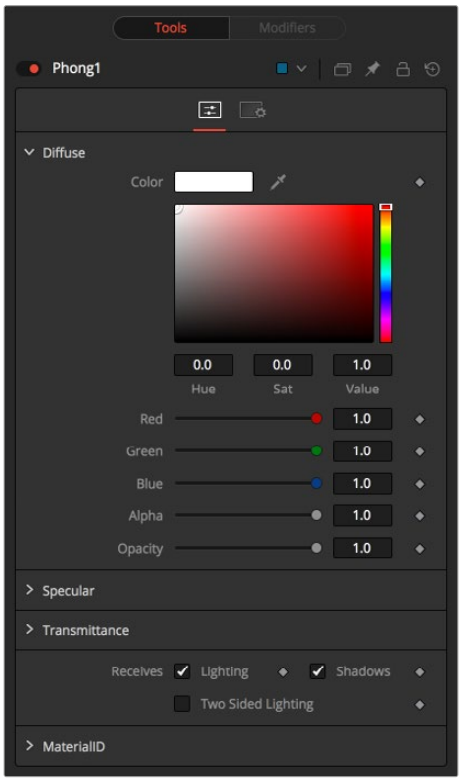
Basic Node Setup

The output of a Phong node is connected to the material input on a 3D scene or 3D geometry node. The Phong node below is taking in a base Color Diffuse input from the Fast Noise node and a bump map texture also generated from a Fast Noise node.



A Phong node with a diffuse color and Bump Map input

Inspector



Phong controls

Controls Tab

The Controls tab contains parameters for adjusting the main color, highlight, and lighting properties of the Phong shader node.

Diffuse

Diffuse describes the base surface characteristics without any additional effects like reflections or specular highlights. Besides defining the base color of an object, the diffuse color also defines the transparency of the object.

The Alpha in a diffuse texture map can be used to make portions of the surface transparent.

Diffuse Color

A material's Diffuse Color describes the base color presented by the material when it is lit indirectly or by ambient light. If a diffuse texture map is provided, then the color value provided here is multiplied by the color values in the texture.

Alpha

This slider sets the material's Alpha channel value. This affects diffuse and specular colors equally and affects the Alpha value of the material in the rendered output. If a diffuse texture map is provided, then the Alpha value set here is multiplied by the Alpha values in the texture map.

Opacity

Reducing the material's Opacity decreases the color and Alpha values of the specular and diffuse colors equally, making the material transparent.

Specular

The parameters in the Specular section describe the look of the specular highlight of the surface. These values are evaluated in a different way for each illumination model.

Specular Color

Specular Color determines the color of light that reflects from a shiny surface. The more specular a material is, the glossier it appears. Surfaces like plastics and glass tend to have white specular highlights, whereas metallic surfaces like gold have specular highlights that inherit their color from the material color. If a specular texture map is provided, then the value provided here is multiplied by the color values from the texture.

Specular Intensity

Specular Intensity controls how strong the specular highlight is. If the specular intensity texture is provided, then this value is multiplied by the Alpha value of the texture.

Specular Exponent

Specular Exponent controls the falloff of the specular highlight. The greater the value, the sharper the falloff, and the smoother and glossier the material appears. If the specular exponent texture is provided, then this value is multiplied by the Alpha value of the texture map.

Transmittance

Transmittance controls the way light passes through a material. For example, a solid blue sphere casts a black shadow, but one made of translucent blue plastic would cast a much lower density blue shadow.

There is a separate Opacity option. Opacity determines how transparent the actual surface is when it is rendered. Fusion allows adjusting both opacity and transmittance separately. At first, this might be a bit counterintuitive to those who are unfamiliar with 3D software. It is possible to have a surface that is fully opaque but transmits 100% of the light arriving upon it, effectively making it a luminous/emissive surface.

Attenuation

Attenuation determines how much color is passed through the object. For an object to have transmissive shadows, set the attenuation to (1, 1, 1), which means 100% of green, blue, and red light passes through the object. Setting this color to RGB (1, 0, 0) means that the material transmits 100% of the red arriving at the surface but none of the green or blue light. This can be used to create “stained glass”-styled shadows.

Alpha Detail

When the Alpha Detail slider is set to 0, the Alpha channel of the object is ignored and the entire object casts a shadow. If it is set to 1, the Alpha channel determines what portions of the object cast a shadow.

Color Detail

The Color Detail slider modulates light passing through the surface by the diffuse color + texture colors. Use this to throw a shadow that contains color details of the texture applied to the object. Increasing the slider from 0 to 1 brings in more diffuse color + texture color into the shadow. Note that the Alpha and opacity of the object are ignored when transmitting color, allowing an object with a solid Alpha to still transmit its color to the shadow.

Saturation

The Saturation slider controls the saturation of the color component transmitted to the shadow. Setting this to 0.0 results in monochrome shadows.

Receives Lighting/Shadows

These checkboxes control whether the material is affected by lighting and shadows in the scene. If turned off, the object is always fully lit and/or unshadowed.

Two-Sided Lighting

This effectively makes the surface two sided by adding a second set of normals facing the opposite direction on the backside of the surface. This is normally off to increase rendering speed, but it can be turned on for 2D surfaces or for objects that are not fully enclosed, to allow the reverse or interior surfaces to be visible as well.

Normally, in a 3D application, only the front face of a surface is visible and the back face is culled, so that if a camera were to revolve around a plane in a 3D application, when it reached the backside, the plane would become invisible. Making a plane two sided in a 3D application is equivalent to adding another plane on top of the first but rotated by 180 degrees so the normals are facing the opposite direction on the backside. Thus, when you revolve around the back, you see the second image plane, which has its normals facing the opposite way.

Fusion does exactly the same thing as 3D applications when you make a surface two sided. The confusion about what two-sided lighting does arises because Fusion does not cull back-facing polygons by default. If you revolve around a one-sided plane in Fusion, you still see it from the backside (but you are seeing the frontside duplicated through to the backside as if it were transparent). Making the plane two sided effectively adds a second set of normals to the backside of the plane.

> **NOTE:** This can become rather confusing once you make the surface transparent, as the same rules still apply and produce a result that is counterintuitive. If you view from the frontside a transparent two-sided surface illuminated from the backside, it looks unlit.

Material ID

This slider sets the numeric identifier assigned to this material. This value is rendered into the MatID auxiliary channel if the corresponding option is enabled in the renderer.

Common Controls

Settings Tab

The Settings tab in the Inspector is duplicated in other 3D nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Reflect [3Rr]



The Reflect node

Reflect Node Introduction

The Reflect node is used to add environment map reflections and refractions to materials.

Control is offered over the face on and glancing strength, falloff, per channel refraction indexes, and tinting. Several texture map inputs can modify the behavior of each parameter.

Environment mapping is an approximation that assumes an object's environment is infinitely distant from the object. It's best to picture this as a cube or sphere with the object at the center. Specifically, this infinite distance assumption means that objects cannot interact with themselves (e.g., the reflections on the handle of a teapot do not show the body of the teapot but rather the infinite environment map). It also means that if you use the same cube map on multiple objects in the scene, those objects do not inter-reflect each other (e.g., two neighboring objects would not reflect each other). If you want objects to reflect each other, you need to render a cube map for each object.

> For more information, *see Chapter 25, "3D Compositing Basics,"* in the Fusion Reference Manual.

Inputs

There are five inputs on the Reflect node that accept 2D images or 3D materials. These inputs control the overall color and image used for the 3D object as well as controlling the color and texture used in the reflective highlights.

Background Material: The orange Background material input accepts a 2D image or a 3D material. If a 2D image is provided, the node treats it as a diffuse texture map applied to a basic material.

Reflection Color Material: The white Reflection Color material input accepts a 2D image or a 3D material. The RGB channels are used as the reflection texture, and the Alpha is ignored.

Reflection Intensity Material: The white Reflection Intensity material input accepts a 2D image or a 3D material. The Alpha channel of the texture is multiplied by the intensity of the reflection.

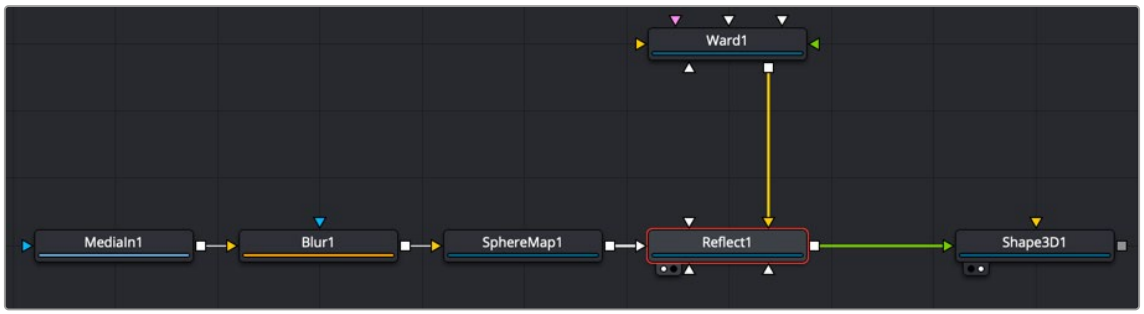
Refraction Tint Material: The white Refraction Tint material input accepts a 2D image or a 3D material. The RGB channels are used as the refraction texture.

Bump Map Texture: The white Bump Map texture input accepts only a 3D material. Typically, you connect the texture into a Bump Map node, and then connect the Bump Map node to this input. This input uses the RGB information as texture-space normals.

When nodes have as many inputs and some using the same color as this one does, it is often difficult to make connections with any precision. Hold down the Option or Alt key while dragging the output from another node over the node tile, and keep holding Option or Alt when releasing the left mouse button. A small drop-down menu listing all the inputs provided by the node appears. Click on the desired input to complete the connection.

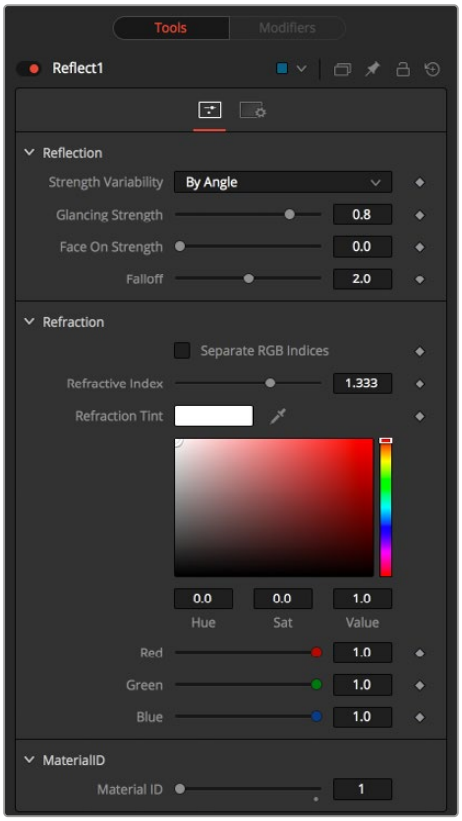
Basic Node Setup

The Reflection node can be the main shader for an object as it is in the example below, or it can be used to feed the diffuse material input of a Ward, Blinn, Phong, or other material node. Usually, a Sphere Map node is used as the source of the Reflect node's reflection color input.



A Reflect node used to create a highly reflective surface of a Shape 3D node

Inspector



Reflect controls

Controls Tab

The Controls tab contains parameters for adjusting the reflective strength based on the orientation of the object, as well as the tint color of the Reflect shader node.

Reflection

Reflection Strength Variability

This multi-button control can be set to Constant or By Angle for varying the reflection intensity, corresponding to the relative surface orientation to the viewer. The following three controls are visible only when this control is set to By Angle.

Glancing Strength

[By Angle] Glancing Strength controls the intensity of the reflection for those areas of the geometry where the reflection faces away from the camera.

Face On Strength

[By Angle] Face On Strength controls the intensity of the reflection for those parts of the geometry that reflect directly back to the camera.

Falloff

[By Angle] Falloff controls the sharpness of the transition between the Glancing and Face On Strength regions. It can be considered similar to applying gamma correction to a gradient between the Face On and Glancing values.

Constant Strength

[Constant Angle] This control is visible only when the reflection strength variability is set to Constant. In this case, the intensity of the reflection is constant despite the incidence angle of the reflection.

Refraction

If the incoming background material has a lower opacity than 1, then it is possible to use an environment map as refraction texture, and it is possible to simulate refraction effects in transparent objects.

Separate RGB Refraction Indices

When this checkbox is enabled, the Refraction Index slider is hidden, and three sliders for adjusting the refraction index of the Red, Green, and Blue channels appear in its place. This allows simulation of the spectral refraction effects commonly seen in thick imperfect glass, for example.

Refraction Index

This slider controls how strongly the environment map is deformed when viewed through a surface. The overall deformation is based on the incidence angle. Since this is an approximation and not a simulation, the results are not intended to model real refractions accurately.

Refraction Tint

The refraction texture is multiplied by the tint color for simulating color-filtered refractions. It can be used to simulate the type of coloring found in tinted glass, as seen in many brands of beer bottles, for example.

Common Controls

Settings Tab

The Settings tab in the Inspector is duplicated in other 3D nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Stereo Mix [3SMM]



The Stereo Mix node

Stereo Mix Node Overview

This node is used to swap the left and right material inputs. It is often used to output to the left and right eye of the 3D Render.

Inputs

This node has two inputs that are both required for this node to work. Both inputs accept either a 2D image or a 3D material.

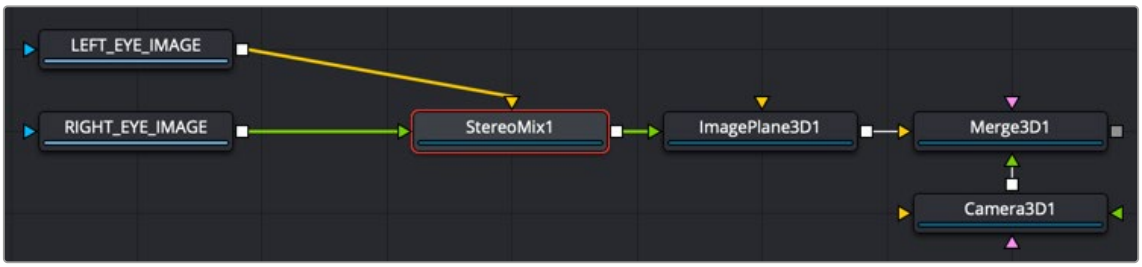
LeftMaterial: The orange left material input accepts a 2D image or a 3D material to be used as the material for the left eye rendering. If a 2D image is used, it is converted to a diffuse texture map using the basic material type.

RightMaterial: The green right material input accepts a 2D image or a 3D material to be used as the material for the right eye rendering. If a 2D image is used, it is converted to a diffuse texture map using the basic material type.

While the inputs can be either 2D images or 3D materials, the output is always a material.

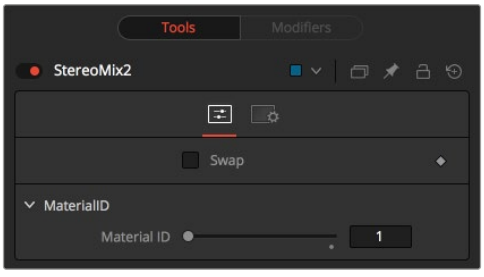
Basic Node Setup

The Stereo Mix node can be used with either stereo images or materials. The example below shows two images combined in the Stereo Mix node causing the output to be a stereo anaglyph material.



A Stereo Mix node used to combine left and right images into a single stereo material

Inspector



Stereo Mix controls

Controls Tab

The Controls tab contains a single switch that swaps the left and right material inputs.

Swap

This option swaps both inputs of the node.

Material ID

This slider sets the numeric identifier assigned to this material. This value is rendered into the MatID auxiliary channel if the corresponding option is enabled in the renderer.

Common Controls

Settings Tab

The Settings tab in the Inspector is duplicated in other 3D nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Ward [3Wd]



The Ward node

Ward Node Introduction

The Ward node is a basic illumination material that can be applied to geometry in the 3D scene. It describes how the object responds to light and provides multiple texture map inputs to allow fine control over the diffuse, specular, and bump map components of the material.

Specifically, the Ward node is ideal for simulating brushed metal surfaces, as the highlight can be elongated along the U or V directions of the mapping coordinates. This is known as an anisotropic highlight.

The Ward node outputs a 3D Material that can be connected to the material inputs on any 3D geometry node.

Inputs

There are six inputs on the Ward node that accept 2D images or 3D materials. These inputs control the overall color and image used for the 3D object as well as controlling the color and texture used in the specular highlight. Each of these inputs multiplies the pixels in the texture map by the equivalently named parameters in the node itself. This provides an effective method for scaling parts of the material.

Diffuse Material: The orange Diffuse material input accepts a 2D image or a 3D material to be used as a main color and texture of the object.

Specular Color Material: The green Specular Color material input accepts a 2D image or a 3D material to be used as a highlight color and texture of the object.

Specular Intensity Material: The magenta Specular Intensity material input accepts a 2D image or a 3D material to be used as an intensity map for the material's highlights. When the input is a 2D image, the Alpha channel is used to create the map, while the color channels are discarded.

Spread U Material: The white Spread U material input accepts a 2D image or a 3D material. The value of the Spread U option in the node's controls is multiplied against the pixel values in the material's Alpha channel.

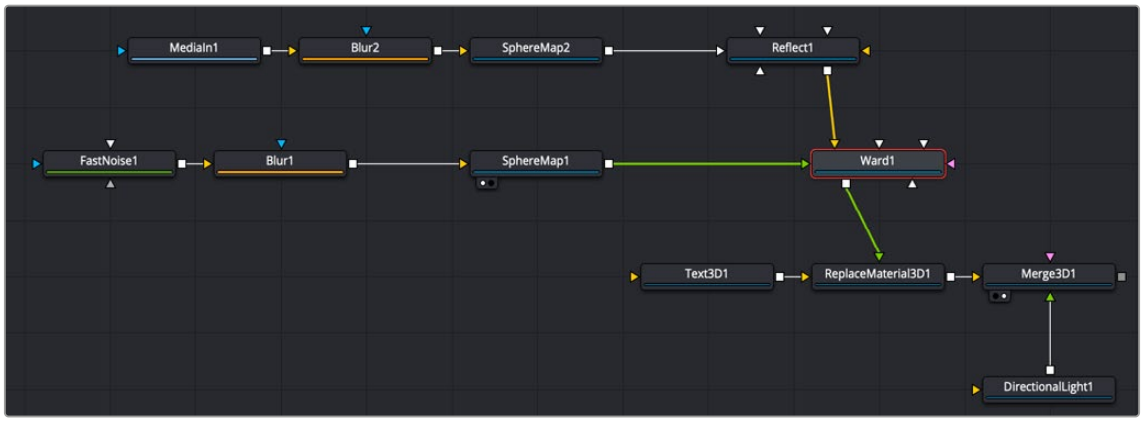
Spread V Material: The white Spread V material input accepts a 2D image or a 3D material. The value of the Spread V option in the node's controls is multiplied against the pixel values in the material's Alpha channel.

Bump Map Material: The white Bump Map material input accepts only a 3D material. Typically, you connect the texture into a Bump Map node, and then connect the Bump Map node to this input. This input uses the RGB information as texture-space normals.

When nodes have as many inputs and some using the same color as this one does, it is often difficult to make connections with any precision. Hold down the Option or Alt key while dragging the output from another node over the node tile, and keep holding Option or Alt when releasing the left mouse button. A small drop-down menu listing all the inputs provided by the node appears. Click on the desired input to complete the connection.

Basic Node Setup

The Ward node is used to make a shiny glass surface and replace the 3D text material in the example below. A diffuse color material comes from Reflect node, and the specular color is altered by a gradient color Fast Noise node.



A Ward node used with a diffuse connection and specular color connection

Inspector



Ward controls

Controls Tab

The Controls tab contains parameters for adjusting the main color, highlight, and lighting properties of the Ward shader node.

Diffuse

Diffuse describes the base surface characteristics without any additional effects like reflections or specular highlights. Besides defining the base color of an object, the diffuse color also defines the transparency of the object. The Alpha in a diffuse texture map can be used to make portions of the surface transparent.

Diffuse Color

A material's Diffuse Color describes the base color presented by the material when it is lit indirectly or by ambient light. If a diffuse texture map is provided, then the color value provided here is multiplied by the color values in the texture.

Alpha

This slider sets the material's Alpha channel value. This affects diffuse and specular colors equally and affects the Alpha value of the material in the rendered output. If a diffuse texture map is provided, then the Alpha value set here is multiplied by the Alpha values in the texture map.

Opacity

Reducing the material's Opacity decreases the color and Alpha values of the specular and diffuse colors equally, making the material transparent.

Specular

The parameters in the Specular section describe the look of the specular highlight of the surface. These values are evaluated in a different way for each illumination model.

Specular Color

Specular Color determines the color of light that reflects from a shiny surface. The more specular a material is, the glossier it appears. Surfaces like plastics and glass tend to have white specular highlights, whereas metallic surfaces like gold have specular highlights that inherit their color from the material color. If a specular texture map is provided, then the value provided here is multiplied by the color values from the texture.

Specular Intensity

Specular Intensity controls how strong the specular highlight is. If the specular intensity texture is provided, then this value is multiplied by the Alpha value of the texture.

Spread U

Spread U controls the falloff of the specular highlight along the U-axis in the UV map of the object. The smaller the value, the sharper the falloff, and the smoother and glossier the material appears in this direction. If the Spread U texture is provided, then this value is multiplied by the Alpha value of the texture.

Spread V

Spread V controls the falloff of the specular highlight along the V-axis in the UV map of the object. The smaller the value, the sharper the falloff, and the smoother and glossier the material appear in this direction. If the Spread V texture is provided, then this value is multiplied by the Alpha value of the texture.

Transmittance

Transmittance controls the way light passes through a material. For example, a solid blue sphere casts a black shadow, but one made of translucent blue plastic would cast a much lower density blue shadow.

There is a separate Opacity option. Opacity determines how transparent the actual surface is when it is rendered. Fusion allows adjusting both opacity and transmittance separately. At first, this might be a bit counterintuitive to those who are unfamiliar with 3D software. It is possible to have a surface that is fully opaque but transmits 100% of the light arriving upon it, effectively making it a luminous/emissive surface.

Attenuation

Attenuation determines how much color is passed through the object. For an object to have transmissive shadows, set the attenuation to (1, 1, 1), which means 100% of green, blue, and red light passes through the object. Setting this color to RGB (1, 0, 0) means that the material transmits 100% of the red arriving at the surface but none of the green or blue light. This can be used to create “stained glass”-styled shadows.

Alpha Detail

When the Alpha Detail slider is set to 0, the Alpha channel of the object is ignored, and the entire object casts a shadow. If it is set to 1, the Alpha channel determines what portions of the object cast a shadow.

Color Detail

The Color Detail slider modulates light passing through the surface by the diffuse color + texture colors. Use this to throw a shadow that contains color details of the texture applied to the object. Increasing the slider from 0 to 1 brings in more diffuse color + texture color into the shadow. Note that the Alpha and opacity of the object are ignored when transmitting color, allowing an object with a solid Alpha to still transmit its color to the shadow.

Saturation

The Saturation slider controls the saturation of the color component transmitted to the shadow. Setting this to 0.0 results in monochrome shadows.

Receives Lighting/Shadows

These checkboxes control whether the material is affected by lighting and shadows in the scene. If turned off, the object is always fully lit and/or unshadowed.

Two-Sided Lighting

This effectively makes the surface two sided by adding a second set of normals facing the opposite direction on the backside of the surface. This is normally off to increase rendering speed, but it can be turned on for 2D surfaces or for objects that are not fully enclosed, to allow the reverse or interior surfaces to be visible as well.

Normally, in a 3D application, only the front face of a surface is visible and the back face is culled, so that if a camera were to revolve around a plane in a 3D application, when it reached the backside, the plane would become invisible. Making a plane two sided in a 3D application is equivalent to adding another plane on top of the first but rotated by 180 degrees so the normals are facing the opposite direction on the backside. Thus, when you revolve around the back, you see the second image plane, which has its normals facing the opposite way.

Fusion does exactly the same thing as 3D applications when you make a surface two sided. The confusion about what two-sided lighting does arises because Fusion does not cull back-facing polygons by default. If you revolve around a one-sided plane in Fusion you still see it from the backside (but you are seeing the frontside duplicated through to the backside as if it were transparent). Making the plane two sided effectively adds a second set of normals to the backside of the plane.

> **NOTE:** This can become rather confusing once you make the surface transparent, as the same rules still apply and produce a result that is counterintuitive. If you view from the frontside a transparent two-sided surface illuminated from the backside, it looks unlit.

Material ID

This slider sets the numeric identifier assigned to this material. This value is rendered into the MatID auxiliary channel if the corresponding option is enabled in the renderer.

Common Controls

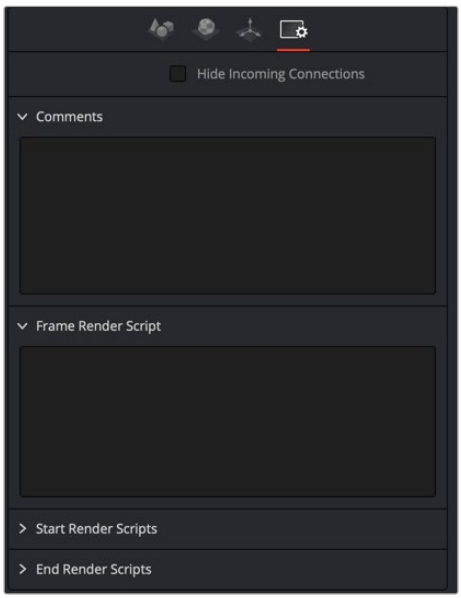
Settings Tab

The Settings tab in the Inspector is duplicated in other 3D nodes. These common controls are described in detail in the following “The Common Controls” section.

The Common Controls

Nodes that handle 3D geometry share a number of identical controls in the Inspector. This section describes controls that are common among 3D Material nodes.

Settings Tab



Common Settings 3D controls

Common Settings tab can be found on most tools in Fusion The following controls are specific settings for 3D nodes

Hide Incoming Connections

Enabling this checkbox can hide connection lines from incoming nodes, making a node tree appear cleaner and easier to read. When enabled, fields for each input on a node are displayed. Dragging a connected node from the node tree into the field hide that incoming connection line as long as the node is not selected in the node tree. When the node is selected in the node tree, the line reappears.

Comment Tab

The Comment tab contains a single text control that is used to add comments and notes to the tool. When a note is added to a tool, a small red dot icon appears next to the setting's tab icon, and a text bubble appears on the node. To see the note in the Node Editor, hold the mouse pointer over the node for a moment. The contents of the Comments tab can be animated over time, if required.

Scripting Tab

The Scripting tab is present on every tool in Fusion. It contains several edit boxes used to add scripts that process when the tool is rendering. For more details on the contents of this tab, please consult the scripting documentation.

Chapter 32

3D Texture Nodes

This chapter details the 3D Texture nodes available when creating 3D composites in Fusion. The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Bump Map [3Bu] 868
Catcher [3Ca] 871
CubeMap [3Cu] 873
Falloff [3Fa] 876
Fast Noise Texture [3FN] 878
Gradient 3D [3Gd] 881
Sphere Map [3SpM] 884
Texture 2D [3Tx] 886
Texture Transform [3TT] 888
The Common Controls 891

Bump Map [3Bu]



The Bump Map node

Bump Map Node Overview

The Bump Map node is used to convert a grayscale (height map) image into a bump map and takes an input directly from a bump map created by the Create Bump Map node. The node outputs a material.

Inputs

The Bump Map node includes a single orange input for connecting a 2D image you want to use as the bump map texture, or it can accept the output of the Create Bump Map node.

ImageInput: The orange Image input is used to connect a 2D RGBA image for the bump calculation or an existing bump map from the Create Bump map node.

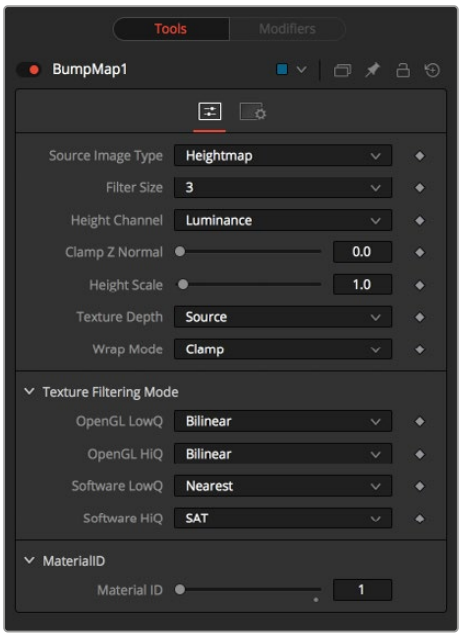
Basic Node Setup

The Bump Map node is connected to the Bump Map material input on any one of the material shader nodes. Below, the example uses a Fast Noise node to generate an image that connects to the Bump Map node. The output of the Bump Map node connects to the Bump Map material input on a Ward node.



A Bump Map is connected to the Bump Map material input on a material node.

Inspector



Bump Map controls

Controls Tab

The Controls tab contains all parameters for modifying the input source and the appearance of the bump map.

Source Image Type

Toggle between Height Map, which creates a bump map similar to the Create Bump Map node, and Bump Map, which expects a bump map created by the Create Bump Map node.

Filter Size

A custom filter generates the bump information. The drop-down menu sets the filter size.

Height Channel

Sets the channel from which to extract the grayscale information.

Clamp Z Normal

Clips the lower values of the blue channel in the resulting bump texture.

Height Scale

Changes the contrast of the resulting values in the bump map. Increasing this value yields a more visible bump map.

Texture Depth

Optionally converts the resulting bump map texture into the desired bit depth.

Wrap Mode

Wraps the image at the borders, so the filter produces correct result when using seamless tile textures.

Common Controls

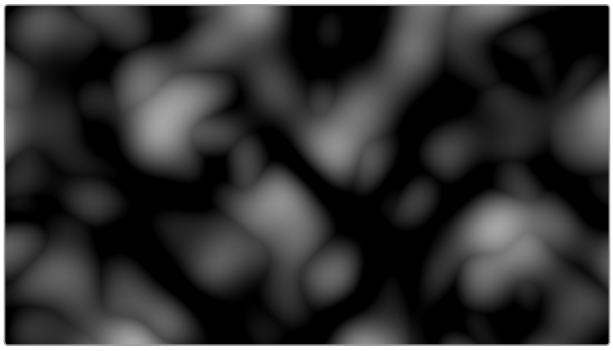
Settings Tab

The Settings tab in the Inspector is duplicated in other 3D nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Notes on Bump Maps

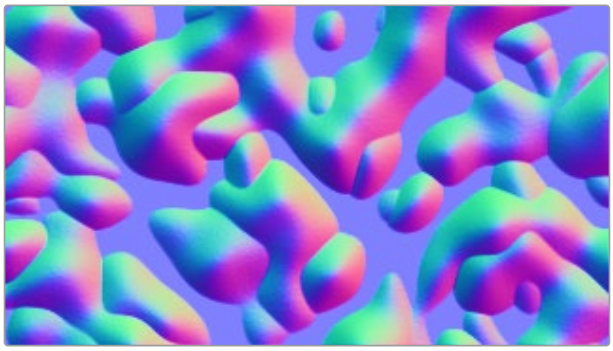
There is some confusion of terminology with bump mapping, depending on where you get your information. Here are Fusion conventions:

Height Map



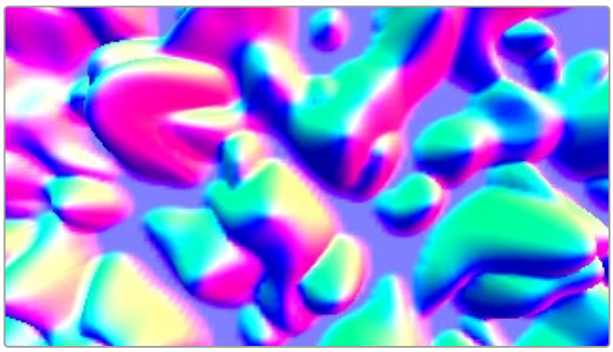
A grayscale image containing a height value per pixel

Bump Map



An image containing normals stored in the RGB channels used for modifying the existing normals (usually given in tangent space)

Normals Map



An image containing normals stored in the RGB channels used for replacing the existing normals (usually given in tangent or object space)

Catcher [3Ca]



The Catcher node

Catcher Node Overview

The Catcher material is used to “catch” texture-mode projections cast from Projector 3D and Camera 3D nodes. The intercepted projections are converted into a texture map and applied by the Catcher material to the geometry to which it is connected.

To understand the Catcher node, it helps to understand the difference between light-based projections and texture-based projections. Choosing Light from the projection mode menu on the Projector 3D or Camera 3D nodes simply adds the values of the RGB channels in the projected image to the diffuse texture of any geometry that lies within the projection cone. This makes it impossible to clip away geometry based on the Alpha channel of an image when using light mode projections.

Imagine a scenario where you want to project an image of a building onto an image plane as part of a set extension shot. You first rotoscope the image to mask out the windows. This makes it possible to see the geometry of the rooms behind the wall in the final composite. When this image is projected as light, the Alpha channel is ignored, so the masked windows remain opaque.

By connecting the Catcher to the diffuse texture map of the material applied to the image plane, and then switching the projection mode menu in the Projector 3D or Camera 3D node from Light or Ambient Light mode to Texture mode, the projected image is applied as a texture map. When using this technique for the example above, the windows would become transparent, and it would be possible to see the geometry behind the window.

The main advantages of this approach over light projection are that the Catcher can be used to project Alpha onto an object, and it doesn‘t require lighting to be enabled. Another advantage is that the Catcher is not restricted to the diffuse input of a material, making it possible to project specular intensity maps, or even reflection and refraction maps.

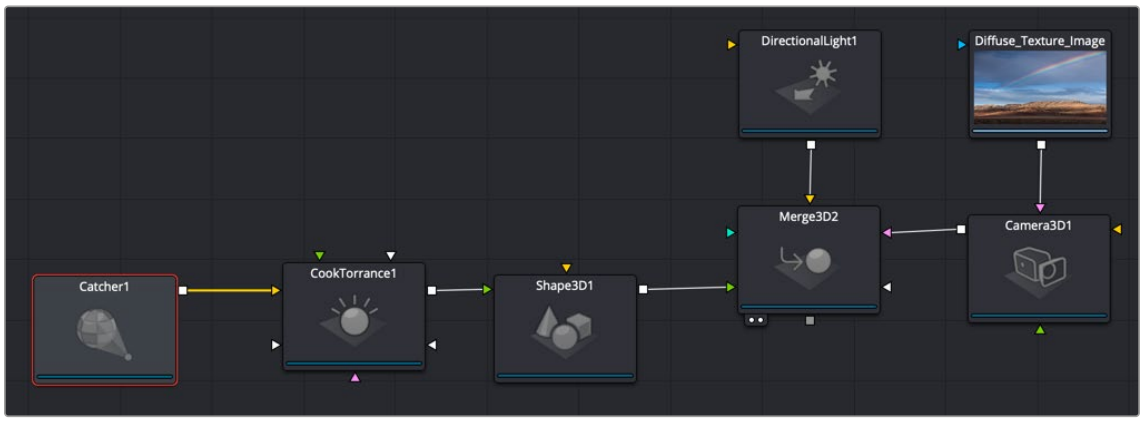
NOTE: The Catcher material requires a Projector 3D or Camera 3D node in the scene, set to project an image in Texture mode on the object to which the Catcher is connected. Without a projection, or if the projection is not set to Texture mode, the Catcher simply makes the object transparent and invisible.

Inputs

The Catcher node has no inputs. The output of the node is connected to the diffuse color material input of the Blinn, Cook Torrance, or other material node applied to the 3D geometry.

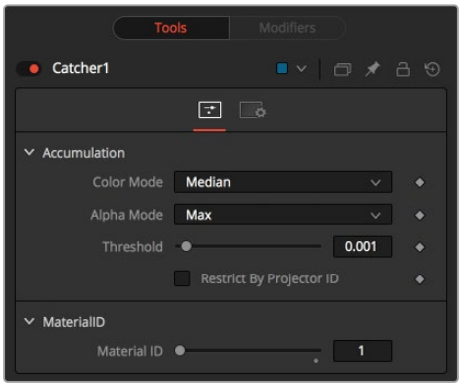
Basic Node Setup

The output of a Catcher node should be connected to the material input of your 3D geometry node. A camera is set up as a proctor with an image connected to the camera's image input. When the camera is set to texture projection mode, the Catcher node is used to determine which geometry receives the texture.



A Catcher node output is connected to the input of the geometry node that receives the texture projection

Inspector



Catcher controls

Controls Tab

The Options in the Controls tab determine how the Catcher handles the accumulation of multiple projections.

Enable

Use this checkbox to enable or disable the node. This is not the same as the red switch in the upper-left corner of the Inspector. The red switch disables the tool altogether and passes the image on without any modification. The Enable checkbox is limited to the effect part of the tool. Other parts, like scripts in the Settings tab, still process as normal.

Color Mode

The Color mode menu is used to control how the Catcher combines the light from multiple projectors. It has no effect on the results when only one projector is in the scene. This control is designed to work with the software renderer in the Renderer 3D node and has no effect when using the OpenGL renderer.

Alpha Mode

The Alpha mode is used to control how the Catcher combines the Alpha channels from multiple projectors. It has no effect on the results when only one projector is in the scene. This control is designed to work with the software renderer in the Renderer 3D node and has no effect when using the OpenGL renderer.

Threshold

The Threshold can be used to exclude certain low values from the accumulation calculation. For example, when using the Median Accumulation mode, a threshold of 0.01 would exclude any pixel with a value of less than 0.01 from the median calculation.

Restrict by Projector ID

When active, the Catcher only receives light from projectors with a matching ID. Projectors with a different ID are ignored.

Material ID

This slider sets the numeric identifier assigned to this material. This value is rendered into the MatID auxiliary channel if the corresponding option is enabled in the Renderer 3D node.

Common Controls

Settings Tab

The Settings tab in the Inspector is duplicated in other 3D nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

CubeMap [3Cu]



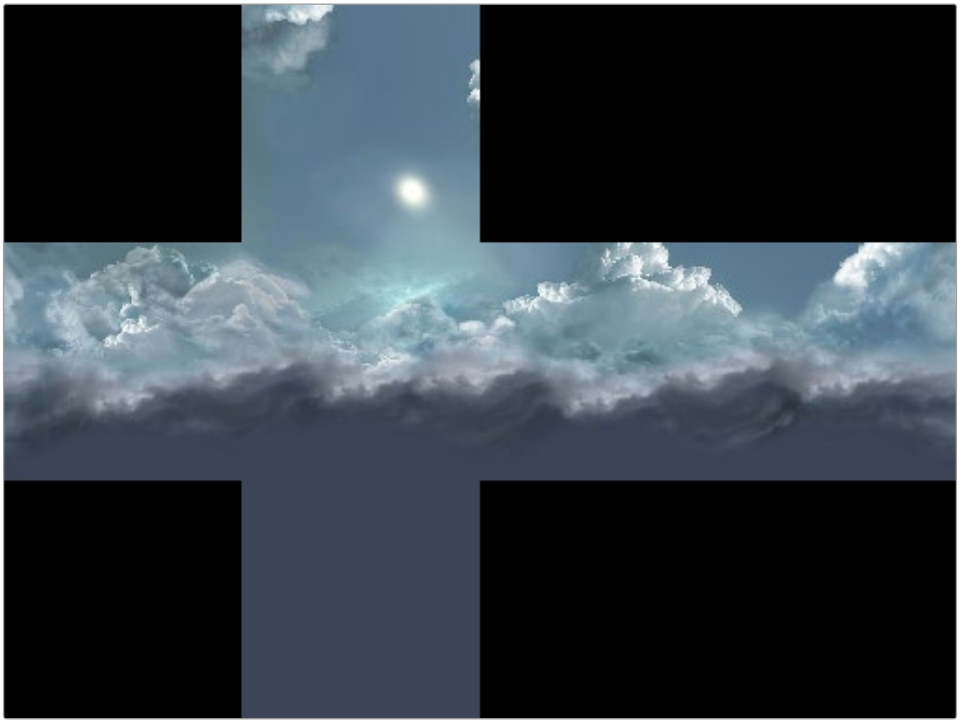
The Cube Map node

Cube Map Node Overview

The Cube Map node creates texture maps using separate images for each face of the cube. It can also extract the individual faces of the cube from a single image containing an unfolded cube in the Vertical or Horizontal Cross layouts.

A cube map is produced by mounting six cameras at 90 degrees angle of views to point up, down, left, right, front, and back.

The node provides options to set the reference coordinate system and rotation for the resulting texture map. The Cube Map node is typically used to produce environment maps for distant areas (such as skies or horizons) or reflection and refraction maps.



Sample cube map

Inputs

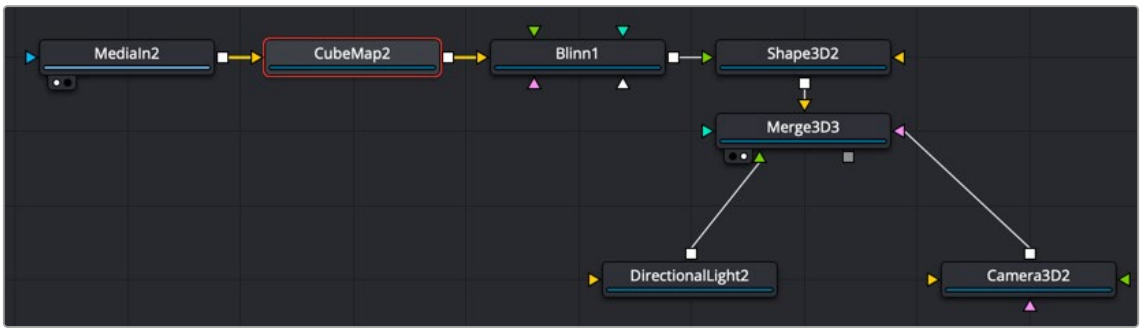
The Inputs on this node change based on the settings of the Layout menu in the Inspector. The single input uses a 2D image for the entire cube, while six inputs can handle a different 2D image for each side of a cube.

CrossImage: The orange Cross Image input is visible by default or when the Layout menu in the Inspector is set to either Vertical Cross or Horizontal Cross. The input accepts a 2D image.

CubeMap.[DIRECTION]: These six multi-colored inputs are visible only when the Layout menu in the Inspector is set to Separate Images. Each input accepts an image aligned to match the left, right, top, bottom, front, and back faces.

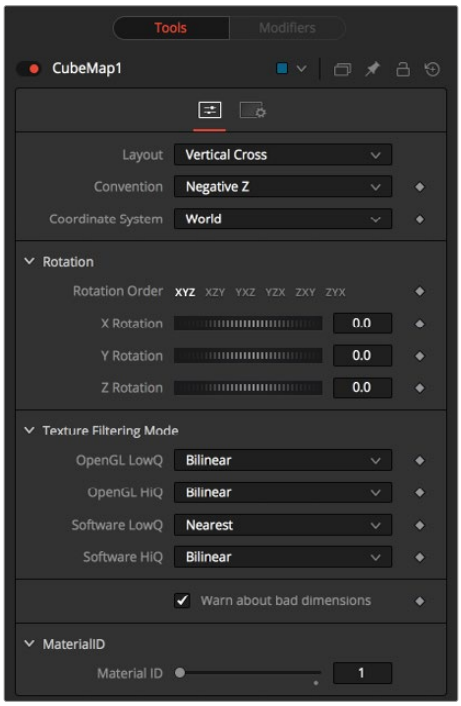
Basic Node Setup

The Cube Map node uses a vertical or horizontal cross image represented by MediaIn2 node connected into the orange cross image input. The Cube Map node is used similarly to the Sphere Map node. It creates an environment that surrounds the geometry connected to a Shader node.



A Cube Map node receives a cross image input, creating an environment for the Shape 3D

Inspector



Cube Map controls

Controls Tab

Layout

The Layout menu determines the type and number of inputs for the cube map texture.
Valid options are:

- **Separate Images:** This option exposes six inputs on the node, one for each face of the cube. If the separate images are not square or not of the same size, they are rescaled into the largest 1:1 image that can contain all of them.
- **Vertical Cross:** This option exposes a single input on the node. The image should be an unwrapped texture of a cube containing all the faces organized into a Vertical Cross formation, where the height is larger than the width. If the image aspect of the cross image is not 3:4, the CubeMap node crops it down so it matches the applicable aspect ratio.
- **Horizontal Cross:** This option exposes a single input on the node. The image should be an unwrapped texture of a cube containing all the faces organized into a Horizontal Cross formation, where the width is larger than the height. If the image aspect of the cross image is not 4:3, the CubeMap node crops it down so that matches the applicable aspect ratio.

Coordinate System

The coordinate system menu sets the position values used when converting the image into a texture.

- **Model:** This option orients the texture along the object local coordinate system.
- **World:** This option orients the resulting texture using the global or world coordinate system.
- **Eye:** This option aligns the texture map to the coordinate system of the camera or viewer.

Rotation

The rotation controls are divided into buttons that select the order of rotation along each axis of the texture. For example, XYZ would apply the rotation to the X axis first, followed by the Y axis, and finally the Z axis. The other half of the rotation controls are dials that rotate the texture around its pivot point.

Warn About Bad Dimensions

Selecting this checkbox displays a warning message on the console if the dimensions of the image provided did not meet the requirements of the selected orientation mode.

Material ID

This slider sets the numeric identifier assigned to this material. This value is rendered into the MatID auxiliary channel if the corresponding option is enabled in the renderer.

Common Controls

Settings Tab

The Settings tab in the Inspector is duplicated in other 3D nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

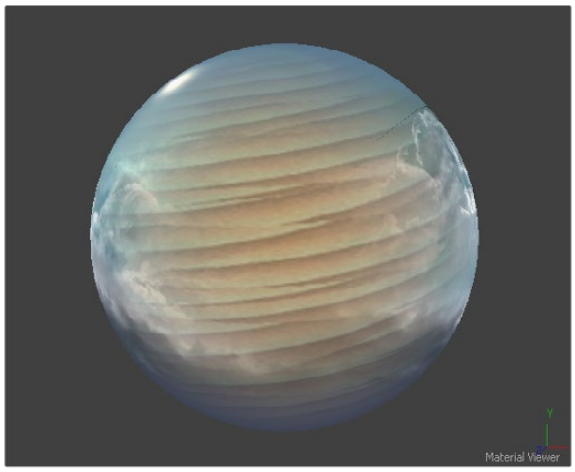
Falloff [3Fa]



The Falloff node

Falloff Node Overview

The Falloff node blends two materials or textures together based on the incidence angle between the object to which the material is applied and the camera. This is useful when you wish to use one material for portions of the geometry that would reflect light directly back to the camera and a different material for parts that reflect light back into the scene.



Falloff example

Inputs

The two Inputs on the Falloff node are used to connect two images or materials. One is used to reflect back at the camera, while the other reflects away from the camera and into the scene.

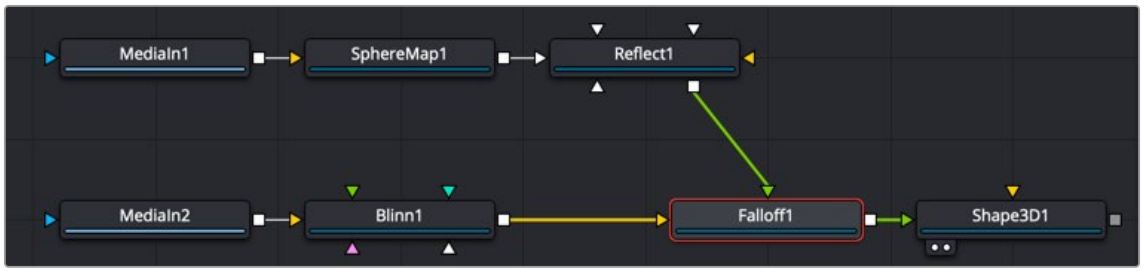
Face On Material: The orange Face On material input accepts a 2D image or a 3D material. If a 2D image is provided, it is turned into a diffuse texture map using the basic material shader. This input is used for the material that is reflecting directly back to the camera

Glancing Material: The green Glancing material input accepts a 2D image or a 3D material. If a 2D image is provided, it is turned into a diffuse texture map using the basic material shader. This input is used for the material that is reflecting away from the camera and into the scene.

While the inputs for this node can be images, the output is always a material.

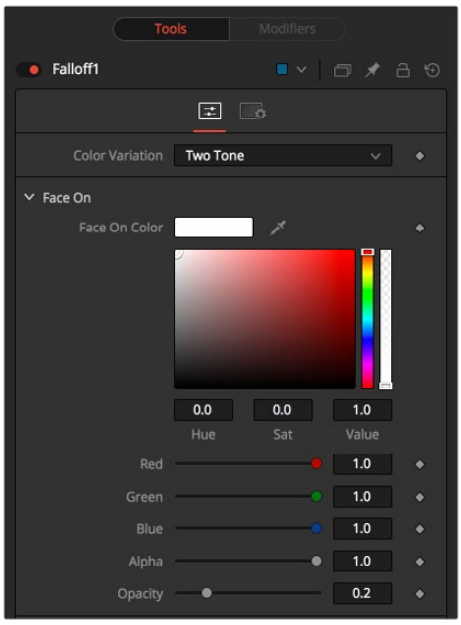
Basic Node Setup

The Falloff node below is used to control the strength of the Blinn material and the Reflect material. You connect the Face On input of the Falloff node to the material you want shown for the sides of the object that face the camera and connect the Glance input to the material you want shown for the sides not directly facing the camera.



The Falloff node uses one input for the material facing the camera and one for the material not directly facing the camera.

Inspector



Falloff controls

Controls Tab

The parameters in the Controls tabs modify the tint and opacity of the Face On material and the Glancing material. A Falloff slider controls the blending between the two.

Color Variation

- **Two Tone:** Two regular Color controls define the colors for Glancing and Face On.
- **Gradient:** A Gradient control defines the colors for Glancing and Face On. This can be used for a multitude of effects, like creating Toon Shaders, for example.

Face On Color

The Face On Color defines the color of surface parts facing the camera. If the Face On texture map is provided, then the color value provided here is multiplied by the color values in the texture.

Reducing the material's opacity decreases the color and Alpha values of the Face On material, making the material transparent.

Glancing Color

The Glancing Color defines the color of surface parts more perpendicular to the camera. If the Glancing material port has a valid input, then this input is multiplied by this color.

Reducing the material's opacity decreases the color and Alpha values of the Glancing material, making the material transparent.

Falloff

This value controls the transition between Glancing and Face On strength. It is very similar to a gamma operation applied to a gradient, blending one value into another.

Material ID

This slider sets the numeric identifier assigned to this material. This value is rendered into the MatID auxiliary channel if the corresponding option is enabled in the renderer.

Common Controls

Settings Tab

The Settings tab in the Inspector is duplicated in other 3D nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Fast Noise Texture [3FN]



The Fast Noise Texture node

Fast Noise Texture Node Overview

The Fast Noise Texture node is the procedural resolution-independent version of the 2D Fast Noise node. It creates a noise texture directly as a material for usage with 3D nodes. It offers a 3D volumetric mode for creating seamless textures in conjunction with nodes providing UVW texture coordinates (Similar to the UV Map 3D node set to XYZ-to-UVW or Camera).

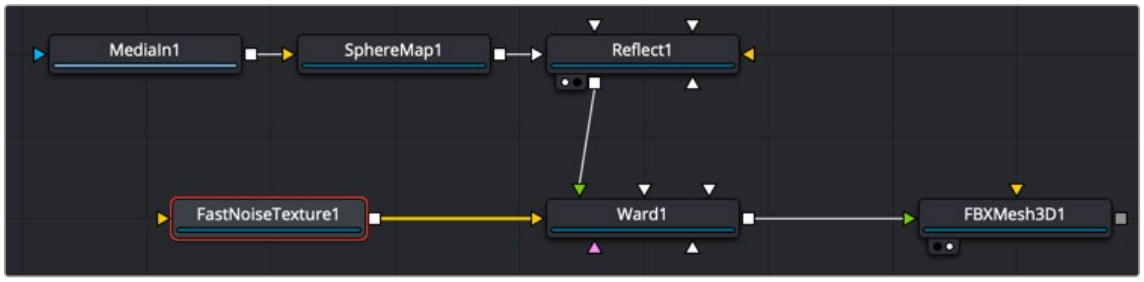
Inputs

The Fast Noise Texture node includes an optional input that can be used to connect a 2D image or material.

SourceMaterial: The Source Materials input accepts a 2D image or a 3D material. The image is then altered by the noise pattern.

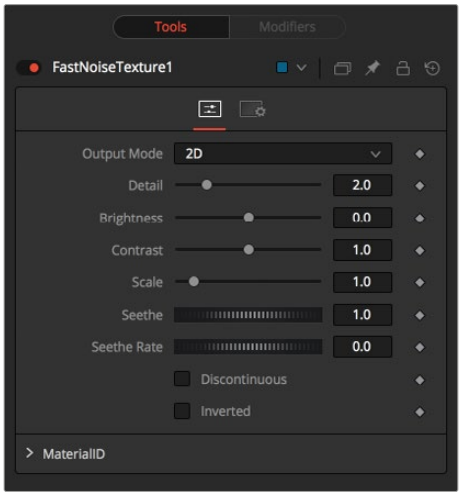
Basic Node Setup

The Fast Noise Texture node below is used to generate a resolution-independent 3D texture for an FBX imported model.



A Fast Noise Texture node generates a seamless texture, taking advantage of UVW coordinates.

Inspector



Fast Noise Texture controls

Controls Tab

The parameters of the Fast Noise Texture node control the appearance and, for 2D, the animation of the noise.

Output Mode

- **2D:** Calculates the noise texture based on 2D texture coordinates (UV). This setting allows smoothly varying the noise pattern with animation.
- **3D:** Calculates the noise texture based on 3D texture coordinates (UVW). Nodes like Shape 3D automatically provide a third texture coordinate; otherwise, a 3D texture space can be created using the UV Map node. The 3D setting does not support animation of the noise pattern.

Detail

Increase the value of this slider to produce a greater level of detail in the noise result. Larger values add more layers of increasingly detailed noise without affecting the overall pattern. High values take longer to render but can produce a more natural result (not all graphics cards support higher detail levels in hardware).

Brightness

This control adjusts the overall Brightness of the noise map.

Contrast

This control increases or decreases the overall Contrast of the noise map. It can exaggerate the effect of the noise.

Scale

The scale of the noise map can be adjusted using the Scale slider, changing it from gentle variations over the entire image to a tighter overall texture effect. This value represents the scale along the UV axis.

Scale Z

(3D only) The Scale Z value scales the noise texture along the W-axis in texture space. W represents a direction perpendicular to the UV plane for a 3D texture map.

Seethe

(2D only) The Seethe control smoothly varies the 2D noise pattern.

Seethe Rate

(2D only) As with the Seethe control above, the Seethe Rate also causes the noise map to evolve and change. The Seethe Rate defines the rate at which the noise changes each frame, causing an animated drift in the noise automatically, without the need for spline animation.

Discontinuous

Normally, the noise function interpolates between values to create a smooth continuous gradient of results. You can enable the Discontinuous checkbox to create hard discontinuity lines along some of the noise contours. The result is a dramatically different effect.

Invert

Enable the Invert checkbox to invert the noise, creating a negative image of the original pattern. This is most effective when Discontinuous is also enabled.

Material ID

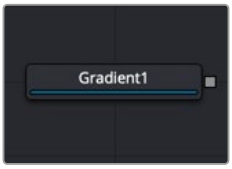
This slider sets the numeric identifier assigned to this material. This value is rendered into the MatID auxiliary channel if the corresponding option is enabled in the renderer.

Common Controls

Settings Tab

The Settings tab in the Inspector is duplicated in other 3D nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Gradient 3D [3Gd]



The Gradient node

Gradient Node Overview

The Gradient 3D node is used to texture objects with a variety of gradient types. It offers many of the same controls as the Background node. While it is not possible to transform the gradient directly in 3D space, it is orientable using the following nodes:

Texture Transform Node: The Texture Transform node can be used to adjust the mapping per pixel.

UV Map Node: The UV Map node can be used to adjust the mapping per vertex (use the XYZtoUVW mode). This has onscreen controls, so you can see what the gradient is doing. Using this node is recommended because it is faster to evaluate.

The gradient defaults to a linear gradient that goes from -1 to +1 along the Z-axis. All primitives in the Shape 3D node can output a third texture coordinate for UVW mapping.

Inputs

The Gradient node has no Inputs. The output of the node is connected to a material input on 3D geometry.

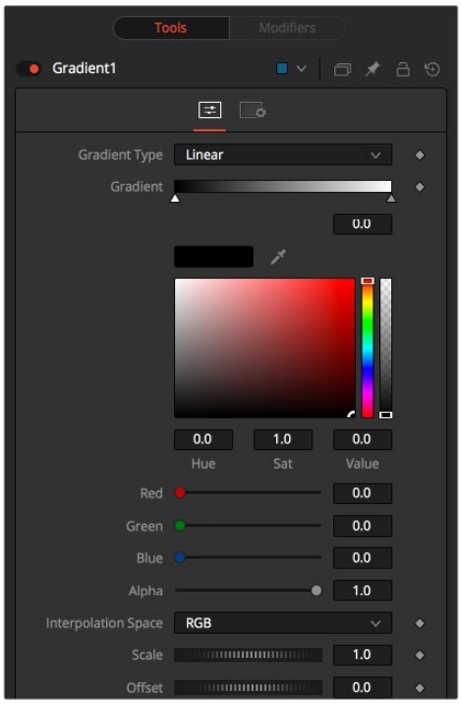
Basic Node Setup

The Gradient 3D node below is used to generate a resolution-independent 3D texture for an FBX imported model. Positioning in UVW space is easiest to do using a UV Map tool placed after the geometry.



A Gradient 3D node generates a resolution-independent gradient texture positioned by the UV Map tool

Inspector



Gradient 3D controls

Controls Tab

The Controls tab for the Gradient node control the pattern and colors used for the gradient texture.

Gradient Type

Determines the type or pattern used for the gradient.

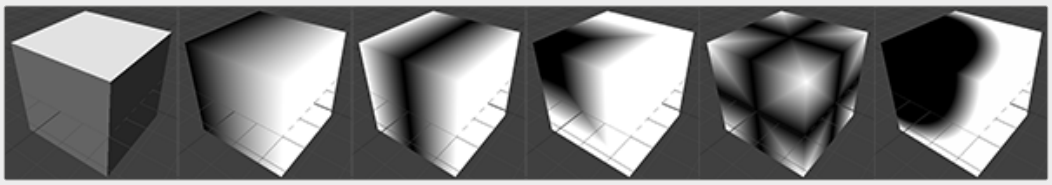
Linear: A simple linear gradient.

Reflect: Based on the Linear mode, this gradient is mirrored at the middle of the textured range.

Square: The gradient is applied using a square pattern.

Cross: Similar to the Reflect mode, but Cross uses two axes to apply the gradient.

Radial: The Radial mode uses a circular pattern to apply the gradient.



Gradient 3D modes

Gradient Bar

The Gradient control consists of a bar where it is possible to add, modify, and remove color stops of the gradient. Each triangular color stop on the Gradient bar represents a color in the gradient. It is possible to animate the color as well as the position of the point. Furthermore, a From Image modifier can be applied to the gradient to evaluate it from an image.

Interpolation Space

The gradient is linearly interpolated from point to point in RGB color space by default. This can sometimes lead to unwanted colors. Choosing another color space may provide a better result.

Scale

Allows sizing of the gradient.

Offset

Allows panning through the gradient.

Repeat

Defines how the left and right borders of the gradient are treated.

Once: When using the Gradient Offset control to shift the gradient, the border colors keep their values. Shifting the default gradient to the left results in a white border on the left, while shifting it to the right results in a black border on the right.

Repeat: When using the Gradient Offset control to shift the gradient, the border colors wrap around. Shifting the default gradient to the left results in a sharp jump from white to black, while shifting it to the right results in a sharp jump from black to white.

Ping Pong: When using the Gradient Offset control to shift the gradient, the border colors ping-pong back and forth. Shifting the default gradient to the left results in the edge fading from white back to black, while shifting it to the right results in the edge fading from black back to white.



Gradients set to Once, Repeat, and Ping Pong from top to bottom, respectively, and shifting the gradient to the left

Sub Pixel

Determines the accuracy with which the gradient is created.

Material ID

This slider sets the numeric identifier assigned to this material. This value is rendered into the MatID auxiliary channel if the corresponding option is enabled in the renderer.

Common Controls

Settings Tab

The Settings tab in the Inspector is duplicated in other 3D nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Sphere Map [3SpM]



The SphereMap node

Sphere Map Node Overview

The Sphere Map node can be used to create simulated environment mapping, also called reflection mapping. Ray trace rendering a reflective scene can be very time consuming, but sphere map-based reflection mapping can generate 360-degree reflections faster with little loss of accuracy. For example, when creating a reflective environment, a sphere map is created, large enough to surround the 3D object in your scene. The sphere is mapped with the environment you want reflected and connected to the Reflection Color input on a Reflect node.

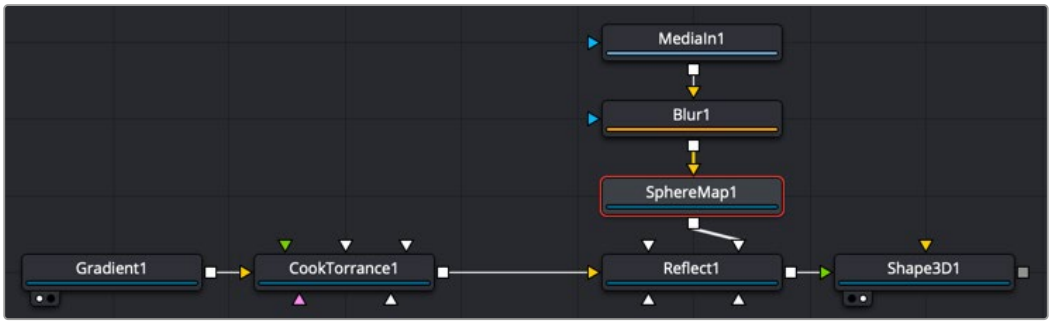
Inputs

The single image input on the Sphere Map node accepts a 2D image texture in an equirectangular format (where the X-axis represents 0–360 degrees longitude, and the Y-axis represents –90 to +90 degrees latitude.)

ImageInput: The orange Image input accepts a 2D RGBA image. Preferably, this is an equirectangular image that shows the entire vertical and horizontal angle of view up to 360 degrees.

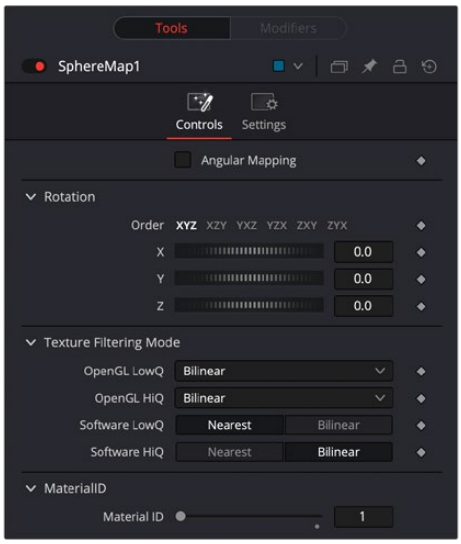
Basic Node Setup

The Sphere Map node below is mapped with a spherical image to generate the environment reflected on the Shape 3D. It is connected to the Reflection Color input on a Reflect node.



A Sphere Map node generates a reflective environment when connected to a Reflect node Reflection Color input.

Inspector



Sphere Map controls

Controls Tab

The Controls tab in the Inspector modifies the mapping of the image input to the sphere map.

Angular Mapping

Adjusts the texture coordinate mapping so the poles are less squashed and areas in the texture get mapped to equal areas on the sphere. It turns the mapping of the latitude lines from a hemispherical fisheye to an angular fisheye. This mapping attempts to preserve area and makes it easier to paint on or modify a sphere map since the image is not as compressed at the poles.

Rotation

Offers controls to rotate the texture map.

Material ID

This slider sets the numeric identifier assigned to this material. This value is rendered into the MatID auxiliary channel if the corresponding option is enabled in the renderer.

The node expects an image with an aspect ratio of 2:1. Otherwise, the image is clamped according to the following rules:

- **2 * width > height:** The width is fitted onto the sphere, and the poles display clamped edges.
- **2 * width < height:** The height is fitted onto the sphere, and there is clamping about the 0-degree longitude line.

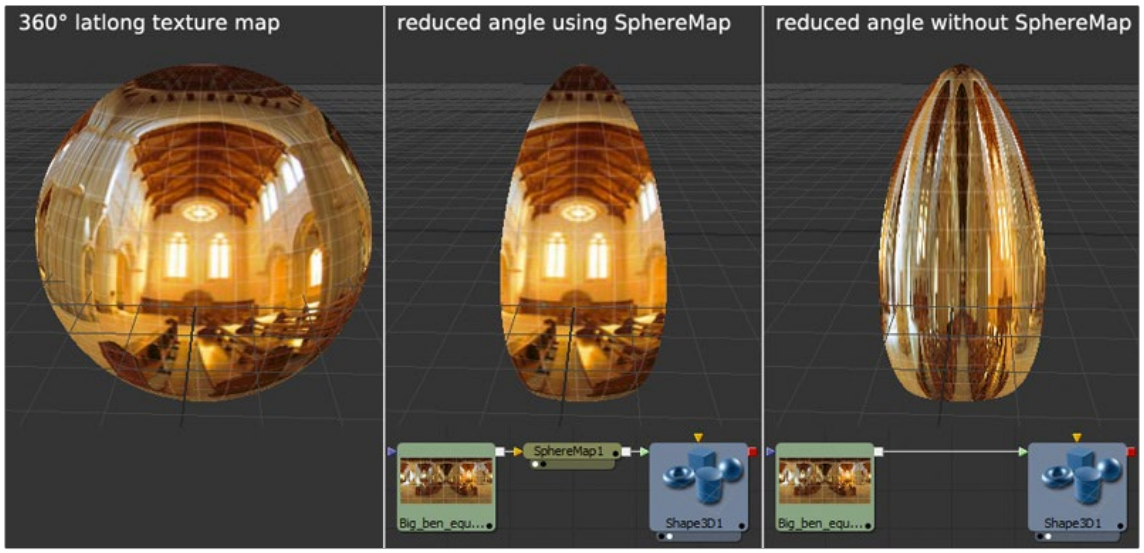
Common Controls

Settings tab

The Settings tab in the Inspector is duplicated in other 3D nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Sphere Map vs. Connecting the Texture to a Sphere Directly

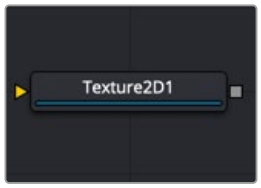
You can connect an equirectangular texture map directly to a sphere instead of piping it through the Sphere Map node first. This results in a different rendering if you set the start/end angle and latitude to less than 360°/180°. In the first case, the texture is squashed. When using the Sphere Map node, the texture is cropped. Compare:



Spherical mapping differences

NOTE: If you pipe the texture directly into the sphere, it is also mirrored horizontally. You can change this by using a Transform node first.

Texture 2D [3Tx]



The Texture 2D node

Texture Node Overview

The Texture 2D node sets metadata of an image being used for a texture map. By default, an image will be (0,0) to (1,1) UV, but that can be changed. The Texture node relies on the presence of U and V Map channels in 3D rendered images. If these channels are not present, this node has no effect.

NOTE: Background pixels may have U and V values of 0.0, which set those pixels to the color of the texture's corner pixel. To restrict texturing to specific objects, use an effect mask based on the Alpha of the object, or its Object or Material ID channel.

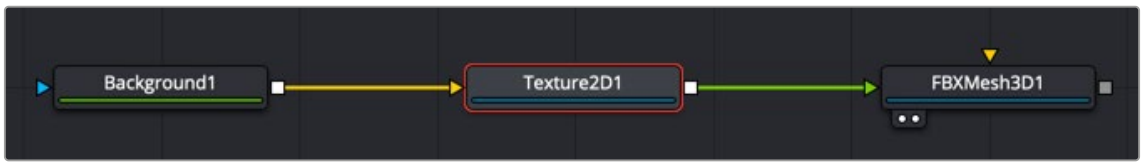
For more information, *see Chapter 18, "Understanding Image Channels,"* in the Fusion Reference Manual.

Inputs

Image Input: The orange image input expects a 2D image.

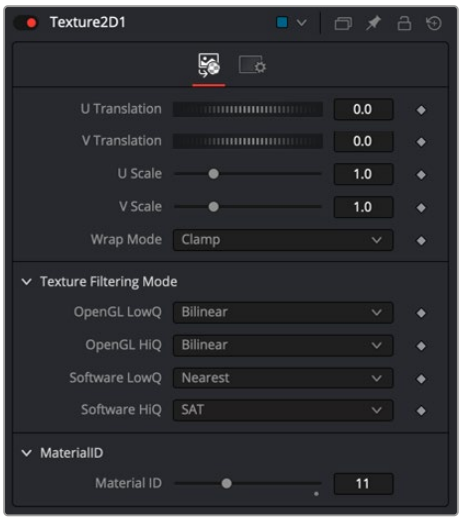
Basic Node Setup

The Texture 2D node below takes a 2D gradient from the Background node and sets the UV metadata for it. The texture is then applied to the FBX geometry based on that metadata. If you have the option to use the UV Map tool, it is recommended because it may be faster and has onscreen controls.



A Texture 2D node is used to set the 3D texture metadata for the input image.

Inspector



Texture 2D controls

Controls Tab

The Controls tab of the Inspector includes the following options.

U/V Offset

These sliders can be used to offset the texture along the U and V coordinates.

U/V Scale

These sliders can be used to scale the texture along the U and V coordinates.

Wrap Mode

If a texture is transformed in the texture space (using the controls below or the UV Map node), then it's possible that areas beyond the image borders will be mapped on the object. The Wrap Mode determines how the image is applied in these areas.

- **Wrap:** This wraps the edges of the image around the borders of the image.
- **Clamp:** The color at the edges of the images is used for texturing. This mode is similar to the Duplicate mode in the Transform node.
- **Black:** The image is clipped along its edges. A black color with Alpha = 0 is used instead.
- **Mirror:** The image is mirrored in both X and Y.

Texture Filtering Mode

The texture can be filtered differently depending on whether you are using the Software Renderer or OpenGL renderer in the Renderer 3D node. Within the two render engines, you can choose between high-quality anti-aliasing or low quality. The texture filtering mode provides different filtering options for the two render engines and the two anti-aliasing settings.

- **Nearest:** The simplest filtering technique is very fast but can cause artifacts when scaling textures.
- **Bilinear:** A standard isotropic filtering technique for scaling textures into multiple resolutions. Works well for magnification of textures.
- **Trilinear:** An extension of Bilinear filtering. Trilinear tends to be a better option when scaling down textures
- **Anisotropic:** The highest-quality filtering method that takes the camera orientation and polygon perspective into account.
- **SAT:** SAT (Summed Area Table) is a method of performing high-quality filtering, but it can require more memory than other options. Works very well on smaller bitmaps.

Common Controls

Settings Tab

The Settings tab in the Inspector is duplicated in other 3D nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Texture Transform [3TT]



The Texture Transform node

Texture Transform Node Overview

The Texture Transform node can be used to translate, rotate, and scale the UVW texture coordinates of a 3D object. While the input can also be an image, the output is always a material.

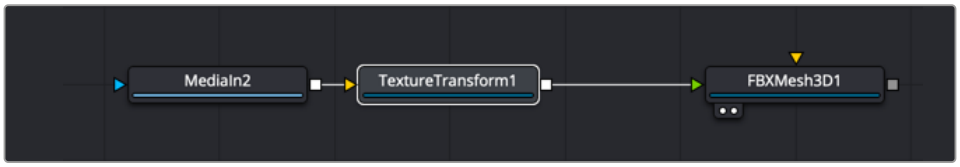
Inputs

The Texture Transform node includes a single input that is used to connect the image or material you want to transform.

Material Input: The orange Material input accepts a 2D image or 3D material whose texture coordinates are transformed using the controls in the Inspector.

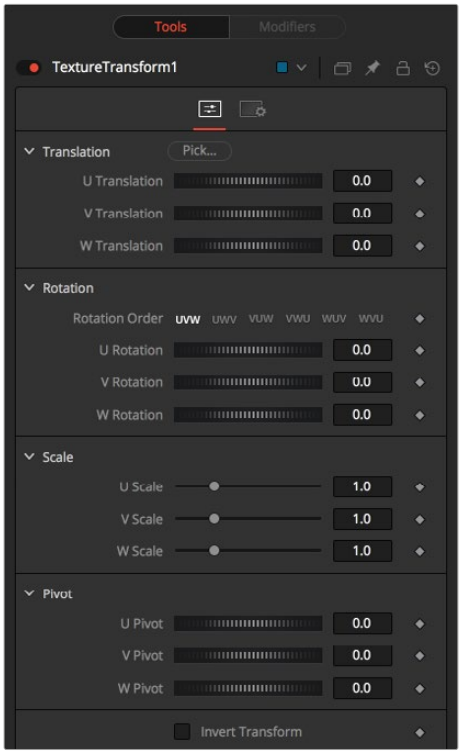
Basic Node Setup

The Texture Transform node below is used to take in a 2D image, transform it, and output a material to be used on 3D geometry.



A Texture Transform node transforms a texture applied to 3D geometry.

Inspector



Texture Transform controls

NOTE: Not all Wrap modes are supported by all graphics cards.

FUSION

Controls Tab

The Controls tab for the Texture Transform node includes many common transform controls that are used to transform the texture using UVW coordinates.

Translation

The U, V, W translation sliders shift the texture along U, V, and W axes.

Rotation

Rotation Order buttons set the order in which the rotation is applied. In conjunction with the buttons, the UVW dials define the rotation around the UVW axes.

Scale

U, V, W sliders scale the texture along the UVW axes.

Pivot

U, V, W Pivot sets the reference point for rotation and scaling.

Material ID

This slider sets the numeric identifier assigned to this material. This value is rendered into the MatID auxiliary channel if the corresponding option is enabled in the renderer.

Common Controls

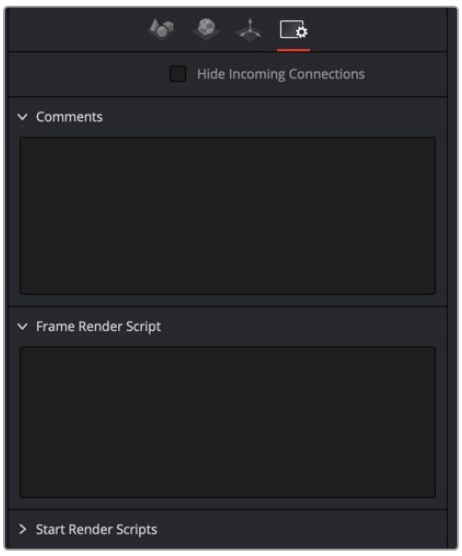
Settings Tab

The Settings tab in the Inspector is duplicated in other 3D nodes. These common controls are described in the following “The Common Controls” section.

The Common Controls

Nodes that handle 3D geometry share a number of identical controls in the Inspector. This section describes controls that are common among 3D Texture nodes.

Settings Tab



Common Settings 3D controls

The Common Settings tab can be found on most tools in Fusion. The following controls are specific settings for 3D nodes.

Hide Incoming Connections

Enabling this checkbox can hide connection lines from incoming nodes, making a node tree appear cleaner and easier to read. When enabled, fields for each input on a node are displayed. Dragging a connected node from the node tree into the field hides that incoming connection line as long as the node is not selected in the node tree. When the node is selected in the node tree, the line reappears.

Comment Tab

The Comment tab contains a single text control that is used to add comments and notes to the tool. When a note is added to a tool, a small red dot icon appears next to the setting's tab icon, and a text bubble appears on the node. To see the note in the Node Editor, hold the mouse pointer over the node for a moment. The contents of the Comments tab can be animated over ime, if required.

Scripting Tab

The Scripting tab is present on every tool in Fusion. It contains several edit boxes used to add scripts that process when the tool is rendering. For more details on the contents of this tab, please consult the scripting documentation.

Chapter 33

Blur Nodes

This chapter details the Blur nodes available in Fusion. The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Blur [Blur] 893
Defocus [Dfo] 896
Directional Blur [DrBl] 898
Glow [Glo] 900
Sharpen [Shrp] 903
Soft Glow [SGlo] 905
Unsharp Mask [UsM] 908
Vari Blur [VBl] 910
Vector Motion Blur [VMB] 912
The Common Controls 914

Blur [Blur]



The Blur node

Blur Node Introduction

The Blur node does exactly what its name implies – it blurs the input image. This is one of the most commonly used image-processing operations.

Inputs

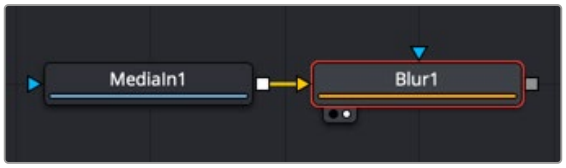
The two inputs on the Blur node are used to connect a 2D image and an effect mask that can be used to limit the blurred area.

Input: The orange input is used for the primary 2D image that is blurred.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the blur to only those pixels within the mask. An effect mask is applied to the tool after the tool is processed.

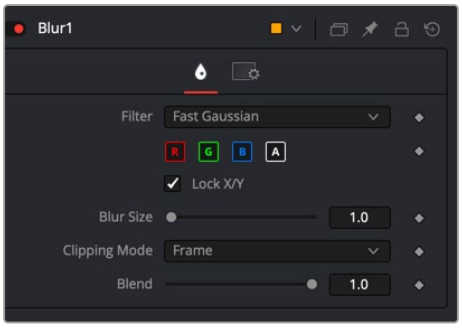
Basic Node Setup

The Blur node, like many 2D image-processing nodes, receives a 2D image like the MediaIn1 shown below. The output continues the node tree by connecting to another 2D image-processing node or a Merge node.



A Blur node applied to a MediaIn1 node

Inspector



Blur controls

NOTE: Since a perfect Gaussian filter would require examining an infinite number of pixels, all practical Gaussians are, of necessity, approximations. The algorithm Fusion uses is a highly-optimized approach that has many strengths, but can create visible ringing around the edges in certain extreme cases. This ringing appears only when blurring float-depth images and is normally far below the limits of visibility, especially in final renders or HiQ mode, but may appear in subsequent processing. If you experience this, selecting the Multi-box filter may be a good choice.

Controls Tab

The Controls tab contains the primary controls necessary for customizing the blur operation, including five filter algorithms.

Filter

The Filter menu is where you select the type of filter used to create the blur.

- **Box Blur:** This option is faster than the Gaussian blur but produces a lower-quality result.
- **Bartlett:** This option is a more subtle, anti-aliased blur filter.
- **Multi-box:** Multi-box uses a Box filter layered in multiple passes to approximate a Gaussian shape. With a moderate number of passes (e.g., four), a high-quality blur can be obtained, often faster than the Gaussian filter and without any ringing.
- **Gaussian:** Gaussian applies a smooth, symmetrical blur filter, using a sophisticated constant-time Gaussian approximation algorithm.
- **Fast Gaussian:** Gaussian applies a smooth, symmetrical blur filter, using a sophisticated constant-time Gaussian approximation algorithm. This mode is the default filter method.

Color Channels (RGBA)

The filter defaults to operating on R, G, B, and A channels. Selective channel filtering is possible by clicking each channel button to make them active or inactive.

NOTE: This is not the same as the RGBA checkboxes found under the common controls. The node takes these selections into account before it processes the image, so deselecting a channel causes the node to skip that channel when processing, speeding up the rendering of the effect. In contrast, the channel controls under the Common Controls tab are applied after the node has processed.

Lock X/Y

Locks the X and Y Blur sliders together for symmetrical blurring. This is enabled by default.

Blur Size

Sets the amount of blur applied to the image. When the Lock X and Y control is deselected, independent control over each axis is provided.

Clipping Mode

This option determines how edges are handled when performing domain-of-definition rendering. This is profoundly important for nodes like Blur, which may require samples from portions of the image outside the current domain.

- **Frame:** The default option is Frame, which automatically sets the node's domain of definition to use the full frame of the image, effectively ignoring the current domain of definition. If the upstream DoD is smaller than the frame, the remaining area in the frame is treated as black/transparent.
- **Domain:** Setting this option to Domain respects the upstream domain of definition when applying the node's effect. This can have adverse clipping effects in situations where the node employs a large filter.
- **None:** Setting this option to None does not perform any source image clipping at all. This means that any data required to process the node's effect that would normally be outside the upstream DoD is treated as black/transparent.

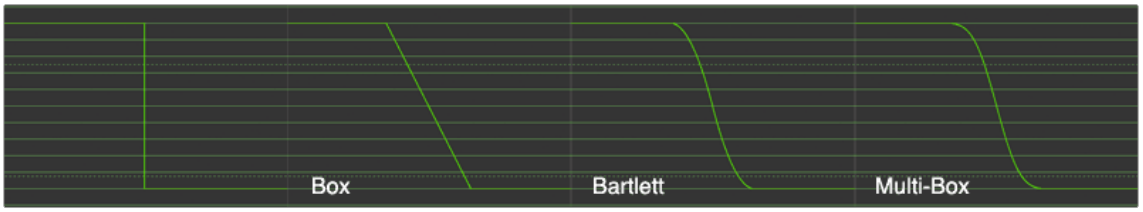
Blend

The Blend slider determines the percentage of the affected image that is mixed with original image. It blends in more of the original image as the value gets closer to 0.

This control is a cloned instance of the Blend slider in the Common Controls tab. Changes made to this control are simultaneously made to the one in the common controls.

Examples

Following is a comparison of Blur filters visualized as "cross-sections" of a filtered edge. As you can see, Box creates a linear ramp, while Bartlett creates a somewhat smoother ramp. Multi-box and Gaussian are indistinguishable unless you zoom in really close on the slopes. They both lead to even smoother ramps, but as mentioned above, Gaussian overshoots slightly and may lead to negative values if used on floating-point images.



Blur filters visualized as "cross sections" of a filtered edge

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Blur nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Defocus [Dfo]



The Defocus node

Defocus Node Introduction

The Defocus node simulates the effects of an out-of-focus camera lens, including blooming and image flaring. It provides a fast Gaussian mode, as well as a more realistic but slower Lens mode.

Inputs

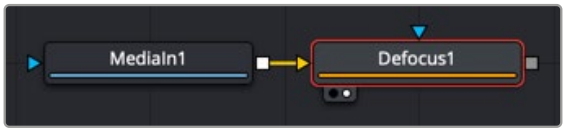
The two inputs on the Defocus node are for connecting a 2D image and an effect mask that can be used to limit the simulated defocused area.

Input: The orange input is used for the primary 2D image for defocusing.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the defocus to only those pixels within the mask. An effect mask is applied to the tool after it is processed.

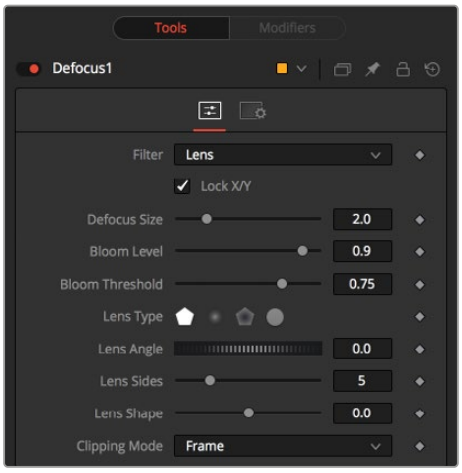
Basic Node Setup

The Defocus node receives a 2D image like the MediaIn1 shown below. The output continues the node tree by connecting to another 2D image-processing node or a Merge node.



A Defocus node applied to a MediaIn1 node

Inspector



Defocus controls

Controls Tab

The Controls tab contains all the primary controls necessary for customizing the defocus operation.

Filter

Use this menu to select the exact method applied to create the defocus. Gaussian applies a simplistic effect, while Lens mode creates a more realistic defocus. Lens mode takes significantly longer than Gaussian.

Lock X/Y

When Lock X/Y is selected, this performs the same amount of defocusing to both the X- and Y-axis of the image. Deselect to obtain individual control.

Defocus Size

The Defocus Size control sets the size of the defocus effect. Higher values blur the image by greater amounts and produce larger blooms.

Bloom Level

The Bloom Level control determines the intensity and size of the blooming applied to pixels that are above the bloom threshold.

Bloom Threshold

Pixels with values above the set Bloom Threshold are defocused and have a glow applied (blooming). Pixels below that value are only defocused.

The following four lens options are available only when the Filter is set to Lens.

- **Lens Type:** The basic shape used to create the “bad bokeh” effect. This can be refined further with the Angle, Sides, and Shape sliders.
- **Lens Angle:** Defines the rotation of the shape. Best visible with NGon lens types. Because of the round nature of a circle, this slider has no visible effect when the Lens Type is set to Circle.
- **Lens Sides:** Defines how many sides the NGon shapes have. Best visible with NGon lens types. Because of the round nature of a circle, this slider has no visible effect when the Lens Type is set to Circle.
- **Lens Shape:** Defines how pointed the NGons are. Higher values create a more pointed, starry look. Lower values create smoother NGons. Best visible with NGon lens types and Lens Sides between 5 and 10. Because of the round nature of a circle, this slider has no visible effect when the Lens Type is set to Circle.

Clipping Mode

This option determines how edges are handled when performing domain-of-definition rendering. This is profoundly important for nodes like Blur, which may require samples from portions of the image outside the current domain.

- **Frame:** The default option is Frame, which automatically sets the node’s domain of definition to use the full frame of the image, effectively ignoring the current domain of definition. If the upstream DoD is smaller than the frame, the remaining area in the frame is treated as black/transparent.
- **Domain:** Setting this option to Domain respects the upstream domain of definition when applying the node’s effect. This can have adverse clipping effects in situations where the node employs a large filter.
- **None:** Setting this option to None does not perform any source image clipping at all. This means that any data required to process the node’s effect that would normally be outside the upstream DoD is treated as black/transparent.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Blur nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Directional Blur [DrBl]



The Directional Blur node

Directional Blur Node Introduction

This node is used to create Directional and Radial blurs. It is useful for creating simulated motion blur and light ray-type effects. Directional Blur affects all channels (RGBA).

Inputs

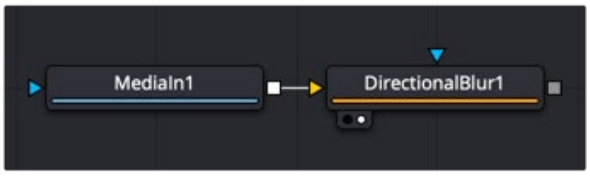
The two inputs on the Directional Blur node are used to connect a 2D image and an effect mask which can be used to limit the blurred area.

Input: The orange input is used for the primary 2D image that has the directional blur applied.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the directional blur to only those pixels within the mask. An effect mask is applied to the tool after it is processed.

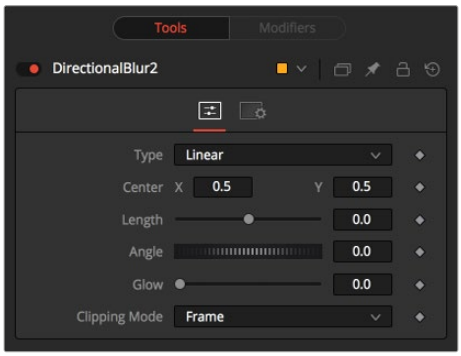
Basic Node Setup

The Directional Blur node receives a 2D image like the MediaIn1 shown below. The output continues the node tree by connecting to another 2D image-processing node or a Merge node.



A Directional Blur node applied to a MediaIn1 node

Inspector



Directional Blur controls

Controls Tab

The Controls tab contains all the primary controls necessary for customizing the directional blur operation.

Type

This menu is used to select the type of directional blur to be applied to the image.

- **Linear:** Linear distorts the image in a straight line, resembling the scenery that appears in the window of a speeding train.
- **Radial:** Radial creates a distortion that originates at some arbitrary center, radiating outward the way that a view would appear if one were at the head of the train looking forward.
- **Centered:** The Centered button produces a similar result to linear, but the blur effect is equally distributed on both sides of the original.
- **Zoom:** Zoom creates a distortion in the scale of the image smear to simulate the zoom streaking of a camera filming with a slow shutter speed.

Center X and Y

This coordinate control and related viewer crosshair affects the Radial and Zoom Motion Blur types only. It is used to position where the blurring effect starts.

Length

Length adjusts the strength and heading of the effect. Values lower than zero cause blurs to head opposite the angle control. Values greater than the slider maximum may be typed into the slider's edit box.

Angle

In both Linear and Center modes, this control modifies the direction of the directional blur. In the Radial and Zoom modes, the effect is similar to the camera spinning while looking at the same spot. If the setting of the length slider is other than zero, the effect creates a whirlpool effect.

Glow

This adds a Glow to the directional blur, which can be used to duplicate the effect of increased camera exposure to light caused by longer shutter speeds.

Clipping Mode

This option determines how edges are handled when performing domain-of-definition rendering. This is profoundly important for nodes like Blur, which may require samples from portions of the image outside the current domain.

- **Frame:** The default option is Frame, which automatically sets the node's domain of definition to use the full frame of the image, effectively ignoring the current domain of definition. If the upstream DoD is smaller than the frame, the remaining area in the frame is treated as black/transparent.
- **Domain:** Setting this option to Domain respects the upstream domain of definition when applying the node's effect. This can have adverse clipping effects in situations where the node employs a large filter.
- **None:** Setting this option to None does not perform any source image clipping at all. This means that any data required to process the node's effect that would normally be outside the upstream DoD is treated as black/transparent.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Blur nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Glow [Glo]



The Glow nozzde

Glow Node Introduction

A Glow is created by blurring an image, and then brightening the blurred result and mixing it back with the original. The Glow node provides a variety of variations on this theme. For example, a Bartlett glow is a high-quality glow with a smoother drop-off; however, it is more processor-intensive at larger sizes.

Inputs

The Glow node has three inputs: an orange one for the primary 2D image input, a blue one for an effect mask, and a third white input for a Glow mask.

Input: The orange input is used for the primary 2D image that has the glow applied.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input restricts the source of the glow to only those pixels within the mask. An effect mask is applied to the tool after it is processed.

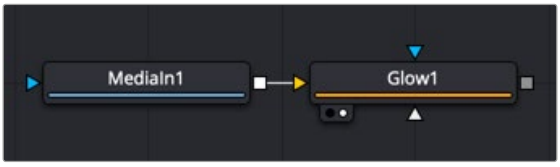
Glow Mask: The Glow node supports pre-masking using the white glow mask input. A Glow pre-mask filters the image before applying the glow. The glow is then merged back over the original image. This is different from a regular effect mask that clips the rendered result.

The Glow mask allows the glow to extend beyond the borders of the mask, while restricting the source of the glow to only those pixels within the mask.

> Glow masks are identical to Effect masks in every other respect.

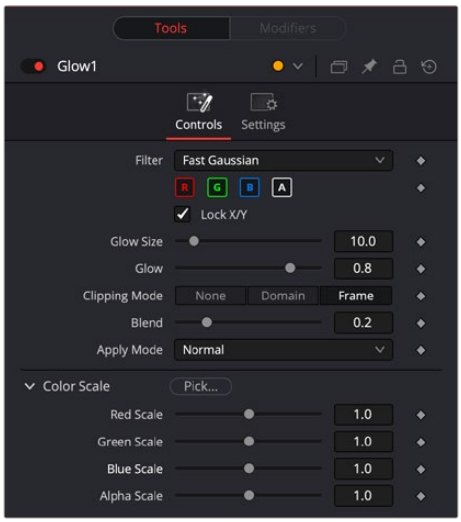
Basic Node Setup

The Glow node receives a 2D image like the MediaIn1 shown below. The output continues the node tree by connecting to another 2D image-processing node or a Merge node.



A Glow node applied to a MediaIn1 node in DaVinci Resolve

Inspector



Glow controls

Controls Tab

The Controls tab contains all the primary controls necessary for customizing the glow operation. A Color Scale section at the bottom of the Inspector can be used for tinting the glow.

Filter

Use this menu to select the method of Blur used in the filter. The selections are described below.

- **Box:** A simple but very fast Box filter.
- **Bartlett:** Bartlett adds a softer, subtler glow with a smoother drop-off but may take longer to render than Box.
- **Multi-box:** Multi-box uses a Box filter layered in multiple passes to approximate a Gaussian shape. With a moderate number of passes (e.g., four), a high-quality blur can be obtained, often faster than the Gaussian filter, and without any ringing.
- **Gaussian:** Gaussian adds a soft glow, blurred by the Gaussian algorithm.
- **Fast Gaussian:** Fast Gaussian adds a soft glow, blurred by the Gaussian algorithm. This is the default method.
- **Blend:** Blend adds a nonlinear glow that is evenly visible in the whites and blacks.
- **Hilight:** Hilight adds a glow without creating a halo in the surrounding pixels.
- **Solarize:** Solarize adds a glow and solarizes the image.

Color Channels (RGBA)

This filter defaults to operating on R, G, B, and A channels. Selective channel filtering is possible by clicking each channel to make them active or inactive.

> **NOTE:** This is not the same as the RGBA checkboxes found under the common controls. The node takes these selections into account before it processes the image, so deselecting a channel causes the node to skip that channel when processing, speeding up the rendering of the effect. In contrast, the channel controls under the Common Controls tab are applied after the node has processed.

Lock X/Y

When Lock X/Y is checked, both the horizontal and vertical glow amounts are locked. Otherwise, separate amounts of glow may be applied to each axis.

Glow Size

Glow Size determines the size of the glow effect. Larger values expand the size of the glowing highlights of the image.

Num Passes

Only available in Multi-box mode. Larger values lead to a smoother distribution of the effect, but also increase render times. It's good to find the line between desired quality and acceptable render times.

Glow

The Glow slider determines the intensity of the glow effect. Larger values tend to completely blow the image out to white.

Clipping Mode

This option determines how edges are handled when performing domain-of-definition rendering. This is profoundly important for nodes like Blur, which may require samples from portions of the image outside the current domain.

- **Frame:** The default option is Frame, which automatically sets the node's domain of definition to use the full frame of the image, effectively ignoring the current domain of definition. If the upstream DoD is smaller than the frame, the remaining area in the frame is treated as black/transparent.
- **Domain:** Setting this option to Domain respects the upstream domain of definition when applying the node's effect. This can have adverse clipping effects in situations where the node employs a large filter.
- **None:** Setting this option to None does not perform any source image clipping at all. This means that any data required to process the node's effect that would normally be outside the upstream DoD is treated as black/transparent.

Blend

The Blend slider determines the percentage of the affected image that is mixed with original image. It blends in more of the original image as the value gets closer to 0.

This control is a cloned instance of the Blend slider in the Common Controls tab. Changes made to this control are simultaneously made to the one in the common controls.

Apply Mode

Three Apply Modes are available when it comes to applying the glow to the image.

- **Normal:** Default. This mode simply adds the glow directly over top of the original image.
- **Merge Under:** Merge Under places the glow beneath the image, based on the Alpha channel. Threshold mode permits clipping of the threshold values.
- **Threshold:** This control clips the effect of the glow. A new range slider appears. Pixels in the glowed areas with values below the low value are pushed to black. Pixels with values greater than high are pushed to white.
- **High-Low Range Control:** Available only in Threshold mode. Pixels in the glowed areas with values below the low value are pushed to black. Pixels with values greater than high are pushed to white.

Color Scale (RGBA)

These Scale sliders can be used to adjust the amount of glow applied to each color channel individually, by tinting the glow.

By click and holding on the Pick button, then dragging the pointer over the viewer, you can select a specific color from the image.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Blur nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Sharpen [Shrp]



The Sharpen node

Sharpen Node Introduction

The Sharpen node uses a convolution filter to enhance detail in an image overall or to an individual channel.

Inputs

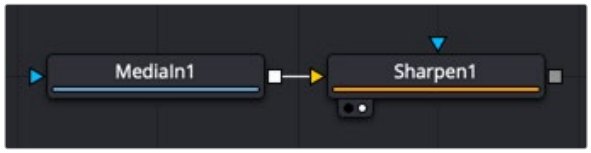
The two inputs on the Sharpen node are used to connect a 2D image and an effect mask that can limit the area affected by the sharpen.

Input: The orange input is used for the primary 2D image for sharpening.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the sharpen to only those pixels within the mask. An effect mask is applied to the tool after it is processed.

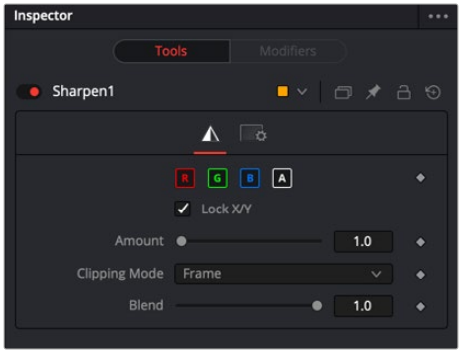
Basic Node Setup

The Sharpen node receives a 2D image like the MediaIn1 shown below. The output continues the node tree by connecting to another 2D image-processing node or a Merge node.



A Sharpen node applied to a MediaIn1 node in DaVinci Resolve

Inspector



Sharpen controls

Controls Tab

The Controls tab contains all the primary controls necessary for customizing the sharpen operation.

Color Channels (RGBA)

This filter defaults to operating on R, G, B, and A channels. Selective channel filtering is possible by clicking the channel buttons to make them active or inactive.

> **NOTE:** This is not the same as the RGBA checkboxes found under the common controls. The node takes these selections into account before it processes the image, so deselecting a channel causes the node to skip that channel when processing, speeding up the rendering of the effect. In contrast, the channel controls under the Common Controls tab are applied after the node has processed.

Lock X/Y

This locks the X and Y Sharpen sliders together for symmetrical sharpening. This is checked by default.

Amount

This slider sets the amount of sharpening applied to the image. When the Lock X/Y control is deselected, independent control over each axis is provided.

Clipping Mode

This option determines how edges are handled when performing domain-of-definition rendering. This is profoundly important for nodes like Blur, which may require samples from portions of the image outside the current domain.

- **Frame:** The default option is Frame, which automatically sets the node's domain of definition to use the full frame of the image, effectively ignoring the current domain of definition. If the upstream DoD is smaller than the frame, the remaining area in the frame is treated as black/transparent.
- **Domain:** Setting this option to Domain respects the upstream domain of definition when applying the node's effect. This can have adverse clipping effects in situations where the node employs a large filter.
- **None:** Setting this option to None does not perform any source image clipping at all. This means that any data required to process the node's effect that would normally be outside the upstream DoD is treated as black/transparent.

Blend

The Blend slider determines the percentage of the affected image that is mixed with original image. It blends in more of the original image as the value gets closer to 0.

This control is a cloned instance of the Blend slider in the Common Controls tab. Changes made to this control are simultaneously made to the one in the common controls.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Blur nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Soft Glow [SGlo]



The Soft Glow node

Soft Glow Node Introduction

The Soft Glow node is similar to the Glow node but performs additional processing of the image to create a much softer, more natural glow.

This node is perfect for atmospheric haze around planets, skin tones, and simulating dream like environments.

Inputs

Like the Glow node, Soft Glow also has three inputs: an orange one for the primary image input, a blue one for an effect mask, and a third white input for a Glow mask.

Input: The orange input is used for the primary 2D image for the soft glow.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the soft glow to only those pixels within the mask. An effect mask is applied to the tool after it is processed.

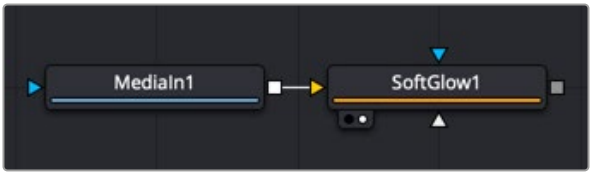
Glow Mask: The Soft Glow node supports pre-masking using the white glow mask input. A Glow pre-mask filters the image before applying the soft glow. The soft glow is then merged back over the original image. This is different from a regular effect mask that clips the rendered result.

The Glow mask allows the soft glow to extend beyond the borders of the mask, while restricting the source of the soft glow to only those pixels within the mask.

Glow masks are identical to effect masks in every other respect.

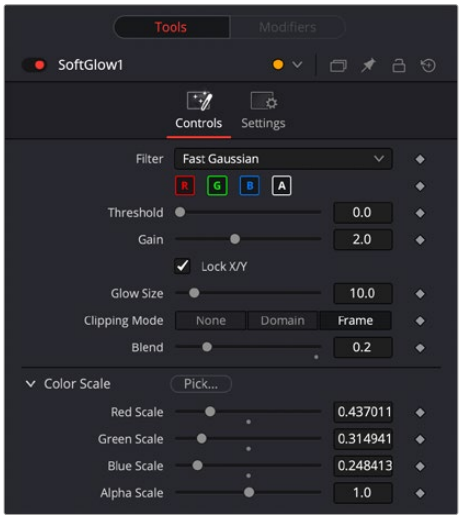
Basic Node Setup

The Soft Glow node receives a 2D image like the MediaIn1 shown below. The output continues the node tree by connecting to another 2D image-processing node or a Merge node.



A Soft Glow node applied to a MediaIn1 node

Inspector



Soft Glow controls

Controls Tab

The Controls tab contains all the primary controls necessary for customizing the soft glow operation. A color scale section at the bottom of the Inspector can be used for tinting the soft glow.

Filter

Use this menu to select the method of Blur used in the filter. The selections are described below.

- **Box:** A simple but very fast Box filter.
- **Bartlett:** Bartlett adds a softer, subtler glow with a smoother drop-off but may take longer to render than Box.
- **Multi-box:** Multi-box uses a Box filter layered in multiple passes to approximate a Gaussian shape. With a moderate number of passes (e.g., four), a high-quality blur can be obtained, often faster than the Gaussian filter and without any ringing.
- **Gaussian:** Gaussian adds a soft glow, blurred by the Gaussian algorithm. This is the default method.

Color Channels (RGBA)

The filter defaults to operating on R, G, B, and A channels. Selective channel filtering is possible by clicking the channel buttons to make them active or inactive.

> **NOTE:** This is not the same as the RGBA checkboxes found under the common controls. The node takes these selections into account before it processes the image, so deselecting a channel causes the node to skip that channel when processing, speeding up the rendering of the effect. In contrast, the channel controls under the Common Controls tab are applied after the node has processed.

Threshold

This control is used to limit the effect of the soft glow. The higher the threshold, the brighter the pixel must be before it is affected by the glow.

Gain

The Gain control defines the brightness of the glow.

Lock X/Y

When Lock X/Y is checked, both the horizontal and vertical glow amounts are locked. Otherwise, separate amounts of glow may be applied to each axis of the image.

Glow Size

This amount determines the size of the glow effect. Larger values expand the size of the glowing highlights of the image.

Num Passes

Available only in Multi-box mode. Larger values lead to a smoother distribution of the effect, but also increase render times. It's good to find the line between desired quality and acceptable render times.

Clipping Mode

This option determines how edges are handled when performing domain-of-definition rendering. This is profoundly important for nodes like Blur, which may require samples from portions of the image outside the current domain.

- **Frame:** The default option is Frame, which automatically sets the node's domain of definition to use the full frame of the image, effectively ignoring the current domain of definition. If the upstream DoD is smaller than the frame, the remaining area in the frame is treated as black/transparent.
- **Domain:** Setting this option to Domain respects the upstream domain of definition when applying the node's effect. This can have adverse clipping effects in situations where the node employs a large filter.
- **None:** Setting this option to None does not perform any source image clipping at all. This means that any data required to process the node's effect that would normally be outside the upstream DoD is treated as black/transparent.

Blend

The Blend slider determines the percentage of the affected image that is mixed with original image. It blends in more of the original image as the value gets closer to 0.

This control is a cloned instance of the Blend slider in the Common Controls tab. Changes made to this control are simultaneously made to the one in the common controls.

Color Scale (RGBA)

These Scale sliders are used to adjust the amount of glow applied to each color channel individually, by tinting the glow.

By click and holding on the Pick button, then dragging the pointer over the viewer, you can select a specific color from the image.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Blur nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Unsharp Mask [UsM]



The Unsharp Mask node

Unsharp Mask Introduction

Unsharp masking is a technique used to sharpen only the edges within an image. This node is most often used to correct for blurring and loss of detail in low-contrast images; for example, to extract useful detail from long exposure shots of faraway galaxies.

This filter extracts a range of frequencies from the image and blurs them to reduce detail. The blurred result is then compared to the original images. Pixels with a significant difference between the original and the blurred image are likely to be an edge detail. The pixel is then brightened to enhance it.

Inputs

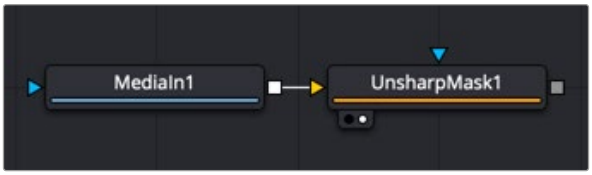
The two inputs on the Unsharp Mask node are used to connect a 2D image and an effect mask for limiting the effect.

Input: The orange input is used for the primary 2D image for the Unsharp Mask.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the Unsharp Mask to only those pixels within the mask. An effect mask is applied to the tool after it is processed.

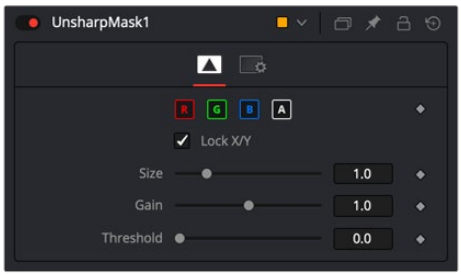
Basic Node Setup

The Unsharp Mask node receives a 2D image like the MediaIn1 shown below. The output continues the node tree by connecting to another 2D image-processing node or a Merge node.



An Unsharp mask node applied to a MediaIn1 node

Inspector



Unsharp Mask controls

Controls Tab

The Controls tab contains all the primary controls necessary for customizing the Unsharp Mask operation.

Color Channels (RGBA)

The filter defaults to operating on R, G, B, and A channels. Selective channel filtering is possible by clicking the channel buttons to make them active or inactive.

NOTE: This is not the same as the RGBA checkboxes found under the common controls. The node takes these selections into account before it processes the image, so deselecting a channel causes the node to skip that channel when processing, speeding up the rendering of the effect. In contrast, the channel controls under the Common Controls tab are applied after the node has processed.

Lock X/Y

When Lock X/Y is checked, both the horizontal and vertical sharpen amounts are locked. Otherwise, separate amounts of glow may be applied to each axis of the image.

Size

This control adjusts the size of blur filter applied to the extracted image. The higher this value, the more likely it is that pixels are identified as detail.

Gain

The Gain control adjusts how much gain is applied to pixels identified as detail by the mask. Higher values create a sharper image.

Threshold

This control determines the frequencies from the source image to be extracted. Raising the value eliminates lower-contrast areas from having the effect applied.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Blur nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Vari Blur [VBl]



The Vari Blur node

Vari Blur Node Introduction

The Vari Blur node gives a true per-pixel variable blur, using a second image to control the amount of blur for each pixel. It is somewhat similar in effect to the Depth Blur node but uses a different approach for frequently cleaner results.

Inputs

There are two inputs on the Vari Blur node for the primary image: the blur map image, and an effect mask.

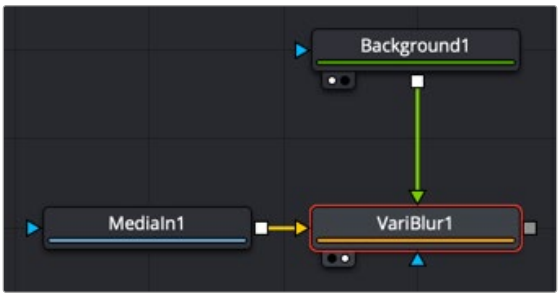
Input: The gold image input is a required connection for the primary image you wish to blur.

Blur Image: The green input is also required, but it can accept a spline shape, text object, still image, or movie file as the blur map image. Once connected, you can choose red, green, blue, Alpha, or luminance channel to create the shape of the blur.

Effect Mask: The optional blue effect mask input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the Vari Blur to only those pixels within the mask. An effect mask is applied to the tool after it is processed.

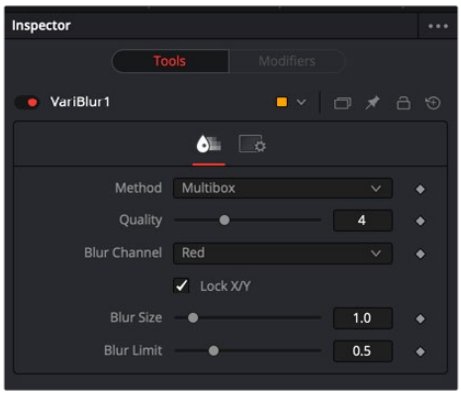
Basic Node Setup

The Vari Blur node receives a 2D image like the MediaIn1 shown below. A gradient Background tool connects to the Blur image input to control the areas affected by the blur. The output continues the node tree by connecting to another 2D image-processing node or a Merge node.



A Vari Blur node applied to a MediaIn1 node and a gradient background directing the blurred regions

Inspector



Vari Blur controls

Controls Tab

The Controls tab contains all the primary controls necessary for customizing the Vari Blur operation.

Method

Use this menu to select the method of Blur used in the filter. The selections are described below.

- **Soften:** This method varies from a simple Box shape to a Bartlett triangle to a decent-looking Smooth blur as Quality is increased. It is a little better at preserving detail in less-blurred areas than Multi-box.
- **Multi-box:** Similar to Soften, this gives a better Gaussian approximation at higher Quality settings.
- **Defocus:** Produces a flat, circular shape to blurred pixels that can approximate the look of a defocus.

Quality

Increasing Quality gives smoother blurs, at the expense of speed. Quality set to 1 uses a very fast but simple Box blur for all Method settings. A Quality of 2 is usually sufficient for low Blur Size values. A Quality of 4 is generally good enough for most jobs unless Blur Size is particularly high.

Blur Channel

This selects which channel of the Blur Image controls the amount of blurring applied to each pixel.

Lock X/Y

When selected, only a Blur Size control is shown, and changes to the amount of blur are applied to both axes equally. If the checkbox is cleared, individual controls appear for both X and Y Blur Size.

Blur Size

Increasing this control increases the overall amount of blur applied to each pixel. Those pixels where the Blur image is black or nonexistent are blurred, despite the Blur Size.

Blur Limit

This slider limits the useable range from the Blur image. Some Z-depth images can have values that go to infinity, which skew blur size. The Blur Limit is a way to keep values within range.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Blur nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Vector Motion Blur [VMB]



The Vector Motion Blur node

Vector Motion Blur Introduction

This node is used to create directional blurs based on a Motion Vector map or AOV (Arbitrary Output Variable) channels exported from 3D-rendering software like Arnold, Renderman, or VRay. You can also generate motion vectors using the Optical Flow node in Fusion.

The vector map is typically two floating-point images: one channel specifies how far the pixel is moving in X, and the other specifies how far the pixel is moving in Y. These channels may be embedded in OpenEXR or RLA/RPF images, or may be provided as separate images using the node’s Vectors input.

The vector channels should use a float16 or float32 color depth, to provide + and – values.

A value of 1 in the X channel would indicate that pixel has moved one pixel to the right, while a value of –10 indicates ten pixels of movement to the left.

Inputs

The Vector Motion Blur node has three inputs for a 2D image, a motion vector pass, and an effect mask.

Input: The required orange input is for a 2D image that receives the motion blur.

Vectors: The green input is also required. This is where you connect a motion vector AOV rendered from a 3D application or an EXR file generated from the Optical Flow node in Fusion.

Vector Mask: The white Vector Mask input is an optional input that masks the image before processing.

Effect Mask: The common blue input is used for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input restricts the source of the motion blur to only those pixels within the mask. An effect mask is applied to the tool after it is processed.

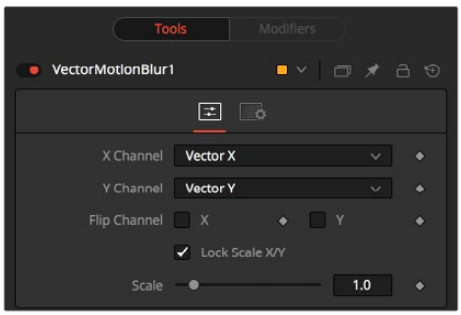
Basic Node Setup

The Vector Motion Blur node receives a 2D image like the IMAGE shown below. A MediaIn or Loader node containing motion vectors is connected to the Vector's input. The output continues the node tree by connecting to another 2D image-processing node or a Merge node.



A Vector Motion Blur node applied to a MediaIn or Loader node with motion vectors connected to the Vectors input

Inspector



Vector Motion Blur node

Controls Tab

The Controls tab contains all the primary controls necessary for customizing the Vector Motion Blur operation.

X Channel

Use this menu to select which channel of the image provides the vectors for the movement of the pixels along the X-axis.

Y Channel

Use this menu to select which channel of the image provides the vectors for the movement of the pixels along the Y-axis.

Flip Channel

These checkboxes can be used to flip, or invert, the X and Y vectors. For instance, a value of 5 for a pixel in the X-vector channel would become –5 when the X checkbox is enabled.

Lock Scale X/Y

Selecting this checkbox provides access to separate sliders for X and Y Scale. By default, only a single Scale slider is provided.

Scale

The X and Y vector channel values for a pixel are multiplied by the value of this slider. For example, given a scale of 2 and a vector value of 10, the result would be 20. This slider splits to show Scale X and Scale Y if the Lock Scale X/Y checkbox is not enabled.

Common Controls

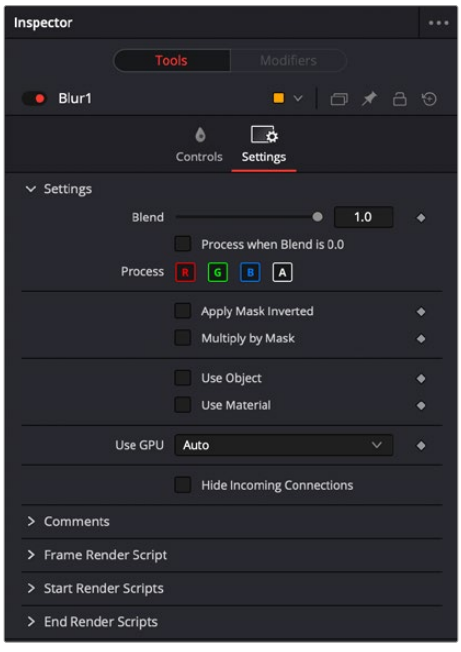
Settings Tab

The Settings tab in the Inspector is also duplicated in other Blur nodes. These common controls are described in the following “The Common Controls” section.

The Common Controls

Nodes that handle blur operations share several identical controls in the Inspector. This section describes controls that are common among Blur nodes.

Inspector



Common Blur settings

Settings Tab

The Settings tab in the Inspector can be found on every tool in the Blur category. The Settings controls are even found on third-party Blur-type plugin tools. The controls are consistent and work the same way for each tool.

Blend

The Blend control is used to blend between the tool's original image input and the tool's final modified output image. When the blend value is 0.0, the outgoing image is identical to the incoming image. Normally, this causes the tool to skip processing entirely, copying the input straight to the output.

Process When Blend Is 0.0

The tool is processed even when the input value is zero. This can be useful if processing of this node is scripted to trigger another task, but the value of the node is set to 0.0.

Red/Green/Blue/Alpha Channel Selector

These four buttons are used to limit the effect of the tool to specified color channels. This filter is often applied after the tool has been processed.

For example, if the Red button on a Blur tool is deselected, the blur is first applied to the image, and then the red channel from the original input is copied back over the red channel of the result.

There are some exceptions, such as tools where deselecting these channels causes the tool to skip processing that channel entirely. Tools that do this generally possess a set of identical RGBA buttons on the Controls tab in the tool. In this case, the buttons in the Settings and the Controls tabs are identical.

Apply Mask Inverted

Enabling the Apply Mask Inverted option inverts the complete mask channel for the tool. The mask channel is the combined result of all masks connected to or generated in a node.

Multiply by Mask

Selecting this option causes the RGB values of the masked image to be multiplied by the mask channel's values. This causes all pixels of the image not in the mask (i.e. set to 0) to become black/transparent.

Use Object/Use Material (Checkboxes)

Some 3D software can render to file formats that support additional channels. Notably, the EXR file format supports object and material ID channels, which can be used as a mask for the effect. These checkboxes determine whether the channels are used, if present. The specific material ID or object ID affected is chosen using the next set of controls.

Correct Edges

This checkbox appears only when the Use Object or Use Material checkboxes are selected. It toggles the method used to deal with overlapping edges of objects in a multi-object image. When enabled, the Coverage and Background Color channels are used to separate and improve the effect around the edge of the object. If this option disabled (or no Coverage or Background Color channels are available), aliasing may occur on the edge of the mask.

> For more information on the Coverage and Background Color channels, *see Chapter 18, "Understanding Image Channels,"* in the Fusion Reference Manual.

Object ID/Material ID (Sliders)

Use these sliders to select which ID is used to create a mask from the object or material channels of an image. Use the Sample button in the same way as the Color Picker to grab IDs from the image displayed in the viewer. The image or sequence must have been rendered from a 3D software package with those channels included.

Use GPU

The GPU menu has three settings. Disable turns off GPU hardware accelerated rendering. Enabled uses the GPU hardware for rendering the node. Auto uses a capable GPU if one is available and falls back to software rendering when a capable GPU is not available.

Motion Blur

- **Motion Blur:** This toggles the rendering of Motion Blur on the tool. When this control is toggled on, the tool's predicted motion is used to produce the motion blur caused by the virtual camera's shutter. When the control is toggled off, no motion blur is created.
- **Quality:** Quality determines the number of samples used to create the blur. A Quality setting of 2 causes Fusion to create two samples to either side of an object's actual motion. Larger values produce smoother results but increase the render time.
- **Shutter Angle:** Shutter Angle controls the angle of the virtual shutter used to produce the motion blur effect. Larger angles create more blur but increase the render times. A value of 360 is the equivalent of having the shutter open for one whole frame exposure. Higher values are possible and can be used to create interesting effects.
- **Center Bias:** Center Bias modifies the position of the center of the motion blur. This allows the creation of motion trail effects.
- **Sample Spread:** Adjusting this control modifies the weighting given to each sample. This affects the brightness of the samples.

Comments

The Comments field is used to add notes to a tool. Click in the field and type the text. When a note is added to a tool, a small red square appears in the lower-left corner of the node when the full tile is displayed, or a small text bubble icon appears on the right when nodes are collapsed. To see the note in the Node Editor, hold the mouse pointer over the node to display the tooltip.

Scripts

Three Scripting fields are available on every tool in Fusion from the Settings tab. They each contain edit boxes used to add scripts that process when the tool is rendering. For more details on scripting nodes, please consult the Fusion scripting documentation.

Chapter 34

Color Nodes

This chapter details the Color nodes available in Fusion. The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

> For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

ColorFX

ACESTransform [ATr] 918

Chromatic Adaptation [CRa] 920

Color Space Transform [CSt] 922

Gamut Limiter [GMl] 925

Gamut Mapping [GMp] 927

Revival

Chromatic Aberration Removal [CAr] 930

Color

Auto Gain [AG] 933

Brightness Contrast [BC] 935

Channel Booleans [Bol] 938

Color Corrector [CC] 941

Color Curves [CCv] 952

Color Gain [Clr] 955

Color Matrix [CMx] 960

Color Space [CS] 964

Copy Aux [CpA] 966

Gamut [Gmt] 969

Hue Curves [HCv] 972

OCIO CDL Transform [OCD] 975

OCIO Color Space [OCC] 978

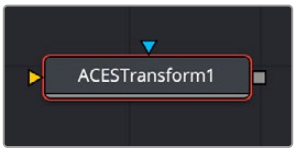
OCIO Display [OCD] 980

OCIO File Transform [OCF] 981
Set Canvas Color [SCv] 983
White Balance [WB] 985
The Common Controls 989

ColorFX

A sub category of the Color tools in Fusion that gives you access to various color processing tools.

ACESTransform [ATr]



The ACES Transform node

ACES Transform Node Introduction

A simple node that lets you perform color transforms using the ACES Input Device Transform and ACES Output Device Transform parameters.

Inputs

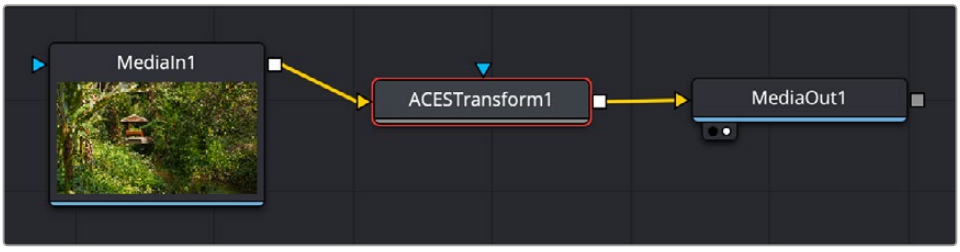
The two inputs on the ACES Transform node are the input and effect mask.

Input: The orange input connects the primary 2D image for the auto gain.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the adjustment to only those pixels within the mask. An effect mask is applied to the tool after the tool is processed.

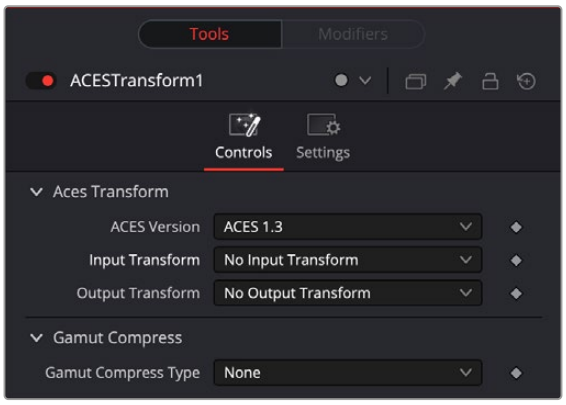
Basic Node Setup

The ACES Transform node, like many 2D image-processing nodes, receives a 2D image, such as a Loader node or the MediaIn1 shown below. The output continues the node tree by connecting to another 2D image-processing node or a Merge node.



An ACES Transform node applied to a MediaIn1 node

Inspector



ACES Transform controls

Controls Tab

The Controls tab contains the few primary controls necessary for performing the ACES transform on the input.

ACES Version

Lets you choose which version of ACES you want to use.

Input Transform

This menu lets you choose which IDT (Input Device Transform) to use to process the image. Typically you would pick the camera that your footage was shot with here.

Output Transform

This menu lets you choose an ODT (Output Device Transform) with which to transform the image data to the color space for your exported timeline.

Gamut Compress Type

This lets you choose the type of ACES Gamut compression applied to the node. Choosing either option can prevent monochromatic highly saturated light sources (LEDs, neon signs, tail lights, etc.) from clipping the gamut.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Color nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Chromatic Adaptation [CRa]



The Chromatic Adaptation node

Chromatic Adaptation Node Introduction

A simple node that lets you precisely transform an image that has been lit or processed assuming one specific color temperature to another user-selectable color temperature. This transformation alters the appearance of all colors in the image as perceived by the human vision system in the same way that a new illuminant would, whether that illuminant is a light in the environment being recorded or the color temperature of a display on which the image is shown. This node is useful for performing specific color temperature transformations as part of color management workflows or for setting up precise color temperature adjustments.

Inputs

The two inputs on the Chromatic Adaptation node are the input and effect mask.

Input: The orange input connects the primary 2D image for the auto gain.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the adjustment to only those pixels within the mask. An effect mask is applied to the tool after the tool is processed.

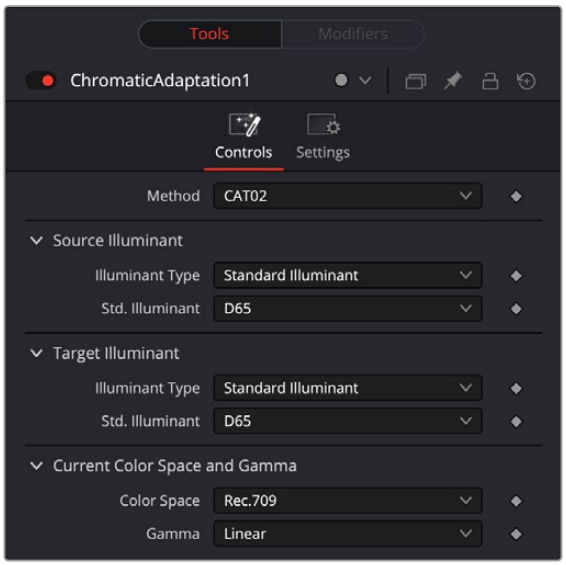
Basic Node Setup

The Chromatic Adaptation node, like many 2D image-processing nodes, receives a 2D image, such as a Loader node or the MediaIn1 shown below. The output continues the node tree by connecting to another 2D image-processing node or a Merge node.



A Chromatic Adaptation node applied to a MediaIn1 node

Inspector



ACES Transform controls

Controls Tab

The Controls tab contains the primary controls necessary for adjusting the Chromatic Adaptation parameters.

Method

The Method pop-up menu provides a variety of different transform methods to choose from, defaulting to CAT02. Each option in the Method pop-up menu uses different measurement datasets to create individual CAT matrixes to guide this transformation. As a result, each method prioritizes different levels of accuracy for different sets of colors. For example:

- CAT02 has a non-linear component that compensates for the tendency of extremely saturated blues to go purple, a typical weakness of other methods. It usually gives the best result for the widest variety of measured data sets and works best for emissive sources (displays) and dim viewing environments.
- Bradford Linear is also a commonly used method, albeit one in which extremely saturated blues will go purple during the transform. It works well for both emissive sources in dim environments and for reflective sources (theater screens) and dark environments.
- Von Kries is one of the oldest methods in common use, although it's also one in which extremely saturated blues will go purple during the transform. This, as well as all other methods, are available if you need to match work done in another image processing application.

NOTE: Be aware that all methods listed will match neutral colors perfectly; the only differences lie in how different ranges of saturated color are transformed.

Source/Target Illuminant

You control this transformation by using pop-up menus to define the Illuminant Type of the source (typically the color temperature the camera was set to) and that of the target that you want to transform the image to. For both sets of controls, you can choose a Standard Illuminant from a list, a Color Temperature via a slider, or CIE 1931 xy coordinate values. This image transform is extremely precise because the image is first transformed from the Timeline Color Space to XYZ, and then it's transformed to match the LMS (long, medium, short) color space that models the cone response of the human eye to colors lit by different illuminants.

Current Color Space and Gamma

This node also takes into account the current color space and gamma of the clip, which default to those set for the current timeline. If you wish to change these values, you can do so using the Color Space and Gamma menus.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Color nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Color Space Transform [CSt]



The Color Space Transform node

Color Space Transform Node Introduction

A simple node that lets you perform the kind of color transforms that LUTs do, but instead of using lookup tables, this plugin uses the same math used by Resolve Color Management (RCM) in order to do extremely clean color transforms without clipping.

Inputs

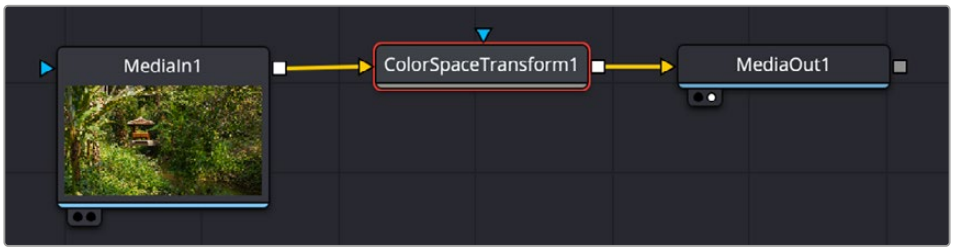
The two inputs on the Color Space Transform node are the input and effect mask.

Input: The orange input connects the primary 2D image for the auto gain.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the adjustment to only those pixels within the mask. An effect mask is applied to the tool after the tool is processed.

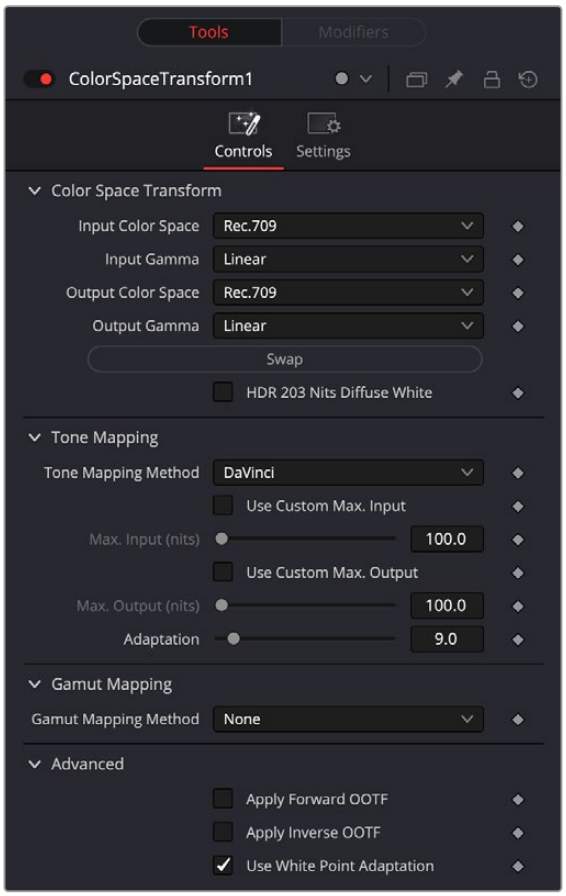
Basic Node Setup

The Color Space Transform node, like many 2D image-processing nodes, receives a 2D image like a Loader node or the MediaIn1 shown below. The output continues the node tree by connecting to another 2D image-processing node or a Merge node.



A Color Space Transform node applied to a MediaIn1 node

Inspector



Color Space Transform controls

Controls Tab

The Controls tab contains the primary controls necessary for performing the Color Space Transform on the input.

Color Space Transform

Exposes four pop-up menus that let you set an Input Colorspace, Input Gamma, Output Colorspace, and Output Gamma, in order to do controlled transforms from the Input settings to the Output settings, right within a node. The Swap button lets you quickly reverse the color and gamma space conversion.

Tone Mapping

Tone Mapping lets you enable tone mapping to accommodate workflows where you need to transform one color space into another with a dramatically larger or smaller dynamic range by automating an expansion or contraction of image contrast in such a way as to give a pleasing result with no clipping.

- **None:** This setting disables Input DRT Tone Mapping. No tone mapping is applied to the Input to Timeline Color Space conversion at all, resulting in a simple 1:1 mapping to the Timeline Color Space.
- **Clip:** Hard clips all out-of-bounds values.
- **Simple:** Uses a simple curve to perform this transformation, compressing or expanding the highlights and/or shadows of the timeline dynamic range to better fit the output dynamic range. Note that the “Simple” option maps between approximately 5500 nits and 100 nits, so if you’re mapping from an HDR source with more than 5500 nits to an SDR destination there may still be some clipping of the highlights above 5500 nits.

- **Luminance Mapping:** Same as DaVinci, but more accurate when the Input Color Space of all your media is in a single standards-based color space, such as Rec. 709 or Rec. 2020.
- **DaVinci:** This option tone maps the transform with a smooth luminance roll-off in the shadows and highlights, and controlled desaturation of image values in the very brightest and darkest parts of the image. This setting is particularly useful for wide-gamut camera media and is a good setting to use when mixing media from different cameras.
- **Saturation Preserving:** This option has a smooth luminance roll-off in the shadows and highlights but does so without desaturating dark shadows and bright highlights, so this is an effective option for colorists who like to push color harder. However, because over-saturation in the highlights of the image can look unnatural, two parameters are exposed to provide some user-adjustable automated desaturation.

  Sat. Rolloff Start: Lets you set a threshold, in nits (cd/m2), at which saturation will roll off along with highlight luminance. Beginning of the rolloff.

  Sat. Rolloff Limit: Lets you set a threshold, in nits (cd/m2), at which the image will be totally desaturated. End of the rolloff.

  Use Custom Max Input/Output: Checking these boxes and adjusting the slider below allows you to specify the minimum and maximum luminance of the input image in nits. Using these two sliders together, you can set which value from the Input Gamma is mapped to which value of the Output Gamma.
- **Adaptation:** Used to compensate for large differences in the viewer's state of visual adaptation when viewing a bright image on an HDR display versus seeing that same image on an SDR display. For most "average" images this setting works best set between 0–10. However, when you're converting very bright images (for example, a snow scene at noon), then using a higher value will yield more image detail within the highlights.

Gamut Mapping

Gamut Mapping Method accommodates workflows where you need to transform one color space into another with a dramatically larger or smaller gamut by helping to automate an expansion or contraction of image saturation in such a way as to give a pleasing and naturalistic result with no clipping.

- **None:** results in no Gamut mapping at all.
- **Saturation Mapping:** Enables saturation mapping to fit the range of saturation values from the Input Color Space and Gamma into the Output Color Space and Gamma. It enables the Saturation Knee and Saturation Max. controls.

  The Saturation Knee slider sets the image level at which saturation mapping begins. Below this level, no remapping is applied. All saturation values from this level on up are remapped according to the Saturation Max. slider. A value of 1.0 is maximum saturation in the currently selected output color space.

  The Saturation Max. slider sets the new maximum level to which you want to either raise or lower all saturation values that are above the Saturation Knee setting. A value of 1.0 is maximum saturation in the currently selected output color space.
- **Clip:** Hard clips all out-of-gamut values.

NOTE: While this node has ACES settings, it does transforms to the ACES color space colormetrically, which is not actually correct for ACES workflows. For actual ACES workflows, use the ACES Transform node, which uses transforms specified by the Academy.

Advanced

This drop-down menu exposes the advanced features of the Color Space Transform node.

- **Apply Forward OOTF:** Check this box to convert the image from scene referred to display referred color management.
- **Apply Inverse OOTF:** Check this box to convert the image from display referred to scene referred color management.
- **Use White Point Adaptation:** Applies a chromatic adaptation transform to account for different white points between color spaces.

Uncheck this box if you simply want to view the input color space's white point unaltered in the output color space. For example, wanting to use a P3-D60 mastered clip inside a P3-D65 timeline for reference purposes.

Check this box to apply the chromatic adaptation transform to convert the input white point to match the output color space's white point. For example, wanting a P3-D60 mastered clip to cut in with other clips mastered in a P3-D65 timeline.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Color nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Gamut Limiter [GMI]



The Gamut Limiter node

Gamut Limiter Node Introduction

Lets you limit the gamut to a specified standard. Useful in situations where the delivery color space is a large gamut such as Rec. 2020, but the QC specification requires limiting to a smaller gamut such as P3, in order to limit the amount of saturation allowable in the final output. This is a limiting operation, so out of bounds values are hard clipped. Because it's a limiter, it should probably be one of the last operations in any node tree to prevent useful image data from being clipped.

Inputs

The two inputs on the Gamut Limiter node are the input and effect mask.

Input: The orange input connects the primary 2D image for the auto gain.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the adjustment to only those pixels within the mask. An effect mask is applied to the tool after the tool is processed.

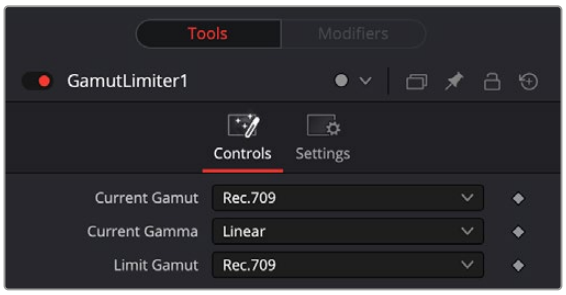
Basic Node Setup

The Gamut Limiter node, like many 2D image-processing nodes, receives a 2D image, such as a Loader node or the MediaIn1 shown below. The output continues the node tree by connecting to another 2D image-processing node or a Merge node.



A Gamut Limiter node applied to a MediaIn1 node

Inspector



Gamut Limiter controls

Controls Tab

The Controls tab contains the primary controls necessary for adjusting the Gamut Limiter parameters.

Current Gamut

Choose the timeline gamut currently being used by the image.

Current Gamma

Choose the timeline gamma currently being used by the image.

Limit Gamut

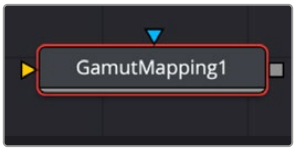
Choose the gamut you want to restrict the image to here.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Color nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Gamut Mapping [GMp]



The Gamut Mapping node

Gamut Mapping Node Introduction

The Color Space Transform node provides Gamut Mapping controls to accommodate workflows where you need to transform one color space into another that has a dramatically larger or smaller gamut. These controls are identical to those found in the Color Space Transform node's Gamut Mapping group and are similar to those found in the Color Management panel of the Project Settings.

Inputs

The two inputs on the Gamut Mapping node are the input and effect mask.

Input: The orange input connects the primary 2D image for the auto gain.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the adjustment to only those pixels within the mask. An effect mask is applied to the tool after the tool is processed.

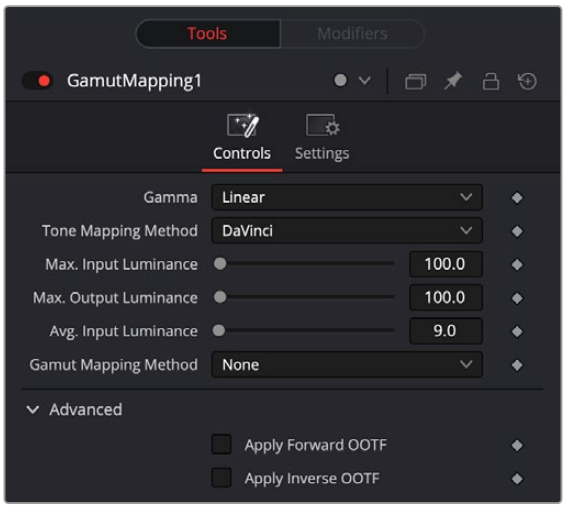
Basic Node Setup

The Gamut Mapping node, like many 2D image-processing nodes, receives a 2D image, such as a Loader node or the MediaIn1 shown below. The output continues the node tree by connecting to another 2D image-processing node or a Merge node.



A Gamut Map node applied to a MediaIn1 node

Inspector



Gamut Mapping controls

Controls Tab

The Controls tab contains the primary controls necessary for adjusting the Gamut Limiter parameters.

Gamma

A pop-up menu lets you specify what type of gamma the clip is supposed to have, so set this to whatever matches that image (this may match the Timeline Color Space, but it depends on how you're working).

Tone Mapping Method

Lets you enable tone mapping to accommodate workflows where you need to transform one color space into another with a dramatically larger or smaller dynamic range, by automating an expansion or contraction of image contrast in such a way as to give a pleasing result with no clipping.

- **None:** This setting disables Input DRT Tone Mapping. No tone mapping is applied to the Input to Timeline Color Space conversion at all, resulting in a simple 1:1 mapping to the Timeline Color Space.
- **Clip:** Hard clips all out-of-bounds values.
- **Simple:** Uses a simple curve to perform this transformation, compressing or expanding the highlights and/or shadows of the timeline dynamic range to better fit the output dynamic range. Note that the "Simple" option maps between approximately 5500 nits and 100 nits, so if you're mapping from an HDR source with more than 5500 nits to an SDR destination there may still be some clipping of the highlights above 5500 nits.
- **Luminance Mapping:** Same as DaVinci, but more accurate when the Input Color Space of all your media is in a single standards-based color space, such as Rec. 709 or Rec. 2020.
- **DaVinci:** This option tone maps the transform with a smooth luminance roll-off in the shadows and highlights, and controlled desaturation of image values in the very brightest and darkest parts of the image. This setting is particularly useful for wide-gamut camera media and is a good setting to use when mixing media from different cameras.

- **Saturation Preserving:** This option has a smooth luminance roll-off in the shadows and highlights but does so without desaturating dark shadows and bright highlights, so this is an effective option for colorists who like to push color harder. However, because over-saturation in the highlights of the image can look unnatural, two parameters are exposed to provide some user-adjustable automated desaturation.

  Sat. Rolloff Start: Lets you set a threshold, in nits (cd/m2), at which saturation will roll off along with highlight luminance. Beginning of the rolloff.

  Sat. Rolloff Limit: Lets you set a threshold, in nits (cd/m2), at which the image will be totally desaturated. End of the rolloff.

Max Input/Output Luminance

Adjusting the sliders allows you to specify the minimum and maximum luminance of the input image in nits. Using these two sliders together, you can set which value from the Input Gamma is mapped to which value of the Output Gamma.

Average Input Luminance

Used to compensate for large differences in the viewer's state of visual adaptation when viewing a bright image on an HDR display versus seeing that same image on an SDR display. For most "average" images this setting works best set between 0–10. However, when you're converting very bright images (for example, a snow scene at noon), then using a higher value will yield more image detail within the highlights.

Gamut Mapping Method

Accommodates workflows where you need to transform one color space into another with a dramatically larger or smaller gamut by helping to automate an expansion or contraction of image saturation in such a way as to give a pleasing and naturalistic result with no clipping. Choosing Saturation Mapping lets you remap the saturation values of the image. It enables the Saturation Knee and Saturation Max. controls.

- The Saturation Knee slider sets the image level at which saturation mapping begins. Below this level, no remapping is applied. All saturation values from this level on up are remapped according to the Saturation Max. slider. A value of 1.0 is maximum saturation in the currently selected output color space.
- The Saturation Max. slider sets the new maximum level to which you want to either raise or lower all saturation values that are above the Saturation Knee setting. A value of 1.0 is maximum saturation in the currently selected output color space.

Advanced

- **Apply Forward OOTF:** Check this box to convert the image from scene referred to display referred color management.
- **Apply Inverse OOTF:** Check this box to convert the image from display referred to scene referred color management.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Color nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Revival

A sub category of the Color tools in Fusion that pertains to fixing common technical errors with programs being finished, remastered, or restored.

Chromatic Aberration Removal [CAr]



The Chromatic Aberration Removal node

Chromatic Aberration Removal Node Introduction

A node that lets you manually correct the slight color fringing that results from chromatic aberration in a lens. Show Estimated Fringes checkboxes for the RGB channels display an “alignment guide” that visually isolates each of the types of fringing against gray.

Inputs

The two inputs on the Chromatic Aberration Removal node are the input and effect mask.

Input: The orange input connects the primary 2D image for the auto gain.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the adjustment to only those pixels within the mask. An effect mask is applied to the tool after the tool is processed.

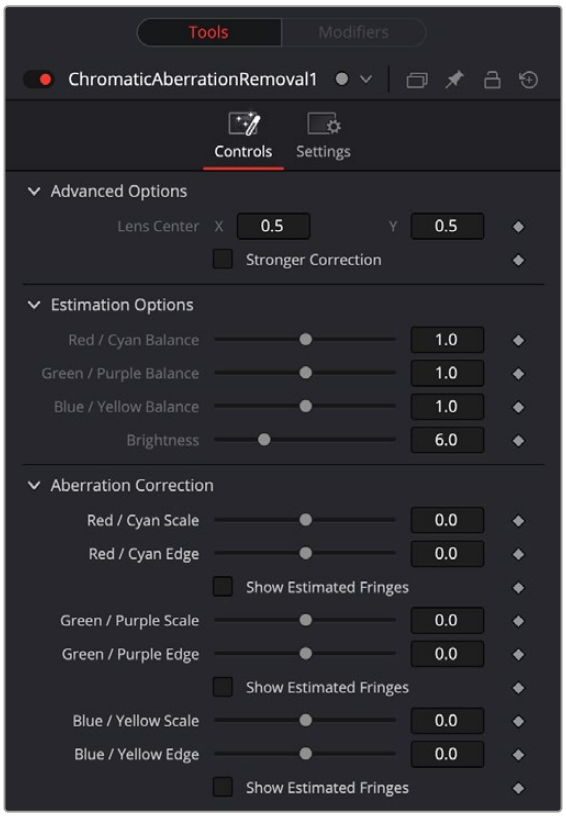
Basic Node Setup

The Chromatic Aberration Removal node, like many 2D image-processing nodes, receives a 2D image, such as a Loader node or the MediaIn1 shown below. The output continues the node tree by connecting to another 2D image-processing node or a Merge node.



A Chromatic Aberration Removal node applied to a MediaIn1 node

Inspector



Chromatic Aberration Removal controls

Controls Tab

The Controls tab contains the primary controls necessary for adjusting the Chromatic Aberration Removal parameters.

Advanced Options

Advanced Options provide additional parameters for problem shots.

- **Lens Center:** Center X and Y parameters let you offset the center of the lens if you're dealing with a reframed and re-rendered shot.
- **Stronger Correction:** Shows the location of image features that resemble fringing.

Estimation Options

The Estimation Options are used solely to highlight and identify areas of fringing. They do not affect the final output of the image. They only become available when one of the Show Estimated Fringes boxes is checked in the Aberration Correction section.

- **R/C, G/P, B/Y Balance:** These tools adjust the incoming balance between their respective colors to better identify hard to see fringing.
- **Brightness:** Magnifies the fringe indicators that are displayed when you turn on either of the Estimated Fringes checkboxes.

Aberration Correction

These tools let you make manual adjustments to correct aberration issues.

- **R/C, G/P, B/Y Scale:** Adjust these sliders to eliminate the fringing from their respective colors.
- **R/C, G/P, B/Y Edge:** Adjusts to compensate for the difference in fringing due to the curvature of the lens' edges
- **Show Estimated Fringes:** Checking this box will show just the estimated fringes over a gray background. It will also activate the Estimation Options sliders, which will let you further highlight and identify areas of fringing.

  This makes it easy to make manual adjustments to correct the problem, using the Scale and Edge controls to individually adjust Red/Cyan, Green/Purple, and Blue/Yellow fringing.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Color nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Color

A sub category of the Color tools in Fusion that gives you access to various color processing tools.

Auto Gain [AG]



The Auto Gain node

Auto Gain Node Introduction

The Auto Gain node automatically adjusts the tonal range of an image, setting the darkest and brightest pixels to user-selected values. By default, the darkest pixels get pushed to black, the brightest pixels get pushed to white, and pixels in between get stretched to cover the tonal range evenly.

This can be useful when compensating for variations in lighting, dealing with low-contrast images, or visualizing the full color range of float images (although the viewer's View Normalized Image option is generally more suitable for this).

Inputs

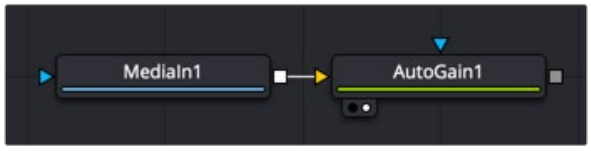
The two inputs on the Auto Gain node are the input and effect mask.

Input: The orange input connects the primary 2D image for the auto gain.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the auto gain adjustment to only those pixels within the mask. An effect mask is applied to the tool after the tool is processed.

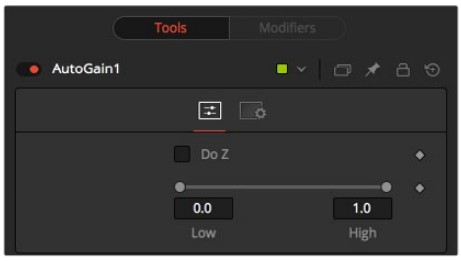
Basic Node Setup

The Auto Gain node, like many 2D image-processing nodes, receives a 2D image like a Loader node or the MediaIn1 shown below. The output continues the node tree by connecting to another 2D image-processing node or a Merge node.



An Auto Gain node applied to a MediaIn1 node

Inspector



Auto Gain controls

Controls Tab

The Controls tab contains the few primary controls necessary for customizing the AutoGain operation.

NOTE: Variations over time in the input image can cause corresponding variations in the levels of the result. For example, if a bright object moves out of an otherwise dark shot, the remaining scene gets suddenly brighter, as the remaining darker values get stretched to white. This also applies to sudden depth changes when Do Z is applied; existing objects may be pushed forward or backward when a near or far object enters or leaves the scene.

Do Z

Select the Do Z checkbox to apply the Auto Gain effect to the Z or Depth channels. This can be useful for matching the ranges of one Z-channel to another, or to view a float Z-channel in the RGB values.

Range

This Range control sets the black point and white point in the image. All tonal values in the image rescale to fit within this range.

EXAMPLE Create a horizontal gradient with the Background node. Set one color to dark gray (RGB Values 0.2). Set the other color to light gray (RGB Values 0.8).

Add an Auto Gain node and set the Low value to 0.0 and the High value to 0.5. This causes the brightest pixels to be pushed down to 0.5, and the darkest pixels get pushed to black. The remainder of the pixel values scale between those limits.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Color nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Brightness Contrast [BC]



The Brightness Contrast node

Brightness Contrast Node Introduction

The Brightness Contrast node adjusts the gain, brightness, contrast, gamma, and saturation of an image. The order of the controls represents the order in which the operations are applied. For example, gamma gets applied before contrast but after gain. The Brightness Contrast is also reversible using the Forward and Reverse buttons. So color corrections, once applied, can be reversed further downstream.

For this to work best, image processing should operate in 32-bit floating point.

Inputs

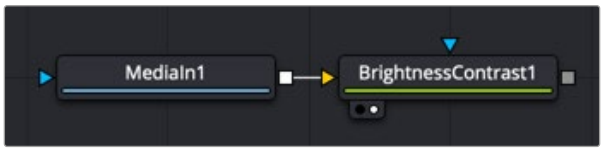
The two inputs on the Brightness Contrast node are the input and effect mask.

Input: The orange input connects the primary 2D image for the brightness contrast.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the Brightness Contrast adjustment to only those pixels within the mask. An effect mask is applied to the tool after the tool is processed.

Basic Node Setup

The Brightness Contrast node, like many 2D image-processing nodes, receives a 2D image like a Loader node or the MediaIn1 shown below. The output continues the node tree by connecting to another 2D image-processing node or a Merge node.



A Brightness Contrast node applied to a MediaIn1 node

Inspector



Brightness Control controls

Controls Tab

The Controls tab contains all the primary controls necessary for customizing the brightness, contrast operations.

Color Channels (RGBA)

The filter defaults to operating on R, G, B, and A channels. Selective channel filtering is possible by clicking each channel button to make them active or inactive.

> **NOTE:** This is not the same as the RGBA checkboxes found under the common controls. The node takes these selections into account before it processes the image, so deselecting a channel causes the node to skip that channel when processing, speeding up the rendering of the effect. In contrast, the channel controls under the Common Controls tab get applied after the node has processed.

Gain

The gain slider is a multiplier of the pixel value. A Gain of 1.2 makes a pixel that is R0.5 G0.5 B0.4 into R0.6 G0.6, B0.48 (i.e., 0.4 * 1.2 = 0.48) while leaving black pixels unaffected. Gain affects higher values more than it affects lower values, so the effect is most influential in the midrange and top range of the image.

Lift

While Gain scales the color values around black, Lift scales the color values around white. The pixel values get multiplied by the value of this control. A Lift of 0.5 makes a pixel that is R0.0 G0.0 B0.0 into R0.5 G0.5, B0.5 while leaving white pixels unaffected. Lift affects lower values more than it affects higher values, so the effect is most influential in the midrange and low range of the image.

Gamma

Values higher than 1.0 raise the Gamma (mid-gray), whereas lower values decrease it. The effect of this node is not linear, and existing black or white points are not affected at all. Pure gray colors are affected the most.

Contrast

Contrast is the range of difference between the light to dark areas. Increasing the value of this slider increases the contrast, pushing color from the midrange toward black and white. Reducing the contrast causes the colors in the image to move toward midrange, reducing the difference between the darkest and brightest pixels in the image.

Brightness

The value of the Brightness slider gets added to the value of each pixel in the image. This control's effect on an image is linear, so the effect is applied identically to all pixels regardless of value.

Saturation

Use this control to increase or decrease the amount of Saturation in the image. A saturation of 0 has no color, reducing the image to grayscale.

Low and High

This range control is similar to the Gain control in some respects. If Low gets anchored at 0.0 and the High value gets reduced from 1.0, the effect is identical to increasing the gain. High values get multiplied by the inverse of the high value. (e.g., if high is 0.75, each pixel is multiplied by 1/0.75, or 1.3333).

Leaving the high anchored at 1.0 and increasing the low is the same as inverting the image colors and increasing the gain and inverting it back again. This pushes more of the image toward black without affecting the whites at all.

Direction

Forward applies all values normally. Reverse effectively inverts all values.

Clip Black/White

The Clip Black and Clip White checkboxes clip out-of-range color values that can appear in an image when processing in floating-point color depth. Out-of-range colors are below black (0.0) or above white (1.0). These checkboxes have no effect on images processed at 8-bit or 16-bit per channel, as such images cannot have out-of-range values.

Pre-Divide/Post-Multiply

Selecting the Pre-Divide/Post-Multiply checkbox causes the image pixel values to be divided by the Alpha values before the color correction, and then re-multiplied by the Alpha value after the correction.

This helps to prevent the creation of illegally additive images when color correcting images with premultiplied Alpha channels.

Common Controls

Settings Tab

The Settings tab in the Inspector appears in other Color nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Channel Booleans [Bol]



The Channel Booleans node

Channel Booleans Node Introduction

The Channel Booleans node applies a variety of mathematical and logical operations on the channels in an image. This node works by using one image's channels to modify another image's channels. If a foreground input is not available, selecting options that use color channels from the foreground ends up using the background input's color channels instead.

NOTE: Be aware of another similarly named Channel Boolean (3Bol), which is a 3D node used to remap and modify channels of 3D materials. When modifying 2D channels, use the Channel Booleans (with an "s") node (Bol).

Inputs

There are four inputs on the Channel Booleans node in the Node Editor, but only the orange Background input is required.

Background: This orange input connects a 2D image that gets adjusted by the foreground input image.

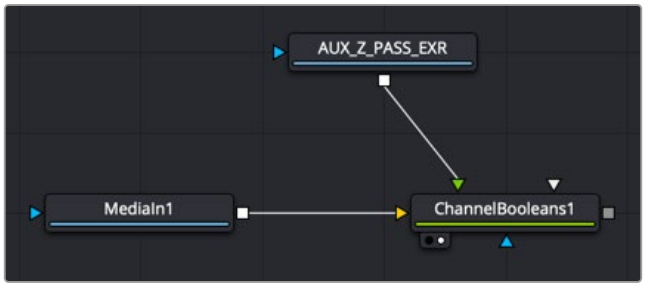
Effect Mask: The blue effect mask input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the channel booleans adjustment to only those pixels within the mask.

Foreground: The green foreground input connects a 2D image that is used to adjust the background input image.

Matte: The white matte input can be used to combine external mattes with the foreground and background operations.

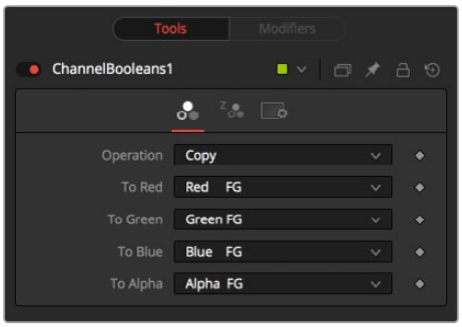
Basic Node Setup

The Channel Booleans node is an extremely flexible tool used in many different ways. The example below copies the z-depth channel from the foreground input (green) into the background image (orange).



A Channel Booleans set to copy from foreground to background

Inspector



Channel Booleans controls

Color Channel Tab

On the Color Channels tab, the controls are divided into two columns.

On the left side are target channels for the image connected into the orange background input. The drop-down menu to the right lets you choose whether you want to modify the BG image with its channels (suffix BG after list name) or with the channels from an image connected into the green foreground input on the node (suffix FG in the drop-down list).

Operation

This menu selects the mathematical operation applied to the selected channels.
The options are as follows:

Copy

Copy the value from one color channel to another. For example, copy the foreground red channel into the background's Alpha channel to create a matte.

- **Add:** Add the color values from one color channel to another channel.
- **Subtract:** Subtract the color values of one color channel from another color channel.
- **And:** Perform a logical AND on the color values from color channel to color channel. The foreground image generally removes bits from the color channel of the background image.
- **Or:** Perform a logical OR on the color values from color channel to color channel. The foreground image generally adds bits from the color channel of the background image.
- **Exclusive Or:** Perform a logical XOR on the color values from color channel to color channel. The foreground image generally flips bits from the color channel of the background image.
- **Multiply:** Multiply the values of a color channel. This gives the appearance of darkening the image as the values scale from 0 to 1. White has a value of 1, so the result would be the same. Gray has a value of 0.5, so the result would be a darker image or, in other words, an image half as bright.
- **Divide:** Divide the values of a color channel. This gives the appearance of lightening the image as the values scale from 0 to 1.
- **Maximum:** Compare the two images and take the maximum, or brightest, values from each image.
- **Minimum:** Compare the two images and take the minimum, or darkest, values from each image.
- **Negative:** Invert the FG input to make a negative version of the image.
- **Solid:** Solid sets a channel to a full value of 255. This is useful for setting the Alpha to full value.
- **Clear:** Clear sets a channel to a value of zero. This is useful for clearing the Alpha.

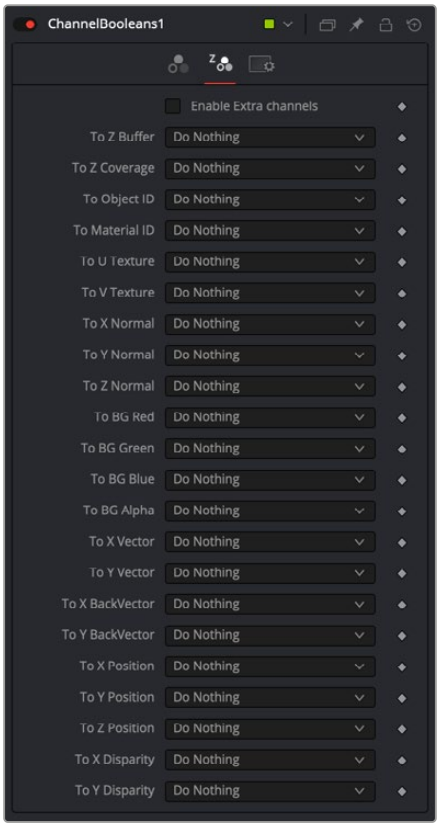
- **Difference:** Difference subtracts the greater color values of one color channel from the lesser values of another color channel.
- **Signed Add:** Signed Add subtracts areas that are lower than mid-gray and adds areas that are higher than mid-gray, which is useful for creating effects with embossed gray images.

To Red, To Green, To Blue, To Alpha

These menus represent the four color channels of the output image. Use the drop-down menu to select which channel from the source images produces the output channel.

The default setting copies the channels from the foreground channel. Select any one of the four color channels, as well as several auxiliary channels like Z-buffer, saturation, luminance, and hue.

Inspector



Aux Channel Inspector

Aux Channel Tab

This tab includes a series of menus where you select a source for the auxiliary channels of the output image.

Enable Extra Channels

When the Enable Extra Channels checkbox is selected, the Channel Booleans node can output images with channels beyond the usual RGBA. Once enabled, the remaining controls in the Aux Channels tab can copy data into the auxiliary channels.

EXAMPLES To copy the Alpha channel of one image to its color channels, set the red, green, and blue channels to Alpha BG. Set the Operation to Copy.

To copy the Alpha channel from another image, set operation type to Alpha FG.

To replace the existing Alpha channel of an image with the Alpha of another image, choose “Do Nothing” for To Red, To Green, and To Blue and “Alpha FG” for To Alpha. Pipe the image containing the Alpha into the foreground input on the Channel Booleans node. Set Operation: “Copy.” The same operation is available in the Matte Control node.

To combine any mask into an Alpha channel of an image, choose “Do Nothing” for To Red, To Green, and To Blue and “Matte” for To Alpha. Pipe the mask into the foreground input on the Channel Booleans node. Set Operation: “Copy.”

To subtract the red channel’s pixels of another image from the blue channel, choose “Do Nothing” for To Red and To Green and “Red FG” for To Blue. Pipe the image containing the red channel to subtract into the foreground input on the Channel Booleans node. Set Operation: “Subtract.”

Common Controls

Settings Tab

The Settings tab in the Inspector appears in other Color nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Color Corrector [CC]



The Color Corrector node

Color Corrector Node Introduction

The Color Corrector node is a comprehensive color node with histogram matching, and equalization, hue shifting, tinting, and color suppression.

Controls in the Color Corrector node are separated into four tabs: Correction, Ranges, Options, and Settings.

Inputs

The Color Corrector node includes four inputs in the Node Editor.

Input: This orange input is the only required connection. It connects a 2D image for color correction.

Effect Mask: The optional blue input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the color corrector adjustment to only those pixels within the mask. An effect mask is applied to the tool after the tool is processed.

Match Reference: The green input is used to connect an image that can be a reference for histogram matching.

Match Mask: This optional white input accepts any mask much like an effect mask. However, this mask defines of the area to match during a Histogram Match. It offers more flexibility in terms of shape than the built-in Match rectangle in the Inspector.

Basic Node Setup

The Color Corrector node, like many 2D image-processing nodes, receives a 2D image like a Loader node or the MediaIn1 shown below. The output continues the node tree by connecting to another 2D image-processing node or a Merge node.



A Color Corrector node applied to a MediaIn1 node

Inspector



Color Corrector controls

Correction Tab Colors Menu

The main Correction tab is further separated into four types of correction methods: colors, levels, histogram, and suppress. Selecting one from the menu at the top of the Correction tab causes that method's controls to appear. The Color method is described in detail below.

Range

This menu determines the tonal range affected by the color correction controls in this tab. The menu can be set to Shadows, Midtones, Highlights, and Master, where Master is the default affecting the entire image.

The selected range is maintained throughout the Colors, Levels, and Suppress sections of the Color Corrector node.

Adjustments made to the image in the Master channel are applied to the image after any changes made to the Highlight, Midtone, and Shadow ranges.

> **NOTE:** The controls are independent for each color range. For example, adjusting the Gamma control while in Shadows mode does not change or affect the value of the Gamma control for the Highlights mode. Each control is independent and applied separately.

Color Wheel

The color wheel provides a visual representation of adjustments made to Hue and Saturation, as well as any tinting applied to the image. Adjustments can be made directly by dragging the color indicator, or by entering values in the numeric boxes under the color wheel.

The tinting is represented in the color wheel color indicator that shows the color and strength of the tint. The Highlight setting uses a black outline for the color indicator. The Midtones and Shadows use gray color indicators. The Master color indicator is also black, but it has a white M in the center to distinguish it from the others.

The mouse can position the color indicator for each range only when the applicable range is selected. For example, the Highlight color indicator cannot be moved when the Master range is selected.

Holding down the Command or Ctrl key while dragging this indicator allows you to make finer adjustments by reducing the control's sensitivity to mouse movements. Holding down the Shift key limits the movement of the color indicator to a single axis, allowing you to restrict the effect to either tint or strength.

Tint Mode

This menu is used to select the speed and quality of the algorithm used to apply the hue and saturation adjustments. The default is Better, but for working with larger images, it may be desirable to use a faster method.

Hue

This slider is a clone of the Hue control located under the color wheel. The slider makes it easier to make small adjustments to the value with the mouse. The Hue control provides a method of shifting the hue of the image (or selected color range) through the color spectrum. The control value has an effective range between -0.1 and 1.0, which represents the angle of rotation in a clockwise direction. A value of 0.25 would be 90 degrees (90/360) and would have the effect of shifting red toward blue, green to red, and so on.

Hue shifting can be done by dragging the slider, entering a value directly into the text control, or by placing the mouse above the outer ring of the color wheel and dragging the mouse up or down. The outer ring always shows the shifted colors compared to the original colors shown in the center of the wheel.

Saturation

This slider is a clone of the Saturation control located under the color wheel. The slider makes it easier to make small adjustments to the value with the mouse. The Saturation control is used to adjust the intensity of the color values. A saturation of 0 produces gray pixels without any color component, whereas a value of 1.0 produces no change in the chroma component of the input image. Higher values generate oversaturated values with a high color component.

Saturation values can be set by dragging the slider, entering a value directly into the text control, or by dragging the mouse to the left and right on the outer ring of the color wheel control.

Channel

This menu is set for the Histogram, Color, and Levels sections of the Color Corrector node. When the red channel is selected, the controls in each mode affect the red channel only, and so on.

The controls are independent, so switching to blue does not remove or eliminate any changes made to red, green, or Master. The animation and adjustments made to each channel are separate. This menu simply determines what controls to display.

Contrast

Contrast is the range of difference between the light to dark areas. Increasing the value of this slider increases the contrast, pushing color from the midrange toward black and white. Reducing the contrast causes the colors in the image to move toward midrange, reducing the difference between the darkest and brightest pixels in the image.

Gain

The Gain slider is a multiplier of the pixel value. A gain of 1.2 makes a pixel that is R0.5 G0.5 B0.4 into R0.6 G0.6, B0.48 (i.e., 0.4 * 1.2 = 0.48), while leaving black pixels totally unaffected. Gain affects higher values more than it affects lower values, so the effect is strongest in the midrange and top range of the image.

Lift

While Gain scales the color values around black, Lift scales the color values around white. The pixel values are multiplied by the value of this control. A Lift of 0.5 makes a pixel that is R0.0 G0.0 B0.0 into R0.5 G0.5, B0.5, while leaving white pixels totally unaffected. Lift affects lower values more than it affects higher values, so the effect is strongest in the midrange and low range of the image.

Gamma

Values higher than 1.0 raise the Gamma (mid gray), whereas lower values decrease it. The effect of this node is not linear, and existing black or white points are not affected at all. Pure gray colors are affected the most.

Brightness

The value of the Brightness slider is added to the value of each pixel in your image. This control's effect on an image is linear, so the effect is applied identically to all pixels despite value.

Reset All Color Changes

Selecting this button returns all color controls in this section to their default values.

Correction Tab Levels Menu

The main Correction tab is further separated into four types of correction methods: colors, levels, histogram, and suppress. When Levels is selected from the menu, you can remap the white and black points of an image, with a Gamma control to adjust midtones. A histogram provides a view of the tonal distribution in the image to help guide your adjustments. The Level method is described in detail below.



Color Corrector Levels controls

Range

Identical to the Range menu when Color is selected in the Menu, the Range menu determines the tonal range affected by the color correction controls in this tab. The menu can be set to Shadows, Midtones, Highlights, and Master, where Master is the default affecting the entire image.

The selected range is maintained throughout the Colors, Levels, and Suppress sections of the Color Corrector node.

Adjustments made to the image in the Master channel are applied to the image after any changes made to the Highlights, Midtones, and Shadows ranges.

> **NOTE:** The controls are independent for each color range. For example, adjusting the Gamma control while in Shadows mode does not change or affect the value of the Gamma control for the Highlights mode. Each control is independent and applied separately.

Channel

This menu is used to select and display the histogram for each color channel or for the Master channel.

Histogram Display

A histogram is a chart that represents the distribution of color values in the scene. The chart reads from left to right, with the leftmost values representing the darkest colors in the scene and the rightmost values representing the brightest. The more pixels in an image with the same or similar value, the higher that portion of the chart is.

Luminance is calculated per channel; therefore, the red, green, and blue channels all have their own histogram, and the combined result of these comprises the Master Histogram.

To scale the histogram vertically, place the mouse pointer inside the control and drag the pointer up to zoom in or down to zoom out.

Display Selector Toolbar

The Display Selector toolbar at the top of the histogram provides a method of enabling and disabling components of the histogram display. Hold the mouse pointer over the button to display a tooltip that describes the button's function.

- **Input Histogram:** This enables or disables the display of the input image's histogram.
- **Reference Histogram:** This enables or disables the display of the reference image's histogram.
- **Output Histogram:** This enables or disables the display of the histogram from the post-color-corrected image.
- **Corrective Curve:** This toggles the display of a spline used to visualize exactly how auto color corrections applied using a reference image are affecting the image. This can be useful when equalizing luminance between the input and reference images.

Histogram Controls

These controls along the bottom of the histogram display are used to adjust the input image's histogram, compressing or shifting the ranges of the selected color channel.

The controls can be adjusted by dragging the triangles beneath the histogram display to the left and right.

Shifting the High value toward the left (decreasing the value) causes the histogram to slant toward white, shifting the image distribution toward white. The Low value has a similar effect in the opposite direction, pushing the image distribution toward black.

Output Level

The Output Level control can apply clipping to the image, compressing the histogram. Decreasing the High control reduces the value of pixels in the image, sliding white pixels down toward gray and gray pixels toward black.

Adjusting the Low control toward High does the opposite, sliding the darkest pixels toward white.

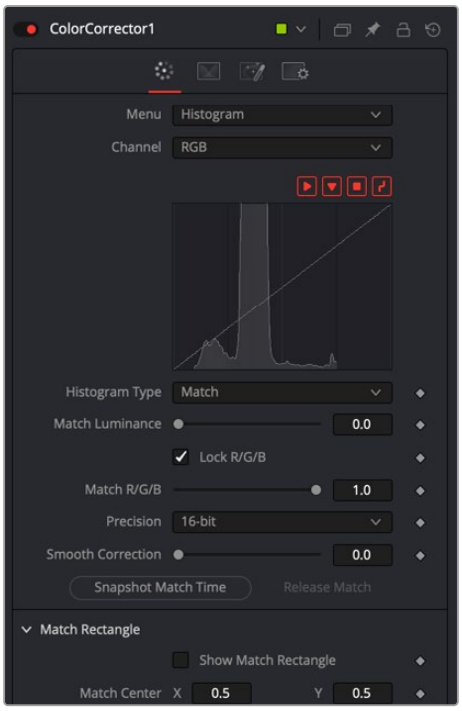
If the low value were set to 0.1, pixels with a value of 0.0 would be set to 0.1 instead, and other values would increase to accommodate the change. The best way to visualize the effect is to observe the change to the output histogram displayed above.

Reset All Levels

Clicking this button resets all the controls in the Levels section to their defaults.

Correction Tab Histogram Menu

When the menu is set to Histogram, a histogram display is produced of the input image. If a reference image is also provided, the histogram for the reference image is also displayed. The controls in this tab are primarily used to match one image to another, using either the Equalize or Match modes of the Color Corrector.



Color Correction in Histogram mode

Channel

This menu is used to select and display the histogram for each color channel or for the Master channel.

Histogram Display

A histogram is a chart that represents the distribution of color values in the scene. The chart reads from left to right, with the leftmost values representing the darkest colors in the scene and the rightmost values representing the brightest. The more pixels in an image with the same or similar value, the higher that portion of the chart is.

Luminance is calculated per channel; therefore, the red, green, and blue channels all have their own histogram, and the combined result of these comprises the Master Histogram.

To scale the histogram vertically, place the mouse pointer inside the control and drag the pointer up to zoom in or down to zoom out.

Display Selector Toolbar

The Display Selector toolbar at the top of the histogram provides a method of enabling and disabling components of the histogram display. Hold the mouse pointer over the button to display a tooltip that describes the button's function.

- **Input Histogram:** This enables or disables the display of the input image's histogram.
- **Reference Histogram:** This enables or disables the display of the reference image's histogram.
- **Output Histogram:** This enables or disables the display of the histogram from the post-color-corrected image.

- **Corrective Curve:** This toggles the display of a spline used to visualize exactly how auto color corrections applied using a reference image are affecting the image. This can be useful when equalizing luminance between the input and reference images.

Float Images and Histogram Equalization or Matching

By using the Histogram Match or Equalize methods on a float image, the color depth of the output image is converted to 16-bit integer. Two-dimensional histograms are not well suited to working with the extreme dynamic range of float images, so these operations always revert to 16-bit integer processing.

Histogram Type

Each of these menu options enables a different type of color correction operation.

- **Keep:** Keep produces no change to the image, and the reference histogram is ignored.
- **Equalize:** Selecting Equalize adjusts the source image so that all the color values in the image are equally represented—in essence, flattening the histogram so that the distribution of colors in the image becomes more even.
- **Match:** The Match mode modifies the source image based on the histogram from the reference image. It is used to match two shots with different lighting conditions and exposures so that they appear similar.

  When selected, the Equalize and Match modes reveal the following controls.

- **Match/Equalize Luminance:** This slider affects the degree that the Color Corrector node attempts to affect the image based on its luminance distribution. When this control is zero (the default), matching and equalization are applied to each color channel independently, and the luminance, or combined value of the three color channels, is not affected.

  If this control has a positive value when equalizing the image, the input image's luminance distribution is flattened before any color equalization is applied.

  If this control has a positive value when the correction mode is set to Match, the luminance values of the input are matched to the reference before any correction is applied to the R, G, and B channels.

  The Luminance and RGB controls can have a cumulative effect, and generally they are not both set to full (1.0) simultaneously.

- **Lock R/G/B:** When this checkbox is selected, color matching is applied to all color channels equally. When the checkbox is not selected, individual controls for each channel appear.

Equalize/Match R/G/B

The name of this control changes depending on whether the Equalize or Match modes have been selected. The slider can be used to reduce the correction applied to the image to equalize or match it. A value of 1.0 causes the full effect of the Equalize or Match to be applied, whereas lower values moderate the result.

Precision

This menu determines the color fidelity used when sampling the image to produce the histogram. 10-bit produces higher fidelity than 8-bit, and 16-bit produces higher fidelity than 10-bit.

Smooth Correction

Often, color equalization and matching operations introduce posterization in an image, which occurs because gradients in the image have been expanded or compressed so that the dynamic range between colors is not sufficient to display a smooth transition. This control can be used to smooth the correction curve, blending some of the original histogram back into the result for a more even transition.

Snapshot Match Time

Click this button to take a freeze snapshot of the current reference histogram, storing its current state as a snapshot in memory. If the reference histogram is not snapshot, the reference histogram is updated from frame to frame. This can cause flickering and phasing of the correction as the node tries to match a changing source to a changing reference.

Release Match

Click this button to release the current snapshot of the histogram and return to using the live reference input.

Reset All Histogram Changes

Selecting this button removes all changes made to the histogram, returning the controls to default and setting the mode back to Keep.

Correction Tab Suppress Menu

Color Suppression provides a mechanism for removing an unwanted color component from the image. The Color Wheel control is similar to that shown in the Colors section of the node, but this one is surrounded by six controls, each representing a specific color along the wheel.

To suppress a color in the selected range, drag the control that represents that color toward the center of the color wheel. The closer the control is to the center, the more that color is suppressed from the image.



Color Corrector Suppression controls

Suppression Angle

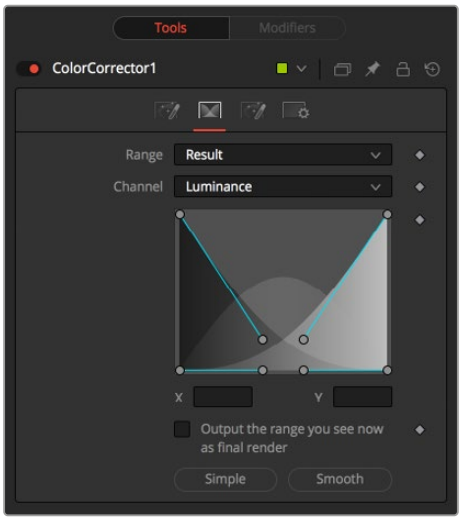
Use the Suppression Angle control to rotate the controls on the suppression wheel and zero in on a specific color.

Reset All Suppression

Clicking this control resets the suppression colors to 1.0, the default value.

Ranges Tab

The Ranges tab contains the controls used to specify which pixels in an image are considered to be shadows and which are considered to be highlights. The midrange is always calculated as pixels not already included in the shadows or the highlights.



Color Corrector Luminance controls

Range

This menu is used to select the tonal range displayed in the viewers. They help to visualize the pixels in the range. When the Result menu option is selected, the image displayed by the color corrector in the viewers is that of the color corrected image. This is the default.

Selecting one of the other menu options switches the display to a grayscale image showing which pixels are part of the selected range. White pixels represent pixels that are considered to be part of the range, and black pixels are not in the range. For example, choosing Shadows would show pixels considered to be shadows as white and pixels that are not shadows as black. Mid gray pixels are only partly in the range and do not receive the full effect of any color adjustments to that range.

Channel

The Channel menu in this tab can be used to examine the range of a specific color channel. By default, Fusion displays the luminance channel when the color ranges are examined.

Spline Display

The ranges are selected by manipulating the spline handles. There are four spline points, each with one Bézier handle. The two handles at the top represent the start of the shadow and highlight ranges, whereas the two at the bottom represent the end of the range. The Bézier handles are used to control the falloff.

The midtones range has no specific controls since its range is understood to be the space between the shadow and the highlight ranges.

The X and Y text controls below the spline display can be used to enter precise positions for the selected Bézier point or handle.

Output the Range You See Now as Final Render

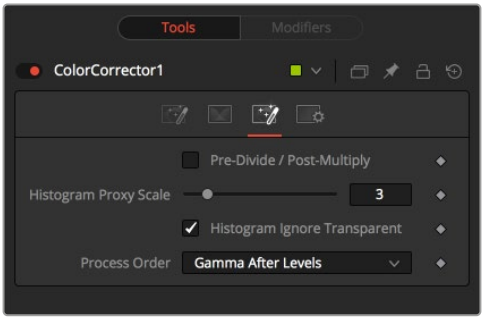
Selecting this checkbox causes the monochrome display of the range shown in the viewers to be output as the final render. Normally, the Color node outputs the full RGBA image, even if the node were left to display one of the color ranges in the view instead. This control makes it possible to use the Color Corrector node to generate a range's matte for use as an effect mask in other nodes.

Preset Simple/Smooth Ranges

These two buttons can be used to return the spline ranges to either Smooth (default) or Simple (linear) settings.

Options Tab

The Options tab includes a few very important processing operations including a simple solution when color correcting premultiplied Alpha channels.



Color Corrector controls

Pre-Divide/Post-Multiply

Selecting this option divides the color channels by the value of the Alpha before applying the color correction. After the color correction, the color values are re-multiplied by the Alpha to produce a properly additive image. This is crucial when performing an additive merge or when working with CG images generated with premultiplied Alpha channels.

Histogram Proxy Scale

The Histogram Proxy Scale determines the precision used when creating and calculating histograms. Lower values represent higher precision, and higher values produce a rougher, generalized histogram.

Process Order

This menu is used to select whether adjustments to the image's gamma are applied before or after any changes made to the images levels.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Color nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Color Curves [CCv]



The Color Curves node

Color Curves Node Introduction

The Color Curves node is a spline-based node for performing Lookup table (LUT) color manipulations. A separate spline is provided for each color channel. The effect can be animated or dissolved and can be applied to the image using RGB, YUV, YIQ, CMY, or HLS color spaces.

The LUT view in the Color Corrector can be scaled using the + and - keys on the numeric keypad. The color curves LUT fully supports out-of-range values—i.e., pixels with color values above 1.0 or below 0.0.

The splines shown in this LUT view are also available from the Spline Editor, should greater precision be required when adjusting the controls.

Inputs

The Color Curves node includes three inputs in the Node Editor.

Input: This orange input is the only required connection. It connects a 2D image that is adjusted by the color curves.

Effect Mask: The optional effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the color curves adjustment to only those pixels within the mask. An effect mask is applied to the tool after it is processed.

Reference Image: The optional green input is used to connect a second 2D image that can be used for reference matching.

Match Mask: This optional white input accepts any mask much like an effect mask. However, this mask defines of the area to match during a Match. It offers more flexibility in terms of shape than the built-in Match reference rectangle in the Inspector.

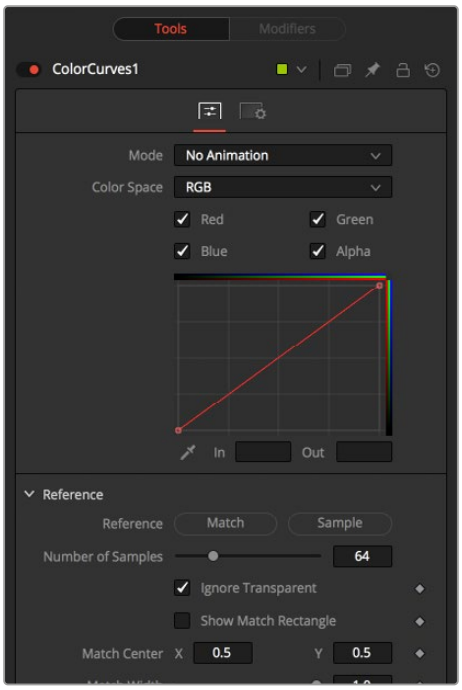
Basic Node Setup

The Color Curves node, like many 2D image-processing nodes, receives a 2D image like a Loader node or the MediaIn1 shown below. The output continues the node tree by connecting to another 2D image-processing node or a Merge node.



A Color Curves node applied to a MediaIn1 node

Inspector



Color Curves controls

Controls Tab

The Controls tab for the color curves is divided into two sections. The top half of the Inspector includes the curves and LUT controls. The bottom half is dedicated primarily to matching the reference image.

Mode

The Mode options change between Animated and Dissolve modes. The default mode is No Animation, where adjustments to the curves are static. Setting the mode provides a change spline for each channel, allowing the color curve to be animated over time.

Dissolve mode is essentially obsolete and is included for compatibility reasons only.

Color Space

The splines in the LUT view represent color channels from a variety of color spaces. The default is Red, Green, and Blue. The options in this menu allow an alternate color space to be selected.

A detailed description of the color spaces available here are below:

- **RGB (Red, Green, Blue):** Fusion uses the RGB color space, and most nodes and displays interpret the primary channels of an image as Red, Green, and Blue.
- **YUV (Luma, Blue Chroma, and Red Chroma):** The YUV color space is used in the analog broadcast of PAL video. Historically, this format was often used to color correct images, because of its familiarity to a large percentage of video engineers. Each pixel is described in terms of its Luminance, Blue Chroma, and Red Chroma components.
- **HLS (Hue, Luminance, and Saturation):** Each pixel in the HLS color space is described in terms of its Hue, Luminance, and Saturation components.

- **YIQ (Luma, In Phase, and Quadrature):** The YIQ color space is used in the analog broadcast of NTSC video. This format is much rarer than YUV and almost never seen in production. Each pixel is described in terms of its Luminance, Chroma (in-phase or red-cyan channel) and Quadrature (magenta-green) components.
- **CMY (Cyan, Magenta, and Yellow):** Although more common in print, the CMY format is often found in computer graphics from other software packages. Each pixel is described in terms of its Cyan, Magenta, and Yellow components. CMY is nonlinear.

Color Channels (RGBA)

Use the Color Channel controls to select which channel's spline is currently active for editing. The labels of these controls change to reflect the names of the channels for the current color space. Normally, they are read as Red, Green, and Blue. If the Color Curves node is operating in YUV color space, they are read as Y, U, and V instead.

These controls do not restrict the effect of the node to a specific channel. They only select whether the spline for that channel is editable. These controls are most often used to ensure that adding or moving points on one channel's spline do not unintentionally affect a different channel's spline.

Spline Window

The Spline Window displays a standard curve editor for each RGBA channel. These splines can be edited individually or as a group, depending on the color channels selected above.

The spline defaults to a linear range, from 0 in/0 out at the bottom left to the 1 in/1 out at the top right. At the default setting, a color processes to the same value as the output. If a point is added in the middle at 0.5 in/0.5 out, and the point is moved up, this raises the mid color of the image brighter.

The spline curves allow precise control over color ranges, so specific adjustments can be made without affecting other color values.

In and Out

Use the In and Out controls to manipulate the precise values of a selected point. To change a value, select a point and enter the in/out values desired.

Eyedropper (Pick)

Click the Eyedropper icon, also called the Pick button, and select a color from an image in the display to automatically set control points on the spline for the selected color. The new points are drawn with a triangular shape and can only be moved vertically (if point is locked, only the Out value can change).

Points are only added to enabled splines. To add points only on a specific channel, disable the other channels before making the selection.

One use for this technique is white balancing an image. Use the Pick control to select a pixel from the image that should be pure gray. Adjust the points that appear so that the Out value is 0.5 to change the pixel colors to gray.

Use the contextual menu's Locked Pick Points option to unlock points created using the Pick option, converting them into normal points.

Reference

The Reference section includes controls that handle matching to sample areas of the connected reference image.

- **Match Reference:** The Match Reference button adds points on the curve to match an image connected to the green reference image input. The number of points used to match the image is based on the Number of Samples slider below.
- **Sample Reference:** Clicking the Sample Reference button samples the center scanline of the background image and creates a LUT of its color values. The number of points used to match the samples scanline is based on the Number of Samples slider below.
- **Number of Samples:** This slider determines how many points are used to match the curve to the range in the reference image.
- **Show Match Rectangle:** Enabling this checkbox displays a rectangle in the viewer showing the area on the reference image used during the match process. The match rectangle affects only the result of the Match Reference operation. The Sample reference is always done from the center scaling of the image.
- **Match Center:** The X and Y parameters allow you to reposition the match rectangle to sample a different area when matching.
- **Match Width:** Width controls the width of the match rectangle.
- **Match Height:** Heigh controls the height of the match rectangle.
- **Pre-Divide/Post-Multiply:** Selecting this checkbox causes the image's pixel values to be divided by the Alpha values prior to the color correction, and then re-multiplied by the Alpha value after the correction. This helps to avoid the creation of illegally additive images, particularly around the edges of a blue/green key or when working with 3D-rendered objects.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Color nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Color Gain [Clr]



The Color Gain node

Color Gain Node Introduction

The Color Gain node contains options for adjusting the gain, gamma, saturation, and hue of the image. Many controls provided by the Color Gain node are also found in the Color Corrector node, but this simpler node may render more quickly. One feature that distinguishes the Color Gain node from the Color Corrector is its balance tab controls. These can be used to adjust the tinting of the colors in the highs, mids, and lows.

Inputs

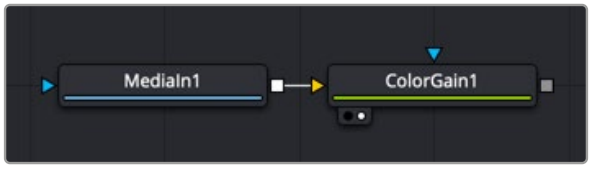
The Color Gain node includes two inputs: one for the main image and the other for an effect mask.

Input: This orange input is the only required connection. It connects a 2D image that gets adjusted by the color gain.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the color gain adjustment to only those pixels within the mask. An effect mask is applied to the tool after it is processed.

Basic Node Setup

The Color Gain node, like many 2D image-processing nodes, receives a 2D image like a Loader node or the MediaIn1 shown below. The output continues the node tree by connecting to another 2D-image processing node or a Merge node.



A Color Gain node applied to a MediaIn1 node

Inspector



Color Gain controls

Gain Tab

The Gain tab provides control of individual RGBA Lift/Gamma/Gain parameters. These controls can quickly enable you to fix irregular color imbalances in specific channels.

Lock R/G/B

When selected, the Red, Green, and Blue channel controls for each effect are combined into one slider. Alpha channel effects remain separate.

Gain RGBA

The Gain RGBA controls multiply the values of the image channel in a linear fashion. All pixels are multiplied by the same factor, but the effect is larger on bright pixels and smaller on dark pixels. Black pixels do not change because multiplying any number times 0 is always 0.

Lift RGBA

While Gain scales the color values around black, Lift scales the color values around white. The pixel values are multiplied by the value of this control. A Lift of 0.5 makes a pixel that is R0.0 G0.0 B0.0 into R0.5 G0.5, B0.5, while leaving white pixels totally unaffected. Lift affects lower values more than it affects higher values, so the effect is strongest in the midrange and low range of the image.

Gamma RGBA

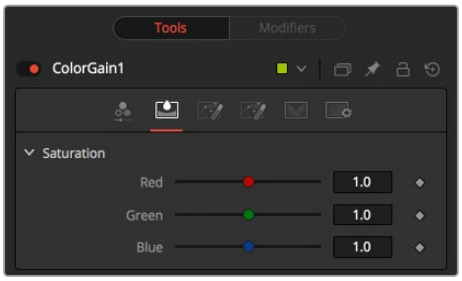
The Gamma RGBA controls affect the brightness of the midrange in the image. The effect of this node is nonlinear. White and black pixels in the image are not affected when gamma is modified, whereas pure grays are affected most by changes to this parameter. Large changes to this control tend to push midrange pixels into black or white, depending on the value used.

Pre-Divide/Post-Multiply

Selecting this checkbox causes the image pixel values to be divided by the Alpha values prior to the color correction, and then re-multiplied by the Alpha value after the correction. This helps when attempting to color correct images with premultiplied Alpha channels.

Saturation Tab

This Setting tab includes controls for the intensity of the colors in the individual RGB channels.



Color Gain Saturation setting tab

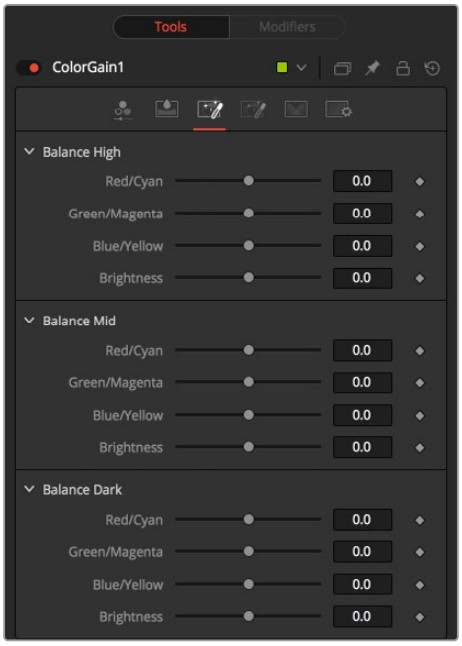
RGB Saturation

When adjusting an individual channel, a value of 0.0 strips out all that channel's color. Values greater than one intensify the color in the scene, pushing it toward the primary color.

Balance Tab

This tab in the Color Gain node offers controls for adjusting the overall balance of a color channel. Independent color and brightness controls are offered for the High, Mid, and Dark ranges of the image.

Colors are grouped into opposing pairs from the two dominant color spaces. Red values can be pushed toward Cyan, Green values toward Magenta and Blue toward Yellow. Brightness can be raised or lowered for each of the channels.



Color Gain Balance tab

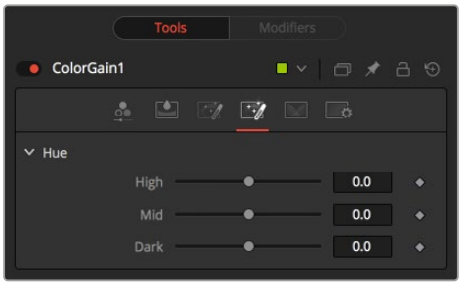
CMY Brightness Highs/Mids/Darks

By default, the Balance sliders can be adjusted by -1 to +1, but values outside this range can be entered manually to increase the effect. A value of 0.0 for any slider indicates no change to the image channel. Positive and negative values indicate that the balance of the image channel has been pushed toward one color or the other in the pair.

Hue Tab

Use the Hue tab of the Color Gain node to shift the overall hue of the image, without affecting the brightness, or saturation. Independent controls of the High, Mid, and Dark ranges are offered by three sliders.

The following is the order of the hues in the RGB color space: Red, Yellow, Green, Cyan, Blue, Magenta and Red.



Color Gain Hue tab

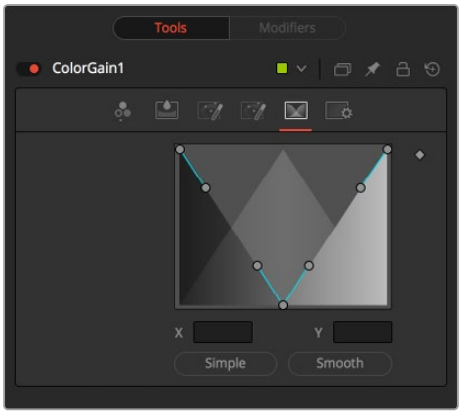
High/Mid/Dark Hue

Values above 0 push the hue of the image toward the right (red turns yellow). Values below 0 push the hue toward the left (red turns magenta). At -1.0 or 1.0, the hue completes the cycle and returns to its original value.

The default range of the hue sliders is -1.0 to +1.0. Values outside this range can be entered manually.

Ranges Tab

The Ranges tab contains the controls used to specify which pixels in an image are considered to be shadows and which are considered to be highlights. The midrange is always calculated as pixels not included in either the shadows or the highlights.



Color Gain Ranges tab

Spline Display

The ranges are selected by manipulating the spline handles. There are four spline points, each with one Bézier handle. The two handles at the top represent the start of the shadow and highlight ranges, whereas the two at the bottom represent the end of the range. The Bézier handles are used to control the falloff.

The midtones range has no specific controls since its range is understood to be the space between the shadow and the highlight ranges. The X and Y text controls below the Spline display can be used to enter precise positions for the selected Bézier point or handle.

Preset Simple/Smooth Ranges

These two buttons can be used to return the spline ranges to either Smooth (default) or Simple (linear) settings.

Settings Tab

The Settings tab in the Inspector is also duplicated in other Color nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Color Matrix [CMx]



The Color Matrix node

Color Matrix Node Introduction

The ColorMatrix allows a vast number of operations to modify values individually in the different color channels.

Inputs

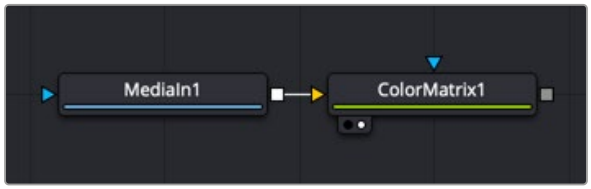
The Color Matrix node includes two inputs: one for the main image and the other for an effect mask.

Input: This orange input is the only required connection. It connects a 2D image that is adjusted by the color matrix.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the color matrix adjustment to only those pixels within the mask. An effect mask is applied to the tool after it is processed.

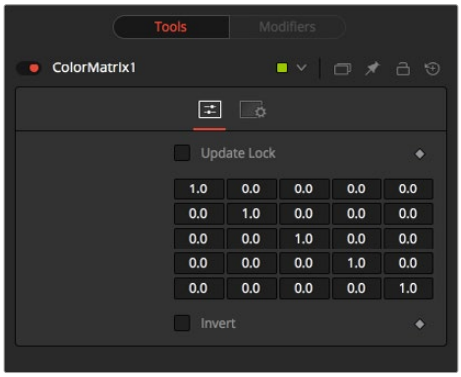
Basic Node Setup

The Color Matrix node, like many 2D image-processing nodes, receives a 2D image like a Loader node or the MediaIn1 shown below. The output continues the node tree by connecting to another 2D image-processing node or a Merge node.



A Color Matrix node applied to a MediaIn1 node

Inspector



Color Matrix controls

Controls Tab

Color Matrix multiplies the RGBA channels based on the values entered in a 4 x 4 grid. The fifth column/row is an Add column.

Update Lock

When this control is selected, Fusion does not render the node. This is useful for setting up each value of the node, and then turning off Update Lock to render it.

Matrix

This defines what type of operation actually takes place. The horizontal rows define the output values of the node. From left to right, they are R, G, B, A, and Add. The vertical columns define the input values. From top to bottom, they are R, G, B, A, and Add. The Add column allows simple adding of values to the individual color channels.

By default, the output values are identical to the input values.

- 1.0 means 100% of the Red channel input is copied to the Red channel output.
- 1.0 means 100% of the Green channel input is copied to the Green channel output.
- 1.0 means 100% of the Blue channel input is copied to the Blue channel output.
- 1.0 means 100% of the Alpha channel input is copied to the Alpha channel output.

Written as mathematical equations, the default settings of the matrix would appear as follows:

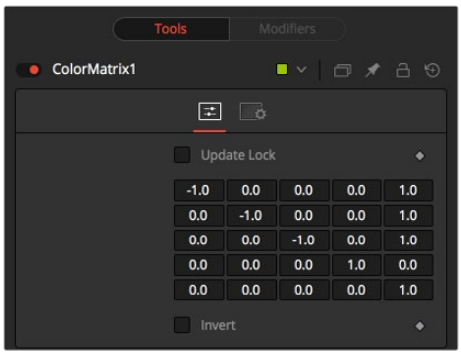
```
[R out] = 1 * [R in] + 0 * [G in] + 0 * [B in] + 0 * [A in] + 0
[G out] = 0 * [R in] + 1 * [G in] + 0 * [B in] + 0 * [A in] + 0
[B out] = 0 * [R in] + 0 * [G in] + 1 * [B in] + 0 * [A in] + 0
[A out] = 0 * [R in] + 0 * [G in] + 0 * [B in] + 1 * [A in] + 0
```

Invert

Enabling this option inverts the Matrix. Think of swapping channels around, doing other operations with different nodes, and then copying and pasting the original ColorMatrix and setting it to Invert to get your channels back to the original.

Example 1: Invert

If you want to do a simple invert or negative of the color values, but leave the Alpha channel untouched, the matrix would look like this:

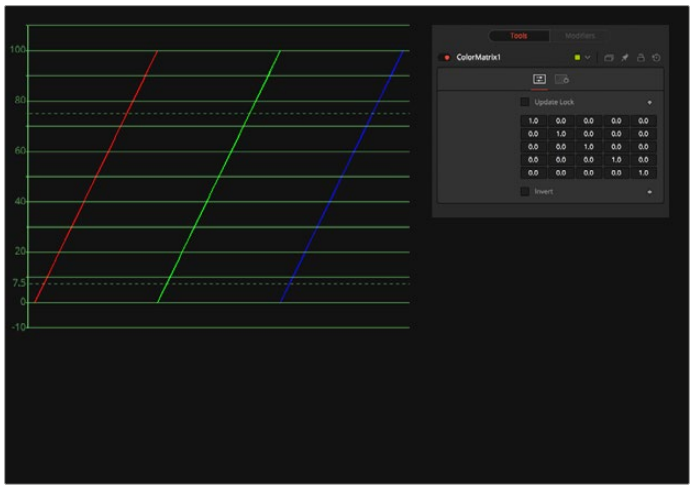


Color Matrix example

Observe the fact that we have to add 1 to each channel to push the inverted values back into the positive numbers.

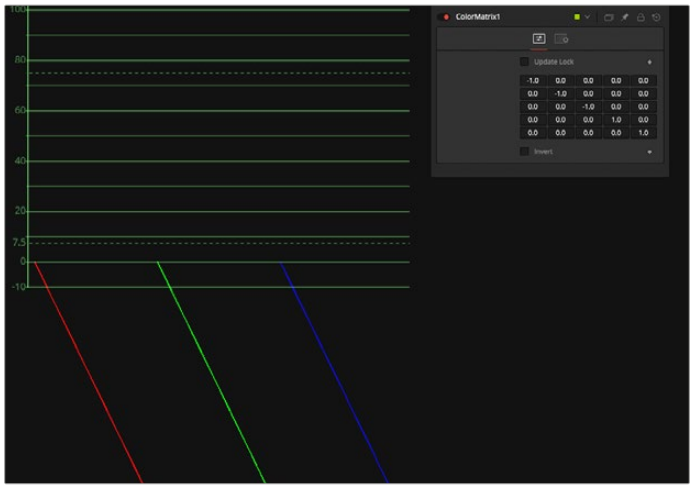
Let's follow this example step by step by viewing the waveform of a 32-bit grayscale gradient.

1 The original grayscale.



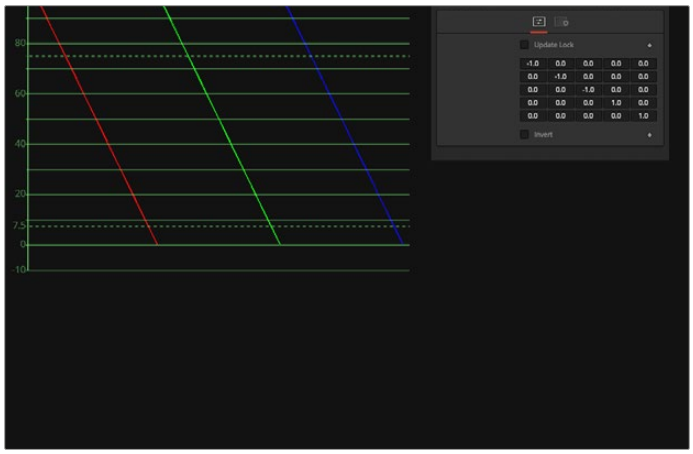
Original Grayscale

2 RGB set to -1. The values get inverted but fall below 0.



RGB set to -1

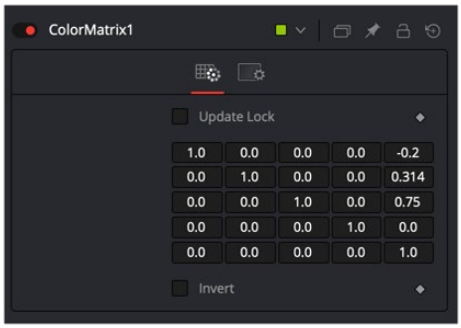
3 Adding 1 to each channel keeps the inversion but moves the values back into a positive range.



Adding 1 to each channel

Example 2: Brightness per Channel

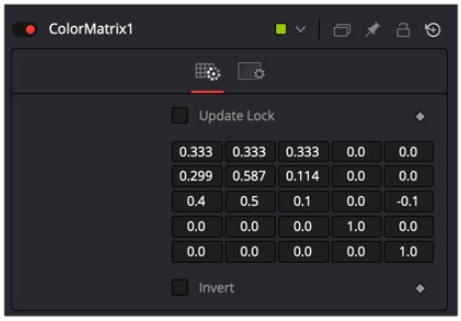
This example influences the brightness of each channel individually. This subtracts 0.2 from the red channel, adds 0.314 to the green channel, and adds 0.75 to the blue channel, while keeping Alpha as it is.



Brightness per channel example

Example 3: Copying Values

You can also copy color values back and forth between individual channels. In this example, the red channel contains the luminance values of the image based on thirds, and the green channel contains the luminance values based on the proper black-and-white conversion method, whereas the blue channel uses a third method based on getting more information from red and less from blue. The blue channel's brightness is also lowered by 0.1, and the Alpha channel is replaced with the original blue channel.



Copying Values example

Settings Tab

The Settings tab in the Inspector is also duplicated in other Color nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Color Space [CS]



The Color Space node

Color Space Node Introduction

The Color Space node provides the ability to work on an image in a variety of alternate color space formats. By default, Fusion uses the RGB color space, and most nodes and displays interpret the primary channels of an image as Red, Green, and Blue.

Changing the color space from RGB causes most images to look odd, as Fusion's viewers still interpret the primary channels as Red, Green, and Blue. For example, viewing an image converted to YUV in one of the viewers shows the Y channel as Red, the U channel as Green, and the V channel as Blue.

Several common elements of the Fusion interface refer to the RGB channels directly. The four buttons commonly found on the Inspector's Settings tab to restrict the effect of the node to a single color channel are one example. When a conversion is applied to an image, the labels of these buttons remain R, G, and B, but the values they represent are from the current color space. (For example, Red is Hue, Green is Luminance, and Blue is Saturation for an RGB to HLS conversion. The Alpha value is never changed by the color space conversion.)

Inputs

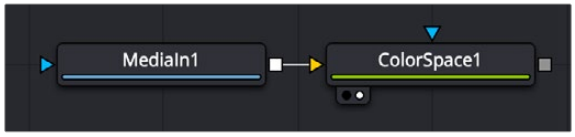
The Color Space node includes two inputs: one for the main image and the other for an effect mask.

Input: This orange input is the only required connection. It connects a 2D image that is converted by the color space operation.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the color space adjustment to only those pixels within the mask. An effect mask is applied to the tool after it is processed.

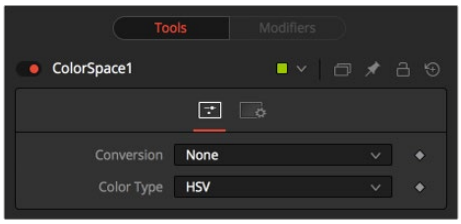
Basic Node Setup

The Color Space node, like many 2D image-processing nodes, receives a 2D image like a Loader or the MediaIn1 shown below, processes the image, and then extends the node tree by connecting to another 2D image-processing node or a Merge node.



A Color Space node applies a conversion to a MediaIn1 node

Inspector



Color Space controls

Controls Tab

The Controls tab in the Color Space node consists of two menus. The top Conversion menu determines whether you are converting an image to RGB or from RGB. The bottom menu selects the alternative color space you are either converting to or from.

Conversion

This menu has three options. The None option has no effect on the image. When To Color is selected, the input image is converted to the color space selected in the Color Type control found below. When To RGB is selected, the input image is converted back to the RGB color space from the type selected in the Color Type menu (for example, YUV to RGB).

Color Type

This menu is used to select the color space conversion applied when the To Color conversion is selected. When the To RGB option is selected in the Conversion menu, the Color Type option should reflect the input image's current color space. There are eight color space options to choose from.

- **HSV (Hue, Saturation, and Value):** Each pixel in the HSV color space is described in terms of its Hue, Saturation, and Value components. Value is defined as the quality by which we distinguish a light color from a dark one or brightness. Decreasing saturation roughly corresponds to adding white to a paint chip on a palette. Increasing Value is roughly similar to adding black.
- **YUV (Luma, Blue Chroma, and Red Chroma):** The YUV color space is used in the analog broadcast of PAL video. Historically, this format was often used to color correct images because of its familiarity to a large percentage of video engineers. Each pixel is described in terms of its Luminance, Blue Chroma, and Red Chroma components.
- **YIQ (Luma, In Phase, and Quadrature):** The YIQ color space is used in the analog broadcast of NTSC video. This format is much rarer than YUV and almost never seen in production. Each pixel is described in terms of its Luminance, Chroma (in-phase or red-cyan channel), and Quadrature (magenta-green) components.
- **CMY (Cyan, Magenta, and Yellow):** Although more common in print, the CMY format is often found in computer graphics from other software packages. Each pixel is described in terms of its Cyan, Magenta, and Yellow components. CMY is nonlinear.
- **HLS (Hue, Luminance, and Saturation):** Each pixel in the HLS color space is described in terms of its Hue, Luminance, and Saturation components. The differences between HLS and HSV color spaces are minor.
- **XYZ (CIE Format):** This mode is used to convert a CIE XYZ image to and from RGB color spaces. CIE XYZ is a weighted space, instead of a nonlinear one, unlike the other available color spaces. Nonlinear in this context means that equal changes in value at different positions in the color space may not necessarily produce the same magnitude of change visually to the eye.

  Expressed simply, the CIE color space is a perceptual color system, with weighted values obtained from experiments where subjects were asked to match an existing light source using three primary light sources.

This color space is most often used to perform gamut conversion and color space matching between image display formats because it contains the entire gamut of perceivable colors.

- **Negative:** The color channels are inverted. The color space remains RGBA.
- **BW:** The image is converted to black and white. The contribution of each channel to the luminance of the image is adjustable via slider controls that appear when this option is selected. The default values of these sliders represent the usual perceptual contribution of each channel to an image's luminance. The color space of the image remains RGBA.

Settings Tab

The Settings tab in the Inspector is also duplicated in other Color nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Copy Aux [CpA]



The Copy Aux node

Copy Aux Node Introduction

The Copy Aux node is used to shuffle channels between visible channels and auxiliary data channels in a single 2D image. Typically, these auxiliary channels are rendered from 3D applications. Auxiliary channels supported in the Copy Aux node include background color, z-depth, texture coordinates, coverage, object ID, material ID, normals, vectors, back vectors, and world position.

The Copy Aux node is mostly a convenience node, as the copying can also be accomplished with more effort using a Channel Booleans node. Where Channel Booleans deals with individual channels, Copy Aux deals with channel groups. By default, the Copy Aux node automatically promotes the depth of its output to match the depth of the aux channel.

Copy Aux also supports static normalization ranges. The advantage of static normalization versus the dynamic normalization that Fusion's viewers do is that colors remain constant over time. For example, if you are viewing Z or WorldPos values for a ball, you see a smooth gradient from white to black. Now imagine that some other 3D object is introduced into the background at a certain time. Dynamic normalization turns the ball almost completely white while the background object is now the new black. Dynamic normalization also causes flicker problems while viewing vector/disparity channels, which can make it difficult to compare the aux channels of two frames at different times visually.

Inputs

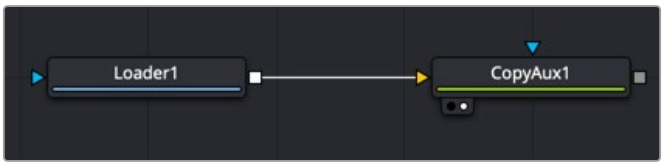
The Copy Aux node includes two inputs: one for the main image and the other for an effect mask.

Input: This orange input is the only required connection. It connects a 2D image for the Copy Aux node operation.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the Copy Aux operation to only those pixels within the mask. An effect mask is applied to the tool after the tool is processed.

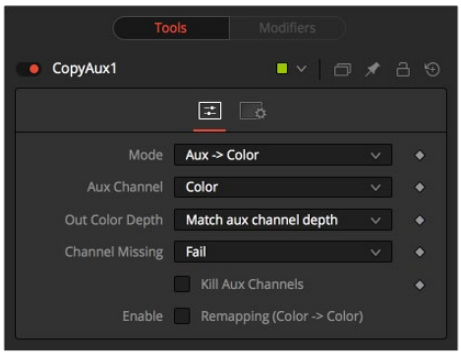
Basic Node Setup

The Copy Aux node, like many 2D image-processing nodes, receives a 2D image like a Loader node or the MediaIn1 shown below. The output continues the node tree by connecting to another 2D image-processing node or a Merge node.



A Copy Aux node applied to a MediaIn1 node

Inspector



Copy Aux controls

Controls Tab

The Controls tab is used to copy auxiliary channel groups into RGBA channels. Although Copy Aux has quite a few options, most of the time you select only the channel to copy and ignore the remaining functionality.

Mode

The Mode menu determines whether the auxiliary channel is copied into the RGBA color channel (Aux to Color) or vice versa (Color to Aux). Using this option, you can use one Copy Aux node to bring an auxiliary channel into color, do some compositing operations on it, and then use another Copy Aux node to write the color back into the auxiliary channel. When the Mode is set to Color to Aux, all the options in the Controls tab except the Aux Channel menu are hidden.

Aux Channel

The Aux Channel menu selects the auxiliary channel to be copied from or written to depending on the current mode. When the aux channel abcd has one valid component, it is copied as aaa1, two valid components as ab01, three valid components as abc1, and four components as abcd. For example, the Z-channel is copied as zzz1, texture coordinates as uv01, and normals as nxnynz1.

Out Color Depth

Out Color Depth controls the color depth of the output image. Most aux channels contain float values or, if they are integer valued, they can contain values beyond 255. When you copy float values into an int8 or int16 image, this can be a problem since negative values and values over 1.0 can get clipped. In addition, precision can be lost. This option determines what happens if the depth of RGBA channels of the input image is insufficient to contain the copied aux channel.

Be careful when copying float channels into integer image formats, as they can get clipped if you do not set up Copy Aux correctly. For this node, all aux channels are considered to be float32 except ObjectID or MaterialID, which are considered to be int16.

- **Match Aux Channel Depth:** The bit depth of the RGBA channels of the output image is increased to match the depth of the aux channel. Specifically, this means that the RGBA channels of the output image are either int16 or float32. Be careful when using this option because, for example, if you normally have int8 color channels, you are now using 2x or 4x more memory for the color channels. Particularly, the Z, Coverage, TextureCoordinate, Normal, Vector, BackVector, WorldPosition, and Disparity channels are always output as float, and the Material/ObjectID channels are output as int16.
- **Match Source Color Depth:** The bit depth of the RGBA channels of the output image is the same as the input image. This can have some unexpected consequences. For example, if your input image is int8, the XYZ components of normals that are floating-point numbers in the [-1, 1] range are clipped to non-negative numbers [0, 1] range. As a more extreme example, consider what happens to Z values. Z values are floating-point numbers stored in the [-1e30, 0] range, and they all get truncated to the [0, 1] range, which means your Z-channel is full of zeroes.
- **Force Float32:** The bit depth of the RGBA channels of the output image is always float32.

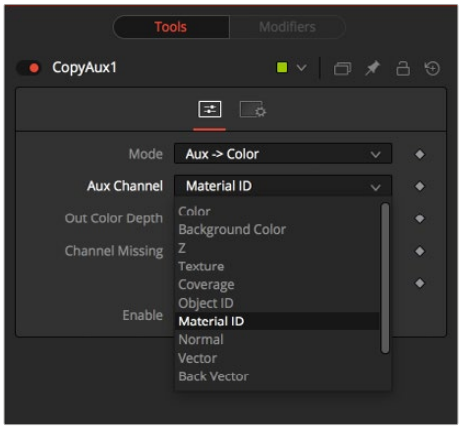
Channel Missing

Channel Missing determines what happens if a channel is not present. For example, this determines what happens if you chose to copy Disparity to Color and your input image does not have a Disparity aux channel.

- **Fail:** The node fails and prints an error message to the console.
- **Use Default Value:** This fills the RGBA channels with the default value of zero for everything except Z, which is -1e30.

Kill Aux Channels

When this is checked, Copy Aux copies the requested channel to RGBA and then outputs a resulting image that is purely RGBA with other channels being killed. This is useful if you want to increase the number of frames of Copy Aux that can be cached for playback—for example, to play back a long sequence of disparity. A handy tip is that you can use the Kill Aux feature also with just Color to Aux > Color for a longer color playback.



Kill Aux channels

Enable Remapping

When remapping is enabled, the currently selected aux channel is rescaled, linearly mapping the range according to the From and To slider selections as explained below. The Remapping options are applied before the conversion operation. This means you could set the From > Min-Max values to -1, 1 to rescale your normals into the [0, 1] range, or set them to [-1000, 0] to rescale your Z values from [-1000, 0] into the [0, 1] range before the clipping occurs.

Note that the Remapping options are per channel options. That means the default scale for normals can be set to [-1, +1] > [0, 1] and for Z it can be set [-1000, 0] > [0, 1]. When you flip between normals and Z, both options are remembered. One way this could be useful is that you can set up the remapping ranges and save this as a setting that you can reuse. The remapping can be useful to squash the aux channels into a static [0, 1] range for viewing or, for example, if you wish to compress normals into the [0, 1] range to store them in an int8 image.

- **From > Min:** This is the value of the aux channel that corresponds to To > Min.
- **From > Max:** This is the value of the aux channel that corresponds to To > Max. It is possible to set the max value less than the min value to achieve a flip/inversion of the values.
- **Detect Range:** This scans the current image to detect the min/max values and then sets the From > Min/ From > Max Value controls to these values.
- **Update Range:** This scans the current image to detect the min/max values and then enlarges the current [From > Min, From > Max] region so that it contains the min/max values from the scan.
- **To > Min:** This is the minimum output value, which defaults to 0.
- **To > Max:** This is the maximum output value, which defaults to 1.
- **Invert:** After the values have been rescaled into the [To > Min, To > Max] range, this inverts/flips the range.

Settings Tab

The Settings tab in the Inspector is also duplicated in other Color nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Gamut [Gmt]



The Gamut node

Gamut Node Introduction

The Gamut node has controls to transform one color space to another and remove/add gamma curves. This node, along with the Cineon Log node, is primarily used to linearize incoming images and then reapply the applicable output gamma curve at the end of a node tree.

Inputs

The Gamut node includes two inputs: one for the main image and the other for an effect mask to limit the conversion area.

Input: This orange input is the only required connection. It connects a 2D image output that is the source of the gamut conversion.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the Gamut operation to only those pixels within the mask. An effect mask is applied to the tool after the tool is processed.

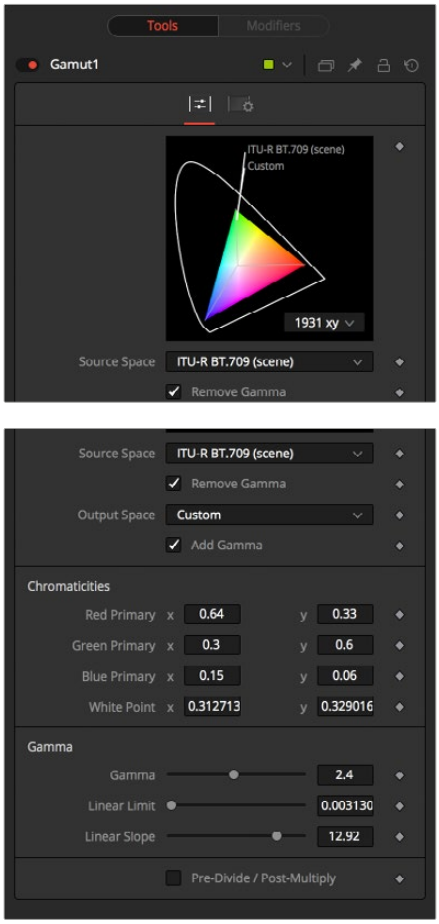
Basic Node Setup

A Gamut node is most often placed directly after the MediaIn node in DaVinci Resolve or a Loader node in Fusion Studio. Another Gamut node is usually placed at the end of a node tree before a MediaOut node in DaVinci Resolve or a Saver node in Fusion Studio.



A Gamut node applied to a MediaIn1 node

Inspector



Gamut controls

Controls Tab

The Controls tab is where all the conversion operations take place. It has a section for incoming images and a section for the node's output. Which section you use depends on whether you are stripping an image of a gamma curve to make it linear or converting a linear image to a specific color space and gamma curve for output.

Source Space

Source Space determines the input color space of the image. When placed directly after a Loader node in Fusion or a MediaIn node in DaVinci Resolve, you would select the applicable color space based on how the image was created and check the Remove Gamma checkbox. The output of the node would be a linearized image. You leave this setting at No Change when you are adding gamma using the Output Space control and placing the node directly before the Saver node in Fusion or a MediaOut node in DaVinci Resolve.

DCI-P3

The DCI-P3 color space is most commonly used in association with DLP projectors. It is frequently provided as a color space available with DLP projectors and as an emulation mode for 10-bit LCD monitors such as the HP Dreamcolor and Apple's Pro Display XDR. This color space is defined in the SMPTE-431-2 standard.

Custom

The Custom gamut allows you to describe the color space according to CIE 1931 primaries and white point, which are expressed as XY coordinates, as well as by gamma, limit, and slope. For example, the DCI-P3 gamut mentioned above would have the following values if described as a Custom color space.

Red Primary	0.68	0.32
Green Primary	0.265	0.69
Blue Primary	0.15	0.06
White Point	0.314	0.351
Gamma	2.6	–
Linear Limit	0.0313	–

To understand how these controls work, you could view the node attached to a gradient background in Waveform mode and observe how different adjustments modify the output.

Output Space

Output Space converts the gamut to the desired color space. For instance, when working with linearized images in a composite, you place the Gamut node just before the Saver node and use the Output Space to convert to the gamut of your final output file. You leave this setting at No Change when you want to remove gamma using the Source Space control.

NOTE: When outputting to HD specification Rec. 709, Fusion uses the term Scene to refer to a gamma of 2.4 and the term Display for a gamma of 2.2.

Remove/Add Gamma

Select these checkboxes to do the gamut conversion in a linear or nonlinear gamma, or simply remove or add the applicable gamma values without changing the color space.

Pre-Divide/Post-Multiply

Selecting this checkbox causes the image's pixel values to be divided by the Alpha values prior to the color correction, and then re-multiplied by the Alpha value after the correction. This helps to avoid the creation of illegally additive images, particularly around the edges of a blue/green key or when working with 3D-rendered objects.

Settings Tab

The Settings tab in the Inspector is also duplicated in other Color nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Hue Curves [HCv]



The Hue Curves node

Hue Curves Node Introduction

The Hue Curves node allows you to adjust the color in an image using a series of spline curves. Splines are provided to control the image's hue, saturation, and luminance as well as each individual color channel. An additional set of curves allows you to apply suppression to individual color channels.

The advantage of the Hue Curves node over other color correction nodes in Fusion is that the splines can be manipulated to restrict the node's effect to a very narrow portion of the image, or expanded to include a wide-ranging portion of the image. Additionally, these curves can be animated to follow changes in the image over time. Since the primary axis of the spline is defined by the image's hue, it is much easier to isolate a specific color from the image for adjustment.

Inputs

The Hue Curves node includes two inputs: one for the main image and the other for an effect mask to limit the color correction area.

Input: This orange input is the only required connection. It connects a 2D image for the Hue Curves color correction.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the Hue Curves operation to only those pixels within the mask. An effect mask is applied to the tool after it is processed.

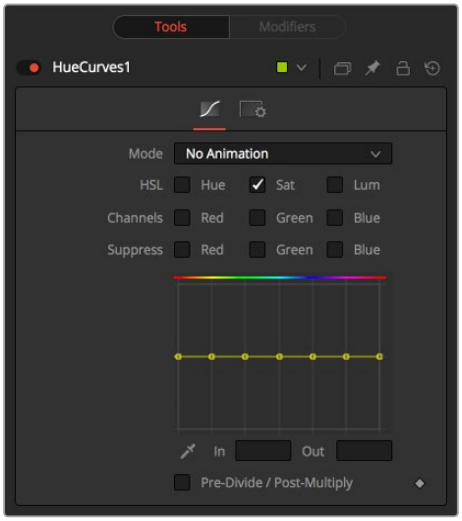
Basic Node Setup

The Hue Curves node, like many 2D image-processing nodes, receives a 2D image like a Loader node or the MediaIn1 shown below. The output continues the node tree by connecting to another 2D image-processing node or a Merge node.



A Hue Curves node applied to a Loader node in Fusion Studio

Inspector



Hue Curves controls

Controls Tab

The Controls tab consists of color attribute checkboxes that determine which splines are displayed in the Spline window. The spline graph runs horizontally across with control points placed horizontally at each of the primary colors. You can manipulate these control points to change the selected color attribute.

Mode

The Mode options change between No Animation and Animated Points modes. The default mode is No Animation, where adjustments to the curves are applied consistently over time. Setting the Mode to Animated Points or Dissolve allows the color curve to be animated over time.

Dissolve mode is essentially obsolete and is included for compatibility reasons only.

Color Channel Checkboxes

These checkboxes define which splines are editable and are included when using the Eyedropper to pick a color in the image.

Any number of activated splines can be edited simultaneously; however it's more convenient to have only the currently modified spline active to avoid unwanted changes to other splines.

When using the Eyedropper icon, a point is created on all active splines, representing the selected color.

Spline Window

This graph display is the main interface element of the Hue Curves node, which hosts the various splines. In appearance, the node is very similar to the Color Curves node, but here the horizontal axis represents the image's hue, while the vertical axis represents the degree of adjustment. The Spline window shows the curves for the individual channels. It is a miniature Spline Editor. In fact, the curves shown in this window can also be found and edited in the Spline Editor.

The spline curves for all components are initially flat, with control points placed horizontally at each of the primary colors. From left to right, these are: Red, Yellow, Green, Cyan, Blue, and Magenta. Because of the cyclical design of the hue gradient, the leftmost control point in each curve is connected to the rightmost control point of the curve.

Right-clicking in the graph displays a contextual menu containing options for resetting the curves, importing external curves, adjusting the smoothness of the selected control points, and more.

In and Out

Use the In and Out controls to manipulate the precise values of a selected point. To change a value, select a point and enter the In/Out values desired.

Eyedropper

Left-clicking and dragging from the Eyedropper icon changes the current mouse cursor to an Eyedropper. While still holding down the mouse button, drag the cursor to a viewer to pick a pixel from a displayed image. This causes control points, which are locked on the horizontal axis, to appear on the currently active curves. The control points represent the position of the selected color on the curve. Use the contextual menu's Lock Selected Points toggle to unlock points and restore the option of horizontal movement.

Points are only added to enabled splines. To add points only on a specific channel, disable the other channels before making the selection.

Pre-Divide/Post-Multiply

Selecting this checkbox causes the image's pixel values to be divided by the Alpha values prior to the color correction, and then re-multiplied by the Alpha value after the correction. This helps when color correcting images that include a premultiplied Alpha channel.

Settings Tab

The Settings tab in the Inspector is also duplicated in other Color nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

OCIO CDL Transform [OCD]



The OCIO CDL Transform node

OCIO CDL Transform Node Introduction

Fusion supports the Open Color IO color management workflow by way of three OCIO nodes.

- The OCIO CDL Transform node allows you to create, save, load, and apply a Color Decision List (CDL) grade.
- The OCIO Color Space allows sophisticated color space conversions, based on an OCIO config file.
- The OCIO File Transform allows you to load and apply a variety of Lookup tables (LUTs).

Generally, the OCIO color pipeline is composed from a set of color transformations defined by OCIO-specific config files, commonly named with a ".ocio" extension. These config files allow you to share color settings within or between facilities. The path to the config file to be used is normally specified by a user-created environment variable called "OCIO," although some tools allow overriding this. If no other *.ocio config files are located, the DefaultConfig.ocio file in Fusion's LUTs directory is used.

For in-depth documentation of the format's internals, please refer to the official pages on opencolorio.org.

Inputs

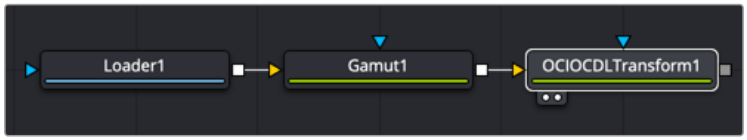
The OCIO CDL Transform node includes two inputs: one for the main image and the other for an effect mask to limit the area where the CDL is applied.

Input: This orange input is the only required connection. It connects a 2D image output for the CDL grade.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the CDL grade to only those pixels within the mask. An effect mask is applied to the tool after it is processed.

Basic Node Setup

The OCIO CDL Transform node is often applied after a Gamut node converts the Loader to linear color in Fusion Studio.



A OCIO CDL Transform node applied to a Loader node after a Gamut conversion to linear in Fusion Studio

Inspector



OCIO Transform controls

Controls Tab

The Controls tab for the OCIO CDL Transform contains primary color grading color correction controls in a format compatible with CDLs. You can make R, G, B adjustments based on the Slope, Offset, and Power. There is also overall Saturation control. You can also use the Controls tab to import and export the CDL compatible adjustments.

Operation

This menu switches between File and Controls. In File mode, standard ASC-CDL files can be loaded. In Controls mode, manual adjustments can be made to Slope, Offset, Power, and Saturation, and the CDL file can be saved.

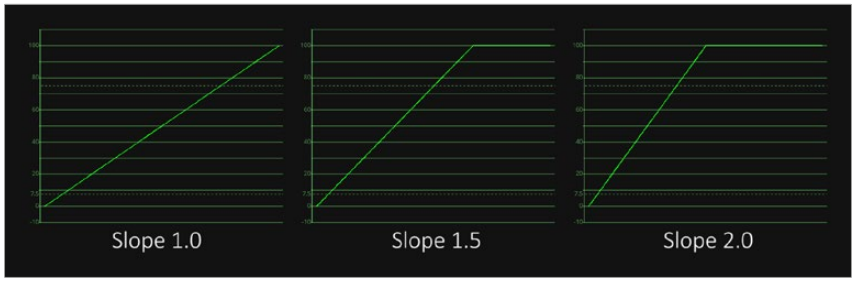
NOTE: Using DaVinci Resolve terminology, slope is similar to gain. It controls mids-to-high contrast. Offset is the overall offset of color balance and exposure. Power is very similar to contrast with a raised pivot, giving you control over shadow contrast.

Direction

Toggles between Forward and Reverse. Forward applies the corrections specified in the node, while Reverse tries to remove those corrections. Keep in mind that not every color correction can be undone.

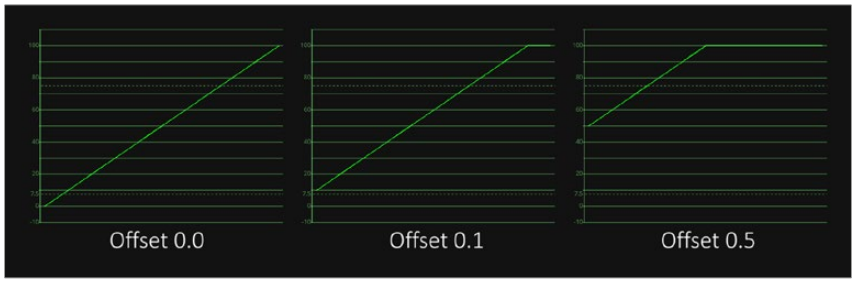
Imagine that all slope values have been set to 0.0, resulting in a fully black image. Reversing that operation is not possible, neither mathematically nor visually.

Slope



Slope multiplies the color values; this is the same as Gain in the Brightness Contrast node.

Offset



Offset adds to the color values; this is the same as Brightness in the Brightness Contrast node.

Power

Applies a Gamma Curve. This is an inverse of the Gamma function of the Brightness Contrast node.

Saturation

Enhances or decreases the color saturation. This works the same as Saturation in the Brightness Contrast node.

Export File

Allows the user to export the settings as a CDL file.

Settings Tab

The Settings tab in the Inspector is also duplicated in other Color nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

OCIO Color Space [OCC]



The OCIO Color Space node

OCIO Color Space Node Introduction

Fusion supports the Open Color IO color management workflow by way of three OCIO nodes.

- The OCIO CDL Transform node allows you to create, save, load, and apply a Color Decision List (CDL) grade.
- The OCIO Color Space allows sophisticated color space conversions, based on an OCIO config file.
- The OCIO File Transform allows you to load and apply a variety of Lookup tables (LUTs).

Generally, the OCIO color pipeline is composed from a set of color transformations defined by OCIO-specific config files, commonly named with a ".ocio" extension. These config files allow you to share color settings within or between facilities. The path to the config file to be used is normally specified by a user-created environment variable called "OCIO," though some tools allow overriding this. If no other *.ocio config files are located, the DefaultConfig.ocio file in Fusion's LUTs directory is used.

For in-depth documentation of the format's internals, please refer to the official pages on opencolorio.org.

Sample configs can be obtained from https://opencolorio.readthedocs.io/en/latest/configurations/_index.html#configurations

The functionality of the OCIO Color Space node is also available as a View LUT node from the View LUT menu.

Inputs

The OCIO Color Space node includes two inputs: one for the main image and the other for an effect mask to limit the area where the color space conversion is applied.

Input: This orange input is the only required connection. It connects a 2D image for the color space conversion.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the color space conversion to only those pixels within the mask. An effect mask is applied to the tool after it is processed.

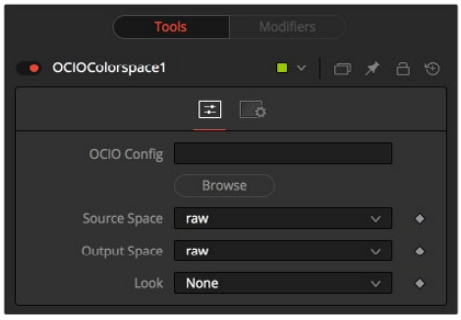
Basic Node Setup

The OCIO Color Space node is typically placed directly after a MediaIn node in DaVinci Resolve or a Loader node in Fusion Studio. Another OCIO Color Space node is placed just before a Media Out node in DaVinci Resolve or a Saver node in Fusion Studio.



An OCIO Color Space node applied to a Loader node and a Saver node in Fusion Studio

Inspector



OCIO Color Space controls

Controls Tab

The Controls tab for the OCIO Color Space node allows you to convert an image from one color space to another based on an OCIO config file. By default, it uses the config file included with Fusion; however, the Controls tab does allow you to load your own config file as well.

OCIO Config

Displays a File > Open dialog to load the desired config file.

Source Space

Based on the config file, the available source color spaces are listed here.

The content of this list is based solely on the loaded profile and hence can vary immensely. If no other OCIO config file is loaded, the DefaultConfig.ocio file in Fusion's LUTs directory is used to populate this menu.

Output Space

Based on the config file, the available output color spaces are listed here.

The content of this list is based solely on the loaded profile and hence can vary immensely. If no other OCIO config file is loaded, the DefaultConfig.ocio file in Fusion's LUTs directory is used to populate this menu.

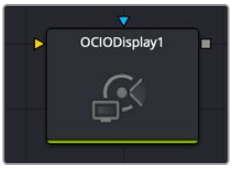
Look

Installed OCIO Color Transform Looks appear in this menu. If no looks are installed, this menu has only None listed as an option.

Settings Tab

The Settings tab in the Inspector is also duplicated in other Color nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

OCIO Display [OCD]



The OCIO Display node

OCD Display Node Introduction

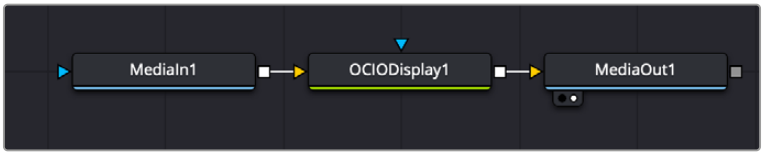
The OCIO Display node allows you to load an OpenColorIO (OCIO) Display transform (.ocio) to an input.

Inputs

The single input on the OCD Display node is used to connect the node you want to apply the OCIO Display transform to, and an effect mask, which can be used to limit the display transform area.

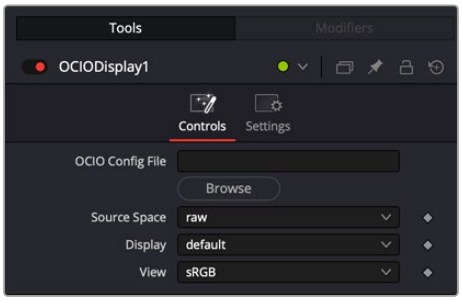
Basic Node Setup

Connect the media or tool that you wish to apply the OCIO Display transform to the input. The output of the node then gets passed to the remainder of the node tree.



A simple node tree showing the OCIO Display node taking the Media In Node, performing the Display Transform, and passing it to Media Out

Inspector



The OCIO Display controls

Controls

The OCIO Display controls lets you load and apply an .ocio configuration file and apply its color space and display transform to to the node's input.

- **OCIO Config File:** Click the Browse button to open up a file system prompt to locate the .ocio file.
- **Source Space:** Choose the color space of the source material from the list.
- **Display:** Choose which monitor your viewer is on.
- **View:** Choose which color space to view.

OCIO File Transform [OCF]



The OCIO File Transform node

OCIO File Transform Node Introduction

Fusion supports the Open Color IO color management workflow by way of three OCIO nodes.

- The OCIO CDL Transform node allows you to create, save, load, and apply a Color Decision List (CDL) grade.
- The OCIO Color Space allows sophisticated color space conversions, based on an OCIO config file.
- The OCIO File Transform allows you to load and apply a variety of Lookup tables (LUTs).

Generally, the OCIO color pipeline is composed from a set of color transformations defined by OCIO-specific config files, commonly named with a ".ocio" extension. These config files allow you to share color settings within or between facilities. The path to the config file to be used is normally specified by a user-created environment variable called "OCIO," though some tools allow overriding this. If no other *.ocio config files are located, the DefaultConfig.ocio file in Fusion's LUTs directory is used.

For in-depth documentation of the format's internals, please refer to the official pages on opencolorio.org.

The functionality of the OCIO File Transform node is also available as a View LUT node from the View LUT menu.

Inputs

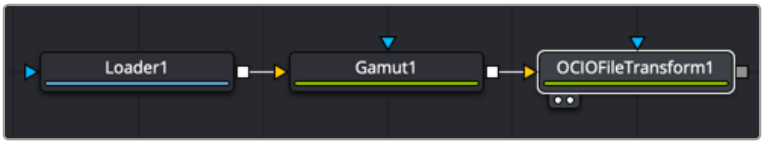
The OCIO File Transform node includes two inputs: one for the main image and the other for an effect mask to limit the area where the color space conversion is applied.

Input: This orange input is the only required connection. It connects a 2D image for the LUT.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the applied LUT to only those pixels within the mask. An effect mask is applied to the tool after it is processed.

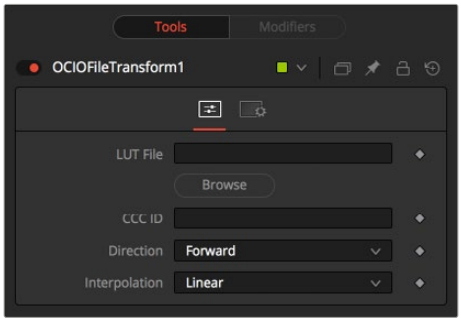
Basic Node Setup

The OCIO File Transform node is often applied after a Gamut node converts the Loader to linear color in Fusion Studio.



An OCIO File Transform node applied to a Loader node after a Gamut node conversion to linear color space

Inspector



OCIO File Transform controls

Controls Tab

The Controls tab for the OCIO File Transform node includes options to import the LUT, invert the transform, and select the color interpolation method.

LUT File

Displays a File > Open dialog to load the desired LUT.

CCC ID

This is the ID key used to identify the specific file transform located within the ASC CDL color correction XML file.

Direction

Toggles between Forward and Reverse. Forward applies the corrections specified in the node, while Reverse tries to remove those corrections. Keep in mind that not every color correction can be undone. Imagine that all slope values have been set to 0.0, resulting in a fully black image. Reversing that operation is not possible, neither mathematically nor visually.

Interpolation

Allows the user to select the color interpolation to achieve the best quality/render time ratio. Nearest is the fastest interpolation, while Best is the slowest.

Settings Tab

The Settings tab in the Inspector is also duplicated in other Color nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Set Canvas Color [SCv]



The Set Canvas Color node

Set Canvas Color Node Introduction

Set Canvas Color is used to set the color of the area outside the domain of definition (DoD). This is the workspace area beyond the raster by default, which is invisible since outside the raster is not rendered. However, the DoD can be within the raster as well. This can occur when compositing images smaller than the raster, or with transforms. By default, the canvas color used is black/no Alpha (transparent). However, since some nodes may change an image's canvas color—for example, inverting a mask changes the mask's canvas from black to white—the Set Canvas Color allows you to control the color of the canvas to whatever you require.

The Set Canvas Color node sets the color of the workspace outside the domain of definition (DOD). For example, if you create a circular gradient, the DoD is a square around the circular gradient in the viewer. Everything outside the DoD is understood to be black and therefore does not have to be rendered. To change the area outside the DoD, attach the Set Canvas Color node after the background and change the color.

> **NOTE:** Position the mouse pointer in a black area outside the raster to view the RGB canvas color in the status bar at the bottom left of the Fusion window.

Inputs

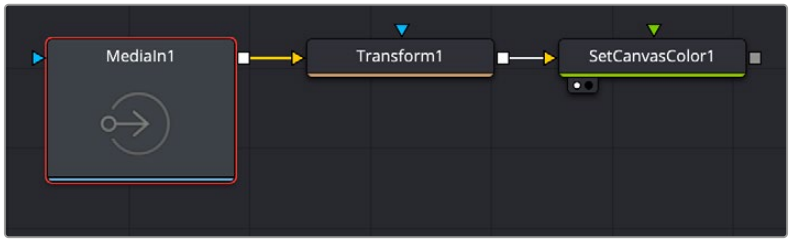
The Set Canvas Color node includes two inputs: one for the main image and a second for a foreground.

Input: This orange input is the only required connection. It accepts a 2D image that reveals the canvas color if the image's DoD is smaller than the raster.

Foreground: The optional green foreground input allows the canvas color to be sampled from an image connected to this input.

Basic Node Setup

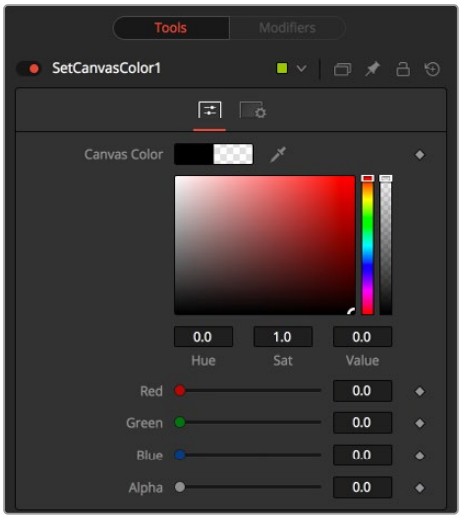
The Set Canvas Color node is placed after the image is transformed to reveal part of the raster outside the domain of definition.



A Set Canvas Color node applied to a transformed MediaIn 1 node

In this example, the footage from MediaIn1 has been shrunk in size using the Transform1 node, leaving empty space around the image in the frame. Feeding that output into the SetCanvasColor node allows you to change the color of the empty space around the frame.

Inspector



Set Canvas Color controls

Controls Tab

The Controls tab for the Set Canvas Color is used for simple color selection. When the green foreground is connected, the tab is empty.

Color Picker

Use these controls to adjust the Color and the Alpha value for the image's canvas. It defaults to black with zero Alpha.

Settings Tab

The Settings tab in the Inspector is also duplicated in other Color nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

White Balance [WB]



The White Balance node

White Balance Node Introduction

The White Balance node can be used to automatically remove color casts in the image caused by the incorrect setup of a camera or bad lighting conditions.

Correction can be done by selecting a color temperature or by choosing a neutral color from the original image that exhibits the color cast to be corrected.

IMPORTANT: When picking neutral colors using the Custom method, make sure you are picking from the source image, not the results of the White Balance node. This ensures that the image doesn't change while you are still picking, and that the White Balance node gets an accurate idea of the original colors it needs to correct.

Inputs

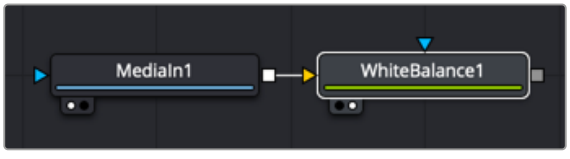
The White Balance node includes two inputs: one for the main image and the other for an effect mask to limit the area where the white balance is applied.

Input: This orange input is the only required connection. It connects a 2D image for the white balance.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the white balance to only those pixels within the mask. An effect mask is applied to the tool after it is processed.

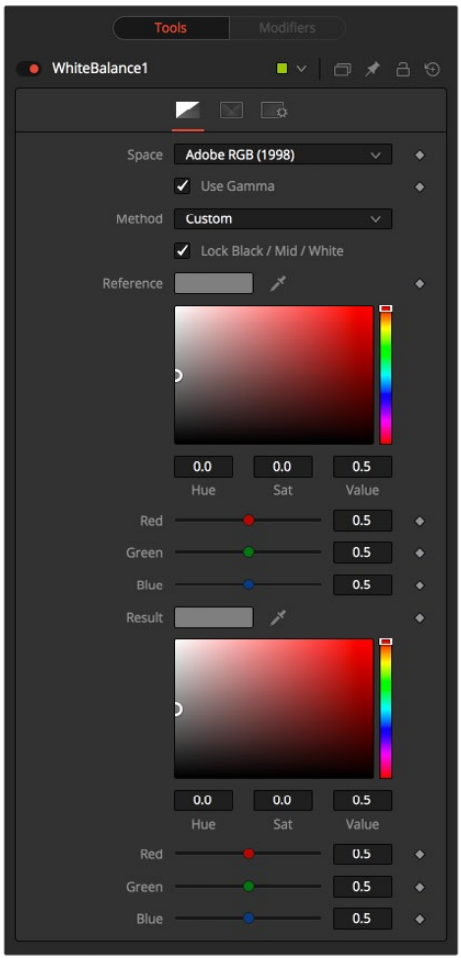
Basic Node Setup

The White Balance node, like many 2D image-processing nodes, receives a 2D image like the MediaIn1 shown below. The output continues the node tree by connecting to another 2D image-processing node or a Merge node.



A White Balance node applied to a MediaIn1 node

Inspector



White Balance controls

Balance Tab

Space

Use this menu to select the color space of the source image, if it is known. This can make the correction more accurate since the node can take the natural gamma of the color space into account as part of the correction. If the color space that the image uses is unknown, leave this menu at its default value.

Method

The White Balance node can operate using one of two methods: a Custom method or a color Temperature method.

- **Custom:** The Custom method requires the selection of a pixel from the scene that should have been pure gray. The node uses this information to calculate the color correction required to convert the pixel so that it actually is gray. When the correction is applied without an effect mask connected and the LockBlack/Mid/White checkbox enabled, the node white balances the entire shot.
- **Temperature:** The color Temperature method requires that the actual color temperature of the shot be specified.

Lock Black/Mid/White

This checkbox locks the Black, Midtones, and White points together so that the entire image is affected equally. Unchecking the control provides individual controls for white balancing each range separately. This control affects both methods equally.

Black/Mid/White Reference

These controls appear only if the Custom method is selected. They are used to select a color from a pixel in the source image. The White Balance node color corrects the image so that the selected color is transformed to the color set in the Result Color Picker below. Generally, this is gray. A color that is supposed to be pure gray but is not truly gray for one reason or another should be selected.

If the Lock Black/Mid/White checkbox is deselected, different references can be selected for each color range.

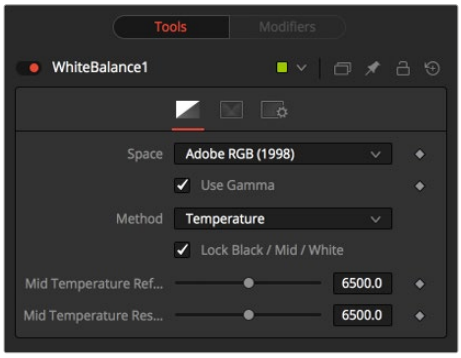
For example, try to select a pixel for the black and white references that are not clipped in any of the color channels. In the high end, an example would be a pixel that is light pink with values of 255, 240, 240. The pixel is saturated/clipped in the red, although the color is not white. Similarly, a really dark blue-gray pixel might be 0, 2, 10. It is clipped in red as well, although it is not black.

Neither example would be a good choice as a reference pixel because there would not be enough headroom left for the White Balance node.

Black/Mid/White Result

These controls appear only if the Custom method is selected. They are used to select the color that the node uses to balance the reference color. This generally defaults to pure, midrange gray.

If the Lock Black/Mid/White checkbox is deselected, different results can be selected for each color range.



White Balance Temperature controls

Temperature Reference

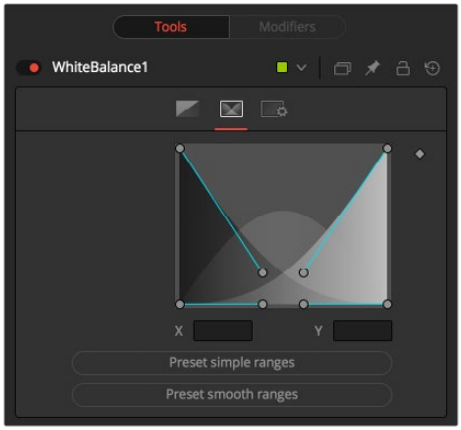
When the Method menu is set to Temperature, the Temperature reference control is used to set the color temperature of the source image. If the Lock Black/ Mid/White checkbox is deselected, different references can be selected for each color range.

Temperature Result

Use this control to set the target color temperature for the image. If the Lock Black/Mid/White checkbox is deselected, different results can be selected for each color range.

Use Gamma

This checkbox selects whether the node takes the gamma of the image into account when applying the correction, using the default gamma of the color space selected in the Space menu at the top of the tab.



White Balance range tab

Ranges Tab

The Ranges tab can be used to customize the range of pixels in the image considered to be shadows, midtones, and highlights by the node.

Spline Display

The ranges are selected by manipulating the spline handles. There are four spline points, each with one Bézier handle. The two handles at the top represent the start of the shadow and highlight ranges, whereas the two at the bottom represent the end of the range. The Bézier handles are used to control the falloff.

The midtones range has no specific controls since its range is understood to be the space between the shadow and the highlight ranges.

The X and Y text controls below the Spline display can be used to enter precise positions for the selected Bézier point or handle.

Preset Simple/Smooth Ranges

These two buttons can be used to return the spline ranges to either Smooth (default) or Simple (linear) settings.

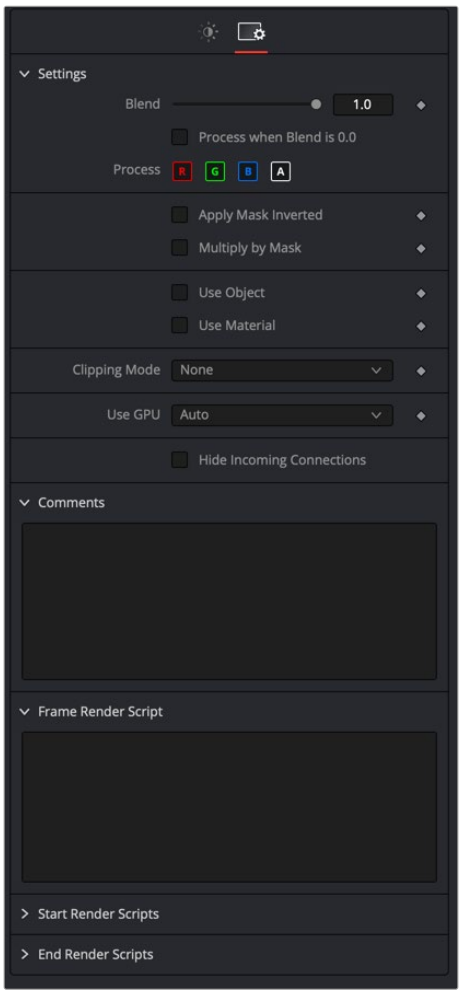
Settings Tab

The Settings tab in the Inspector is also duplicated in other Color nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

The Common Controls

Nodes that handle Color adjustment operations share several identical controls in the Inspector. This section describes controls that are common among color nodes.

Inspector



Color Settings Inspector

Settings Tab

The Settings tab in the Inspector can be found on every tool in the Color category. The Settings controls are even found on third-party color type plugin tools. The controls are consistent and work the same way for each tool, although some tools do include one or two individual options that are also covered here.

Blend

The Blend control is used to blend between the tool's original image input and the tool's final modified output image. When the blend value is 0.0, the outgoing image is identical to the incoming image. Normally, this causes the tool to skip processing entirely, copying the input straight to the output.

Process When Blend Is 0.0

The tool is processed even when the input value is zero. This can be useful if processing of this node is scripted to trigger another task, but the value of the node is set to 0.0.

Red/Green/Blue/Alpha Channel Selector

These four buttons are used to limit the effect of the tool to specified color channels. This filter is often applied after the tool has been processed.

For example, if the red button on a blur tool is deselected, the blur is first applied to the image, and then the red channel from the original input is copied back over the red channel of the result.

There are some exceptions, such as tools for which deselecting these channels causes the tool to skip processing that channel entirely. Tools that do this generally possess a set of identical RGBA buttons on the Controls tab in the tool. In this case, the buttons in the Settings and the Controls tabs are identical.

Apply Mask Inverted

Enabling the Apply Mask Inverted option inverts the complete mask channel for the tool. The mask channel is the combined result of all masks connected to or generated in a node.

Multiply by Mask

Selecting this option causes the RGB values of the masked image to be multiplied by the mask channel's values. This causes all pixels not included in the mask (i.e., set to 0) to become black/transparent.

Use Object/Use Material (Checkboxes)

Some 3D software can render to file formats that support additional channels. Notably, the EXR file format supports Object and Material ID channels, which can be used as a mask for the effect. These checkboxes determine whether the channels are used, if present. The specific Material ID or Object ID affected is chosen using the next set of controls.

Correct Edges

This checkbox appears only when the Use Object or Use Material checkboxes are selected. It toggles the method used to deal with overlapping edges of objects in a multi-object image. When enabled, the Coverage and Background Color channels are used to separate and improve the effect around the edge of the object. If this option disabled (or no Coverage or Background Color channels are available), aliasing may occur on the edge of the mask.

> For more information on the Coverage and Background Color channels, *see Chapter 18, "Understanding Image Channels,"* in the Fusion Reference Manual.

Object ID/Material ID (Sliders)

Use these sliders to select which ID is used to create a mask from the object or material channels of an image. Use the Sample button in the same way as the Color Picker: to grab IDs from the image displayed in the viewer. The image or sequence must have been rendered from a 3D software package with those channels included.

Clipping Mode

This option determines how edges are handled when performing domain of definition rendering. This is mostly important for nodes like Blur, which may require samples from portions of the image outside the current domain.

- **Frame:** The default option is Frame, which automatically sets the node's domain of definition to use the full frame of the image, effectively ignoring the current domain of definition. If the upstream DoD is smaller than the frame, the remaining area in the frame is treated as black/transparent.
- **Domain:** Setting this option to Domain respects the upstream domain of definition when applying the node's effect. This can have adverse clipping effects in situations where the node employs a large filter.
- **None:** Setting this option to None does not perform any source image clipping at all. This means that any data required to process the node's effect that would normally be outside the upstream DoD is treated as black/transparent.

Use GPU

The Use GPU menu has three settings. Setting the menu to Disable turns off hardware-accelerated rendering using the graphics card in your computer. Enabled uses the hardware. Auto uses a capable GPU if one is available and falls back to software rendering when a capable GPU is not available.

Motion Blur

- **Motion Blur:** This toggles the rendering of Motion Blur on the tool. When this control is toggled on, the tool's predicted motion is used to produce the motion blur caused by the virtual camera's shutter. When the control is toggled off, no motion blur is created.
- **Quality:** Quality determines the number of samples used to create the blur. A quality setting of 2 causes Fusion to create two samples to either side of an object's actual motion. Larger values produce smoother results but increase the render time.
- **Shutter Angle:** Shutter Angle controls the angle of the virtual shutter used to produce the motion blur effect. Larger angles create more blur but increase the render times. A value of 360 is the equivalent of having the shutter open for one whole frame exposure. Higher values are possible and can be used to create interesting effects.
- **Center Bias:** Center Bias modifies the position of the center of the motion blur. This allows the creation of motion trail effects.
- **Sample Spread:** Adjusting this control modifies the weighting given to each sample. This affects the brightness of the samples.

Comments

The Comments field is used to add notes to a tool. Click in the field and type the text. When a note is added to a tool, a small red square appears in the lower-left corner of the node when the full tile is displayed, or a small text bubble icon appears on the right when nodes are collapsed. To see the note in the Node Editor, hold the mouse pointer over the node to display the tooltip.

Scripts

Three Scripting fields are available on every tool in Fusion from the Settings tab. They each contain edit boxes used to add scripts that process when the tool is rendering. For more details on scripting nodes, please consult the Fusion scripting documentation.

Chapter 35

Composite Nodes

This chapter details the Dissolve and Merge nodes available in Fusion.

The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Dissolve [Dx] 993

Merge [Mrg] 996

MultiMerge [MMrg] 1005

The Common Controls 1009

Dissolve [Dx]



The Dissolve node

Dissolve Node Introduction

The Dissolve node is typically used to mix two images together, providing a gradual transition between two clips. A Background/Foreground slider controls the mix between the foreground and background images. Dissolves are commonly used to transition between one clip and another and are a very common effect in editing. However, you can also use the extreme left and right positions of the Background/Foreground slider to switch between inputs. Unlike other nodes in Fusion, the Dissolve node does not require you to connect an image to the background but lets you output either the background or foreground according to the setting of the Background/Foreground slider.

This quality makes it possible for you use the Dissolve node as an automatic layer switching tool when connected to background and foreground clips with different durations. Simply connect each clip to the background and foreground inputs, respectively, and set the Background/Foreground slider to the input of shorter duration, to determine which is "on top." After the last frame of that clip has ended, the Dissolve node automatically switches to the clip that's connected to the other input.

Besides the default dissolve, the Gradient Wipe setting of the Operation menu allows you to create arbitrary animated dissolve patterns based on the luminance of an image connected to the optional Gradient Wipe input. You can use this capability with images of geometric shapes or gradients of different kinds, movie clips of fire, water ripples, or rain, the Fast Noise node, or even particle systems you create within the Fusion page to create a variety of unique and creative transitions. Soft-edged effect masks may also be used to add to the possible effects.

Ultimately, animating the Background/Foreground control allows you to control the transition that's being used to switch from the foreground input to the background, or vice versa.

Inputs

The Dissolve node provides three image inputs, all of which are optional:

Background: The first of two images you want to switch between or mix. Unlike most other nodes, it is unnecessary to connect the background input before connecting the foreground input.

Foreground: The second of two images you want to switch between or mix. The Dissolve node works best when both foreground and background inputs are connected to images with the same resolution.

Gradient Map: (Optional) The Gradient Map is required only when Gradient Wipe is selected.

Basic Node Setup

Dissolve nodes are typically connected in the following way, with two input images connected to the background and foreground inputs, and the output connected to the next node in the composition.



A typical dissolve node structure in Fusion

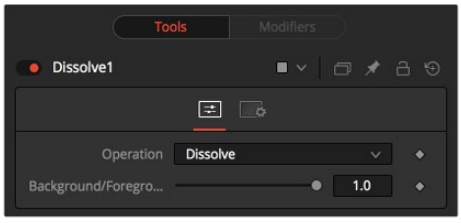
Resolution Handling

It is recommended to make sure that all images connected to the foreground, background, and gradient map inputs of the Dissolve node have the same resolution and the same pixel aspect. This is not required, however. But, the result if you mix resolutions depends on how you set the Background/Foreground slider.

- If the input images are different sizes, but the Foreground/Background slider is set to full Foreground (all the way to the right) or full Background (all the way to the left), then the output resolution will be identical to the image resolution of the corresponding node input.
- If input images of different sizes are mixed by setting the Background/Foreground slider somewhere between, the output resolution will be set to the larger of the two input resolutions to make sure there's enough room to contain both images. In this case, you may experience undesirable resolution changes when the slider moves from full foreground or background to somewhere in between.

For example, if you try to dissolve between a 4K image (connected to the background) and an 8K image (connected to the foreground), the output of the Dissolve node will be 4K when the slider is set to full Background, but will suddenly jump to 8K when set to full Foreground or when mixed somewhere between the foreground and background.

Inspector



Dissolve controls

Controls Tab

These are the main controls that govern the Dissolve node's behavior.

- **Operation Pop-Up:** The Operation menu contains one of seven different methods for mixing the Foreground and Background inputs. The two images are mixed using the value of the Background/Foreground slider to determine the percentage each image contributes.

  Dissolve: The standard Dissolve mode is the equivalent of a cross dissolve: one clip fades out as another clip fades in.

  Additive Dissolve: Similar in look to a standard film dissolve, an Additive dissolve adds the second clip and then fades out the first one.

  Erode: The Erode method transitions between the two images by growing the darkest areas of the background image to reveal the foreground image. The effect appears similar to a filmstrip burning out.

  Random Dissolve: A randomly generated dot pattern is used to perform the mix of the images.

  Random Noise Dissolve: A moving random dot pattern is used to perform the mix of the images.

  Gradient Wipe: The dissolve is controlled by the luminance values of the image in the Gradient Map input. The edges of this dissolve can be softened. The density and the color of the border can be adjusted independently.

  SMPTE Wipe: The SMPTE wipe is similar to the basic effect wipes found on many video effects switchers. There is a horizontal wipe and a vertical wipe provided. The wipes can have soft edges and borders added. The density and the color of the border can be adjusted independently.

- **Background/Foreground Slider:** Defaults to Foreground. This control determines whether the output is the background image, the foreground image, or a mix between the two. The type of mix is determined by the Operation control. If one of the input images is not currently available, the other one will be output despite the setting of this slider.

Gradient/SMPTE Wipe Controls

The following controls appear only when Gradient Wipe or SMPTE Wipe are selected.

- **Wipe Style:** (SMPTE Wipe only) This drop-down list allows the selection of two wipe styles: Horizontal - Left to Right and Vertical - Top to Bottom. The direction of the wipes can be reversed by using the Invert Wipe checkbox.
- **Invert Wipe:** (SMPTE Wipe only) When checked, the direction of the wipe will be reversed.
- **Softness:** Use this control to soften the edge of the transition.
- **Border:** Select the Border to enable coloring of the transition's edge and to reveal the associated controls. The effect is to create a border around the transition edge.
- **Border Softness:** (Appears only when Border is turned on) The Border Softness slider controls the width and density of the border. Higher values will create a denser border, and lower values will create a thinner one.
- **Border Color:** (Appears only when Border is turned on) Use Border Color to select the color used in the border.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in both the Dissolve and Merge nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Merge [Mrg]



The Merge node

Merge Node Introduction

The Merge node combines two images based on the Alpha (opacity) channel associated with the one in front. This node takes two inputs: a background and a foreground image. The Operation mode determines which method is used to combine the foreground and background images, supporting the standard Over, In, Held Out, Atop, and XOr methods for compositing images. Meanwhile, an Apply Mode pop-up lets you use different composite modes, transfer modes, or blend modes (whichever is your preferred terminology) to combine the foreground against the background in different ways. This includes such standard modes as screen, dissolve, multiply, overlay, as well as many others.

The Merge node can perform both additive (premultiplied) and subtractive (non-premultiplied) compositing, depending on how your compositions and media are set up. However, you also have the flexibility of using the Additive/Subtractive slider to blend between additive and subtractive composite results, which has the bonus of providing solutions for problem edges in some cases.

Ordinarily, the foreground and background input connections determine the layer order of images composited with this node. However, you can also enable Z-Depth compositing if Z-channels are available in the input images. Z-merging compares the depth value of each pixel in each layer to determine which pixels should be in front and which should be behind.

Inputs

The Merge node provides three image inputs, all of which are optional:

Background: The orange background input is for the first of two images you want to composite together. You should connect the background input before connecting the foreground input. If you connect an image to the background without connecting anything to the foreground input, the Merge node will output the background image.

Foreground: The green foreground input is for the second of two images you want to composite together, which is typically a foreground subject that should be in front of the background. If you connect an image to the foreground input without connecting anything to the background input first, the Merge node won't output anything.

Effect Mask: (Optional) The effect mask input lets you mask a limited area of the output image to be merged where the mask is white (where the foreground image shows in front of the background), letting the background image show through by itself where the mask is black.

Basic Node Setup

Merge nodes are typically connected in the following way, with two input images connected to the background and foreground inputs, and the output connected to the next node in the composition. In this example, the effect mask input is not used, as this is not typical.



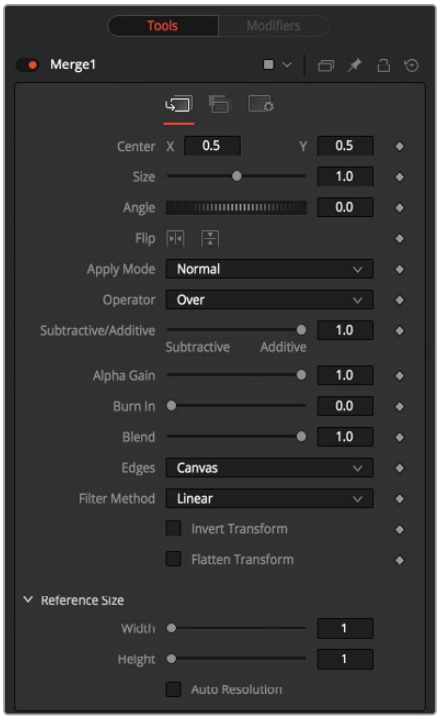
A typical Merge node structure in DaVinci Resolve

Resolution Handling

While you can connect images of any resolution to the background and foreground inputs of the Merge node, the image that's connected to the background input determines the resolution of the output.

TIP: If you want to change the resolution of the image connected to the background, you can use the Crop node to change the "canvas" resolution of the image without changing the size of the original image, or you can use the Resize node to change both the resolution and the size of the image.

Inspector



Merge node controls

Merge Tab

The Merge tab contains most of the controls necessary for customizing most merge operations.

Foreground Sizing Controls

These controls let you adjust the sizing of the image connected to the foreground input, making it unnecessary to use a separate Transform node to fit the foreground layer to match the background layer in simple compositions.

- **Center X and Y:** This control determines the position of the foreground image in the composite. The default is 0.5, 0.5, which centers the foreground image in the exact center of the background image. The value shown is always the actual position in normalized coordinates, multiplied by the reference size. See below for a description of the reference size controls.
- **Size:** Use this control to increase or decrease the size of the foreground image before it is composited over the background. The range of values for this slider is 0.0 to 5.0, but any value greater than 0 can be entered manually. A size of 1.0 gives a pixel-for-pixel composition, where a single pixel in the foreground is the same size as a single pixel in the background.
- **Angle:** Use this control to rotate the foreground image before it is combined with the background.

Compositing Mode and Adjustment Controls

The next six parameters control how the background and foreground input images are combined to create a single output image.

- **Apply Modes:** The Apply Mode setting determines the math used when blending or combining the foreground and background pixels.

  Normal: The default Normal merge mode uses the foreground's Alpha channel as a mask to determine which pixels are transparent and which are not. When this is active, another menu shows possible operations, including Over, In, Held Out, Atop, and XOr.

  Screen: Screen merges the images based on a multiplication of their color values. The Alpha channel is ignored, and layer order becomes irrelevant. The resulting color is always lighter. Screening with black leaves the color unchanged, whereas screening with white will always produce white. This effect creates a similar look to projecting several film frames onto the same surface. When this is active, another menu shows possible operations, including Over, In, Held Out, Atop, and XOr.

  Dissolve: Dissolve mixes two image sequences together. It uses a calculated average of the two images to perform the mixture.

  Darken: Darken looks at the color information in each channel and selects the background or foreground image's color value, whichever is darker, as the result color. Pixels lighter than the merged colors are replaced, and pixels darker than the merged color do not change.

  Multiply: Multiplies the values of a color channel. This will give the appearance of darkening the image as the values are scaled from 0 to 1. White has a value of 1, so the result would be the same. Gray has a value of 0.5, so the result would be a darker image, or an image half as bright.

  Color Burn: Color Burn uses the foreground's color values to darken the background image. This is similar to the photographic dark room technique of burning by increasing the exposure of an area of a print.

  Linear Burn: Linear Burn decreases the brightness of the base color based on the value of the blend color. The result is darker than Multiply but less saturated than Color Burn. Linear Burn also produces the most contrast in darker colors than any of the other blending modes in the Darker group.

Darker Color: The Darker Color blending mode is very similar to Darken. This blending mode does not blend pixels. It only compares the base and blend colors, and it keeps the darkest of the two. The difference is that Darker Color looks at the composite of all the RGB channels, whereas Darken looks at each RGB channel individually to come up with a final blend.

Lighten: Lighten looks at the color information in each channel and selects the background or foreground image's color values, whichever is lighter, as the result color value. Pixels darker than the merged color are replaced, and pixels lighter than the merged color do not change.

Color Dodge: Color Dodge uses the foreground's color values to brighten the background image. This is similar to the photographic dark room technique of dodging by reducing the exposure of an area of a print.

Linear Dodge: Linear Dodge (Add) produces similar but stronger results than Screen or Color Dodge. This blending mode looks at the color information in each channel and brightens the base color to reflect the blend color by increasing the brightness. Blending with black produces no change. Linear Dodge (Add) blends differently when Fill Opacity is adjusted, compared to when Opacity is adjusted.

Lighter Color: Lighter Color is very similar to Lighten. This blending mode does not blend pixels. It only compares the base and blend colors, and it keeps the brightest of the two. The difference is that Lighter Color looks at the composite of all the RGB channels, whereas Lighten looks at each RGB channel to come up with a final blend.

Overlay: Overlay multiplies or screens the color values of the foreground image, depending on the color values of the background image. Patterns or colors overlay the existing pixels while preserving the highlights and shadows of the color values of the background image. The background image is not replaced but is mixed with the foreground image to reflect the original lightness or darkness of the background image.

Soft Light: Soft Light darkens or lightens the foreground image, depending on the color values of the background image. The effect is similar to shining a diffused spotlight on the image.

Hard Light: Hard Light multiplies or screens the color values of the foreground image, depending on the color values of the background image. The effect is similar to shining a harsh spotlight on the image.

Vivid Light: Vivid Light is an extreme version of Overlay and Soft Light. Anything darker than 50% gray is darkened, and anything lighter than 50% gray is lightened. Vivid Light is one of those blending modes where you may want to adjust the opacity, since 100% opacity is generally too strong. Vivid Light is the fifth blending mode of eight that give you different results when you reduce the fill compared to opacity.

Linear Light: Linear Light uses a combination of the Linear Dodge blending on lighter pixels and a Linear Burn on darker pixels. Typically, the resulting colors are extreme, and you may want to use the Opacity or Fill sliders to adjust it. Linear Lights blends differently when Fill Opacity is adjusted, compared to when Opacity is adjusted.

Pin Light: Pin Light is an extreme blending mode that performs a Darken and Lighten blending mode simultaneously. It can result in patches or blotches, and it completely removes all mid-tones.

Difference: Difference looks at the color information in each channel and subtracts the foreground color values from the background color values or the background from the foreground, depending on which has the greater brightness value. Merging with white inverts the color. Merging with black produces no change.

Exclusion: Exclusion creates an effect similar to, but lower in contrast than, the Difference mode. Merging with white inverts the base color values. Merging with black produces no change.

Hue: Hue creates a result color with the luminance and saturation of the background color values and the hue of the foreground color values.

Saturation: Saturation creates a result color with the luminance and hue of the base color and the saturation of the blend color.

Color: Color creates a result color with the luminance of the background color value and the hue and saturation of the foreground. This preserves the gray levels in the image and is useful for coloring monochrome images.

Luminosity: Luminosity creates a result color with the hue and saturation of the background color values and the luminance of the foreground color values. This mode creates an inverse effect from that of the Color mode.

Hypotenuse: This blend mode is good for HDR images that have out of range colors above 1, It uses a square root of each color squared and added to blend the color.

```
Out = sqrt (Fc*Fc +Bc*Bc)
```

Geometric: This blend mode is good for HDR images that have out of range colors above 1, For values above zero the result is 2 times the foreground times the background color divided by the foreground plus background color.

```
Out = 2*Fc*Bc / (Fc+Bc)
```

- **Operator Modes:** This menu is used to select the Operation mode of the merge. Changing the Operation mode changes how the foreground and background are combined to produce a result. This pop-up menu is visible only when the Merge node's Apply mode is set to either Normal or Screen.

  For an excellent description of the math underlying the Operation modes, read Compositing Digital Images, Porter, Thomas, and T. Duff, ACM SIGGRAPH Computer Graphics proceedings, 1984, pages 253-259. Essentially, the math is as described below. Note that some modes not listed in the Operator drop-down menu (Under, In, Held In, Below) are easily obtained by swapping the foreground and background inputs (with Command-T or Ctrl-T) and choosing a corresponding mode. The formula used to combine pixels in the merge is always (fg * x) + (bg * y). The different operations determine exactly what x and y are, as shown in the description for each mode.

  The Operator modes are as follows:

  Over: The Over mode adds the foreground layer to the background layer by replacing the pixels in the background with the pixels from the Z wherever the foreground's Alpha channel is greater than 1.

  ```
  x = 1, y = 1-[foreground Alpha]
  ```

  In: The In mode multiplies the Alpha channel of the background input against the pixels in the foreground. The color channels of the foreground input are ignored. Only pixels from the foreground are seen in the final output. This essentially clips the foreground using the mask from the background.

  ```
  x = [background Alpha], y = 0
  ```

  Held Out: Held Out is essentially the opposite of the In operation. The pixels in the foreground image are multiplied against the inverted Alpha channel of the background image. Accomplish exactly the same result using the In operation and a Matte Control node to invert the matte channel of the background image.

  ```
  x = 1-[background Alpha], y = 0
  ```

  Atop: Atop places the foreground over the background only where the background has a matte.

```
x = [background Alpha], y = 1–[foreground Alpha]
```

XOr: XOr combines the foreground with the background wherever either the foreground or the background has a matte, but never where both have a matte.

```
x = 1–[background Alpha], y = 1–[foreground Alpha]
```

Conjoint: The Conjoint mode will make a decision based combination of Alpha channels of the foreground and background images; this is helpful in soft edge and motion blurred Alpha where Alpha is not solid.

```
X= 1, Y= X+Y(1–af)/ab, if af>ab
```

Disjoint: The Disjoint mode will make a decision based combination of Alpha channels of the foreground and background images; this is helpful in combining layers as to not get out of range Alpha, and premultiplied edges get the correct Alpha combination.

```
X= X+Y(1–af)/ab, Y= X+Y if af+ab<1
```

Mask: The Mask mode will output the background image multiplied by the foreground Alpha.

```
X = X * af, Y =0
```

Stencil: The Stencil mode will output the background image multiplied by the inverse foreground Alpha.

```
X = X * (1–af), Y =0
```

Under: The Under mode is the same operation as the Over mode but will swap foreground and background images in the operations.

```
X = Y, Y =X *(1–af)
```

- **Subtractive/Additive slider:** This slider controls whether Fusion performs an Additive merge, a Subtractive merge, or a blend of both. This slider defaults to Additive merging for most operations, assuming the input images are premultiplied (which is usually the case). If you don't understand the difference between Additive and Subtractive merging, here's a quick explanation.
  - An Additive merge is necessary when the foreground image is premultiplied, meaning that the pixels in the color channels have been multiplied by the pixels in the Alpha channel. The result is that transparent pixels are always black, since any number multiplied by 0 always equals 0. This obscures the background (by multiplying with the inverse of the foreground Alpha), and then simply adds the pixels from the foreground.
  - A Subtractive merge is necessary if the foreground image is not pre-multiplied. The compositing method is similar to an additive merge, but the foreground image is first multiplied by its own Alpha to eliminate any background pixels outside the Alpha area.

  While the Additive/Subtractive option could easily have been a checkbox to select one mode or another, the Merge node lets you blend between the Additive and Subtractive versions of the merge operation—an operation that is occasionally useful for dealing with problem composites with edges that are calling attention to themselves as too bright or too dark.

  For example, using Subtractive merging on a premultiplied image may lead to darker edges, whereas using Additive merging with a non-premultiplied image will cause any non-black area outside the foreground's Alpha to be added to the result, thereby lightening the edges. By blending between Additive and Subtractive, you can tweak the edge brightness to be just right for your situation.

- **Alpha Gain slider:** Alpha Gain linearly scales the values of the foreground's Alpha channel. In Subtractive merges, this controls the density of the composite, similarly to Blend. In Additive merges, this effectively reduces the amount that the background is obscured, thus brightening the overall result. In an Additive merge with Alpha Gain set to 0.0, the foreground pixels are simply added to the background.
- **Burn In slider:** The Burn In control adjusts the amount of Alpha used to darken the background, without affecting the amount of foreground added in. At 0.0, the merge behaves like a straight Alpha blend, whereas at 1.0, the foreground is effectively added onto the background (after Alpha multiplication if in Subtractive mode). This gives the effect of the foreground image brightening the background image, as with Alpha Gain. For Additive merges, increasing the Burn In gives an identical result to decreasing Alpha Gain.
- **Blend slider:** This is a cloned instance of the Blend slider in the Common Controls tab. Changes made to this control are simultaneously made to the one in the common controls. The Blend slider mixes the result of the node with its input, blending back the effect at any value less than 1.0. In this case, it will blend the background with the merged result.

Additional Controls

The remaining controls let you fine-tune the results of the above settings.

- **Filter Method:** For input images that are being resized, this setting lets you choose the filter method used to interpolate image pixels when resizing clips. The default setting is Linear. Different settings work better for different kinds of resizing. Most of these filters are useful only when making an image larger. When shrinking images, it is common to use the Linear filter; however, the Catmull-Rom filter will apply some sharpening to the results and may be useful for preserving detail when scaling down an image.

  Nearest Neighbor: This skips or duplicates pixels as needed. This produces the fastest but crudest results.

  Box: This is a simple interpolation resize of the image.

  Linear: This uses a simplistic filter, which produces relatively clean and fast results.

  Quadratic: This filter produces a nominal result. It offers a good compromise between speed and quality.

  Cubic: This produces better results with continuous-tone images. If the images have fine detail in them, the results may be blurrier than desired.

  Catmull-Rom: This produces good results with continuous-tone images that are resized down. Produces sharp results with finely detailed images.

  Gaussian: This is very similar in speed and quality to Bi-Cubic.

  Mitchell: This is similar to Catmull-Rom but produces better results with finely detailed images. It is slower than Catmull-Rom.

  Lanczos: This is very similar to Mitchell and Catmull-Rom but is a little cleaner and also slower.

  Sinc: This is an advanced filter that produces very sharp, detailed results; however, it may produce visible "ringing" in some situations.

  Bessel: This is similar to the Sinc filter but may be slightly faster.

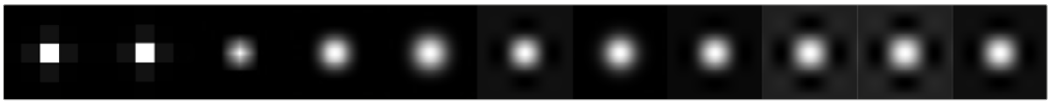
  Window Method: Some filters, such as Sinc and Bessel, require an infinite number of pixels to calculate exactly. To speed up this operation, a windowing function is used to approximate the filter and limit the number of pixels required. This control appears when a filter that requires windowing is selected.

  Hanning: This is a simple tapered window.

  Hamming: Hamming is a slightly tweaked version of Hanning.

Blackman: A window with a more sharply tapered falloff.

Kaiser: A more complex window, with results between Hamming and Blackman.



Resize Filters from left to right: Nearest Neighbor, Box, Linear, Quadratic, Cubic, Catmull-Rom, Gaussian, Mitchell, Lanczos, Sinc, and Bessel

- **Edges Buttons:** Four buttons let you choose how to handle the space around images that are smaller than the current DoD of the canvas as defined by the resolution of the background image.

  Canvas: The area outside the frame is set to the current color/opacity of the canvas. If you want to change this value, you can attach a Set Canvas Color node between the image connected to the foreground input and the foreground input itself, using Set Canvas Color to choose a color and/or transparency setting with which to fill the canvas.

  Wrap: Creates a "video wall" effect by duplicating the foreground image as a grid.

  Duplicate: Duplicates the outermost pixels along the edge of the foreground image, duplicating them to stretch up, down, left, and right from each side to reach the end of the DoD.

  Mirror: Similar to duplicate, except every other iteration of the foreground image is flipped and flopped to create a repeating pattern.

- **Invert Transform:** Select the Invert Transform control to invert any position, rotation, or scaling transformation. This option is useful when connecting the merge to the position of a tracker for match moving.
- **Flatten Transform:** The Flatten Transform option prevents this node from concatenating its transformation with subsequent nodes. The node may still concatenate transforms from its input, but it will not concatenate its transformation with the node at its output.
- **Reference Size:** The controls under Reference Size do not directly affect the image. Instead, they allow you to control how Fusion represents the position of the Merge node's center.

  Normally, coordinates are represented as values between 0 and 1, where 1 is a distance equal to the full width or height of the image. This allows resolution independence, because the size of the image can be changed without having to change the value of the center.

  One disadvantage to this approach is that it complicates making pixel-accurate adjustments to an image. To demonstrate, imagine an image that is 100 x 100 pixels in size. To move the center of the foreground element to the right by 5 pixels, we would change the X value of the merge center from 0.5, 0.5 to 0.55, 0.5. We know the change must be 0.05 because 5/100 = 0.05.

  If you specify the dimensions of the background image in the Reference Size controls, this changes the way the center control values are displayed so that it shows the actual pixel positions in its X and Y fields.

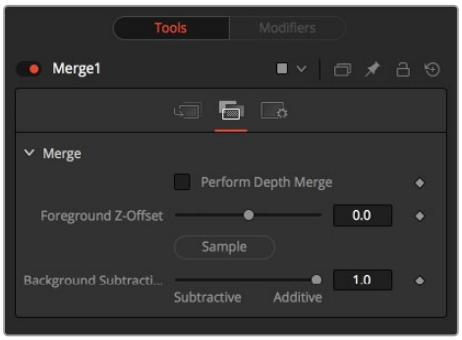
  Extending the example, set the width and height to 100 each and the center will now be shown as 50, 50, and we would move it 5 pixels toward the right by entering 55, 50.

  Internally, the Merge node still stores this value as a number between 0 to 1 and, if the center control's value was to be queried via scripting or the center control was to be published for use by other nodes, the original normalized value would be retrieved. The change is only visible in the value shown for merge center in the node control.

  Use Frame Format Settings: Select this to force the merge to use the composition's current frame format settings to set the reference width and reference height values.

  Width and Height: Set these sliders to the width and height of the image to change the way that Fusion displays the values of the Merge node's center control.

Channels Tab



Merge node Channels tab

The Channels tab has controls that let the Merge node use Z-channels embedded within each image to define what's in front and what's behind during a Merge operation. The following controls let you customize the result.

- **Perform Depth Merge:** Off by default. When turned on, the Z-channel of both images will be used to determine the composite order. Alpha channels are still used to define transparency, but the values of the Z-Depth channels will determine the ordering of image elements, front to back. If a Z-channel is not available for either image, the setting of this checkbox will be ignored, and no depth compositing will take place. If Z-Depth channels are available, turning this checkbox off disables their use within this operation.
- **Foreground Z-Offset:** This slider sets an offset applied to the foreground image's Z value. Click the Pick button to pick a value from a displayed image's Z-channel, or enter a value using the slider or input boxes. Raising the value causes the foreground image's Z-channel to be offset further away along the Z-axis, whereas lowering the value causes the foreground to move closer.
- **Subtractive/Additive:** When Z-compositing, it is possible for image pixels from the background to be composited in the foreground of the output because the Z-buffer for that pixel is closer than the Z of the foreground pixel. This slider controls whether these pixels are merged in an Additive or a Subtractive mode, in exactly the same way as the comparable slider in the Merge tab.

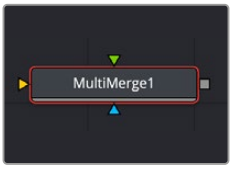
  When merged over a background of a different color, the original background will still be visible in the semitransparent areas. An Additive merge will maintain the transparencies of the image but will add their values to the background.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in both the Dissolve and Merge nodes. These common controls are described in the following "The Common Controls" section.

MultiMerge [MMrg]



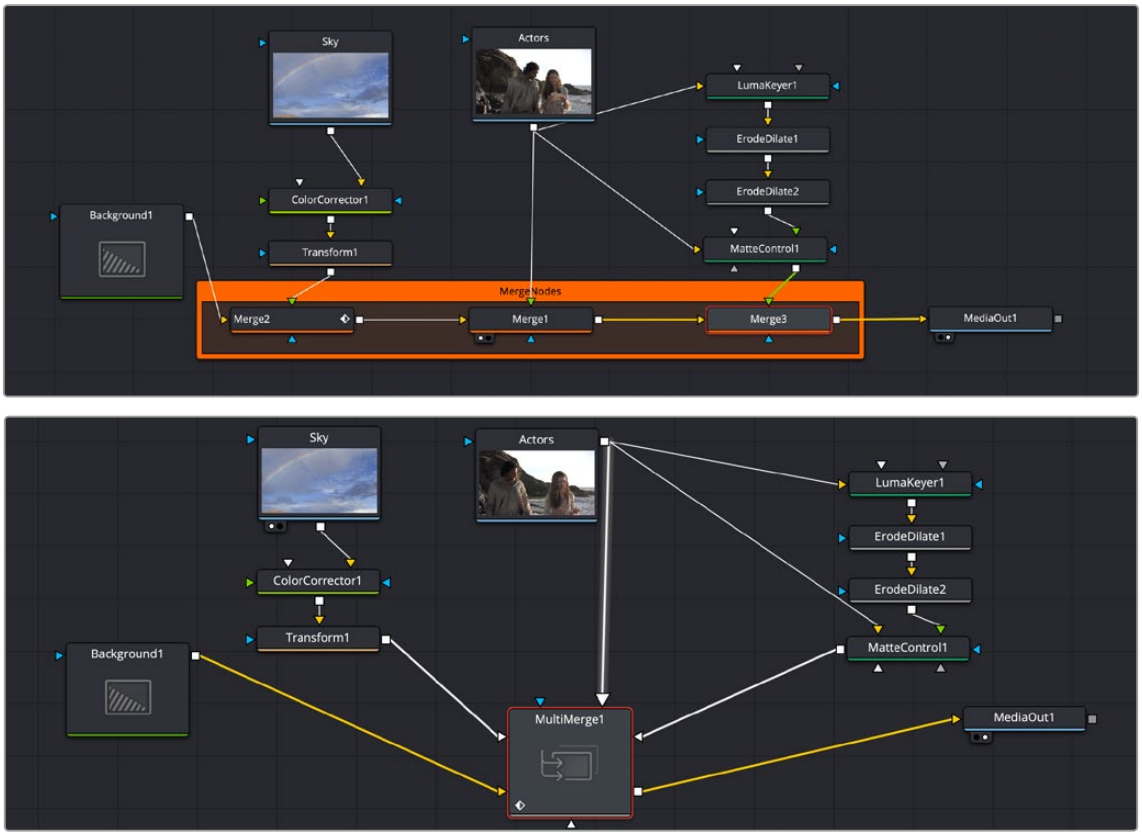
The MultiMerge node

MultiMerge Node Introduction

The MultiMerge node allows you to combine several sources together in a layer-based structure. Unlike the standard Merge tool, which only accepts one foreground input, MultiMerge creates a new consecutive foreground input each time a new source is dragged to the tool.

These inputs are added to the Layer List in the Inspector and are displayed in a hierarchy from top to bottom as a stack. The layers are arranged so the layer closest to the foreground is at the top of the list, and the one closest to the background is at the bottom. Each layer has its own separate and independent Merge controls (see the Merge node above) that appear in the lower half of the Inspector, allowing you to individually adjust the position, size, composite modes, etc., for each source input.

As your composite increases in complexity, you can end up with a large amount of standard Merge nodes scattered throughout the node tree. MultiMerge can be used to combine these Merge nodes together, both as an organizational tool and allowing you to control the order of operations in a layer-based, rather than a node-based environment. You could conceivably composite and organize hundreds of separate layers using one MultiMerge node on your node tree.



A sky replacement composite requiring three separate Merge nodes (previous page) is replaced by one single Multi Merge node (above). The pipe of the currently selected layer in the Layer List has a slight glow applied (Actors node to MultiMerge node).

Inputs

The MultiMerge node provides these image inputs:

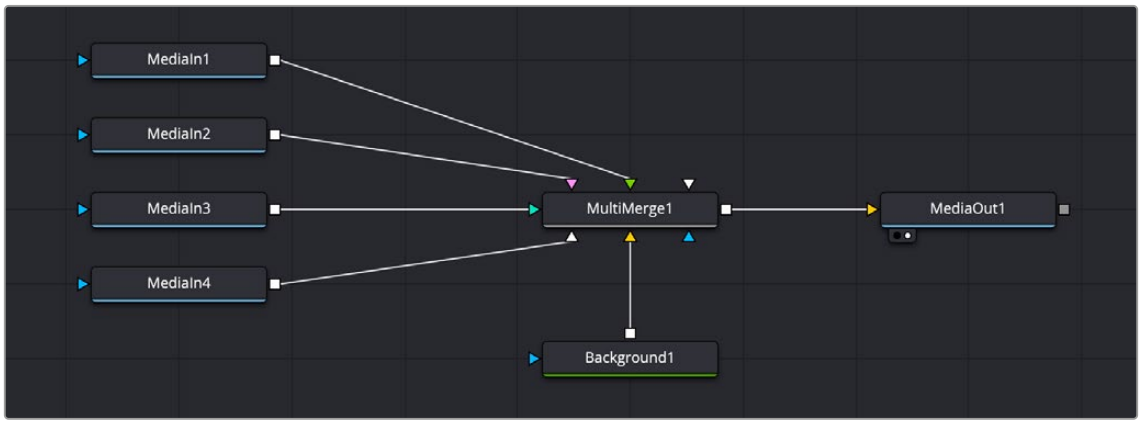
Background: The orange background input is for the first of the images you want to composite together. You should connect the background input before connecting the foreground input.

Foreground: The foreground inputs are for each subsequent image you want to composite together, which is typically a foreground subject that should be in front of the background. Connecting a pipe from a new source onto the MultiMerge node automatically creates a new foreground input on the node and a new layer in the Layer List. Foreground inputs are all colored white.

Effect Mask: (Optional) The effect mask input lets you mask a limited area of the output image to be merged where the mask is white (where the foreground image shows in front of the background), letting the background image show through by itself where the mask is black.

Basic Node Setup

MultiMerge nodes are typically connected in the following way, with multiple input images connected to the foreground inputs, a background input, and the output connected to the next node in the composition. In this example, the effect mask input is not used, as this is not typical.

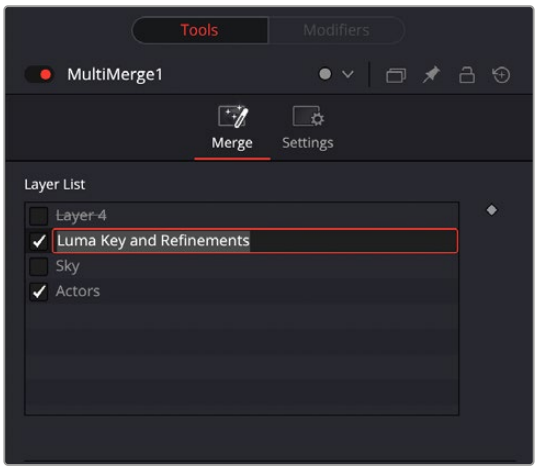


A typical MultiMerge node structure in DaVinci Resolve

Resolution Handling

While you can connect images of any resolution to the background and foreground inputs of the MultiMerge node, the image that's connected to the background input determines the resolution of the output.

Inspector



The Layer List of the MultiMerge node. The input to Layer 4 has been disconnected (strike through the name) and disabled (unchecked). The Luma Key layer is having its name changed, the Sky layer is disabled (unchecked), while the Actors layer is enabled (checked).

Layer List

For each foreground input connected to the MultiMerge node, a new layer is created sequentially in the Layer List. The Layer List is hierarchically sorted from top to bottom, with layers on top being above any layers below them in the image. You can customize this layer structure in several ways.

To Select a Layer:

Click on the layer in the Layer List to make that layer active. The active layer's Merge controls will appear at the bottom of the Inspector, and the pipe connecting to the layer's input on the MultiMerge node will glow slightly in the node tree.

You can multi-select layers by holding the Shift key down while clicking to select a range of layers, or the Command key to select multiple non-adjacent layers.

Right-clicking on any layer opens up a contextual menu with the following options:

- **Go To Source:** Selects the tool that is connected to the layer's input on the MultiMerge node and opens the tool's parameters in the Inspector.
- **Split Here:** Automatically creates another connected MultiMerge node and splits the existing layers between the two. The layer that was selected for the split, and all layers above it, will be spun off into a new MultiMerge node. The original MultiMerge node will be connected to the background input of the new MultiMerge node.
- **Rename Layer X:** Opens up a dialog box allowing you to rename the selected layer.
- **Rename Layer to Match Source:** Automatically changes the name of the layer to the name of the tool it's connected to.
- **Set Layer X to Default:** This resets the layer's merge settings below, back to the default.
- **Delete Layer:** This deletes the layer entirely from the Layer List and disconnects the pipe from the original tool. It does not delete the original tool.

To Move a Layer up or down the List:

Click and drag on the Layer name and move it up or down the list. A grey line will highlight where in the list the layer will now be placed. Let go of the mouse button to make the change.

To Rename a Layer:

By default each layer is named based in the input tool's name. To change the Layer's name to something more meaningful, double-click on the Layer's name and type a new name in the text field. You can also right-click on a Layer's name and select Rename Layer from the contextual menu.

To Disable/Enable a Layer:

Checking or unchecking the box next to the Layer's name will enable or disable that layer respectively. Disabling the layer turns off that input's contribution to the overall composite in the MultiMerge but does not delete it.

To Replace a Layer:

By design, if you disconnect the pipe from a the MultiMerge's input, the layer will still remain in its place in the Layer List but with a strike-through its name to let you know that nothing is currently connected. This allows you to quickly iterate and audition several input sources (clips, graphics, etc.) to the same layer without having to constantly rearrange the order of your Layer List.

Keyframing the Layer List:

The Layer List itself is keyframable, using the standard keyframe controls to the right of the list. Parameters that can be keyframed are layer order, and layer enable and disable.

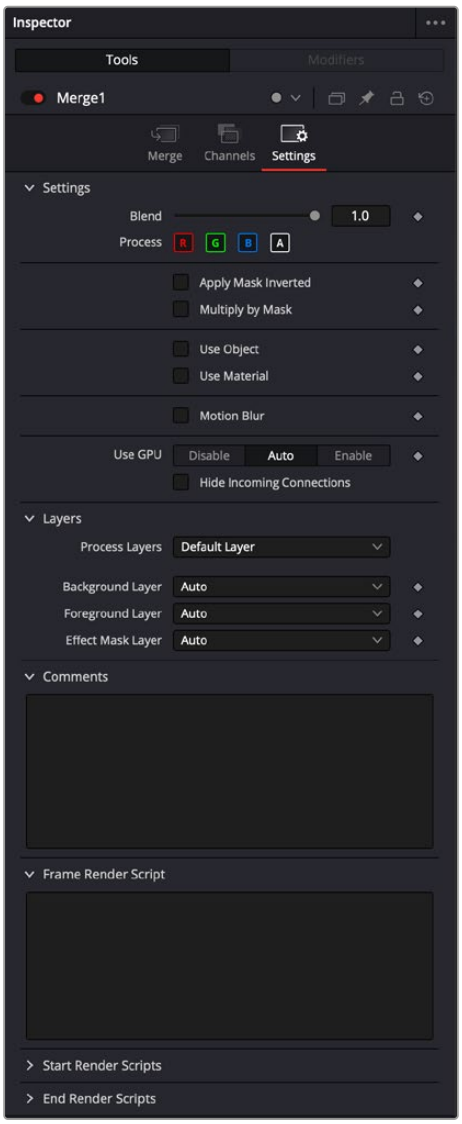
Merge Controls

Each layer has a separate and independent set of Merge controls that appear here when the layer is selected. For detailed information on how the Merge controls work, see the previous section in this chapter titled *Merge [Mrg]*.

The Common Controls

The Merge, MultiMerge, and Dissolve nodes share several identical controls in the Inspector. This section describes controls that are common among those two nodes.

Inspector



Merge Setting Inspector

Settings Tab

The Settings tab in the Inspector can be found on both tools in the Composite category. The Settings controls are even found on third-party color type plugin tools. The controls are consistent and work the same way for each tool, although some tools do include one or two individual options that are also covered here.

Blend

The Blend control is used to blend between the tool's original image input and the tool's final modified output image. When the blend value is 0.0, the outgoing image is identical to the incoming image. Normally, this will cause the tool to skip processing entirely, copying the input straight to the output.

Process When Blend Is 0.0

The tool is processed even when the input value is zero. This can be useful if processing of this node is scripted to trigger another task, but the value of the node is set to 0.0.

Red/Green/Blue/Alpha Channel Selector

These four buttons are used to limit the effect of the tool to specified color channels. This filter is often applied after the tool has been processed.

For example, if the red button on a blur tool is deselected, the blur will first be applied to the image, then the red channel from the original input will be copied back over the red channel of the result.

There are some exceptions, such as tools for which deselecting these channels causes the tool to skip processing that channel entirely. Tools that do this will generally possess a set of identical RGBA buttons on the Controls tab in the tool. In this case, the buttons in the Settings and the Controls tabs are identical.

Apply Mask Inverted

Enabling the Apply Mask Inverted option inverts the complete mask channel for the tool. The mask channel is the combined result of all masks connected to or generated in a node.

Multiply by Mask

Selecting this option will cause the RGB values of the masked image to be multiplied by the mask channel's values. This will cause all pixels of the image not in the mask (i.e., set to 0) to become black/transparent.

Use Object/Use Material (Checkboxes)

Some 3D software can render to file formats that support additional channels. Notably, the EXR file format supports Object and Material ID channels, which can be used as a mask for the effect. These checkboxes determine whether the channels will be used, if present. The specific Material ID or Object ID affected is chosen using the next set of controls.

Correct Edges

This checkbox appears only when the Use Object or Use Material checkboxes are selected. It toggles the method used to deal with overlapping edges of objects in a multi-object image. When enabled, the Coverage and Background Color channels are used to separate and improve the effect around the edge of the object. If this option disabled (or no Coverage or Background Color channels are available), aliasing may occur on the edge of the mask.

> For more information on the Coverage and Background Color channels, *see Chapter 18, "Understanding Image Channels,"* in the Fusion Reference Manual.

Object ID/Material ID (Sliders)

Use these sliders to select which ID will be used to create a mask from the object or material channels of an image. Use the Sample button in the same way as the Color Picker: to grab IDs from the image displayed in the view. The image or sequence must have been rendered from a 3D software package with those channels included.

Use GPU

The Use GPU menu has three settings. Setting the menu to Disable turns off hardware-accelerated rendering using the graphics card in your computer. Enabled uses the hardware. Auto uses a capable GPU if one is available and falls back to software rendering when a capable GPU is not available

Motion Blur

- **Motion Blur:** This toggles the rendering of Motion Blur on the tool. When this control is toggled on, the tool's predicted motion is used to produce the motion blur caused by the virtual camera's shutter. When the control is toggled off, no motion blur is created.
- **Quality:** Quality determines the number of samples used to create the blur. A quality setting of 2 will cause Fusion to create two samples to either side of an object's actual motion. Larger values produce smoother results but increase the render time.
- **Shutter Angle:** Shutter Angle controls the angle of the virtual shutter used to produce the motion blur effect. Larger angles create more blur but increase the render times. A value of 360 is the equivalent of having the shutter open for one whole frame exposure. Higher values are possible and can be used to create interesting effects.
- **Center Bias:** Center Bias modifies the position of the center of the motion blur. This allows the creation of motion trail effects.
- **Sample Spread:** Adjusting this control modifies the weighting given to each sample. This affects the brightness of the samples.

Layers

- **Process Layers:** Select the layer to process for effects from:
  - Default Layer (default) - To apply the effect on the base/unnamed layer, not affecting other layers.
  - All Layers - To apply the effect on all available layers.
  - Custom - Choice of other available layers to apply the effect on, not affecting other layers.
- **Input Layer:** The Input Layer lists all layers available from the upstream node of the default input. The selected layer can be used by the current tool or farther down the node tree in isolation.
- **Effect Mask Layer:** If the tool has an Effect Mask input, this control lists the available layers that can provide mattes for the tool.

Comments

The Comments field is used to add notes to a tool. Click in the field and type the text. When a note is added to a tool, a small red square appears in the lower left corner of the node when the full tile is displayed or a small text bubble icon appears on the right when nodes are collapsed. To see the note in the node editor, hold the mouse pointer over the node to display the tooltip.

Scripts

Three Scripting fields are available on every tool in Fusion from the Settings tab. They each contain edit boxes used to add scripts that process when the tool is rendering. For more details on scripting nodes, please consult the Fusion scripting documentation.

Chapter 36

Deep Image Nodes

This chapter describes Fusion's Deep Image Nodes for compositing objects in 3D space.

The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Deep Image Compositing Toolset (Studio Version Only) 1013
dCrop [dCr] 1013
Deep to Image [DtI] 1015
Deep to Points [DtP] 1017
dHoldout [dHld] 1021
dMerge [dMg] 1023
dRecolor [dRc] 1024
dResize [dRz] 1025
dTransform [dXf] 1026
Image to Deep [ItD] 1028
The Common Controls 1029

Deep Image Compositing Toolset (Studio Version Only)

Fusion supports a subset of tools to load, utilize, and save deep image data. Deep image data, particularly in formats like OpenEXR, stores multiple samples per pixel to represent data for objects at different depths in the same pixel. The Fusion Deep Image toolset can utilize this valuable information to layer complex 3D renders, particularly where multiple objects overlap, and integrate them with rendered 3D scenes.

The tools can be found in the Deep Image subcategory with a d-prefix.

Additionally:

- Renderer3D nodes can output deep image data.
- Loader and the MediaIn node can import deep image data EXRs.
- Saver can export deep image EXRs.
- All deep compositing tools (outside of Image to Deep and Deep to Image) can automatically convert 2D image inputs to a deep image.

NOTE: Deep compositing requires a linear colorspace. Once merged, users can convert the results into their desired colorspace.

dCrop [dCr]



The dCrop node

dCrop Node Introduction

This tool crops Deep Image samples using specified min and max depth values. It also crops pixels with an input mask.

Inputs

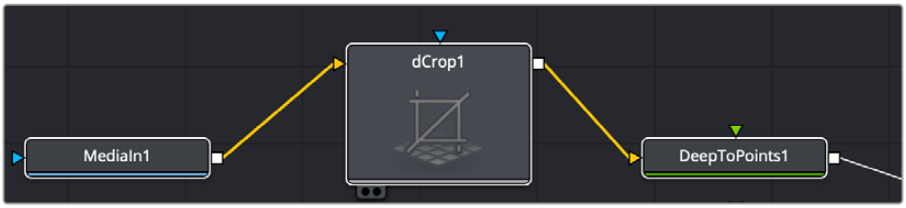
The dCrop tool uses one input.

Input: This orange input is the only required connection. It is connected from a node with deep image data.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools.

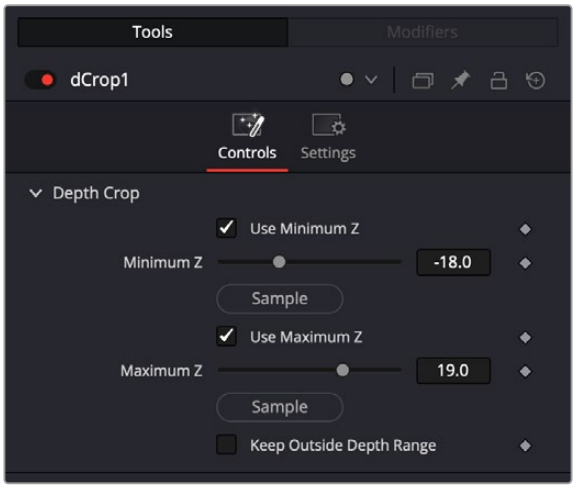
Basic Node Setup

The dCrop node receives the media in containing the deep image data. It then lets you refine the depth boundaries with which to pass to the next tool.



dCrop accepts deep image data

Inspector



dCrop controls

Controls Tab

The Controls tab includes parameters for adjusting the amount of depth information to include or exclude from being passed on the the rest of the node tree.

Depth Crop

Lets you crop the Z parameters of the deep image data.

- **Use Minimum Z:** Check this box to use a Minimum Z depth for the crop.
- **Use Minimum Z:** Adjust the Minimum Z amount using this slider.
- **Sample:** Click and drag from the sample button into the Viewer to choose a value.
- **Use Maximum Z:** Check this box to use a Maximum Z depth for the crop.
- **Keep Outside Depth Range:** Use this control to invert the cropped selection. After setting a minimum and/or maximum value in the sliders above, this control flips the selection so that the cropped region is now visible.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Deep Image nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Deep to Image [Dtl]



The Deep to Image node

Deep to Image Node Introduction

This tool converts deep data into a 2D image by flattening its depth samples. This outputs the deep composite ready for traditional 2D compositing.

Inputs

The Deep to Image tool uses one input.

Input: This orange input is the only required connection. It is connected from a node with deep image data.

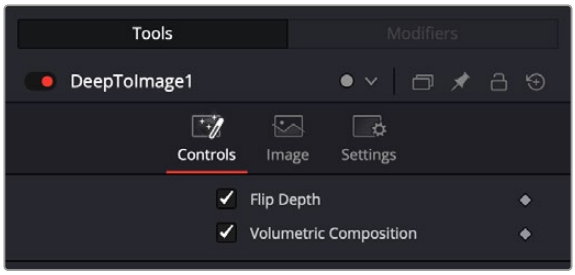
Basic Node Setup

The Deep to Image node receives the media in containing the deep image data. It then flattens the image for further use in normal 2D compositing.



Deep to Image accepts deep image data, flattens it, then sends the 2D image down the node tree.

Inspector



Deep to Image controls

Controls Tab

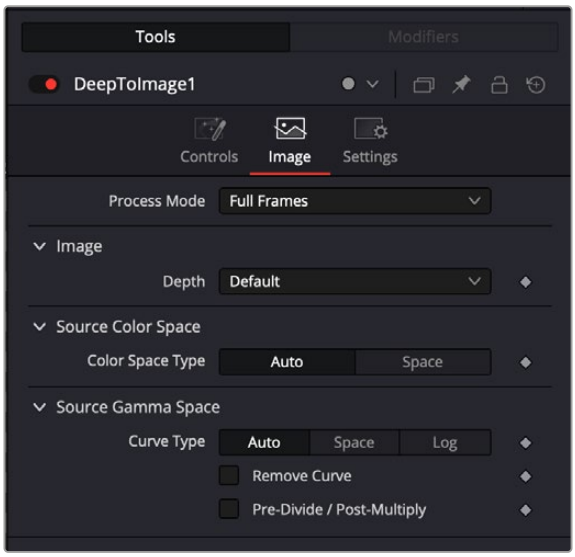
The Controls tab includes parameters for adjusting the depth information.

- **Flip Depth:** Inverts the depth information.

- **Volumetric Composition:** The Volumetric Composition option blends overlapping deep samples as semi-transparent volumes rather than flat surfaces, simulating accurate interaction through elements such as fog or VDBs. This results in a more realistic compositing of semi-transparent elements.

Image Tab

The controls in this tab are used to set the resolution, color depth, and pixel aspect of the rendered image produced by the node.



The Deep to Image Image tab

Process Mode

Use this menu control to select the Fields Processing mode used by Fusion to render changes to the image. The default option is determined by the Has Fields checkbox control in the Frame Format preferences.

Depth

The Depth menu is used to set the pixel color depth of the particles. 32-bit pixels require 4X the memory of 8-bit pixels but have far greater color accuracy. Float pixels allow high dynamic range values outside the normal 0...1 range for representing colors that are brighter than white or darker than black.

Source Color Space

You can use the Source Color Space menu to set the Color Space of the footage to help achieve a linear workflow. Unlike the Gamut tool, this doesn‘t perform any actual color space conversion, but rather adds the source space data into the metadata, if that metadata doesn‘t exist. The metadata can then be used downstream by a Gamut tool with the From Image option, or in a Saver, if explicit output spaces are defined there. There are two options to choose from:

- **Auto:** Automatically reads and passes on the metadata that may be in the image.
- **Space:** Displays a Color Space Type menu where you can choose the correct color space of the image.

Source Gamma Space

Using the Curve type menu, you can set the Gamma Space of the footage and choose to remove it by way of the Remove Curve checkbox when working in a linear workflow. There are three choices in the Curve type menu:

- **Auto:** Automatically reads and passes on the metadata that may be in the image.
- **Space:** Displays a Gamma Space Type menu where you can choose the correct gamma curve of the image.
- **Log:** Brings up the Log/Lin settings, similar to the Cineon tool.

For more information, *see Chapter 39, "Film Nodes,"* in the Fusion Reference Manual.

Remove Curve

Depending on the selected Gamma Space or on the Gamma Space found in Auto mode, the Gamma Curve is removed from, or a log-lin conversion is performed on, the material, effectively converting it to a linear output space.

Pre-Divide/Post-Multiply

Lets you convert "straight" Alpha channels into pre-multiplied Alpha channels, when necessary.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Deep Image nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Deep to Points [DtP]



The Deep to Points node

Deep to Points Node Introduction

This tool converts Deep images into a 3D point cloud. This tool is useful for visualizing depth samples and offers a way to see elements in Z when incorporating elements.

Inputs

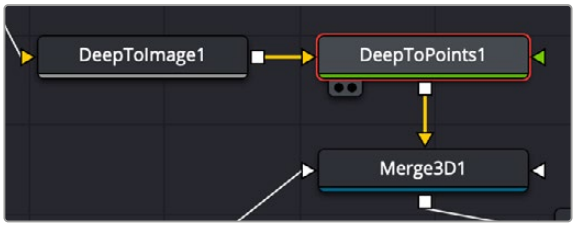
The Deep to Points tool uses one input.

Input: This orange input is the only required connection. It is connected from a node with deep image data.

Camera: Connecting a scene's camera to the green camera input precisely positions the point cloud relative to the original 3D scene. The camera's perspective is taken into account so that points are accurately placed within the 3D Viewer.

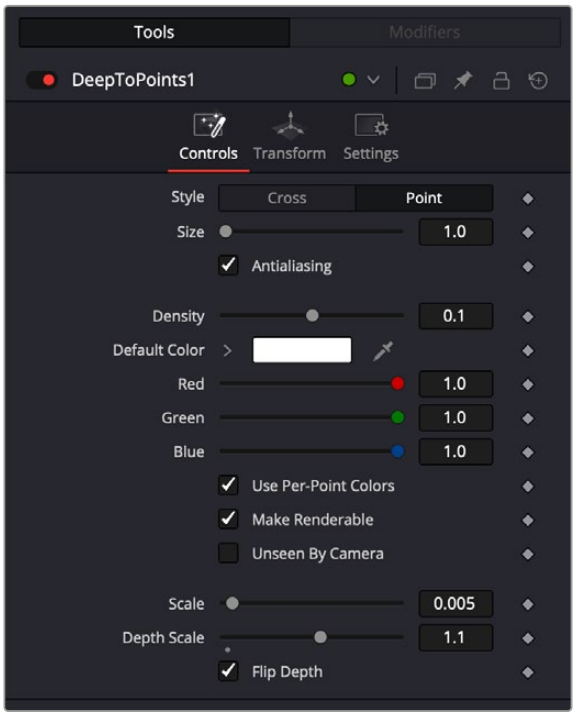
Basic Node Setup

The Deep to Points node receives the media in containing the deep image data. It then creates a 3D point cloud from the image to make it easier to visualize elements in Z space.



Deep to Points accepts deep image data and creates a Point Cloud before passing that visualization off to the rest of the node tree.

Inspector



Deep to Points controls

Controls Tab

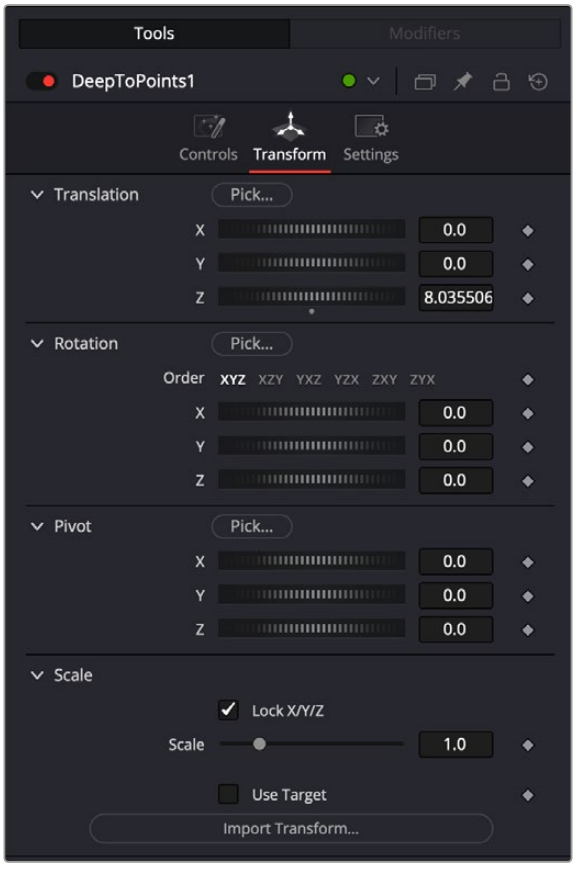
The Controls tab includes parameters for adjusting the point cloud.

- **Style:** Style allows you to display the points as cross hairs or points in the viewer.
- **Size:** This slider can be used to increase the size of the onscreen crosshairs or points
- **Antialiasing:** When enabled, points will appear rounded and smooth, resulting in softer edges.

- **Density:** This slider allows you to control the overall concentration of points in the viewer. Increasing the value will build up the intensity of points whereas decreasing this value will have the opposite effect.
- **Default Color:** When unchecking the Use Per-Point Colors checkbox, each point will inherit a color specified in this control.
- **Use Per-Point Colors:** When enabled, each point will inherit its color from the connected source clip.
- **Make Renderable:** Determines whether the point cloud visible in the 3D viewer will be renderable in a Renderer3D node. Enabling this checkbox will allow the point cloud to be rendered out into a flat 2D image.
- **Unseen by Camera:** This checkbox control appears when the Make Renderable option is selected. If the Unseen by Cameras checkbox is selected, the point cloud is visible in the viewers but not rendered into the output image by the Renderer 3D node.
- **Scale:** This slider allows you to increase or decrease the overall size of the point cloud.
- **Depth Scale:** The slider allows you to expand or contract the overall depth of the point cloud. Increasing the depth scale will amplify the overall distance of depth; spacing out the points giving you a fuller representation of the point cloud.
- **Flip Depth:** Use this checkbox to invert the depth data, flipping the position of the points from negative to positive

Transform Tab

The controls in this tab are used to set the position and scaling parameters of the point cloud.



The Deep to Points Transform tab

Translation

X, Y, Z Offset

These controls can be used to position the point cloud.

Rotation

Rotation Order

Use these buttons to select which order is used to apply rotation along each axis of the object. For example, XYZ would apply the rotation to the X axis first, followed by the Y axis and then finally the Z axis.

X, Y, Z Rotation

Use these controls to rotate the object around its pivot point. If the Use Target checkbox is selected, then the rotation is relative to the position of the target; otherwise, the global axis is used.

Pivot

X, Y, Z Pivot

A Pivot point is the point around which an object rotates. Normally, an object rotates around its own center, which is considered to be a pivot of 0,0,0. These controls can be used to offset the pivot from the center.

Scale

X, Y, Z Scale

If the Lock X/Y/Z checkbox is checked, a single Scale slider is shown. This adjusts the overall size of the object. If the Lock checkbox is unchecked, individual X, Y, and Z sliders are displayed to allow individual scaling in each dimension. Note: If the Lock checkbox is checked, scaling of individual dimensions is not possible, even when dragging specific axes of the Transformation Widget in scale mode.

Use Target

Selecting the Use Target checkbox enables a set of controls for positioning an XYZ target. When target is enabled, the object always rotates to face the target. The rotation of the object becomes relative to the target.

Import Transform

Opens a file browser where you can select a scene file saved or exported by your 3D application. It supports the following file types:

LightWave Scene	`.lws`
Max Scene	`.ase`
Maya Ascii Scene	`.ma`

The Import Transform button imports only transformation data. For 3D geometry, lights, and cameras, consider using the File > FBX Import option.

Onscreen Transformation Controls

Most of the controls in the Transform tab are represented in the viewer with onscreen controls for transformation, rotation, and scaling. To change the mode of the onscreen controls, select one of the three buttons in the toolbar in the upper left of the viewer. The modes can also be toggled using the keyboard shortcut Q for translation, W for rotation, and E for scaling. In all three modes, individual axes of the control may be dragged to affect just that axis, or the center of the control may be dragged to affect all three axes.

The scale sliders for most 3D tools default to locked, which causes uniform scaling of all three axes. Unlock the Lock X/Y/Z Scale checkbox to scale an object on a single axis only.



Viewer Transform controls

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Deep Image nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

dHoldout [dHld]



The dHoldout node

dHoldout Node Introduction

This tool uses a foreground Deep Image to occlude background Deep Image samples, either by removing them or computing the value after occlusion

Inputs

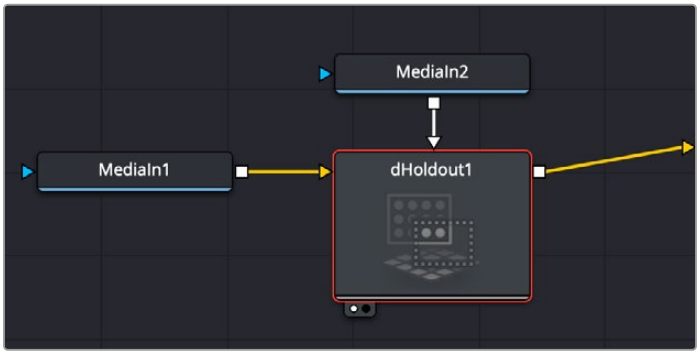
The dHoldout tool uses one background input and a foreground input.

Background: This orange input is connected from a node with deep image data that you want to use for the background.

Foreground: The white foreground input is for the second of two deep images which is typically a foreground subject that should be used to occlude the background.

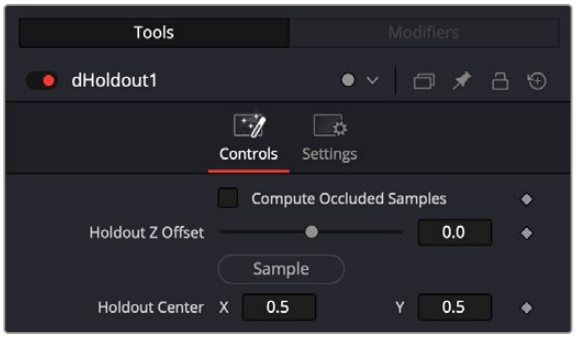
Basic Node Setup

The dHoldout node receives the background media containing the deep image data. It then lets you refine the depth boundaries of the foreground deep image to occlude the background.



dHoldout accepts deep image data from two sources and uses the foreground source to occlude the background.

Inspector



dHolodout controls

Controls Tab

The Controls tab includes parameters for adjusting the amount of depth information to include or exclude from occluding the background deep image.

- **Compute Occluded Samples:** Use this checkbox to remove samples from a semi-transparent foreground input. Visible background pixels will be recolored as if they were seen through a partially transparent foreground.
- **Holdout Z Offset:** This slider allows you to reposition the holdout depth samples forward or backwards. Decreasing this value will push the holdout forward, whereas increasing the value will have the opposite effect.
- **Sample:** Click and drag from the sample button into the viewer to choose a value.
- **Holdout Center:** This allows you to reposition the holdout along the the X and Y axis.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Deep Image nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

dMerge [dMg]



The dMerge node

dMerge Node Introduction

This tool combines multiple data streams, merging their depth samples for seamless compositing. This tool has numerous inputs similar to the merge 3D tool.

Inputs

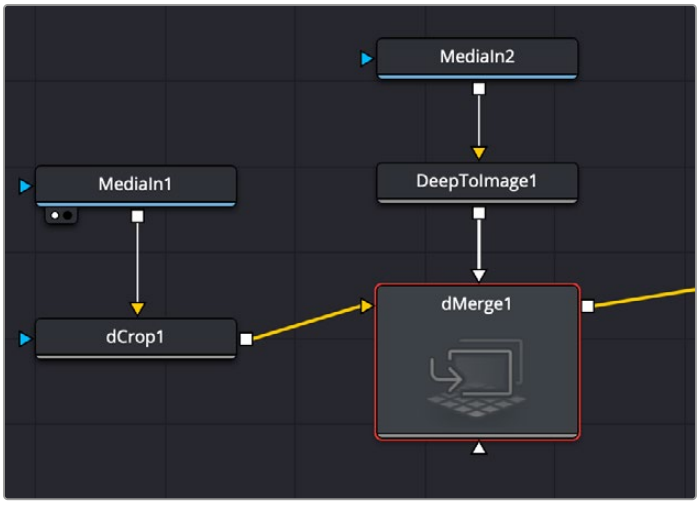
The dMerge tool uses one background input and multiple foreground inputs.

Background: This orange input is connected from a node with deep image data that you want to use for the background.

Foreground: The white foreground input is for connecting other deep image nodes that you want to combine with the background. There can be multiple foreground inputs to this tool just by dragging additional connections between the nodes.

Basic Node Setup

The dMerge node receives the background containing the deep image data. It then lets you merge additional deep image nodes to combine the data.



dMerge accepts deep image data from multiple sources and combines them.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Deep Image nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

dRecolor [dRc]



The dRecolor node

dRecolor Node Introduction

This tool combines Deep Image data with an RGB (multi-layer) image. This node has two inputs, one for Deep/depth and the other for Color. Use this node to recolor Deep Image samples with 2D RGB values

Inputs

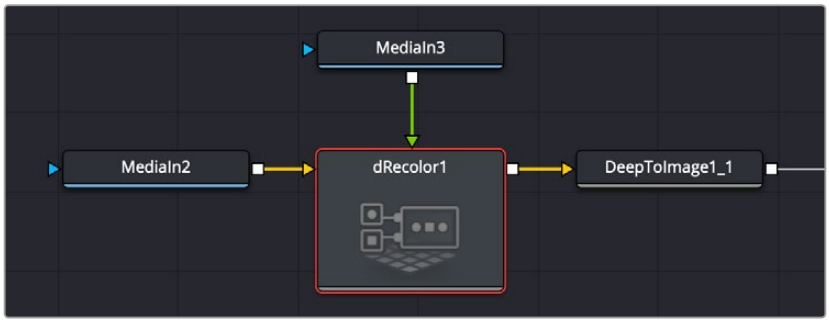
The dRecolor tool uses two inputs, one for deep color images and the other for RGB images.

Depth: This orange input is connected from a node with deep image data that you want to recolor.

Color: The green input is for the RGB image that you want to use to combine the Depth input with.

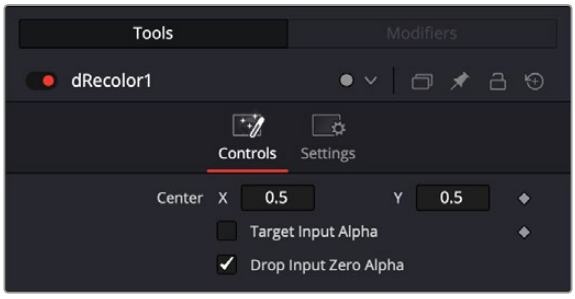
Basic Node Setup

The dRecolor node receives media containing deep image data. It then lets you combine it with a normal 2D image using the RGB input.



dRecolor accepts deep image data and lets you combine it with a 2D RGB image.

Inspector



dRecolor controls

Controls Tab

The Controls tab includes parameters for blending the 2D RGB image into the background deep image.

- **Center X/Y:** Reposition the 2D RGB image over the deep image.
- **Target Input Alpha:** Select this checkbox to enable the Alpha channel from the Color input, influencing the transparency of deep samples. When this option is turned off, only RGB will be used to recolor the deep image.
- **Drop Input Zero Alpha:** When this is enabled, pixels in the RGBA source with zero alpha are ignored and won't influence the deep image.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Deep Image nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

dResize [dRz]



The dResize node

dResize Node Introduction

This tool resizes a Deep Image's width and height.

Inputs

The dResize tool uses a single input.

Image: This orange input is connected from a node with deep image data that you want to resize.

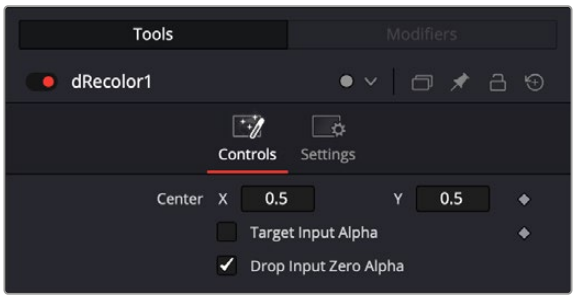
Basic Node Setup

The dResize node receives media containing deep image data. It then lets you resize it before passing it onto the rest of the node tree..



dResize accepts deep image data and lets you resize it.

Inspector



dResize controls

Controls Tab

The Controls tab includes parameters for resizing the deep image.

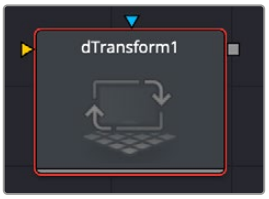
- **Width:** Scales the image's width.
- **Height:** Scales the image's height.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Deep Image nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

dTransform [dXf]



The dTransform node

dTransform Node Introduction

This tool allows you to adjust spatial positioning of deep image data while preserving its depth samples. You can scale and translate Deep Images, including the sample Z values.

Inputs

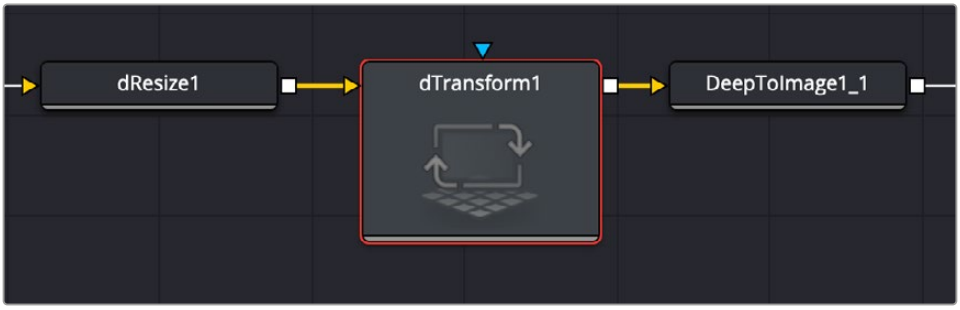
The dTransform tool uses a single input and an effect mask.

Image: This orange input is connected from a node with deep image data that you want to transform.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools.

Basic Node Setup

The dTransform node receives media containing deep image data. It then lets you transform it before passing it onto the rest of the node tree.



dTransform accepts deep image data and lets you transform it.

Inspector



dTransform controls

Controls Tab

The Controls tab includes parameters for transforming the deep image.

- **Scale Z:** The slider allows you to expand or contract the overall depth samples of the source. Increasing the scale will amplify the overall distance of depth. Decreasing will compress the depth range.
- **Translate Z:** Translate Z will allow you to move an object forward or backward along its depth channel. This gives you the control to relocate an object closer or farther away from camera.
- **Sample:** Click and drag from the sample button into the viewer to choose a value.
- **Center X/Y:** Move the center position of the deep image.
- **Pivot X/Y:** Move the pivot point of the deep image.
- **Size X/Y:** Resize the deep image on the selected axis.
- **Use Size and Aspect:** Check this box to activate the parameters below.
- **Size:** Scales the image equally along both the X and Y axes
- **Aspect:** Lets you set the aspect ratio of the image.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Deep Image nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Image to Deep [ItD]



The Image to Deep node

Image to Deep Node Introduction

This tool has an input for a 2D image. When connected, the image is given deep image data and can be composited in a deep environment via the dMerge node.

Inputs

The Image to Deep tool uses one input.

Input: This orange input is the only required connection. It is connected from a node with 2D image data that you want to use with a Deep Image.

Basic Node Setup

The Image to Deep node receives the media in containing the 2D image data. It then adds Deep Image information to it, allowing it to be used in deep image compositing with a dMerge node.

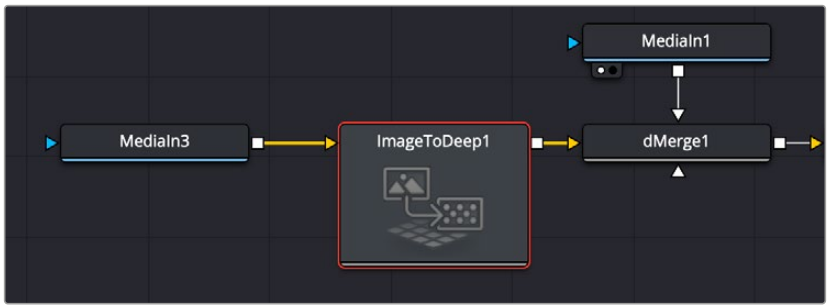


Image to Deep accepts a 2D image (MediaIn3) then adds Deep Image data, allowing it to be composited with Deep Image media (MediaIn1) via a dMerge node.

Inspector

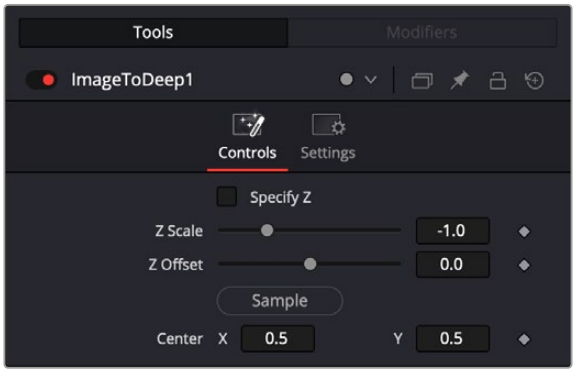


Image to Deep controls

Controls Tab

The Controls tab includes parameters for adjusting the depth information.

- **Specify Z:** Check this box to specify the current Z depth. If unchecked, you can adjust it using the Z Scale control below.
- **Z Scale:** The slider allows you to expand or contract the overall depth samples of the source. Increasing the scale will amplify the overall distance of depth. Decreasing will compress the depth range.
- **Z Offset:** Will allow you to move an object forward or backward along its depth channel. This gives you the control to relocate an object closer or farther away from camera.
- **Sample:** Click and drag from the sample button into the viewer to choose a value.
- **Center X/Y:** Reposition the 2D image along the X/Y axes.

Common Controls

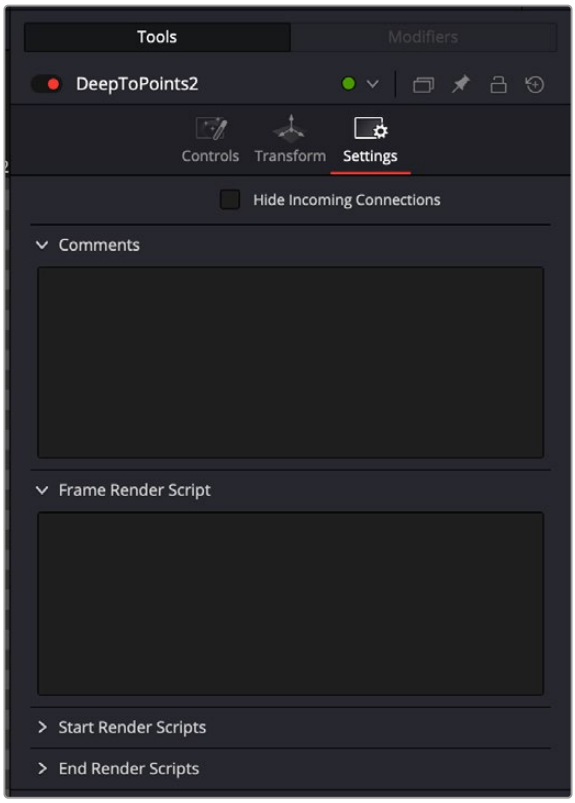
Settings Tab

The Settings tab in the Inspector is also duplicated in other Deep Image nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

The Common Controls

Settings Tab

The Common Settings tab can be found on most tools in Fusion. The following controls are specific settings for Deep Image nodes.



The Deep Image Settings tab

Hide Incoming Connections

Enabling this checkbox can hide connection lines from incoming nodes, making a node tree appear cleaner and easier to read. When enabled, empty fields for each input on a node are displayed in the Inspector. Dragging a connected node from the node tree into the field hides that incoming connection line as long as the node is not selected in the node tree. When the node is selected in the node tree, the line reappears.

Comments Tab

The Comments tab contains a single text control that is used to add comments and notes to the tool. When a note is added to a tool, a small red dot icon appears next to the setting's tab icon and a text bubble appears on the node. To see the note in the Node Editor, hold the mouse pointer over the node for a moment. The contents of the Comments tab can be animated over time, if required.

Scripting Tabs

The Scripting tab is present on every tool in Fusion. It contains several edit boxes used to add scripts that process when the tool is rendering. For more details on the contents of this tab, please consult the scripting documentation.

Chapter 37

Deep Pixel Nodes

This chapter details Deep Pixel nodes found in Fusion. Deep Pixel nodes are capable of handling AOVs (Arbitrary Output Variables) from 3D-rendered files.

The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Ambient Occlusion [SSAO] 1032
Depth Blur [DBl] 1035
Fog [Fog] 1037
Shader [Shd] 1039
Texture [Txr] 1042
The Common Controls 1044

Ambient Occlusion [SSAO]



The Ambient Occlusion node

Ambient Occlusion Node Introduction

Ambient Occlusion (AO) is the lighting caused when a scene is surrounded by a uniform diffuse spherical light source. Think of the scene as being surrounded by a humongous sphere that uniformly emits light from its surface. AO captures the low frequency lighting. It does not capture sharp shadows or Diffuse or Specular lighting. So, AO is usually combined with Diffuse and Specular lighting to create a full lighting solution.

The Ambient Occlusion node generates global lighting effects in 3D-rendered scenes as a post effect. It quickly approximates computationally expensive ray-traced global illumination. Being a post effect, it exposes similar aliasing issues like the Shader, Texture, and Volume Fog nodes. Hence, artifacts may appear in certain situations.

Usage

The AO node rarely works out of the box, and requires some tweaking. The setup process involves adjusting the Kernel Radius and Number Of Samples to get the desired affect.

The Kernel Radius depends on the natural “scale” of the scene. Initially, there might appear to be no AO at all. In most cases, the Kernel Radius is too small or too big, and working values must be found.

Inputs

There are three inputs on the AO node. The standard effect mask is used to limit the AO effect. The Input and Camera connections are required. If either of these is not supplied, the node does not render an image on output.

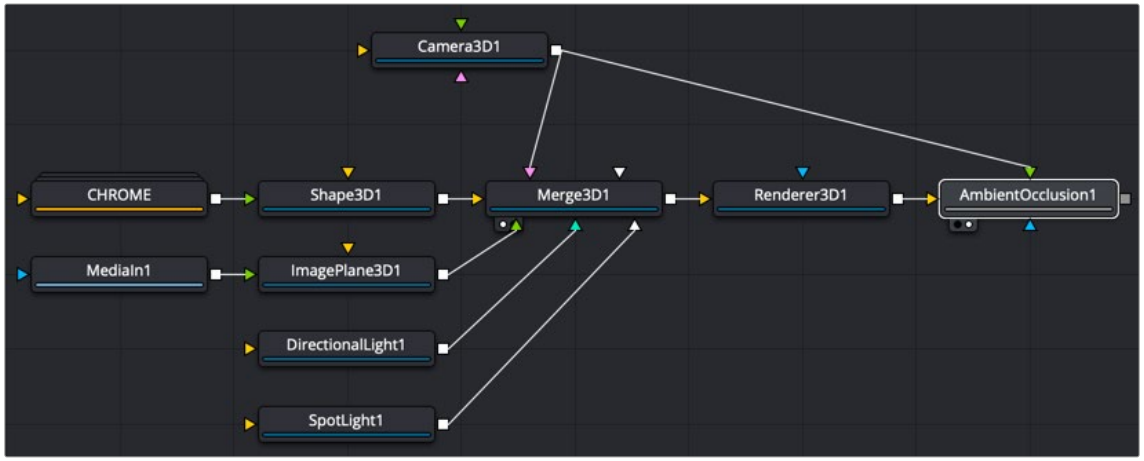
Input: This orange input accepts a 2D RGBA image, Z-Depth, and Normals.

Camera: The green camera input can take either a 3D Scene or a 3D Camera that rendered the 2D image.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the Ambient Occlusion to only those pixels within the mask. An effects mask is applied to the tool after the tool is processed.

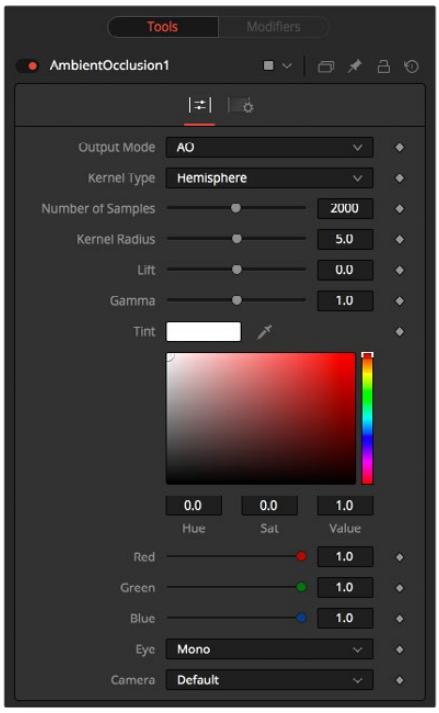
Basic Node Setup

The Ambient Occlusion node is typically placed after a Renderer 3D node. The Renderer 3D node must have Z-Depth and Normals enabled in its output channels. The Camera3D that is rendered by the Renderer3D node is then connected to the camera input on the AO node.



Ambient occlusion as a 2D post process

Inspector



Ambient Occlusion controls

Controls Tab

The controls tab includes all the main controls for compositing with AO. It controls the quality and appearance of the effect.

Output Mode

- **Color:** Using the Color menu option combines the incoming image with Ambient Occlusion applied.
- **AO:** This option outputs the pure Ambient Occlusion as a grayscale image. White corresponds to regions in the image that should be bright, while black correspond to regions that should be darker. This allows you to create a lighting equation by combining separate ambient/diffuse/specular passes. Having the AO as a separate buffer allows creative freedom to combine the passes in various ways.

Kernal Type

To determine the AO, rays are cast from a point on the surface being shaded, outward to a large enclosed sphere.

The AO factor is determined by the unoccluded rays that reach the sphere.

- **Hemisphere:** Rays are cast toward a hemisphere oriented to the surfaces normal. This option is more realistic than Sphere and should be used unless there is a good reason otherwise. Flat surfaces receive 100% ambient intensity, while other parts are darkened.
- **Sphere:** Rays are cast toward a sphere centered about the point being shaded. This option is provided to produce a stylistic effect. Flat surfaces receive 50% ambient intensity, while other parts are made darker or brighter.

Number of Samples

Increase the samples until artifacts in the AO pass disappear. Higher values can generate better results but also increase render time.

Kernel Radius

The Kernel Radius controls the size of the filter kernel in 3D space. For each pixel, it controls how far one searches in 3D space for occluders. The Filter Kernel should be adjusted manually for each individual scene.

If made too small, nearby occluders can be missed. If made too large, the quality of the AO decreases and the samples must be increased dramatically to get the quality back.

This value is dependent on the scene Z-depth. That means with huge Z values in the scene, the kernel size must be large as well. With tiny Z values, a small kernel size like 0.1 should be sufficient.

Lift/Gamma/Tint

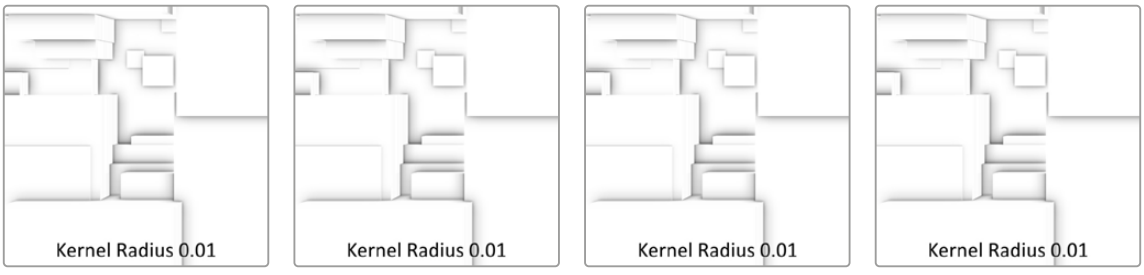
You can use the lift, gamma, and tint controls to adjust the AO for artistic effects.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Deep Pixel nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

TIP: Combining multiple AO passes with different kernel radii can produce better effects.



AO Tips and Limitations

Transparency/Translucency: AO is designed to work with opaque objects. There are known limitations with transparent receivers and those with transparent occluders. You can work around some of these limitations by splitting out the transparent/translucent objects into separate scenes and only computing AO on the opaque objects.

Particles: Because of the transparency/translucency limitations, do not use AO on particles, unless the particles are solid opaque geometry. Anti-aliased edges are another form of transparency, so they also cause problems with AO.

Supersampling: To render anti-aliasing with Ambient Occlusion, enable HiQ for the Z and Normals pass in the Renderer 3D.

Viewer Dependence: AO methods work in viewer space, and the results are viewer dependent. This means the amount of darkening can vary depending on the view location, when in reality it should be constant. If at a point on an object the AO is 0.5, moving the camera could change it to 0.4.

Baking of AO: The OpenGL UV renderer can be used to bake AO into the textures on models.

Depth Blur [DBl]



The Depth Blur node

Depth Blur Node Introduction

The Depth Blur node is primarily used to create focal length or depth-of-field effects. It blurs 3D-rendered images based on included Z-channel values, and can also be used for general per-pixel blurring effects through the Blur Channel controls.

Inputs

The Depth Blur node includes three inputs: one for the main image, one for a blur image, and another for an effect mask to limit the area where the depth blur is applied.

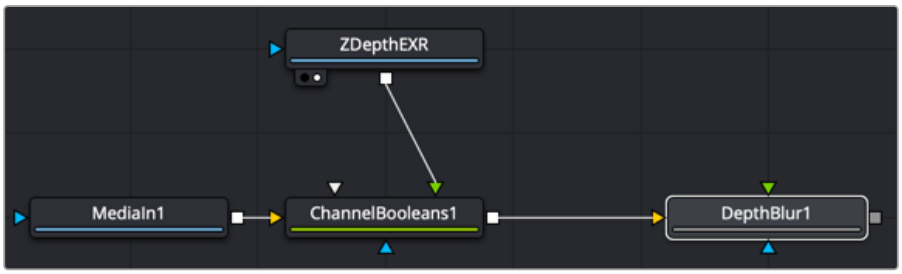
Input: This orange input is the only required connection. It accepts a 2D image that includes a Z channel. The Z channel is used to determine the blur amount in different regions of the image.

Blur Image: If the Blur Image input is connected, channels from the image are used to control the blur. This allows general 2D per-pixel blurring effects.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the depth blur to only those pixels within the mask. An effects mask is applied to the tool after the tool is processed.

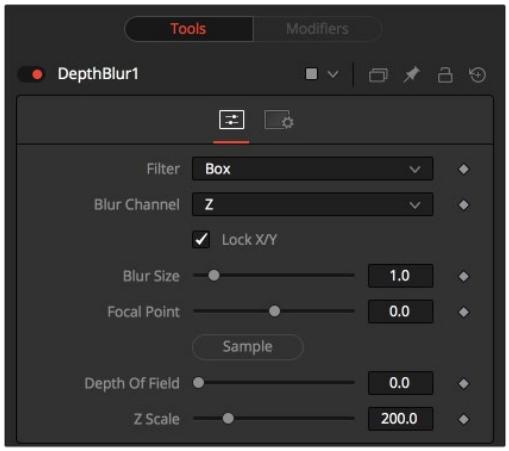
Basic Node Setup

The Depth Blur node receives the image containing the Z channel. Below, the Z depth channel is provided in a separate image file and combined with the RGB (beauty) image using a Channels Booleans tool. Channel Booleans sets the Z buffer channel to copy into the luminance foreground.



Depth Blur is applied to the beauty image based on a Z depth channel

Inspector



Depth Blur controls

Controls Tab

The Controls tab includes parameters for adjusting the amount of blur applied and the depth of the blurred area. It also includes options for selecting channels other than the Z channel for the blur map.

Filter

This menu selects the filter used for the blur.

- **Box:** This applies a basic depth-based box blur effect to the image.
- **Soften:** This applies a depth-based general softening filter effect.
- **Super Soften:** This applies a depth-based high-quality softening filter effect.

Blur Channel

Select one of these options to determine the channel used to control the level of blur applied to each pixel. The channel from the main image input is used, unless an image is connected to the node's green Blur Image input.

Lock X/Y

When toggled on, this control locks the X and Y Blur sliders together for symmetrical blurring.

Blur Size

This slider is used to set the strength of the horizontal and vertical blurring.

Focal Point

This control is visible only when the Blur channel menu is set to use the Z channel.

Use this control to select the distance of the simulated point of focus. Lowering the value causes the Focal Point to be closer to the camera; raising the value causes the Focal Point to be farther away.

Depth of Field

This control is used to determine the depth of the area in focus. The focal point is positioned in the middle of the region, and all pixels with a Z-value within the region stay in focus. For example, if the focal point were selected from the image and set to a value of 300, and the depth of field is set to 200, any pixel with a Z-value between 200 and 400 would remain in focus.

Z Scale

Scales the Z-buffer value by the selected amount. Raising the value causes the distances in the Z-channel to expand. Lowering the value causes them to contract. This is useful for exaggerating the depth effect. It can also be used to soften the boundaries of the blur. Some images with small depth values may require the Z-scale to be set quite low, below 1.0.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Deep Pixel nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Fog [Fog]



The Fog node

Fog Node Introduction

The Fog node is used to create simulated fog effects on 3D-rendered images that contain a valid Z-buffer channel. The fog can be placed in front of or behind various elements of a rendered image based on the selected Z-channel planes.

Inputs

The Fog node includes three inputs: one for the main image with a Z channel, one for a blur image, and another for an effect mask to limit the area where the depth blur is applied.

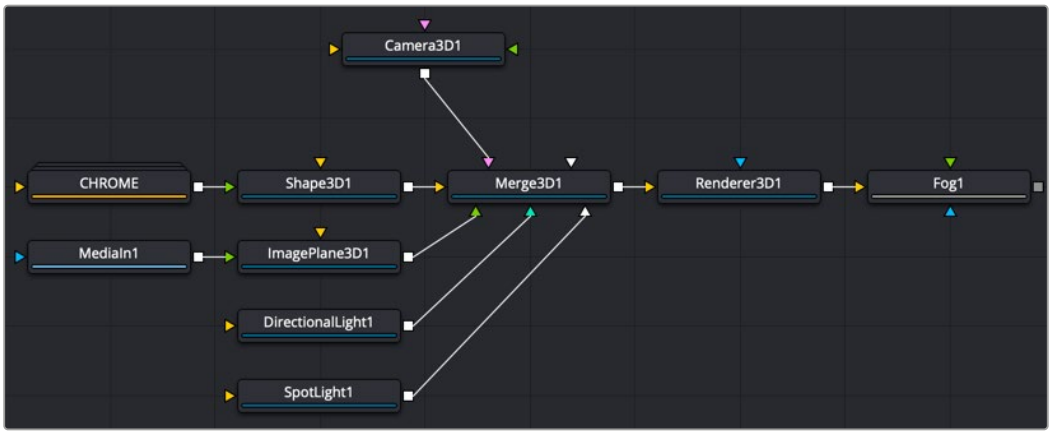
Input: This orange input is the only required connection. It accepts a 2D image that includes a Z channel. The Z channel is used to determine the fog amount in different regions of the image.

Blur Image: The green second image input connects an image that is used as the source of the fog. If no image is provided, the fog consists of a single color. Generally, a noise map of some sort is connected here.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the fog to only those pixels within the mask. An effects mask is applied to the tool after the tool is processed.

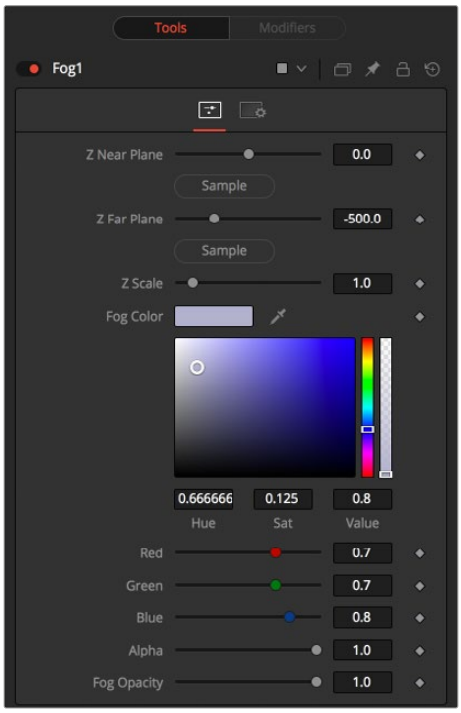
Basic Node Setup

The Fog node is typically placed after a Renderer 3D node. The Renderer 3D node must have Z-Depth enabled in its output channels.



A Fog node added as a 2D post process after the Renderer 3D node

Inspector



Fog controls

Controls Tab

The Controls tab includes parameters for adjusting the density and color of the fog.

Z-Buffer Near Plane and Far Plane

These controls are used to select the extents of the fog within the scene. To pick a value, drag the Pick button to an area on the image being viewed where the plane is to be located.

The Near Plane is used to select the depth where the fog thins out to nothing. The Far Plane is used to select the depth at which the fog becomes opaque.

Z Depth Scale

This option scales the Z-buffer values by the selected amount. Raising the value causes the distances in the Z-channel to expand, whereas lowering the value causes the distances to contract. This is useful for exaggerating the fog effect.

Fog Color

This option displays and controls the current fog color. Alpha adjusts the fog's transparency value.

Fog Opacity

Use this control to adjust the opacity on all channels of the fog.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Deep Pixel nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Shader [Shd]



The Shader node

Shader Node Introduction

The Shader node can control the lighting, reflection mapping, and 3D shading of elements in a rendered image. The node relies on the presence of the normal map channel in a rendered image. If this channel is not present, this node has no effect.

Inputs

The Shader node includes three inputs: one for the main image with normal map channels, one for a reflection map, and another for an effect mask to limit the area where the depth blur is applied.

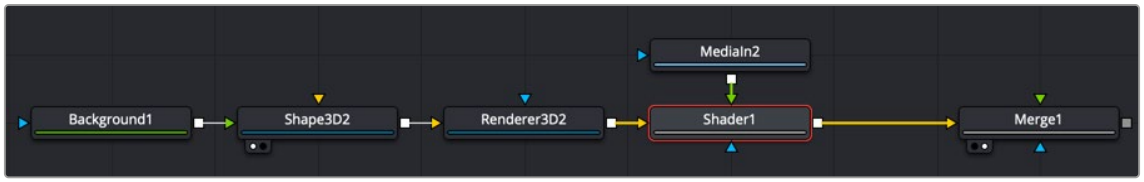
Input: This orange input is the only required connection. It accepts a 2D image that includes a normals channel.

Reflection Map Image: The green reflection map image input projects an image onto all elements in the scene or to elements selected by the Object and Material ID channels in the Common Controls. Reflection maps work best as 32-bit floating point, equirectangular formatted images

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the shader to only those pixels within the mask. An effects mask is applied to the tool after the tool is processed.

Basic Node Setup

The Shader node is inserted after a 2D image that contains a Normals channel. Below, a Renderer 3D is used to add a Normals channel to an image. The Shader node uses the normals for refining the surface appearance with a reflection map connected.



The Shader node using normals from a Renderer 3D and a reflection input

Inspector



Shader controls

Controls Tab

The Controls tab for the Shader node includes parameters for adjusting the overall surface reaction to light sources. You can modify the ambient, diffuse, specular, and reflection properties of the image connected to the orange image input.

Light Tab

The Controls tab includes parameters for basic lighting brightness and reflections.

Ambient

Ambient controls the Ambient color present in the scene or the selected object. This is a base level of light added to all pixels, even in completely shadowed areas.

Diffuse

This option controls the Diffuse color present in the scene or for the selected object. This is the normal color of the object, reflected equally in all directions.

Specular

This option controls the Specular color present in the scene or for the selected object. This is the color of the glossy highlights reflected toward the eye from a light source.

Reflection

This option controls the Reflection contribution in the scene or for the selected object. High levels make objects appear mirrored, while low levels overlay subtle reflections giving a polished effect. It has no effect if no reflection map is connected.

Reflection Type

This menu determines the type of reflection mapping used to project the image in the second input.

Screen: Screen causes the reflection map to appear as if it were projected on to a screen behind the point of view.

Spherical: Spherical causes the reflection map to appear as if it were projected on to a huge sphere around the whole scene.

Refraction: Refraction causes the reflection map to appear as if it were refracting or distorting according to the geometry in the scene.

Equator Angle

Equator Angle controls the left to right angle of the light generated and mapped by the Shader node for the scene or the selected object.

Polar Height

Polar Height controls the top to bottom angle of the light generated and mapped by the Shader node for the scene or the selected object.



Shader tab controls

Shader Tab

The Shader tab is used to adjust the falloff of the Diffuse and Specular light and the tint color of the specular highlight.

Diffuse and Specular

When enabled, these checkboxes allow you to edit the Diffuse and/or Specular Shader curves in the Shader spline window.

In and Out

These options are used to display and edit point values on the spline.

Specular Color

Use the Diffuse curve to manipulate the diffuse shading and the Specular curve to affect the specular shading. Drag a box over several points to group-select them. Right-clicking displays a menu with options for adjusting the spline curves.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Deep Pixel nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Texture [Txr]



The Texture node

Texture Node Introduction

The Texture node controls the texture mapping of elements in a rendered image. The Texture node relies on the presence of U and V Map channels in a 3D-rendered image connected to the main Image input. If these channels are not present, this node has no effect.

NOTE: Background pixels may have U and V values of 0.0, which set those pixels to the color of the texture’s corner pixel. To restrict texturing to specific objects, use an effect mask based on the Alpha of the object or its Object or Material ID channel.

Inputs

The Texture node includes three inputs: one for the main image with UV map channels, one for a texture map image, and another for an effect mask to limit the area where the replace texture is applied.

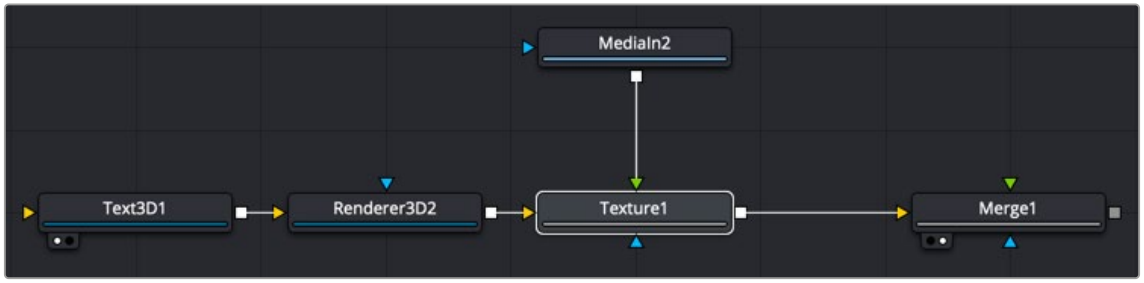
Input: This orange input accepts a 2D image that includes UV channels. If the UV channels are not in the images, this node has no effect.

Texture: The green texture map input provides the texture that is wrapped around objects, replacing the current texture.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the texture to only those pixels within the mask. An effects mask is applied to the tool after the tool is processed.

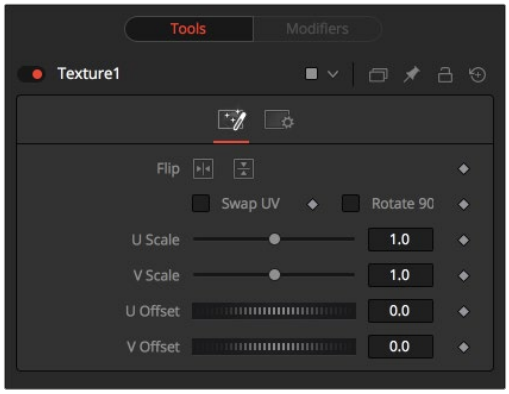
Basic Node Setup

The Texture node is inserted after a 2D image that contains a Texture UV channel. Below, a Renderer 3D is used to add texture coordinates to 3D text. The Texture node can then be used to manipulate those coordinates using a new texture connected to the green texture input.



A Texture node used to manipulate the texture coordinates and add texture to a Text 3D node

Inspector



Texture controls

Texture Tab

The Texture tab controls allow you to flip, swap, scale, and offset the UV texture image connected to the texture input.

Flip Horizontal and Vertical

Use these two buttons to flip the texture image horizontally and/or vertically.

Swap UV

When this checkbox is selected, the U and V channels of the source image are swapped.

Rotate 90

The texture map image is rotated 90 degrees when this checkbox is enabled.

U and V Scale

These controls change the scaling of the U and V coordinates used to map the texture. Changing these values effectively enlarges and shrinks the texture map as it is applied.

U and V Offset

Adjust these controls to offset the U and V coordinates. Changing the values causes the texture to appear to move along the geometry of the object.

Common Controls

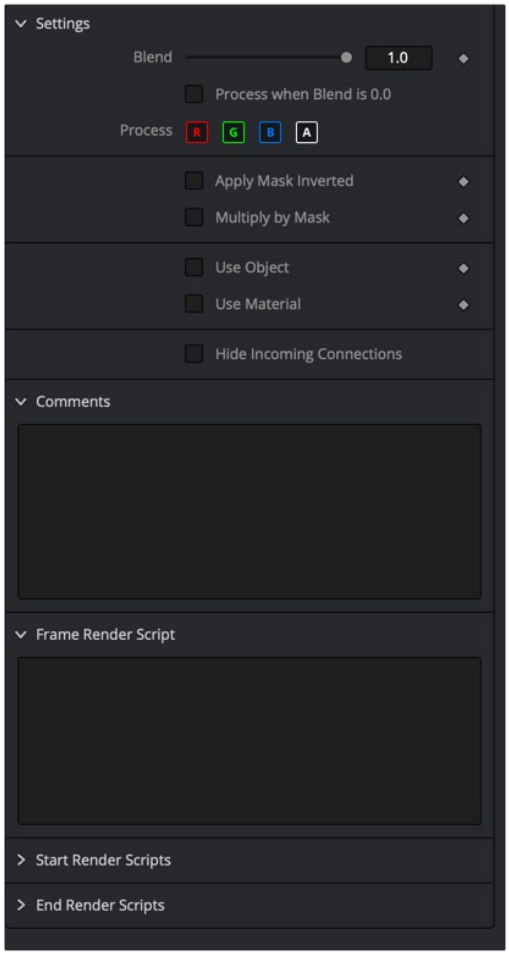
Settings Tab

The Settings tab in the Inspector is also duplicated in other Deep Pixel nodes. These common controls are described in the following “The Common Controls” section.

The Common Controls

Nodes that handle Deep Pixel compositing operations share several identical controls in the Inspector. This section describes controls that are common among Deep Pixel nodes.

Inspector



Common Settings tab in Deep Pixel Nodes

Settings Tab

The Settings tab in the Inspector can be found on every tool in the Deep Pixel category. The Settings controls are even found on third-party Deep Pixel-type plugin tools. The controls are consistent and work the same way for each tool although some tools do include one or two individual options that are also covered here.

Blend

The Blend control is used to blend between the tool's original image input and the tool's final modified output image. When the blend value is 0.0, the outgoing image is identical to the incoming image. Normally, this causes the tool to skip processing entirely, copying the input straight to the output.

Process When Blend Is 0.0

The tool is processed even when the input value is zero. This can be useful if processing of this node is scripted to trigger another task, but the value of the node is set to 0.0.

Red/Green/Blue/Alpha Channel Selector

These four buttons are used to limit the effect of the tool to specified color channels. This filter is often applied after the tool has been processed.

For example, if the red button on a Blur tool is deselected, the blur is first applied to the image, and then the red channel from the original input is copied back over the red channel of the result.

There are some exceptions, such as tools for which deselecting these channels causes the tool to skip processing that channel entirely. Tools that do this generally possess a set of identical RGBA buttons on the Controls tab in the tool. In this case, the buttons in the Settings and the Controls tabs are identical.

Apply Mask Inverted

Enabling the Apply Mask Inverted option inverts the complete mask channel for the tool. The mask channel is the combined result of all masks connected to or generated in a node.

Multiply by Mask

Selecting this option causes the RGB values of the masked image to be multiplied by the mask channel's values. This causes all pixels of the image not in the mask (i.e., set to 0) to become black/transparent.

Use Object/Use Material (Checkboxes)

Some 3D software can render to file formats that support additional channels. Notably, the EXR file format supports Object and Material ID channels, which can be used as a mask for the effect. These checkboxes determine whether the channels are used if present. The specific Material ID or Object ID affected is chosen using the next set of controls.

Correct Edges

This checkbox appears only when the Use Object or Use Material checkboxes are selected. It toggles the method used to deal with overlapping edges of objects in a multi-object image. When enabled, the Coverage and Background Color channels are used to separate and improve the effect around the edge of the object. If this option is disabled (or no Coverage or Background Color channels are available), aliasing may occur on the edge of the mask.

For more information, *see Chapter 18, "Understanding Image Channels,"* in the Fusion Reference Manual.

Object ID/Material ID (Sliders)

Use these sliders to select which ID is used to create a mask from the object or material channels of an image. Use the Sample button in the same way as the Color Picker: to grab IDs from the image displayed in the viewer. The image or sequence must have been rendered from a 3D software package with those channels included.

Use GPU

The Use GPU menu has three settings. Setting the menu to Disable turns off hardware-accelerated rendering using the graphics card in your computer. Enabled uses the hardware. Auto uses a capable GPU if one is available, but falls back to software rendering when a capable GPU is not available.

Motion Blur

- **Motion Blur:** This toggles the rendering of Motion Blur on the tool. When this control is toggled on, the tool's predicted motion is used to produce the motion blur caused by the virtual camera's shutter. When the control is toggled off, no motion blur is created.
- **Quality:** Quality determines the number of samples used to create the blur. A quality setting of 2 causes Fusion to create two samples to either side of an object's actual motion. Larger values produce smoother results but increase the render time.
- **Shutter Angle:** Shutter Angle controls the angle of the virtual shutter used to produce the motion blur effect. Larger angles create more blur but increase the render times. A value of 360 is the equivalent of having the shutter open for one whole frame exposure. Higher values are possible and can be used to create interesting effects.
- **Center Bias:** Center Bias modifies the position of the center of the motion blur. This allows the creation of motion trail effects.
- **Sample Spread:** Adjusting this control modifies the weighting given to each sample. This affects the brightness of the samples.

Comments

The Comments field is used to add notes to a tool. Click in the field and type the text. When a note is added to a tool, a small red square appears in the lower-left corner of the node when the full tile is displayed, or a small text bubble icon appears on the right when nodes are collapsed. To see the note in the Node Editor, hold the mouse pointer over the node to display the tooltip.

Scripts

Three Scripting fields are available on every tool in Fusion from the Settings tab. They each contain edit boxes used to add scripts that process when the tool is rendering. For more details on scripting nodes, please consult the Fusion scripting documentation.

Chapter 38

Effect Nodes

This chapter details the Effect nodes in Fusion.

The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Duplicate [Dup] 1048
Highlight [HIL] 1054
Hot Spot [Hot] 1057
Object Removal [ORm] 1063
Pseudo Color [PsCl] 1067
Rays [CIR] 1069
Shadow [SH] 1070
Trails [TRLS] 1073
TV [TV] 1078
The Common Controls 1081

Duplicate [Dup]



The Duplicate node

Duplicate Node Introduction

Similar to the Duplicate 3D node, the Duplicate node can be used to quickly duplicate any 2D image, applying a successive transformation to each, and creating repeating patterns and complex arrays of objects. The options in the Jitter tab allow for non-uniform transformations, such as random positioning or sizes.

Inputs

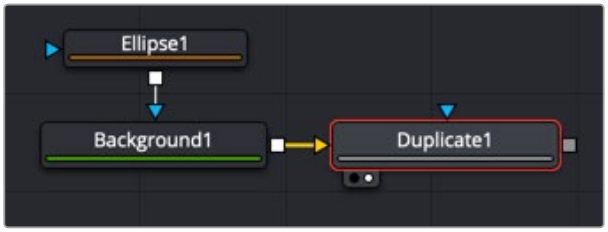
The two inputs on the Duplicate node are used to connect a 2D image and an effect mask, which can be used to limit the area where duplicated objects appear.

Input: The orange input is used for the primary 2D image that is duplicated.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the duplicated objects to appear only those pixels within the mask. An effects mask is applied to the tool after the tool is processed.

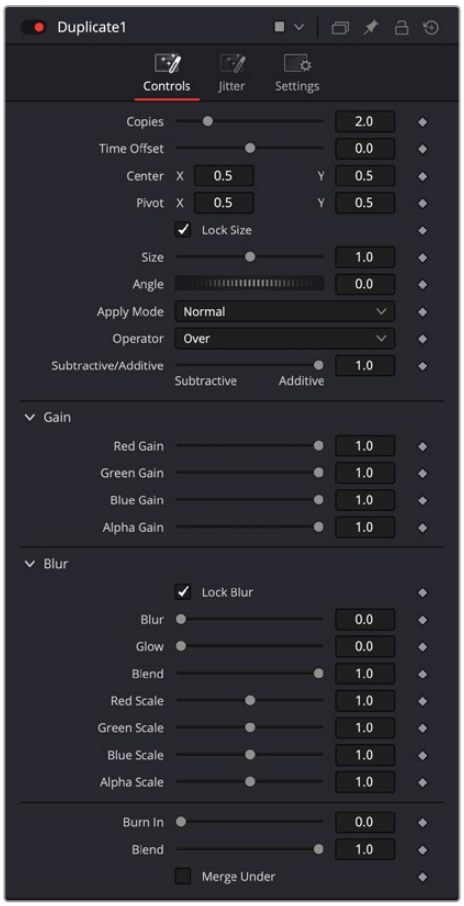
Basic Node Setup

The Duplicate node can be used in a variety of different ways and with a variety of different inputs. Below, to create motion graphics, a masked Background node creates a circular shape that is duplicated in the Duplicate node.



A Duplicate node used to create a repeating circular object

Inspector



Duplicate controls

Controls Tab

The Controls tab includes all the parameters you can use to create, offset, and scale copies of the object connected to the input on the node.

Copies

Use this slider to set the number of copies made. Each copy is a copy of the last copy. So, when set to 5, the parent is copied, then the copy is copied, then the copy of the copy is copied, and so on. This allows for some interesting effects when transformations are applied to each copy using the following controls.

Time Offset

Use the Time Offset slider to offset any animations that are applied to the original image by a set amount per copy. For example, set the value to -1.0 and use a square set to rotate on the Y-axis as the source. The first copy shows the animation from a frame earlier. The second copy shows animation from a frame before that, and so forth. This can be used with great effect on textured planes, for example, where successive frames of a clip can be shown.

Center

The X and Y Center controls set the offset position applied to each copy. An X offset of 1 would offset each copy 1 unit along the X-axis from the last copy.

Pivot

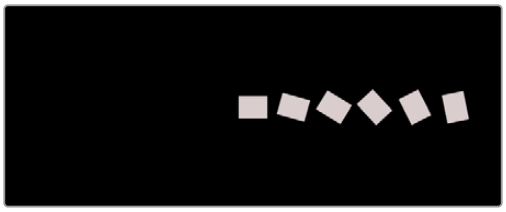
The Pivot controls determine the position of the pivot point used when changing the size, position, or angle of each copy. The pivot does not move with the original object or the duplicated array. To have the pivot follow the army, you must modify the pivot controls.

Size

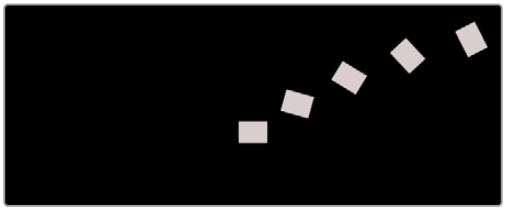
The Size control determines how much scaling to apply to each copy.

Angle

The Angle control sets the amount of Z rotation applied to each copy. The angle adjustment is linear based on the location of the pivot point.



Angle adjustment with centered pivot



Angle adjustment with offset pivot

Apply Mode

The Apply Mode setting determines the math used when blending or combining duplicated objects that overlap.

- **Normal:** The default mode uses the foreground object's Alpha channel as a mask to determine which pixels are transparent and which are not. When this is active, another menu shows possible operations, including Over, In, Held Out, Atop, and XOr.
- **Screen:** Screen blends the objects based on a multiplication of their color values. The Alpha channel is ignored, and layer order becomes irrelevant. The resulting color is always lighter. Screening with black leaves the color unchanged, whereas screening with white always produces white. This effect creates a similar look to projecting several film frames onto the same surface. When this is active, another menu shows possible operations, including Over, In, Held Out, Atop, and XOr.
- **Dissolve:** Dissolve mixes overlapping objects. It uses a calculated average of the objects to perform the mixture.
- **Multiply:** Multiplies the values of a color channel. This gives the appearance of darkening the object as the values are scaled from 0 to 1. White has a value of 1, so the result would be the same. Gray has a value of 0.5, so the result would be a darker object or, in other words, an object half as bright.
- **Overlay:** Overlay multiplies or screens the color values of the foreground object, depending on the color values of the object behind. Patterns or colors overlay the existing pixels while preserving the highlights and shadows of the color values of the objects behind the foreground objects. Objects behind other objects are not replaced but mixed with the front objects to reflect the original lightness or darkness of the objects behind.

- **Soft Light:** Soft Light darkens or lightens the foreground object, depending on the color values of the objects behind them. The effect is similar to shining a diffused spotlight on the image.
- **Hard Light:** Hard Light multiplies or screens the color values of the foreground object, depending on the color values of the objects behind them. The effect is similar to shining a harsh spotlight on the image.
- **Color Dodge:** Color Dodge uses the foreground object's color values to brighten the objects behind them. This is similar to the photographic practice of dodging by reducing the exposure of an area of a print.
- **Color Burn:** Color Burn uses the foreground object's color values to darken the objects behind them. This is similar to the photographic practice of burning by increasing the exposure of an area of a print.
- **Darken:** Darken looks at the color information in each channel and selects the object's foreground or background's color value, whichever is darker, as the result color. Pixels lighter than the blended colors are replaced, and pixels darker than the blended color do not change.
- **Lighten:** Lighten looks at the color information in each channel and selects the object's foreground or background's color values, whichever is lighter, as the result color value. Pixels darker than the blended color are replaced, and pixels lighter than the blended color do not change.
- **Difference:** Difference looks at the color information in each channel and subtracts the foreground object's color values from the background object's color values or the behind object's values from the foreground object's values, depending on which has the higher brightness value. Blending with white inverts the color. Blending with black produces no change.
- **Exclusion:** Exclusion creates an effect similar to but lower in contrast than the Difference mode. Blending with white inverts the base color values. Blending with black produces no change.
- **Hue:** Hue creates a result color with the luminance and saturation of the background objects color values and the hue of the foreground object's color values.
- **Saturation:** Saturation creates a result color with the luminance and hue of the base color and the saturation of the blend color.
- **Color:** Color creates a result color with the luminance of the background object's color value and the hue and saturation of the objects in the foreground. This preserves the gray levels in the image and is useful for colorizing monochrome objects.
- **Luminosity:** Luminosity creates a color using the hue and saturation of the background object and the luminance of the foreground object. This mode creates an inverse effect from that of the Color mode.

Operator

This menu is used to select the Operation mode used when the duplicate objects overlap. Changing the Operation mode changes how the overlapping objects are combined. This drop-down menu is visible only when the Apply mode is set to Normal.

The formula used to combine pixels in the Duplicate node is always (fg object * x) + (bg object * y). The different operations determine what x and y are, as shown in the description for each mode.

The Operator Modes are as follows:

- **Over:** The Over mode adds the foreground object to the background object by replacing the pixels in the background with the pixels from the Z wherever the foreground object's Alpha channel is greater than 1.

  ```
  x = 1, y = 1 - [foreground object Alpha]
  ```

- **In:** The In mode multiplies the Alpha channel of the background object against the pixels in the foreground object. The color channels of the foreground object are ignored. Only pixels from the

foreground object are seen in the final output. This essentially clips the foreground object using the mask from the background object.

```
x = [background Alpha], y = 0
```

- **Held Out:** Held Out is essentially the opposite of the In operation. The pixels in the foreground object are multiplied against the inverted Alpha channel of the background object.

```
x = 1 - [background Alpha], y = 0
```

- **Atop:** Atop places the foreground object over the background object only where the background object has a matte.

```
x = [background Alpha], y = 1 - [foreground Alpha]
```

- **XOr:** XOr combines the foreground object with the background object wherever either the foreground or the background have a matte, but never where both have a matte.

```
x = 1 - [background Alpha], y = 1-[foreground Alpha]
```

Subtractive/Additive

This slider controls whether Fusion performs an Additive composite, a Subtractive composite, or a blend of both when the duplicate objects overlap. This slider defaults to Additive assuming the input image's Alpha channel is premultiplied (which is usually the case). If you don't understand the difference between Additive and Subtractive compositing, here's a quick explanation.

An Additive blend operation is necessary when the foreground image is premultiplied, meaning that the pixels in the color channels have been multiplied by the pixels in the Alpha channel. The result is that transparent pixels are always black since any number multiplied by 0 always equals 0. This obscures the background (by multiplying with the inverse of the foreground Alpha), and then adds the pixels from the foreground.

A Subtractive blend operation is necessary if the foreground image is not premultiplied. The compositing method is similar to an additive composite, but the foreground image is first multiplied by its Alpha, to eliminate any background pixels outside the Alpha area.

While the Additive/Subtractive option is often an either/or mode in most other applications, the Duplicate node lets you blend between the Additive and Subtractive versions of the compositing operation. This can be useful for dealing with problem composites with bright or dark edges.

For example, using Subtractive merging on a premultiplied image may result in darker edges, whereas using Additive merging with a non-premultiplied image causes any non-black area outside the foreground's Alpha to be added to the result, thereby lightening the edges. By blending between Additive and Subtractive, you can tweak the edge brightness to be just right for your situation.

Gain

The Gain RGB controls multiply the values of the image channel linearly. All pixels are multiplied by the same factor, but the effect is larger on bright pixels and smaller on dark pixels. Black pixels are not changed since multiplying any number times 0 always equals 0.

Alpha Gain linearly scales the Alpha channel values of objects in front. This effectively reduces the amount that the objects in the background are obscured, thus brightening the overall result. When the Subtractive/Additive slider is set to Additive with Alpha Gain set to 0.0, the foreground pixels are simply added to the background.

When Subtractive/Additive slider is set to Subtractive, this controls the density of the composite, similarly to Blend.

All Gain values will compound based on the number of duplications.

Blur

Adds a blurring effect to the duplicated layers.

- **Lock Blur:** Locks the X and Y Blur sliders together for symmetrical blurring. This is enabled by default. When the Lock Blur control is deselected, independent control over each axis is provided.
- **Blur:** Sets the amount of blur applied to the duplicated layers in the tool. The Blur amount will not compound based on the number of duplications.
- **Glow:** Adds a glow effect to the blur of the duplicated layers.
- **Blend:** The Blend slider determines the percentage of the affected image that is mixed with original image. It blends in more of the original image as the value gets closer to 0.
- **RGBA Scale:** Allows adjusting the strength of the individual Red, Green, Blue, and Alpha channels to the blur of the duplicated layers.

Burn In

The Burn In control adjusts the amount of Alpha used to darken the objects that fall behind other objects, without affecting the amount of foreground objects added. At 0.0, the blending behaves like a straight Alpha blend, in contrast to a setting of 1.0 where the objects in the front are effectively added on to the objects in the back (after Alpha multiplication if in Subtractive mode). This gives the effect of the foreground objects brightening the objects in the back, as with Alpha Gain. In fact, for Additive blends, increasing the Burn In gives an identical result to decreasing Alpha Gain.

Blend

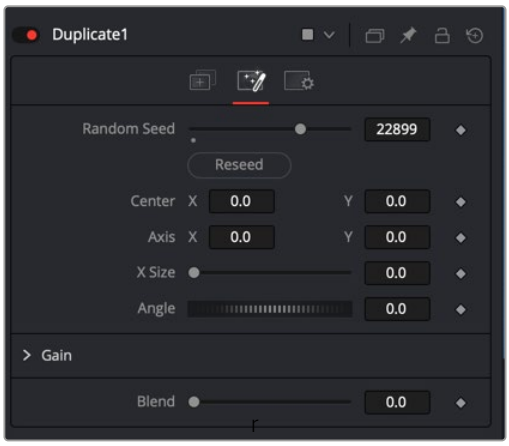
This blend control is different from the Blend slider in the Common Settings tab. Changes made to this control apply the blend between objects. The Blend slider fades the results of the last object first, the penultimate after that, and so on. The blending is divided between 0 and 1, with 1 being all objects are fully opaque and 0 being only the original object showing.

Merge Under

This checkbox reverses the layer order of the duplicated elements, making the last copy the bottommost layer and the first copy the topmost layer.

Jitter Tab

The options in the Jitter tab allow you to randomize the position, rotation, size, and color of all the copies created in the Controls tab.



Duplicate Jitter tab

Random Seed

The Random Seed slider and Reseed button are used to generate a random starting point for the amount of jitter applied to the duplicated objects. Two Duplicate nodes with identical settings but different random seeds produce two completely different results.

Center X and Y

Use these two controls to adjust the amount of variation in the X and Y position of the duplicated objects.

Axis X and Y

Use these two controls to adjust the amount of variation in the rotational pivot center of the duplicated objects. This affects only the additional jitter rotation, not the rotation produced by the Rotation settings in the Controls tab.

X Size

Use this control to adjust the amount of variation in the Scale of the duplicated objects.

Angle

Use this dial to adjust the amount of variation in the Z rotation of the duplicated objects.

Gain

The Gain RGBA controls randomly multiply the values of the image channel linearly.

Blend

Changes made to this control randomize the blend between objects.

Common Controls

Settings Tab

The Settings tab controls are common to all Effect nodes, so their descriptions can be found in "The Common Controls" section at the end of this chapter.

Highlight [HIL]



The Highlight node

Highlight Node Introduction

The Highlight filter creates star-shaped highlights or glints in bright regions of the image, similar to a lens star filter effect.

Inputs

There are three Inputs on the Highlight node: one for the image, one for the effects mask, and another for a highlight mask.

Input: The orange input is used for the primary 2D image that gets the highlight applied.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input restricts the highlight to be within the pixels of the mask. An effects mask is applied to the tool after the tool is processed.

Highlight Mask: The Highlight node supports pre-masking using the white highlight mask input. The image is filtered before the highlight is applied. The highlight is then merged back over the original image. Unlike regular effect masks, it does not crop off highlights from source pixels when the highlight extends past the edges of the mask.

Highlight masks are identical to effects masks in every other respect.

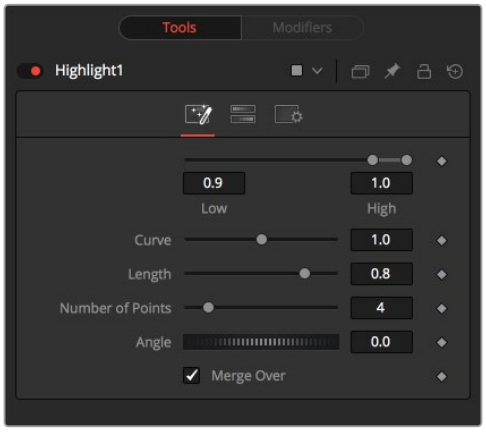
Basic Node Setup

The Highlight node below is used to create glint-type highlights on an incoming image. The highlight mask is used to limit the area where the effect is applied.



A Highlight node applied to an image, with a highlight mask limiting the area of the effect

Inspector



Highlight controls

Controls Tab

The Controls tab includes parameters for the highlight style except for color, which is handled in the Color Scale tab.

Low and High

This range control designates the range of Luminance values in the image that generates highlights. Values less than the Low value do not receive highlights. Values above the High value receive the full highlight effect.

Curve

The Curve value changes the drop-off over the length of the highlight. Higher values cause the brightness of the flares to drop off closer to the center of the highlight, whereas lower values drop off farther from the center.

Length

This designates the length of the flares from the highlight.

Number of Points

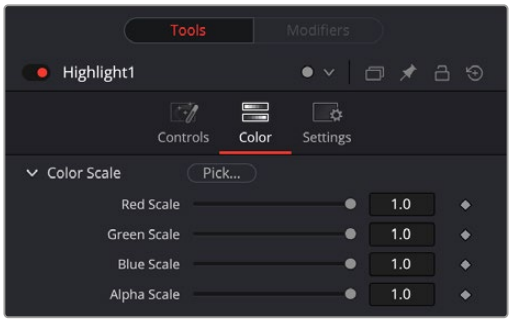
This determines the number of flares emanating from the highlight.

Angle

Use this control to rotate the highlights.

Merge Over

When enabled, the effect is overlaid on the original image. When disabled, the output is the highlights only. This is useful for downstream color correction of the highlights.



Highlight Color Scale controls

Color Scale Tab

The Color Scale tab controls the color of the highlight.

By click and holding on the Pick button, then dragging the pointer over the viewer, you can select a specific color from the image.

Red, Green, and Blue Scale

Moving the sliders of one or all of these channels down changes the falloff color of the highlight.

Alpha Scale

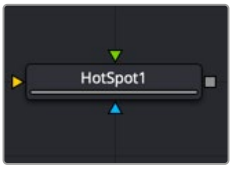
Moving the Alpha slider down makes highlight falloff more transparent.

Common Controls

Setting Tab

The Settings tab controls are common to all Effect nodes, so their descriptions can be found in “The Common Controls” section at the end of this chapter.

Hot Spot [Hot]



The Hot Spot node

Hot Spot Node Introduction

The Hot Spot node is used to create lens flare, spotlight, and burn/dodge effects of various types.

In the real world, lens flares occur when extremely bright light sources in the scene by the reflections are reflected off elements inside the lens of the camera. One might see lens flares in a shot when viewing a strong light source through a camera lens, like the sun or another bright star.

Inputs

There are three inputs on the Hot Spot node: one for the image, one for the effects mask, and another for an Occlusion image.

Input: The required orange input is used for the primary 2D image that gets the hot spot applied.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input restricts the hot spot to be within the pixels of the mask. An effects mask is applied to the tool after the tool is processed.

Occlusion: The green Occlusion input accepts an image to provide the occlusion matte. The matte is used to block the hot spot, causing it to “wink.” The white pixels in the image occlude the hot spot. Gray pixels partially suppress the hot spot.

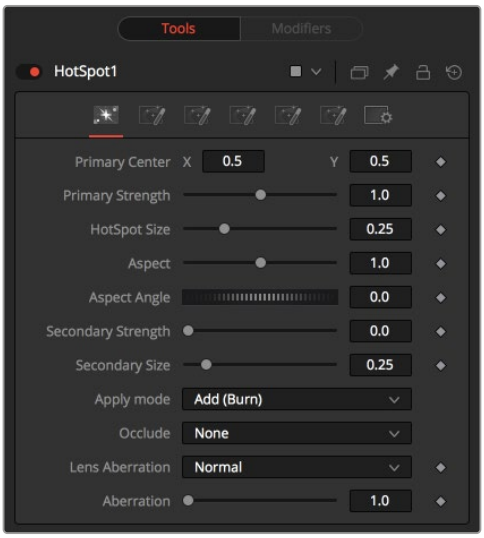
Basic Node Setup

The Hot Spot node is not a stand-alone generator, so it must have an image input that gets the hot spot applied.



Hot Spot node applied to an image

Inspector



Hot Spot controls

Hot Spot Tab

The Hotspot tab is used to control the primary and secondary hot spots. You can adjust their position, size, strength, angle, and apply mode.

Primary Center X and Y

This is the position of the primary hot spot within the scene. Secondary lens elements and reflections are positioned relative to the position of the primary hot spot.

Primary Strength

This control determines the brightness of the primary hot spot.

Hot Spot Size

This control determines the diameter of the primary hot spot. A value of 1.0 represents a circle the full width of the image.

Aspect

This controls the aspect of the spot. A value of 1.0 produces a perfectly circular hot spot. Values above 1.0 elongate the circle horizontally, and values below 1.0 elongate the circle vertically.

Aspect Angle

This control can be used to rotate the primary hot spot.

Secondary Strength

This control determines the strength, which is to say the brightness, of the secondary hot spot. The secondary hot spot is a reflection of the primary hot spot. It is always positioned on the opposite side of the image from the primary hot spot.

Secondary Size

This determines the size of the secondary hot spot.

Apply Mode

This control determines how the hot spot affects the underlying image.

- **Add (Burn):** This causes the spots created to brighten the image.
- **Subtract (Dodge):** This causes the spots created to dim the image.
- **Multiply (Spotlight):** This causes the spots created to isolate a portion of the image with light and to darken the remainder of the image.

Occlude

This menu is used to select which channel of the image connected to the Hot Spot node's Occlusion input is used to provide the occlusion matte. Occlusion can be controlled from Alpha or R, G, or B channels of any image connected to the Occlusion input on the node's tile.

Lens Aberration

Aberration changes the shape and behavior of the primary and secondary hot spots.

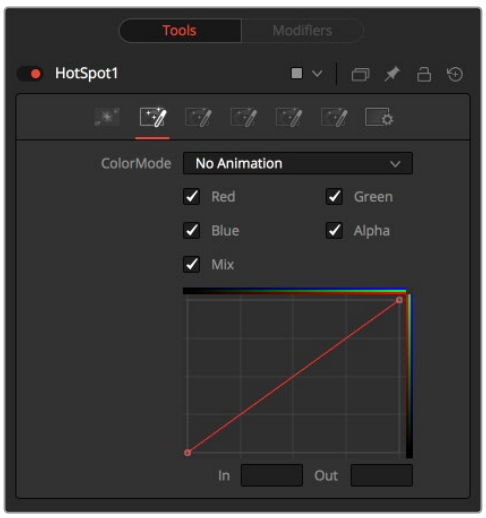
- **In and Out Modes:** Elongates the shape of the hot spot into a flare. The hot spot stretches toward the center when set to In mode and stretches toward the corners when set to Out mode.
- **Flare In and Flare Out Modes:** This option is a lens distortion effect that is controlled by the movement of the lens effect. Flare In causes the effect to become more severe, the closer the hot spot gets to the center. Flare Out causes the effect to increase as the hot spot gets closer to the edges of the image.
- **Lens:** This mode emulates a round, ringed lens effect.

Aberration

The Aberration slider controls the overall strength of the lens aberration effect.

Color Tab

The Color tab is used to modify the color of the primary and secondary hot spots.



Hot Spot color controls

Color Mode

This menu allows you to choose between animated or static color modifications using the small curves editor in the Inspector.

- **None:** The default None setting retains a static curve adjustment for the entire range.
- **Animated Points:** This setting allows the color curves in the spline area to be animated over time. Once this option is selected, moving to the desired frame and making a change in the Spline Editor sets a keyframe.
- **Dissolve mode:** Dissolve mode is mostly obsolete and is included for compatibility reasons only.

Color Channel and Mix

When selected, these checkboxes enable the editing of the chosen splines in the small Inspector Spline Editor. The Mix checkbox enables the Mix Spline, which is used to determine the influence of the controls that the Radial tab has along the radius of the hot spot.

Red, Green, Blue, and Alpha Splines

The Spline Window shows the curves for the individual channels. It is a miniature Spline Editor. The Red, Green, Blue, and Alpha splines are used to adjust the color of the spotlight along the radius of the hot spot.

The vertical axis represents the intensity or strength of the color channel. The horizontal axis represents the hot spot position along the radius, from the left outside edge to the inside right edge.

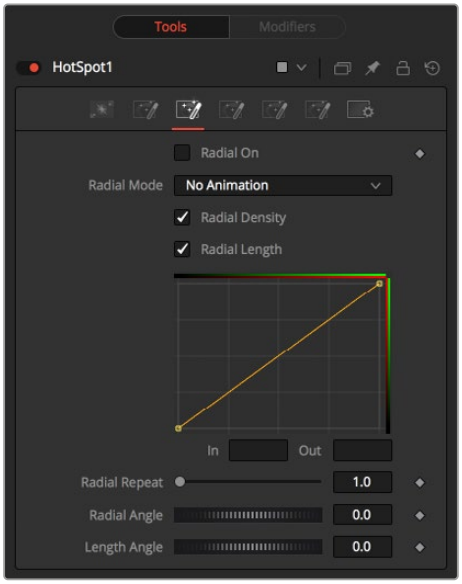
The default curve indicates that the red, green, blue, and Alpha channels all have a linear falloff.

Mix Spline

The Mix spline is used to determine the influence that the Radial controls have along the radius of the hot spot. The horizontal axis represents the position along the circle's circumference, with 0 being 0 degrees and 1.0 being 360 degrees. The vertical axis represents the amount of the radial hot spot to blend with the color hot spot. A value of 0 is all radial hot spot, while a value of 1.0 is all color hot spot.

NOTE: Right-clicking in the LUT displays a contextual menu with options related to modifying spline curves.

For more information on the LUT Editor, *see Chapter 7, "Using Viewers."* in the Fusion Reference Manual.



Hot Spot Radial tab

Radial Tab

Radial On

This control enables the Radial splines. Otherwise, the radial matte created by the splines is not applied to the hot spot, and the Mix spline in the color controls does not affect the hot spot.

Radial Mode

Similar to the Color mode menu, this menu allows you to choose between animated or static radial hot spot modifications using the small curves editor in the Inspector.

- **No Animation:** The default setting retains a static curve adjustment for the entire range.
- **Animated Points:** This setting allows the radial curves in the spline area to be animated over time. Once this option is selected, moving to the desired frame and making a change in the Spline Editor sets a keyframe.

The Interpolated Values option is mostly obsolete and is included for compatibility reasons only.

Radial Length and Radial Density Splines

The Spline window shows curves for the Length and Density of the hot spot. It is a miniature Spline Editor. The key to these splines is realizing that the horizontal axis in Inspector's Spline Editor represents a position around the circumference of the hot spot. A value of 0.0 is 0 degrees, and 1.0 is 360 degrees. With that in mind, the length determines the radius of light making up the hot spot along the circumference. The density represents how bright the light is along the circumference.

Radial Repeat

This control repeats the effect of the radial splines by x number of times. For example, a repeat of 2.0 causes the spline to take effect between 0 and 180 degrees instead of 0 and 360, repeating the spline between 180 and 360.

Length Angle

This control rotates the effect of the Radial Length spline around the circumference of the hot spot.

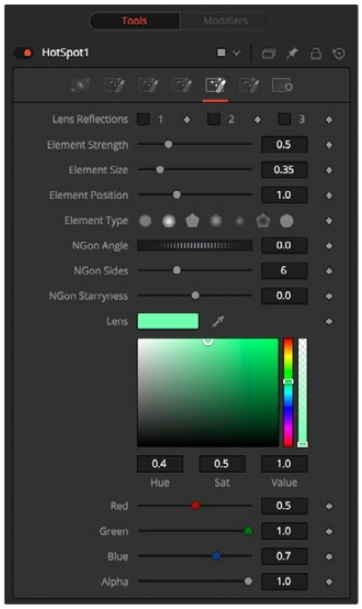
Density Angle

This control rotates the effect of the Radial Density spline around the circumference of the hot spot.

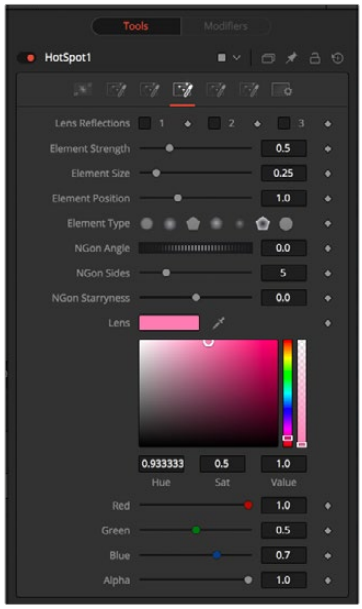
NOTE: Right-clicking in the spline area displays a contextual menu containing options related to modifying spline curves.

A complete description of LUT Editor controls and options, *see Chapter 47, "LUT Nodes,"* in the Fusion Reference Manual.

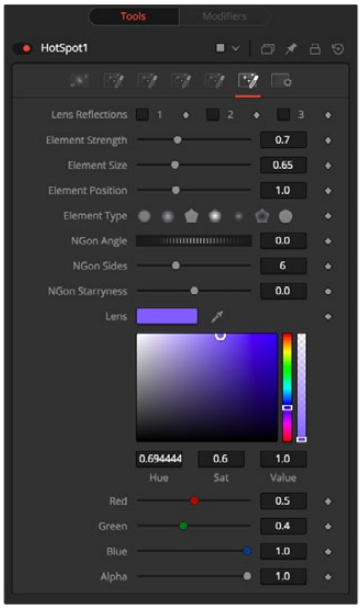
L1, L2, and L3 Tab



L1 tab



L2 tab



L3 tab

Lens Reflect Tabs

The three Lens Reflect tabs are used to enable and design additional lens flare elements beyond the primary and secondary hot spots.

Lens Reflect 1-3

Each of these three checkboxes enables a pair of lens reflection elements that you can modify using the controls in this tab. The parameters affect all the enabled Lens reflection elements in this tab.

Element Strength

This determines the brightness of element reflections.

Element Size

This determines the size of element reflections.

Element Position

This determines the distance of element reflections from the axis. The axis is calculated as a line between the hot spot position and the center of the image.

Element Type

Use this group of buttons to choose the shape and density of the element reflections. The presets available are described below.

- **Circular:** This creates slightly soft-edged circular shaped reflections.
- **Soft Circular:** This creates very soft-edged circular shaped reflections.
- **Circle:** This creates a hard-edged circle shape.
- **NGon Solid:** This creates a filled polygon with a variable number of sides.
- **NGon Star:** This creates a very soft-edged star shape with a variable number of sides.
- **NGon Shaded Out:** This creates soft-edged circular shapes.
- **NGon Shaded In:** This creates a polygon with a variable number of sides, which has a very soft reversed (dark center, bright radius) circle.
- **NGon Angle:** This control is used to determine the angle of the NGon shapes.
- **NGon Sides:** This control is used to determine the number of sides used when the Element Type is set to Ngon Star, Ngon Shaded Out, and Ngon Shaded In.
- **NGon Starriness:** This control is used to bend polygons into star shapes. The higher the value, the more star-like the shape.

Lens Color Controls

These controls determine the color of the lens that affects the colors of the reflections. To choose a lens color, pick one from a displayed image or enter RGBA values using the sliders or input boxes.

Common Controls

Settings Tab

The Settings tab controls are common to all Effect nodes, so their descriptions can be found in “The Common Controls” section at the end of this chapter.

Object Removal [ORm]



The Object Removal node

Object Removal Node Introduction

Object Removal uses the DaVinci Neural Engine to attempt to remove an object in the frame as automatically as possible. This node works best when removing a moving object that passes over a temporally stable background, or dirt on the lens of a shot where the camera is in motion. Smaller objects get better results than larger objects, but your results really depend on the footage.

Inputs

You can add masks and clean plates to the Object Removal node.

Input: The orange input is used for the primary footage you wish to remove an object from.

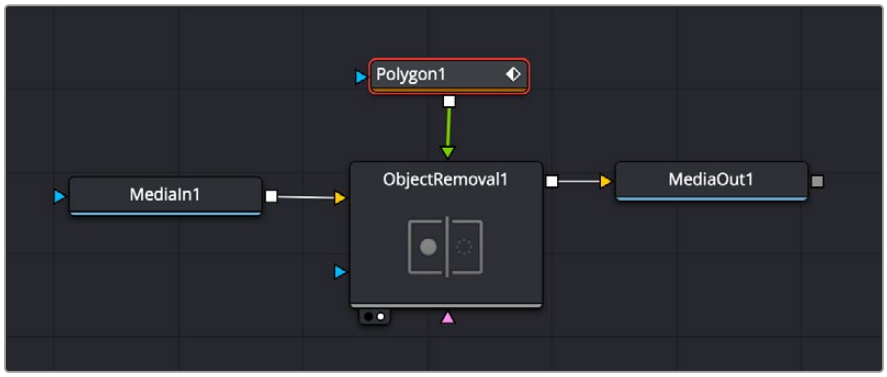
Clean Plate: The magenta input accepts a clean plate image source, if selected in the controls.

Mask: The green input accepts a mask shape (i.e., polygon, ellipse, rectangle, etc.) that is drawn around the object you want to remove from the scene.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the filter to only those pixels within the mask. An effects mask is applied to the tool after the tool is processed.

Basic Node Setup

Connect the media you wish to remove the object from to the yellow Input, and then either a polygon shape to the green Mask input, or a frame without the object to the clean plate input.



The Object Removal tool node structure. You will need either a polygon or another shape attached to the green Mask input, or a clean plate attached to the pink clean plate input.

To setup the Object Removal:

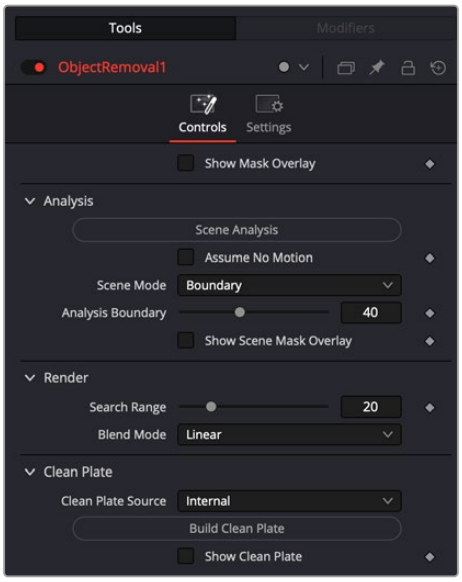
1. Draw a polygon around the object that you want to remove from the scene in the Polygon node.
2. Track the polygon either manually, or using a tracker, to make sure that the object you want to remove is always inside its boundaries through the duration of the clip.
3. Press the Scene Analysis button in the Object Removal tool controls.

These steps should accommodate most simple object removal tasks, but if there are issues, you can adjust the parameters of the tool in the Inspector. The Object Removal tool is not appropriate for all scenes and works best with small objects that move through the frame, not in front of or behind other objects.



Drawing and tracking a polygon around the cyclist, then pressing the Scene Analysis button in the Object Removal controls, removes the cyclist from the background.

Inspector



The Object Removal controls help you remove unwanted objects from an image.

Controls

The Object Removal controls let you adjust the parameters of the object analysis.

- **Show Mask Overlay:** Lets you see the actual mask used, highlighted in red, for the Object Removal.

Analysis

- **Scene Analysis:** Press this button to analyze the clip using the selected parameters and replace the object with the appropriate background. You will need to press this button again if you change any of the parameters in this section.
- **Assume No Motion:** If the camera has been locked off for the shot while the object moves through it, check this box to greatly simplify the background analysis.
- **Scene Mode:** Provides different methods of analyzing the scene, improving the analysis of how the area that needs to be replaced moves to best determine how to fill the hole left by the object being removed.

  Background: Analyzes the entire image except for the object region.

  Boundary: Analyzes the boundary area surrounding the object region.

  Object: For analyzing an object that moves with the background, like a sticker that's on a window while the camera moves.
- **Analysis Boundary:** Sets the object's boundary size for analysis.
- **Show Scene Mask Overlay:** Toggles the scene mask overlay on or off, letting you see what's behind the mask.

Render

- **Search Range:** If you notice, while playing through the analyzed result, that the object removal mask has gray fringing on some frames, you can try adjusting the "Search Range" slider, which is the distance, in frames, from the current frame that the Object Removal plugin is searching for replacement image detail. For example, if the Search Range is 20, it searches +/-20 frames from the current location, or 40 frames total. The allowance of 10 frames means we look at every fourth frame. You will generally get the best results for the smallest range that gives an acceptable result.
- **Blend Mode:** If the patch is successfully filled, but the result isn't blending well with the background, you can try changing the Blend mode. The default is Linear, which is a simple cloning operation, but you can also choose Adaptive Blend, which can provide better results except in certain situations where the edges of the replacement patch have a different color or brightness than the background.

Clean Plate

- **Clean Plate Source:** Lets you set the source of the clean plate. Options are Gray Image, which is effectively no source; Internal, which takes a "best guess" approach to generating a background with which to fill the frame and integrates this with frames that could be successfully filled in; and External, which lets you connect your own plate to the pink input from another node.
- **Build Clean Plate:** Executes the building of a new clean plate background based on the Clean Plate Source parameter set.
- **Show Clean Plate:** Checking this box lets you see behind the mask to the clean plate behind it.

Pseudo Color [PsCl]



The Pseudo Color node

Pseudo Color Node Introduction

The Pseudo Color node provides the ability to produce variations of an image's color based on waveforms generated by the node's controls. Static or animated variances of the original image can be produced.

Inputs

There are two Inputs on the Pseudo Color node: one for an image and one for an effects mask.

Input: The orange input is used for the primary 2D image that gets its color modified.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input restricts the pseudo color to be within the pixels of the mask. An effects mask is applied to the tool after the tool is processed.

Basic Node Setup

The Pseudo Color node is not a stand-alone generator, so it must have an image input that it uses to generate variations in colors.



Pseudo Color node applied to an image

Inspector



Pseudo Color RGBA controls

Red/Green/Blue/Alpha Tabs

The node's controls are separated into four identical tabs, one for each of the RGBA color channels.

Color Checkbox

When enabled, the Pseudo Color node affects this color channel.

Wrap

When enabled, waveform values that exceed allowable parameter values are wrapped to the opposite extreme.

High and Low

High and Low determine the range to be affected by the node in a specific color channel.

Soft Edge

This slider determines the soft edge of color transition.

Waveform

This selects the type of waveform to be created by the generator. Four waveforms are available: Sine, Triangle, Sawtooth, and Square.

Frequency

This controls the frequency of the waveform selected. Higher values increase the number of occurrences of the variances.

Phase

This modifies the Phase of the waveform. Animating this control produces color cycling effects.

Mean

This determines the level of the waveform selected. Higher values increase the overall brightness of the channel until the allowed maximum is reached.

Amplitude

Amplitude increases or decreases the overall power of the waveform.

Common Controls

Settings Tab

The Settings tab controls are common to all Effect nodes, so their descriptions can be found in "The Common Controls" section at the end of this chapter.

Rays [CIR]



The Rays node

Rays Node Introduction

Rays is a modified zoom blur effect that radiates through an object from a specified point.

Inputs

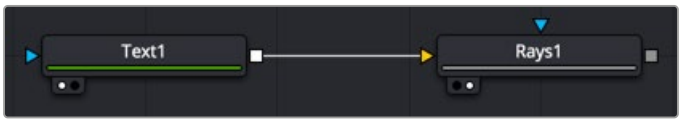
There are two inputs on the Rays node: one for the image and one for the effects mask.

Input: The orange input is used for the primary 2D image that gets the rays applied to it.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input restricts the rays to be within the pixels of the mask. An effects mask is applied to the tool after the tool is processed.

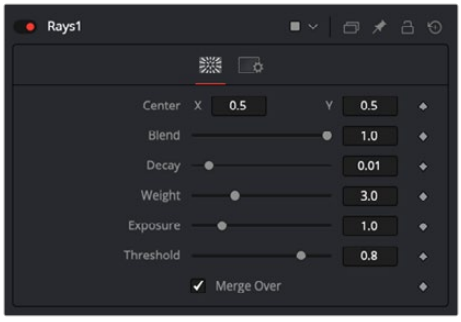
Basic Node Setup

The Rays node works best when the image or graphic connected to the orange input includes an Alpha channel from which the rays emit.



Rays node set up to emit from a line of text

Inspector



Rays node controls

Controls Tab

The Controls tab contains all the primary controls necessary for customizing the rays.

Center X and Y

This coordinate control and related viewer crosshair set the center point for the light source.

Blend

Sets the percentage of the original image that's blended with the light rays.

Decay

Sets the length of the light rays.

Weight

Sets the falloff of the light rays.

Exposure

Sets the intensity level of the light rays.

Threshold

Sets the luminance limit at which the light rays are produced.

Common Controls

Settings Tab

The Settings tab controls are common to all Effect nodes, so their descriptions can be found in the "The Common Controls" section at the end of this chapter.

Shadow [SH]



The Shadow node

Shadow Node Introduction

Shadow is a versatile node used in the creation of a drop shadow, based on the Alpha channel in an image. Optionally, a second image can be used as a depth matte to distort the shadow based on the varying depth in a background image.

Input

The three inputs on the Shadow node are used to connect a 2D image that causes the shadow. A depth map input and an effect mask can be used to limit the area where trails appear. Typically, the output of the shadow is then merged over the actual background in the composite.

Input: The orange input is used for the primary 2D image with Alpha channel that is the source of the shadow.

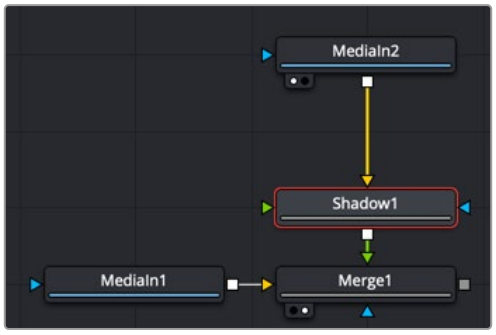
Depth: The green Depth map input takes a 2D image as its input and extracts a depth matte from a selected channel. The light Position and Distance controls can then be used to modify the appearance of the shadow based on depth.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the area where the shadow appears. An effects mask is applied to the tool after the tool is processed.

NOTE: The Shadow node is designed to create simple 2D drop shadows. Use a Spot Light node and an Image Plane 3D node for full 3D shadow casting.

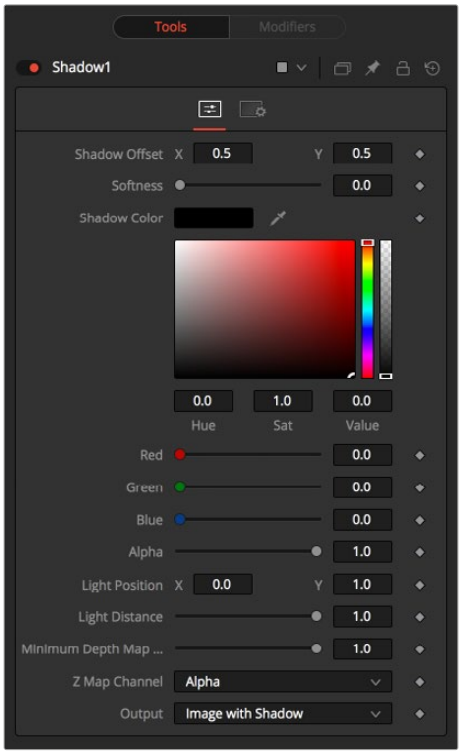
Basic Node Setup

Below, the Shadow node uses the output of an image with Alpha and connects to the foreground of a Merge. The shadow is shown over the background input to the Merge.



A Shadow node is generated from the MediaIn2 and shown over MediaIn1

Inspector



Shadow node controls

Controls Tab

The Controls tab contains all the primary controls necessary for customizing the shadow appearance.

Shadow Offset

This control sets the X and Y position of the shadow. When the Shadow node is selected, you can also adjust the position of the Shadow Offset using the crosshair in the viewer.

Softness

Softness controls how blurry the shadow's edges appear.

Shadow Color

Use this control to select the color of the shadow. The most realistic shadows are usually not totally black and razor sharp.

Light Position

This control sets the position of the light relative to the shadow-casting object. The Light Position is only taken into consideration when the Light Distance slider is not set to infinity (1.0).

Light Distance

This slider varies the apparent distance of the light between infinity (1.0) and zero distance from the shadow-casting object. The advantage of setting the Light Distance is that the resulting shadow is more realistic-looking, with the further parts of the shadow being longer than those that are closer.

Minimum Depth Map Light Distance

This control is active when an image is connected to the shadow's Depth Map input. The slider is used to control the amount that the depth map contributes to the Light Distance. Dark areas of a depth map make the shadow deeper. White areas bring it closer to the camera.

Z Map Channel

This menu is used to select which color channel of the image connected to the node's Depth Map input is used to create the shadow's depth map. Selections exist for the RGB and A, Luminance, and Z-buffer channels.

Output

This menu determines if the output image contains the image with shadow applied or the shadow only.

The shadow only method is useful when color correction, perspective, or other effects need to be applied to the resulting shadow before it is merged back with the object.

Common Controls

Settings Tab

The Settings tab controls are common to all Effect nodes, so their descriptions can be found in "The Common Controls" section at the end of this chapter.

Trails [TRLS]



The Trails node

Trails Node Introduction

The Trails node is used to create a ghost-like after-trail of the image. This creates an interesting effect when applied to moving images with an Alpha channel. Unlike a directional blur, only the preceding motion of an image is displayed as part of the effect. Since the trail effect is based on an image buffer, it requires you to play or activate the pre-roll for some number of frames before you see the effect.

Input

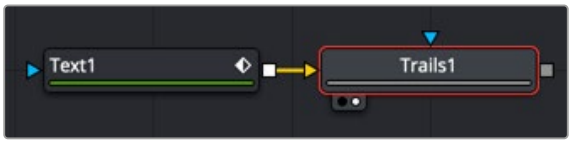
The two inputs on the Trails node are used to connect a 2D image and an effect mask that can be used to limit the area where trails appear.

Input: The orange input is used for the primary 2D image that receives the trails applied.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the area where the trails effect appears. An effects mask is applied to the tool after the tool is processed.

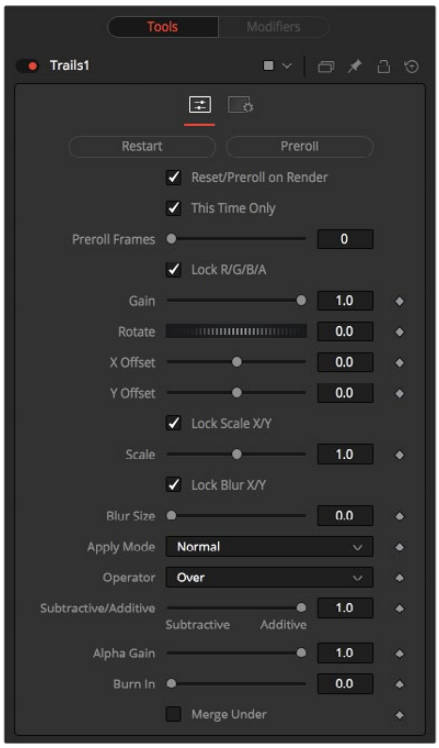
Basic Node Setup

The output of an animated Text node is connected to the input of the Trails node. Trails are generated based on the motion of the text. The Reset button must be pressed in the Inspector between each preview, or the trails will accumulate.



A Trails node generates trails for the animation in the Text node

Inspector



Trails node controls

Controls Tab

The Controls tab contains all the primary controls necessary for customizing the trails.

Restart

This control clears the image buffer and displays a clean frame, without any of the ghosting effects.

Preroll

This makes the Trails node pre-render the effect by the number of frames on the slider.

Reset/Preroll on Render

When this checkbox is enabled, the Trails node resets itself when a preview or final render is initiated. It pre-rolls the designated number of frames.

This Time Only

Selecting this checkbox makes the pre-roll use this current frame only and not the previous frames.

Preroll Frames

This determines the number of frames to pre-roll.

Lock RGBA

When selected, this checkbox allows the Gain of the color channels to be controlled independently. This allows for tinting of the Trails effect.

Gain

The Gain control affects the overall intensity and brightness of the image in the buffer. Lower values in this parameter create a much shorter, fainter trail, whereas higher values create a longer, more solid trail.

Rotate

The Rotate control rotates the image in the buffer before the current frame is merged into the effect. The offset is compounded between each element of the trail. This is different than each element of the trail rotating on its pivot point. The pivot remains over the original object.

Offset X/Y

These controls offset the image in the buffer before the current frame is merged into the effect. Control is given over each axis independently. The offset is compounded between each element of the trail.

Lock Scale X/Y

When selected, this checkbox allows the X- and Y-axis scaling of the image buffer to be manipulated separately for each axis.

Scale

The Scale control resizes the image in the buffer before the current frame is merged into the effect. The size is compounded between each element of the trail.

Lock Blur X/Y

When selected, this checkbox allows the blurring of the image buffer to be controlled separately for each axis.

Blur Size

The Blur Size control applies a blur to the trails in the buffer before the current frame is merged into the effect. The blur is compounded between each element of the trail.

Apply Mode

The Apply Mode setting determines the math used when blending or combining the trailing objects that overlap.

- **Normal:** The default mode uses the foreground object's Alpha channel as a mask to determine which pixels are transparent and which are not. When this is active, another menu shows possible operations, including Over, In, Held Out, Atop, and XOr.
- **Screen:** Screen blends the objects based on a multiplication of their color values. The Alpha channel is ignored, and layer order becomes irrelevant. The resulting color is always lighter. Screening with black leaves the color unchanged, whereas screening with white always produces white. This effect creates a similar look to projecting several film frames onto the same surface. When this is active, another menu shows possible operations, including Over, In, Held Out, Atop, and XOr.
- **Dissolve:** Dissolve mixes overlapping objects. It uses a calculated average of the objects to perform the mixture.
- **Multiply:** Multiplies the values of a color channel. This gives the appearance of darkening the object as the values are scaled from 0 to 1. White has a value of 1, so the result would be the same. Gray has a value of 0.5, so the result would be a darker object or, in other words, an object half as bright.
- **Overlay:** Overlay multiplies or screens the color values of the foreground object, depending on the color values of the background object. Patterns or colors overlay the existing pixels while preserving the highlights and shadows of the color values of the objects behind the foreground objects. The objects behind the foreground objects are not replaced but mixed with the foreground objects to reflect the original lightness or darkness of the background objects.
- **Soft Light:** Soft Light darkens or lightens the foreground object, depending on the color values of the objects behind them. The effect is similar to shining a diffused spotlight on the image.

- **Hard Light:** Hard Light multiplies or screens the color values of the foreground object, depending on the color values of the objects behind them. The effect is similar to shining a harsh spotlight on the image.
- **Color Dodge:** Color Dodge uses the foreground object's color values to brighten the objects behind them. This is similar to the photographic practice of dodging by reducing the exposure of an area of a print.
- **Color Burn:** Color Burn uses the foreground object's color values to darken the objects behind them. This is similar to the photographic practice of burning by increasing the exposure of an area of a print.
- **Darken:** Darken looks at the color information in each channel and selects the color value from the object in front or behind, whichever is darker. Pixels lighter than the blended colors are replaced, and pixels darker than the blended color do not change.
- **Lighten:** Lighten looks at the color information in each channel and selects the color value from the object in front or behind, whichever is lighter. Pixels darker than the blended color are replaced, and pixels lighter than the blended color do not change.
- **Difference:** Difference looks at the color information in each channel and subtracts the foreground object's color values from the background object's color values or vice versa, depending on which has the higher brightness value. Blending with white inverts the color. Blending with black produces no change.
- **Exclusion:** Exclusion creates an effect similar to but lower in contrast than the Difference mode. Blending with white inverts the base color values. Blending with black produces no change.
- **Hue:** Hue creates color with the luminance and saturation of the background object's color and the hue of the foreground object's color.
- **Saturation:** Saturation creates color with the luminance and hue of the base color and the saturation of the blend color.
- **Color:** Color creates color with the luminance of the background object's color and the hue and saturation of the object in front. This preserves the gray levels in the image and is useful for colorizing monochrome objects.
- **Luminosity:** Luminosity creates color with the hue and saturation of the background object's color and the luminance of the foreground object's color. This mode creates an inverse effect from that of the Color mode.

Operator

This menu is used to select the Operation mode used when the trailing objects overlap. Changing the Operation mode changes how the overlapping objects are combined to produce a result. This drop-down menu is visible only when the Apply mode is set to Normal.

The formula used to combine pixels in the trails node is always (fg object * x) + (bg object * y). The different operations determine what x and y are, as shown in the description for each mode.

The Operator Modes are as follows:

- **Over:** The Over mode adds the foreground object to the background object by replacing the pixels in the background with the pixels from the Z wherever the foreground object's Alpha channel is greater than 1.

```
x = 1, y = 1 - [foreground object Alpha]
```

- **In:** The In mode multiplies the Alpha channel of the background object against the pixels in the foreground object. The color channels of the foreground object are ignored. Only pixels from the foreground object are seen in the final output. This essentially clips the foreground object using the mask from the background object.

```
x = [background Alpha], y = 0
```

- **Held Out:** Held Out is essentially the opposite of the In operation. The pixels in the foreground object are multiplied against the inverted Alpha channel of the background object.

  ```
  x = 1 - [background Alpha], y = 0
  ```

- **Atop:** Atop places the foreground object over the background object only where the background object has a matte.

  ```
  x = [background Alpha], y = 1 - [foreground Alpha]
  ```

- **XOr:** XOr combines the foreground object with the background object wherever either the foreground or the background have a matte, but never where both have a matte.

  ```
  x = 1 - [background Alpha], y = 1 - [foreground Alpha]
  ```

Subtractive/Additive

This slider controls whether Fusion performs an Additive composite, a Subtractive composite, or a blend of both when the trailing objects overlap. This slider defaults to Additive assuming the input image's Alpha channel is premultiplied (which is usually the case). If you don't understand the difference between Additive and Subtractive compositing, below is a quick explanation.

> **NOTE:** An Additive blend operation is necessary when the foreground image is premultiplied, meaning that the pixels in the color channels have been multiplied by the pixels in the Alpha channel. The result is that transparent pixels are always black since any number multiplied by 0 always equals 0. This obscures the background (by multiplying with the inverse of the foreground Alpha), and then adds the pixels from the foreground.
>
> A Subtractive blend operation is necessary if the foreground image is not premultiplied. The compositing method is similar to an additive composite, but the foreground image is first multiplied by its Alpha, to eliminate any background pixels outside the Alpha area.
>
> Although the Additive/Subtractive option is often an either/or checkbox in other software, the Trails node lets you blend between the Additive and Subtractive versions of the compositing operation. This can be useful when dealing with problem edges that are too bright or too dark.

For example, using Subtractive merging on a premultiplied image may result in darker edges, whereas using Additive merging with a non-premultiplied image causes any non-black area outside the foreground's Alpha to be added to the result, thereby lightening the edges. By blending between Additive and Subtractive, you can tweak the edge brightness to be just right for your situation.

Alpha Gain

Alpha Gain linearly scales the Alpha channel values of the trailing objects in front. This effectively reduces the amount that the trailing objects in the background are obscured, thus brightening the overall result. When the Subtractive/Additive slider is set to Additive with Alpha Gain set to 0.0, the foreground pixels are added to the background.

When the Subtractive/Additive slider is set to Subtractive, this controls the density of the composite, similar to Blend.

Burn In

The Burn In control adjusts the amount of Alpha used to darken the objects that trail under other objects, without affecting the amount of foreground objects added. At 0.0, the blending behaves like a straight Alpha blend. At 1.0, the objects in the front are effectively added onto the objects in the back (after Alpha multiplication if in Subtractive mode). This gives the effect of the foreground objects

brightening the objects in the back, as with Alpha Gain. In fact, for Additive blends, increasing the Burn In gives an identical result to decreasing Alpha Gain.

Merge Under

When enabled, the current image is placed under the generated trail, rather than the usual, over top operation. The layer order of the trailing elements is also reversed, making the last trail the topmost layer.

Common Controls

Settings Tab

The Settings tab controls are common to all Effect nodes, so their descriptions can be found in “The Common Controls” section at the end of this chapter.

TV [TV]



The TV node

TV Node Introduction

The TV node is a simple node designed to mimic some of the typical flaws seen in analog television broadcasts and screens. This Fusion-specific node is mostly obsolete when using DaVinci Resolve because of the more advanced Analog Damage ResolveFX.

Input

The two inputs on the TV node are used to connect a 2D image and an effect mask, which can be used to limit the area where the TV effect appears.

Input: The orange input is used for the primary 2D image that gets the TV distortion applied.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the area where the TV effect to appears. An effects mask is applied to the tool after the tool is processed.

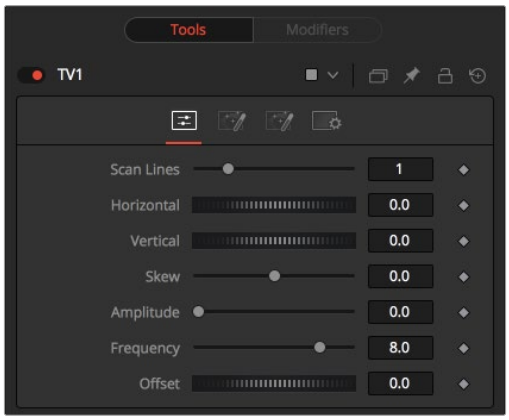
Basic Node Setup

The output of an image is connected to the input of the TV node. The style of TV interference is then customized using the controls in the Inspector.



The TV node simulates TV-style flaws in the image connected to the orange input

Inspector



TV node controls

Controls Tab

The Controls tab is the first of three tabs used to customize the analog TV distortion. The Controls tab modifies the scan lines and image distortion of the effect.

Scan Lines

This slider is used to emulate the interlaced look by dropping lines out of the image. Setting it to black, with a transparent Alpha, drops a line. A value of 1 (default) drops every second line. A value of 2 shows one line, and then drops the second and third and repeats. A value of zero turns off the effect.

Horizontal

Use this slider to apply a simple Horizontal offset to the image.

Vertical

Use this slider to apply a simple Vertical offset to the image.

Skew

This slider is used to apply a diagonal offset to the image. Positive values skew the image to the top left. Negative values skew the image to the top right. Pixels pushed off frame wrap around and reappear on the other side of the image.

Amplitude

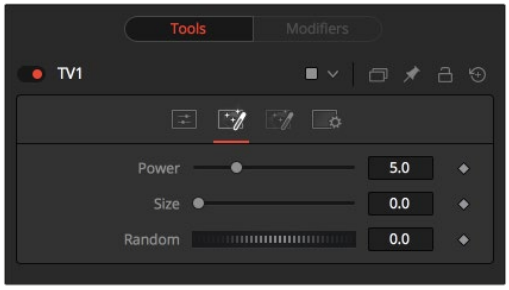
The Amplitude slider can be used to introduce smooth sine wave-type deformation to the edges of the image. Higher values increase the intensity of the deformation. Use the Frequency control to determine how often the distortion is repeated.

Frequency

The Frequency slider sets the frequency of the sine wave used to produce distortion along the edges of the image when the amplitude control is greater than 1.

Offset

Use Offset to adjust the position of the sine wave, causing the deformation applied to the image via the Amplitude and Frequency controls to see across the image.



The TV Noise tab

Noise Tab

The Noise tab is the second of three tabs used to customize the analog TV distortion. The Noise tab modifies the noise in the image to simulate a weak analog antenna signal.

Power

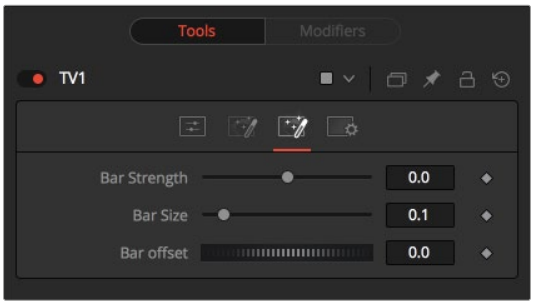
Increase the value of this slider above 0 to introduce noise into the image. The higher the value, the stronger the noise.

Size

Use this slider to scale the noise map larger.

Random

If this thumbwheel control is set to 0, the noise map is static. Change the value over time to cause the static to change from frame to frame.



The TV Roll Bar tab

Roll Bar Tab

The Roll Bar tab is the third of three tabs used to customize the analog TV distortion. The Roll Bar tab animates the bar.

Bar Strength

At the default value of 0, no bar is drawn. The higher the value, the darker the area covered by the bar becomes.

Bar Size

Increase the value of this slider to make the bar taller.

Bar Offset

Animate this control to scroll the bar across the screen.

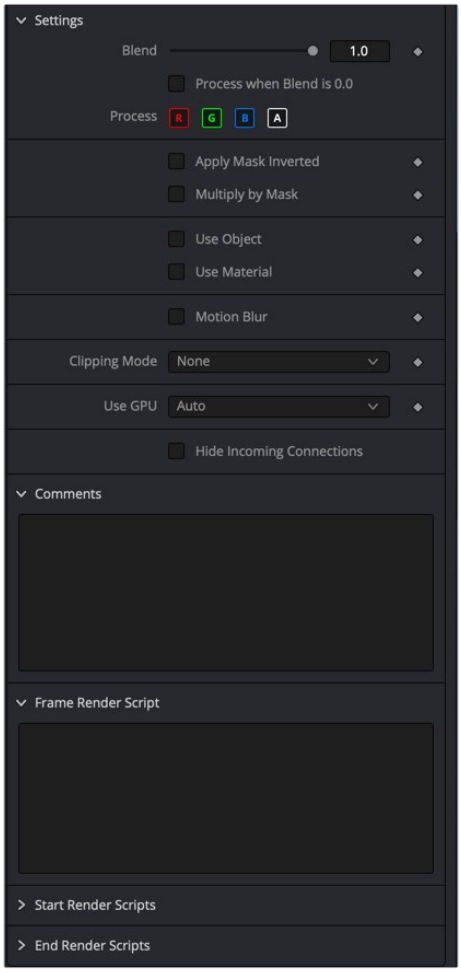
Settings Tab

The Settings tab controls are common to all Effect nodes, so their descriptions can be found in the following “The Common Controls” section.

The Common Controls

Effect nodes share several identical controls in the Inspector. This section describes controls that are common among Effect nodes.

Inspector



The Common Effects Settings tab

Settings Tab

The Settings tab in the Inspector can be found on every tool in the Effects category. The Settings controls are even found on third-party Effects-type plugin tools. The controls are consistent and work the same way for each tool, although some tools do include one or two individual options, which are also covered here.

Blend

The Blend control is used to blend between the tool's original image input and the tool's final modified output image. When the blend value is 0.0, the outgoing image is identical to the incoming image. This causes the tool to skip processing entirely, copying the input straight to the output.

Process When Blend Is 0.0

The tool is processed even when the input value is zero. This is useful when this node is scripted to trigger another task, but the blend is set to 0.0.

Red/Green/Blue/Alpha Channel Selector

These four buttons are used to limit the effect of the tool to specified color channels. This filter is often applied after the tool has been processed.

For example, if the red button on a Blur tool is deselected, the blur is first applied to the image, and then the red channel from the original input is copied back over the red channel of the result.

There are some exceptions, such as tools for which deselecting these channels causes the tool to skip processing that channel entirely. Tools that do this possess a set of like RGBA buttons on the Controls tab in the tool. In this case, the buttons in the Settings and the Control tabs are identical.

Apply Mask Inverted

Enabling the Apply Mask Inverted option inverts the complete mask channel for the tool. The mask channel is the combined result of all masks connected to or generated in a node.

Multiply by Mask

Selecting this option causes the RGB values of the masked image to be multiplied by the mask channel's values. This causes all pixels of the image not included in the mask (i.e., set to 0) to become black/transparent.

Use Object/Use Material (Checkboxes)

Some 3D software can render to file formats that support additional channels. Notably, the EXR file format supports Object and Material ID channels, which can be used as a mask for the effect. These checkboxes determine whether the channels are used, if present. The specific Material ID or Object ID affected is chosen using the next set of controls.

Correct Edges

This checkbox appears only when the Use Object or Use Material checkboxes are selected. It toggles the method used to deal with overlapping edges of objects in a multi-object image. When enabled, the Coverage and Background Color channels are used to separate and improve the effect around the edge of the object. If this option is disabled (or no Coverage or Background Color channels are available), aliasing may occur on the edge of the mask.

> For more information on coverage and background channels, *see Chapter 18, "Understanding Image Channels,"* in the Fusion Reference Manual.

Object ID/Material ID (Sliders)

Use these sliders to select which ID is used to create a mask from the object or material channels of an image. Use the Sample button in the same way as the Color Picker: to grab IDs from the image displayed in the view. The image or sequence must have been rendered from a 3D software package with those channels included.

Clipping Mode

This option determines how the domain of definition rendering handles edges. The Clipping mode is most important when blur or softness is applied, which may require samples from portions of the image outside the current domain.

- **Frame:** The default option is Frame, which automatically sets the node's domain of definition to use the full frame of the image, effectively ignoring the current domain of definition. If the upstream DoD is smaller than the frame, the remaining area in the frame is treated as black/transparent.
- **None:** Setting this option to None does not perform any source image clipping. Any data required to process the node's effect that would usually be outside the upstream DoD is treated as black/transparent.

Use GPU

The Use GPU menu has three settings. Setting the menu to Disable turns off hardware-accelerated rendering using the graphics card in your computer. Enabled uses the hardware. Auto uses a capable GPU if one is available and falls back to software rendering when a capable GPU is not available

Motion Blur

- **Motion Blur:** This toggles the rendering of Motion Blur on the tool. When this control is toggled on, the tool's predicted motion is used to produce the motion blur caused by the virtual camera's shutter. When the control is toggled off, no motion blur is created.
- **Quality:** Quality determines the number of samples used to create the blur. A quality setting of 2 causes Fusion to create two samples to either side of an object's actual motion. Larger values produce smoother results but increase the render time.
- **Shutter Angle:** Shutter Angle controls the angle of the virtual shutter used to produce the motion blur effect. Larger angles create more blur but increase the render times. A value of 360 is the equivalent of having the shutter open for one full frame exposure. Higher values are possible and can be used to create interesting effects.
- **Center Bias:** Center Bias modifies the position of the center of the motion blur. This allows for the creation of motion trail effects.
- **Sample Spread:** Adjusting this control modifies the weighting given to each sample. This affects the brightness of the samples.

Comments

The Comments field is used to add notes to a tool. Click in the empty field and type the text. When a note is added to a tool, a small red square appears in the lower-left corner of the node when the full tile is displayed, or a small text bubble icon appears on the right when nodes are collapsed. To see the note in the Node Editor, hold the mouse pointer over the node to display the tooltip.

Scripts

Three Scripting fields are available on every tool in Fusion from the Settings tab. They each contain edit boxes used to add scripts that process when the tool is rendering. For more details on scripting nodes, please consult the Fusion scripting documentation.

Chapter 39

Film Nodes

This chapter details the Film nodes in Fusion.

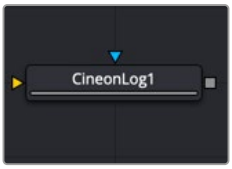
The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Cineon Log [Log] 1085
Film Grain [FGr] 1089
Grain [Grn] 1093
Light Trim [LT] 1096
Remove Noise [RN] 1098
The Common Controls 1100

Cineon Log [Log]



The Cineon Log node

Cineon Log Node Introduction

The Cineon Log node is used to convert several different log camera formats to linear gamma and back again. Although the name implies that it should be used with Cineon files, it handles “log” gamma from many different digital cinema sources such as Blackmagic Design, Arri, and Red cameras.

Input

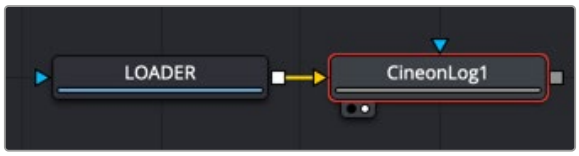
There are two Inputs on the Cineon Log node: one for the log image and one for the effects mask.

Input: The orange input is used for the primary 2D image that gets the highlight applied.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input restricts the log conversion to be within the pixels of the mask. An effects mask is applied to the tool after the tool is processed.

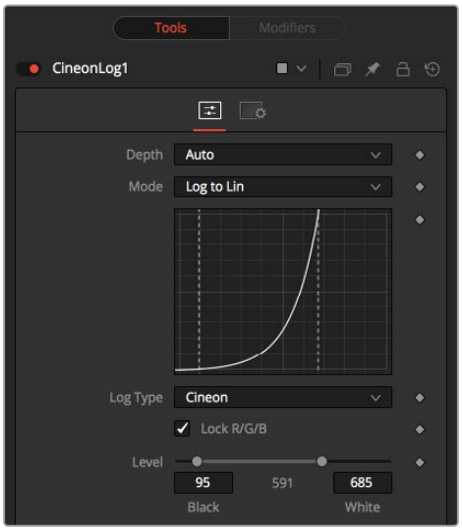
Basic Node Setup

The Cineon Log node is placed directly after a MediaIn node in DaVinci Resolve or a Loader node in Fusion Studio. It is also commonly placed before a MediaOut or Saver node to convert back to a Log-encoded image.



A Cineon Log node placed after a Loader node in Fusion Studio

Inspector



Cineon Log controls

Controls Tab

The Controls tab includes settings for converting from log gamma to linear or from linear to log. You first select the Mode and then the Log Type. For instance, choose Log to Lin from the Mode menu, and then select BMD Film if you are compositing with a RAW clip from a Blackmagic Design camera. Those settings output a linear image ready for compositing.

Depth

The Depth menu is used to select the color depth used to process the input image. The default option is Auto. Auto determines the color depth based on the file format loaded. For example, JPEG files automatically process at 8 bit because the JPEG file format does not store color depths greater than 8. Blackmagic RAW files load at Float, etc. If the color depth of the format is undetermined, the default depth defined in the Frame Format preferences is used.

Mode

The Mode menu offers two options: one for converting log images to linear and one for converting linear images to logarithmic.

Log Type

The Log Type menu allows you to select the source of the file. Typically, you select the camera used to create the image, although the Josh Pines option is specific to film scan workflows. This menu contains the following camera log types:

- Cineon
- Arri Log C
- BMD Film
- Canon Log
- Nikon N Log
- Panalog
- Panasonic V-Log
- Red Log Film
- Sony S-Log
- Viper Film Stream
- ACESlog

Lock RGB

When enabled, the settings in this tab affect all color channels equally.

Disable this control to convert the red, green, and blue channels of the image using separate settings for each channel.

Level

Use this range control to set the black level and white level in the log image before converting. The left handle adjusts the black level, while the right handle adjusts the white level. Pixels with values in log space below the black level become out-of-range values below 0.0. Pixels with values above the white level become out-of-range values above 1.0 after conversion.

When processing in floating-point color space, both negative and high out-of-range values are preserved. When using 16-bit or 8-bit mode, the out-of-range values are clipped.

Soft Clip (Knee)

The Soft Clip control is used to draw values that are out of range back into the image. This is done by smoothing the conversion curve at the top and bottom of the curve, allowing more values to be represented.

Applying a soft clip of any value other than 1 causes the node to process at 16-bit integer, eliminating all out-of-range values that do not fit within the soft clip.

Film Stock Gamma, Conversion Gamma, and Conversion Table

These controls are used to set the response curves of the logarithmic data during conversion. In addition to the settings above, a custom ASCII file Lookup Table (LUT) can be created with specific conversion values. The ASCII LUT file can be loaded using the Browse button.

Common Controls

Settings Tab

The Settings tab controls are common to all Film nodes, so their descriptions can be found in "The Common Controls" section at the end of this chapter.

Black Rolloff

Since a mathematical log() operation on a value of zero or lower results in invalid values, Fusion clips values below 1e-38 (0 followed by 38 zeros) to 0 to ensure correct results. This is almost never an issue, since values that small have no visual impact on an image. To see such tiny values, you would have to add three Brightness Contrast nodes, each with a gain set to 1,000,000. Even then, the values would hover very close to zero.

We have seen processes where instead of cropping these minimal values, they are instead scaled. So values between 0.0 and 1e-16 are scaled between 1e-18 and 1e-16. The idea is to crush the majority of the visual range in a float image into values very near to zero, then expand them again, forcing a gentle ramp to produce a small ramp in the extreme black values. Should you find yourself facing a color pipeline using this process, here is how you can mimic it with the help of a Custom node.

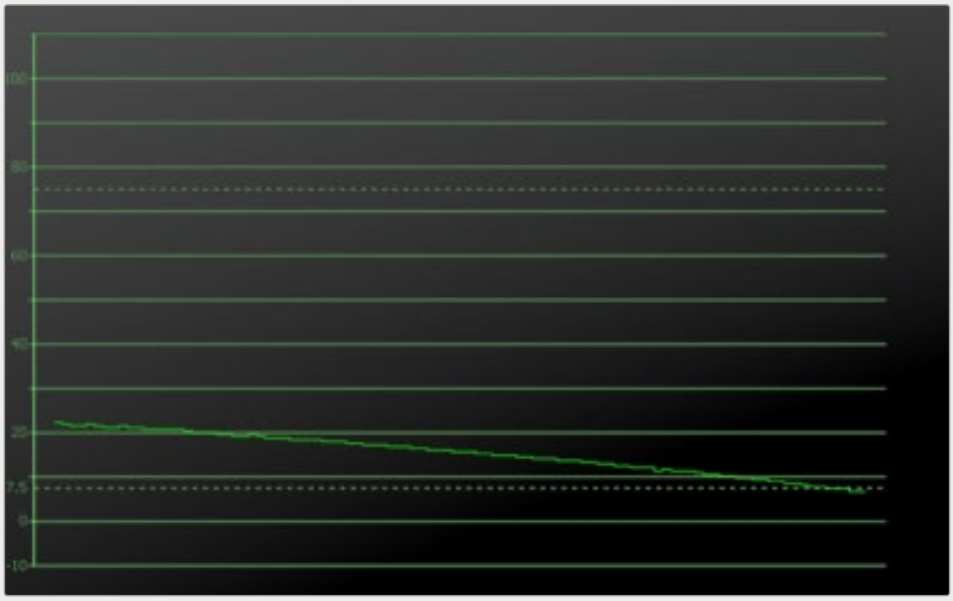
The process involves converting the log image to linear with a very small gamma and a wider than normal black level to white level (e.g., conversion gamma of 0.6, black of 10, white of 1010). This crushes most of the image's range into very small values. This is followed by a Custom node (described below), and then by a linear to log conversion that reverses the process but uses a slightly higher black level. The difference between the black levels defines the falloff range.

Since this lifts the blacks, the image is usually then converted back to linear one more time, using more traditional values (i.e., 95-685) to reset the black point.

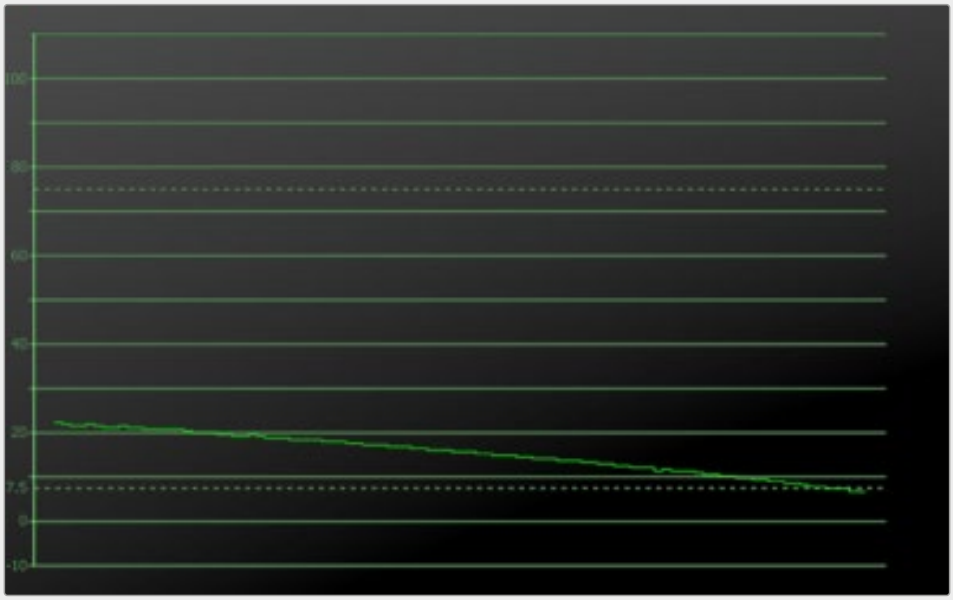
The Custom node should use the following equation in the red, green, and blue expressions:

```
if (c1< 1e-16, 1e-18 + (c1/1e-16)*(1e-16 - 1e-18), c1)
```

Falloff Comparison



The black falloff from the native Fusion process



Virtually identical black falloff from the ramped clipping process

Film Grain [FGr]



The Film Grain node

Film Grain Node Introduction

The Film Grain node adds generated film grain to an image. Typically, when compositing with different elements shot on film and in digital, the grain is removed for compositing operations and then reapplied to the final composite. This helps create the appearance that all elements were shot as a single scene with the same film stock.

NOTE: Although more accurate, the Film Grain node does not replace the older Grain node, which is still provided to allow older compositions to load and render, but in almost every case, it is better to use the Film Grain node.

Input

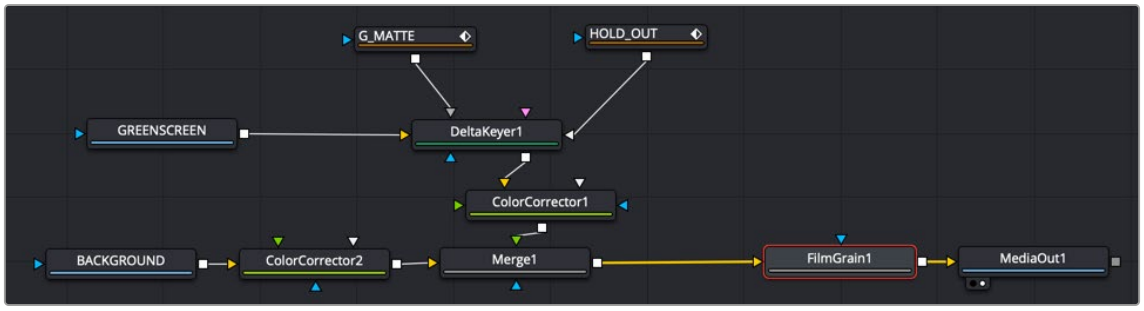
There are two inputs on the Film Grain node: one for the image and one for the effects mask.

Input: The orange input is used for the primary 2D image that gets the grain applied.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the grain to be within the pixels of the mask. An effects mask is applied to the tool after the tool is processed.

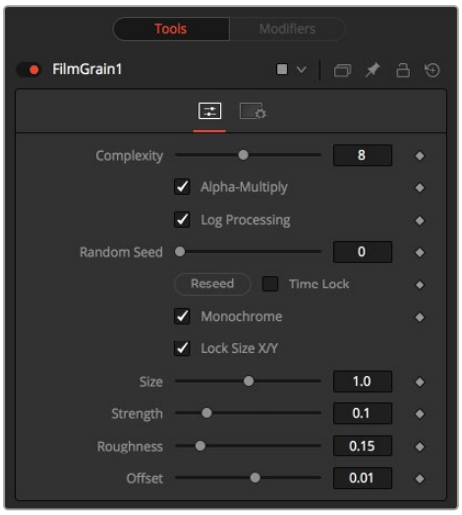
Basic Node Setup

The Film Grain node is commonly used before a MediaOut node in DaVinci Resolve or before a Saver node in Fusion Studio. Since grain is often removed from source material to achieve cleaner keys and composite, film grain is added back before the output to create a more realistic composite.



Film Grain node adding grain back before output

Inspector



Film Grain controls

Controls Tab

The Controls tab includes all the parameters for modifying the appearance of the film grain.

Complexity

The Complexity setting indicates the number of “layers” of grain applied to the image. With a complexity of 1, only one grain layer is calculated and applied to the image. When complexity is set to 4, the node calculates four separate grain layers and applies the mean combined result of each pass to the final image. Higher complexities produce visually more sophisticated results, without the apparent regularity often perceivable in digitally-produced grain.

Alpha Multiply

When the Alpha Multiply checkbox is enabled, the Film Grain node multiplies its results by the source image’s Alpha channel. This is necessary when working with post-multiplied images to ensure that the grain does not affect areas of the image where the Alpha is 0.0 (transparent).

> **NOTE:** Since it is impossible to say what the final value of semitransparent pixels in the image are until after they are composited with their background, you should avoid applying log-processed grain to the elements until after they have been composited. This ensures that the strength of the grain is accurate.

Log Processing

When this checkbox is enabled (default), the grain applied to the image has its intensity applied nonlinearly to match the grain profile of most film. Roughly speaking, the intensity of the grain increases exponentially from black to white. When this checkbox is disabled, the grain is applied uniformly, regardless of the brightness of the affected pixel.

One of the primary features of grain in film is that the appearance of the grain varies radically with the exposure so that there appears to be minimal grain present in the blacks, with the amount and deviation of the grain increasing as the pixels exposure increases. In a film negative, the darkest portions of the developed image appear entirely opaque, and this obscures the grain. As the negative becomes progressively clearer, more of the grain becomes evident in the result. Chemical differences

in the R, G, B, layer's response to light also cause each color component of the film to present a different grain profile, typically with the blue channel presenting the most significant amount of grain.

As a result, an essential control in the Film Grain node is the Log Processing checkbox, which should be enabled when matching film, and disabled when working with images that require a more linear grain response. Having this checkbox enabled closely mimics the results of preceding the old Grain node with a Linear to Log conversion and following with a Log to Linear conversion immediately after.

Seed

The Seed slider and Reseed button are presented whenever a Fusion node relies on a random result. Two nodes with the same seed values produce the same random results. Click on the Reseed button to randomly select a new seed value, or adjust the slider to select a new seed value manually.

Time Lock

Enabling Time Lock stops the random seed from generating new grain on every frame.

Monochrome

When the Monochrome checkbox is enabled (default), the grain is applied to the red, green, and blue color channels of the image equally. When deselected, individual control over the Size, Strength, and Roughness of the grain in each channel becomes possible.

Lock Size X/Y

Deselect the Lock Size X/Y checkbox to control the size of the grain along the X- and Y-axis individually.

Size

The grain size is calculated relative to the size of a pixel. Consequently, changing the resolution of the image does not impact the relative appearance of the grain. The default grain size of 1.0 produces grain kernels that cover roughly 2 pixels.

Strength

Grain is expressed as a variation from the original color of a pixel. The stronger the grain's strength, the wider the possible variation from the original pixel value. For example, given a pixel with an original value of p, and a Grain node with complexity = 1 size = 1; roughness = 0; log processing = off; the grain produces an output value of p +/- strength. In other words, a pixel with a value of 0.5 with a grain strength of 0.02 could end up with a final value between 0.48 and 0.52.

Once again, that's a slight oversimplification, especially when the complexity exceeds 1. Enabling the Log Processing checkbox also causes that variation to be affected such that there is less variation in the blacks and more variation in the whites of the image.

> **NOTE:** When visualizing the effect of the grain on the image, the more mathematically inclined may find it helps to picture a sine wave, where each lobe of the sine wave covers 1 pixel when the Grain Size is 1.0. The Grain Size controls the frequency of the sine wave, while the Grain Strength controls its amplitude. Again, this is something of an oversimplification.

Roughness

The Roughness slider applies low frequency variation to give the impression of clumping in the grain. Try setting the roughness to 0, and observe that the grain produced has a very even luminance variation across the whole image. Increase the roughness to 1.0 and observe the presence of "cellular" differences in the luminance variation.

Offset

The Offset control helps to match the intensity of the grain in the deep blacks by offsetting the values before the intensity (strength) of the grain is calculated. So an offset of 0.1 would cause a pixel with a value of 0.2 to receive grain as if its value was 0.2.

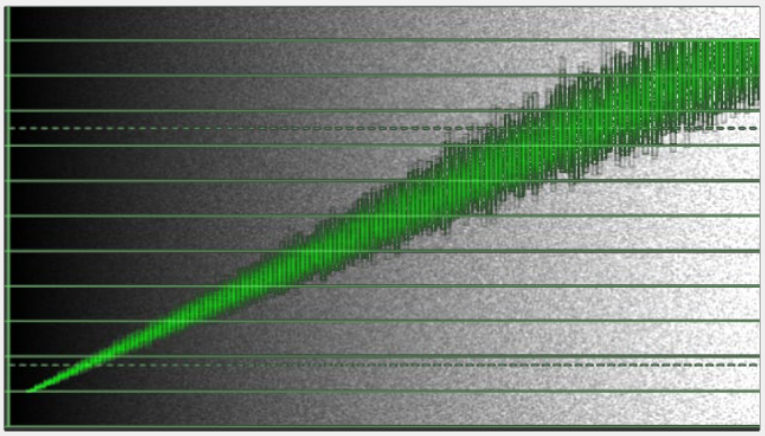
Common Controls

Settings Tab

The Settings tab controls are common to all Film nodes, so their descriptions can be found in "The Common Controls" section at the end of this chapter.

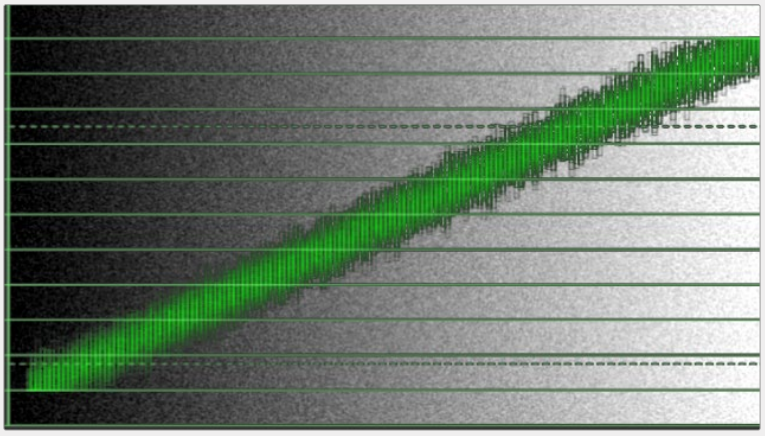
Processing Examples

Log Processing On



In the default setting, the different amounts of Grain are applied to the blacks and the whites of the image

Log Processing Off



When Log processing is off, the Grain is applied evenly to the entire image, as shown here

Grain [Grn]



The Grain node

Grain Node Introduction

The Grain node offers an older film grain emulation than the more modern and accurate Film Grain node. The Grain node is still provided to allow older compositions to load and render, but in almost every case, it is better to use the Film Grain node.

Input

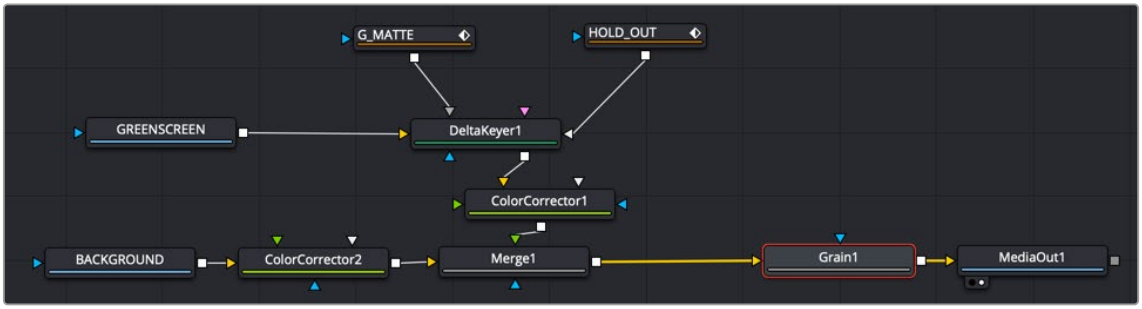
There are two Inputs on the Grain node: one for the 2D image and one for the effects mask.

Input: The orange input is used for the primary 2D image that gets the grain applied.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the grain to be within the pixels of the mask. An effects mask is applied to the tool after the tool is processed.

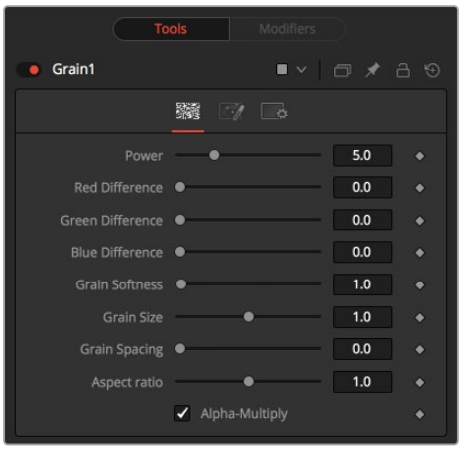
Basic Node Setup

The Grain node is commonly placed before a MediaOut node in DaVinci Resolve, or before a Saver node in Fusion Studio. It adds back grain previously removed in a composite.



A Grain node used to add grain back for a more realistic composite

Inspector



Grain controls

Controls Tab

The Controls tab includes all the parameters for modifying the appearance of the grain.

Power

This slider determines the strength of the grain. A higher value increases visibility, making the grain more prevalent.

RGB Difference

Separate Red, Green, and Blue sliders are used to modify the strength of the effect on a per channel basis.

Grain Softness

This slider controls the blurriness or fuzziness of the grain. Smaller values cause the grain to be more sharp or coarse.

Grain Size

This slider determines the size of the grain particles. Higher values increase the grain size.

Grain Spacing

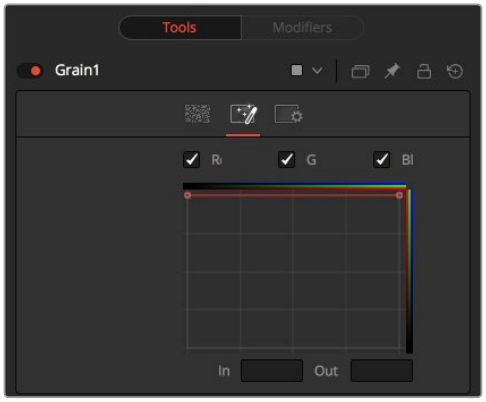
This slider determines the density or amount of grain per area. Higher values cause the grain to appear more spaced out.

Aspect Ratio

This slider adjusts the aspect of the grain so that it can be matched with anamorphic images.

Alpha-Multiply

When enabled, this checkbox multiplies the image by the Alpha, clearing the black areas of any grain effect.



Grain Spread controls

Spread Tab

The Spread tab uses curves for the red, green, and blue channels to control the amount of grain over each channel's tonal range.

RGB Checkboxes

The red, green, and blue checkboxes enable each channel's custom curve, allowing you to control how much grain appears in each channel. To mimic usual film responses, more grain would appear in the blue channel than the red, and the green channel would receive the least. Right-clicking in the spline area displays a contextual menu containing options related to modifying spline curves.

For more information on the LUT Editor's controls see *Chapter 47, "LUT Nodes,"* in the Fusion Reference Manual.

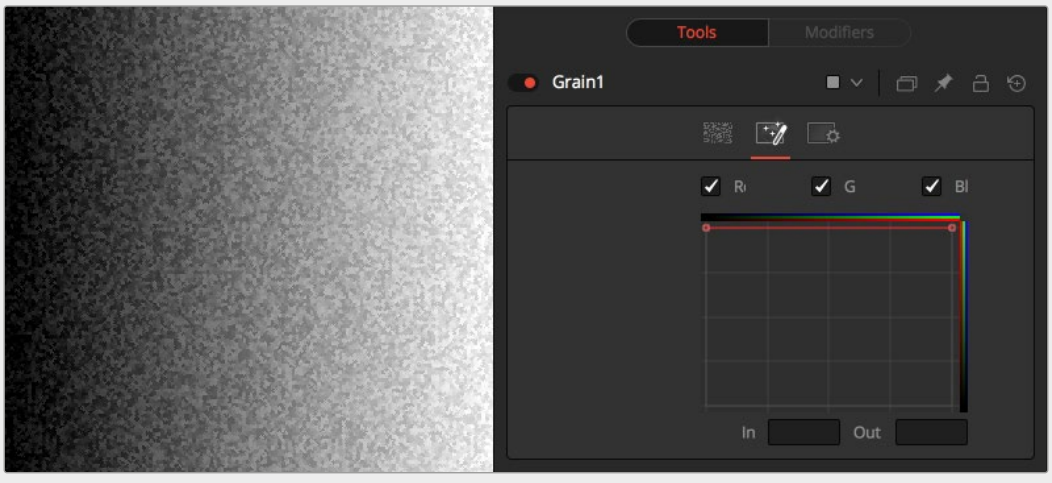
In and Out

This control provides direct editing of points on the curve by setting In/Out point values.

Examples

Default Spread

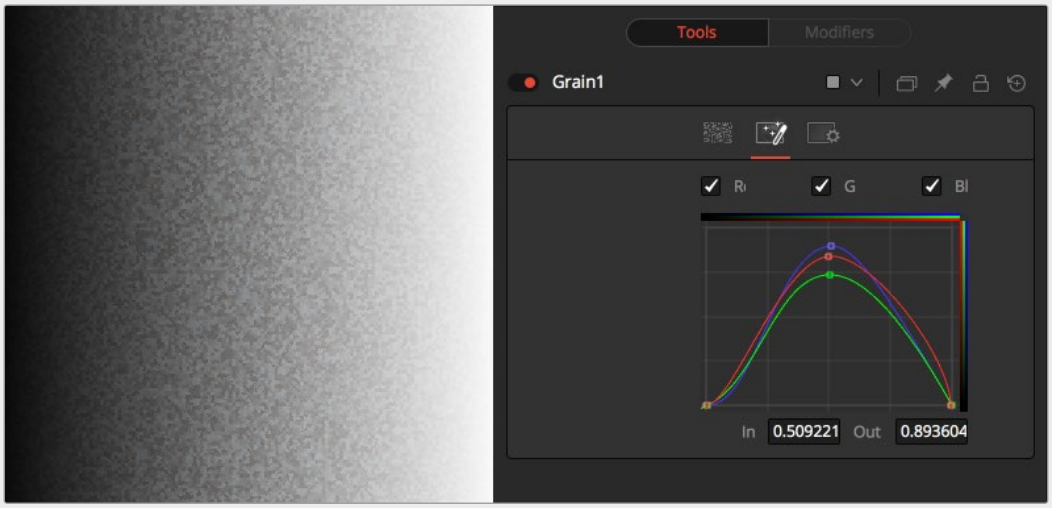
In the default setting, the grain is applied evenly to the entire image, as shown here. However, film often shows a different amount of grain in the blacks, mids, and whites.



Bell-Shaped Spread

Setting a bell shape is often a good starting point to create a more realistic-looking grain. Here we have a non-uniform distribution with different amounts of grain in the red, green, and blue channels.

In both examples, the grain's power has been exaggerated to show the effect a bit better.



Common Controls

Setting Tab

The Settings tab controls are common to all Film nodes, so their descriptions can be found in "The Common Controls" section at the end of this chapter.

Light Trim [LT]



The Light Trim node

Light Trim Node Introduction

This node emulates film scanner light trims. By design, this node works best with logarithmic data, such as the images stored by Cineon, Arri, or Blackmagic RAW files. When logarithmic data is provided, the Light Trim node can be used to increase or decrease the apparent exposure level of the image.

Inputs

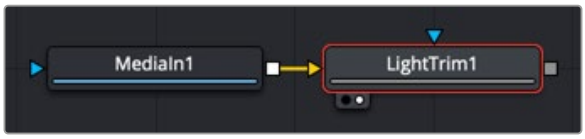
There are two Inputs on the Light Trim node: one for the 2D image and one for the effects mask.

Input: The orange input is used for the primary Log 2D image that gets its exposure adjusted.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the exposure change to be within the pixels of the mask. An effects mask is applied to the tool after the tool is processed.

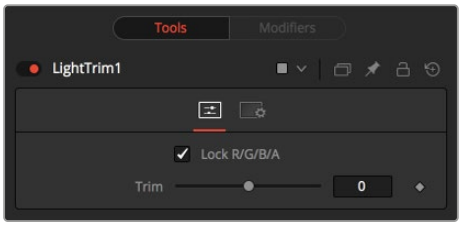
Basic Node Setup

The Light Trim node is placed after a LOG clip but before the Log clip is converted by a Cineon LOG node.



A Light Trim node used to adjust exposure on a LOG clip

Inspector



Light Trim controls

Controls Tab

The Controls tab includes a single slider that adjusts the exposure of the image.

Lock RGBA

When selected, the Lock RGBA control collapses control of all image channels into one slider. This selection is on by default. To manipulate the various color channels independently, deselect this checkbox.

Trim

This slider shifts the color in film, optical printing, and lab printing points. 8 points equals one stop of exposure.

Common Controls

Settings Tab

The Settings tab controls are common to all Film nodes, so their descriptions can be found in “The Common Controls” section at the end of this chapter.

Remove Noise [RN]



The Remove Noise node

Remove Noise Node Introduction

The Remove Noise node provides simple noise management. The basic operation is that the node blurs the image channels, and then compares the blurred image to the original to extract the noise. A sharpness is then applied to the image, except where noise was detected.

To use this node, view the image and look at the red channel. Then increase the Red Softness until the grain appears to be gone. Next, increase the sharpness until the detail reappears, but stop before the grain reappears. Repeat for the green and blue channels.

Inputs

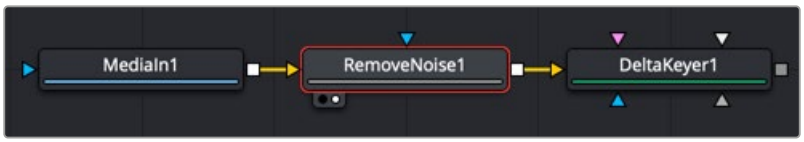
There are two inputs on the Remove Noise node: one for the 2D image and one for the effects mask.

Input: The orange input is used for the primary 2D image that gets noise removed.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the noise removal change to be within the pixels of the mask. An effects mask is applied to the tool after the tool is processed.

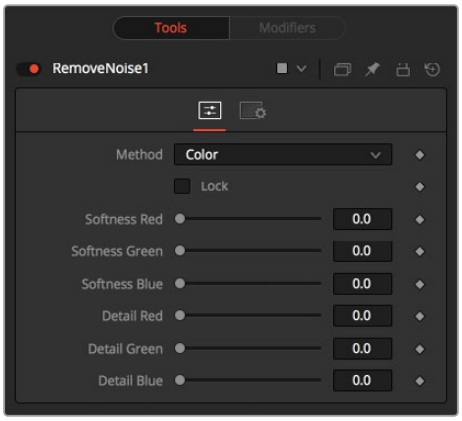
Basic Node Setup

The Remove Noise node can be used on any clip with noise. For example, it is used below to remove noise prior to keying the clip using the DeltaKeyer.



A Remove Noise node used to remove noise prior to keying

Inspector



Remove Noise controls

Controls Tab

The Controls tab switches the noise removal between two methods: Color and Chroma. When the Method is set to Color, the Controls tab adjusts the amount of blur and sharpness individually for each RGB channel. When the Method is set to Chroma, the blur and sharpness is adjusted based on Luma and Chroma controls.

Method

This menu is used to choose whether the node processes color using the Color or Chroma method. This also gives you a different set of control sliders.

Lock

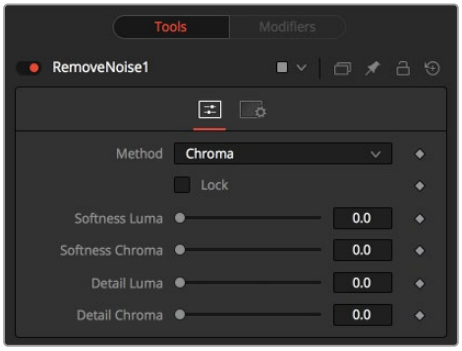
This checkbox links the Softness and Detail sliders of each channel together.

Softness Red, Green, and Blue

The Softness sliders determine the amount of blur applied to each channel of the image. In Chroma mode, you have sliders for the softness in the Luminance and Chrominance channels, respectively.

Detail Red, Green, and Blue

The Sharpness sliders determine how much detail is reintroduced into each channel after each channel is softened. In Chroma mode, you have sliders for Luminance and Chrominance channels, respectively.



Chroma Method controls

Common Controls

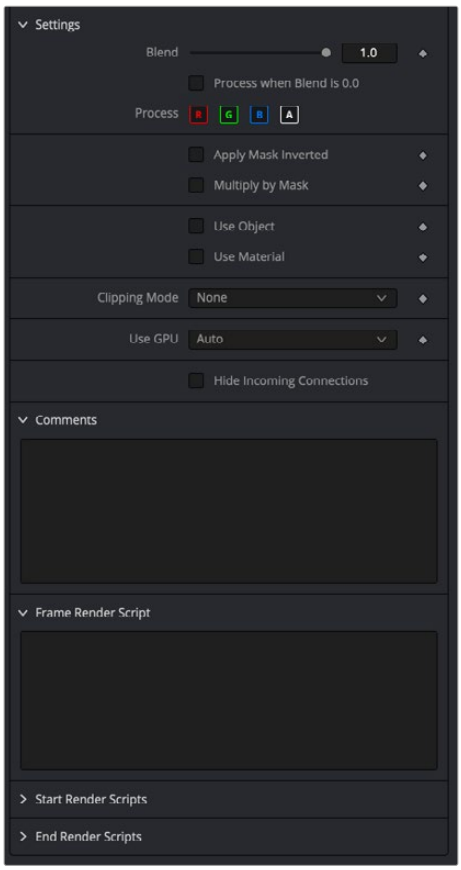
Settings Tab

The Settings tab controls are common to all Film nodes, so their descriptions can be found in the following “The Common Controls” section.

The Common Controls

Film nodes share a number of identical controls in the Inspector. This section describes controls that are common among Film nodes.

Inspector



The Common Film settings inspector

Settings Tab

The Settings tab in the Inspector can be found on every tool in the Film category. The Settings controls are even found on third-party Film-type plugin tools. The controls are consistent and work the same way for each tool, although some tools do include one or two individual options, which are also covered here.

Blend

The Blend control is used to blend between the tool's original image input and the tool's final modified output image. When the blend value is 0.0, the outgoing image is identical to the incoming image. Commonly, this causes the tool to skip processing entirely, copying the input straight to the output.

Process When Blend Is 0.0

The tool is processed even when the input value is zero. This can be useful if processing of this node is scripted to trigger another task, but the value of the node is set to 0.0.

Red/Green/Blue/Alpha Channel Selector

These four buttons are used to limit the effect of the tool to specified color channels. This filter is often applied after the tool has been processed.

For example, if the red button on a Blur tool is deselected, the blur is first applied to the image, and then the red channel from the original input is copied back over the red channel of the result.

There are some exceptions, such as tools for which deselecting these channels causes the tool to skip processing that channel entirely. Tools that do this generally possess a set of identical RGBA buttons on the Controls tab in the tool. In this case, the buttons in the Settings and the Controls tabs are identical.

Apply Mask Inverted

Enabling the Apply Mask Inverted option inverts the complete mask channel for the tool. The mask channel is the combined result of all masks connected to or generated in a node.

Multiply by Mask

Selecting this option causes the RGB values of the masked image to be multiplied by the mask channel's values. This causes all pixels not included in the mask (i.e., set to 0) to become black/transparent.

Use Object/Use Material (Checkboxes)

Some 3D software can render to file formats that support additional channels. Notably, the EXR file format supports Object and Material ID channels, which can be used as a mask for the effect. These checkboxes determine whether the channels are used, if present. The specific Material ID or Object ID affected is chosen using the next set of controls.

Correct Edges

This checkbox appears only when the Use Object or Use Material checkboxes are selected. It toggles the method used to deal with overlapping edges of objects in a multi-object image. When enabled, the Coverage and Background Color channels are used to separate and improve the effect around the edge of the object. If this option disabled (or no Coverage or Background Color channels are available), aliasing may occur on the edge of the mask.

> For more information on Coverage and Background Color channels, *see Chapter 18, "Understanding Image Channels,"* in the Fusion Reference Manual.

Object ID/Material ID (Sliders)

Use these sliders to select which ID is used to create a mask from the object or material channels of an image. Use the Sample button in the same way as the Color Picker: to grab IDs from the image displayed in the view. The image or sequence must have been rendered from a 3D software package with those channels included.

Clipping Mode

This option determines how edges are handled when performing domain of definition rendering. This is mostly important for nodes like Blur, which may require samples from portions of the image outside the current domain.

- **Frame:** The default option is Frame, which automatically sets the node's domain of definition to use the full frame of the image, effectively ignoring the current domain of definition. If the upstream DoD is smaller than the frame, the remaining area in the frame is treated as black/transparent.
- **Domain:** Setting this option to Domain respects the upstream domain of definition when applying the node's effect. This can have adverse clipping effects in situations where the node employs a large filter.
- **None:** Setting this option to None does not perform any source image clipping at all. This means that any data required to process the node's effect that would normally be outside the upstream DoD is treated as black/transparent.

Use GPU

The Use GPU menu has three settings. Setting the menu to Disable turns off hardware-accelerated rendering using the graphics card in your computer. Enabled uses the hardware. Auto uses a capable GPU if one is available and falls back to software rendering when a capable GPU is not available

Comments

The Comments field is used to add notes to a tool. Click in the empty field and type the text. When a note is added to a tool, a small red square appears in the lower-left corner of the node when the full tile is displayed, or a small text bubble icon appears on the right when nodes are collapsed. To see the note in the Node Editor, hold the mouse pointer over the node to display the tooltip.

Scripts

Three Scripting fields are available on every tool in Fusion from the Settings tab. They each contain edit boxes used to add scripts that process when the tool is rendering. For more details on scripting nodes, please consult the Fusion scripting documentation.

Chapter 40

Filter Nodes

This chapter details the Filter nodes available in Fusion.

The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Create Bump Map [CBu] 1104
Custom Filter Node [CFlt] 1106
Erode Dilate Node [ErDl] 1110
Filter Node [Fltr] 1112
Rank Filter Node [RFlt] 1115
The Common Controls 1117

Create Bump Map [CBu]



The Create Bump Map node

Create Bump Map Node Introduction

The Create Bump Map node converts a grayscale (height map) image into a bump map. Unlike the Bump Map node that turns an image into a 3D material, the Create Bump Map node creates bump vector data and provides the output as an RGB image so other image-processing operations can be applied.

Input

The Create Bump Map node includes two inputs: one for the main image and the other for an effect mask to limit the area where the bump map is created.

Input: The orange input takes the RGBA channels from an image to calculate the bump map.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the creation of the bump map to only those pixels within the mask. An effects mask is applied to the tool after the tool is processed.

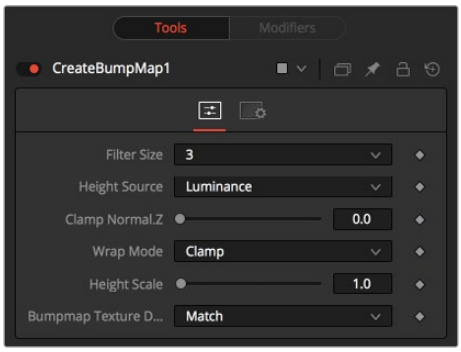
Basic Node Setup

The Create Bump Map node accepts a 2D grayscale image like a fast noise, which can then go through various 2D image-processing filters to create the bump map texture.



A Create Bump Map node produces a bump map as an RGB image for further image processing

Inspector



Create Bump Map controls

Controls Tab

The Controls tab contains all parameters for creating the bump map.

Filter Size

This menu sets the filter size for creating the bump map. You can set the filter size at 3 x 3 pixels or 5 x 5 pixels, thus determining the radius of the pixels sampled. The larger the size, the more time it takes to render.

Height Source

The Height Source menu selects the channel for extracting the grayscale information.

Clamp Normal.Z

This slider clips the lower values of the blue channel in the resulting bump texture.

Wrap Mode

This menu determines how the image wraps at the borders, so the filter produces a correct result when using seamless tiling textures.

Height Scale

The height scale menu modifies the contrast of the resulting values in the bump map. Increasing this value yields in a more visible bump map.

Bump Map Texture Depth

This menu matches or converts the resulting bump texture into the desired bit depth.

NOTE: The below definitions are provided to clarify some of the terminology used in the Create Bump Map node and other similar types of nodes.

Height Map: A grayscale image containing a height value per pixel.

Bump Map: An image containing normals stored in the RGB channels used for modifying the existing normals (usually given in tangent space).

Normal Map: An image containing normals stored in the RGB channels used for replacing the existing normals (usually given in tangent or object space).

Common Controls

Settings Tab

The Settings tab controls are common to all Filter nodes, so their descriptions can be found in “The Common Controls” section at the end of this chapter.

Custom Filter Node [CFlt]



The Custom Filter node

Custom Filter Node Introduction

The Custom Filter node is used to apply custom convolution filters to images. A custom convolution filter can give a wide variety of image effects. For example, emboss, relief, sharpen, blurring, and edge detection are all convolution filters. There are many supplied custom filters in the Filters directory that can be loaded by right-clicking on the control header and selecting Settings > Load from the contextual menu.

The Custom filter uses an array (or grid) of either 3 x 3, 5 x 5, or 7 x 7 values. (Note: The array in the Inspector always shows a 7 x 7 grid; however, setting the Matrix Size to 3 x 3 uses only the center 9 cells.) The center of the array represents the current pixel, and entries nearby represent adjacent pixels. A value of 1 applies the full value of the pixel to the filter. A value of 0 ignores the pixel’s value. A value greater than 1 multiplies the pixel’s effect on the result. Negative values can also be entered, where the value of the pixel is subtracted from the average. Only integer values can be entered; 0.x is not valid.

Input

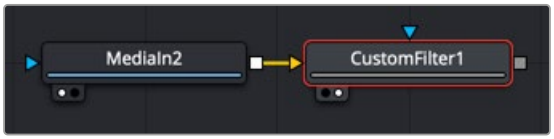
The Custom Filter node includes two inputs: one for the main image and the other for an effect mask to limit the area where the custom filter is applied.

Input: The orange input takes the RGBA channels from an image to calculate the custom filter.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the custom filter to only those pixels within the mask. An effects mask is applied to the tool after the tool is processed.

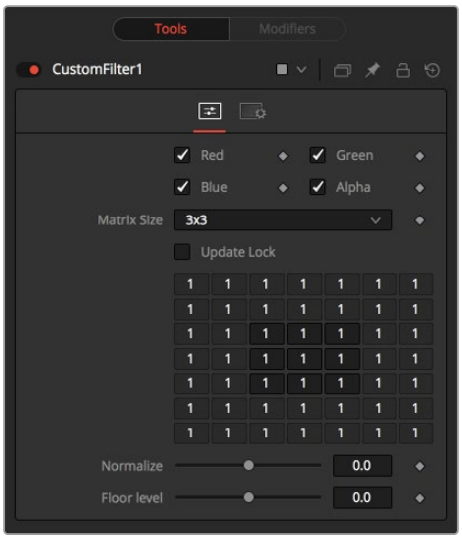
Basic Node Setup

The Custom Filter node can be inserted after an image, mask, or any node that needs a custom convolution filter applied.



A Custom Filter node placed after a MediaIn node in DaVinci Resolve

Inspector



Custom Filter controls

Controls Tab

The Controls tab is used to set the filter size and then use the filter matrix to enter convolution filter values.

Color Channels (RGBA)

The custom filter defaults to operating on R, G, B, and A channels. Selective channel editing is possible by enabling or disabling the checkboxes beside each channel.

This is not the same as the RGBA checkboxes found under the Common Controls. The node takes these controls into account before it processes. Deselecting a channel causes the node to skip that channel when processing, speeding up the rendering of the effect. In contrast, these controls under the Common Controls tab are applied after the node has processed.

Matrix Size

This menu is used to set the size of the filter at 3 x 3 pixels, 5 x 5 pixels, or 7 x 7 pixels, thus setting the radius of the pixels sampled. The larger the size, the more time it takes to render.

Update Lock

When this control is selected, Fusion does not render the filter. This is useful for setting up each value of the filter, and then turning off Update Lock and rendering the filter.

Filter Matrix

The Filter Matrix control is a 7 x 7 grid of text boxes where a number is entered to represent how much influence each pixel has on the overall convolution filter. The text box in the center represents the pixel that is processed by the filter. The text box to the left of the center represents the pixel to the immediate left, and so forth.

The default Matrix size is 3 x 3. Only the pixels immediately adjacent to the current pixel are analyzed. If a larger Matrix size is set, more of the text boxes in the grid are enabled for input.

Normalize

This controls the amount of filter normalization that is applied to the result. Zero gives a normalized image. Positive values brighten or raise the level of the filter result. Negative values darken or lower the level.

Floor Level

This adds or subtracts a minimum, or Floor Level, to the result of the filtered image. Zero does not add anything to the image. Positive values add to the filtered image, and negative values subtract from the image.

Examples

Original Image Example

For example, a filter with the values...

0 0 0

0 1 0

0 0 0

...has zero effect from its neighboring pixels, and the resulting image would be unchanged.



Original image

Softening Example

A slight softening effect would be...

1 1 1

1 1 1

1 1 1

...where the neighboring pixels are averaged with the center.



Before and after averaging neighboring pixels to soften the image

Emboss Example

The example below subtracts five times the value from the top left and adds five times the value from the lower right.

```
-5  0  0
0  1  0
0  0  5
```

If parts of the processed image are very smooth in color, the neighboring values are very similar.

In parts of the image where the pixels are different (e.g., an edge), the results are different and tend to highlight or emboss edges in the image.



A Custom Filter adding and subtracting neighboring pixels to create an embossed image

Exposure Example

Using the values...

```
1  1  1
1  1  1
1  1  1
```

...and adjusting Normalize to a positive value makes the image brighter or glow, simulating film overexposure.



The Custom Filter Normalize slider used to change exposure

Relief Example

Using the values...

```
-1  0  0
0  0  0
0  0  1
```

... and adjusting Floor Level to a positive value creates a Relief filter.



Custom Filter Floor Level slider used to create a relief

Common Controls

Settings Tab

The Settings tab controls are common to all Filter nodes, so their descriptions can be found in “The Common Controls” section at the end of this chapter.

Erode Dilate Node [ErDl]



The Erode Dilate node

Erode Dilate Node Introduction

The Erode Dilate node contracts or expands the image, depending on whether the Amount slider is set to a negative or positive value.

Inputs

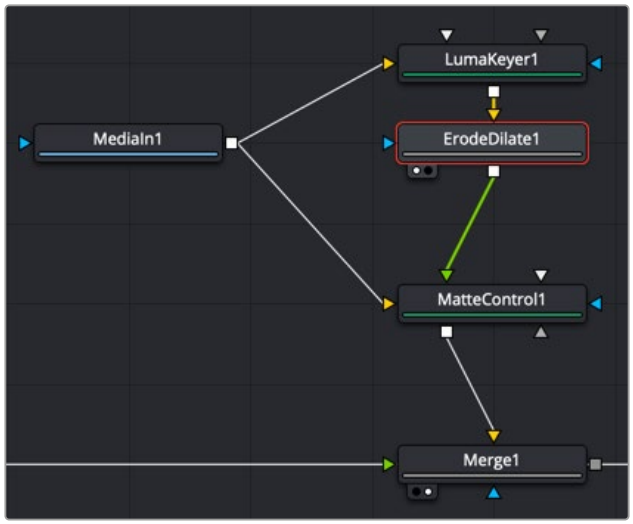
The Erode Dilate node includes two inputs: one for the main image and the other for an effect mask to limit the area where the erode or dilate is applied.

Input: The orange input takes the RGBA channels from an image to calculate the custom filter.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the erode or dilate to only those pixels within the mask. An effects mask is applied to the tool after the tool is processed.

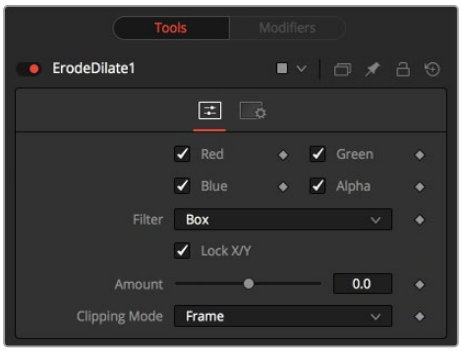
Basic Node Setup

The Erode Dilate node is commonly used to contract or expand mattes. Below, a Luma Keyer is connected to the Erode Dilate and passes the modified key to a Matte Control where it is embedded into the image.



An Erode Dilate node placed after a Luma Keyer to operate on a Matte Control

Inspector



Erode Dilate controls

Controls Tab

The Controls tab includes the main Amount slider that determines whether you are performing an erode by entering a negative value or a dilate by entering a positive value.

Color Channels (RGBA)

The Erode Dilate node defaults to operating on R, G, B, and A channels. Selective channel editing is possible by enabling or disabling the checkboxes beside each channel.

This is not the same as the RGBA checkboxes found under the Common Controls. The node takes these controls into account before it processes. Deselecting a channel causes the node to skip that channel when processing, speeding up the rendering of the effect. In contrast, the channel controls under the Common Controls tab are applied after the node has processed.

Lock X/Y

The Lock X/Y checkbox is used to separate the Amount slider into amount X and amount Y, allowing a different value for the effect on each axis.

Amount

A negative value for Amount causes the image to erode. Eroding simulates the effect of an underexposed frame, shrinking the image by growing darker areas of the image so that they eat away at brighter regions.

A positive value for Amount causes the image to dilate, similar to the effect of overexposing a camera. Regions of high luminance and brightness grow, eating away at the darker regions of the image. Both techniques eradicate fine detail in the image and tend to posterize fine gradients.

The Amount slider scale is based on the input image width. An amount value of 1 = image width. So, if you want to erode or dilate by exactly 1 pixel on an HD image, you would enter 1/1920, or 0.00052083.

Common Controls

Settings Tab

The Settings tab controls are common to all Filter nodes, so their descriptions can be found in "The Common Controls" section at the end of this chapter.

Filter Node [Fltr]



The Filter node

Filter Node Introduction

The Filter node contains several standard convolution filters, easily selectable from a list. This node enables a variety of effects, from radically changing the look of an image to adding subtle randomly-generated film grain. The Sobel and Laplacian settings are often used for edge detection.

Inputs

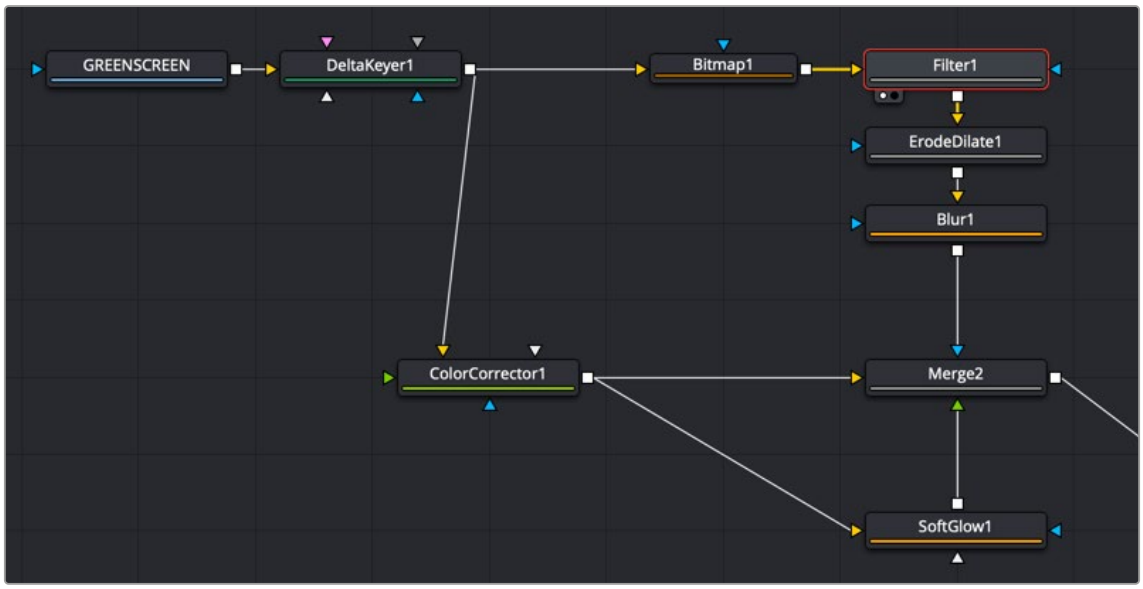
The Filter node includes two inputs: one for the main image and the other for an effect mask to limit the area where the filter is applied.

Input: The orange input is used for the primary 2D image that gets the filter applied.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the filter to only those pixels within the mask. An effects mask is applied to the tool after the tool is processed.

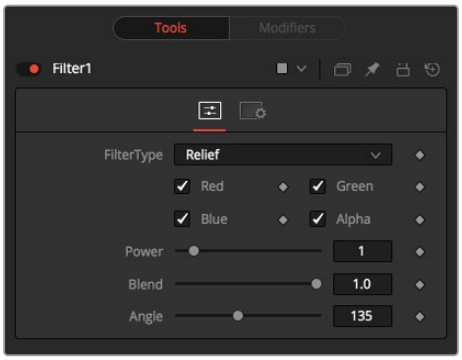
Basic Node Setup

The Filter node can be inserted after an image, mask, or any node that needs a filter applied. Below, it is used to create an edge matte, which is then used to mask the soft glow around the keyed foreground.



A Filter node using the Sobel setting to extract an edge matte

Inspector



Filter controls

Controls Tab

The Controls tab is used to set the filter type, the channels the filter is applied to, and the amount it blends with the original image.

Filter Type

The Filter Type menu provides a selection of filter types described below.

- **Relief:** This appears to press the image into metal, such as an image on a coin. The image appears to be bumped and overlaid on gray.
- **Emboss Over:** Embosses the image over the top of itself, with adjustable highlight and shadow height and direction.

- **Noise:** Uniformly adds noise to images. This is often useful for 3D computer-generated images that need to be composited with live action, as it reduces the squeaky-clean look that is inherent in rendered images. The frame number acts as the random generator seed. Therefore, the effect is different on each frame and is repeatable.
- **Defocus:** This filter type blurs the image.
- **Sobel:** Sobel is an advanced edge detection filter. Used in conjunction with a Glow filter, it creates impressive neon light effects from live-action or 3D-rendered images.
- **Laplacian:** Laplacian is a very sensitive edge detection filter that produces a finer edge than the Sobel filter.
- **Grain:** Adds noise to images similar to the grain of film (mostly in the midrange). This is useful for 3D computer-generated images that need to be composited with live action as it reduces the squeaky-clean look that is inherent in rendered images. The frame number acts as the random generator seed. Therefore, the effect is different on each frame and is repeatable.

Color Channels (RGBA)

The Filter node defaults to operating on R, G, B, and A channels. Selective channel editing is possible by enabling or disabling the checkboxes beside each channel.

Power

Values range from 1 to 10. Power proportionately increases the amount by which the selected filter affects the image. This does not apply to the Sobel or Laplacian filter type.

Angle

This control has a range from 0 to 315 degrees and changes the effect in increments of 45 degrees. This applies only to the Relief and Emboss filters.

Median

Depending on which Filter Type is selected, the Median control may appear. It varies the Median filter's effect. A value of 0.5 produces the true median result, as it finds the middle values. A value of 0.0 finds the minimums, and 1.0 finds the maximums. This applies to the Median setting only.

Seed

This control is visible only when applying the Grain or Noise filter types. The Seed slider can be used to ensure that the random elements of the effect are seeded with a consistent value. The randomizer always produces the same result, given the same seed value.

Animated

This control is visible only when applying the Grain or Noise filter types. Select the checkbox to cause the noise or grain to change from frame to frame. To produce static noise, deselect this checkbox.

Common Controls

Settings Tab

The Settings tab controls are common to all Filter nodes, so their descriptions can be found in "The Common Controls" section at the end of this chapter.

Rank Filter Node [RFlt]



The Rank Filter node

Rank Filter Node Introduction

The Rank Filter examines nearby pixels, sorts the pixels by value, and then replaces the color of the examined pixels with the color of the pixel with the selected rank.

Inputs

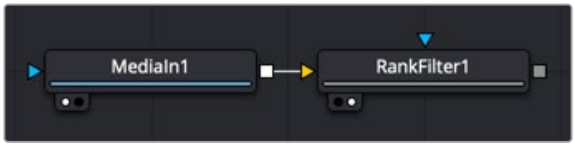
The Rank Filter node includes two inputs: one for the main image and the other for an effect mask to limit the area where the filter is applied.

Input: The orange input is used for the primary 2D image that gets the Rank filter applied.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the rank filter to only those pixels within the mask. An effects mask is applied to the tool after the tool is processed.

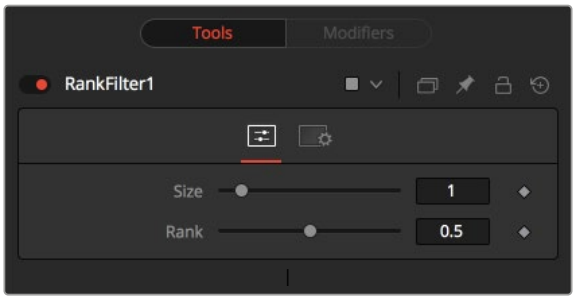
Basic Node Setup

The Rank Filter node can be placed anywhere in a node tree to apply an effect to an image.



A Rank Filter node placed after a MediaIn node in DaVinci Resolve

Inspector



Rank Filter controls

Controls Tab

The Controls tab is used to set the size and rank value of the filter.

Size

This control determines the size in pixels of the area sampled by the filter. A value of 1 samples 1 pixel in each direction, adjacent to the center pixel. This produces a total of 9 pixels, including the center sampled pixel. Larger values sample from a larger area.

Low Size settings are excellent for removing salt and pepper style noise, while larger Size settings produce an effect similar to watercolor paintings.

Rank

The Rank slider determines which value from the sampled pixels is chosen. A value of 0 is the lowest value (darkest pixel), and 1 is the highest value (brightest pixel).

Example

Below is a before and after example of a Rank filter with Size set to 7 and a Rank of 0.7 to create a watercolor effect.

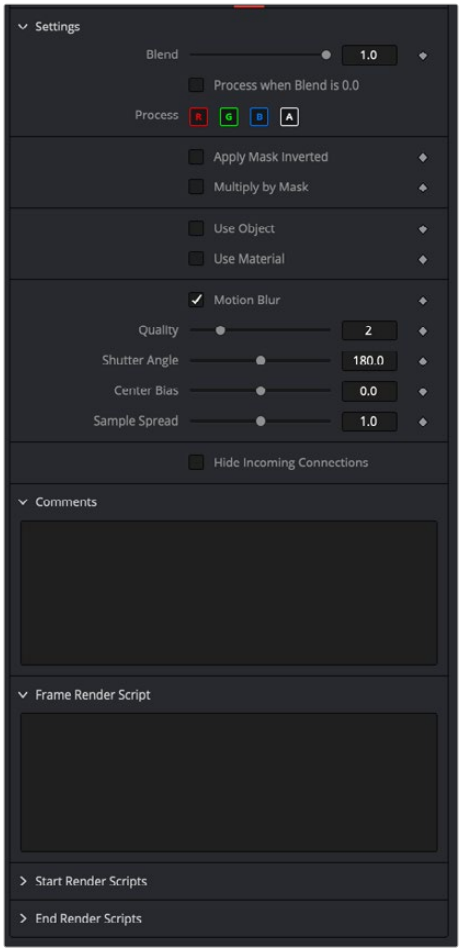


Before and after of a Rank Filter producing a watercolor-style effect

The Common Controls

Filter nodes share a number of identical controls in the Inspector. This section describes controls that are common among Filter nodes.

Inspector



Common Filter settings Inspector

Settings Tab

The Settings tab in the Inspector can be found on every tool in the Filter category. The Settings controls are even found on third-party filter-type plugin tools. The controls are consistent and work the same way for each tool, although some tools do include one or two individual options, which are also covered here.

Blend

The Blend control is used to blend between the tool's original image input and the tool's final modified output image. When the blend value is 0.0, the outgoing image is identical to the incoming image. Normally, this causes the tool to skip processing entirely, copying the input straight to the output.

Process When Blend Is 0.0

The tool is processed even when the input value is zero. This can be useful if processing of this node is scripted to trigger another task, but the value of the node is set to 0.0.

Red/Green/Blue/Alpha Channel Selector

These four buttons are used to limit the effect of the tool to specified color channels. This filter is often applied after the tool has been processed.

For example, if the red button on a Blur tool is deselected, the blur is first applied to the image, and then the red channel from the original input is copied back over the red channel of the result.

There are some exceptions, such as tools for which deselecting these channels causes the tool to skip processing that channel entirely. Tools that do this generally possess a set of identical RGBA buttons on the Controls tab in the tool. In this case, the buttons in the Settings and the Controls tabs are identical.

Apply Mask Inverted

Enabling the Apply Mask Inverted option inverts the complete mask channel for the tool. The mask channel is the combined result of all masks connected to or generated in a node.

Multiply by Mask

Selecting this option causes the RGB values of the masked image to be multiplied by the mask channel's values. This causes all pixels of the image not included in the mask (i.e., set to 0) to become black/transparent.

Use Object/Use Material (Checkboxes)

Some 3D software can render to file formats that support additional channels. Notably, the EXR file format supports Object ID and Material ID channels, which can be used as a mask for the effect. These checkboxes determine whether the channels are used, if present. The specific Material ID or Object ID affected is chosen using the next set of controls.

Correct Edges

This checkbox appears only when the Use Object or Use Material checkboxes are selected. It toggles the method used to deal with overlapping edges of objects in a multi-object image. When enabled, the Coverage and Background Color channels are used to separate and improve the effect around the edge of the object. If this option disabled (or no Coverage or Background Color channels are available), aliasing may occur on the edge of the mask.

> For more information on the Coverage and Background Color channels, *see Chapter 18, "Understanding Image Channels,"* in the Fusion Reference Manual.

Object ID/Material ID (Sliders)

Use these sliders to select which ID is used to create a mask from the object or material channels of an image. Use the Sample button in the same way as the Color Picker: to grab IDs from the image displayed in the view. The image or sequence must have been rendered from a 3D software package with those channels included.

Use GPU

The Use GPU menu has three settings. Setting the menu to Disable turns off hardware-accelerated rendering using the graphics card in your computer. Enabled uses the hardware. Auto uses a capable GPU if one is available and falls back to software rendering when a capable GPU is not available

Motion Blur

- **Motion Blur:** This toggles the rendering of Motion Blur on the tool. When this control is toggled on, the tool's predicted motion is used to produce the motion blur caused by the virtual camera's shutter. When the control is toggled off, no motion blur is created.
- **Quality:** Quality determines the number of samples used to create the blur. A quality setting of 2 causes Fusion to create two samples to either side of an object's actual motion. Larger values produce smoother results but increase the render time.
- **Shutter Angle:** Shutter Angle controls the angle of the virtual shutter used to produce the motion blur effect. Larger angles create more blur but increase the render times. A value of 360 is the equivalent of having the shutter open for one whole frame exposure. Higher values are possible and can be used to create interesting effects.
- **Center Bias:** Center Bias modifies the position of the center of the motion blur. This allows for the creation of motion trail effects.
- **Sample Spread:** Adjusting this control modifies the weighting given to each sample. This affects the brightness of the samples.

Comments

The Comments field is used to add notes to a tool. Click in the empty field and type the text. When a note is added to a tool, a small red square appears in the lower-left corner of the node when the full tile is displayed, or a small text bubble icon appears on the right when nodes are collapsed. To see the note in the Node Editor, hold the mouse pointer over the node to display the tooltip.

Scripts

Three Scripting fields are available on every tool in Fusion from the Settings tab. They each contain edit boxes used to add scripts that process when the tool is rendering. For more details on scripting nodes, please consult the Fusion scripting documentation.

Chapter 41

Flow Nodes

This chapter details the Sticky Note and Underlay features available in Fusion.

The abbreviations next to each feature name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Sticky Note [Nte] 1121

Sticky Note Introduction 1121

Usage 1121

Underlay [Und] 1122

Underlay Introduction 1122

Usage 1122

Sticky Note [Nte]



The Sticky Note

Sticky Note Introduction

A Sticky Note is not a node at all. It is a useful way of attaching notes, comments, and history to a specific area of a comp. By changing their size and color, they can provide unobtrusive comments or important notices, as required. Sticky Notes make an excellent complement to the Comments tab in the Inspector.

Usage



A Sticky Note added to a node tree can provide a description for other people or as a reminder to yourself

To create a Sticky Note, click in an empty area of the Node Editor where you want a Sticky Note to appear. Then, from the Effects Library, click the Sticky Note effect located in the Tools > Flow category or press Shift-Spacebar and search for the Sticky Note in the Select Tool window.

Like Groups, Sticky Notes are created in a smaller, collapsed form. They can be expanded by double-clicking on them. Once expanded, they can be resized using any side or corner of the note or moved by dragging on the name header. To collapse the Sticky Note again, click the icon in the top-left corner.

To rename, delete, copy, or change the color of the note, right-click over the note and choose from the contextual menu. Using this menu, you can also lock the note to prevent editing.

To edit the text in a Sticky Note, first expand it by double-clicking anywhere on the note, and then click below its title bar. If the note is not locked, you can edit the text.

Underlay [Und]

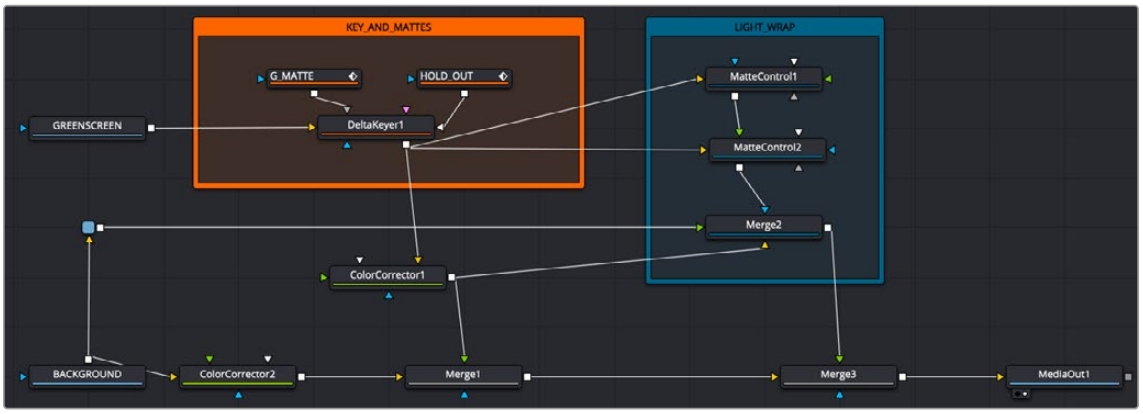


The Underlay

Underlay Introduction

Underlays are a convenient method of visually organizing areas of a composition. As with Groups, Underlays can improve the readability of a comp by separating it into labeled functional blocks. While Groups are designed to streamline the look of a comp by collapsing complex layers down to single nodes, Underlays highlight, rather than hide, and do not restrict outside connections.

Usage



Underlay node structure

As with Sticky Notes , an Underlay can be added to a comp by selecting it from the Flow category in the Effects Library or searching for it in the Select Tool window. The Underlay to the Node Editor with its title bar is centered on the last-clicked position.

Underlays can be resized using any side or corner. This will not affect any nodes.

Underlays can also be used as simple selection groups. Activating an Underlay, by clicking its title, will select all the tools contained wholly within it as well, allowing the entire set to be moved, duplicated, passed through, and so on.

To rename an Underlay, first ensure that nodes contained within the Underlay are not selected. Then, Option-click on the Underlay title to select the Underlay without selecting the nodes it contains. Once selected, right-click over the title and choose Rename. Underlays can be assigned a color using the same right-click contextual menu.

Chapter 42

Flow Organizational Nodes

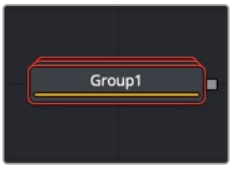
This chapter details the Groups, Macro, and Pipe Router nodes, which are designed to help organize your compositions, making the node tree easier to see and understand.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Groups 1124
Groups Introduction 1124
Usage 1124
Macro 1125
Macro Introduction 1125
Usage 1125
Macro Editor 1125
The Final Macro 1126
Pipe Router 1126
Router Introduction 1126
Usage 1127
Router 1127

Groups



The Group node

Groups Introduction

Groups are used to keep complex node trees organized. You can select any number of nodes in the node tree and then group them to create a single node icon in the Node Editor. Groups are non-destructive and can be opened at any time.

Usage

- To group nodes, select them in the Node Editor, and then right-click over any of the selected nodes and choose Group from the contextual menu.



Selecting nodes to group

- To edit the individual nodes in a group, right-click and choose Expand Group from the contextual menu. All individual nodes contained in the group are displayed in a floating node tree window. When opened, groups hover over existing elements, allowing editing of the enclosed nodes.



An opened node group

- To remove or decompose a group and retain the individual nodes, right-click the group and choose Ungroup.

Macro

Macros are not technically a node. Instead, they are a group of nodes that act as a single node.

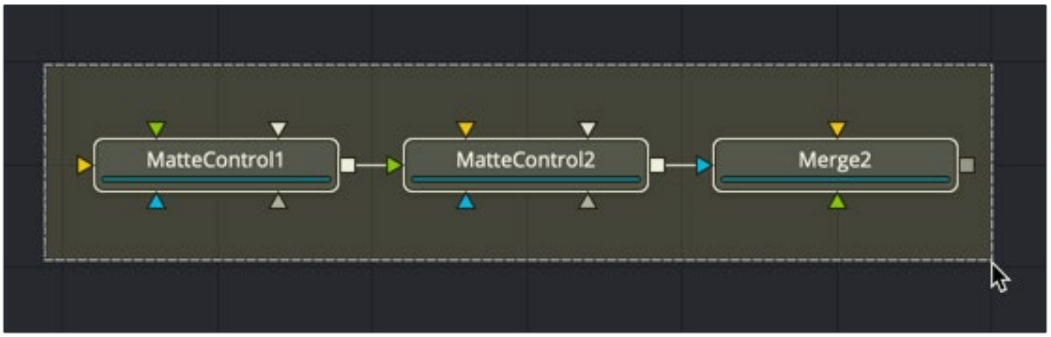
Macro Introduction

Macros can be used to combine multiple nodes and expose a user-definable set of controls.

They are meant as a fast and convenient way of building custom nodes.

Usage

To create a Macro, select the nodes intended for the macro. The order in which the nodes are selected becomes the order in which they are displayed in the Macro Editor. Right-click on any of the selected nodes and choose Macro > Create Macro from the contextual menu.

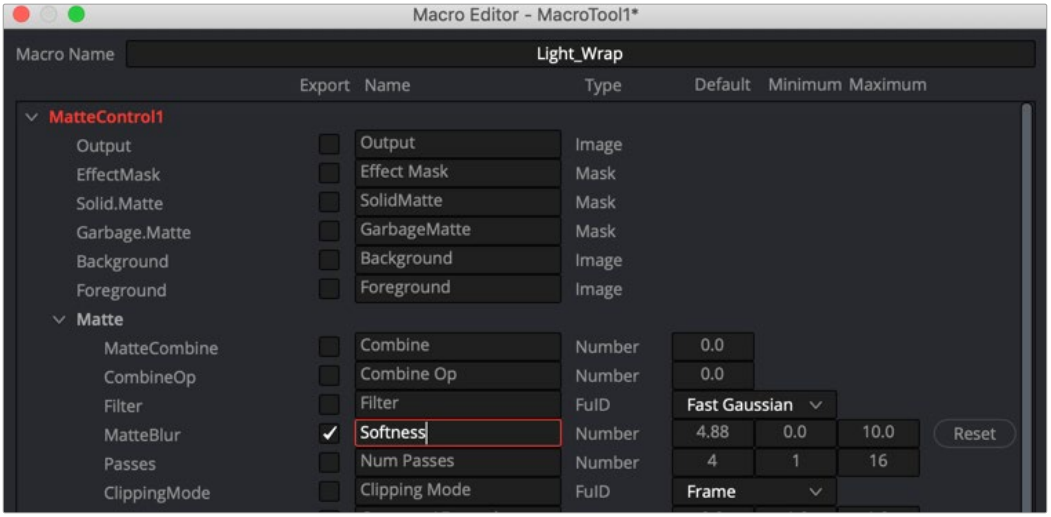


Select the nodes to include in the macro

Macro Editor

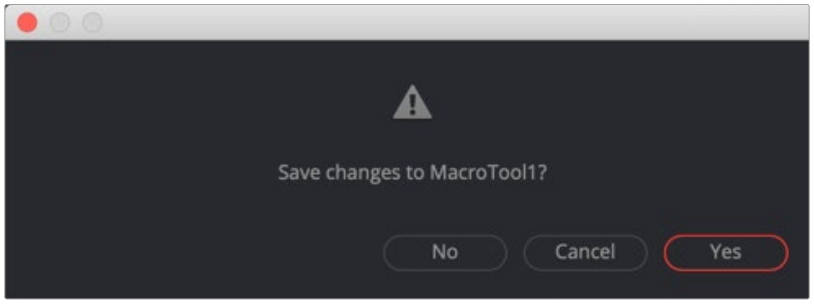
The Macro Editor allows you to specify and rename the controls that are exposed in the final macro tool.

In the example below, the tool is named Light_Wrap at the top. The Blur slider for Matte Control 1 is enabled and renamed to Softness, as it will appear in the Inspector.



The Macro Editor

After setting up the macro to your liking, click the Close button in the lower-right corner of the dialog. Then, in the Save dialog, click Yes to save the macro, click No to leave Macro Editor without saving the changes, or click Cancel to return to the Macro Editor.



The Macro Editor save dialog

To add the macro to your node tree, right-click anywhere on the node tree and select Macro > [NameOfYourMacro] from the contextual menu.

Saving a Macro as a Title Template in the Edit Page (DaVinci Resolve)

When using DaVinci Resolve, macros are available only in the Fusion page. However, if the macro is a title animation, you can save it to the Titles Templates folder and have it appear in the Edit page Effects Library.

To save a title macro so it appears in the Edit page Effects Library, save the macro to:

macOS: Users > UserName > Library > Application Support > Blackmagic Design > DaVinci Resolve > Fusion > Templates > Edit > Titles

Windows: C Drive > Users > UserName > AppData > Roaming > Blackmagic Design > DaVinci Resolve > Support > Fusion > Templates > Edit > Titles

The Final Macro

The final macro looks and behaves just like any other node in Fusion.

As another example, you could take a single Channel Boolean, set it to Add mode, and make it into a macro exposing no controls at all, thus creating the equivalent of an Add Mix node like the one that can be found in programs like Nuke.

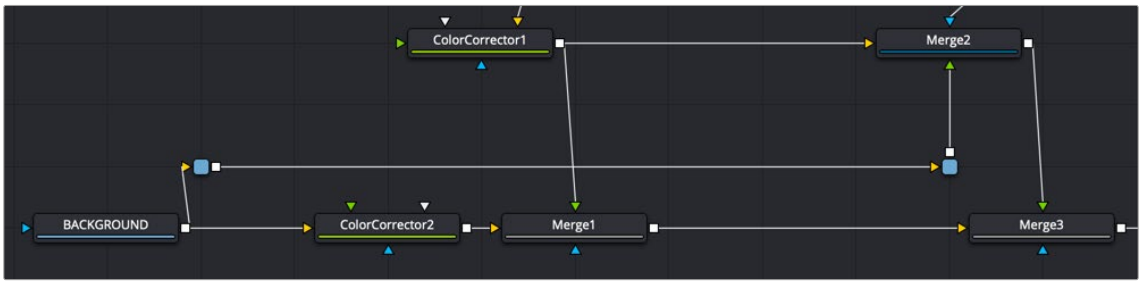
Pipe Router

Pipe Routers are another type of organizational tool you use to improve the layout and appearance of the node tree.

Router Introduction

Routers can be used to neatly organize your comps by creating “elbows” in your node tree, so the connection lines do not overlap nodes, making them easier to understand. Routers do not have any influence on render times.

Usage

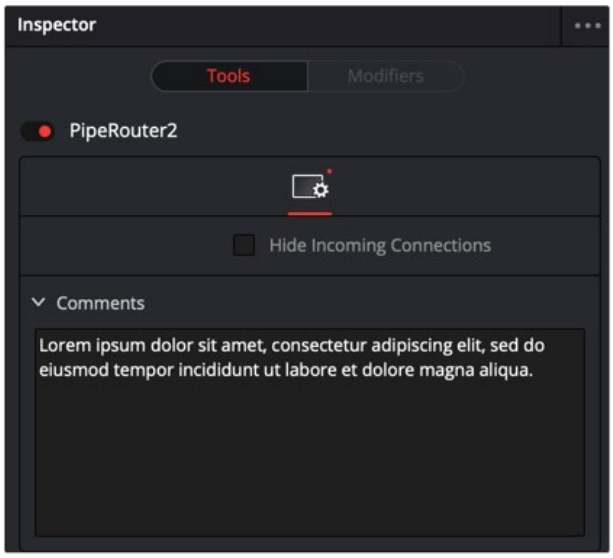


A Pipe Router usage example

Router

To insert a router along a connection line, Option- or Alt-click on the line. The router can then be repositioned to arrange the connections as needed.

Although routers have no actual controls, they still can be used to add comments to a comp.



An example comment in a Router node

Chapter 43

Fuses

This chapter introduces Fuses, which are scriptable plugins that can be used within Fusion.

Contents

Fuses [Fus] 1129

Fuses Introduction 1129

Installing Fuses 1129

Working with Fuses in a Composition 1129

Fuses [Fus]



A Fuse node

Fuses Introduction

Fuses are plugins. The difference between a Fuse and an Open FX plugin is that a Fuse is created using a Lua script. Fuses can be edited within Fusion or DaVinci Resolve, and the changes you make compile on-the-fly.

Using a Lua script makes it easy for even non-programmers to prototype and develop custom nodes. A new Fuse can be added to a composition, edited and reloaded, all without having to close the current composition. They can also be used as modifiers to manipulate parameters, curves, and text very quickly. ViewShader Fuses can make use of the GPU for faster performance. This makes Fuses much more convenient than an Open FX plugin that uses Fusion's OFX SDK. However, this flexibility comes at a cost. Since a Fuse is compiled on-the-fly, it can be significantly slower than the identical node created using the Open FX SDK.

As an example, Fuses could generate a mask from the over-exposed areas of an image, or create initial particle positions based on the XYZ position stored within a text file.

Please contact Blackmagic Design for access to the SDK (Software Developer Kit) documentation.

Installing Fuses

Fuses are installed in the Fusion:\Fuses path map. By default this folder is located at Users/ User_Name/Library Application Support/Blackmagic Design/Fusion (or DaVinci Resolve)/Fuses on macOS or C:\Users\User_Name\AppData\Roaming\Blackmagic Design\Fusion (or DaVinci Resolve)\Fuses, on Windows. Files must use the extension .fuse, or they will be ignored by Fusion.

Working with Fuses in a Composition

Fuses can be designed to appear in any category in the Effects Library. Once installed, they are added to a composition exactly as any native or third-party plugin node. However, since a Fuse is just a text document, it can be edited by clicking the Edit button that appears at the top of the Inspector when the Fuse node is selected. This will open the Fuse in the default script editor specified in the Global Preferences/Scripting panel.

NOTE: Any changes made to a Fuse's script do not immediately affect other copies of the same Fuse node already added to a composition. To use the updated Fuse script on all similar Fuses in the composition, either close and reopen the composition, or click on the Reload button in each Fuse's Inspector.

When a composition containing a Fuse node is opened, the currently saved version of the Fuse script is used. The easiest way to ensure that a composition is running the current version of a Fuse is to close and reopen the composition.

Chapter 44

Generator Nodes

This chapter details the Generator nodes available in Fusion.

The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Background [Bg] 1131
Day Sky [DS] 1134
Fast Noise [FN] 1137
Mandelbrot [MAN] 1140
Plasma [PLAS] 1143
Text+ [Txt+] 1145
On-screen Controls 1145
Text+ Modifiers 1159
Character Level Styling 1159
Comp Name 1160
Follower 1160
Text Scramble 1162
Text Timer 1163
Time Code 1164
The Common Controls 1165

Background [Bg]



The Background node

Background Node Introduction

The Background node can be used to produce anything from a solid color background to complex looped gradients.

Inputs

There is one input on the Background node for an effect mask input.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the background color to only those pixels within the mask.

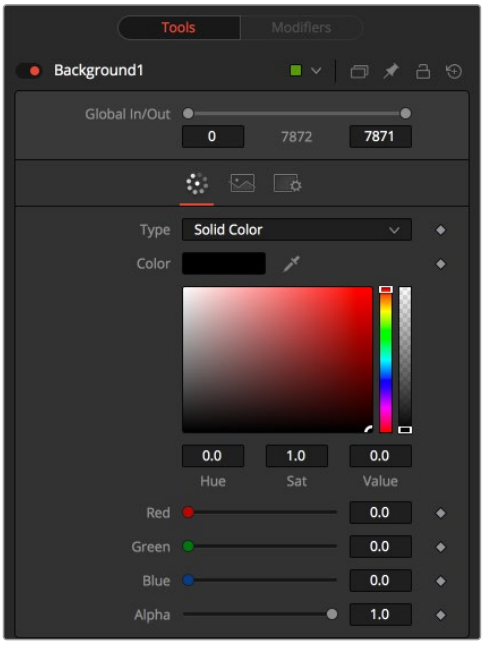
Basic Node Setup

The Background node generates a color and passes the color off to another node like the Paint node below.



A Background node generating a background for the Paint tool

Inspector



The Background node controls

Color Tab

The Color tab is used to control the style and color(s) of the generated background.

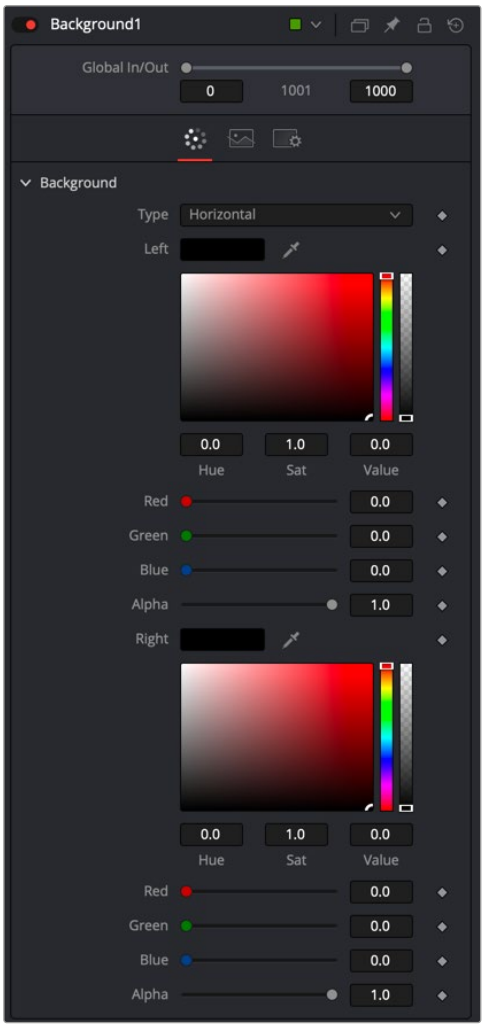
Type

This control is used to select the style of background generated by the node. Four selections are available:

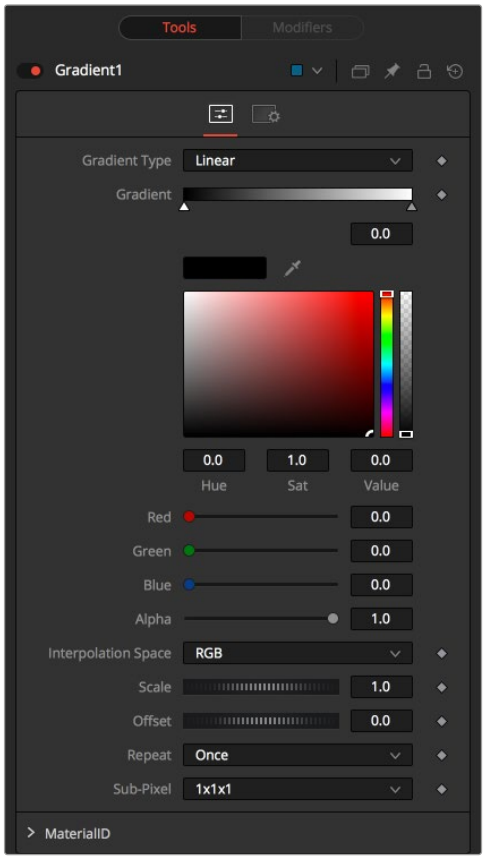
- **Solid Color:** This default creates a single-color image.
- **Horizontal:** This creates a two-color horizontal gradation.
- **Vertical:** This creates a two-color vertical gradation.
- **Four Corner:** This creates a four-color corner gradation.

Horizontal/Vertical/Four Point

When the Type menu is set to Horizontal, Vertical, or Four Corner, two- or four-color swatches are displayed where the left/right, top/bottom, or four corners of the gradient colors can be set



Horizontal gradient controls



Gradient controls

Gradient

When the Type menu is set to Gradient, additional controls are displayed where the gradient colors' direction can be customized.

Gradient Type

This menu selects the form used to draw the gradient. There are six choices:

- **Linear:** Draws the gradient along a straight line from the starting color stop to the ending color stop.
- **Reflect:** Draws the gradient by mirroring the linear gradient on either side of the starting point.
- **Square:** Draws the gradient by using a square pattern when the starting point is at the center of the image.
- **Cross:** Draws the gradient using a cross pattern when the starting point is at the center of the image.
- **Radial:** Draws the gradient in a circular pattern when the starting point is at the center of the image.
- **Angle:** Draws the gradient in a counterclockwise sweep when the starting point is at the center of the image.

Viewer Start and End Position

The Start and End positions in the viewer are represented by two red control points connected by a green line. They determine where the gradient begins and ends.

Gradient Colors

This gradient color bar is used to select the colors for the gradient. The default two color stops set the start and end colors. You can change the colors used in the gradient by selecting the color stop, and then using the Eyedropper or color swatch to set a new color.

You can add, move, copy, and delete color from the gradient using the gradient bar.

To modify one of the colors, select the triangle below the color on the bar.

To add a color stop to the gradient bar:

1. Click anywhere along the bottom of the gradient bar.
2. Use the Eyedropper or color swatch to select the color for the color stop.

To move a color stop on the gradient bar:

Drag a color stop left or right along the bar.

To copy a color stop:

Hold Command (macOS) or Ctrl (Windows) while you drag a color stop.

To delete a color stop:

Drag the color stop up past the gradient bar.

Interpolation Space

This menu determines what color space is used to calculate the colors between color stops.

Offset

The Offset control is used to offset the position of the gradient relative to the start and end markers. This control is most useful when used in conjunction with the repeat and ping-pong modes described below.

Repeat

This menu includes three options used to set the behavior of the gradient when the Offset control scrolls the gradient past its start and end positions. Selecting Once keeps the color continuous

for offset. Selecting Repeat loops around to the start color when the offset goes beyond the end color. Selecting Ping-pong repeats the color pattern in reverse.

Sub-Pixel

The Sub-Pixel menu controls the sub-pixel precision used when the edges of the gradient become visible in repeat mode, or when the gradient is animated. Higher settings will take significantly longer to render but are more precise.

Gradient Contextual Menu

Gradients have their own contextual menu that you can bring up by right-clicking on the gradient bar. The Gradient contextual menu includes options for animating, publishing, and connecting one gradient to another. There is also a gradient-specific modifier that builds a custom gradient by sampling colors from the output of a node.

Common Controls

Image and Settings Tabs

The Image and Settings tabs in the Inspector are duplicated in many Generator nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Day Sky [DS]



The Day Sky node

Day Sky Node Introduction

The Day Sky node aims to produce a simulation of the daylight produced at a specific time and location on the earth, and generates a high dynamic range image that represents a map of that light. It is not a sky generator, although it could be combined with a cloud generator or Noise node to produce one.

NOTE: This generator is a practical implementation of the research paper, *A Practical Analytical Model for Daylight*, by Preetham, Shirley, and Smits. A copy of the original paper can be found at the website for the Visual Simulation Group at the University of Utah [https://www.cs.utah.edu/~shirley/papers/sunsky/sunsky.pdf].

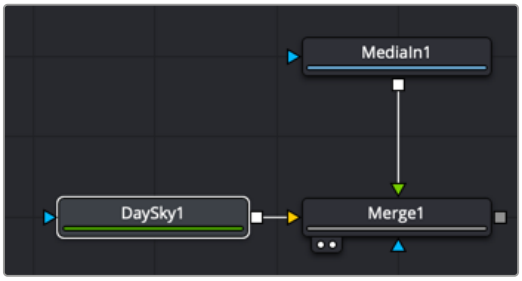
Inputs

There is a single input on the Day Sky node for an effect mask to limit the area where the day sky simulation occurs is applied.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the Day Sky to only those pixels within the mask.

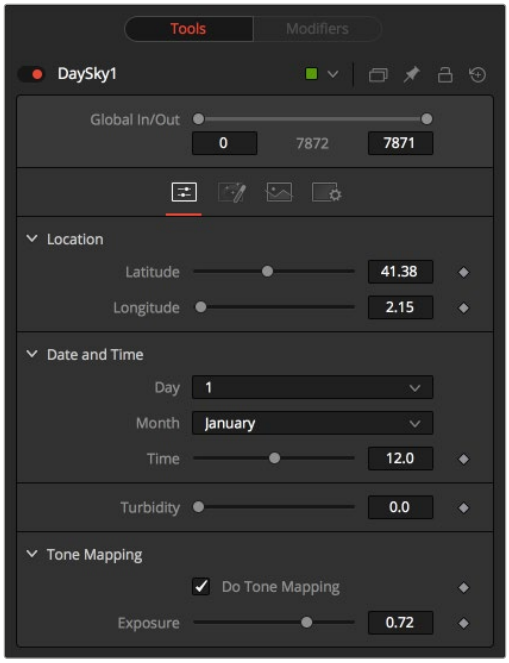
Basic Node Setup

The Day Sky node is a generator, so it typically starts the branch of a node tree and connects to some other node, like a Merge node.



A Day Sky connected as the background to a Merge node

Inspector



Day Sky controls

Controls Tab

The Controls tab is used to set the location and time of the daylight simulation. This will determine the overall look that is generated.

Location

The Latitude and Longitude sliders are used to specify the location used to create the Day Sky simulation.

Date and Time

The Day, Month, and Time controls are used to select the specific time for the Day Sky simulation.

Turbidity

Turbidity causes light to be scattered and absorbed instead of transmitted in straight lines through the simulation. Increasing the turbidity will give the sky simulation a murky feeling, as if smoke or atmospheric haze were present.

Do Tone Mapping

Since the simulation is calculated in 32-bit floating-point color space, it generates color values well above 1.0 and well below 0.0. Tone mapping is a process that takes the full dynamic range of the resulting simulation and compresses the data into the desired exposure range while attempting to preserve as much detail from the highlights and shadows as possible. Deselect this checkbox to disable any tone mapping applied to the simulation.

Generally, this option should be deselected only if the resulting image will later be color corrected as part of a floating-point color pipeline.

Exposure

Use this control to select the exposure used for tone mapping.



Day sky Advance tab controls

Advanced Tab

The Advanced tad provides more specific controls over the brightness and width of the different ranges in the generated sky.

Horizon Brightness

Use this control to adjust the brightness of the horizon relative to the sky.

Luminance Gradient

Use this control to adjust the width of the gradient separating the horizon from the sky.

Circumsolar Region Intensity

Use this control to adjust the intensity or brightness of the sky nearest to the sun.

Circumsolar Region Width

Use this control to adjust the width or size of the area in the sky affected by the sun.

Backscattered Light

Use this control to increase or decrease the backscatter light in the simulation.

Common Controls

Image and Settings Tabs

The Image and Settings tabs in the Inspector are duplicated in many Generator nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Fast Noise [FN]



The Fast Noise node

Fast Noise Node Introduction

The Fast Noise node is a very fast and flexible Perlin Noise generator. It can be useful for a wide range of effects, from clouds, to swirling fog, waves, water caustics, stylized fire, and smoke, and other organic textures. It is also invaluable as a noise source for other effects, such as heat shimmers, particle systems, and dirtiness maps.

Inputs

The two map inputs on the Fast Noise node allow you to use masks to control the value of the noise detail and brightness controls for each pixel. These two optional inputs can allow some interesting and creative effects. There is also a standard effect mask input for limiting the Fast Noise size.

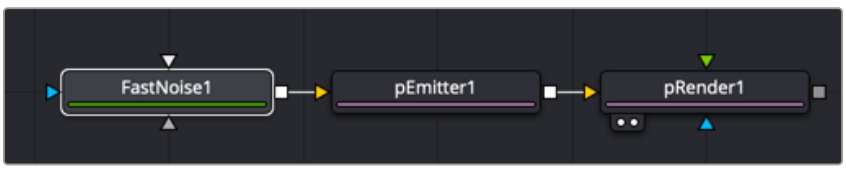
Noise Detail Map: A soft-edged mask connected to the gray Noise Detail Map input will give a flat noise map (zero detail) where the mask is black, and full detail where it is white, with intermediate values smoothly reducing in detail. It is applied before any gradient color mapping. This can be very helpful for applying maximum noise detail in a specific area, while smoothly falling off elsewhere.

Noise Brightness Map: A mask connected to this white input can be used to control the noise map completely, such as boosting it in certain areas, combining it with other textures, or if Detail is set to 0, replacing the Perlin Noise map altogether.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the Fast Noise to only those pixels within the mask.

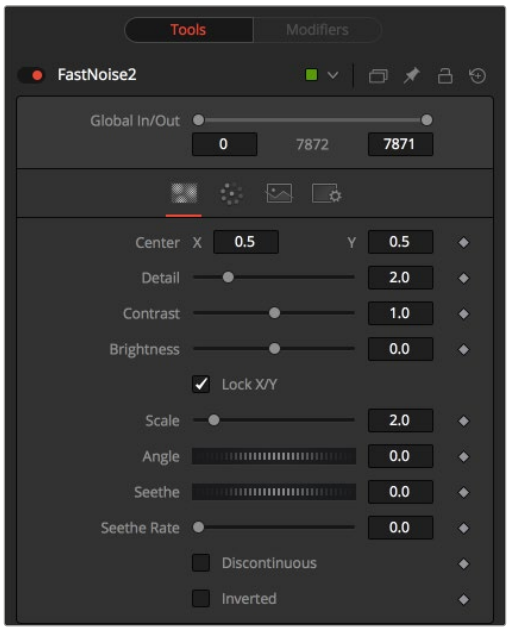
Basic Node Setup

The Fast Noise node is used to generate images for other nodes to take advantage of. For example, below the Fast Noise node is used as a bitmap source in the Particle Emitter.



A Fast Noise node used as a bitmap source for a Particle Emitter

Inspector



Fast Noise controls

Noise Tab

The Noise tab controls the shape and pattern of the noise for the Fast Noise node.

Discontinuous

Normally, the noise function interpolates between values to create a smooth continuous gradient of results. Enable this checkbox to create hard discontinuity lines along some of the noise contours. The result will be a dramatically different effect.

Inverted

Select this checkbox to invert the noise, creating a negative image of the original pattern. This is most effective when Discontinuous is also enabled.

Center

Use the Center coordinate control to pan and move the noise pattern.

Detail

Increase the value of this slider to produce a greater level of detail in the noise result. Larger values add more layers of increasingly detailed noise without affecting the overall pattern. High values take longer to render but can produce a more natural result.

Brightness

This control adjusts the overall brightness of the noise map, before any gradient color mapping is applied. In Gradient mode, this has a similar effect to the Offset control.

Contrast

This control increases or decreases the overall contrast of the noise map, prior to any gradient color mapping. It can exaggerate the effect of the noise and widen the range of colors applied in Gradient mode.

Lock and Scale X/Y

The size of the noise map can be adjusted using the Scale slider, changing it from gentle variations over the whole image to a tighter overall texture effect. The Scale slider can be separated into independent X- and Y-axis scale sliders by clicking on the Lock X/Y checkbox immediately above, which can be useful for a brushed-metal effect.

Angle

Use the Angle control to rotate the noise pattern.

Seethe

Adjust this thumbwheel control to interpolate the noise map against a different noise map. This will cause a crawling shift in the noise, as if it was drifting or flowing. This control must be animated to affect the gradient over time, or you can use the Seethe Rate control below.

Seethe Rate

As with the Seethe control above, the Seethe Rate also causes the noise map to evolve and change. The Seethe Rate defines the rate at which the noise changes each frame, causing an animated drift in the noise automatically, without the need for spline animation.



Fast Noise node Color tab

Color Tab

The Color tab allows you to adjust the gradient colors used in the generated noise pattern.

Two Color

A simple two-color gradient is used to color the noise map. The noise function will smoothly transition from the first color into the second.

Gradient

The Advanced Gradient control in Fusion is used to provide more control over the color gradient used with the noise map.

Common Controls

Image and Settings Tabs

The Image and Settings tabs in the Inspector are duplicated in many Generator nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Mandelbrot [MAN]



The Mandelbrot node

Mandelbrot Node Introduction

This node creates an image pattern based on the Mandelbrot fractal theory set.

Inputs

The one input on the Mandelbrot node is for an effect mask to limit the area where the fractal noise is applied.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the fractals to only those pixels within the mask.

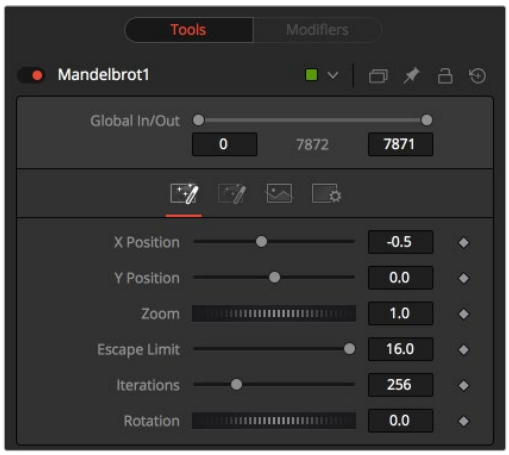
Basic Node Setup

The Mandelbrot node generates images that can be used for motion graphics and sci-fi effects. Below, it is used to generate motion graphics for the background of a title.



A Mandelbrot node applied as a background to the Merge node

Inspector



Mandelbrot node controls

Noise Tab

The Noise tab controls the shape and pattern of the noise for the Mandelbrot node.

Position X and Y

This chooses the image's horizontal and vertical position or seed point.

Zoom

Zoom magnifies the pattern in or out. Every magnification is recalculated so that there is no practical limit to the zoom.

Escape Limit

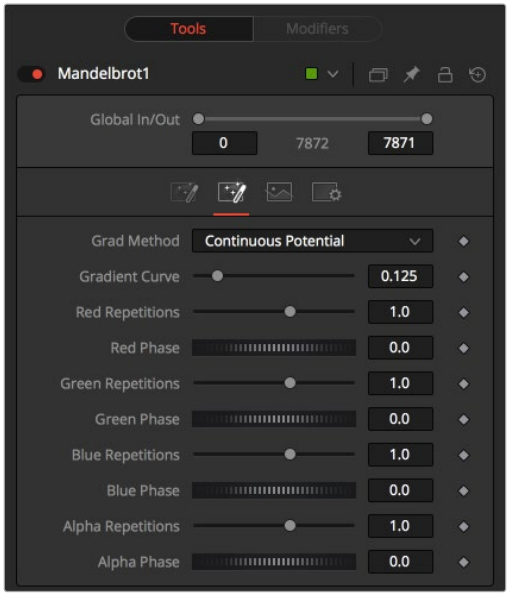
Defines a point where the calculation of the iteration is aborted. Low values lead to blurry halos.

Iterations

This determines the repetitiveness of the set. When animated, it simulates a growing of the set.

Rotation

This rotates the pattern. Every new angle requires recalculation of the image.



The Mandelbrot node Color tab

Color Tab

The Color tab allows you to adjust the gradient and repetition of the gradient colors for the generated pattern.

Grad Method

Use this control to determine the type of gradation applied at the borders of the pattern.

Continuous Potential

This causes the edges of the pattern to blend to the background color.

Iterations

This causes the edges of the pattern to be solid.

Gradient Curve

This affects the width of the gradation from the pattern to the background color.

R/G/B/A Phase and Repetitions

These controls set the color values of the pattern generators.

Common Controls

Image and Settings Tabs

The Image and Settings tabs in the Inspector are duplicated in other generator nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Plasma [PLAS]



The Plasma node

Plasma Node Introduction

The Plasma node is a background generation node that uses four circular patterns to generate images similar to plasma. It is useful as a deforming pattern for the Shadow and Deform nodes and to create a variety of other useful shapes and patterns. It is similar to the Fast Noise node.

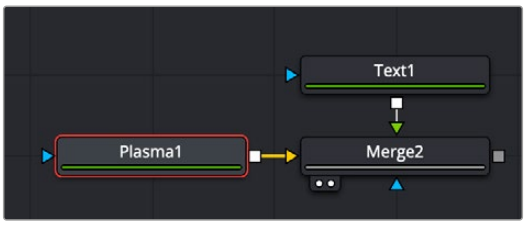
Inputs

The one input on the Plasma node is for an effect mask to limit the area where the plasma pattern is applied.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the plasma to only those pixels within the mask.

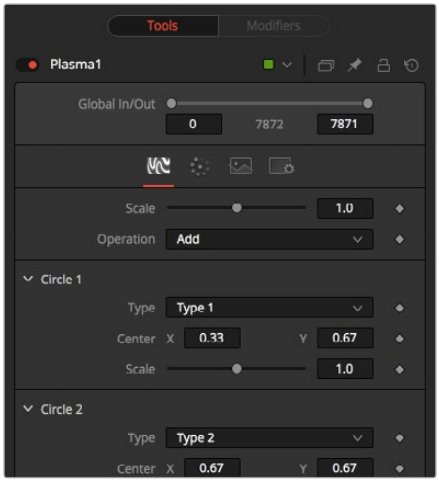
Basic Node Setup

Like the Mandelbrot node, the Plasma node generates images that can be used for motion graphics and sci-fi effects. Below, it is used to generate motion graphics for the background of a title.



A Plasma node applied as a background to a Merge node

Inspector



Plasma node Circles tab

Circles Tab

The Circles tab controls the shape and pattern generated by the Plasma node.

Scale

The Scale control is used to adjust the size of the pattern created.

Operation

The options in this menu determine the mathematical relationship among the four circles whenever they intersect.

Circle Type

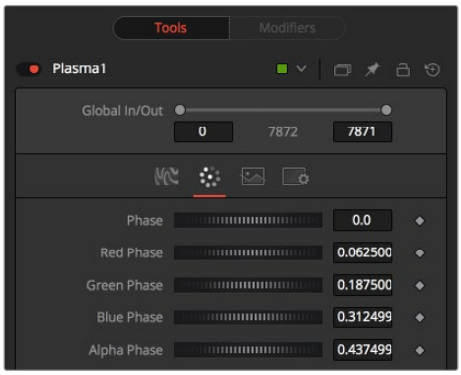
Select the type of circle to be used.

Circle Center

Report and change the position of the circle center.

Circle Scale

Determines the size of the circle to be used for the pattern.



Plasma node Color tab

Color Tab

The Color tab allows you to adjust the colors and location within the pattern of the colors for the generated plasma.

Phase

Phase changes the color phase of the entire image. When animated, this creates psychedelic color cycles.

R/G/B/A Phases

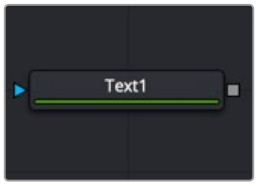
Changes the phase of the individual color channels and the Alpha. When animated, this creates color cycling effects.

Common Controls

Image and Settings Tabs

The Image and Settings tabs in the Inspector are duplicated in many Generator nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Text+ [Txt+]



The Text+ node

Text+ Node Introduction

Fusion's Text+ node is an advanced character generator capable of multiple styles, 3D transformations, and several layers of shading. Text can be laid out to a user-defined frame, circle, or along a path.

Any TrueType, OpenType, or PostScript 1 font installed on the computer can be used to create text. Support for multibyte and Unicode characters allows text generation in any language, including right to left and vertically oriented text.

This node generates a 2D image. To produce extruded 3D text with optional beveling, see the Text 3D node.

Inputs

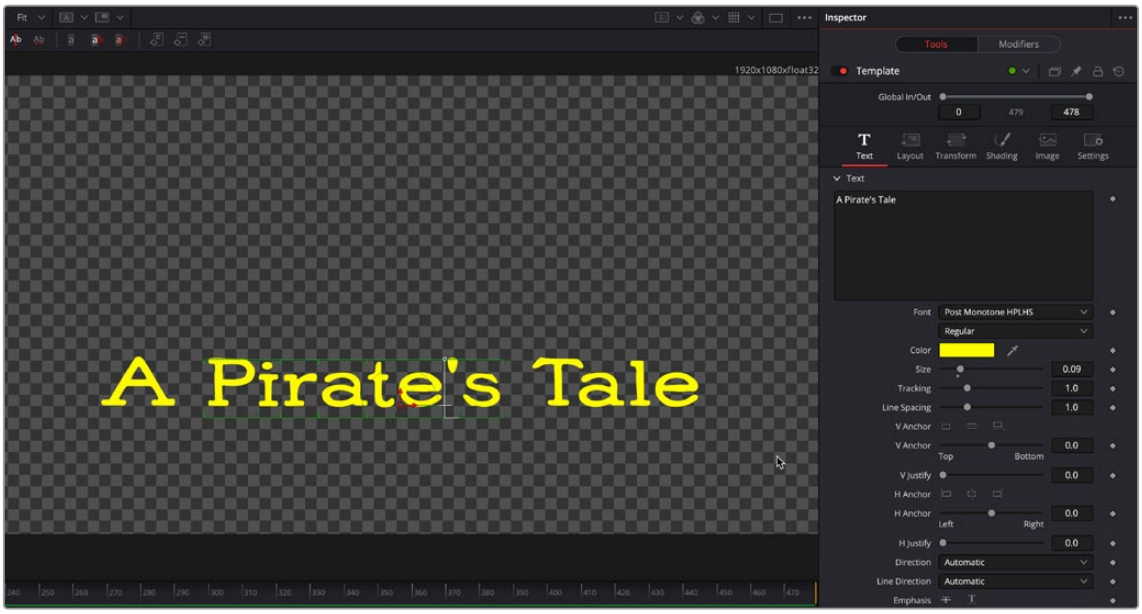
The one input on the Text+ node is for an effect mask to crop the text.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the text to only those pixels within the mask.

On-screen Controls

You can perform various editing actions directly on Text+ in the Viewer overlay, such as editing text, displaying a cursor, and selecting text.

To activate this feature, hover over the popup toolbar within the Fusion viewer and activate it, or alternatively, right-click on the View menu. This functionality is intuitive; when Text+ is selected, double-clicking anywhere in the Fusion viewer activates the in-viewer text editing mode.



The Text+ on screen viewer editing controls in the Fusion viewer.

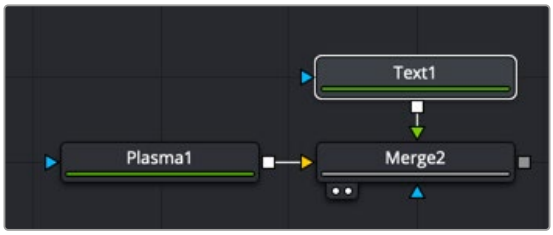
Within the Fusion viewer, you'll also find pivot points for your Character, Word, and Line transforms. For easy access, a Pivot Points icon has been added to the popup toolbar, allowing you to toggle their visibility when adjusting rotation and scaling transformations.

The Z position-based sorting mode is a powerful addition to our toolkit. In contrast to the Distance mode, which sorts based on absolute distance from the camera, this mode sorts solely on the Z distance from the camera, irrespective of the relative orientation. With Z position, you can flexibly arrange overlapping text elements, ensuring precise layering and depth management.

Character-level styling has been refined to facilitate more intuitive offsets and translation transformations. This enhancement reduces the need for manual positioning and ensures a consistent application across multiple separate character selections.

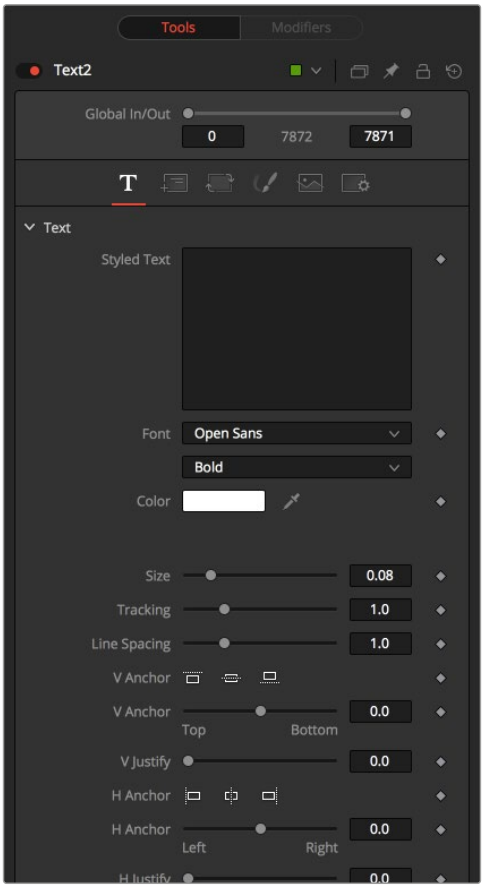
Basic Node Setup

The Text+ node can be combined with many other nodes to create motion graphics. Below it is connected as the foreground of a Merge node.



Text+ node connected as the foreground to a Merge node

Inspector



Text+ Text tab controls

Text Tab

The Text tab in the Inspector is divided into three sections: Text, Advanced Controls, and Tab Spacing. The Text section includes parameters that will be familiar to anyone who has used a word processor. It includes commonly used text formatting options. The Advanced controls are used for kerning options, and the Tab Spacing is used to define the location and alignment of tabs in the layout.

Styled Text

The edit box in this tab is where the text to be created is entered. Any common character can be typed into this box. The common OS clipboard shortcuts (Command-C or Ctrl-C to copy, Command-X or Ctrl-X to cut, Command-V or Ctrl-V to paste) will also work; however, right-clicking in the edit box displays a custom contextual menu. More information on these modifiers can be found at the end of this section.

The Styled Text contextual menu includes the following options:

- **Animate:** Used to animate the text over time.
- **Character Level Styling:** Used to change the font, color, size and transformations of individual characters or words through the Modifiers tab.
- **Comp Name:** Places the name of the composition in the Styled text box for creating slates.
- **Follower:** A text modifier used to ripple animation across each character of the text.
- **Publish:** Publishes the text for connection to other text nodes.
- **Text Scramble:** A text modifier used to randomize the characters in the text.
- **Text Timer:** A text modifier used to display a countdown or the current date and time.
- **Time Code:** A text modifier used to display Time Code for the current frame.
- **Connect To:** Used to connect the text to the published output of another node.

Font

Two Font menus are used to select the font family and typeface, such as Regular, Bold, and Italic.

Color

Sets the basic fill color of the text. This is the same control displayed in the Shading tab color swatch.

Size

This control is used to increase or decrease the size of the text. This is not like selecting a point size in a word processor. The size is relative to the width of the image.

Tracking

The Tracking parameter adjusts the uniform spacing between each character of text.

Line Spacing

Line Spacing adjusts the distance between each line of text. This is sometimes called leading in word-processing applications.

V Anchor

The vertical anchor controls consist of three buttons and a slider. The three buttons are used to align the text vertically to the top of the text, middle of the text, or bottom baseline. The slider can be used to customize the alignment. Setting the vertical anchor will affect how the text is rotated as well as the location for line spacing adjustments. This control is most often used when the Layout type is set to Frame in the Layout tab.

V Justify

The vertical justify slider allows you to customize the vertical alignment of the text from the V Anchor setting to full justification so it is aligned evenly along the top and bottom edges. This control is most often used when the Layout type is set to Frame in the Layout tab.

H Anchor

The horizontal anchor controls consist of three buttons and a slider. The three buttons justify the text alignment to the left edge, middle, or right edge of the text. The slider can be used to customize the justification. Setting the horizontal anchor will affect how the text is rotated as well as the location for tracking (leading) spacing adjustments. This control is most often used when the Layout type is set to Frame in the Layout tab.

H Justify

The horizontal justify slider allows you to customize the justification of the text from the H Anchor setting to full justification so it is aligned evenly along the left and right edges. This control is most often used when the Layout type is set to Frame in the Layout tab.

Direction

This menu provides options for determining the Direction in which the text is to be written, either horizontally or vertically in either direction. This allows certain Asian languages to flow properly during animation.

Line Direction

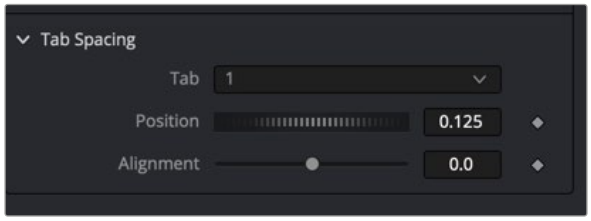
These menu options are used to determine the text flow from top to bottom, bottom to top, left to right, or right to left.

Underline and Strikeout

These buttons apply the addition of emphasis styles to the text.

Write On

This range control is used to quickly apply simple Write On and Write Off effects to the text. To create a Write On effect, animate the End portion of the control from 1 to 0 over the length of time required. To create a Write Off effect, animate the Start portion of the range control from 0 to 1.



Tab Spacing

Tab Spacing

The controls in the Tabs section are used to configure the horizontal screen positions of eight separate tab stops. Any tab characters in the text will conform to these positions.

You can add tabs directly in the Styled Text input as you type. You can also add tabs by copying from another document, such as Text on macOS or Notepad on Windows, and paste it into the text box.

Position

This control is used to set the horizontal position of the tab in the frame. The values range from -0.5 to 0.5, where 0 is the center. The position of the tab will be indicated in the viewer by a thin vertical white line when the Text node is selected. At the top of each tab line in the viewer is a handle. The handle can be used to position the tab manually.

Alignment

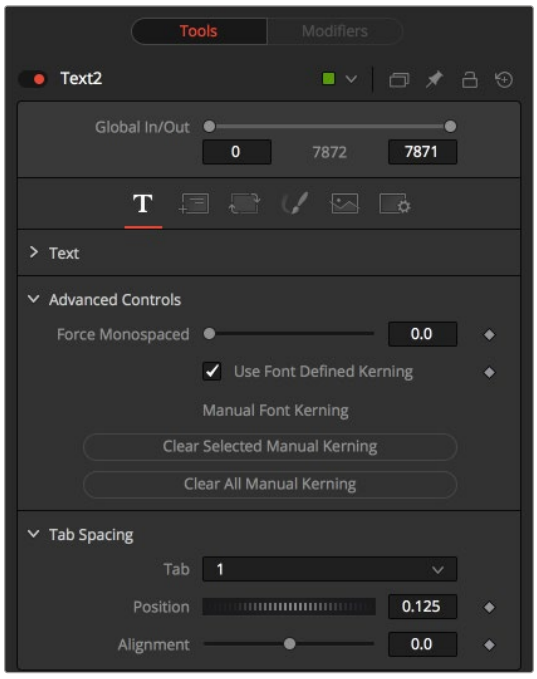
Each tab can be set either left aligned, right aligned, or centered. This slider ranges from -1.0 to 1.0, where -1.0 is a left-aligned tab, 0.0 is a centered tab and 1.0 is a right-aligned tab. Clicking the tab handles in the viewer will toggle the alignment of the tab among the three states.

Reading Direction

These options allow you to set the reading direction of the text, either automatically or manually. You can specify Left to Right languages like English, German, etc. or Right to Left languages like Arabic and Hebrew.

Force Monospaced

This slider control can be used to override the kerning (spacing between characters) defined in the font. Setting this slider to zero (the default value) will cause Fusion to rely entirely on the kerning defined with each character. A value of one will cause the spacing between characters to be completely even, or monospaced.



Advanced font controls

Use Font Defined Kerning

This enables kerning as specified in the TrueType font and is on by default.

Use Ligatures

If your font supports ligatures, you can activate them here by choosing All Scripts. Ligatures combine individual letters into single glyphs, like **ff** and **fl**. If you're animating individual text letters, often you want the ligature letters separated individually rather than as a single glyph, so None is the default for

Latin characters. Ligatures are required to render some languages like Arabic correctly, and use the Non-Latin setting.

Style Can Split Ligatures

Check this box If your font supports the option to split ligatures.

Stylistic Set

If your font includes stylistic sets, you can select them in the drop-down menu.

Font Features

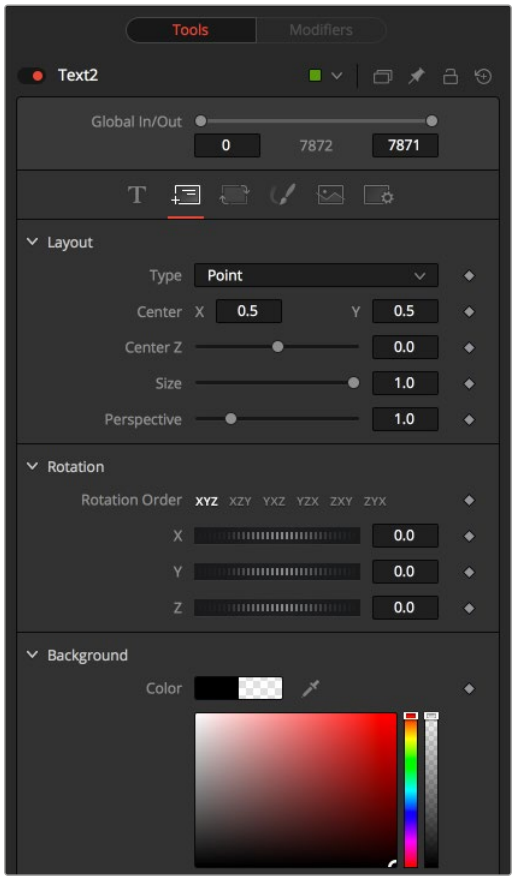
This allows you to set OpenType 4 letter tags to activate certain font features. For example, "smcp" will show **SMALL CAPITALS**, and "frac" will display fractions as ½ instead of 1/2. Not all features are supported by a particular font. A full list of OpenType feature codes can be found here: https://docs.microsoft.com/en-us/typography/opentype/spec/featurelist

Manual Font Kerning/Placement

Right-clicking on this label will display a contextual menu that can be used to animate the kerning of the text. See the “Toolbar” section of this node later in this chapter for details on manual kerning.

Layout Tab

The controls used to position the text are located in the Layout tab. One of four layout types can be selected using the Type drop-down menu.



Text+ Layout tab controls

- **Point:** Point layout is the simplest of the layout modes. Text is arranged around an adjustable center point.
- **Frame:** Frame layout allows you to define a rectangular frame used to align the text. The alignment controls are used for justifying the text vertically and horizontally within the boundaries of the frame.
- **Circle:** Circle layout places the text around the curve of a circle or oval. Control is offered over the diameter and width of the circular shape. When the layout is set to this mode, the Alignment controls determine whether the text is positioned along the inside or outside of the circle's edge, and how multiple lines of text are justified.
- **Path:** Path layout allows you to shape your text along the edges of a path. The path can be used simply to add style to the text, or it can be animated using the Position on Path control that appears when this mode is selected.

Center X, Y, and Z

These controls are used to position the center of the layout element in space. X and Y are onscreen controls, and Center Z is a slider in the node controls.

Size

This slider is used to control the scale of the layout element.

Perspective

This slider control is used to add or remove perspective from the rotations applied by the Angle X, Y, and Z controls.

Rotation

Rotation consists of a series of buttons allowing you to select the order in which 3D rotations are applied to the text. Angle dials can be used to adjust the angle of the Layout element along any axis.

Width and Height

The Width control is visible when the Layout mode is set to Circle or Frame. The Height control is visible only when the Layout mode is set to Frame. Width and Height controls are used to adjust the dimensions and aspect of the Layout element.

Fit Characters

This menu control is visible only when the Layout type is set to Circle. This menu is used to select how the characters are spaced to fit along the circumference.

Position on Path

The Position on Path control is used to control the position of the text along the path. Values less than 0 or greater than 1 will cause the text to move beyond the path in the same direction as the vector of the path between the last two keyframes.

Background Color

The text generated by this node is normally rendered with a transparent background. This Color Picker control can be used to set a background color.

Right-Click Here for Shape Animation

This label appears only when the Layout type is set to Path. It is used to provide access to a contextual menu that provides options for connecting the path to other paths in the node tree, and animating the shape of the path over time.

For more information, *see Chapter 11, "Animating with Motion Paths,"* in the Fusion Reference Manual.

Transform Tab

The Transform tab is used to move, rotate, shear and scale text based on a character, word, or line.



Text+ Transform tab controls

Transform

The Transform menu is used to determine the portion of the text affected by the transformations applied in this tab. Transformations can be applied to line, word, and character levels simultaneously. This menu is only used to keep the visible controls to a reasonable number.

- **Characters:** Each character of text is transformed along its own center axis.
- **Words:** Each word is transformed separately on the word's center axis.
- **Lines:** Each line of the text is transformed separately on that line's center axis.

Spacing

The Spacing slider is used to adjust the space between each line, word, or character. Values less than 1 will usually cause the characters to begin overlapping.

Pivot X, Y, and Z

This provides control over the exact position of the axis. By default, the axis is positioned at the calculated center of the line, word, or character. The Axis control works as an offset, such that a value of 0.1, 0.1 in this control would cause the axis to be shifted downward and to the right for each of the text elements. Positive values in the Z-axis slider will move the axis of rotation further along the axis (away from the viewer). Negative values will bring the axis of rotation closer.

Rotation

These buttons are used to determine the order in which transforms are applied. X, Y, and Z would mean that the rotation is applied to X, then Y, and then Z.

X, Y, and Z

These controls can be used to adjust the angle of the text elements in any of the three dimensions.

Shear X and Y

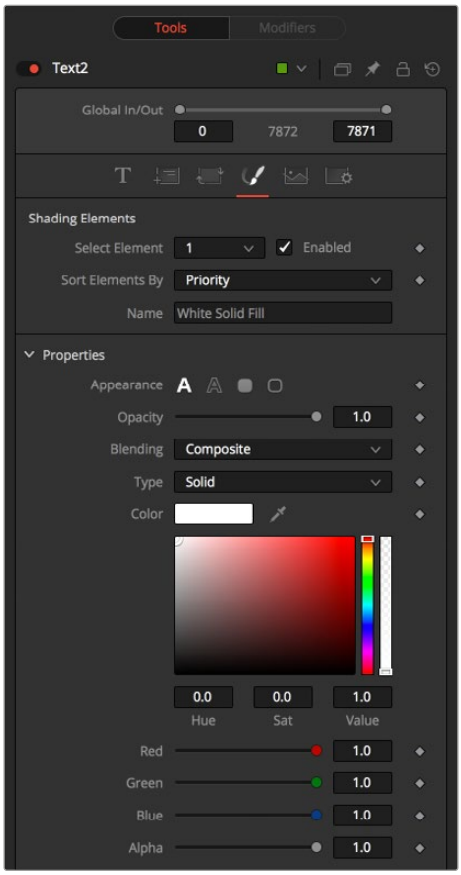
Adjust these sliders to modify the slanting of the text elements along the X- and Y-axis.

Size X and Y

Adjust these sliders to modify the size of the text elements along the X- and Y-axis.

Shading Tab

The Shading tab provides controls to adjust the shading, texture, and softness of the text. Transformations can be controlled from this tab as well, applying additional transformations to as many as eight separate text shading elements independently.



Text+ Shading tab controls

Shading Element

The eight number values in the menu are used to select the element affected by adjustments in this tab.

Enabled

Select this checkbox to enable or disable each layer of shading elements. Element 1, which is the fill color, is enabled by default. The controls for a shading element will not be displayed unless this checkbox is selected.

Sort By

This menu allows you to sort the shading elements by number priority, with 1 being the topmost element and 8 being the bottommost element, or Z depth, based on the Z Position parameter.

Name

This text label can be used to assign a more descriptive name to each shading element you create.

Appearance

The four Appearance buttons determine how the shading element is applied to the text. Different controls will appear below depending on the appearance type selected.

- **Text Fill:** The shading element is applied to the entire text. This is the default mode.
- **Text Outline:** The shading element is drawn as an outline around the edges of the text.
- **Border Fill:** The shading element fills a border surrounding the text. Five additional controls are provided with this shading mode.
- **Border Outline:** The Border Outline mode draws an outline around the border that surrounds the text. It offers several additional controls.

Opacity

The Opacity slider controls the overall transparency of the shading element. It is usually better to assign opacity to a shading element than to adjust the Alpha of the color applied to that element.

Blending

This menu is used to select how the renderer deals with an overlap between two characters in the text.

- **Composite:** Merges the shading over the top of itself.
- **Solid:** Sets the pixels in the overlap region to opaque.
- **Transparent:** Sets the pixels in the overlap region to transparent.

Thickness

(Outline only) Thickness adjusts the thickness of the outline. Higher values equal thicker outlines.

Adapt Thickness to Perspective

(Outline only) Selecting this checkbox will cause your outline to become thinner where the text is farther away from the camera, and thicker where it is closer. This will create a more realistic outline for text transformed in 3D but takes significantly longer to render.

Outside Only

(Outline only) Selecting this checkbox will cause the outline to be drawn only on the outside edge of the text. By default, the outline is centered on the edge and partially overlaps the text.

Join Style

(Outline only) These buttons provide options for how the corners of the outline are drawn. Options include Sharp, Rounded, and Beveled.

Line Style

(Outline only) This menu offers additional options for the style of the line. Besides the default solid line, a variety of dash and dot patterns are available.

Level

(Border Fill only) This is used to control the portion of the text border filled.

- **Text:** This draws a border around the entire text.
- **Line:** This draws a border around each line of text.
- **Word:** This draws a border around each word.
- **Character:** This draws a border around each character.

Extend Horizontal and Extend Vertical

(Border only) Use this slider to change the dimensions of each border.

Round

(Border Fill and Border Outline only) This slider is used to round off the edges of the border.

Color Types

Besides solid shading, it is also possible to use a gradient fill or map an external image onto the text. This menu is used to determine if the color of the shading element is derived from a user-selected color or gradient, or if it comes from an external image source. Different controls will be displayed below depending on the Color Type selected.

- **Solid:** When the Type menu is set to Solid mode, color selector controls are provided to select the color of the text.
- **Image:** The output of a node in the node tree will be used to texture the text. The node used is chosen using the Color Image control that is revealed when this option is selected.
- **Gradient:** When the Type menu is set to Gradient, additional controls are displayed where the gradient colors can direction can be customized.

Using the Gradient Color Bar

The gradient color bar is used to select the colors for the gradient. The default two color stops set the start and end colors. You can change the colors used in the gradient by selecting the color stop, and then using the Eyedropper or color swatch to set a new color.

You can add, move, copy and delete color using the gradient bar.

To modify one of the colors, select the triangle below the color on the bar.

To add a color stop to the gradient bar:

1. Click anywhere along the bottom of the gradient bar.
2. Use the Eyedropper or color swatch to select the color for the color stop.

To move a color stop on the gradient bar:

Drag a color stop left or right along the bar.

To copy a color stop: Hold Command (macOS) or Ctrl (Windows) while you drag a color stop.
To delete a color stop: Drag the color stop up past the gradient bar.

Image Source

(Image Mode only) The Image Source menu includes three options for acquiring the image used to fill the text.

- **Tool:** Displays a Color image text field where you can add a tool from the node tree as the fill for text.
- **Clip:** Provides a Browse button to select a media file from your hard drive as the fill for text.
- **Brush:** Displays a Color Brush menu where you can select one of Fusion's paint brush bitmaps as the fill for text.

Color Image/Color Brush

(Image Mode only) The Color Image text box and Color Brush menu are used to select the tool or paint brush that will fill the text. In the Color Image text box, enter the name of the node in the node tree that will provide the image. You can type the name in with the keyboard, drag the node from the node tree into the text box, or right-click and select Connect To from the contextual menu to select the image to be used. For the Color Brush, select the Brush from the menu.

Image Sampling

(Image Mode only) This menu is used to select the sampling type for shading rendering and transformations. The default of Pixel shading is sufficient for 90% of tasks. To reduce detectable aliasing in the text, set the sampling type to Area. This is slower but may produce better-quality results. A setting of None will render faster, but with no additional sampling applied so the quality will be lower.

Image Edges

(Image Mode only) This menu is used to choose how transformations applied to image shading elements are handled when they wrap off the text's edges.

Shading Mapping

(Image Mode only) This menu is used to select whether the entire image is stretched to fill the text or scaled to fit, maintaining the aspect ratio but cropping part of the image as needed.

Mapping Angle

(Image and Gradient Modes only) This control rotates the image or gradient on the Z-axis.

Mapping Size

(Image and Gradient Modes only) This control scales the image or gradient.

Mapping Aspect

(Image and Gradient Modes only) This control vertically stretches or shrinks the image or gradient.

Mapping Level

(Image and Gradient Modes only) The Mapping Level menu is used to select how the image is mapped to the text.

- **Full Image:** Applies the entire image to the text.
- **Text:** Applies the image to fit the entire set of text.
- **Line:** Applies the image per line of text.
- **Word:** Applies the image per each word of text.
- **Character:** Applies the image per individual character.

Softness X and Y

These sliders control the softness of the text outline used to create the shading element. Control is provided for the X- and Y-axis independently.

Apply Softness to Fill Color

Selecting this checkbox will cause blur (softness) to be applied to the shading element. The effect is best seen when applied to a shading element using an external image.

Softness Glow

This slider will apply a glow to the softened portion of the shading element.

Softness Blend

This slider controls the amount that the result of the softness control is blended back with the original. It can be used to tone down the result of the soften operation.

Priority Back/Front

Only enabled when the Sort By menu is set to Priority, this slider overrides the priority setting and determines the layer's order for the shading elements. Slide the control to the right to bring an element closer to the front. Move it to the left to tuck one shading element behind another.

Offset X, Y, and Z

These controls are used to apply offset from the text's global center (as set in the Layout tab) for the shading elements. A value of X0.0, Y0.1 in the coordinate controls would place the shading element centered with 10 percent of the image further down the screen along the Y-axis. Positive values in the Z-Offset slider control will push the center further away from the camera, while positive values will bring it closer to the camera.

Pivot X, Y, and Z

These controls are used to set the exact position of the axis for the currently selected shading element. By default, the axis is positioned at the calculated center of the line, word, or character. The axis control works as an offset, such that a value of 0.1, 0.1 in this control would cause the axis to be shifted downward and to the right for the shading element. Positive values in the Z-axis slider will move the axis of rotation further along the axis (away from the viewer). Negative values will bring the axis of rotation closer.

Rotation X, Y, and Z

These controls are used to adjust the angle of the currently selected shading element in any of the three dimensions.

Shear X and Y

Adjust these sliders to modify the slanting of the currently selected shading element along the X and Y axis.

Size X and Y

Adjust these sliders to modify the size of the currently selected shading element along the X and Y axis.

Common Controls

Image and Settings Tabs

The Image and Settings tabs in the Inspector are duplicated in many Generator nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Text+ Toolbar

When the Text node is selected, a toolbar will appear in the viewer. Each button is described below from left to right.



Text+ toolbar

Allow Typing in Viewer

When selected, you can type and edit text directly in the viewer. Clicking on the text in the viewer will move the insert cursor within the text. Left and Right Arrow keys will move the insert cursor between characters. Using the Up and Down Arrows will move between text lines.

Allow Manual Kerning

The Manual Kerning button overrides the automatic kerning normally applied to text. Clicking the small red handle at the bottom of any character selects the character for kerning. You can also draw a selection rectangle around the characters you want to kern. Once any characters are selected, hold down the Option (macOS) or Alt key (Windows) while pressing the Left or Right Arrow key to make small adjustments to the kerning of the selected characters. Hold the Alt + Shift (Windows) or Option + Shift (macOS) keys down while pressing arrow keys to move the character in larger increments.

To animate the position of each character, right-click on the control label Manual Font Kerning in the Inspector’s Advanced Controls and select Animate from the contextual menu. A new key will be set on the animation spline each time a character is moved. All characters are animated with the same spline, as with polyline mask animation.

No Text Outline

When this button is selected, it disables the drawing of any outline around the edges of the text. The outline is not a part of the text; it is an onscreen control used to help identify the position of the text. This is a three-way toggle with the Text Outline Outside Frame Only, and Show Always Text Outline buttons.

Text Outline Outside Frame Only

This button draws an outline around the edges of text, which is outside the visible frame. This is useful for locating text that has moved offscreen and is no longer rendering a visible result. This is a three-way toggle with the No Text Outline, and Show Always Text Outline buttons.

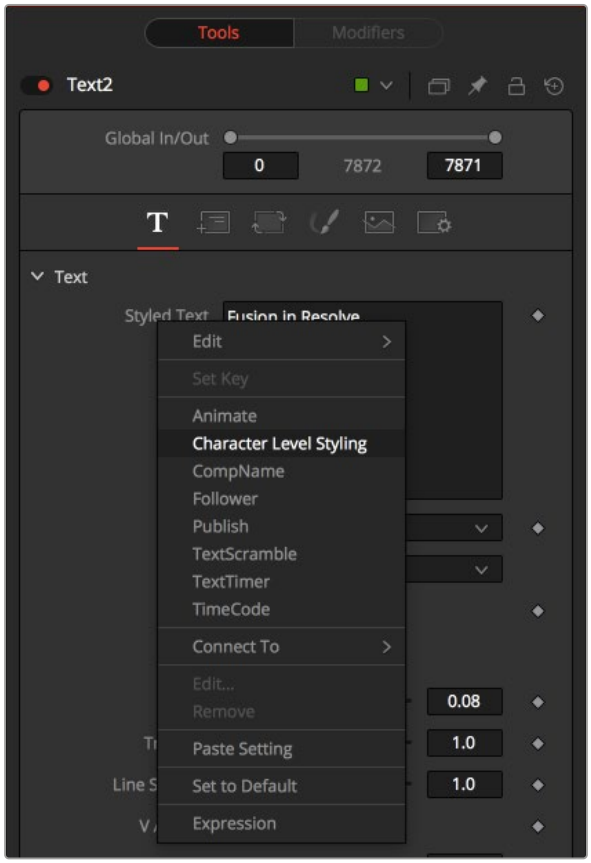
Show Always Text Outline

This button draws an outline around the edges of text, whether or not the text is visible within the frame. This is a three-way toggle with the No Text Outline, and Text Outline Outside Frame Only buttons.

Image and Settings Tabs

The Image and Settings tabs in the Inspector are duplicated in many Generator nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Text+ Modifiers



Text+ modifiers

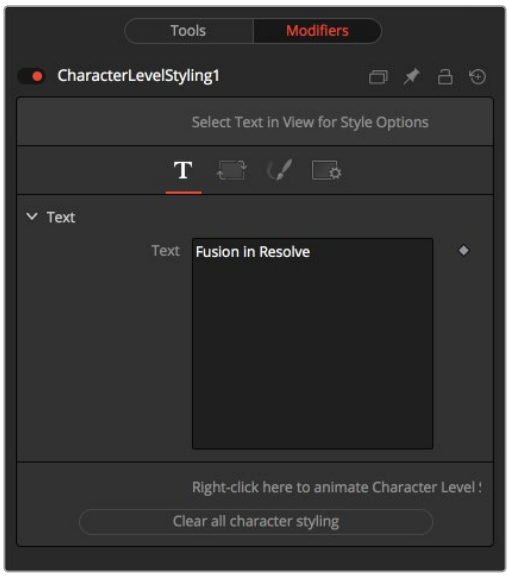
Text modifiers can be assigned by right-clicking in the Styled Text box and selecting a modifier from the contextual menu. Once a modifier is selected, its controls are found in the Modifiers tab at the top of the Inspector.

Character Level Styling

The Character Level Styling modifier works only on Text+ nodes. Once applied, individual characters can be selected directly in the viewer, and different text attributes can be applied to them using the controls in the Modifiers tab.

NOTE: Character Level Styling can only be directly applied to Text+ nodes, not to Text 3D nodes. However, styled text from a Text+ node can be applied to a Text 3D node by copying the Text+, right-clicking on the Text 3D, and choosing Paste Settings.

Inspector



Character Level Styling Modifiers tab

Text Tab

The Styled Text box in the Modifiers tab displays the same text in the Tools tab of the Text+ Inspector. However, individual characters you want to modify cannot be selected in the Styled Text box; they must be selected in the viewer. Once text is selected in the viewer, the Text tab includes familiar text formatting options that will apply only to the selected characters.

Clear Character Styling on Selection

All changes made to the currently selected characters will be reset.

Clear All Character Styling

All character attributes will be reset to their original values.

Transform Tab and Shading Tab

For details on these Text+ tabs, see the “Text+” section above.

Comp Name

The Comp Name sets the styled text to become the current Composition Name. This is quite useful to automate burn-ins for daily renderings. See also the TimeCode modifier. It can be applied by right-clicking in the Styled Text field of a Text+ node and selecting Comp Name.

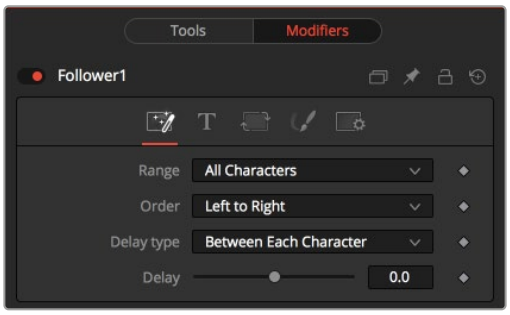
Controls

This modifier has no controls.

Follower

The Follower modifier allows sequencing text animations. The modifier is applied by right-clicking in the Styled Text field of a Text+ node and selecting Follower. In the Modifiers tab, you start by animating the parameters of the text (note that changing any parameter in the Modifiers tab will not be visible unless a keyframe is added.) Then, in the Timing tab you set the animation’s delay between characters.

Inspector



Follower Timing tab

Timing Tab

Once the text is animated using the controls in the Modifiers tab, the Timing tab is used to choose the direction and offset of the animation.

Range

The Range menu is used to select whether all characters should be influenced or only a selected range. To set the range, you can drag-select over the characters directly in the viewer.

Order

The Order menu determines in which direction the characters are influenced. Notice that empty spaces are counted as characters as well. Available options include:

- **Left to right:** The animation ripples from left to right through all characters.
- **Right to left:** The animation ripples from right to left through all characters.
- **Inside out:** The animation ripples symmetrically from the center point of the characters toward the margin.
- **Outside in:** The animation ripples symmetrically from the margin toward the center point of the characters.
- **Random but one by one:** The animation is applied to randomly selected characters but only influences one character at a time.
- **Completely random:** The animation is applied to randomly selected characters, influencing multiple characters at a time.
- **Manual curve:** The affected characters can be specified by sliders.

Delay Type

Determines what sort of delay is applied to the animation. Available options include:

- **Between Each Character:** The more characters there are in your text, the longer the animation will take to the end. A setting of 1 means the first character starts the animation, and the second character starts 1 frame later, the third character starts 1 frame after the second, and so on.
- **Between First and Last Character:** No matter how many characters are in your text, the animation will always be completed in the selected amount of time.

Clear All Character Styling

All character attributes will be reset to their original values.

Text Controls, Alignment, Transform, and Shading Tabs

In these tabs, the actual animation for the characters is done. Observe that simply changing a value in these tabs will have no influence at all. The value must be animated for the effect to show.

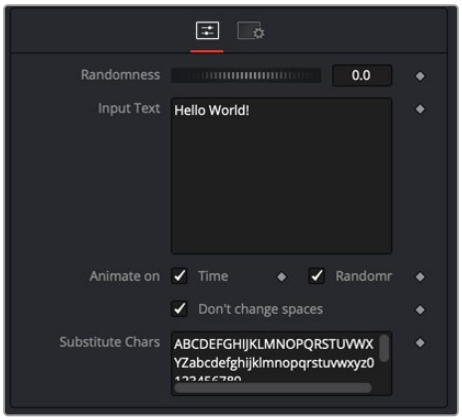
For a detailed description on the various parameters, see the Text+ node documentation.

Text Scramble

The Text Scramble randomly replaces the characters with others from a user definable set. It can be applied by right-clicking into the Styled Text field of a Text+ node and selecting Text Scramble.

The Controls for the Text Scramble are then adjusted in the Modifiers tab.

Inspector



Text Scramble modifier Controls tab

Controls Tab

The Controls tab in the Text Scramble modifier is used to enter text and scramble it using the Randomness control. The scrambled characters are taken from the Substitute Chars field at the bottom of the Inspector.

Randomness

Defines how many characters are exchanged randomly. A value of 0 will change no characters at all. A value of 1 will change all characters in the text. Animating this thumbwheel to go from 0 to 1 will gradually exchange all characters.

Input Text

This reflects the original text in the Text+ Styled Text field. Text can be entered either here or in the Text+ node.

Animate on Time

When enabled, the characters are scrambled randomly on every new frame. This switch has no effect when Randomness is set to 0.

Animate on Randomness

When enabled, the characters are scrambled randomly on every new frame, when the Randomness thumbwheel is animated.

This switch has no effect when Randomness is set to 0.

Don’t Change Spaces

When enabled, spaces are not scrambled, allowing the length of the single words to stay the same, although their characters get scrambled around.

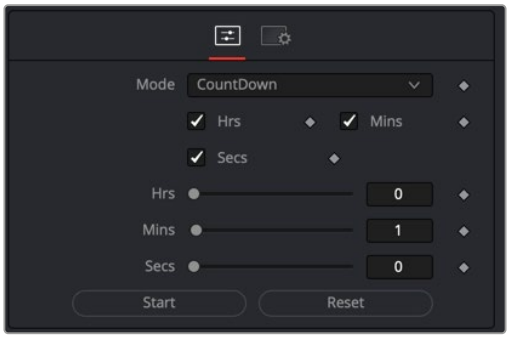
Substitute Chars

This field contains the characters used to scramble the text.

Text Timer

The Text Timer makes the Text+ node either a Countdown, a Timer, or a Clock. This is useful for onscreen real-time displays or to burn in the creation time of a frame into the picture. It can be applied by right-clicking in the Styled Text field of a Text+ node and selecting Text Timer.

Inspector



Text Timer modifier Controls tab

Controls Tab

The Controls tab for the Text Timer modifier is used to set up the type of time display that is generated by this modifier.

Mode

This menu sets the mode the timer is working in. The choices are CountDown, Timer, and Clock. In Clock mode, the current system time will be displayed.

Hrs, Mins, Secs (Checkboxes)

Defines which parts of the clock should be shown onscreen.

Hrs, Mins, Secs (Sliders)

Sets the start time for the CountDown and Timer mode.

Start

Starts the Counter or Timer. Toggles to Stop once the timer is running.

Reset

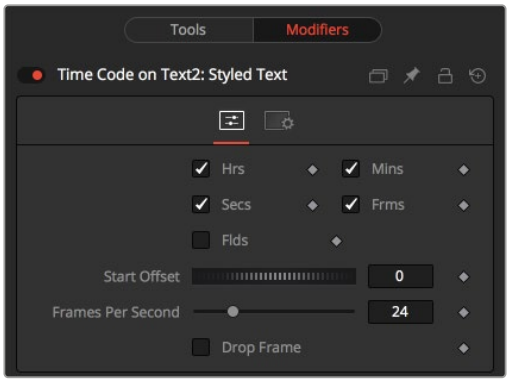
Resets the Counter and Timer to the values set by the sliders.

Time Code

The Time Code only works on Text+ nodes. It sets the Styled text to become a counter based on the current frame. This is quite useful for automating burn-ins for daily renderings.

It can be applied by right-clicking in the Styled Text field of a Text+ node and selecting Time Code.

Inspector



Time Code modifier Controls tab

Controls Tab

The Controls tab for the Time Code modifier is used to set up the time code display that is generated by this modifier.

Hrs, Mins, Secs, Frms, Flds

Activate or deactivate these options to customize the time code display to show hours, minutes, seconds, frames, and fields, respectively. Activating frames will only give you a plain frame counter.

Start Offset

Introduce a positive or negative offset to Fusion's current time to match up with existing time codes.

Frames per Second

This should match with your composition's FPS setting to provide accurate time measurement.

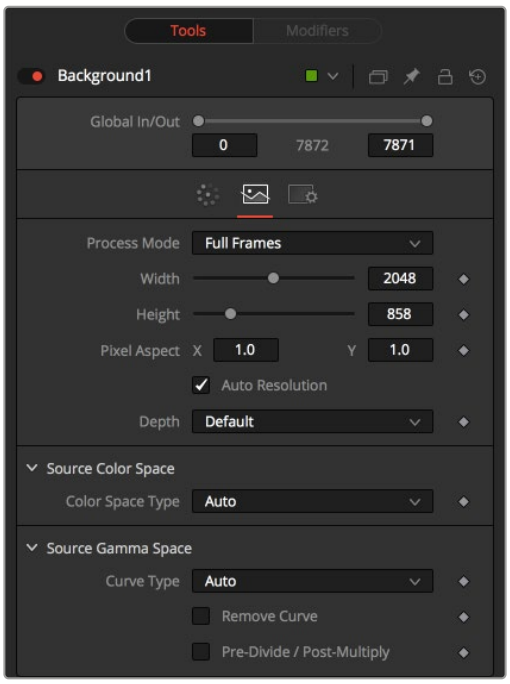
Drop Frame

Activate this checkbox to match the time code with footage that has drop frames—for example, certain NTSC formats.

The Common Controls

Nodes that generate images share a number of identical controls in the Inspector. This section describes controls that are common among Generator nodes.

Inspector



Background node Image tab

Image Tab

The controls in this tab are used to set the resolution, color depth, and pixel aspect of the image produced by the node.

Process Mode

Use this menu control to select the Fields Processing mode used by Fusion to render changes to the image. The default option is determined by the Has Fields checkbox control in the Frame Format preferences.

Global In and Out

Use this control to specify the position of this node within the project. Use Global In to specify which frame that starts the clip and Global Out to specify which frame ends the clip (inclusive) within the project's Global Range.

The node will not produce an image on frames outside this range.

Use Frame Format Settings

When this checkbox is selected, the width, height, and pixel aspect of the image created by the node will be locked to values defined in the composition's Frame Format preferences. If the Frame Format preferences change, the resolution of the image produced by the node will change to match. Disabling this option can be useful to build a composition at a different resolution than the eventual target resolution for the final render.

Width/Height

This pair of controls is used to set the Width and Height dimensions of the image to be created by the node.

Pixel Aspect

This control is used to specify the Pixel Aspect ratio of the created images. An aspect ratio of 1:1 would generate a square pixel with the same dimensions on either side (like a computer display monitor), and an aspect of 0.9:1 would create a slightly rectangular pixel (like an NTSC monitor).

NOTE: Right-click on the Width, Height, or Pixel Aspect controls to display a menu listing the file formats defined in the preferences Frame Format tab. Selecting any of the listed options will set the width, height, and pixel aspect to the values for that format, accordingly.

Depth

The Depth drop-down menu is used to set the pixel color depth of the image created by the Creator node. 32-bit pixels require 4X the memory of 8-bit pixels but have far greater color accuracy. Float pixels allow high dynamic range values outside the normal 0...1 range, for representing colors that are brighter than white or darker than black.

Source Color Space

You can use the Source Color Space menu to set the Color Space of the footage to help achieve a linear workflow. Unlike the Gamut tool, this doesn't perform any actual color space conversion, but rather adds the source space data into the metadata, if that metadata doesn't exist. The metadata can then be used downstream by a Gamut tool with the From Image option, or in a Saver, if explicit output spaces are defined there. There are two options to choose from:

- **Auto:** Automatically reads and passes on the metadata that may be in the image.
- **Space:** Displays a Color Space Type menu where you can choose the correct color space of the image.

Source Gamma Space

Using the Curve type menu, you can set the Gamma Space of the footage and choose to remove it by way of the Remove Curve checkbox when working in a linear workflow. There are three choices in the Curve type menu:

- **Auto:** Automatically reads and passes on the metadata that may be in the image.
- **Space:** Displays a Gamma Space Type menu where you can choose the correct gamma curve of the image.
- **Log:** Brings up the Log/Lin settings, similar to the Cineon tool.

For more information, *see Chapter 39, "Film Nodes,"* in the Fusion Reference Manual.

Remove Curve

Depending on the selected Gamma Space or on the Gamma Space found in Auto mode, the Gamma Curve is removed from, or a log-lin conversion is performed on, the material, effectively converting it to a linear output space.

Fast Noise Image Tab Options

The following controls are specific the Image tab in the Fast Noise Node.

Mask Map Inputs

These external connections allow you to use masks to control the value of the Noise Detail and Brightness controls individually for each pixel. This can allow some interesting and creative effects.

Noise Detail Map

A soft-edged mask connected to the Noise Detail Map will give a flat noise map (zero detail) where the mask is black, and full detail where it is white, with intermediate values smoothly reducing in detail. It is applied before any gradient color mapping. This can be very helpful for applying maximum noise detail in a specific area, while smoothly falling off elsewhere.

Noise Brightness Map

A mask connected to this input can be used to control the noise map completely, such as boosting it in certain areas, combining it with other textures, or if Detail is set to 0, replacing the Perlin Noise map altogether.

Settings Tab

The Settings Tab in the Inspector can be found on every tool in the Color category. The Settings controls are even found on third-party Color-type plugin tools. The controls are consistent and work the same way for each tool, although some tools do include one or two individual options, which are also covered here.



Common Generator settings

Blend

The Blend control is used to blend between the tool's original image input and the tool's final modified output image. When the blend value is 0.0, the outgoing image is identical to the incoming. Normally, this will cause the tool to skip processing entirely, copying the input straight to the output.

Process When Blend Is 0.0

The tool is processed even when the input value is zero. This can be useful if processing of this node is scripted to trigger another task, but the value of the node is set to 0.0.

Red/Green/Blue/Alpha Channel Selector

These four buttons are used to limit the effect of the tool to specified color channels. This filter is often applied after the tool has been processed.

For example, if the red button on a Blur tool is deselected, the blur will first be applied to the image, and then the red channel from the original input will be copied back over the red channel of the result.

There are some exceptions, such as tools for which deselecting these channels causes the tool to skip processing that channel entirely. Tools that do this will generally possess a set of identical RGBA buttons on the Controls tab in the tool. In this case, the buttons in the Settings and the Controls tabs are identical.

Apply Mask Inverted

Enabling the Apply Mask Inverted option inverts the complete mask channel for the tool. The mask channel is the combined result of all masks connected to or generated in a node.

Multiply by Mask

Selecting this option will cause the RGB values of the masked image to be multiplied by the mask channel's values. This will cause all pixels of the image not in the mask (i.e., set to 0) to become black/transparent.

Use Object/Use Material (Checkboxes)

Some 3D software can render to file formats that support additional channels. Notably, the EXR file format supports Object ID and Material ID channels, which can be used as a mask for the effect. These checkboxes determine whether the channels will be used if present. The specific Material ID or Object ID affected is chosen using the next set of controls.

Correct Edges

This checkbox appears only when the Use Object or Use Material checkboxes are selected. It toggles the method used to deal with overlapping edges of objects in a multi-object image. When enabled, the Coverage and Background Color channels are used to separate and improve the effect around the edge of the object. If this option disabled (or no Coverage or Background Color channels are available), aliasing may occur on the edge of the mask.

For more information *see Chapter 18, "Understanding Image Channels,"* in the Fusion Reference Manual.

Object ID/Material ID (Sliders)

Use these sliders to select which ID will be used to create a mask from the object or material channels of an image. Use the Sample button in the same way as the Color Picker: to grab IDs from the image displayed in the view. The image or sequence must have been rendered from a 3D software package with those channels included.

Motion Blur

- **Motion Blur:** This toggles the rendering of Motion Blur on the tool. When this control is toggled on, the tool's predicted motion is used to produce the motion blur caused by the virtual camera's shutter. When the control is toggled off, no motion blur is created.
- **Quality:** Quality determines the number of samples used to create the blur. A quality setting of 2 will cause Fusion to create two samples to either side of an object's actual motion. Larger values produce smoother results but increase the render time.
- **Shutter Angle:** Shutter Angle controls the angle of the virtual shutter used to produce the motion blur effect. Larger angles create more blur but increase the render times. A value of 360 is the equivalent of having the shutter open for one whole frame exposure. Higher values are possible and can be used to create interesting effects.
- **Center Bias:** Center Bias modifies the position of the center of the motion blur. This allows the creation of motion trail effects.
- **Sample Spread:** Adjusting this control modifies the weighting given to each sample. This affects the brightness of the samples.

Use GPU

The user GPU menu has three settings. Setting the menu to Disable turns off hardware-accelerated rendering using the graphics card in your computer. Enabled uses the hardware, and Auto uses a capable GPU if one is available and falls back to software rendering when a capable GPU is not available

Comments

The Comments field is used to add notes to a tool. Click in the field and type the text. When a note is added to a tool, a small red square appears in the lower-left corner of the node when the full tile is displayed, or a small text bubble icon appears on the right when nodes are collapsed. To see the note in the Node Editor, hold the mouse pointer over the node to display the tooltip.

Scripts

Three Scripting fields are available on every tool in Fusion from the Settings tab. They each contain edit boxes used to add scripts that process when the tool is rendering. For more details on scripting nodes, please consult the Fusion scripting documentation.

Chapter 45

I/O Nodes

This chapter details the input and output of media using Loader and Saver nodes within Fusion Studio as well as the MediaIn and MediaOut nodes in DaVinci Resolve.

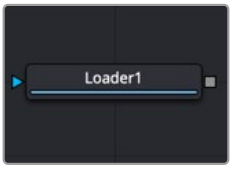
> The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

Contents

Loader Node [Ld] 1171
MediaIn Node [MI] 1179
MediaOut Node [MO] 1183
Saver Node [SV] 1187
The Common Controls 1194

FUSION

Loader Node [Ld]



The Loader node

Loader Node Introduction

NOTE: The Loader node in DaVinci Resolve is only used for importing EXR files.

When using Fusion Studio, the Loader node is the node you use to select and load footage from your hard drive into the Node Editor. There are three ways to add a Loader node, and consequently a clip, to the Node Editor.

- Add the Loader from the Effects Library or toolbar (Fusion Studio only), and then use Loader's file browser to bring a clip into the Node Editor
- Drag clips from an OS window directly into the Node Editor, creating a Loader node in the Node Editor.
- Choose File > Import > Footage (Fusion Studio only), although this method creates a new composition as well as adds the Loader node to the Node Editor.

When a Loader is added to the Node editor, a File dialog is displayed automatically to allow the selection of a clip from your hard drives.

NOTE: You can disable the automatic display of the file browser by disabling Auto Clip Browse in the Global > General Preferences.

Once clips are brought in using the Loader node, the Loader is used for trimming, looping, and extending the footage, as well as setting the field order, pixel aspect, and color depth. The Loader is arguably the most important tool in Fusion Studio.

Inputs

The single input on the Loader node is for an effect mask to crop the image brought in by the Loader.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the loaded image to appear only within the mask. An effects mask is applied to the tool after the tool is processed.

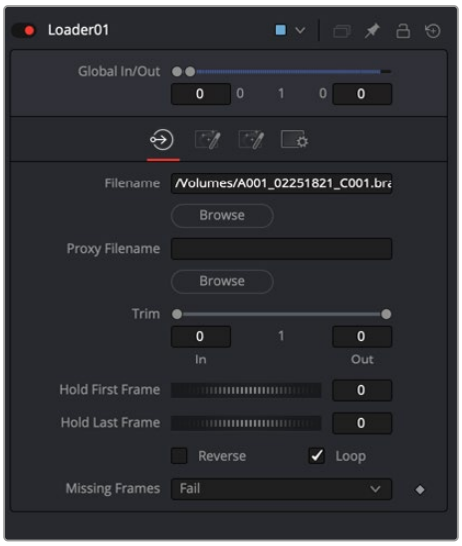
Basic Node Setup

The Loader node is a 2D image of any format supported in Fusion Studio. It is limited to an EXR format in DaVinci Resolve. Below, the LOADER imports an image, which is then masked using an Ellipse matte. The output of the masked LOADER is passed onto 2D image-processing nodes.



A Loader node used for importing images

Inspector



Loader node File tab

File Tab

The File tab for the Loader includes controls for trimming, creating a freeze frame, looping, and reversing the clip. You can also reselect the clip that the Loader links to on your hard drive.

Global In and Out

The Global In and Out handles are used to specify the position of this node within the project. Use Global In to specify the frame on which that the clip starts and use Global Out to specify the frame on which the clip ends within the project's global range. The node does not produce an image on frames outside of this range.

If the Global In and Out values are decreased to the point where the range between the In and Out values is smaller than the number of available frames in the clip, Fusion automatically trims the clip by adjusting the Trim range control. If the Global In/Out values are increased to the point where the range between the In and Out values is larger than the number of available frames in the clip, Fusion automatically lengthens the clip by adjusting the Hold First/Last Frame controls. Extended frames are visually represented in the range control by changing the color of the held frames to green in the control.

To slide the clip in time or move it through the project without changing its length, place the mouse pointer in the middle of the range control and drag it to the new location, or enter the value manually in the Global In value box.

Filename

The Filename field shows the file path of the clip imported to the Node Editor by the Loader node. Clicking on the Browse button opens a standard file browser. The path to the footage can also be typed directly using the field provided. The text box supports filename completion. As the name of a directory or file is typed in the text box, Fusion displays a pop-up that lists possible matches. Use the arrow keys to select the correct match and complete the path.

NOTE: Loading image sequences is common practice for compositing, whether the image sequence comes from a 3D renderer or a digital cinema camera. If the last part of a file's name is a number (not counting the file extension), Fusion automatically scans the directory looking for files that match the sequence. For example, the following filenames would be valid sequences:

```
image.0001.braw , image.0002.braw, image.0003.braw …
```

or

```
image151.exr , image152.exr, image153.exr …
```

The following would not be considered a sequence since the last characters are not numeric.

```
shot.1.fg.jpg, shot.2.fg.jpg, shot.3.fg.jpg
```

It is not necessary to select the first file in the sequence. Fusion searches the entire folder for files matching the sequence in the selected filename. Also, Fusion determines the length of the sequence based on the first and last numeric value in the filenames. Missing frames are ignored. For example, if the folder contains two files with the following names:

```
image.0001.exr, image.0100.exr
```

Fusion sees this as a file sequence with 100 frames, not an image sequence containing two frames. The Missing Frames drop-down menu is used to choose how Fusion handles missing frames.

The Trim In/Trim Out control's context menu can also be used to force a specific clip length or to rescan the folder. Both controls are described in greater detail below.

Occasionally, you want to load only a single frame out of a sequence—e.g., a photograph from a folder containing many other files as well. By default, Fusion detects those as a sequence, but if you hold Shift while dragging the file from the OS window to the Node Editor, Fusion takes only that specific file and disregards any sequencing.

Proxy Filename

The Proxy Filename control only appears once the filename control points to a valid clip. The Proxy Filename can specify a clip that is loaded when the Proxy mode is enabled. This allows smaller versions of the image to be loaded to speed up file I/O from disk, and processing. For example, create a 1/4-scale version of an 8K EXR sequence to use as EXR proxy files. Whenever the Proxy mode of the Composition is enabled, the smaller resolution proxy clip is loaded from disk, and all processing is performed at the lower resolution, significantly improving render times. This is particularly useful when working with large RAW plates stored on a remote file server. Lower-resolution versions of the plates can be stored locally, reducing network bandwidth, interactive render times, and memory usage.

The proxy clip must have the same number of frames as the source clip, and the sequence numbers for the clip must start and end on the same frame numbers. It is strongly suggested that the proxies are the same format as the main files. In the case of formats with options, such as Cineon, DPX, and OpenEXR, the proxies use the same format options as the primary.

Trim

The Trim range control is used to trim frames from the start or end of a clip. Adjust the Trim In to remove frames from the start and adjust Trim Out to specify the last frame of the clip. The values used here are offsets. A value of 5 in Trim In would use the fifth frame in the sequence as the start, ignoring the first four frames. A value of 95 would stop loading frames after the 95th.

Hold First Frame/Hold Last Frame:

The Hold First Frame and Hold Last Frame controls hold the first or last frame of the clip for the specified amount of frames. Held frames are included in a loop if the footage is looped.

Reverse

Select this checkbox to reverse the footage so that the last frame is played first, and the first frame is played last.

Loop

Select this checkbox to loop the footage until the end of the project. Any lengthening of the clip using Hold First/Last Frame or shortening using Trim In/Out is included in the looped clip.

Missing Frames

The Missing Frames menu determines the Loader behavior when a frame is missing or is unable to load for any reason.

- **Fail:** The Loader does not output any image unless a frame becomes available. Rendering is aborted.
- **Hold Previous Output:** The last valid frame is held until a frame becomes available again. This fails if no valid frame has been seen—for example, if the first frame is missing.
- **Output Black:** Outputs a black frame until a valid frame becomes available again.
- **Wait:** Fusion waits for the frame to become available, checking every few seconds. This is useful for rendering a composition simultaneously with a 3D render. All rendering ceases until the frame appears.

The Magic Comp Variable

Loaders and Savers use the absolute file paths for the location of media. However, if you are using Fusion Studio, you can use a file path that is relative to the saved composition location. The Comp variable works for Loaders and Savers and helps you to keep your work organized. Entering Comp:\ in place of the full file path name is a shortcut for the folder where your Fusion composition document is saved.

You can either enter the Comp:\ manually into the filename field of a Loader, or turn on the Enable Reverse Mapping of Paths Preferences checkbox in the Path Map preferences. Enabling the Path Map preference check box will use the Comp:\ automatically.

So as long as all your source footage is stored in subfolders of your Comp folder, Fusion finds that footage regardless of the actual hard drive or network share name.

You could, for example, copy an entire shot from the network to your local drive, set up your Loaders and Savers to use the Comp variable, work all your magic locally (i.e., set up your composition), and then copy just the composition back to the server and issue a net-render. All render slaves automatically find the source footage.

Some examples:

Your composition is stored in:

```
X:\Project\Shot0815\Fusion\Shot0815.comp
```

Your source footage sits in:

```
X:\Project\Shot0815\Fusion\Greenscreen\0815Green_0000.dpx
```

The relative path in the Loader node would then be:

```
Comp:\Greenscreen\0815Green_0000.dpx
```

If your source footage is stored in:

```
X:\Project\Shot0815\Footage\Greenscreen\0815Green_0000.dpx
```

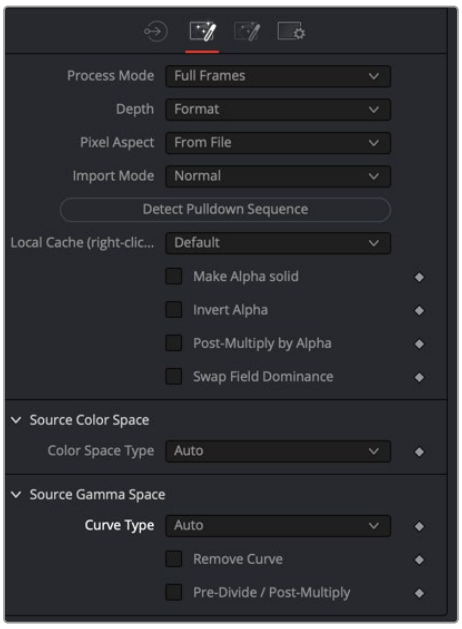
The relative path in the Loader node would then be:

```
Comp:\..\Footage\ Greenscreen\0815Green_0000.dpx
```

Observe how the two dots .. set the directory to go up one folder, much like CD .. in a Command Shell window.

Import Tab

The Import tab includes settings for the frame format and how to deal with fields, pixel aspect, 3:2 pull down/pull up conversion, and removing gamma curve types for achieving a linear workflow.



Loader Import tab

Process Mode

Use this menu to select the Fields Processing mode used by Fusion when loading the image. The Has Fields checkbox control in the Frame Format preferences determines the default option, and the default height as well. Available options include:

- Full frames
- NTSC fields
- PAL/HD fields
- PAL/HD fields (reversed)
- NTSC fields (reversed).

The two reversed options load fields in the opposite order and thus result in the fields being spatially swapped both in time order and in vertical order as well.

Use the Swap Field Dominance checkbox (described below) to swap fields in time only.

Depth

The Depth menu is used to select the color depth used to process footage from this Loader. The default option is Format.

- **Format:** The color depth is determined by the color depth supported in the file format loaded. For example, JPEG files automatically process at 8 bit because the JPEG file format does not store color depths greater than 8. EXR files load at Float. If the color depth of the format is undetermined, the default depth defined in the Frame Format preferences is used. Formats that support multiple color depths are set to the appropriate color depth automatically.
- **Default:** The color depth is determined by the settings in the composition's Frame Format Preferences panel.
- **Int 8 Bit/Int 16 Bit/Float 16/Float 32:** These options set the color depth for processing the image.

Pixel Aspect

This menu is used to determine the image's pixel aspect ratio.

- **From File:** The loader conforms to the image aspect detected in the saved file. There are a few formats that can store aspect information. TIFF, JPEG, and OpenEXR are examples of image formats that may have the pixel aspect embedded in the file's header. When no aspect ratio information is stored in the file, the default frame format method is used.
- **Default:** Any pixel aspect ratio information stored in the header of the image file is ignored. The pixel aspect set in the composition's frame format preferences is used instead.
- **Custom:** Select this option to override the preferences and set a pixel aspect for the clip manually. Selecting this button causes the X/Y Pixel Aspect control to appear.

Custom Pixel Aspect

This control is visible only when Custom is selected from the Pixel Aspect menu. Enter the desired X and Y aspect or right-click on the control to display a menu of common frame formats and their aspects.

Import Mode

This menu provides options for removing pull-up from an image sequence. Pull-up is a reversible method of combining frames used to convert 24 fps footage into 30 fps. It is commonly used to broadcast NTSC versions of films.

- **Normal:** This passes the image without applying pull-up or pull-down 2:3.
- **Pull Up:** This removes existing 3:2 pull-down applied to the image sequence, converting from 30 fps to 24 fps 2:3.
- **Pull Down:** The footage has pull-down applied, converting 24 fps footage to 30 fps by creating five frames out of every four. The process mode of a Loader set to Pull Down should always be Full Frames.

First Frame

This menu appears when the Import Mode is set to either Pull Up or Pull Down. It is used to determine which frame of the 3:2 sequence is used as the first frame of the loaded clip.

Detect Pull-Down Sequence

This button is used to detect and set the pull-up sequence of the footage automatically. It only works if Pull-Up or Pull-Down is first selected from the Import Mode menu. If it succeeds in detecting the order, the First Frame control automatically sets to the correct value.

Make Alpha Solid

When enabled, the original Alpha channel of the clip is cleared and set to solid white (completely opaque).

Invert Alpha

When enabled, the original Alpha channel of the clip is inverted. This may also be used in conjunction with Make Alpha Solid to set the Alpha to pure black (completely transparent).

Post-Multiply by Alpha

Enabling this option causes the color value of each pixel to be multiplied by the Alpha channel for that pixel. This option can be used to convert subtractive (non-premultiplied) images to additive (premultiplied) images.

Swap Field Dominance

When enabled, the field order (dominance) of the image is swapped, so the order in time the fields appear is reversed. Unlike the Process Mode control, this is done without spatially swapping the scanlines in the image

Color Space Type

This menu is used to set the Color Space of the footage to help achieve a linear color space workflow. Unlike the Gamut tool, this doesn't perform any actual color space conversion, but instead adds the source color space data into the metadata, if that metadata doesn't already exist. The metadata can then be used downstream by a Gamut tool with the From Image option, or in a Saver if explicit output spaces are defined there.

- **Auto:** Passes along any metadata that might be in the incoming image.
- **Space:** Allows the user to set the color space based on the recording device used to capture content or software settings used when rendering the content in another application.

Curve Type

This menu is used to determine the gamma curve of the footage. Once the Gamma Curve Type is set, you can choose to remove the curve to help achieve a linear workflow.

- **Auto:** Passes along any metadata that might be in the incoming image.
- **Space:** Allows the user to set the gamma curve based on the recording device used to capture content or software settings used when rendering the content in another application.
- **Log:** Displays the Log/Lin settings, similar to the Cineon Log node.

> For more information on the Log settings, *see Chapter 39, "Film Nodes,"* in the Fusion Reference Manual.

Remove Curve

Depending on the selected Curve Type or on the Gamma Space found in Auto mode, the associated Gamma Curve is removed from, or a log-lin conversion is performed on, the material, effectively converting it to a linear output space.

Format Tab

The Format tab contains file format-specific controls that dynamically change based on the selected Loader and the file it links to. Some formats contain a single control or no controls at all. Others like Camera RAW formats contain RAW-specific debayering controls. A partial format list is provided below for reference.



Format tab

- **OpenEXR:** EXR provides a compact and flexible format to support high dynamic range images (float). The format also supports a variety of extra channels and metadata.The Format tab for OpenEXR files provides a mechanism for mapping any non-RGBA channels to the channels supported natively in Fusion. Using the Format tab, you can enter the name of a channel contained in the OpenEXR file into any of the edit boxes next to the Fusion channel name. A command line utility for dumping the names of the channels can be found at https://www.openexr.com.
- **QuickTime:** QuickTime files can potentially contain multiple tracks. Use the format options to select one of the tracks.

- **Cinema DNG:** CinemaDNG is an open format capable of high-resolution raw image data with a wide dynamic range. It was one of the formats recorded by Blackmagic Design cameras before switching over to the BRAW format.
- **Photoshop PSD Format:** Fusion can load any one of the individual layers stored in the PSD file, or the completed image with all layers. Transformation and adjustment layers are not supported. To load all layers in a PSD file with appropriate blend modes, use File > Import > PSD.

Common Controls

Settings Tab

The Settings tab controls are common to both Loader and Saver nodes, so their descriptions can be found in “The Common Controls” section at the end of this chapter.

MediaIn Node [MI]



The MediaIn node

MediaIn Node Introduction

NOTE: The MediaIn node is only available in DaVinci Resolve.

The MediaIn node is the foundation of every composition you create in DaVinci Resolve’s Fusion page. In most cases, it replaces the Loader node used in Fusion Studio for importing clips. There are four ways to add a MediaIn node to the Node Editor.

- In the Edit or Cut page, position the playhead over a clip in the Timeline, and then click the Fusion page button. The clip from the Edit or Cut page Timeline is represented as a MediaIn node in the Node Editor.
- Drag clips from the Media Pool into the Node Editor, creating a MediaIn node in the Node Editor.
- Drag clips from an OS window directly into the Node Editor, creating a MediaIn node in the Node Editor.
- Choose Fusion > Import> PSD when importing PSD files into the Node Editor. Each PSD layer is imported as a separate MediaIn node.

NOTE: Although a MediaIn tool is located in the I/O section of the Effects Library, it is not used as a method to import clips.

When clips are brought in from the Media Pool, dragged from the OS window, or via the Import PSD menu option, you can use the MediaIn node’s Inspector for trimming, looping, and extending the footage, as well as setting the source’s color and gamma space.

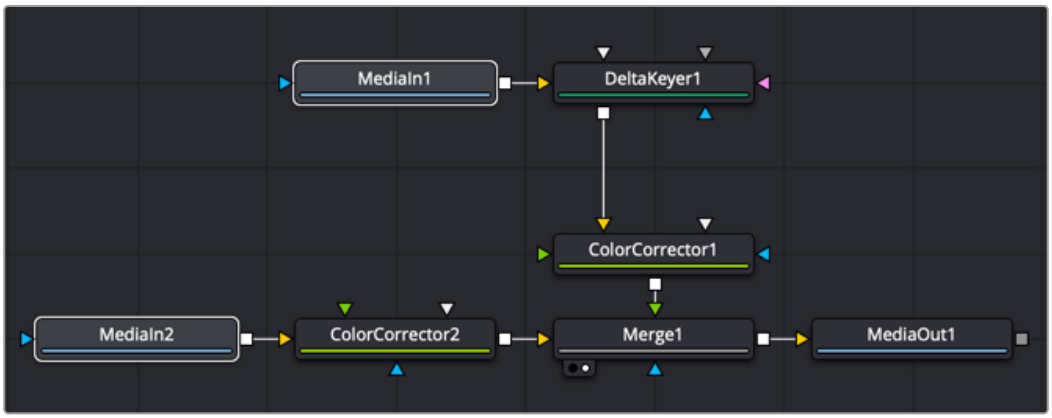
Inputs

The single input on the MediaIn node is for an effect mask to crop the image brought in by the MediaIn.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the source image to appear only within the mask. An effects mask is applied to the tool after the tool is processed.

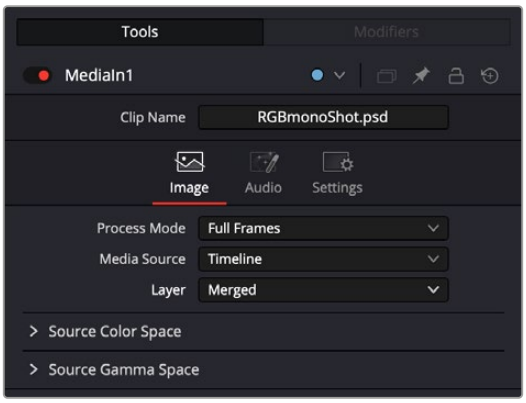
Basic Node Setup

The MediaIn node is typically the starting point for all composites done in the Fusion page of DaVinci Resolve. It contains the clip from the Edit page or Cut page. Any clip brought in from the Media Pool is also added as a MediaIn node.



Two MediaIn nodes: one from the Edit page Timeline and one from the Media Pool

Inspector



MediaIn node Image tab

Image Tab

When brought in from the Media Pool or dragged from an OS window, the MediaIn node's Image tab includes controls for trimming, creating a freeze frame, looping, and reversing the clip. You can also reselect the clip the MediaIn links to on your hard drive. A subset of these controls is available when the MediaIn node is brought in from the Edit or Cut page Timeline.

Global In and Out

Only used when a clip is brought in through the Media Pool or an OS window, the Global In and Out handles are used to specify the start and end of this node within the Fusion effect. Use Global In to specify the frame on which the clip starts and use Global Out to specify the frame on which the clip ends within the comp's global range. The node does not produce an image on frames outside of this range.

If the Global In and Out values are decreased to the point where the range between the In and Out values is smaller than the number of available frames in the clip. Fusion automatically trims the clip by adjusting the Trim range control. If the Global In/Out values are increased to the point where the range between the In and Out values is larger than the number of available frames in the clip, Fusion automatically lengthens the clip by adjusting the Hold First/Last Frame controls. Extended frames are visually represented in the range control by changing the color of the held frames to green in the control.

To slide the clip in time or move it through the project without changing its length, place the mouse pointer in the middle of the range control and drag it to the new location, or enter the value manually in the Global In value box.

- **Clip Name:** Displays the name of that clip.

Process Mode

Use this menu to select the Fields Processing mode used by Fusion when loading the image. The Has Fields checkbox control in the Frame Format preferences determines the default option, and the default height as well. Available options include:

- Full frames
- NTSC fields
- PAL/HD fields
- PAL/HD fields (reversed)
- NTSC fields (reversed).

The two reversed options load fields in the opposite order and thus result in the fields being spatially swapped both in time order and in vertical order as well.

Media Source: This control allows you to choose where media is sourced from; the three available options include: Timeline, Background, and Media Pool.

- **Timeline:** Links the clip from the Edit Timeline (default).
- **Background:** This pulls frames from the lower video tracks in the edit. When toggling the Media Source to Background, the composited result of all the lower tracks is visible.
- **Media Pool:** Links the clip directly from the Media Pool, bypassing the Edit Timeline.

Layer: In the case of Multilayer or Photoshop PSD files, this selects the layer in the file to use. For a Fusion Clip created using multiple video layers on the timeline, this will refer to the index number of each of the layers.

Trim

The Trim range control is used to trim frames from the start or end of a clip. Adjust the Trim In to remove frames from the start and adjust Trim Out to specify the last frame of the clip. The values used here are offsets. A value of 5 in Trim In would use the fifth frame in the sequence as the start, ignoring the first four frames. A value of 95 would stop loading frames after the 95th frame.

Hold First Frame/Hold Last Frame:

The Hold First Frame and Hold Last Frame controls hold the first or last frame of the clip for the specified amount of frames. Held frames are included in a loop if the footage is looped.

Reverse

Select this checkbox to reverse the footage so that the last frame is played first, and the first frame is played last.

Loop

Select this checkbox to loop the footage until the end of the project. Any lengthening of the clip using Hold First/Last Frame or shortening using Trim In/Out is included in the looped clip.

Source Color Space

Lets you choose a color space for the image data output by this MediaIn node.

- **Auto:** uses the Timeline color space, or whichever color space is assigned by Resolve Color Management (RCM) if it's enabled.
- **Space:** Space lets you choose a specific setting from a Color Space drop-down menu, while a visual "horseshoe" graph lets you see a representation of the color space you've selected.

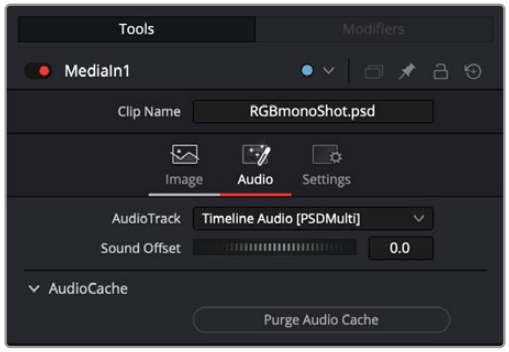
Source Gamma Space

Lets you choose a gamma setting for the image data output by this MediaIn node. Once the gamma curve type is set, you can choose to remove the curve to help achieve a linear workflow.

- **Auto:** Uses the Timeline gamma, or whichever gamma is assigned by Resolve Color Management (RCM) if it's enabled.
- **Space:** Lets you choose a specific setting from a Gamma Space drop-down menu, while a visual graph lets you see a representation of the gamma setting you've selected.
- **Log:** Displays the Log Type drop-down menu where you can choose a specific log encoding profile. A visual graph shows a representation of the log setting chosen from the menu. Additional Lock RGB, Level, Soft Clip, Film Stock Gamma, Conversion Gamma, and Conversion table options are presented to finesse the gamma output.
- **Remove Curve:** The associated gamma curve is removed from, or a log-lin conversion is performed on, the material, effectively converting it to a linear output space.
- **Pre-Divide/Post-Multiply:** Lets you convert "straight" Alpha channels into premultiplied Alpha channels, when necessary.

Audio Tab

The Inspector for the MediaIn node contains an Audio tab, where you can choose to solo the audio from the clip or hear all the audio tracks in the Timeline.



The Audio tab in the MediaIn node is used to select the track for playback, slip the audio timing, and reset the audio cache.

If the audio is out of sync when playing back in Fusion, the Audio tab's Sound Offset wheel allows you to slip the audio in subframe frame increments. The slipped audio is only modified in the Fusion page. All other pages retain the original audio placement.

Audio with Media Pool Clips

Audio from a clip brought in through the Media Pool is muted by default. Hearing the audio from a Media Pool clip is a two step process.

To hear audio from a clip brought in through the Media Pool, do the following:

1. Select the clip in the Node Editor.
2. In the Inspector, click the Audio tab and select the clip name from the Audio Track drop-down menu.

   If more than one MediaIn node exists in the comp, the audio last selected in the Inspector is heard. You can use the Speaker icon in the toolbar to switch between the MediaIn node audio files.
3. Right-click the Speaker icon in the toolbar, then choose the MediaIn for the clip you want to hear.

Purging the Audio Cache

The audio and its settings are cached for faster performance. If you change which audio tracks you want to play back in Fusion, or you use the Sound Offset wheel to slip the audio tracks, you need to purge the audio cache. Also, if you return to the Edit, Cut, or Fairlight page and modify the audio levels, you need to purge the audio cache.

To purge the audio cache after any change to the audio playback:

- Click the Purge Audio Cache button in the Inspector.

The audio will be updated when you next playback the composition.

MediaOut Node [MO]



The MediaOut node

NOTE: The MediaOut node is only available in DaVinci Resolve.

MediaOut Node Introduction

Every composition you create in DaVinci Resolve's Fusion page must include a MediaOut node. The MediaOut node sends the final output back to your Timeline on DaVinci Resolve's Edit or Cut page. In most cases, it replaces the Saver node used in Fusion Studio for exporting clips.

The composition output by the Fusion page's MediaOut node is propagated via the Color page's source inputs, with the sole exception that if you've performed transforms or added plugins to that clip in the Edit or Cut page, then the handoff from the Fusion page to the Color page is as follows:



When using Resolve Color Management or ACES, each MediaOut node converts the output image back to the Timeline color space for handoff to the Color page.

NOTE: Additional MediaOut nodes can be added to the Node Editor from the Effects Library. Additional MediaOut nodes are used to pass mattes to the Color page.

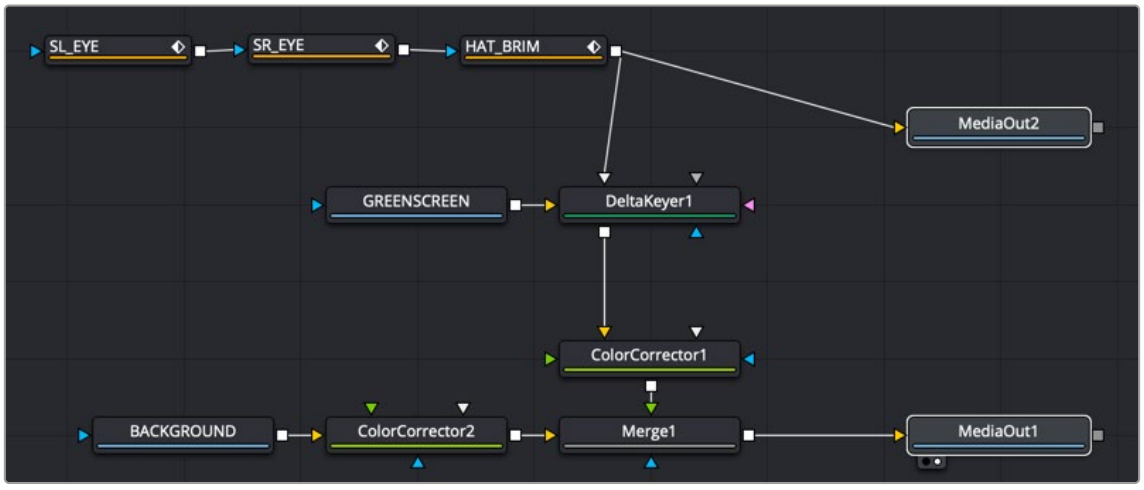
Inputs

The single input on the MediaOut node is where you connect the final composite image you want rendered back into the Edit page.

Input: The orange input is a required input. It accepts any 2D image that you want rendered back into the Edit page.

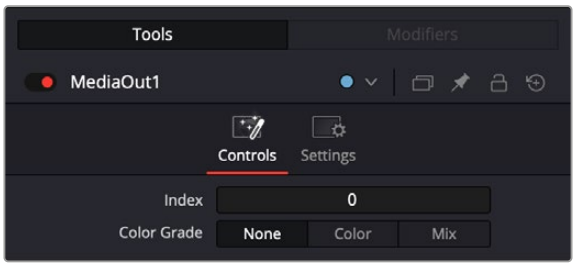
Basic Node Setup

Multiple MediaOut nodes can exist in a single comp. The first MediaOut node always renders the image back to the Edit page Timeline. Additional MediaOut nodes can be used to pass mattes to the Color page.



MediaOut1 node rendering to the Edit page, and MediaOut2 sending mattes to the Color page

Inspector



The MediaOut controls

The MediaOut node has various Color Grade options for viewing in the Fusion page.

- **None:** This is the default view, displaying the output of the Fusion composition only.
- **Color:** This displays the output of the Fusion composition, including Color page adjustments at timeline resolution.
- **Mix:** This displays the output of the Fusion composition including Cut and Edit page effects and Color page adjustments at timeline resolution.

NOTE: These options are previews only, to show your current comp in the context of any other effects which may have been added outside Fusion. The actual output of the MediaOut node itself remains unchanged.



The Color Grade tab set to None

None is displaying the output of the Fusion composition only. Note the display is at the source media/composition resolution.



The Color Grade tab set to Color

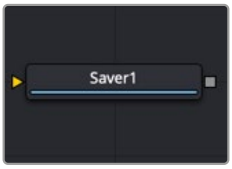
Color is displaying the Fusion composition, including any adjustments applied in the Color page. Note that when the MediaOut view is set to Color, the display will be set to the timeline resolution.



The Color Grade tab set to Mix

Mix is displaying the Fusion composition, including any effects applied in the Cut or Edit page and any adjustments in the Color page. Note that when the MediaOut view is set to Mix, the display will be set to the timeline resolution.

Saver Node [SV]



The Saver node

NOTE: The Saver node in DaVinci Resolve is only used for exporting EXR files.

Saver Node Introduction

The Saver node represents the final composition output from Fusion Studio. It is used to render out movie files or sequential images but can be inserted into a composition at any point to render out intermediate stages of a composition. A composition can contain any number of Saver nodes for rendering different branches of a comp as well as different formats.

The Saver node can also be used to add scratch track audio to your composition, which can be heard during interactive playback.

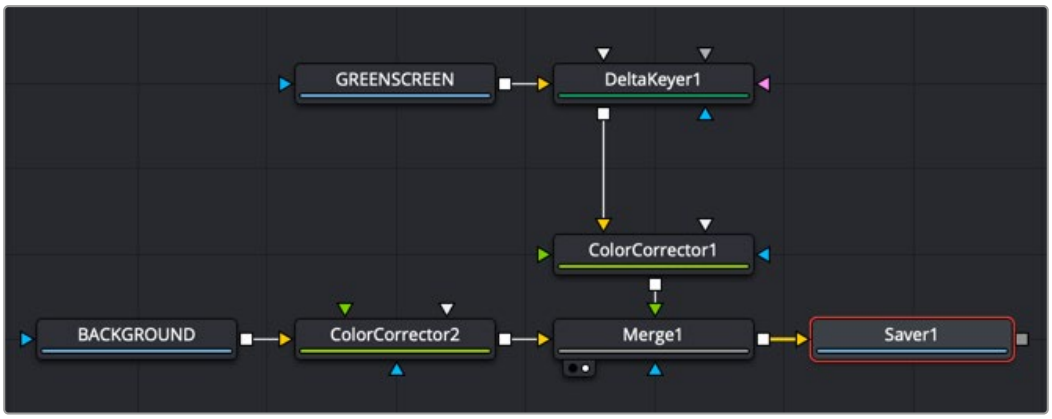
Inputs

The single input on the Saver node is for the final composition you want to render.

Image Input: The orange input is used to connect the resulting image you want rendered.

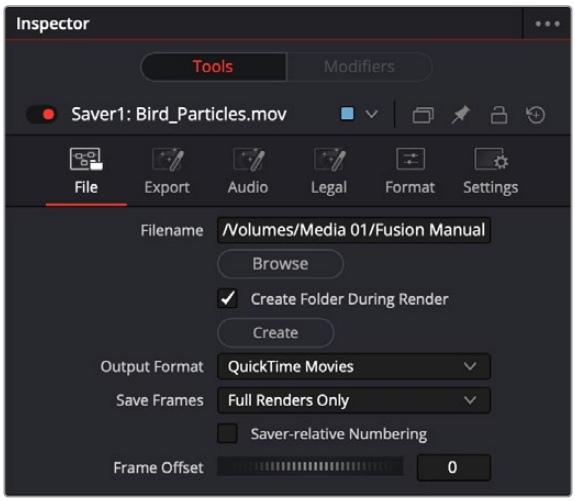
Basic Node Setup

The Saver node is placed at the end of the composition. Multiple Savers can be placed in a comp to render different formats or to render different parts of a composition.



Saver node added to the end of a node tree to render the composition

Inspector



Saver File tab

File Tab

The Saver File tab is used to set the location and output format for the rendered file.

Filename

The Filename dialog is used to select the name and path of the rendered image output. Click on the Browse button to open a file browser and select a location for the output.

Sequence numbering is automatically added to the filename when rendering a sequential image file format. For example, if c\renders\image.exr is entered as the filename and 30 frames of output are rendered, the files are automatically numbered as image0000.tga, image0001.exr, image0003.exr... and so on. Four-digit padding is automatically used for numbers lower than 10000.

You can specify the number of digits to use for padding by explicitly entering the digits into the filename.

For example, image000000.exr would apply 6-digit padding to the numeric sequence, image.001.exr would use 3-digit padding, and image1.exr would use none.

Output Format

This menu is used to select the image format to be saved. Be aware that selecting a new format from this menu does not change the extension used in the filename to match. Modify the filename manually to match the expected extension for that format to avoid a mismatch between name and image format.

Save Frames

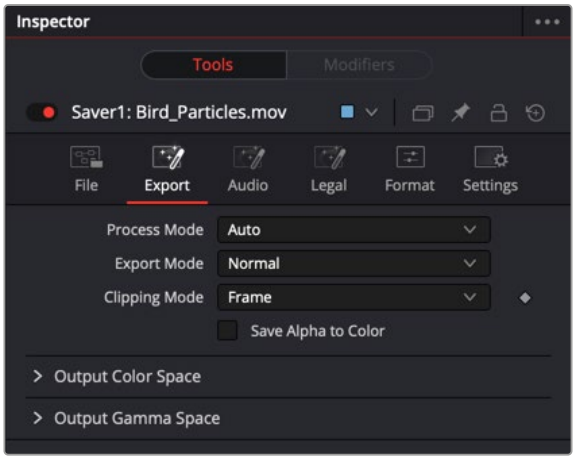
This control selects between two modes of rendering: Full Renders Only or High Quality Interactive.

- **Full Renders Only:** This is the common setting for most situations. Images are saved to disk when a final render is started using the Start Render button in the Time Ruler.
- **High Quality Interactive:** This render mode is designed for real-time rendering when painting and rotoscoping. Fusion saves each frame to disk as it is processed interactively. When used correctly, this feature can eliminate the need to perform a final render after rotoscoping.

NOTE: The High Quality Interactive setting can easily cause confusion when used in conjunction with a node tree that contains spline-animated parameters. If these splines are modified in such a way that frames already saved interactively are changed, the frames already on the disk do not automatically re-render. Either step through each frame again or perform a final render to make certain that the result is correct

Frame Offset

This thumbwheel control can be used to set an explicit start frame for the number sequence applied to the rendered filenames. For example, if Global Start is set to 1 and frames 1-30 are rendered, files are normally numbered 0001-0030. If the Sequence Start Frame is set to 100, the rendered output would be numbered from 100-131.



Saver Export tab

Export Tab

Process Mode

Use this menu to select the Fields Processing mode used by Fusion when saving the images or movie file to disk. The Has Fields checkbox control in the Frame Format preferences determines the default option, and the default height as well. Available options include:

- Full frames
- NTSC fields
- PAL/HD fields
- PAL/HD fields (reversed)
- NTSC fields (reversed).

The two reversed options save fields in the opposite order and thus result in the fields being spatially swapped both in time order and in vertical order as well.

Export Mode

This menu is used to render out the file normally or apply a SMPTE standard 3:2 pulldown to the footage, converting the footage from 24 fps to 30 fps.

Clipping Mode

This menu, sometimes considered source image clipping, defines how the edges of the image should be treated.

- **Frame:** The default Frame setting clips to the parts of the image visible within its visible dimensions. It breaks any infinite-workspace behavior. If the upstream DoD is smaller than the frame, the remaining area in the frame is treated as black/transparent.
- **None:** This setting does not perform any source image clipping at all. This means that any data that would normally be needed outside the upstream DoD is treated as black/transparent. Be aware that this might create humongous images that can consume a considerate amount of disk space. So you should use this option only if really needed.

> For more information about ROI, DoD, and Infinite Workspace, *see Chapter 7, "Using Viewers."* in the Fusion Reference Manual.

Save Alpha to Color

When enabled, this control causes the Alpha channel to be saved into the color channels as a grayscale image. This completely overwrites any existing color information.

Color Space Type

This menu is used to set the Color Space of the output file so you could, for example, color space convert images from linear to Rec709, etc., thereby delivering linear EXRs, 709 Quicktimes, etc. from a single source. The images are not being converted in the Comp, only in the images saved to disk are converted.

- **Auto:** Passes along any metadata that might be in the rendered image.
- **Space:** Allows the user to set the color space based on the output format.

Curve Type

This menu is used to select a Gamma curve of the rendered file. Once the gamma curve type is set, you can choose to apply the curve for output.

- **Auto:** Passes along any metadata that might be in the incoming image.
- **Space:** Allows the user to set the gamma curve based on the selected file format.
- **Log:** Displays the Log/Lin settings, similar to the Cineon Log node.

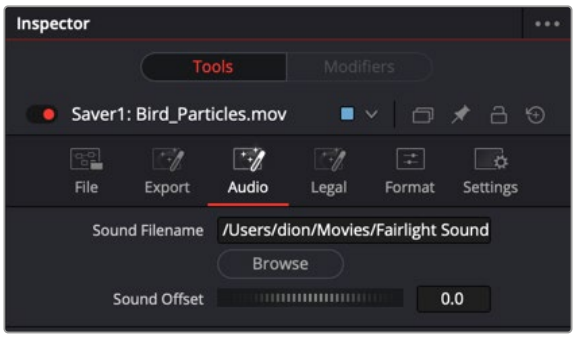
> For more detail on the Log settings, *see Chapter 39, "Film Nodes,"* in the Fusion Reference Manual.

Apply Curve

Depending on the selected Curve Type or on the Gamma Space found in Auto mode, the associated Gamma Curve is applied, effectively converting from a linear working space.

Audio Tab

NOTE: This tab is only available in Fusion Studio's Saver node.



Saver Audio tab

The audio functionality is included in Fusion Studio for scratch track (aligning effects to audio and clip timing) purposes only. Final renders should almost always be performed without audio. The smallest possible audio files should be used, as Fusion loads the entire audio file into memory for efficient display of the waveform in the Timeline. The audio track is included in the saved image if a Quicktime file format is selected. Fusion currently supports playback of WAV audio.

Source Filename

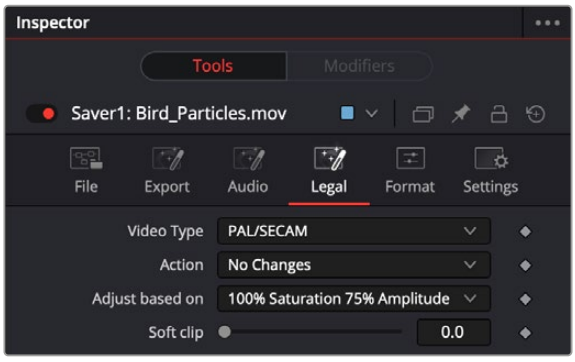
You can enter the file path and name of the audio clip you want to use in the Source Filename field. You can also click the Browse button to open a file browser window and locate the audio scratch track. Select the WAV file of choice, and then in the keyframes panel expand the Saver bar to view the audio waveform. Drag the pointer over the audio wave in the Timeline layout to hear the track.

Sound Offset

Drag the control left or right to slide the Timeline position of the audio clip, relative to other nodes in the Node Editor.

Legal Tab

The Legal tab includes settings for creating "broadcast safe" saturation and video range files for output.



Saver Legal tab

Video Type

Use this menu to select the standard to be used for broadcast legal color correction. NTSC, NHK, or PAL/SECAM can be chosen.

Action

Use this menu to choose how Fusion treats illegal colors in the image.

- **Adjust to Legal:** This causes the images to be saved with legal colors relevant to the Video Type selected.
- **Indicate as Black:** This causes the illegal colors to be displayed as black in the views.
- **Indicate as White:** This causes the illegal colors to be displayed as white in the views.
- **No Changes:** This causes the images to be saved unaffected.

Adjust Based On

This menu is used to choose whether Fusion makes legal the image to 75% or 100% amplitude. Very few broadcast markets permit 100% amplitude, but for the most part this should be left to 75%.

Soft Clip

The Soft Clip control is used to draw values that are out of range back into the image. This is done by smoothing the conversion curve at the top and bottom of the curve, allowing more values to be represented.

Format Tab

The Format tab contains information, options, and settings specific to the image format being saved. The controls for an EXR sequence is entirely different from the ones displayed when a MOV file is saved.



Saver Format tab

EXR is displayed above for reference.

When the Saver node is set to DPX, it's important to understand the reason for the Bypass Conversion > Data is Linear option. When saving log data into a DPX, and not using the Saver's node's own lin-log conversion (that is, Bypass Conversion is checked), the Data Is Linear option should be off. This indicates whether the reason for checking Bypass Conversion is because the data is linear, or whether it's already log.

If Data Is Linear is enabled, then the DPX is marked in its header as containing linear data. In turn, that means that when the DPX is loaded back into Fusion, or into other apps that evaluate the header, those apps think the data is linear and do not perform any log-lin conversion.

Common Controls

Settings Tab

The Settings tab controls are common to both Loader and Saver nodes, so their descriptions can be found in "The Common Controls" section at the end of this chapter.

Clipping Mode

This menu, sometimes considered source image clipping, defines how the edges of the image should be treated.

- **Frame:** The default Frame setting clips to the parts of the image visible within its visible dimensions. It breaks any infinite-workspace behavior. If the upstream DoD is smaller than the frame, the remaining areas in the frame are treated as black/transparent.
- **None:** This setting does not perform any source image clipping at all. This means that any data that would normally be needed outside the upstream DoD is treated as black/transparent. Be aware that this might create humongous images which can consume a considerable amount of disk space. So you should use this option only if really needed.

For more information about ROI, DoD, and Infinite Workspace, *see Chapter 7, "Using Viewers."* in the Fusion Reference Manual.

Save Alpha to Color

When enabled, this control causes the Alpha channel to be saved into the color channels as a grayscale image. This completely overwrites any existing color information.

Color Space Type

This menu is used to set the Color Space of the output file so you could, for example, color space convert images from linear to Rec709, etc., thereby delivering linear EXRs, 709 Quicktimes, etc. from a single source. The images are not being converted in the Comp, only in the images saved to disk are converted.

- **Auto:** Passes along any metadata that might be in the rendered image.
- **Space:** Allows the user to set the color space based on the output format.

Curve Type

This menu is used to select a Gamma curve of the rendered file. Once the gamma curve type is set, you can choose to apply the curve for output.

- **Auto:** Passes along any metadata that might be in the incoming image.
- **Space:** Allows the user to set the gamma curve based on the selected file format.
- **Log:** Displays the Log/Lin settings, similar to the Cineon Log node.

> For more detail on the Log settings, *see Chapter 39, "Film Nodes,"* in the Fusion Reference Manual.

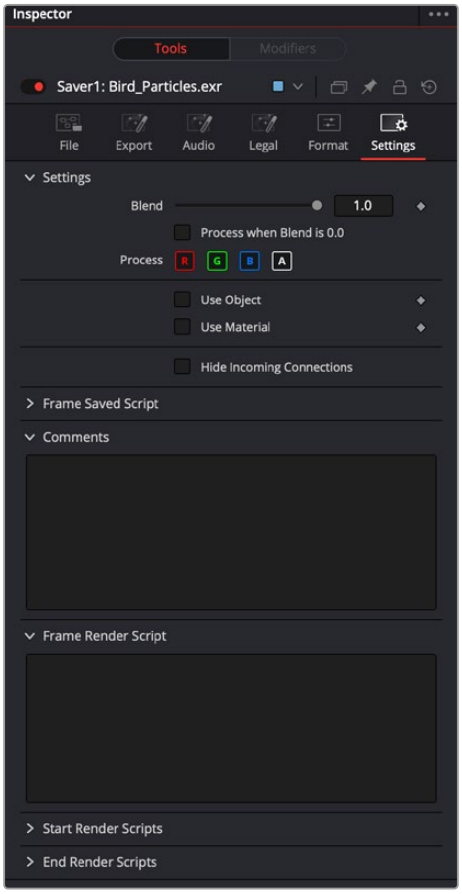
Apply Curve

Depending on the selected Curve Type or on the Gamma Space found in Auto mode, the associated Gamma Curve is applied, effectively converting from a linear working space.

The Common Controls

I/O nodes share a number of identical controls in the Inspector. This section describes controls that are common among I/O nodes.

Inspector



Common Saver settings inspector

Settings Tab

The Settings tab in the Inspector can be found on the Loader, Saver, MediaIn, and MediaOut nodes. The controls are consistent and work the same way for each tool, although some parameters are only available on individual nodes but are covered here.

Blend

The Blend control is used to blend between the tool's original image input and the tool's final modified output image. When the blend value is 0.0, the outgoing image is identical to the incoming image. Normally, this causes the tool to skip processing entirely, copying the input straight to the output.

Process When Blend Is 0.0

The tool is processed even when the input value is zero. This can be useful if processing of this node is scripted to trigger another task, but the value of the node is set to 0.0.

Red/Green/Blue/Alpha Channel Selector

These four buttons are used to limit the effect of the tool to specified color channels. This filter is often applied after the tool has been processed.

For example, if the red button on a Blur tool is deselected, the blur is first applied to the image, and then the red channel from the original input is copied back over the red channel of the result.

There are some exceptions, such as tools for which deselecting these channels causes the tool to skip processing that channel entirely. Tools that do this generally possess a set of identical RGBA buttons on the Controls tab in the tool. In this case, the buttons in the Settings and the Controls tabs are identical.

Apply Mask Inverted

Enabling the Apply Mask Inverted option inverts the complete mask channel for the tool. The mask channel is the combined result of all masks connected to or generated in a node.

Multiply by Mask

Selecting this option causes the RGB values of the masked image to be multiplied by the mask channel's values. This causes all pixels of the image not included in the mask (i.e., set to 0) to become black/transparent.

Use Object/Use Material (Checkboxes)

Some 3D software can render to file formats that support additional channels. Notably, the EXR file format supports Object and Material ID channels, which can be used as a mask for the effect. These checkboxes determine whether the channels are used, if present. The specific Material ID or Object ID affected is chosen using the next set of controls.

Correct Edges

This checkbox appears only when the Use Object or Use Material checkboxes are selected. It toggles the method used to deal with overlapping edges of objects in a multi-object image. When enabled, the Coverage and Background Color channels are used to separate and improve the effect around the edge of the object. If this option disabled (or no Coverage or Background Color channels are available), aliasing may occur on the edge of the mask.

> For more information on Coverage and Background Color channels, *see Chapter 18, "Understanding Image Channels,"* in the Fusion Reference Manual.

Object ID/Material ID (Sliders)

Use these sliders to select which ID is used to create a mask from the object or material channels of an image. Use the Sample button in the same way as the Color Picker: to grab IDs from the image displayed in the view. The image or sequence must have been rendered from a 3D software package with those channels included.

Hide Incoming Connections

Enabling this checkbox can hide connection lines from incoming nodes, making a node tree appear cleaner and easier to read. When enabled, empty fields for each input on a node are displayed in the Inspector. Dragging a connected node from the node tree into the field hides that incoming connection line as long as the node is not selected in the node tree. When the node is selected in the node tree, the line reappears.

Comments

The Comments field is used to add notes to a tool. Click in the empty field and type the text. When a note is added to a tool, a small red square appears in the lower-left corner of the node when the full tile is displayed, or a small text bubble icon appears on the right when nodes are collapsed. To see the note in the Node Editor, hold the mouse pointer over the node to display the tooltip.

Scripts

Three Scripting fields are available on every tool in Fusion from the Settings tab. They each contain edit boxes used to add scripts that process when the tool is rendering. For more details on scripting nodes, please consult the Fusion scripting documentation.

Chapter 46

Layer Nodes

This chapter details the use and modification of multilayer images in your compositions.

The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

Contents

Layer Muxer [LMx] 1198
Layer Regex [LRx] 1199
Layer Remover [LRm] 1202
Swizzler [Swz] 1204
The Common Controls 1210

FUSION

Layer Muxer [LMx]



The Layer Muxer node

Layer Muxer Node Introduction

This tool allows layers from one source to be combined with another source. The tool has two inputs; when multilayer sources are connected, layers from one source are combined with the other.

Inputs

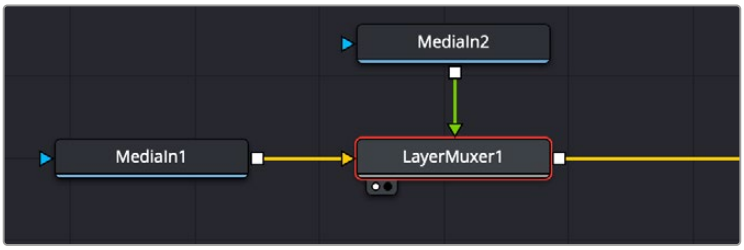
The Layer Muxer tool uses two inputs.

Image 1: This orange input is connected from a node with multilayer data that you want to combine.

Image 2: This green input is connected from a node with multilayer data that you want to combine.

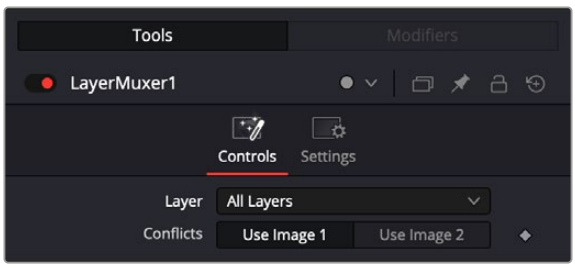
Basic Node Setup

The Layer Muxer node receives two nodes with multilayer data and combines them in various ways based on which layers you want to keep and which you want to ignore.



Layer Muxer combines two multi layer images.

Inspector



Layer Muxer controls

Controls Tab

The Controls tab includes parameters for selecting layers to be used.

- **Layer:** Elect layers from the clip connected to the Image 2 input to combine with the clip connected to Image 1. The pulldown control allows you nominate which layer(s) you'd like to combine. If any option other than All Layers is selected, a small layer icon will appear in the Layer Muxer node to remind you that only certain layers are being passed through.

  Default Layer: If a default layer is used in the Input 2 multilayer media file, it will combine that layer.

  All Layers: Combines all the layers of the Input 2 multilayer image together.

  Custom: Presents a check list off all layers in Input 2 that allows you to turn on and off the multiple layers.

- **Conflicts:** Use this control to switch between either Image 1 or Image 2, the two inputs on this tool; the selected Image will be output as the tool's Default/[Main] layer.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Layer nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Layer Regex [LRx]



The Layer Regex node

Layer Regex Node Introduction

This tool uses complex expressions to filter out layers of multilayer files based on layer naming. Once a filter has been defined, layers that meet the "expression" can be renamed, kept, or removed.

Inputs

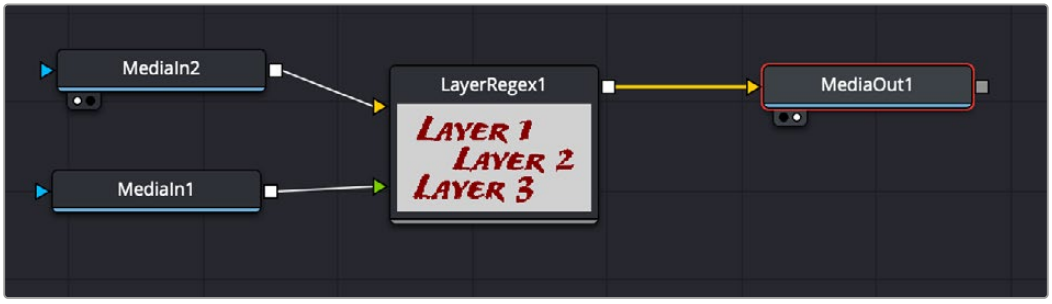
The Layer Regex tool uses two inputs.

Image 1: This orange input is connected from a node with multilayer data that you want to filter.

Image 2: This green input is connected from a node with multilayer data that you want to filter.

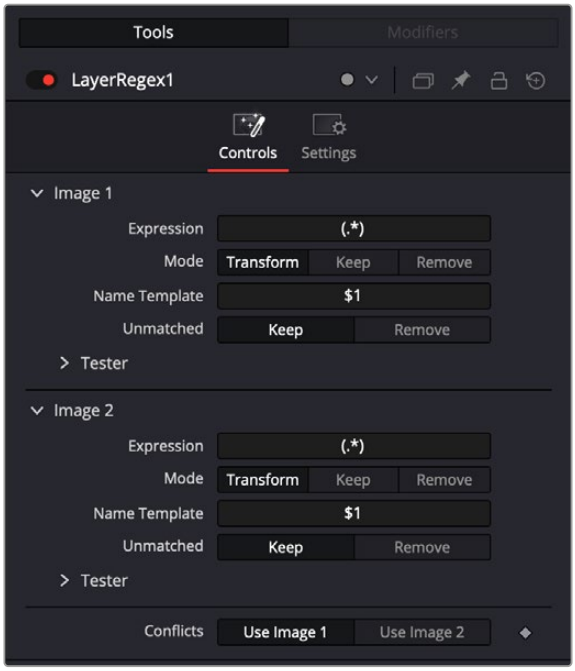
Basic Node Setup

The Layer Regex node receives two nodes with multilayer data and lets you filter them using Regex based on the layer names.



Layer Regex filters layers from two multi layer images.

Inspector



Layer Regex controls

Controls Tab

The Controls tab includes parameters for selecting layers to be used.

- **Expression:** Is the filter used for matching/locating characters.
- **Mode:** Lets you adjust how the expression will be applied.

  Transform: Is the act of renaming/replacing characters.

  Keep: Keeps the layers specified in the Expression rule.

  Remove: Removes the layers specified in the Expression rule.

- **Name Template:** Is the rule applied to the “Expression," allowing you to replace characters or rename layers.
- **Unmatched:** Determines how to deal with layers that don’t match the expression.

  Keep: Any layers that don't match the Expression are kept; all the layers remain.

  Remove: Any layers that don't match the Expression are removed. This isolates the layers that are impacted by the expression.

- **Tester:** A simple tool to view the results of your expression as you change the parameters.
- **Conflicts:** Use this control to switch between either Image 1 or Image 2, the two inputs on this tool; the selected Image will be output as the tool's Default/[Main] layer.

Examples

Add a prefix to existing layer names

Expression: `(.*)`

Name Template: `Right.$1`

This will add Right. at the beginning of each of the layer names, giving you `Right.Layer1` , `Right.Layer2`

NOTE: To add a suffix, change the order in the Name Template, `$1.Right.` This will give you `Layer1.Right` , `Layer2.Right`

Find specific characters/words in the existing layers names, then add a prefix

Expression: `(Right.*)`

Name Template `Stereo-$1`

This will locate the layers which contain the word Right, then add the word `Stereo-` at the beginning of each of the layer names, giving you `Stereo-Right.Layer1`

Find and replace

Expression: `right(.*)`

Name Template: `Left$1`

This will locate any layer names containing right and will then replace that layer name with Left. So, layers labeled `right.1` , `right.2` , `right.3` will be renamed to `Left.1` , `Left.2` , `Left.3` .

Keep in mind, when using this method, any layer name that contains the letters right (e.g., copyright) will be impacted by the expression and will be renamed. To specify a string of characters, use the following expression `^right(.*)`

NOTE: The ^ means match the start of the text. That will make it filter right but not copyright because it has a copy at the start.

Find either word, then remove (or keep)

Expression: `(copy|right)`

Mode: `Remove`

(copylright) locates **copy** and **right** , when the mode is set to Remove, layers containing copy or right, are then removed from the layer list. This expression can be useful for locating variations of a word, for example **(copy|Copy|COPY)**

NOTE: Setting the mode to **Keep** will retain layers that meet the Expression. Using the above example, if the Mode is set to Keep, layers containing copy or right will remain and all other layers are removed.

Find characters, ignoring case

Expression: **(?i)(right)**

This will capture all the characters in the word right, even if they are upper case. This means any capitalized variation of the word right: **Right**, **RIGHT** or **riGth** will be captured.

RegEx Cheat Sheet

For more information on using RegEx https://www.debuggex.com/cheatsheet/regex/pcre

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Layer nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Layer Remover [LRm]



The Layer Remover node

Layer Remover Node Introduction

This tool provides a simple way to disable/remove layers from a connected source.

Inputs

The Layer Remover tool uses one input.

Image 1: This orange input is connected from a node with multilayer data that you want to remove a layer from.

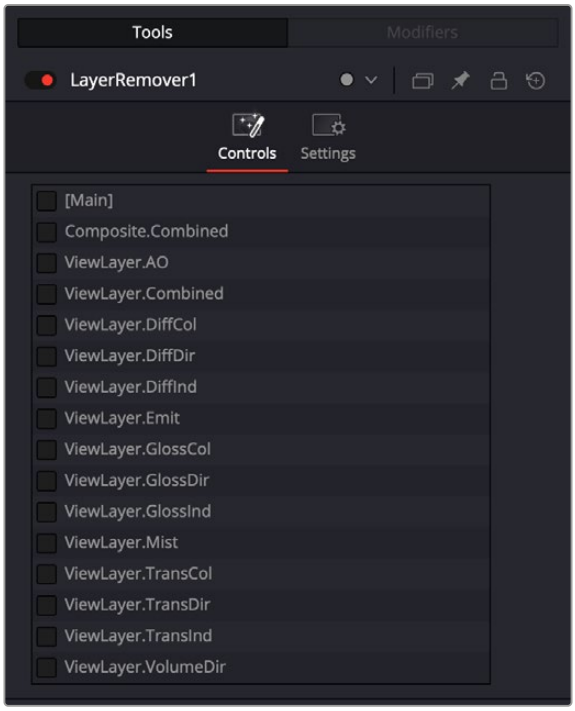
Basic Node Setup

The Layer Remover node receives a node with Multilayer data and lets you strip out specific layers before passing on the image.



Layer Remover lets you disable specific layers of a multilayer file.

Inspector



Layer Remover controls

Controls Tab

The Controls tab presents a simple checklist of all the layers in the multilayer clip. Checking the box will disable/remove layer(s) from the layer list.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Layer nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Swizzler [Swz]



The Swizzler node

Swizzler Node Introduction

This tool enables you to combine one or more image layers and/or channels into a custom-made layer or layers. Similar to Channel Booleans, Swizzler can assign channels from a source and apply them to a destination channel, however, the biggest point of difference is that Swizzler utilizes layers, either targeting layers as a source or creating new layers.

To put it simply, this tool has the functionality to create new layers, push source layers or channels to a layer, or remove channels from a source.

Inputs

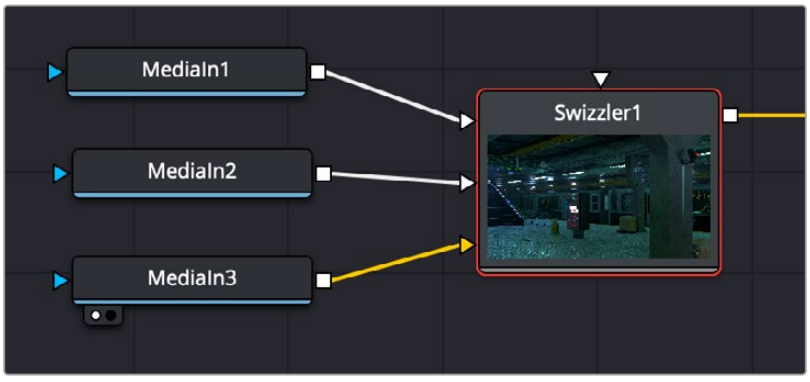
The Swizzler tool uses multiple inputs.

Input 1: This orange input is connected from a node with image data that you want to combine into a new layer.

Input X: The white inputs are connected from additional nodes with image data that you want to combine into a new layer.

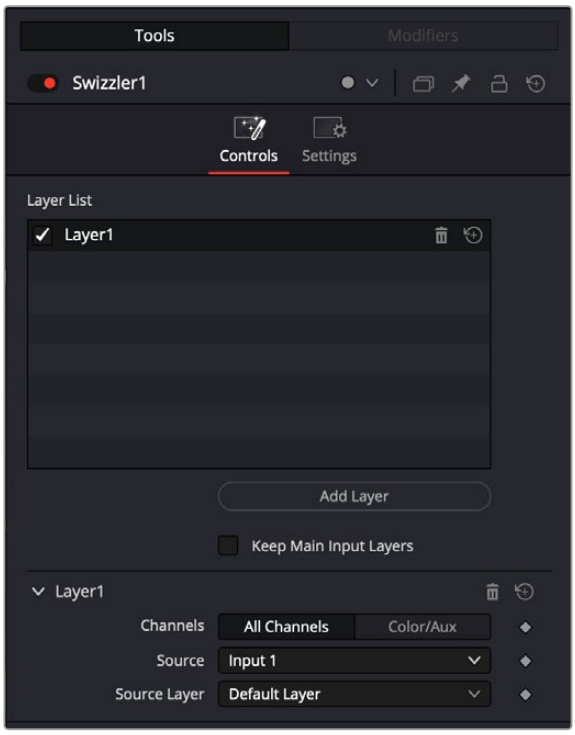
Basic Node Setup

The Swizzler node receives multiple nodes with image data and combines them into new layers.



Swizzler combines multiple images into new layers.

Inspector



Swizzler controls

Controls Tab

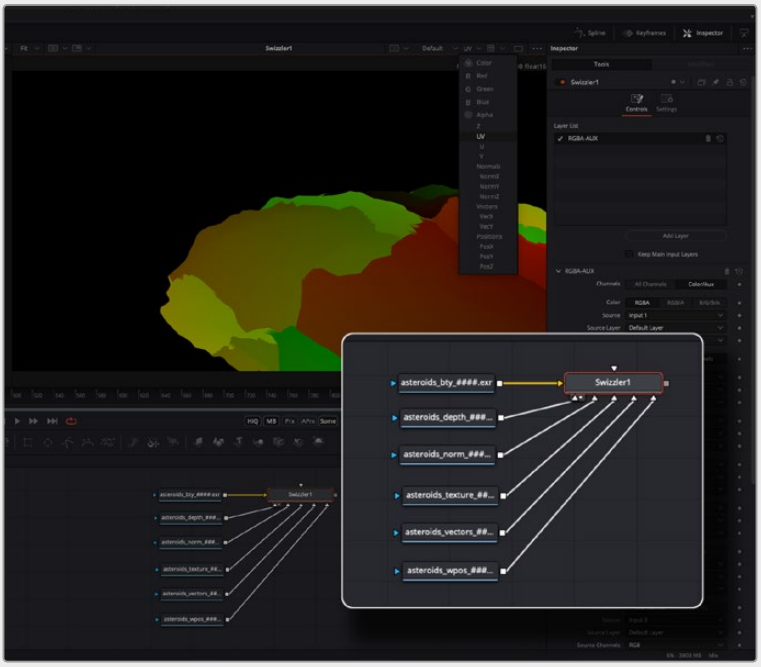
The Controls tab includes parameters for selecting layers to be used.

- **Layer List:** The layer list allows you to view and select custom layers. The button below allows you to add a new layer to this list, introducing a new custom layer. For any layer selected in this list, the options below allow you to choose which channels to include in the layer. Additionally, you can rename layers directly in this list.
- **Add Layer:** This creates a new layer, adding it to the layer list.
- **Keep Main Input Layers:** Checking this box preserves existing layers from the source clip connected to Input1. Having this enabled will pass the layers from Input1 to the Swizzler output, giving you the control to add more layers to an existing multilayer source.
- **Channels:** This control allows you to select which channels will be added to the selected layer in the layer list.
- **All Channels:** Transfers all channels into the selected layer.
- **Color/Aux:** Reveals an additional "Color" and "Aux" options, separating RGBA from AUX channels. This separation allows you to nominate a source for the RGBA and a source for the AUX channels, ultimately combining RGBA from one input and AUX from another input. The other two options, “RGB/A“ and “R/G/B/A,” provide even further granular controls for channels, expanding the source options for individual channels.
- **Source:** Select an input from this drop-down to source layers or channels for your chosen layer. Your chosen layer will be the one selected in the layer list above.
- **Source Layer:** When using a multi-layer file connected to one of the inputs, you can choose a specific layer using this control. All the available layers from the Source Input control will be displayed.
- **Source Channels (Color/Aux mode):** Select which channels from the chosen layer you would like to use for the Aux channel. If the chosen layer doesn't have Aux channels and has been rendered with only RGB, then select the color channels (i.e., RG for UV Texture or RGB for XYZ Normal).

Examples

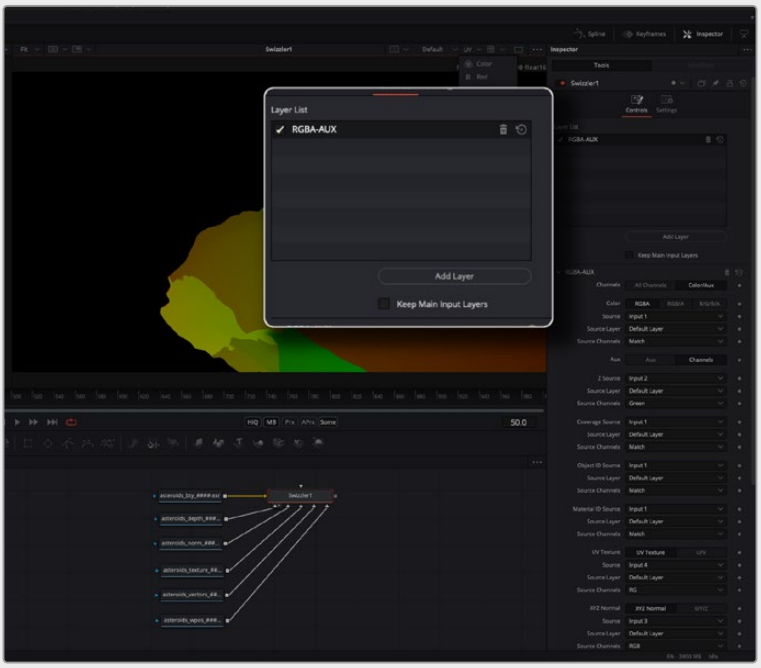
Assigning Aux channels

In the following example we have multiple RGB sources. Each one is a pre-rendered pass which we want to assign to an Aux channel.



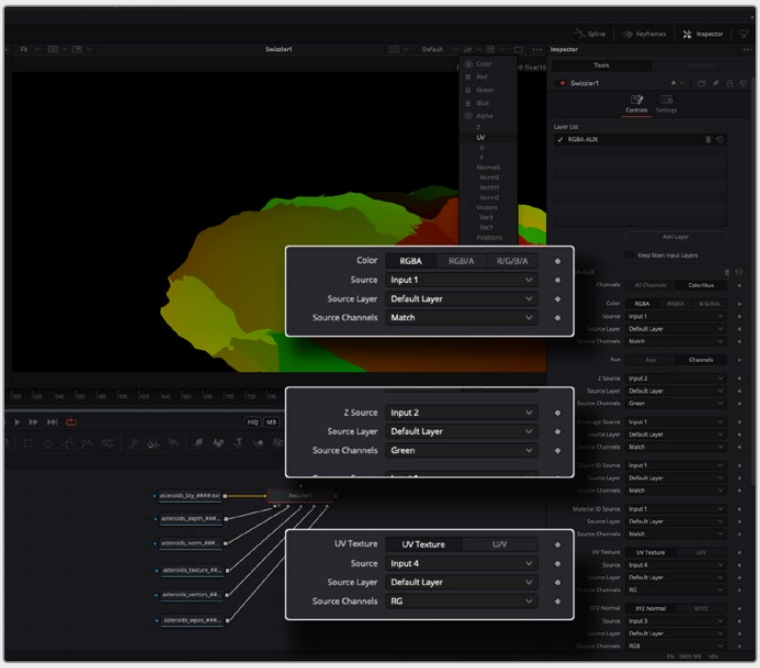
Connecting multiple RGB sources to the Swizzler

Each RGB pass is connected to a new input on the Swizzler.



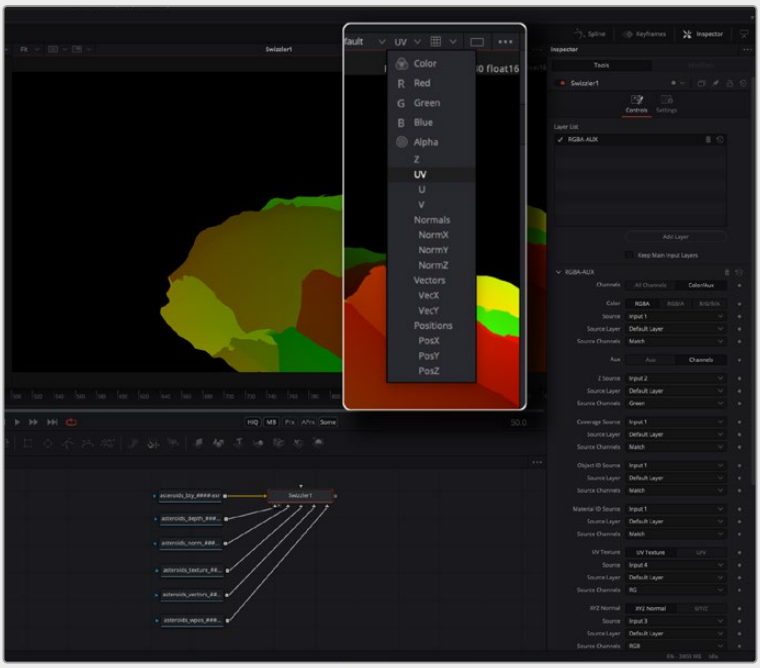
Creating a new layer in the Swizzler

A new layer has been added to the layer list, as there is only one layer, this will become the new ‘default layer’



Assigning sources in the channel controls for the new layer

In the channel controls for the new layer, assign the appropriate source for each of the Aux channels.

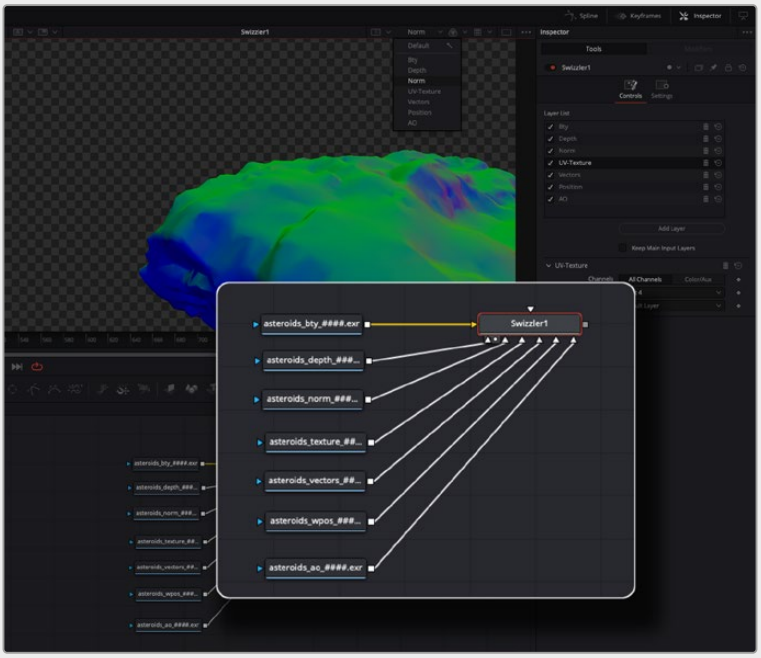


Previewing the new channels in the viewer

Once each of the Aux channels has been assigned an input, the channels can be previewed in the viewer.

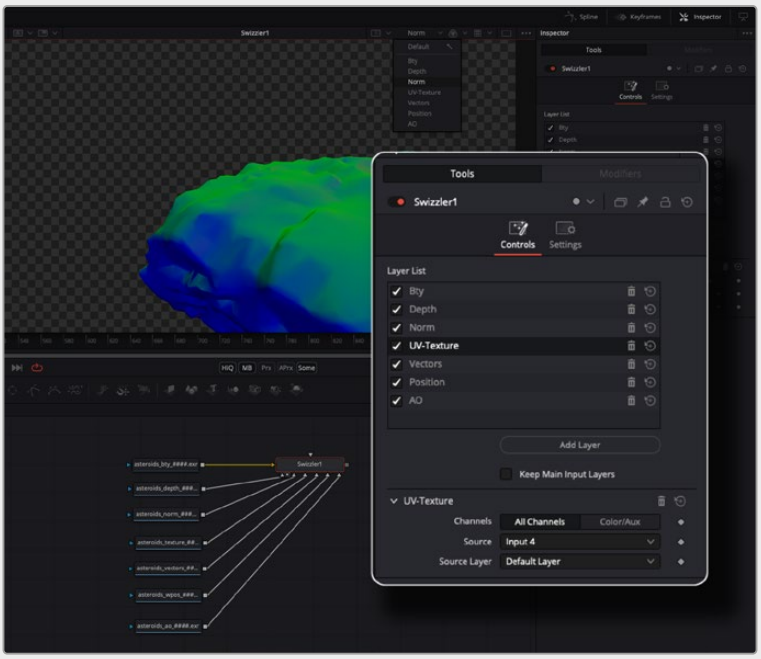
Creating a multilayer clip

In this example, we have multiple RGB sources. Each one is a pre-rendered pass that we want to assign to a new layer, creating a multilayer clip.



Connecting multiple RGB sources to the Swizzler

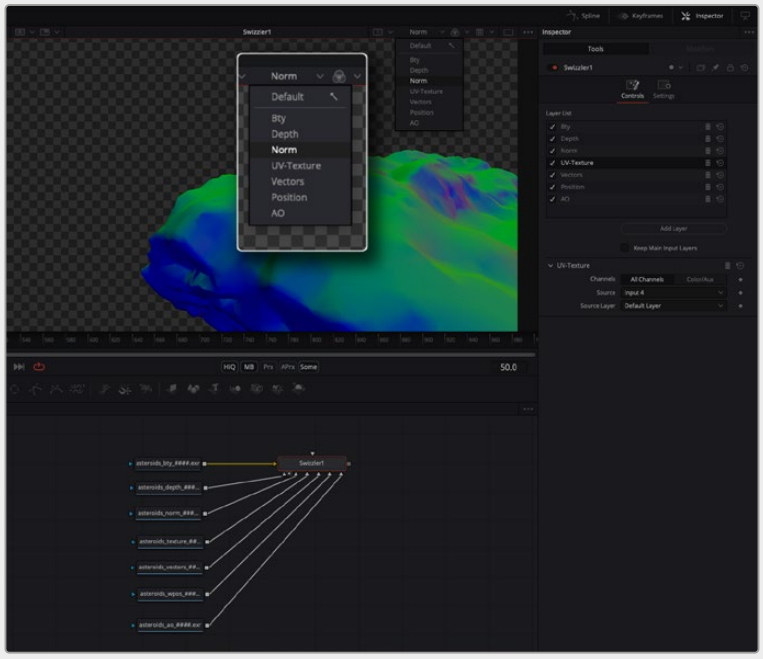
Each RGB pass is connected to a new input on the Swizzler.



Creating new layers in the Swizzler

In the layer list, add a new layer for every pass and give it a custom name. Select a layer in the layer list, and in the controls assign the appropriate source input for that layer.

For example, the UV pass has been connected to Input 4; selecting the custom layer labeled “UV-Texture” and choosing Input 4 from the Input list will push all the channels from Input 4 to the “UV-Texture” layer.



Selecting the new layers in the viewer

Once each of the inputs has been assigned a custom layer, these can be previewed in the viewer and used in the comp.

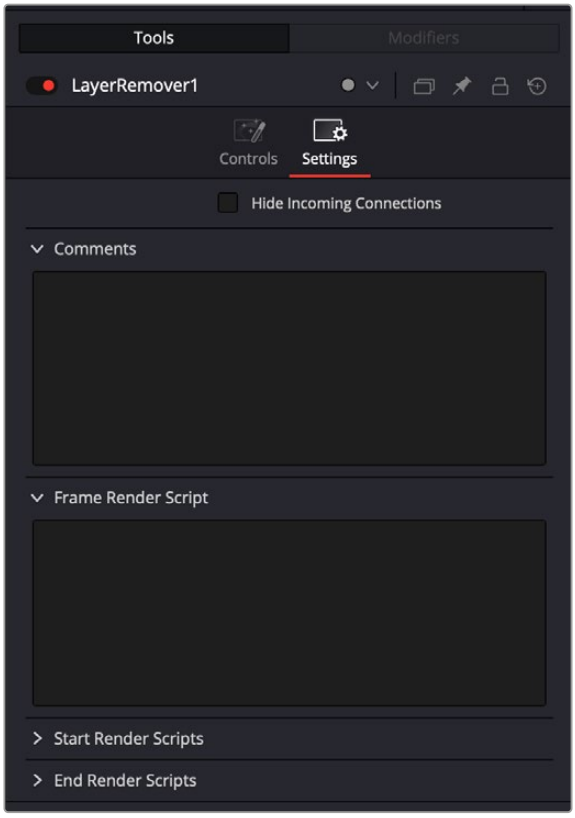
Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Layer nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

The Common Controls

Settings Tab



The Settings tab

The Common Settings tab can be found on most tools in Fusion. The following controls are specific settings for Deep Image nodes.

Hide Incoming Connections

Enabling this checkbox can hide connection lines from incoming nodes, making a node tree appear cleaner and easier to read. When enabled, empty fields for each input on a node are displayed in the Inspector. Dragging a connected node from the node tree into the field hides that incoming connection line as long as the node is not selected in the node tree. When the node is selected in the node tree, the line reappears.

Comments Tab

The Comments tab contains a single text control that is used to add comments and notes to the tool. When a note is added to a tool, a small red dot icon appears next to the setting's tab icon and a text bubble appears on the node. To see the note in the Node Editor, hold the mouse pointer over the node for a moment. The contents of the Comments tab can be animated over time, if required.

Scripting Tabs

The Scripting tab is present on every tool in Fusion. It contains several edit boxes used to add scripts that process when the tool is rendering. For more details on the contents of this tab, please consult the scripting documentation.

Chapter 47

LUT Nodes

This chapter details the LUT nodes available in Fusion.

The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn/MediaOut nodes in DaVinci Resolve are interchangeable with Loader/Saver nodes in Fusion Studio, unless otherwise noted.

Contents

File LUT [FLU] 1212
LUT Cube Analyzer [LCA] 1214
LUT Cube Apply [LCP] 1215
LUT Cube Creator [LCC] 1216
The Common Controls 1218

File LUT [FLU]



The File LUT node

File LUT Node Introduction

The File LUT node applies a Lookup table (LUT) to the image: either a simple 1D LUT or a supported 3D LUT. Unlike the Color Curves node, it does not use a spline-based LUT. Instead, it loads the LUT from a file stored on your computer or server.

This approach has two advantages. The first is that the only part of the LUT stored in the composition is the path to the file. Since LUT files can be large, this can dramatically reduce the file size of a composition when several LUTs are present. The second advantage is that it becomes possible to adjust all File LUT nodes using the same file at the same time, just by changing the contents of the LUT. This can be useful when the same LUT-based color correction is applied in many different compositions.

Inputs

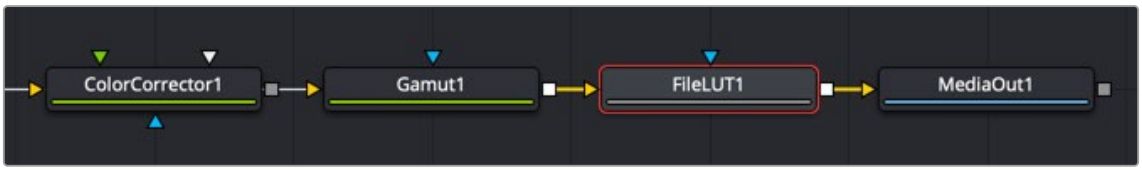
The File LUT node includes two inputs: one for the main image and the other for an effect mask to limit the area where the LUT is applied.

Input: This orange input is the only required connection. It accepts a 2D image output that gets the LUT applied.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the applied LUT to only those pixels within the mask. An effects mask is applied to the tool after the tool is processed.

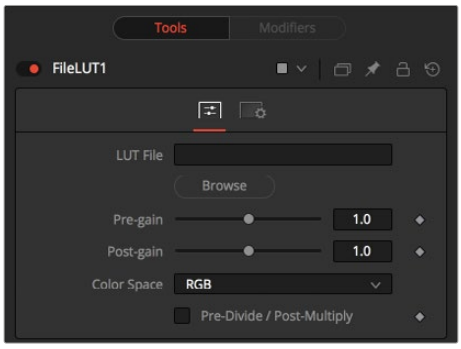
Basic Node Setup

The File LUT node can be placed after a MediaIn node in DaVinci Resolve or a Loader node in Fusion Studio. Sometimes this setup is used to convert the camera-original image to linear color space for compositing. Other times, as in the example below, the File LUT is applied at the end as a grading LUT to apply a look from the colorist.



A File LUT node applied at the end of a node tree as a colorist's look

Inspector



File LUT controls

Controls Tab

The Controls tab includes options for loading a LUT and making adjustments to the gain, color space, and Alpha channel, if one exists.

LUT File

This field is used to enter the path to the LUT file. Clicking the Browse button opens a file browser window to locate the LUT file instead of entering it manually into the LUT File field. Currently, this node supports LUTs exported from Fusion in .LUT and .ALUT formats, DaVinci Resolve's .CUBE format, and several 3D LUT formats. The node fails with an error message on the Console if it is unable to find or load the specified file.

Pre-Gain:

This slider is a gain adjustment before the LUT being applied. This can be useful for pulling in highlights before the LUT clips them.

Post-Gain

This slider is a gain adjustment after the LUT is applied.

Color Space

This menu is used to change the color space the LUT is applied in. The default is to apply the curves described in the LUT to the RGB color space, but options for YUV, HLS, HSV, and others are also available.

Pre-Divide/Post-Multiply

Selecting the Pre-Divide/Post-Multiply checkbox causes the image pixel values to be divided by the Alpha values before applying the LUT, and then re-multiplied by the Alpha value after the correction.

This helps to prevent the creation of illegally additive images, particularly around the edges of a blue/green key or when working with 3D-rendered objects.

Settings Tab

The Settings tab in the Inspector is also duplicated in other LUT nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

LUT Cube Analyzer [LCA]



The LUT Cube Analyzer node

LUT Cube Analyzer Node Introduction

The LUT Cube Analyzer takes an image originated by the LUT Cube Creator as an input and allows the user to create a 3D LUT file in ALUT3, ITX, or 3DL format.

Feeding the original LUT Cube Creator image into the node results in an unaltered, or 1:1, LUT file, and nothing is displayed in the viewer.

You can, however, modify, grade, and color correct the original cube image with as many nodes as you like and feed the result into the LUT Cube Analyzer. This creates a LUT that exactly resembles your color pipeline.

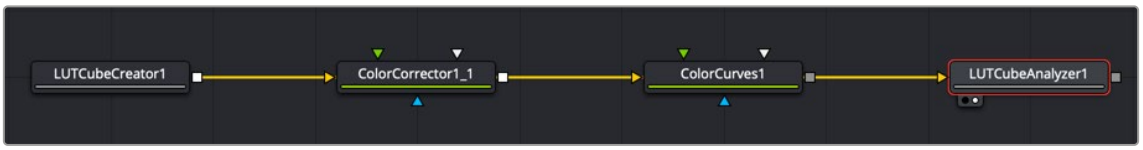
Inputs

The LUT Cube Analyzer includes a single orange input.

Input: The orange input is used to take the output of any node modifying an image that originated with the LUT Cube Creator.

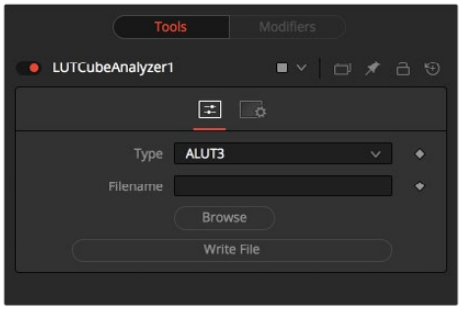
Basic Node Setup

The example below shows a node tree starting with a LUT Cube Creator node and going through two color adjustments. The adjusted image is then connected to a LUT Cube Analyzer, which generates the LUT file.



Generating a LUT starts with the LUT Cube Creator and ends with a LUT Cube Analyzer.

Inspector



LUT Cube Analyzer controls

Controls Tab

The Controls tab for the LUT Cube Analyzer node is used to select the desired LUT output format, specify a filename, and write the 3D LUT to disk.

Type

Select the desired output format of the 3D LUT.

Filename

Enter the path where you want the file saved and enter the name of the LUT file. Alternatively, you can click the Browse button to open a file browser to select the location and filename.

Write File

Press this button to generate the 3D LUT file based on the settings above.

Settings Tab

The Settings tab in the Inspector is also duplicated in other LUT nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

LUT Cube Apply [LCP]



The LUT Cube Apply node

LUT Cube Apply Node Introduction

The LUT Cube Apply node takes an image created by the LUT Cube Creator as the foreground input and applies that LUT to the image connected to the background input.

Feeding the original image into the node would result in an unaltered, or 1:1, output.

You can, however, modify, grade, and color correct the original cube image with as many nodes as you like and feed the result into the LUT Cube Apply. Or, take a LUT image that has been graded beforehand to apply the LUT without having to write an actual 3D LUT using the LUT Cube Analyzer.

Inputs

The LUT Cube Apply has three inputs: a green input where the output of the LUT Cube Creator is connected, an orange input for the image to have the LUT applied, and a blue effect mask input

Input: This orange input accepts a 2D image that gets the LUT applied.

Reference Image: The green input is used to connect the output of the LUT Cube Creator or a node that is modifying the image originating in the LUT Cube Creator.

Effect Mask: The optional effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the LUT Cube Apply to only those pixels within the mask. An effects mask is applied to the tool after the tool is processed.

Basic Node Setup

The example below shows a node tree starting with a LUT Cube Creator node and going through two color adjustments. The adjusted image is then connected to the green Reference input of the LUT Cube Apply. The image you want to apply the LUT to is connected to the orange Input.



The LUT generated by the LUT Cube Creator is applied to an image using the LUT Cube Apply node.

Inspector

There are no controls for the LUT Cube Apply node. The LUT connected to the green foreground input is applied to the image connected to the orange background input without having to write an actual 3D LUT using the LUT Cube Analyzer.

Settings Tab

The Settings tab in the Inspector is also duplicated in other LUT nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

LUT Cube Creator [LCC]



The LUT Cube Creator node

LUT Cube Creator Node Introduction

The LUT Cube Creator generates an image for further use with the LUT Cube Analyzer or LUT Cube Apply nodes. The output can be graded, color corrected, or modified with any tool inside and outside of Fusion. If working outside Fusion, make sure to keep the image in 32-bit floating point to preserve color accuracy.

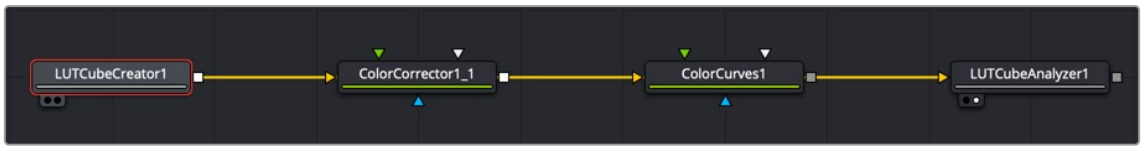
Feeding the original LUT Cube Creator image into the LUT Cube Analyzer node results in an unaltered, or 1:1, LUT file, and nothing is displayed in the viewer.

Inputs

There are no inputs on the LUT Cube Creator. The purpose of the node is to generate an image that can be used to create a LUT.

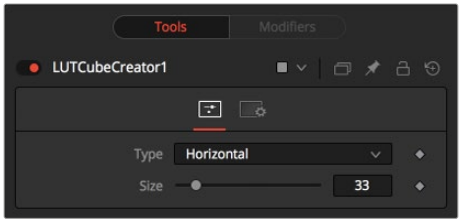
Basic Node Setup

The example below shows a node tree starting with a LUT Cube Creator node and going through two color adjustments. The adjusted image is then connected into a LUT Cube Analyzer that generates the LUT file.



Generating a LUT starts with the LUT Cube Creator and ends with a LUT Cube Analyzer.

Inspector



LUT Cube Creator controls

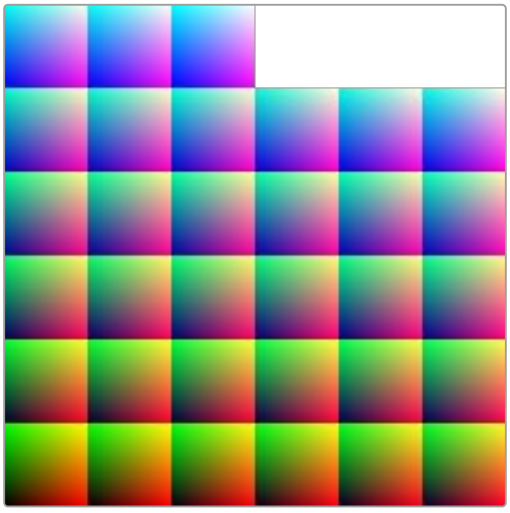
Controls Tab

The Controls tab creates a test pattern of sorts used to create a 3D LUT. The controls here determine the complexity of the pattern used to create a LUT using the LUT Cube Analyzer.

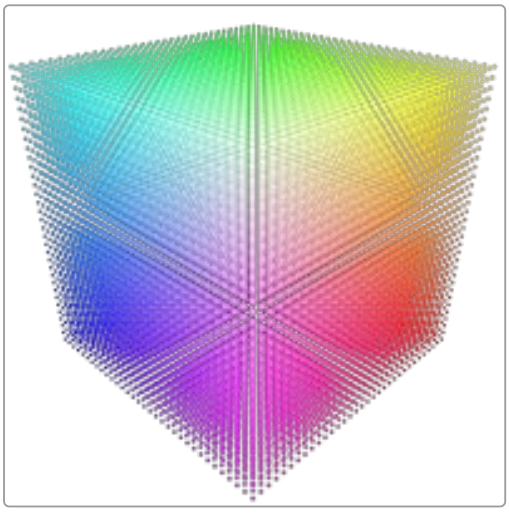
Type:

The Type menu is used to create a pattern of color cubes.

- **Horizontal:** Creates a long, horizontal strip representing a color cube.
- **Vertical:** Creates a long, vertical strip representing a color cube.
- **Rect:** Creates a rectangular image, as depicted below, representing a color cube.



A Cube image created with the Rect type



The resulting color cube

Size

Determines the resolution of the color cube.

Typical Size settings for color cubes are 33 (33 x 33 x 33) or 65 (65 x 65 x 65). These numbers are the samples on each side of the cube. A 33 x 33 x 33 cube has around 35,937 color samples.

> **NOTE:** Higher resolutions yield more accurate results but are also more memory and computationally expensive.

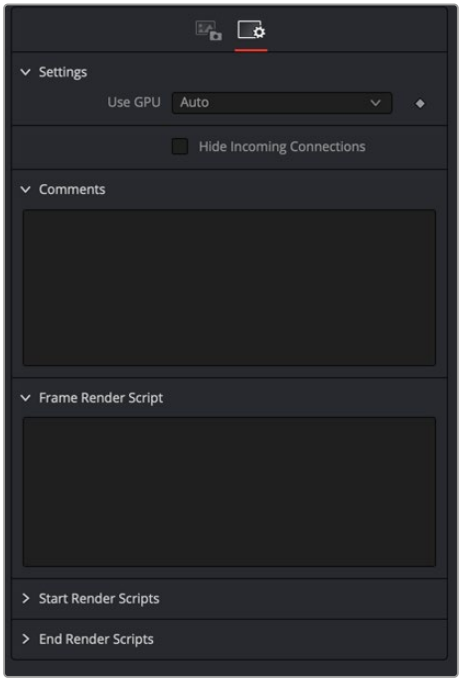
Settings Tab

The Settings tab in the Inspector is also duplicated in other LUT nodes. These common controls are described in the following “The Common Controls” section.

The Common Controls

LUT nodes share a number of identical controls in the Inspector. This section describes controls that are common among LUT nodes.

Inspector



Common LUT Settings Inspector

Settings Tab

The Settings tab in the Inspector can be found on every tool in the LUT category. The controls are consistent and work the same way for each tool, although some tools do include one or two individual options, which are also covered here.

Use GPU

The Use GPU menu has three settings. Setting the menu to Disable turns off hardware-accelerated rendering using the graphics card in your computer. Enabled uses the hardware. Auto uses a capable GPU if one is available and falls back to software rendering when a capable GPU is not available.

Hide Incoming Connections

Enabling this checkbox can hide connection lines from incoming nodes, making a node tree appear cleaner and easier to read. When enabled, empty fields for each input on a node are displayed in the Inspector. Dragging a connected node from the node tree into the field hides that incoming connection line as long as the node is not selected in the node tree. When the node is selected in the node tree, the line reappears.

Comments

The Comments field is used to add notes to a tool. Click in the empty field and type the text. When a note is added to a tool, a small red square appears in the lower-left corner of the node when the full tile is displayed, or a small text bubble icon appears on the right when nodes are collapsed. To see the note in the Node Editor, hold the mouse pointer over the node to display the tooltip.

Scripts

Three Scripting fields are available on every tool in Fusion from the Settings tab. They each contain edit boxes used to add scripts that process when the tool is rendering. For more details on scripting nodes, please consult the Fusion scripting documentation.

Chapter 48

Mask Nodes

This chapter details the Mask nodes available in Fusion.

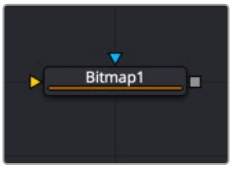
The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Bitmap Mask [Bmp] 1221
B-Spline Mask [BSp] 1226
Ellipse Mask [Elp] 1230
Mask Paint [PnM] 1233
MultiPoly [MPly] 1236
Polygon Mask [Ply] 1238
Ranges Mask [RNG] 1243
Rectangle Mask [Rec] 1248
Triangle Mask [Tri] 1251
Wand Mask [Wnd] 1254
The Common Controls 1258

Bitmap Mask [Bmp]



The Bitmap node

Bitmap Mask Node Introduction

The Bitmap mask node allows images from the node tree to act as masks for nodes and effects. Bitmap masks can be based on values from any of the color, Alpha, hue, saturation, luminance, and auxiliary coverage channels of the image. Nodes can also be masked based on the Object ID or Material ID of a 3D-rendered image (provided those channels were included when the file was rendered).

The Bitmap mask node is not required for effect masks. For effects masks, the Common Settings tab for the masked node displays controls to select which channel of the mask image is used to create the mask.

However, Bitmap mask nodes may still be required to connect to other mask inputs on some nodes, such as Garbage Mattes and Pre-Masks. Also, using a Bitmap mask node between the mask source and the target node provides additional options that would not be available when connecting directly, such as combining masks, blurring the mask, or clipping its threshold.

Inputs

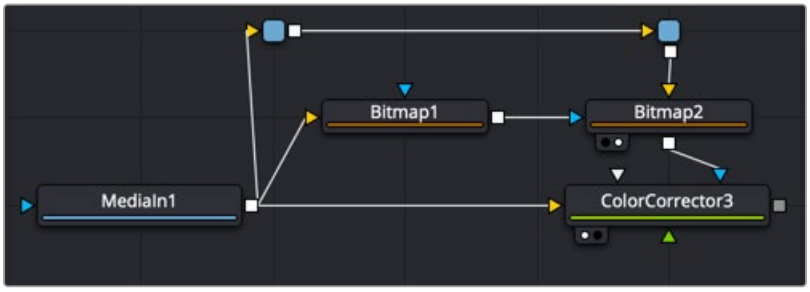
The Bitmap mask node includes two inputs in the Node Editor.

Input: The orange input accepts a 2D image from which the mask will be created.

Effect Mask: The optional blue input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input combines the masks. How masks are combined is handled in the Paint mode menu in the Inspector.

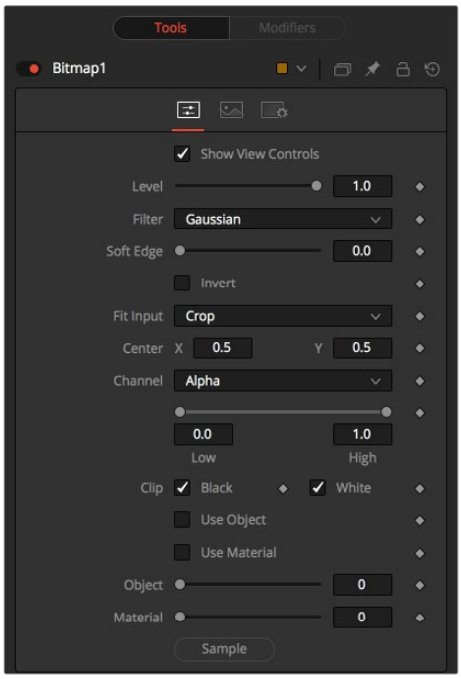
Basic Node Setup

The Bitmap mask node is not required for connecting an image into the effect mask input, but it does provide options that are otherwise unavailable. It allows for selecting channels other than RGBA for the mask, as well as softness and clipping. In the node tree below, the two Bitmap masks are combined using a Paint menu located in the second Bitmap mask node, which allows you to add, subtract, multiply, and perform other operations on the combined mask.



Bitmap nodes can be chained together for more advanced matte operations.

Inspector



Bitmap Mask controls

Controls Tab

The Controls tab is used to refine how the image connected to the orange input converts into the Bitmap mask.

Show View Controls

The Show View Controls checkbox is used to enable/disable the display of the mask's onscreen controls in the viewer. Onscreen controls, including center position, polylines, angles, and others, do not appear when this checkbox is disabled, even when the node is selected.

Level

The Level control sets the transparency level of the pixels in the mask channel. When the value is 1.0, the mask is completely opaque (unless it has a soft edge). Lower values cause the mask to be partially transparent. The result is identical to lowering the Blend control of an effect.

> **NOTE:** Lowering the level of a mask lowers the values of all pixels covered by the mask in the mask channel. For example, if a Circle mask is placed over a Rectangle mask, lowering the level of the Circle mask lowers the values of all the pixels in the mask channel, even though the Rectangle mask beneath it is still opaque.

Filter

This control selects the filtering algorithm used when applying Soft Edge to the mask.

- **Box:** This is the fastest method but at reduced quality. Box is best suited for minimal amounts of blur.
- **Bartlett:** Otherwise known as a Pyramid filter, Bartlett makes a good compromise between speed and quality.
- **Multi-box:** When selecting this filter, the Num Passes slider appears and lets you control the quality. At 1 and 2 passes, results are identical to Box and Bartlett, respectively. At 4 passes and above, results are usually as good as Gaussian, in less time and with no edge "ringing."
- **Gaussian:** The Gaussian filter uses a true Gaussian approximation and gives excellent results, but it is a little slower than the other filters. In some cases, it can produce an extremely slight edge "ringing" on floating-point pixels.

Soft Edge

Use the Soft Edge slider to blur (feather) the mask, using the selected filter. Higher values cause the edge to fade off well beyond the boundaries of the mask. A value of 0.0 creates a crisp, well-defined edge.

Paint Mode

Connecting a mask to the effect mask input displays the Paint mode menu. The Paint mode is used to determine how the incoming mask for the effect mask input and the mask created in the node are combined.

- **Merge:** Merge is the default for all masks. The new mask is merged with the input mask.
- **Add:** The mask's values add to the input mask's values.
- **Subtract:** In the intersecting areas, the new mask values subtract from the input mask's values.
- **Minimum:** Comparing the input mask's values and the new mask, this displays the lowest (minimum) value.
- **Maximum:** Comparing the input mask's values and the new mask, this displays the highest (maximum) value.
- **Average:** This calculates the average (half the sum) of the new mask and the input mask.
- **Multiply:** This multiplies the values of the input mask by the new mask's values.
- **Replace:** The new mask completely replaces the input mask wherever they intersect. Areas that are zero (completely black) in the new mask do not affect the input mask.
- **Invert:** Areas of the input mask that are covered by the new mask are inverted; white becomes black and vice versa. Gray areas in the new mask are partially inverted.
- **Copy:** This mode completely discards the input mask and uses the new mask for all values.
- **Ignore:** This mode completely discards the new mask and uses the input mask for all values.

Invert

Selecting this checkbox inverts the entire mask. Unlike the Invert Paint mode, this checkbox affects all pixels, regardless of whether the new mask covers them.

Fit Input

This menu is used to select how the image source is treated if it does not fit the dimensions of the generated mask.

In the example below, a 720 x 576 image source (yellow) is used to generate a 1920 x 1080 mask (gray).

Crop: If the image source is smaller than the generated mask, it will be placed according to the X/Y controls, masking off only a portion of the mask. If the image source is larger than the generated mask, it will be placed according to the X/Y controls and cropped off at the borders of the mask.



Stretch: The image source will be stretched in X and Y to accommodate the full dimensions of the generated mask. This might lead to visible distortions of the image source.



Inside: The image source will be scaled uniformly until one of its dimensions (X or Y) fits the inside dimensions of the mask. Depending on the relative dimensions of the image source and mask background, either the image source's width or height may be cropped to fit the respective dimensions of the mask.



Width: The image source will be scaled uniformly until its width (X) fits the width of the mask. Depending on the relative dimensions of the image source and mask, the image source's Y dimension might not fit the mask's Y dimension, resulting in either cropping of the image source in Y or the image source not covering the mask's height entirely.



Height: The image source will be scaled uniformly until its height (Y) fits the height of the mask. Depending on the relative dimensions of the image source and mask, the image source's X-dimension might not fit the mask's X-dimension, resulting in either cropping of the image source in X or the image source not covering the mask's width entirely.



Outside: The image source will be scaled uniformly until one of its dimensions (X or Y) fits the outside dimensions of the mask. Depending on the relative dimensions of the image source and mask, either the image source's width or height may be cropped or not fit the respective dimension of the mask.



Center X and Y

These controls adjust the position of the Bitmap mask.

Channel

The Channel menu determines the Channel of the input image used to create the mask. Choices include the red, green, blue, and alpha channels, the hue, luminance, or saturation values, or the auxiliary coverage channel of the input image (if one is provided).

Threshold Low/High

The Threshold range control can be used to clip the bitmap image. Increasing the low range control will clip pixels below the specified value to black (0.0). Decreasing the high range control will force pixels higher than the specified value to white (1.0).

Use Object/Use Material

This control has no effect unless the input image contains a Material or Object ID channel. When toggled on, the Object ID and Material ID are used to create a mask based on the selected object or material. When toggled off, the regular color channels will generate the mask.

Common Controls

Image and Settings Tabs

The Image and Settings tabs in the Inspector are also duplicated in other Mask nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

B-Spline Mask [BSp]



The B-Spline node

A B-Spline mask is identical to a Polygon mask in all respects except one. Where Polygon masks use Bézier splines, this mask node uses B-Splines. Where Bézier splines employ a central point and two handles to manage the smoothing of the spline segment, a B-Spline requires only a single point. This means that a B-Spline shape requires far fewer control points to create a nicely smoothed shape.

When first added to a node, the B-Spline mask consists of only Center control, which is visible onscreen. Points are added to the B-Spline by clicking in the viewer. Each new point is connected to the last one created, but instead of the spline going directly through each control point, B-Spline control points only influence the spline shape. The control point pulls the spline in its direction to create a smooth curve.

Like the Polygon mask tool, the B-Spline mask auto-animates. Adding this node to the Node Editor adds a keyframe to the current frame. Moving to a new frame and changing the shape creates a new keyframe and interpolate between the two defined shapes.

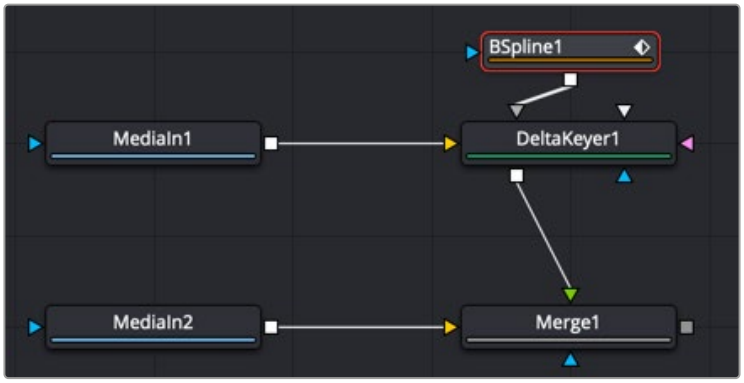
Inputs

The B-Spline mask node includes a single effect mask input.

Effect Mask: The optional blue input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input combines the masks. How masks are combined is handled in the Paint mode menu in the Inspector.

Basic Node Setup

The B-Spline node can be used to generate a single smooth spline shape or combined with other masks for more complex shapes. In the node tree below, the B-Spline mask is used to generate a smooth, curved shape as a garbage matte on the Delta Keyer.



A B-Spline node generates a smooth, curved shape as a garbage matte

Inspector



B-Spline Mask controls

Controls Tab

The Controls tab is used to refine how the B-Spline appears after drawing it in the viewer.

Show View Controls

The Show View Controls checkbox is used to enable/disable the display of the mask's onscreen controls in the viewer. Onscreen controls, including center position, polylines, angles, and others, do not appear when this checkbox is disabled, even when the node is selected.

Level

The Level control sets the transparency level of the pixels in the mask channel. When the value is 1.0, the mask is completely opaque (unless it has a soft edge). Lower values cause the mask to be partially transparent. The result is identical to lowering the blend control of an effect.

> **NOTE:** Lowering the level of a mask lowers the values of all pixels covered by the mask in the mask channel. For example, if a Circle mask is placed over a Rectangle mask, lowering the level of the Circle mask lowers the values of all of the pixels in the mask channel, even though the Rectangle mask beneath it is still opaque.

Filter

This control selects the filtering algorithm used when applying Soft Edge to the mask.

- **Box:** This is the fastest method but at reduced quality. Box is best suited for minimal amounts of blur.
- **Bartlett:** Otherwise known as a Pyramid filter, Bartlett makes a good compromise between speed and quality.

- **Multi-box:** When selecting this filter, the Num Passes slider appears and lets you control the quality. At 1 and 2 passes, results are identical to Box and Bartlett, respectively. At 4 passes and above, results are usually as good as Gaussian, in less time and with no edge "ringing."
- **Gaussian:** The Gaussian filter uses a true Gaussian approximation and gives excellent results, but it is a little slower than the other filters. In some cases, it can produce an extremely slight edge "ringing" on floating-point pixels.

Soft Edge

Use the Soft Edge slider to blur (feather) the mask, using the selected filter. Higher values cause the edge to fade off well beyond the boundaries of the mask. A value of 0.0 creates a crisp, well-defined edge.

Border Width

The Border Width control adjusts the thickness of the mask's edge. When the solid checkbox is toggled on, the border thickens or narrows the mask. When the mask is not solid, the mask shape draws as an outline, and the width uses the Border Width setting.

Paint Mode

Connecting a mask to the effect mask input displays the Paint mode menu. The Paint mode is used to determine how the incoming mask for the effect mask input and the mask created in the node are combined.

- **Merge:** Merge is the default for all masks. The new mask is merged with the input mask.
- **Add:** The mask's values add to the input mask's values.
- **Subtract:** In the intersecting areas, the new mask values subtract from the input mask's values.
- **Minimum:** Comparing the input mask's values and the new mask, this displays the lowest (minimum) value.
- **Maximum:** Comparing the input mask's values and the new mask, this displays the highest (maximum) value.
- **Average:** This calculates the average (half the sum) of the new mask and the input mask.
- **Multiply:** This multiplies the values of the input mask by the new mask's values.
- **Replace:** The new mask completely replaces the input mask wherever they intersect. Areas that are zero (completely black) in the new mask do not affect the input mask.
- **Invert:** Areas of the input mask that are covered by the new mask are inverted; white becomes black and vice versa. Gray areas in the new mask are partially inverted.
- **Copy:** This mode completely discards the input mask and uses the new mask for all values.
- **Ignore:** This mode completely discards the new mask and uses the input mask for all values.

Invert

Selecting this checkbox inverts the entire mask. Unlike the Invert Paint mode, the checkbox affects all pixels, regardless of whether the new mask covers them or not.

Solid

When the Solid checkbox is enabled, the mask is filled to be transparent (white) unless inverted. When disabled, the spline is drawn as just an outline whose thickness is determined by the Border Width slider.

Center X and Y

These controls adjust the position of the B-Spline mask.

Size

Use the Size control to adjust the scale of the B-Spline effect mask, without affecting the relative behavior of the points that compose the mask or setting a keyframe in the mask animation.

X, Y, and Z Rotation

Use these three controls to adjust the rotation angle of the mask along any axis.

Fill Method

The Fill Method menu offers two different techniques for dealing with overlapping regions of a polyline. If overlapping segments in a mask are causing undesirable holes to appear, try switching the setting of this control from Alternate to Non Zero Winding.

Right-Click Here for Shape Animation

By default, all B-Spline masks are animated when they are created. The initial keyframe is set to the current time, and any changes to the shape at different times will create new keys.

Right-clicking on this label will display a contextual menu that offers options for removing or re-adding animation to the mask, or publishing and connecting the masks.

Adding Points

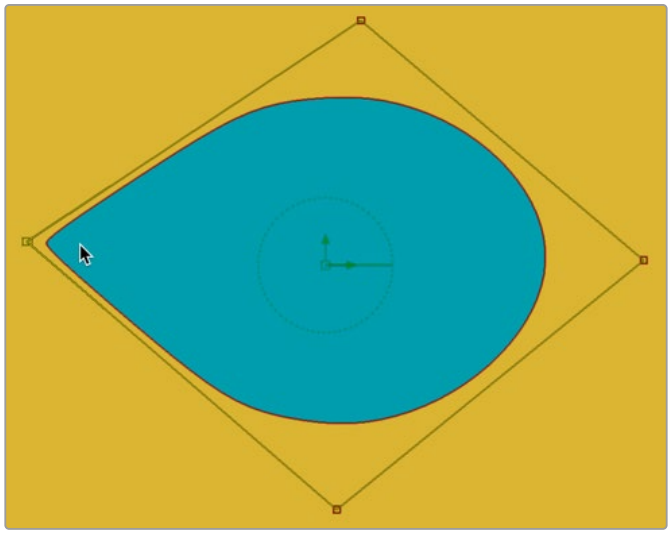
Adding Points to a B-Spline effect mask is relatively simple. Immediately after adding the node to the Node Editor, there are no points, but the tool will be in Click Append mode. Click once in the viewer wherever a point is required for the mask. Continue clicking to draw the shape of the mask.

When the shape is complete, click on the initial point again to close the mask.

When the shape is closed, the mode of the polyline changes to Insert and Modify. This allows you to add and adjust additional points on the mask by clicking the spline segments. To lock down the mask's shape and prevent accidental changes, switch the Polyline mode to Done using the Polyline toolbar or contextual menu.

Adjusting Tension on a B-Spline

The tension of the control point determines the smoothness of a B-Spline. To adjust the tension of a B-Spline's control points, select the point in the viewer, hold down the W key and drag the mouse pointer to the left and right to increase or decrease the tension of the curve through that point.



Adjusting Tension on a B-Spline

B-Spline Toolbar

When a B-Spline mask is selected in the Node Editor, a toolbar appears above the viewer with buttons for easy access to the modes. Position the pointer over any button in the toolbar to display a tooltip that describes that button's function.



B-Spline Mask Polygon toolbar

You can change the way the toolbar is displayed by right-clicking on the toolbar and selecting from the options displayed in the toolbar's contextual menu.

The functions of the buttons in this toolbar are explained in depth in the Polylines section.

Common Controls

Image and Settings Tabs

The Image and Settings tabs in the Inspector are also duplicated in other mask nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Ellipse Mask [Elp]



The Ellipse node

Ellipse Mask Node Introduction

The Ellipse mask is most useful for masking round objects. It is a circle by default, but independent control is offered over the width, height, and angle, providing for a wide variety of ellipsoidal shapes.

Inputs

The Ellipse mask node includes a single effect mask input.

Effect Mask: The optional blue input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input combines the masks. How masks are combined is handled in the Paint mode menu in the Inspector.

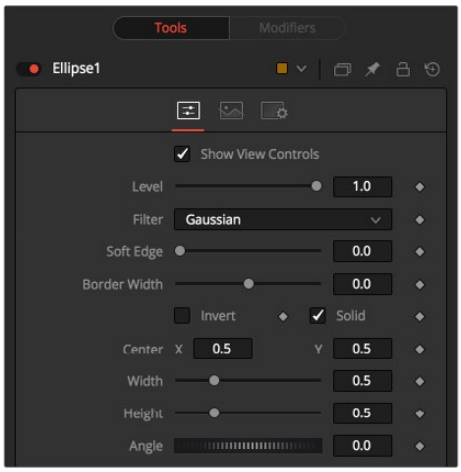
Basic Node Setup

The Ellipse mask node is useful for generating circular masks. Below, the Ellipse mask is used to generate a circular colored shape by cutting a circular shape from a background node.



Ellipse mask node creating a colored circle by masking a background node

Inspector



Ellipse Mask controls

Controls Tab

The Controls tab is used to refine how the ellipse appears after drawing it in the viewer.

Show View Controls

The Show View Controls checkbox is used to enable/disable the display of the mask's onscreen controls in the viewer. Onscreen controls, including center position, polylines, angles, and others, do not appear when this checkbox is disabled, even when the node is selected.

Level

The Level control sets the transparency level of the pixels in the mask channel. When the value is 1.0, the mask is completely opaque (unless it has a soft edge). Lower values cause the mask to be partially transparent. The result is identical to lowering the blend control of an effect.

> **NOTE:** Lowering the level of a mask lowers the values of all pixels covered by the mask in the mask channel. For example, if a Circle mask is placed over a Rectangle mask, lowering the level of the Circle mask lowers the values of all of the pixels in the mask channel, even though the Rectangle mask beneath it is still opaque.

Filter

This control selects the filtering algorithm used when applying Soft Edge to the mask.

- **Box:** This is the fastest method but at reduced quality. Box is best suited for minimal amounts of blur.
- **Bartlett:** Otherwise known as a Pyramid filter, Bartlett makes a good compromise between speed and quality.
- **Multi-box:** When selecting this filter, the Num Passes slider appears and lets you control the quality. At 1 and 2 passes, results are identical to Box and Bartlett, respectively. At 4 passes and above, results are usually as good as Gaussian, in less time and with no edge "ringing."
- **Gaussian:** The Gaussian filter uses a true Gaussian approximation and gives excellent results, but it is a little slower than the other filters. In some cases, it can produce an extremely slight edge "ringing" on floating-point pixels.

Soft Edge

Use the Soft Edge slider to blur (feather) the mask, using the selected filter. Higher values cause the edge to fade off well beyond the boundaries of the mask. A value of 0.0 creates a crisp, well-defined edge.

Border Width

The Border Width control adjusts the thickness of the mask's edge. When the solid checkbox is toggled on, the border thickens or narrows the mask. When the mask is not solid, the mask shape draws as an outline, and the width uses the Border Width setting.

Paint Mode

Connecting a mask to the effect mask input displays the Paint mode menu. The Paint mode is used to determine how the incoming mask for the effect mask input and the mask created in the node are combined.

- **Merge:** Merge is the default for all masks. The new mask is merged with the input mask.
- **Add:** The mask's values add to the input mask's values.
- **Subtract:** In the intersecting areas, the new mask values subtract from the input mask's values.
- **Minimum:** Comparing the input mask's values and the new mask, this displays the lowest (minimum) value.
- **Maximum:** Comparing the input mask's values and the new mask, this displays the highest (maximum) value.
- **Average:** This calculates the average (half the sum) of the new mask and the input mask.
- **Multiply:** This multiplies the values of the input mask by the new mask's values.
- **Replace:** The new mask completely replaces the input mask wherever they intersect. Areas that are zero (completely black) in the new mask do not affect the input mask.
- **Invert:** Areas of the input mask that are covered by the new mask are inverted; white becomes black and vice versa. Gray areas in the new mask are partially inverted.
- **Copy:** This mode completely discards the input mask and uses the new mask for all values.
- **Ignore:** This mode completely discards the new mask and uses the input mask for all values.

Invert

Selecting this checkbox inverts the entire mask. Unlike the Invert Paint mode, the checkbox affects all pixels, regardless of whether the new mask covers them or not.

Solid

When the Solid checkbox is enabled, the mask is filled to be transparent (white) unless inverted. When disabled, the spline is drawn as just an outline whose thickness is determined by the Border Width slider.

Center X and Y

These controls adjust the position of the Ellipse mask.

Width

This control allows independent control of the ellipse mask's Width. In addition to the slider in the mask's controls, interactively drag the width (left or right edge) of the mask on the viewer using the pointer. Any changes will be reflected in this control.

Height

Height allows independent control of the ellipse mask's height. In addition to the slider in the mask's controls, interactively drag the height (top or bottom edge) of the mask on the view using the pointer. Any changes will be reflected in this control.

To change the mask's size without affecting the aspect ratio, drag the onscreen control between the edges (diagonal). This will modify both the width and height proportionately.

Angle

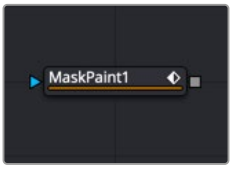
Change the rotational angle of the mask by moving the Angle control left or right. Values can be entered into the number fields provided. Alternately, use the onscreen controls by dragging the little circle at the end of the dashed angle line to interactively adjust the rotation of the ellipse.

Common Controls

Image and Settings Tabs

The Image and Settings tabs in the Inspector are also duplicated in other mask nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Mask Paint [PnM]



The Mask Paint node

Mask Paint Node Introduction

The Mask Paint node allows direct painting of mask images, using the pointer as if it was a paintbrush. In addition to regular paint strokes, it is possible to apply basic primitive shapes and polyline style strokes.

Each stroke can have a duration that lasts for the entire project, a single frame. or an arbitrary number of fields. The strokes can have independent durations in the Keyframes Editor for easy manipulation of time. Alternatively, Multistrokes is a faster but non-editable way for doing many mask clean up paint tasks.

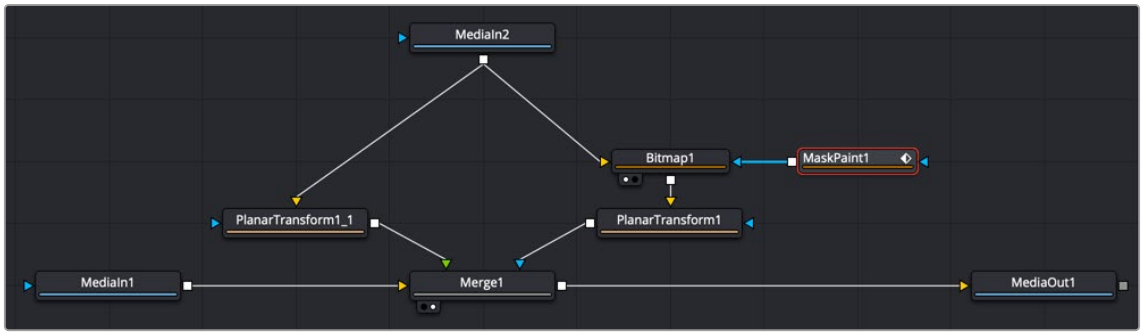
Inputs

The Paint mask node includes a single effect mask input.

Effect Mask: The optional blue input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input combines the masks. How masks are combined is handled in the Paint mode menu in the Inspector.

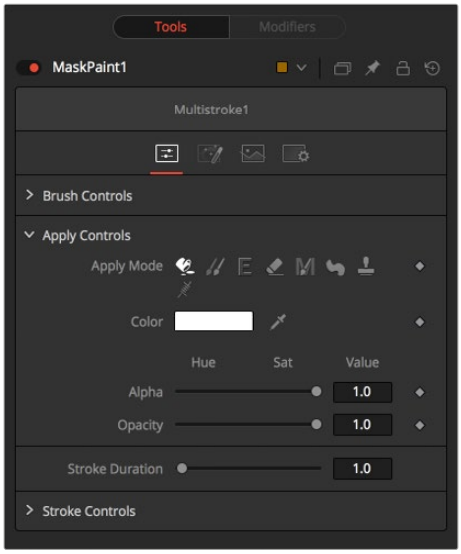
Basic Node Setup

The Mask Paint node is useful for painting masks using a more free hand, pressure sensitive style. In the node tree below, the Mask Paint node is used to patch up holes in a Bitmap mask.



A Mask Paint node can be used to repair problematic areas of a matte.

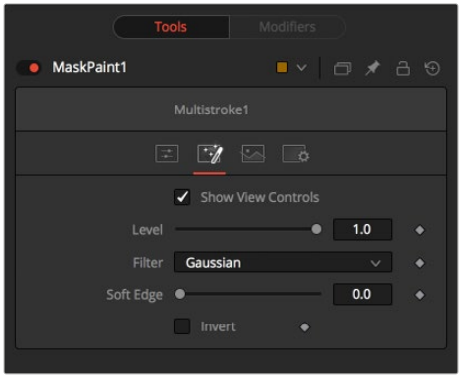
Inspector



Mask Paint controls

As the Controls tab in the Mask Paint node is fundamentally identical to the Paint node, for more detail, *see Chapter 53, "Paint Node,"* in the Fusion Reference Manual or Chapter 113 in the DaVinci Resolve Reference Manual.

The only difference between the two nodes is that, as Mask Paint operates on single-channel mask images, there is no Channel Selector control, and all color controls have only a single Alpha value. The Mask tab, however, includes several parameters that are different from the Paint tool, so they are covered below.



The Mask Paint tab

Mask Tab

The Mask tab is used to refine the basic mask parameters that do not fall into the category of "panting." These include how multiple masks are combined, overall softness control, and level control.

Show View Controls

The Show View Controls checkbox is used to enable/disable the display of the masks onscreen controls in the viewer. Onscreen controls including, center position, polylines, angles, and others, do not appear when this checkbox is disabled, even when the node is selected.

Level

The Level control sets the transparency level of the pixels in the mask channel. When the value is 1.0, the mask is completely opaque (unless it has a soft edge). Lower values cause the mask to be partially transparent. The result is identical to lowering the blend control of an effect.

> **NOTE:** Lowering the level of a mask lowers the values of all pixels covered by the mask in the mask channel. For example, if a Circle mask is placed over a Rectangle mask, lowering the level of the Circle mask lowers the values of all of the pixels in the mask channel, even though the Rectangle mask beneath it is still opaque.

Filter

This control selects the filtering algorithm used when applying Soft Edge to the mask.

- **Box:** This is the fastest method but at reduced quality. Box is best suited for minimal amounts of blur.
- **Bartlett:** Otherwise known as a Pyramid filter, Bartlett makes a good compromise between speed and quality.
- **Multi-box:** When selecting this filter, the Num Passes slider appears and lets you control the quality. At 1 and 2 passes, results are identical to Box and Bartlett, respectively. At 4 passes and above, results are usually as good as Gaussian, in less time and with no edge "ringing."
- **Gaussian:** The Gaussian filter uses a true Gaussian approximation and gives excellent results, but it is a little slower than the other filters. In some cases, it can produce an extremely slight edge "ringing" on floating-point pixels.

Soft Edge

Use the Soft Edge slider to blur (feather) the mask, using the selected filter. Higher values cause the edge to fade off well beyond the boundaries of the mask. A value of 0.0 creates a crisp, well-defined edge.

Paint Mode

Connecting a mask to the effect mask input displays the Paint mode menu. The Paint mode is used to determine how the incoming mask for the effect mask input and the mask created in the node are combined.

- **Merge:** Merge is the default for all masks. The new mask is merged with the input mask.
- **Add:** The mask's values add to the input mask's values.
- **Subtract:** In the intersecting areas, the new mask values subtract from the input mask's values.
- **Minimum:** Comparing the input mask's values and the new mask, this displays the lowest (minimum) value.
- **Maximum:** Comparing the input mask's values and the new mask, this displays the highest (maximum) value.
- **Average:** This calculates the average (half the sum) of the new mask and the input mask.
- **Multiply:** This multiplies the values of the input mask by the new mask's values.
- **Replace:** The new mask completely replaces the input mask wherever they intersect. Areas that are zero (completely black) in the new mask do not affect the input mask.
- **Invert:** Areas of the input mask that are covered by the new mask are inverted; white becomes black and vice versa. Gray areas in the new mask are partially inverted.
- **Copy:** This mode completely discards the input mask and uses the new mask for all values.
- **Ignore:** This mode completely discards the new mask and uses the input mask for all values.

Invert

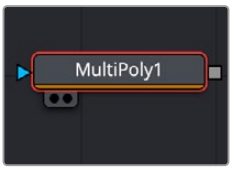
Selecting this checkbox inverts the entire mask. Unlike the Invert Paint mode, the checkbox affects all pixels, regardless of whether the new mask covers them or not.

Common Controls

Image and Settings Tabs

The Image and Settings tabs in the Inspector are also duplicated in other mask nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

MultiPoly [MPly]



The MultiPoly node

Introduction

The MultiPoly tool makes rotoscoping more efficient by allowing you to create multiple masks in a single node. The list of polygon shapes in the inspector allows you to easily modify, organize, and animate shapes through one simple interface.

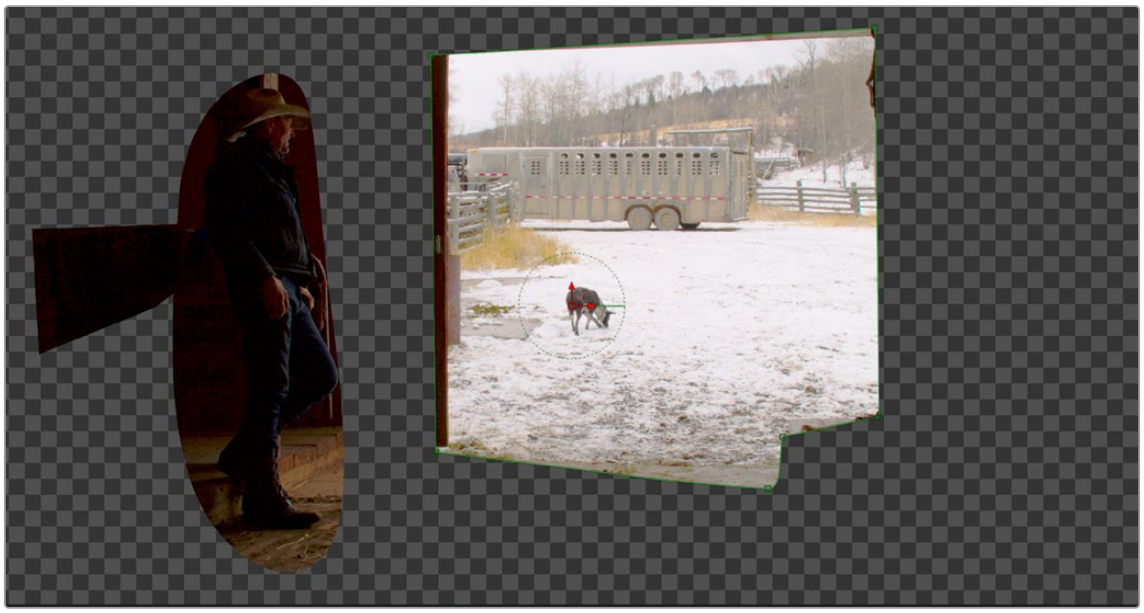
With the multiple masks presented in a scrolling list format, akin to the convenience offered by MultiMerge, you can effortlessly organize, edit, and animate individual shapes. This intuitive interface streamlines the process of drawing intricate masks, empowering you to create polished compositions with ease.

Furthermore, all these advantages are available within this List view, facilitating swift selection and modification of specific shapes, toggling their visibility, and precise adjustment of their parameters.

Usage

After adding the MultiPoly, the tool is ready for you to start drawing your first mask.

To create a mask, use the tools at the top of the viewer to create polygons, splines, and points (these points will make up segments of the polyline). You can close the mask by clicking the initial point again.



The MutliPoly tool showing three creation of masks on the viewer

Inspector



Multipoly controls

In the Inspector, pressing the Polygon or BSpline button will create a new mask type. Each new mask is added to the List view located in the Inspector.

Within the List view, you can rename or reorder any of the masks you've just added. Modify specific shapes by selecting them in the layer list and changing their properties in the lower portion of the Inspector.

To animate any of the drawn masks, select the animation diamond next to the text "Right-click here for shape animation," at the bottom of the Inspector or you can right-click on this text and select Animate.

Within the layer list, right clicking on a mask offers a few options.

Duplicate: Makes a copy of the selected mask and adds a new entry in the layer list.

Split here: Splits the selected mask and any mask below it, copying them to a new fresh MultiPoly node.

Rename: Allows you to rename the selected mask.

Reset to default: Keeps the mask entry in the layer list but removes the mask shape and resets all the controls for that mask.

Delete: Deletes the mask.

Polygon Mask [Ply]



The Polygon node

Polygon Mask Node Introduction

The Polygon mask is most useful for masking objects that do not have a regular shape. When first added to a node, the Polygon mask consists of only Center and Angle controls, which are visible onscreen. Points are added to the polyline by clicking in the viewer. Each new point is connected to the last one created.

Like the B-Spline mask tool, the Polygon mask auto-animates. Adding this node to the Node Editor adds a keyframe to the current frame. Moving to a new frame and changing the shape creates a new keyframe and interpolate between the two defined shapes.

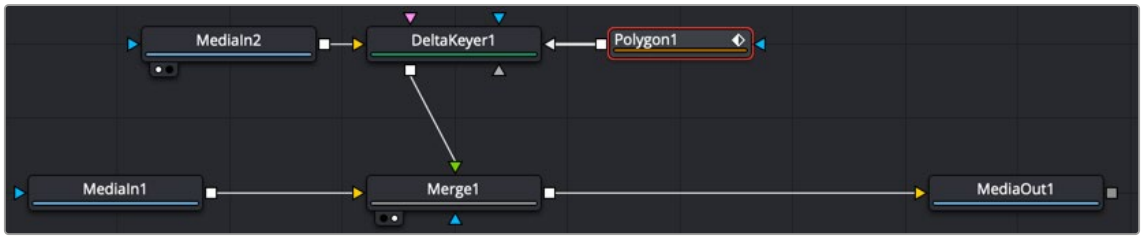
Inputs

The Polygon mask node includes a single effect mask input.

Effect Mask: The optional blue input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input combines the masks. How masks are combined is handled in the Paint mode menu in the Inspector.

Basic Node Setup

The Polygon node can be used to generate a detailed spline shape or combined with other masks for even more complex shapes. In the node tree below, the Polygon mask is used to generate detailed shape as a solid matte on the Delta Keyer.



A Polygon node generates a detailed shape as a Solid matte.

Inspector



Polygon Mask controls

Controls Tab

The Controls tab is used to refine how the polyline appears after drawing it in the viewer.

Show View Controls

The Show View Controls checkbox is used to enable/disable the display of the mask's onscreen controls in the viewer. Onscreen controls, including center position, polylines, angles, and others, do not appear when this checkbox is disabled, even when the node is selected.

Level

The Level control sets the transparency level of the pixels in the mask channel. When the value is 1.0, the mask is completely opaque (unless it has a soft edge). Lower values cause the mask to be partially transparent. The result is identical to lowering the blend control of an effect.

NOTE: Lowering the level of a mask lowers the values of all pixels covered by the mask in the mask channel. For example, if a Circle mask is placed over a Rectangle mask, lowering the level of the Circle mask lowers the values of all of the pixels in the mask channel, even though the Rectangle mask beneath it is still opaque.

Filter

This control selects the filtering algorithm used when applying Soft Edge to the mask.

- **Box:** This is the fastest method but at reduced quality. Box is best suited for minimal amounts of blur.
- **Bartlett:** Otherwise known as a Pyramid filter, Bartlett makes a good compromise between speed and quality.
- **Multi-box:** When selecting this filter, the Num Passes slider appears and lets you control the quality. At 1 and 2 passes, results are identical to Box and Bartlett, respectively. At 4 passes and above, results are usually as good as Gaussian, in less time and with no edge "ringing."
- **Gaussian:** The Gaussian filter uses a true Gaussian approximation and gives excellent results, but it is a little slower than the other filters. In some cases, it can produce an extremely slight edge "ringing" on floating-point pixels.

Soft Edge

Use the Soft Edge slider to blur (feather) the mask, using the selected filter. Higher values cause the edge to fade off well beyond the boundaries of the mask. A value of 0.0 creates a crisp, well-defined edge.

Border Width

The Border Width control adjusts the thickness of the mask's edge. When the solid checkbox is toggled on, the border thickens or narrows the mask. When the mask is not solid, the mask shape draws as an outline, and the width uses the Border Width setting.

Paint Mode

Connecting a mask to the effect mask input displays the Paint mode menu. The Paint mode is used to determine how the incoming mask for the effect mask input and the mask created in the node are combined.

- **Merge:** Merge is the default for all masks. The new mask is merged with the input mask.
- **Add:** The mask's values add to the input mask's values.
- **Subtract:** In the intersecting areas, the new mask values subtract from the input mask's values.
- **Minimum:** Comparing the input mask's values and the new mask, this displays the lowest (minimum) value.
- **Maximum:** Comparing the input mask's values and the new mask, this displays the highest (maximum) value.
- **Average:** This calculates the average (half the sum) of the new mask and the input mask.
- **Multiply:** This multiplies the values of the input mask by the new mask's values.
- **Replace:** The new mask completely replaces the input mask wherever they intersect. Areas that are zero (completely black) in the new mask do not affect the input mask.
- **Invert:** Areas of the input mask that are covered by the new mask are inverted; white becomes black and vice versa. Gray areas in the new mask are partially inverted.
- **Copy:** This mode completely discards the input mask and uses the new mask for all values.
- **Ignore:** This mode completely discards the new mask and uses the input mask for all values.

Invert

Selecting this checkbox inverts the entire mask. Unlike the Invert Paint mode, the checkbox affects all pixels, regardless of whether the new mask covers them or not.

Solid

When the Solid checkbox is enabled, the mask is filled to be transparent (white) unless inverted. When disabled, the spline is drawn as just an outline whose thickness is determined by the Border Width slider.

Center X and Y

These controls adjust the position of the polygon spline mask.

Size

Use the Size control to adjust the scale of the polygon spline effect mask, without affecting the relative behavior of the points that compose the mask or setting a keyframe in the mask animation.

X, Y, and Z Rotation

Use these three controls to adjust the rotation angle of the mask along any axis.

Fill Method

The Fill Method menu offers two different techniques for dealing with overlapping regions of a polyline. If overlapping segments in a mask are causing undesirable holes to appear, try switching the setting of this control from Alternate to Non Zero Winding.

Right-Click Here for Shape Animation

By default, all polygon spline masks are animated when they are created. The initial keyframe is set to the current time, and any changes to the shape at different times create new keys.

Right-clicking on this label displays a contextual menu that offers options for removing or re-adding animation to the mask, or publishing and connecting the masks together.

Common Controls

Image and Settings Tabs

The Image and Settings tabs in the Inspector are also duplicated in other mask nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Adding Points

Adding Points to a polygonal effect mask is relatively simple. Immediately after adding the node to the Node Editor, there are no points, but the tool will be in Click Append mode. Click once in the viewer wherever a point is required for the mask. Continue clicking to draw the shape of the mask. When the shape is complete, click on the initial point again to close the mask.

When the shape is closed, the mode of the polyline will change to Insert and Modify. This allows for the adjusting and adding of additional points to the mask by clicking on segments of the polyline. To lock down the mask’s shape and prevent accidental changes, switch the Polyline mode to Done using the Polyline toolbar or contextual menu.



B-Spline Mask Polygon toolbar

When a Polygon (or B-Spline) mask is added to a node, a toolbar appears above the viewer, offering easy access to modes. Hold the pointer over any button in the toolbar to display a tooltip that describes that button's function.

- **Click:** Click is the default option when creating a polyline (or B-Spline) mask. It is a Bézier style drawing tool. Clicking sets a control point and appends the next control point when you click again in a different location.
- **Draw:** Draw is a freehand drawing tool. It creates a mask similar to drawing with a pencil on paper. You can create a new mask using the Draw tool, or you can extend an existing open spline by clicking the Draw tool and starting to draw from the last control point.
- **Insert:** Insert adds a new control point along the spline.
- **Modify:** Modify allows you to safely move or smooth any exiting point along a spline without worrying about adding new points accidentally.
- **Done:** Prevents any point along the spline from being moved or modified. Also, new points cannot be added. You can, however, move and rotate the entire spline.
- **Closed:** Closes an open spline.
- **Smooth:** Changes the selected control point from a linear to a smooth curve.
- **Linear:** Changes the selected control point from a smooth curve to linear.
- **Select All:** Selects all the control points on the spline.
- **Keys:** Shows or hides the control points along the spline.
- **Handles:** Shows or hides the Bézier handles along the polyline.
- **Shape:** Places a reshape rectangle around the selected spline shape. Using the reshape rectangle, you can deform groups of control points or entire shapes much easier than modifying each point.
- **Delete:** Deletes the selected control point(s).
- **Reduce:** Opens a Freehand precision window that can be used to reduce the number of controls points on a spline. This can make the paint stroke easier to modify, especially if it has been created using the Draw tool.
- **Publish menu:** You can use the publish menu to select between publishing the control points or the path. Publishing is a form of parameter linking, it makes the selected item available for use by other controls. It also allows you to attach a control point to a tracker.
- **Follow Points:** Allows a selected point to follow the path of a published point. The point follows the published point using an offset position.
- **Double Poly:** Allows softening part of the spline curve while keeping other portions of the curve sharp. The double polyline is composed of two shapes, an inner and outer shape. The inner shape is the original shape from the single polyline, whereas the outer shape is used to determine the spread of the softness. The further the outer shape gets from the inner shape, the softer that segment of the shape becomes. Both polylines start with exactly the same shape as the original single polyline, keeping the mask sharp to start. Any animation already applied to the shape remains. To select the outer shape, press the Tab key to cycle between the onscreen controls until the dashed outline is visible, or you can select the outer polyline using the contextual menu's Controls > Outer Polygon menu.
- **Multiframe:** Multiframe is a method of adjusting control points across multiple keyframes. The default setting of none only adjusts the control point of a spline on the current keyframe. Setting the menu to All adjusts the controls point for all keyframes. Prev settings adjust the current and previous keyframe while Next adjusts the current and next keyframe.
- **Onion Skinning:** Enabling onion skinning displays a mix in the viewer of the spline animation. It is useful when aligning spline animation and motion. Selecting Onion Skin Settings from the drop down menu allows you to set the number of overlapping frames.

- **Roto Assist:** Enable the Roto Assist button when you begin painting with the Polyline Stroke tool. The polyline points snap to the closest edge as you click to add points to the shape. A cyan outline indicates the points that have snapped to an edge. There are three main Roto Assist options selectable through the drop down menu:

  Multiple Points: When enabled, a single click on a high contrast edge adds multiple points to define the entire edge, instead of having to add each point individually. This is a one time only click. The second click reverts to single point edge detection.

  Distance 8: Opens a dialog where you can set the pixel range within which searching for an edge takes place.

  Reset: Used for resetting the snap attribute of all snapped points. After resetting, the points become unavailable for tracking.

Change the way the toolbar is displayed by right-clicking on the toolbar and selecting from the options displayed in the toolbar's contextual menu. The functions of the buttons in this toolbar are explained in depth in the Polylines chapter.

Ranges Mask [RNG]



The Ranges node

Ranges Mask Node Introduction

Similar to Bitmap mask, the Ranges mask allows images from the node tree to act as masks for nodes and effects. Instead of creating a simple luminance-based mask from a given channel, Ranges allows spline-based selection of low, mid and high ranges, akin to Color Corrector.

Inputs

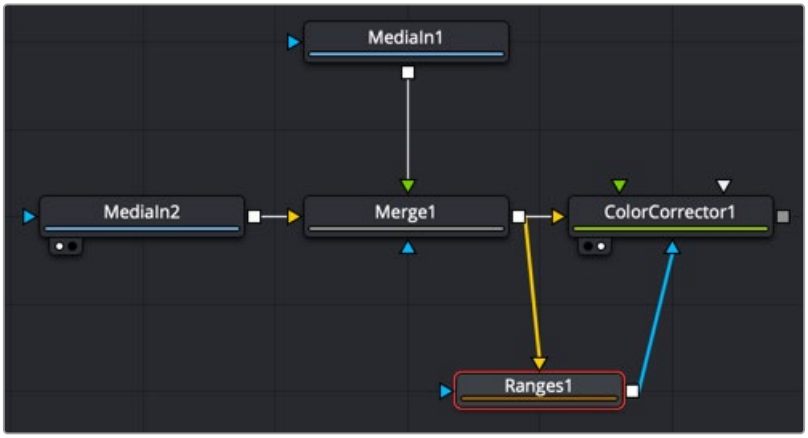
The Ranges mask node includes two inputs in the Node Editor.

Input: The orange input accepts a 2D image from which the mask will be created.

Effect Mask: The optional blue input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input combines the masks. How masks are combined is handled in the Paint mode menu in the Inspector.

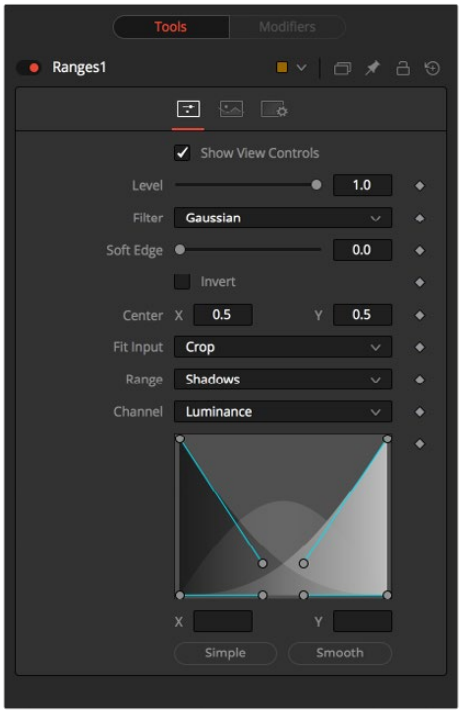
Basic Node Setup

The Ranges node is not required for connecting an image into the effect mask input, but like the Bitmap node, it does provide options that are otherwise unavailable. It allows for selecting channels other than RGBA for the mask, as well as softness and clipping. In the node tree below, the Ranges node takes the composite out of the merge, creating a mask for the color correction.



A Ranges node selects a specific range in the image to create a mask.

Inspector



Rangers Mask controls

Controls Tab

The Controls tab is used to refine how the image connected to the orange input converts into the ranges mask.

Show View Controls

The Show View Controls checkbox is used to enable/disable the display of the mask's onscreen controls in the viewer. Onscreen controls, including center position, polylines, angles, and others, do not appear when this checkbox is disabled, even when the node is selected.

Level

The Level control sets the transparency level of the pixels in the mask channel. When the value is 1.0, the mask is completely opaque (unless it has a soft edge). Lower values cause the mask to be partially transparent. The result is identical to lowering the blend control of an effect.

> **NOTE:** Lowering the level of a mask lowers the values of all pixels covered by the mask in the mask channel. For example, if a Circle mask is placed over a Rectangle mask, lowering the level of the Circle mask lowers the values of all of the pixels in the mask channel, even though the Rectangle mask beneath it is still opaque.

Filter

This control selects the filtering algorithm used when applying Soft Edge to the mask.

- **Box:** This is the fastest method but at reduced quality. Box is best suited for minimal amounts of blur.
- **Bartlett:** Otherwise known as a Pyramid filter, Bartlett makes a good compromise between speed and quality.
- **Multi-box:** When selecting this filter, the Num Passes slider appears and lets you control the quality. At 1 and 2 passes, results are identical to Box and Bartlett, respectively. At 4 passes and above, results are usually as good as Gaussian, in less time and with no edge "ringing."
- **Gaussian:** The Gaussian filter uses a true Gaussian approximation and gives excellent results, but it is a little slower than the other filters. In some cases, it can produce an extremely slight edge "ringing" on floating-point pixels.

Soft Edge

Use the Soft Edge slider to blur (feather) the mask, using the selected filter. Higher values cause the edge to fade off well beyond the boundaries of the mask. A value of 0.0 creates a crisp, well-defined edge.

Paint Mode

Connecting a mask to the effect mask input displays the Paint mode menu. The Paint mode is used to determine how the incoming mask for the effect mask input and the mask created in the node are combined.

- **Merge:** Merge is the default for all masks. The new mask is merged with the input mask.
- **Add:** The mask's values add to the input mask's values.
- **Subtract:** In the intersecting areas, the new mask values subtract from the input mask's values.
- **Minimum:** Comparing the input mask's values and the new mask, this displays the lowest (minimum) value.
- **Maximum:** Comparing the input mask's values and the new mask, this displays the highest (maximum) value.
- **Average:** This calculates the average (half the sum) of the new mask and the input mask.
- **Multiply:** This multiplies the values of the input mask by the new mask's values.
- **Replace:** The new mask completely replaces the input mask wherever they intersect. Areas that are zero (completely black) in the new mask do not affect the input mask.
- **Invert:** Areas of the input mask that are covered by the new mask are inverted; white becomes black and vice versa. Gray areas in the new mask are partially inverted.
- **Copy:** This mode completely discards the input mask and uses the new mask for all values.
- **Ignore:** This mode completely discards the new mask and uses the input mask for all values.

Invert

Selecting this checkbox inverts the entire mask. Unlike the Invert Paint mode, the checkbox affects all pixels, regardless of whether the new mask covers them or not.

Center X and Y

These controls adjust the position of the ranges mask.

Fit Input

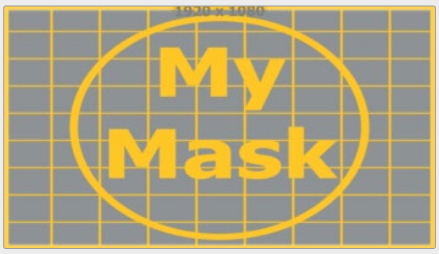
This menu is used to select how the image source is treated if it does not fit the dimensions of the generated mask.

For example, below, a 720 x 576 image source (yellow) is used to generate a 1920 x 1080 mask (gray).

Crop: If the image source is smaller than the generated mask, it is placed according to the X/Y controls, masking off only a portion of the mask. If the image source is larger than the generated mask it is placed according to the X/Y controls and cropped off at the borders of the mask.



Stretch: The image source is stretched in X and Y to accommodate the full dimensions of the generated mask. This might lead to visible distortions of the image source.



Inside: The image source is scaled uniformly until one of its dimensions (X or Y) fits the inside dimensions of the mask. Depending on the relative dimensions of the image source and mask background, either the image source's width or height may be cropped to fit the respective dimension of the mask.



Width: The image source is scaled uniformly until its width (X) fits the width of the mask. Depending on the relative dimensions of the image source and mask, the image source's Y dimension might not fit the mask's Y dimension, resulting in either cropping of the image source in Y or the image source not covering the mask's height entirely.



Height: The image source is scaled uniformly until its height (Y) fits the height of the mask. Depending on the relative dimensions of the image source and mask, the image source's X dimension might not fit the mask's X dimension, resulting in either cropping of the image source in X or the image source not covering the mask's width entirely.



Outside: The image source is scaled uniformly until one of its dimensions (X or Y) fits the outside dimensions of the mask. Depending on the relative dimensions of the image source and mask, either the image source's width or height may be cropped or not fit the respective dimension of the mask.



Channel

The Channel menu determines the Channel of the input image used to create the mask. Choices include the red, green, blue, and alpha channels; the hue, luminance, or saturation values; or the auxiliary coverage channel of the input image (if one is provided).

Shadows/Midtones/Highlights

These buttons are used to select which range is output by the node as a mask. White pixels represent pixels that are considered to be part of the range, and black pixels are not included in the range. For example, choosing Shadows would show pixels considered to be shadows as white, and pixels that are not shadows as black. Mid gray pixels are only partly in the range and do not receive the full effect of any color adjustments to that range.

Channel

The Channel selection buttons shown in this tab can be used to extract a mask from the range of a specific color channel. By default, Fusion uses the luminance channel when the color ranges are examined.

Mini Spline Editor

The extent of the ranges is selected by manipulating the spline handles. There are four spline points, each with one Bézier handle. The two handles at the top represent the start of the shadow and highlight ranges; the two handles at the bottom represent the end of the range. The Bézier handles are used to control the falloff.

The midtones range has no specific control, since its range is understood to be the space between the shadow and highlight ranges. In other words, after low and high masks have been applied, midtones are everything else.

The X and Y text controls below the Mini Spline Editor can be used to enter precise positions for the selected Bézier point or handle.

Presets

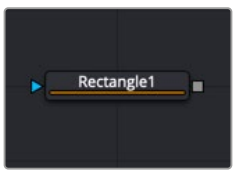
This sets the splines to two commonly-used configurations. The Simple button gives a straightforward linear-weighted selection, while the Smooth button uses a more natural falloff.

Common Controls

Image and Settings Tabs

The Image and Settings tabs in the Inspector are also duplicated in other mask nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Rectangle Mask [Rec]



The Rectangle node

Rectangle Mask Node Introduction

The Rectangle mask creates simple square or rectangular masks. By default, it creates a rectangle in the same aspect ratio as the comp, but independent control is offered over the width, height, and angle, providing for a wide variety of rectangular shapes.

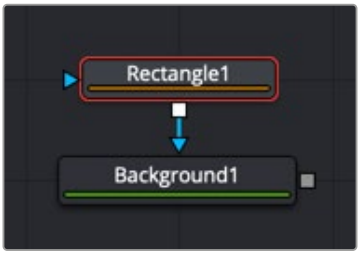
Inputs

The Rectangle mask node includes a single effect mask input.

Effect Mask: The optional blue input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input combines the masks. How masks are combined is handled in the Paint mode menu in the Inspector.

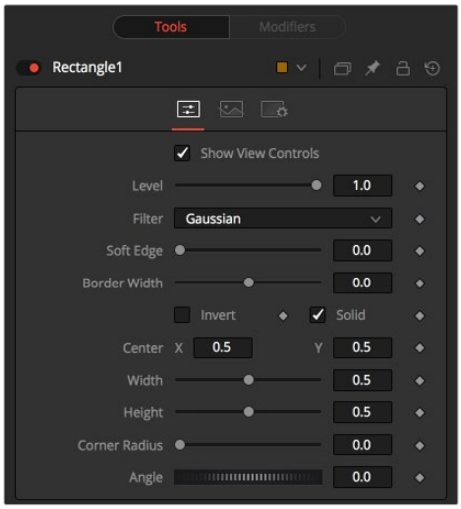
Basic Node Setup

The Rectangle mask node is useful for generating rectangular shapes. Below, the Rectangle mask is used to generate a colored square by cutting a square shape from a background node.



A Rectangle mask node creating a colored square by masking a Background node

Inspector



Rectangle Mask controls

Controls Tab

The Controls tab is used to refine how the rectangle appears after drawing it in the viewer.

Show View Controls

The Show View Controls checkbox is used to enable/disable the display of the mask's onscreen controls in the viewer. Onscreen controls, including center position, polylines, angles, and others, do not appear when this checkbox is disabled, even when the node is selected.

Level

The Level control sets the transparency level of the pixels in the mask channel. When the value is 1.0, the mask is completely opaque (unless it has a soft edge). Lower values cause the mask to be partially transparent. The result is identical to lowering the Blend control of an effect.

NOTE: Lowering the level of a mask lowers the values of all pixels covered by the mask in the mask channel. For example, if a Circle mask is placed over a Rectangle mask, lowering the level of the Circle mask lowers the values of all the pixels in the mask channel, even though the Rectangle mask beneath it is still opaque.

Filter

This control selects the filtering algorithm used when applying Soft Edge to the mask.

- **Box:** This is the fastest method but at reduced quality. Box is best suited for minimal amounts of blur.
- **Bartlett:** Otherwise known as a Pyramid filter, Bartlett makes a good compromise between speed and quality.
- **Multi-box:** When selecting this filter, the Num Passes slider appears and lets you control the quality. At 1 and 2 passes, results are identical to Box and Bartlett, respectively. At 4 passes and above, results are usually as good as Gaussian, in less time and with no edge "ringing."
- **Gaussian:** The Gaussian filter uses a true Gaussian approximation and gives excellent results, but it is a little slower than the other filters. In some cases, it can produce an extremely slight edge "ringing" on floating-point pixels.

Soft Edge

Use the Soft Edge slider to blur (feather) the mask, using the selected filter. Higher values cause the edge to fade off well beyond the boundaries of the mask. A value of 0.0 creates a crisp, well-defined edge.

Border Width

The Border Width control adjusts the thickness of the mask's edge. When the solid checkbox is toggled on, the border thickens or narrows the mask. When the mask is not solid, the mask shape draws as an outline, and the width uses the Border Width setting.

Paint Mode

Connecting a mask to the effect mask input displays the Paint mode menu. The Paint mode is used to determine how the incoming mask for the effect mask input and the mask created in the node are combined.

- **Merge:** Merge is the default for all masks. The new mask is merged with the input mask.
- **Add:** The mask's values add to the input mask's values.
- **Subtract:** In the intersecting areas, the new mask values subtract from the input mask's values.
- **Minimum:** Comparing the input mask's values and the new mask, this displays the lowest (minimum) value.
- **Maximum:** Comparing the input mask's values and the new mask, this displays the highest (maximum) value.
- **Average:** This calculates the average (half the sum) of the new mask and the input mask.
- **Multiply:** This multiplies the values of the input mask by the new mask's values.
- **Replace:** The new mask completely replaces the input mask wherever they intersect. Areas that are zero (completely black) in the new mask do not affect the input mask.
- **Invert:** Areas of the input mask that are covered by the new mask are inverted: white becomes black and vice versa. Gray areas in the new mask are partially inverted.
- **Copy:** This mode completely discards the input mask and uses the new mask for all values.
- **Ignore:** This mode completely discards the new mask and uses the input mask for all values.

Invert

Selecting this checkbox inverts the entire mask. Unlike the Invert Paint mode, this checkbox affects all pixels, regardless of whether the new mask covers them.

Solid

When the Solid checkbox is enabled, the mask is filled to be transparent (white) unless inverted. When disabled, the spline is drawn as just an outline whose thickness is determined by the Border Width slider.

Center X and Y

These controls adjust the position of the Rectangle mask.

Width and Height

Use these controls to change the X or Y scale of the rectangular effect mask independently of each other. Alternatively, drag the edges of the rectangle in the viewer to interactively adjust its size.

Corner Radius

Corner Radius allows the corners of the Rectangle mask to be rounded. A value of 0.0 is not rounding at all, which means that the rectangle has sharp corners. A value of 1.0 applies the maximum amount of rounding to the corners.

Angle

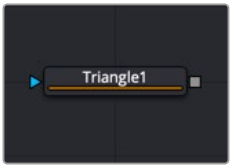
Change the rotation angle of an effect mask by moving the Angle control left or right. Values can be entered in the provided input boxes. Alternatively, use the onscreen controls by dragging the little circle at the end of the dashed angle line to interactively adjust the rotation of the ellipse.

Common Controls

Image and Settings Tabs

The Image and Settings tabs in the Inspector are also duplicated in other mask nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Triangle Mask [Tri]



The Triangle node

Triangle Mask Node Introduction

The Triangle mask creates simple triangular masks. It is unique in that it has no Center, Size, or Angle control. Unlike most other types of masks, all three points of the triangle can attach to a tracker or motion path.

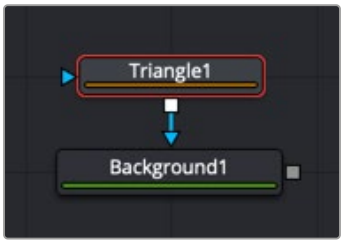
Inputs

The Triangle mask node includes a single effect mask input.

Effect Mask: The optional blue input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input combines the masks. How masks are combined is handled in the Paint mode menu in the Inspector.

Basic Node Setup

The Triangle mask node is useful for generating triangular shapes. Below, the Triangle mask is used to generate a colored triangle by cutting a triangular shape from a background node.



A Triangle mask node creating a colored Triangle shape by masking a Background node

Inspector



Triangle Mask controls

Controls Tab

The Controls tab is used to refine how the triangle appears after drawing it in the viewer.

Show View Controls

The Show View Controls checkbox is used to enable/disable the display of the mask's onscreen controls in the viewer. Onscreen controls, including center position, polylines, angles, and others, do not appear when this checkbox is disabled, even when the node is selected.

Level

The Level control sets the transparency level of the pixels in the mask channel. When the value is 1.0, the mask is completely opaque (unless it has a soft edge). Lower values cause the mask to be partially transparent. The result is identical to lowering the Blend control of an effect.

> **NOTE:** Lowering the level of a mask lowers the values of all pixels covered by the mask in the mask channel. For example, if a Circle mask is placed over a Rectangle mask, lowering the level of the Circle mask lowers the values of all the pixels in the mask channel, even though the Rectangle mask beneath it is still opaque.

Filter

This control selects the filtering algorithm used when applying Soft Edge to the mask.

- **Box:** This is the fastest method but at reduced quality. Box is best suited for minimal amounts of blur.
- **Bartlett:** Otherwise known as a Pyramid filter, Bartlett makes a good compromise between speed and quality.
- **Multi-box:** When selecting this filter, the Num Passes slider appears and lets you control the quality. At 1 and 2 passes, results are identical to Box and Bartlett, respectively. At 4 passes and above, results are usually as good as Gaussian, in less time and with no edge "ringing."
- **Gaussian:** The Gaussian filter uses a true Gaussian approximation and gives excellent results, but it is a little slower than the other filters. In some cases, it can produce an extremely slight edge "ringing" on floating-point pixels.

Soft Edge

Use the Soft Edge slider to blur (feather) the mask, using the selected filter. Higher values cause the edge to fade off well beyond the boundaries of the mask. A value of 0.0 creates a crisp, well-defined edge.

Border Width

The Border Width control adjusts the thickness of the mask's edge. When the solid checkbox is toggled on, the border thickens or narrows the mask. When the mask is not solid, the mask shape draws as an outline, and the width uses the Border Width setting.

Paint Mode

Connecting a mask to the effect mask input displays the Paint mode menu. The Paint mode is used to determine how the incoming mask for the effect mask input and the mask created in the node are combined.

- **Merge:** Merge is the default for all masks. The new mask is merged with the input mask.
- **Add:** The mask's values add to the input mask's values.
- **Subtract:** In the intersecting areas, the new mask values subtract from the input mask's values.
- **Minimum:** Comparing the input mask's values and the new mask, this displays the lowest (minimum) value.
- **Maximum:** Comparing the input mask's values and the new mask, this displays the highest (maximum) value.
- **Average:** This calculates the average (half the sum) of the new mask and the input mask.
- **Multiply:** This multiplies the values of the input mask by the new mask's values.

- **Replace:** The new mask completely replaces the input mask wherever they intersect. Areas that are zero (completely black) in the new mask do not affect the input mask.
- **Invert:** Areas of the input mask that are covered by the new mask are inverted: white becomes black and vice versa. Gray areas in the new mask are partially inverted.
- **Copy:** This mode completely discards the input mask and uses the new mask for all values.
- **Ignore:** This mode completely discards the new mask and uses the input mask for all values.

Invert

Selecting this checkbox inverts the entire mask. Unlike the Invert Paint mode, this checkbox affects all pixels, regardless of whether the new mask covers them.

Solid

When the Solid checkbox is enabled, the mask is filled to be transparent (white) unless inverted. When disabled, the spline is drawn as just an outline whose thickness is determined by the Border Width slider.

Point 1, Point 2, Point 3

These controls show the position coordinates of the three corners of the triangle. Each point can be published, connected to other controls, animated with a path, or attached to trackers. These tasks are performed by right-clicking the Position control or directly on the point in the viewer.

Common Controls

Image and Settings Tabs

The Image and Settings tabs in the Inspector are also duplicated in other mask nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Wand Mask [Wnd]



The Wand node

Wand Mask Node Introduction

The Wand mask masks an image based on a wand-style selection, similar to the Magic Wand tool found in Adobe Photoshop. As with a Bitmap mask, any image in the composition can be the source of the mask. Generally, the default is most useful, where the source image is the input of the node to which the mask is applied.

When adding a Wand mask to a node, a crosshair appears in the viewers. This crosshair should be positioned in the image to select the color used to create the Wand mask. The mask is created by examining the pixel color beneath the selection point and adding that color to the mask. The mask then expands to examine the pixels surrounding the selection point. Surrounding pixels are added to the mask if they are the same color. The mask stops expanding when no connecting pixels fall within the color range of the mask.

Inputs

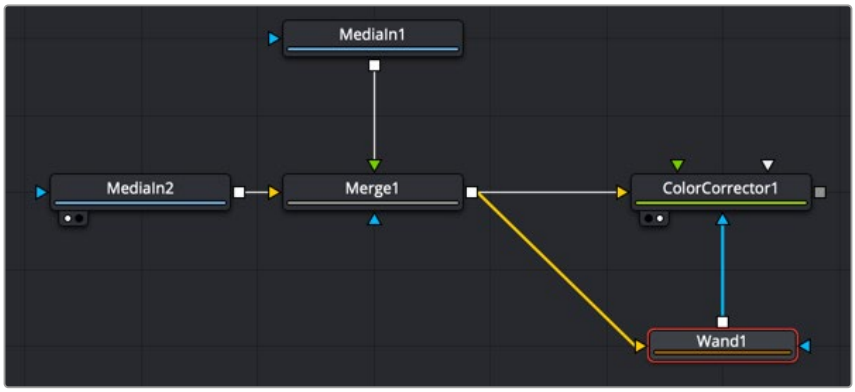
The Wand mask node includes two inputs in the Node Editor.

Input: The orange input accepts a 2D image from which the mask is created.

Effect Mask: The optional blue input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input combines the masks. How masks are combined is handled in the Paint mode menu in the Inspector.

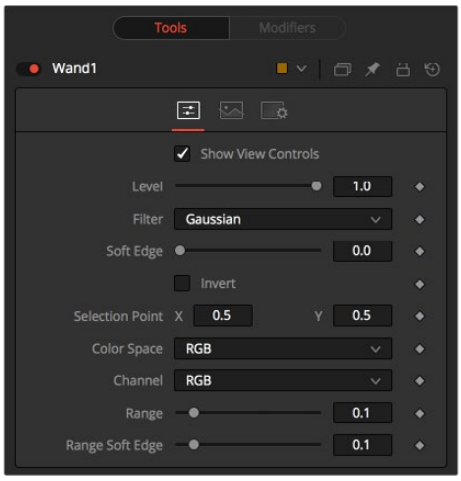
Basic Node Setup

The Wand mask node is not required for connecting an image into the effect mask input, but like the Bitmap node, it does provide options that are otherwise unavailable. It allows for selecting channels other than RGBA for the mask, as well as softness and clipping. In the node tree below, the Wand node takes the composite out of the merge, creating a mask for the color correction.



A Wand node selects a specific area in the image to create a mask.

Inspector



Wand Mask controls

Controls Tab

The Controls tab is used to refine how the mask appears after the Wand makes a selection in the viewer.

Show View Controls

The Show View Controls checkbox is used to enable/disable the display of the mask's onscreen controls in the viewer. Onscreen controls, including center position, polylines, angles, and others, do not appear when this checkbox is disabled, even when the node is selected.

Level

The Level control sets the transparency level of the pixels in the mask channel. When the value is 1.0, the mask is completely opaque (unless it has a soft edge). Lower values cause the mask to be partially transparent. The result is identical to lowering the Blend control of an effect.

> **NOTE:** Lowering the level of a mask lowers the values of all pixels covered by the mask in the mask channel. For example, if a Circle mask is placed over a Rectangle mask, lowering the level of the Circle mask lowers the values of all the pixels in the mask channel, even though the Rectangle mask beneath it is still opaque.

Filter

This control selects the filtering algorithm used when applying Soft Edge to the mask.

- **Box:** This is the fastest method but at reduced quality. Box is best suited for minimal amounts of blur.
- **Bartlett:** Otherwise known as a Pyramid filter, Bartlett makes a good compromise between speed and quality.
- **Multi-box:** When selecting this filter, the Num Passes slider appears and lets you control the quality. At 1 and 2 passes, results are identical to Box and Bartlett, respectively. At 4 passes and above, results are usually as good as Gaussian, in less time and with no edge "ringing."
- **Gaussian:** The Gaussian filter uses a true Gaussian approximation and gives excellent results, but it is a little slower than the other filters. In some cases, it can produce an extremely slight edge "ringing" on floating-point pixels.

Soft Edge

Use the Soft Edge slider to blur (feather) the mask, using the selected filter. Higher values cause the edge to fade off well beyond the boundaries of the mask. A value of 0.0 creates a crisp, well-defined edge.

Paint Mode

Connecting a mask to the effect mask input displays the Paint mode menu. The Paint mode is used to determine how the incoming mask for the effect mask input and the mask created in the node are combined.

- **Merge:** Merge is the default for all masks. The new mask is merged with the input mask.
- **Add:** The mask's values add to the input mask's values.
- **Subtract:** In the intersecting areas, the new mask values subtract from the input mask's values.
- **Minimum:** Comparing the input mask's values and the new mask, this displays the lowest (minimum) value.
- **Maximum:** Comparing the input mask's values and the new mask, this displays the highest (maximum) value.
- **Average:** This calculates the average (half the sum) of the new mask and the input mask.
- **Multiply:** This multiplies the values of the input mask by the new mask's values.

- **Replace:** The new mask completely replaces the input mask wherever they intersect. Areas that are zero (completely black) in the new mask do not affect the input mask.
- **Invert:** Areas of the input mask that are covered by the new mask are inverted; white becomes black and vice versa. Gray areas in the new mask are partially inverted.
- **Copy:** This mode completely discards the input mask and uses the new mask for all values.
- **Ignore:** This mode completely discards the new mask and uses the input mask for all values.

Invert

Selecting this checkbox inverts the entire mask. Unlike the Invert Paint mode, this checkbox affects all pixels, regardless of whether the new mask covers them.

Selection Point

The Selection Point is a pair of X and Y coordinates that determines where in the source image the Wand mask derives its initial color sample. This control is also seen as a crosshair in the viewers. The selection point can be positioned manually, connected to a tracker, path, or other expressions.

Color Space

The Color Space button group determines the color space used when selecting the source color for the mask. The Wand mask can operate in RGB, YUV, HLS, or LAB color spaces.

Channel

The Channel button group is used to select whether the color that is masked comes from all three color channels of the image, the alpha channel, or an individual channel only.

The exact labels of the buttons depend on the color space selected for the Wand mask operation. If the color space is RGB, the options are R, G, or B. If YUV is the color space, the options are Y, U, or V.

Range

The Range slider controls the range of colors around the source color that are included in the mask. If the value is left at 0.0, only pixels of the same color as the source are considered part of the mask. The higher the value, the more that similar colors in the source are considered to be wholly part of the mask.

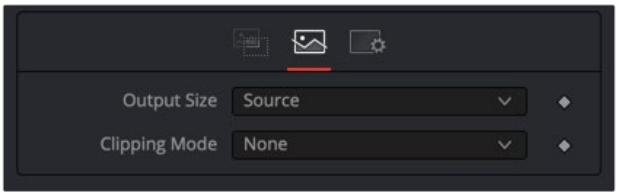
Range Soft Edge

The Range Soft Edge determines the falloff range of the colors selected. Any pixel within the range defined above are treated as 100% within the mask. If the soft range is set to 0.0, no other pixels are considered for the mask. Increasing the soft range increases the number of colors close to, but not quite within, the range included in the mask. These pixels are semitransparent in the mask

The Common Controls

Nodes that create masks share several identical controls in the Inspector. This section describes controls that are common among mask nodes.

Inspector



Mask node Image tab

Image Tab

The controls in this tab set the resolution and clipping method used by the generated mask.

Output Size

The Output size menu sets the resolution of the mask node's output. The three options include the default resolution of the comp, the source input's resolution on nodes that have an input, or a custom resolution.

Custom

When selecting Custom from the Output Size menu, the width, height, and pixel aspect of the mask created are locked to values defined in the composition's Frame Format preferences. If the Frame Format preferences change, the resolution of the mask produced is changed to match. Disabling this option can be useful for building a composition at a different resolution than the eventual target resolution for the final render.

- **Width and Height:** This pair of controls is used to set the Width and Height dimensions of the mask to be created.
- **Pixel Aspect:** This control is used to specify the Pixel Aspect ratio of the created mask. An aspect ratio of 1:1 would generate a square pixel with the same dimensions on either side (like a computer monitor), and an aspect of 0.91 would create a slightly rectangular pixel (like an NTSC monitor).
- **Depth:** The Depth drop-down menu is used to set the pixel color depth of the image created by the mask. 32-bit pixels require four times the memory of 8-bit pixels but have far greater accuracy. Float pixels allow high dynamic range values outside the normal 0..1 range, for representing colors that are brighter than white or darker than black.

NOTE: Right-click on the Width, Height, or Pixel Aspect controls to display a menu listing the file formats defined in the preferences Frame Format tab. Selecting any of the listed options sets the width, height, and pixel aspect to the values for that format.

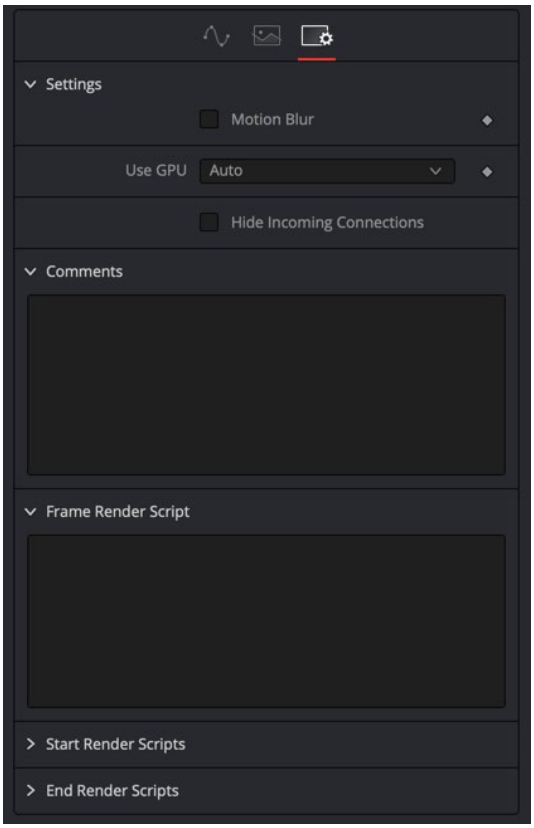
Clipping Mode

This option determines how the domain of definition rendering handles edges. The Clipping mode is most important when blur or softness is applied, which may require samples from portions of the image outside the current domain.

- **Frame:** The default option is Frame, which automatically sets the node's domain of definition to use the full frame of the image, effectively ignoring the current domain of definition. If the upstream DoD is smaller than the frame, the remaining area in the frame is treated as black/transparent.
- **None:** Setting this option to None does not perform any source image clipping. Any data required to process the node's effect that would usually be outside the upstream DoD is treated as black/transparent.

Settings Tab

The Settings tab in the Inspector can be found on every tool in the Mask category. The Settings controls are even found on third-party plugin tools. The controls are consistent and work the same way for each tool, although some tools do include one or two individual options, which are also covered here.



Common Mask Settings controls

Motion Blur

- **Motion Blur:** This toggles the rendering of Motion Blur on the tool. When this control is toggled on, the tool's predicted motion is used to produce the motion blur caused by the virtual camera's shutter. When the control is toggled off, no motion blur is created.
- **Quality:** Quality determines the number of samples used to create the blur. A quality setting of 2 causes Fusion to create two samples to either side of an object's actual motion. Larger values produce smoother results but increase the render time.
- **Shutter Angle:** Shutter Angle controls the angle of the virtual shutter used to produce the motion blur effect. Larger angles create more blur but increase the render times. A value of 360 is the equivalent of having the shutter open for one full frame exposure. Higher values are possible and can be used to create interesting effects.
- **Center Bias:** Center Bias modifies the position of the center of the motion blur. This allows for the creation of motion trail effects.
- **Sample Spread:** Adjusting this control modifies the weighting given to each sample. This affects the brightness of the samples.

Use GPU

The Use GPU menu has three settings. Setting the menu to Disable turns off hardware-accelerated rendering using the graphics card in your computer. Enabled uses the hardware, and Auto uses a capable GPU if one is available and falls back to software rendering when a capable GPU is not available

Comments

The Comments field is used to add notes to a tool. Click in the empty field and type the text. When a note is added to a tool, a small red square appears in the lower-left corner of the node when the full tile is displayed, or a small text bubble icon appears on the right when nodes are collapsed. To see the note in the Node Editor, hold the mouse pointer over the node to display the tooltip.

Scripts

Three Scripting fields are available on every tool in Fusion from the Settings tab. They each contain edit boxes used to add scripts that process when the tool is rendering. For more details on scripting nodes, please consult the Fusion scripting documentation.

Chapter 49

Matte Nodes

This chapter details the Matte nodes available in Fusion.

The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Alpha Divide [ADv] 1262
Alpha Multiply [AMl] 1263
Chroma Keyer [CKy] 1264
Clean Plate 1269
Cryptomatte [Cry] (Studio Version Only) 1272
Delta Keyer 1274
Depth Map [DMp] 1283
Difference Keyer [DfK] 1286
Luma Keyer [LKy] 1289
Magic Mask [MagM] 1293
Matte Control [Mat] 1298
Primatte [Pri] 1303
Relight [RLt] 1318
Ultra Keyer [UKy] 1322
The Common Controls 1329

Alpha Divide [ADv]



The Alpha Divide node

Alpha Divide Node Introduction

As the name gives away, the Alpha Divide's sole purpose is to divide an incoming image's color channels by its alpha channel. When you color correct an image that contains a premultiplied alpha channel, first apply an Alpha Divide node before any color correction node to create a non-premultiplied image. Then you can perform the color correction. After the color correction, add an Alpha Multiply node to return the image to its premultiplied state.

Inputs

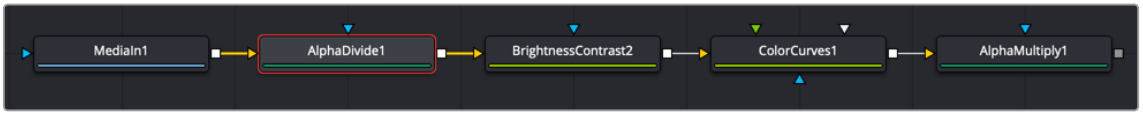
The Alpha Divide node includes two inputs in the Node Editor.

Input: The orange input accepts a 2D image with a premultiplied Alpha.

Effect Mask: The optional blue input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input limits the pixels where the Alpha divide occurs. An effects mask is applied to the tool after the tool is processed.

Basic Node Setup

The Alpha Divide node is placed before any color correction is done to an image with a premultiplied Alpha channel. Here the MediaIn node is assumed to have a premultiplied Alpha channel. The Alpha Divide node is inserted, and then color correction nodes operate on the "straight" Alpha. An Alpha Multiply node is placed at the end of the chain to premultiply the Alpha channel again. If only a single color correction node is used, then the Pre-Divide/Post-Multiply checkbox on the Options tab can be used in place of the Alpha Divide/Alpha Multiple nodes.



An Alpha Divide node is inserted before color correcting an image with premultiplied alpha.

Inspector

This node has no controls.

Alpha Multiply [AMI]



The Alpha Multiply node

Alpha Multiply Node Introduction

As the name gives away, the Alpha Multiply's sole purpose is to multiply an image's color channels by its alpha channel. When you color correct an image that contains a premultiplied alpha channel, first apply an Alpha Divide node before any color correction node to create a non-premultiplied image. Then you can perform the color correction. After the color correction, add an Alpha Multiply node to return the image to its premultiplied state.

Inputs

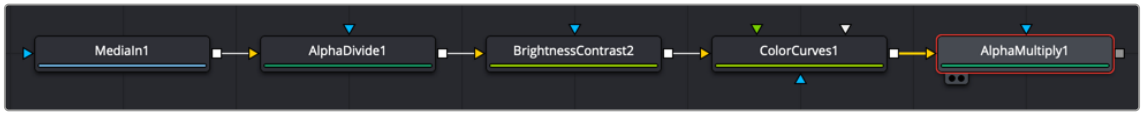
The Alpha Multiply node includes two inputs in the Node Editor.

Input: The orange input accepts a 2D image with a "straight" or non-premultiplied alpha.

Effect Mask: The optional blue input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input limits the pixels where the Alpha multiply occurs. An effects mask is applied to the tool after the tool is processed.

Basic Node Setup

The Alpha Multiply node is placed after color correction is done to an image with a premultiplied Alpha channel. Here the MediaIn node is assumed to have a premultiplied Alpha channel. The Alpha Divide node is inserted, and then color correction nodes operate on the "straight" Alpha. An Alpha Multiply node is placed at the end of the chain to premultiply the Alpha channel again. If only a single color correction node is used, then the Pre-Divide/Post-Multiply checkbox on the Options tab can be used in place of the Alpha Divide/Alpha Multiple nodes.



An Alpha Multiply node is inserted after color correcting an image with premultiplied alpha.

Inspector

This node has no controls.

Chroma Keyer [CKy]



The Chroma Keyer node

Chroma Keyer Node Introduction

The Chroma Keyer node creates an alpha channel (matte) for an image by removing selected colors from the scene. Unlike the Delta Keyer or Primatte, which use specific optimizations for keying from blue and green colors, the Chroma Keyer works equally well with any color.

NOTE: When working with blue- or green-screen shots, it is best to use the Delta Keyer or Primatte node, rather than the more general purpose Chroma Keyer node.

Inputs

The Chroma Keyer node includes four inputs in the Node Editor.

Input: The orange input accepts a 2D image that contains the color you want to be keyed for transparency.

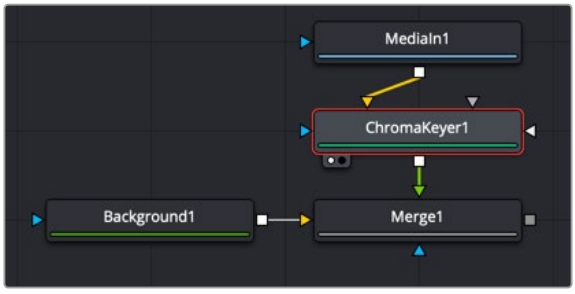
Garbage Matte: The gray garbage matte input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input causes areas of the image that fall within the matte to be made transparent. The garbage matte is applied directly to the alpha channel of the image.

Solid Matte: The white solid matte input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input causes areas of the image that fall within the matte to be fully opaque.

Effect Mask: The optional blue input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input limits the pixels where the alpha multiply occurs. An effects mask is applied to the tool after the tool is processed.

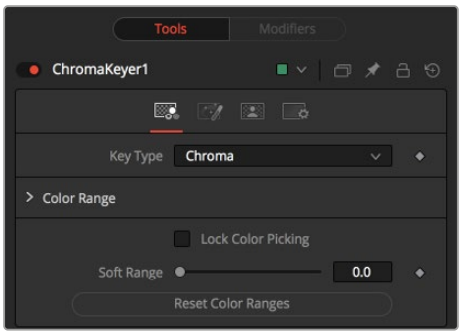
Basic Node Setup

The Chroma Keyer node takes in a 2D image like the MediaIn node below and removes a color that you identify by dragging over it in the viewer. The result is that the selected color is replaced with transparency, allowing you to composite the image as the foreground in a Merge node.



A Chomra Keyer node creating transparency on the MediaIn node.

Inspector



The Chroma Key tab

Chroma Key Tab

The Chroma Key tab is used to make the initial selection of color for keying.

Key Type

The Key Type menu determines the selection method used for the matte creation.

- **Chroma:** The Chroma method creates a matte based on the RGB values of the selected color range.
- **Color:** The Color method creates a matte based on the hue of the selected color range.

Color Range

Colors are made transparent by selecting the Chroma Keyer node in the node tree, and then dragging a selection around the colors in the viewer. The range controls update automatically to represent the current color selection. You can tweak the range sliders slightly, although most often selecting colors in the displays is all that is required.

Lock Color Picking

When this checkbox is activated, selecting colors from the viewer is disabled to prevent accidental addition to the range. It is a good idea to activate this checkbox once you make the color selection for the matte. All other controls in the node remain editable.

Soft Range

This control softens the selected color range, adding additional colors into the matte.

Reset Color Ranges

Clicking this button resets the Chroma Keyer's range controls, discarding all color selections. All other sliders and controls maintain their values.



The Chroma Key Image tab

Image Tab

The Image tab primarily handles removing spill color on the foreground subject. Color spill occurs when light containing the color you are removing is reflected onto the foreground subject.

Spill Color

This menu selects the color used as the base for all spill suppression techniques.

Spill Suppression

This slider sets the amount of spill suppression applied to the foreground subject.

When this slider is set to 0, no spill suppression is applied.

Spill Method

This menu selects the strength of the algorithm used to apply spill suppression to the image.

- **None:** None is selected when no spill suppression is required.
- **Rare:** This removes very little of the spill color and is the lightest of all methods.
- **Medium:** This works best for green screens.
- **Well Done:** This works best for blue screens.
- **Burnt:** This works best for blue. Use this mode only for very troublesome shots. Most likely you will have to add strong color correction after the key to get, for example, your skin tones back.

Fringe Gamma

This control is used to adjust the brightness of the fringe or halo that surrounds the keyed image.

Fringe Size

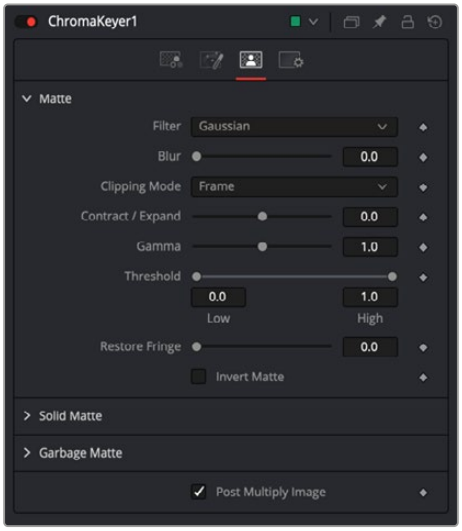
This expands and contracts the size of the fringe or halo surrounding the keyed image.

Fringe Shape

Fringe Shape forces the fringe toward the external edge of the image or toward the inner edge of the fringe. Its effect is most noticeable while the Fringe Size slider's value is large.

Cyan/Red, Magenta/Green, and Yellow/Blue

Use these three controls to color correct the fringe of the image. This is useful for correcting semitransparent pixels that still contain color from the original background to match the new background.



The Chroma Key Matte tab

Matte Tab

The Matte tab refines the softness, density, and overall fit of the resulting matte.

Filter

This control selects the filtering algorithm used when applying blur to the matte.

- **Box:** This is the fastest method but at reduced quality. Box is best suited for minimal amounts of blur.
- **Bartlett:** Otherwise known as a Pyramid filter, Bartlett makes a good compromise between speed and quality.
- **Multi-box:** When selecting this filter, the Num Passes slider appears and lets you control the quality. At 1 and 2 passes, results are identical to Box and Bartlett, respectively. At 4 passes and above, results are usually as good as Gaussian, in less time and with no edge "ringing."
- **Gaussian:** The Gaussian filter uses a true Gaussian approximation and gives excellent results, but it is a little slower than the other filters. In some cases, it can produce an extremely slight edge "ringing" on floating-point pixels.

Blur

Matte Blur blurs the edge of the matte based on the Filter menu setting. A value of zero results in a sharp, cutout-like hard edge. The higher the value, the more blur applied to the matte.

Clipping Mode

This option determines how edges are handled when performing domain of definition rendering. This is profoundly important when blurring the matte, which may require samples from portions of the image outside the current domain.

- **Frame:** The default option is Frame, which automatically sets the node's domain of definition to use the full frame of the image, effectively ignoring the current domain of definition. If the upstream DoD is smaller than the frame, the remaining area in the frame will be treated as black/transparent.
- **Domain:** Setting this option to Domain will respect the upstream domain of definition when applying the node's effect. This can have adverse clipping effects in situations where the node employs a large filter.

- **None:** Setting this option to None will not perform any source image clipping at all. This means that any data required to process the node's effect that would normally be outside the upstream DoD will be treated as black/transparent.

Contract/Expand

This slider shrinks or grows the semitransparent areas of the matte. Values above 0.0 expand the matte, while values below 0.0 contract it.

This control is usually used in conjunction with the Matte Blur to take the hard edge of a matte and reduce fringing. Since this control affects only semitransparent areas, it will have no effect on a matte's hard edge.

Gamma

Matte Gamma raises or lowers the values of the matte in the semitransparent areas. Higher values cause the gray areas to become more opaque, and lower values cause the gray areas to become more transparent. Completely black or white regions of the matte remain unaffected.

Since this control affects only semitransparent areas, it will have no effect on a matte's hard edge.

Threshold

This range slider sets the lower threshold using the handle on the left and sets the upper threshold using the handle on the right.

Any value below the lower threshold setting becomes black or transparent in the matte.

Any value above the upper threshold setting becomes white or opaque in the matte. All values within the range maintain their relative transparency values.

This control is often used to reject salt and pepper noise in the matte.

Restore Fringe

This restores the edge of the matte around the keyed subject. Often when keying, the edge of the subject where you have hair is clipped out. Restore Fringe brings back that edge while keeping the matte solid.

Invert Matte

When this checkbox is selected, the alpha channel created by the keyer is inverted, causing all transparent areas to be opaque and all opaque areas to be transparent.

Solid Matte

Solid Mattes are mask nodes or images connected to the solid matte input on the node. The solid matte is applied directly to the alpha channel of the image. Generally, solid mattes are used to hold out keying in areas you want to remain opaque, such as someone with blue eyes against a blue screen.

Enabling Invert will invert the solid matte, before it is combined with the source alpha.

Garbage Matte

Garbage mattes are mask nodes or images connected to the garbage matte input on the node. The garbage matte is applied directly to the alpha channel of the image. Generally, garbage mattes are used to remove unwanted elements that cannot be keyed, such as microphones and booms. They are also used to fill in areas that contain the color being keyed but that you wish to maintain.

Garbage mattes of different modes cannot be mixed within a single tool. A Matte Control node is often used after a Keyer node to add a garbage matte with the opposite effect of the matte applied to the keyer.

Enabling Invert will invert the garbage matte, before it is combined with the source alpha.

Post-Multiply Image

Select this option to cause the keyer to multiply the color channels of the image against the alpha channel it creates for the image. This option is usually enabled and is on by default.

Deselect this checkbox and the image can no longer be considered premultiplied for purposes of merging it with other images. Use the Subtractive option of the Merge node instead of the Additive option.

> For more information on these Merge node settings, *see Chapter 35, "Composite Nodes,"* in the Fusion Reference Manual.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other matte nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Clean Plate



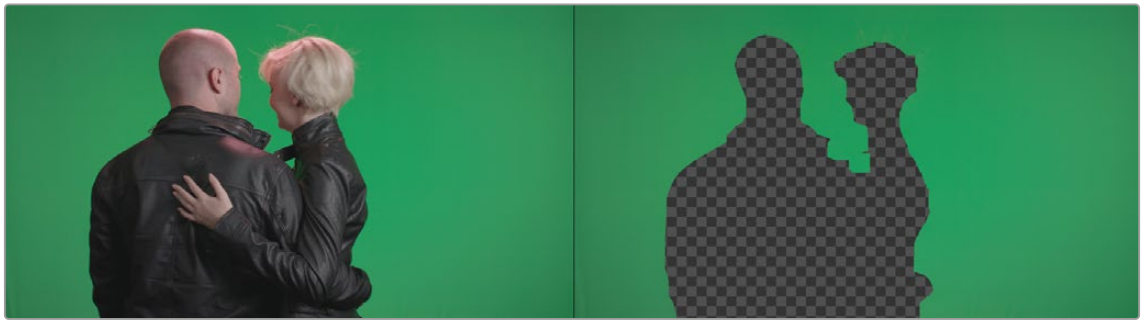
The Clean Plate node

Clean Plate Node Introduction

The Clean Plate tool is a pre-keying node used to generate an image of the green or blue color screen to smooth out the lighting differences. The output of the Clean Plate is later connected to the Clean Plate input on the Delta Keyer so it can key fine detail without choking or clipping the matte.

How to Create a Clean Plate

Creating a clean plate is the opposite of creating a key. When keying, you try to remove the green or blue color. When creating a clean plate, you try to keep as much of the blue- or green-screen as possible. By box selecting areas of the screen color in the viewer, you end up with an image of the green/blue screen. A transparent cutout represents everything that is not part of the blue or green screen.



Once you have the selection, the Erode control expands the pre-matte, removing any small pixels of non-green/blue screen around the edges. Then, growing the pre-matte fills in the holes until you have a solid blue or green image.

Inputs

The Clean Plate node includes three inputs in the Node Editor.

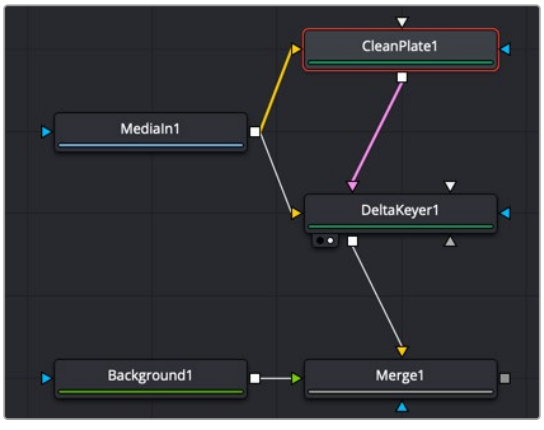
Input: The orange input accepts a 2D image that contains the green or blue screen.

Garbage Matte: The white garbage matte input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input causes areas of the image that fall within the matte to be excluded from the clean plate. For a clean plate, garbage mattes should contain areas that are not part of the blue or green screen.

Effect Mask: The optional blue input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input limits the pixels where the clean plate is generated. An effects mask is applied to the tool after the tool is processed.

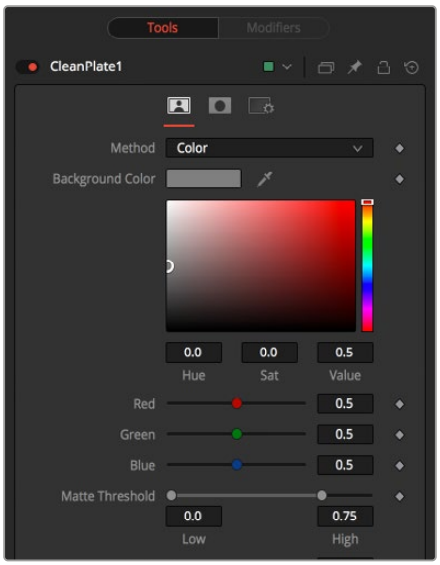
Basic Node Setup

The Clean Plate node and the Delta Keyer are two separate branches stemming from the main image you want to key. The green-screen or blue-screen clip is breached to connect to both the orange image input on the Clean Plate and the orange image input on the Delta Keyer. The output of the clean plate is then connected to the magenta clean plate input on the Delta Keyer. The output of the Delta Keyer is then used as the foreground to a Merge.



A Clean Plate connected to a Delta Keyer

Inspector



The Clean Plate tab

Plate Tab

The Plate tab contains the primary tools for creating a clean plate. Using this tab, you drag over the areas in the viewer, and then use the Erode and Grow Edges sliders to create the clean plate.

Method

The Method menu selects the type of color selection you use when sampling colors in the viewer.

- **Color:** Color uses a difference method to separate the background color. This works well on screen colors that are even.
- **Ranges:** Ranges uses a chroma range method to separate the background color. This is a better option for shadowed screen or screens that have different colors.

Matte Threshold

This range slider sets the lower threshold using the handle on the left and sets the upper threshold using the handle on the right.

Any value below the lower threshold becomes black or transparent in the matte.

Any value above the upper threshold becomes white or opaque in the matte. All values within the range maintain their relative transparency values. This control is often used to reject salt and pepper noise in the matte.

Erode

The Erode slider decreases the size of the screen area. It is used to eat away at small non-screen color pixels that may interfere with creating a smooth green- or blue-screen clean plate.

Crop

Crop trims in from the edges of the image.

Grow Edges

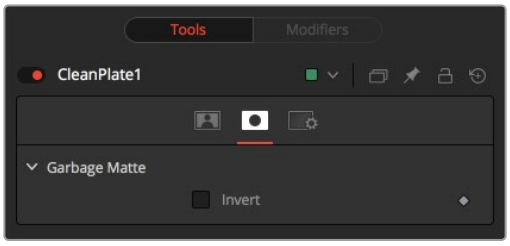
The Grow Edges slider expands the color of the edges to fill in holes until fully green or blue screen is created.

Fill

The Fill checkbox fills in remaining holes with color from the surrounding screen color.

Time Mode

- **Sequence:** Generates a new clean plate every frame.
- **Hold Frame:** Holds the clean plate at a single frame.



The Clean Plate Mask tab

Mask Tab

The Mask tab is used to invert the mask connected to the garbage mask input on the node. The garbage mask can be applied to clear areas before growing edges or filling remaining holes.

Invert

Invert uses the transparent parts of the mask to clear the image.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other matte nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Cryptomatte [Cry] (Studio Version Only)



The Cryptomatte node

Cryptomatte Node Introduction

The Cryptomatte tool in Fusion allows you to read embedded mattes rendered into an EXR. Instead of relying on manually-assigned masks or color keying, Cryptomatte reads encoded IDs and isolates them in a way that allows for easy selection. This tool allows you to create mattes for specific objects or materials in a 3D-rendered scene.

Inputs

The Cryptomatte node has only one input in the Node Editor.

Input: The orange input accepts a 2D EXR image with embedded mattes.

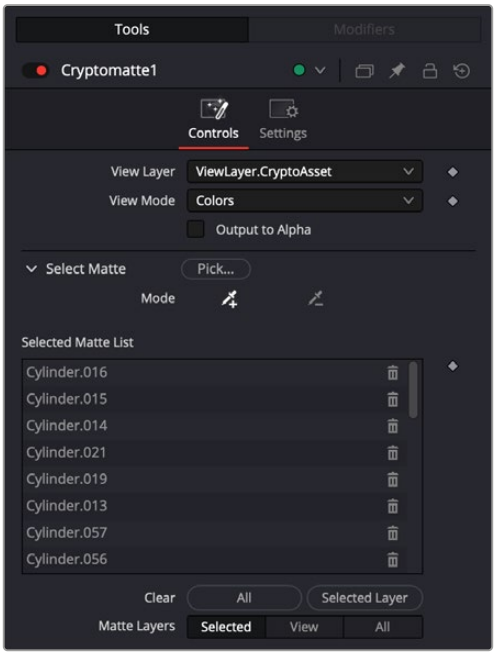
Basic Node Setup

The Cryptomatte node takes an EXR image with embedded mattes, then lets you select which mattes to pass through to the rest of the comp.



Cryptomatte taking in an EXR, selecting its mattes, and then passing them on to the Color Corrector and Bitmap tools

Inspector



The Cryptomatte Controls

These parameters let you control the matte selection of an EXR file.

- **View Layer:** Change the layer and matte type to pick (Object, Material etc). This affects the type of matte you can choose with the Pick control.
- **View Mode:** Changes the output of the Cryptomatte tool.

  Colors: Shows the layer IDs as separate colors.

  Edges: Shows the source EXR with outlined mattes overlayed the image.

  Beauty: Shows the source EXR.

  Matte: Outputs a black and white image of the selected matte.

- **Select Matte:** Select mattes from the viewer or from a list.

  Pick: Click to open a pick window with a matte list selection. Click and drag or use Cmd/Control-clicks to select multiple mattes from the list.

  Mode: Click and drag the control to get an eye dropper cursor to select or remove mattes from the viewer.

- **Selected Matte List:** The Selected Matte List displays the selected mattes in a list view. The selection is highlighted in yellow on the viewer.

  Clear: Click the bin icon to remove the selected mattes.

  Clear All: Click this button to remove all mattes from the Selected Matte List.

  Clear Selected Layer: Click to remove all mattes from the layer.

  Matte Layers: Toggle between only the Selected layers, View, and All.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other matte nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Delta Keyer



The Delta Keyer node

Delta Keyer Node Introduction

The Delta Keyer is an advanced color difference keyer, with many features and controls for tuning the matte and separating the subject from a blue or green screen.

It contains several keying systems: the Key tab is the master difference keyer and Pre-Matte is a built-in clean plate to smooth out screen color. Tuning, Fringe, and Matte finish the keying process. The tabs are laid out in the rough order you tend to use them.

Inputs

The Delta Keyer node includes five inputs in the Node Editor.

Input: The orange input accepts a 2D image that contains the color you want to be keyed for transparency.

Garbage Matte: The gray garbage matte input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input causes areas of the image that fall within the matte to be made transparent. The garbage matte is applied directly to the alpha channel of the image.

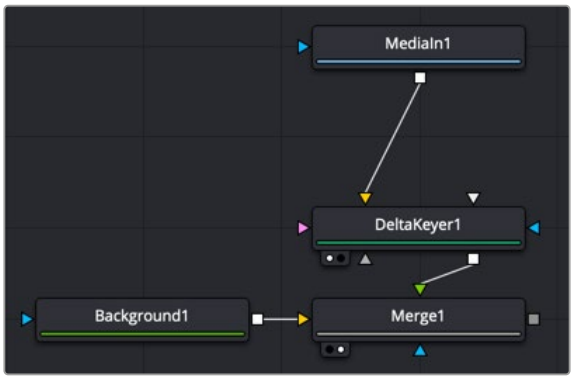
Solid Matte: The white solid matte input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input causes areas of the image that fall within the matte to be fully opaque.

Clean Plate: Accepts the resulting image from the Clean Plate node.

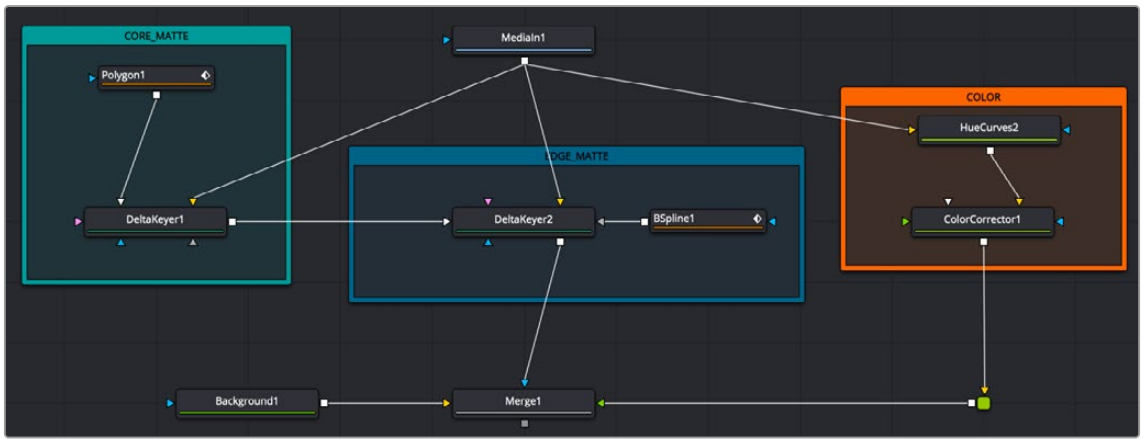
Effect Mask: The optional blue input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input limits the pixels where the keying occurs. An effects mask is applied to the tool after the tool is processed.

Basic Node Setup

The first operation needed to use the Delta Keyer is to open the Delta Keyer in the Inspector, then click and drag the background color eyedropper tool into the Viewer to select the color you want to key out. This color will then turn gray in the Viewer. A single Delta Keyer rarely gets perfect results because most green- or blue-screen shots have problems that the keyer is not made to handle. Keyers often need the help of garbage mattes or solid mattes created with a Polygon or B-Spline node. Shots can also require more than just one keyer to achieve perfect results. That is shown in one of the examples below, where one Delta Keyer is used to create a hard inner white matte, while a second Delta Keyer is used to capture the soft edges. The core matte below is then connected into the solid matte input on the second Delta Keyer. Color correction, including spill suppression, is best handled as a separate branch in your node tree. Separating your matte from your color correction using underlays makes it clear what branch of nodes is performing what operation.

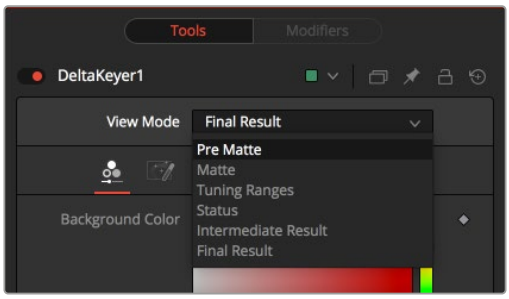


Starting point for a Delta Keyer setup



An expanded key setup with two Delta Keyers, a Polygon and B-Spline nodes for mattes, and a branch for color correction

Inspector



The Delta Keyer View Mode menu

Key Tab View Mode

At the top of the Inspector is the View Mode menu. The default selection shows the final result. You can change the view to see various intermediate stages of the keying process.

- **Pre Matte:** Displays the output of the Pre Matte key.
- **Matte:** Displays the alpha of the key before being combined with solid and garbage masks. When displaying the matte, set the viewer to show the alpha channel.
- **Tuning Ranges:** Displays a false color for Shadow, Midtone, and Highlight ranges of the image. Shadows are in the red channel, Midtones are in the green channel, and Highlights are in the blue channel.
- **Status:** Displays information to indicate areas that are solid, transparent, or in between. It also displays areas that have been affected by matte adjustments, such as thresholding or erode/dilate and areas affected by the solid mask.
- **Intermediate Result:** The untouched source image color channels combined with the final matte. Optionally, combine this output with additional Delta Keyer nodes.
- **Final Result:** The final keyed image with spill suppression, ready to merge onto a scene.



The Delta Keyer Key tab

Key Tab

The Key tab is where most keying begins. It is used to select the screen color.

Background Color

This is the color of the blue or green screen, sometimes called the screen color. To create the key with the Delta Keyer, use the background color Eyedropper to select the screen color from the image.

Pre-Blur

Applies a blur before generating the alpha. This can help with certain types of noise, edge enhancements, and artifacts in the source image.

Gain

Gain increases the influence of the screen color, causing those areas to become more transparent.

Balance

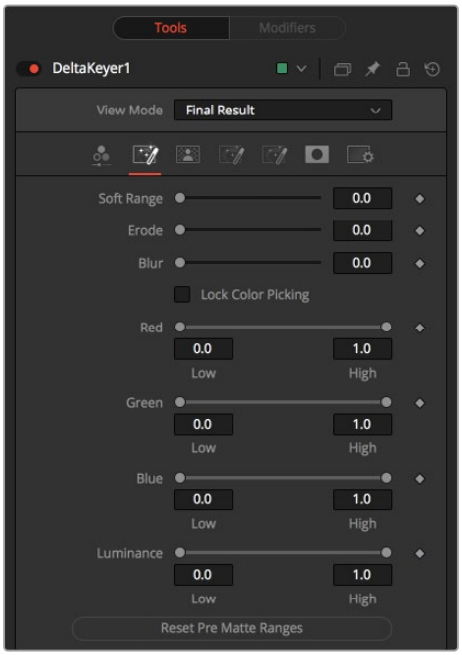
A color difference keyer, like the Delta Keyer, compares the differences between the dominant channel determined by the selected background color and the other two channels. Adjusting balance determines the proportions of the other two channels. A value of 0 uses the minimum of the other two channels, where a value of 1 uses the maximum. A value of 0.5 uses half of each.

Lock Alpha/Spill Removal Color Balance Reference

Unlocking this allows you to use different color references when generating the alpha and when determining how much of the background color to subtract from the image. When enabled, spill and color are combined.

Color Balance Reference

The Color balance reference controls adjust for lighting or white balance that might be reducing background color purity and saturation. A correction is applied based on the reference of a neutral-colored object when generating the key and determining the amount of background color subtraction, without altering the background color that is subtracted.



The Delta Keyer Pre Matte tab

Pre Matte Tab

The Pre Matte tab is in some ways a garbage matte keyer that occurs before the delta keyer process. It is used to even out the screen color before a more delicate key is done. After selecting the initial screen color, you can switch to the Pre Matte tab and view the PreMatte Out from view mode menu at the top of the Inspector. The Pre Matte works on ranges of color, so you start by dragging a bounding box around areas where lighting is different from the initial screen color pick. Use the Soft range to expand, and Erode to cut away the edges.

Once you have a more even screen selection, you can move to the Matte tab.

Soft Range

The Soft Range extends the range of selected color and rolloff of the screen color.

Erode

Erode contracts the edge of the pre matte, so the edge detail does not clip.

Blur

This softens the edges of the pre matte.

Pre Matte Range

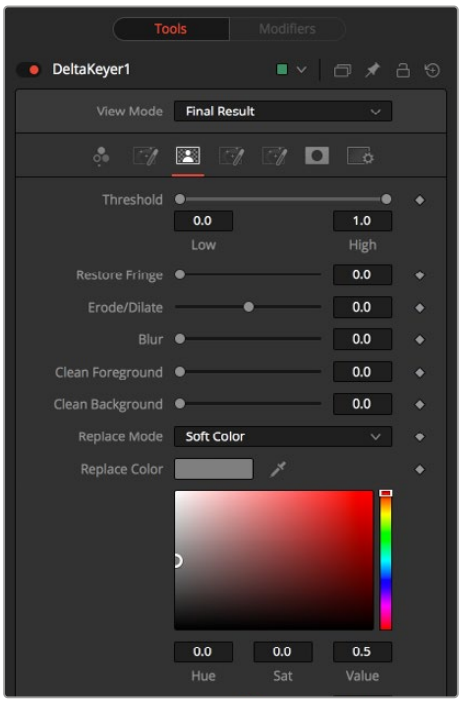
These controls update automatically to represent the current color selection. The controls tweak the selection slightly, although selecting colors in the viewer is all that is required.

Lock Color Picking

When this checkbox is activated, it prevents the accidental growth of the selected range by selecting more colors from the view. It is a good idea to activate this checkbox after selecting the color for the matte. All other controls in the node remain editable.

Reset Pre Matte Ranges

This button discards all color selection by resetting the ranges but maintains all other slider and control values.



The Delta Keyer Matte tab

Matte Tab

The Matte tab refines the alpha of the key, combined with any solid and garbage masks connected to the node. When using the matte tab, set the viewer to display the alpha channel of the Delta Keyer's final output.

Threshold

This range slider sets the lower threshold using the handle on the left and sets the upper threshold using the handle on the right.

Any value below the lower threshold setting becomes black or transparent in the matte.

Any value above the upper threshold setting becomes white or opaque in the matte. All values within the range maintain their relative transparency values.

Restore Fringe

This restores the edge of the matte around the keyed subject. Often when keying, clipping occurs around the edge of the subject where there is hair. Restore Fringe brings back that edge while keeping the matte solid.

Erode/Dilate

Expands or contracts the matte.

Blur

Softens the matte.

Clean Foreground

Fills slightly transparent (light gray) areas of the matte.

Clean Background

Clips the bottom dark range of the matte.

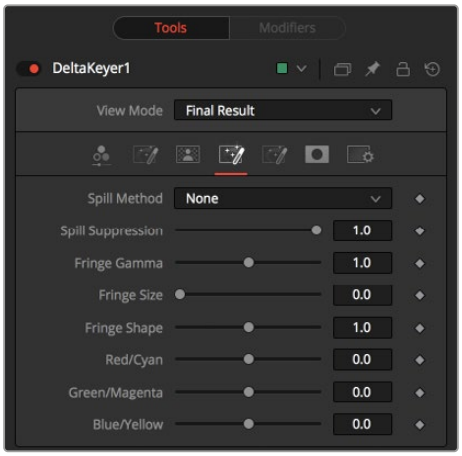
Replace Mode

Determines how matte adjustments restore color to the image.

- **None:** No color replacement. Matte processing does not affect the color.
- **Source:** The color from the original image.
- **Hard Color:** A solid color.
- **Soft Color:** A solid color weighted by how much background color was originally removed.

Replace Color

The color used with the Hard Color and Soft Color replace modes.



The Delta Keyer Fringe tab

Fringe Tab

The Fringe tab handles the majority of spill suppression in the Delta Keyer. Spill suppression is a form of color correction that attempts to remove the screen color from the fringe of the matte.

Spill is the transmission of the screen color through the semitransparent areas of the alpha channel. In the case of blue- or green-screen keying, this usually causes the color of the background to become apparent in the edges of the foreground subject.

Spill Method

This selects the strength of the algorithm used to apply spill suppression to the image.

- **None:** None is selected when no spill suppression is required.
- **Rare:** This removes very little of the spill color and is the lightest of all methods.
- **Medium:** This works best for green screens.
- **Well Done:** This works best for blue screens.
- **Burnt:** This works best for blue screens. Use this mode only for very troublesome shots.

Spill Suppression

When this slider is set to 0, no spill suppression is applied to the image. Increasing the slider increases the strength of the spill method.

Fringe Gamma

This control can be used to adjust the brightness of the fringe or halo that surrounds the keyed image.

Fringe Size

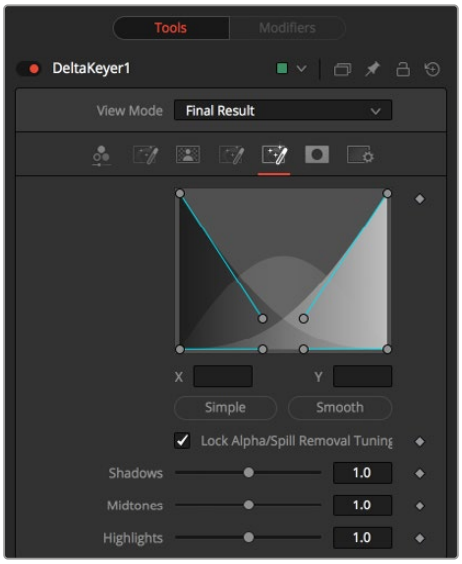
This expands and contracts the size of the fringe or halo surrounding the keyed image.

Fringe Shape

Fringe Shape presses the fringe toward the external edge of the image or pulls it toward the inner edge of the fringe. Its effect is most noticeable while the Fringe Size value is large.

Cyan/Red, Magenta/Green, and Yellow/Blue

Use these three controls to color correct the fringe of the image. This is useful for correcting semitransparent pixels that still contain color from the original background to match the new background.



The Delta Keyer Tuning tab

Tuning Tab

The Tuning tab is an advanced tab that allows you to determine the size of the shadow, midtone, and highlight ranges. By modifying the ranges, you can select the strength of the matte and spill suppression based on tonal values.

Range Controls

The range controls define how much color is considered to fall into the shadows, midtones, and highlights areas of the image. The spline controls allow for easy adjusting of the tonal ranges of each Shadow and Highlight tonal map.

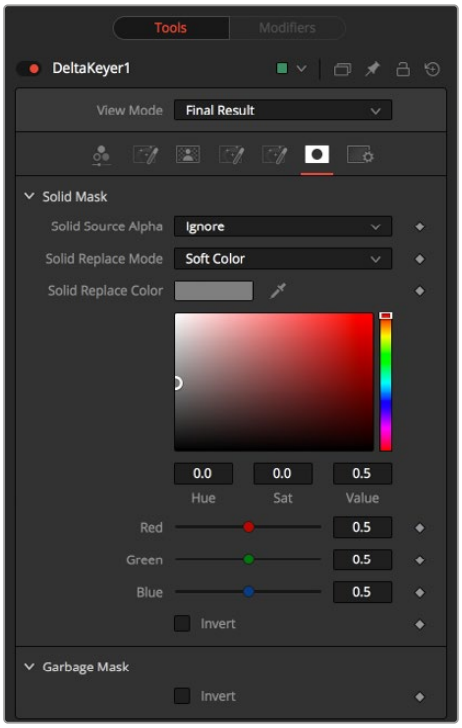
Simple/Smooth

The Simple button sets the range to be linear. The Smooth button sets a smooth tonal gradient for the ranges.

Lock Alpha/Spill Removal Tuning

When this checkbox is disabled, separate tuning controls are used when generating the alpha and when determining how much of the background color to subtract from the image.

- **Shadows:** Adjusts the strength of the key in darker areas of the background.
- **Midtones:** Adjusts the strength of the key in midtone areas of the background.
- **Highlights:** Adjusts the strength of the key in brighter areas of the background.



The Delta Keyer Mask tab

Mask Tab

The Mask tab determines how the solid and garbage mattes are applied to the key.

Solid Source Alpha

Used to combine the existing alpha from the source image with the solid mask.

- **Ignore:** Does not combine the alpha from the source image.
- **Add:** Solid areas of the source image alpha are made solid in the solid mask.
- **Subtract:** Transparent areas of the source image alpha are made transparent in the solid mask.

Solid Replace Mode

This determines how the solid mask restores color to the image.

- **None:** No color replacement. The solid mask does not affect the color.
- **Source:** The color from the original image.
- **Hard Color:** A solid color.
- **Soft Color:** A solid color weighted by how much background color was originally removed.

Solid Replace Color

The color used with the Hard Color and Soft Color replace modes.

- **Invert:** Inverts the solid mask, before it is combined with the source alpha.

Garbage Mask

- **Invert:** Normally, solid areas of a garbage mask remove the image. When inverted, the transparent areas of the mask remove the image.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other matte nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Depth Map [DMp]



The Depth Map node

Depth Map Node Introduction

Depth Map creates an Alpha channel based on the perceived distance of objects in your clip. By being able to isolate a specific depth region, the opportunities to manipulate the resulting image are greatly expanded. For example, combined with a Blur node you could simulate a depth of field effect, but confined only to the distant background of your shot. You could enhance a foreground object, like a person, increasing contrast, saturation, and sharpness, similar to using a qualifier. You can use a Depth Map to fix video issues, for example fixing a color temperature problem by isolating a far window that was shot in daylight, causing background actors to be tinted blue, while leaving the foreground subject that was shot with studio lighting unaffected.

The resulting Depth Map Alpha channel is visualized as a black and white image, with white being the area that is affected by the resulting changes and black areas remaining unchanged.

Inputs

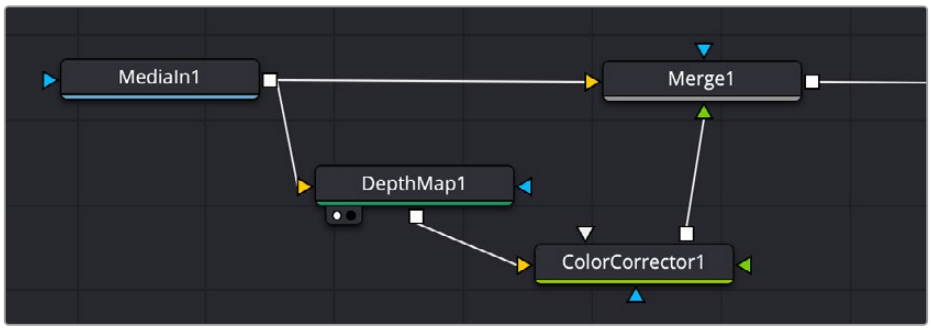
The Depth Map node includes two inputs in the Node Editor.

Input: The yellow input accepts a 2D image that contains the shot you wish to analyze for depth.

Effect Mask: The optional blue input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input limits the pixels where the difference matte occurs. An effects mask is applied to the tool after the tool is processed.

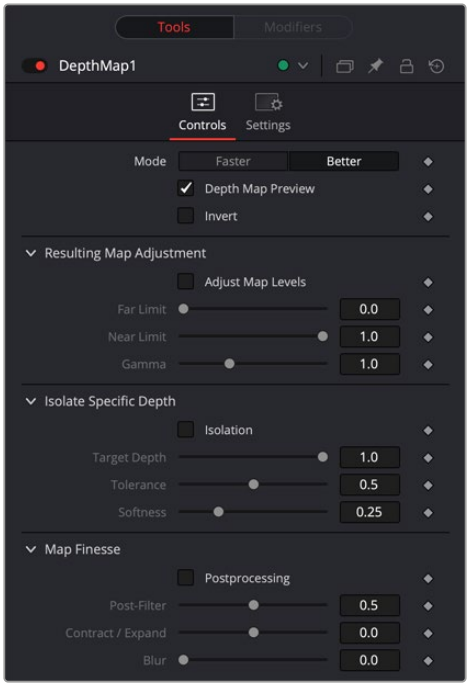
Basic Node Setup

The Depth Map is designed to take an input from an image source and create an Alpha channel based on the perceived distance from the camera. Then it will feed that Alpha to downstream nodes. In the example below, the Depth Map is used to isolate the background of the image and apply a color grade to just the background, while leaving the foreground actors unaffected as they move through the frame.



A sample Depth Map node structure for isolating subjects based on distance from the camera

Inspector



The Depth Map Controls tab

Controls Tab

The Controls tab in the Depth Map contains all the parameters for adjusting the quality of the resulting matte.

- **Mode:** Depth Map is a very computationally intensive effect. The quality setting allows a Faster mode to speed up the responsiveness to adjustments, while the default Better mode gives the best results and should be turned on when the adjustments are finished.

- **Depth Map Preview:** By default this box is checked and shows you the current Depth Map for making adjustments. When this check box is disabled, the resulting Alpha can then be used for grading on other nodes.
- **Invert:** Checking this box reverses the Depth Map, switching its transparent and opaque regions.

Resulting Map Adjustment

These controls let you determine how the Depth Map's contrast is adjusted.

- **Adjust Map Levels:** When deselected (default), all scaling is turned off, allowing you to adjust the full range of the Depth Map. When enabled, this option clips the Depth Map's levels to 0 and 1. This functions as a preview of what will happen to the Depth Map when used as an Alpha channel where the values are always clipped to 0 and 1. Checking this box also activates the tools below.
- **Far Limit:** This control adjusts the black levels of the Depth Map.
- **Near Limit:** This control adjusts the white levels of the Depth Map.
- **Gamma:** This control adjusts the intermediate depth values to be brighter or dimmer compared to the fixed black and white levels.

Isolate Specific Depth

These controls allow you to sweep backwards and forwards in the scene by depth, allowing you to isolate a specific depth range to adjust.

- **Isolation:** This checkbox turns the depth isolation tools on or off.
- **Target Depth:** This controls the specific depth you want to isolate. 1 is fully in the foreground, while 0 is fully in the background.
- **Tolerance:** Sets the range to either side of the Target Depth to include in the Depth Map.
- **Softness:** This sets a subtle ramp-in and ramp-out of the selected range, giving it a more organic selection.

Map Finesse

These controls modify the resulting Depth Map's Alpha channel for use in grading.

- **Post Processing:** This control turns the Map Finesse tools on or off.
- **Post-Filter:** This control blends the map to the smooth areas and edges of the image. It is used to prevent later grading effects from visibly varying within the region.
- **Contract/Expand:** This control dilates or erodes the overall shape at the edges; useful for fine tuning the boundary between the affected and unaffected regions of the map.
- **Blur:** This control softens the boundary of the map, allowing it to blend more smoothly into the resulting image.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Matte nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Difference Keyer [DfK]



The Difference Keyer node

Difference Keyer Node Introduction

Difference keying is a process that produces a matte based on the differences between two images. A difference key uses two input images: one containing the subject with the background and another containing the background without the subject.

Although the process sounds reasonable at first glance, subtle variations in the camera position from shot to shot usually make it difficult to get clean results. Think of the futile attempt of trying to key smoke in front of a brick wall and using a clean plate of the brick wall as your difference input. Part of the wall's structure is always visible in this keying method. Instead, a Difference Keyer is often used to produce a rough matte that is combined with other nodes to produce a more detailed matte.

Inputs

The Difference Keyer node includes four inputs in the Node Editor.

Background: The orange background input accepts a 2D image that contains just the set without your subject.

Foreground: The green foreground input accepts a 2D image that contains the shot with your subject in the frame.

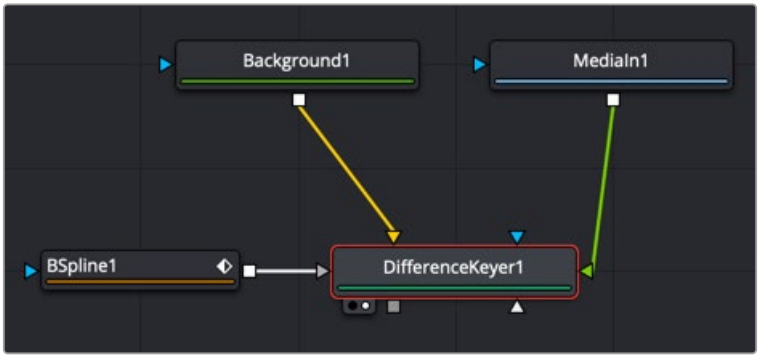
Garbage Matte: The gray garbage matte input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input causes areas of the image that fall within the matte to be made transparent.

Solid Matte: The white solid matte input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input causes areas of the image that fall within the matte to be fully opaque.

Effect Mask: The optional blue input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input limits the pixels where the difference matte occurs. An effects mask is applied to the tool after the tool is processed.

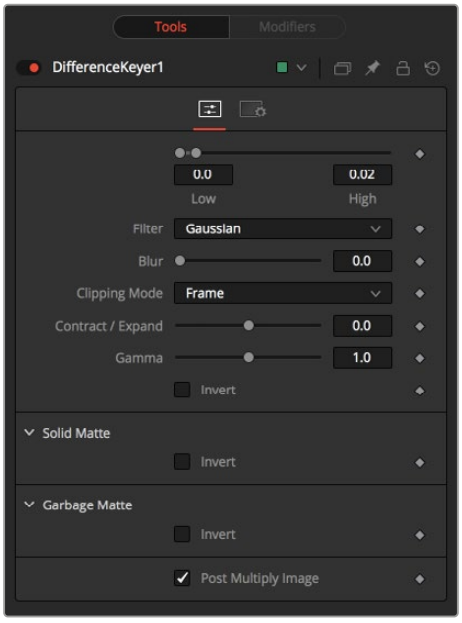
Basic Node Setup

When you do not have content shot on a blue or green screen, the Difference Keyer can be one node in a chain of many used to extract an object from the background. The example below has the MediaIn1 as the main subject and a clean background shot without the subject (Background). A B-Spline is used to limit the area the Difference Keyer must deal with for extraction. The result is a matte that can be used to help but not solve the key.



A Difference Keyer with two inputs: one of the subject on a background and the other of just the background

Inspector



The Difference Keyer Controls tab

Controls Tab

The Controls tab in the Difference Keyer contains all the parameters for adjusting the quality of the matte.

Threshold

This range slider sets the lower threshold using the handle on the left and sets the upper threshold using the handle on the right. Adjusting them defines a range of difference values between the images to create a matte.

A difference below the lower threshold setting becomes black or transparent in the matte.

Any difference above the upper threshold setting becomes white or opaque in the matte.

The difference values in the range in between create a grayscale matte.

Filter

This control selects the filtering algorithm used when applying a blur to the matte.

- **Box:** This is the fastest method but at reduced quality. Box is best suited for minimal amounts of blur.
- **Bartlett:** Otherwise known as a Pyramid filter, Bartlett makes a good compromise between speed and quality.
- **Multi-box:** When selecting this filter, the Num Passes slider appears and lets you control the quality. At 1 and 2 passes, results are identical to Box and Bartlett, respectively. At 4 passes and above, results are usually as good as Gaussian, in less time and with no edge "ringing."
- **Gaussian:** The Gaussian filter uses a true Gaussian approximation and gives excellent results, but it is a little slower than the other filters. In some cases, it can produce an extremely slight edge "ringing" on floating-point pixels.

Blur

This blurs the edge of the matte using the method selected in the Filter menu. A value of zero results in a sharp, cutout-like hard edge. The higher the value, the more blur.

Clipping Mode

This option determines how edges are handled when performing domain of definition rendering. This is profoundly important when blurring the matte, which may require samples from portions of the image outside the current domain.

- **Frame:** The default option is Frame, which automatically sets the node's domain of definition to use the full frame of the image, effectively ignoring the current domain of definition. If the upstream DoD is smaller than the frame, the remaining area in the frame is treated as black/transparent.
- **Domain:** Setting this option to Domain respects the upstream domain of definition when applying the node's effect. This can have adverse clipping effects in situations where the node employs a large filter.
- **None:** Setting this option to None does not perform any source image clipping at all. This means that any data required to process the node's effect that would usually be outside the upstream DoD is treated as black/transparent.

Contract/Expand

This slider shrinks or grows the semitransparent areas of the matte. Values above 0.0 expand the matte, while values below 0.0 contract it.

This control is usually used in conjunction with the blur to take the hard edge of a matte and reduce fringing. Since this control affects only semitransparent areas, it has no effect on a matte's hard edge.

Gamma

Matte Gamma raises or lowers the values of the matte in the semitransparent areas. Higher values cause the gray areas to be more opaque, and lower values cause the gray areas to be more transparent. Wholly black or white regions of the matte remain unaffected.

Invert

Selecting this checkbox inverts the matte, causing all transparent areas to be opaque and all opaque areas to be transparent.

Solid Matte

Solid Mattes are mask nodes or images connected to the solid matte input on the node. The solid matte is applied directly to the alpha channel of the image. Generally, solid mattes are used to hold out keying in areas you want to remain opaque, such as someone with blue eyes against a blue screen.

Enabling Invert, inverts the solid matte before it is combined with the source alpha.

Garbage Matte

Garbage mattes are mask nodes or images connected to the garbage matte input on the node. The garbage matte is applied directly to the alpha channel of the image. Generally, garbage mattes are used to remove unwanted elements that cannot be keyed, such as microphones and booms. They are also used to fill in areas that contain the color being keyed but that you wish to maintain.

Garbage mattes of different modes cannot be mixed within a single tool. A Matte Control node is often used after a Keyer node to add a garbage matte with the opposite effect of the matte applied to the keyer.

Enabling Invert inverts the garbage matte before it is combined with the source alpha.

Post-Multiply Image

Select this option to cause the keyer to multiply the color channels of the image against the alpha channel it creates for the image. This option is usually enabled and is on by default.

Deselect this checkbox and the image can no longer be considered premultiplied for purposes of merging it with other images. Use the Subtractive option of the Merge node instead of the Additive option.

> For more information on these Merge node settings, *see Chapter 35, "Composite Nodes,"* in the Fusion Reference Manual.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Matte nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Luma Keyer [LKy]



The Luma Keyer node

Luma Keyer Node Introduction

The Luma Keyer node uses the overall luminance of an image to create an Alpha channel. The label of this node may seem misleading since it allows pulling mattes from almost any channel. In some respects, it is more accurate to call this node an all-purpose channel keyer, but its primary purpose is for extracting alpha channels based on luminance

Inputs

The Luma Keyer node includes four inputs in the Node Editor.

Input: The orange input accepts a 2D image that contains the luminance values you want to be keyed for transparency.

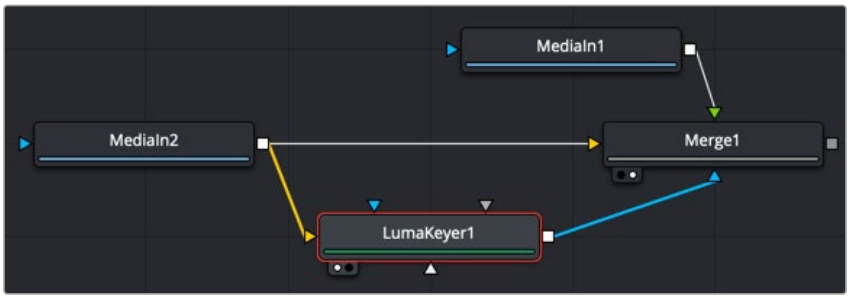
Garbage Matte: The gray garbage matte input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input causes areas of the image that fall within the matte to be made transparent. The garbage matte is applied directly to the alpha channel of the image.

Solid Matte: The white solid matte input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input causes areas of the image that fall within the matte to be fully opaque.

Effect Mask: The optional blue input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input limits the pixels where the luminance key occurs. An effects mask is applied to the tool after the tool is processed.

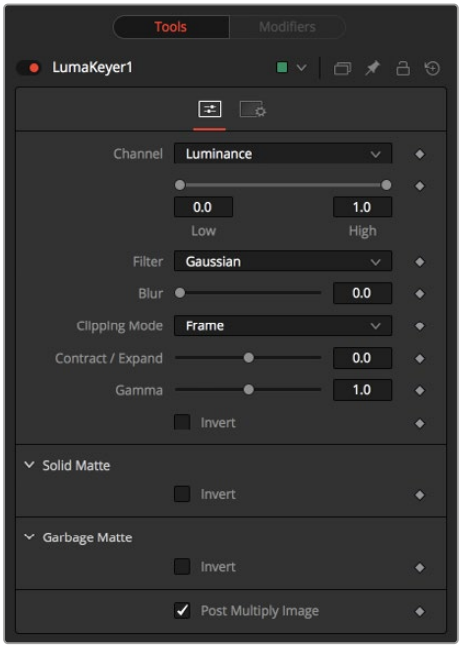
Basic Node Setup

The Luma Keyer takes an input from an image with significant luminance difference to extract a key. You can then use the output of the Luma Keyer into any mask input.



A Luma Keyer output connecting into an effect mask on a Merge node

Inspector



The Luma Keyer Controls tab

Controls Tab

The Controls tab in the Luma Keyer contains all the parameters for adjusting the quality of the matte.

Channel

This menu selects the color channel used for creating the matte. Select from the Red, Green, Blue, Alpha, Hue, Luminance, Saturation, and Depth (Z-buffer) channels.

Threshold

This range slider sets the lower threshold using the handle on the left and sets the upper threshold using the handle on the right. Adjusting them defines a range of luminance values to create a matte.

A value below the lower threshold setting becomes black or transparent in the matte.

Any value above the upper threshold setting becomes white or opaque in the matte.

The values within the range create a grayscale matte.

Filter

This control selects the filtering algorithm used when applying a blur to the matte.

- **Box:** This is the fastest method but at reduced quality. Box is best suited for minimal amounts of blur.
- **Bartlett:** Otherwise known as a Pyramid filter, Bartlett makes a good compromise between speed and quality.
- **Multi-box:** When selecting this filter, the Num Passes slider appears and lets you control the quality. At 1 and 2 passes, results are identical to Box and Bartlett, respectively. At 4 passes and above, results are usually as good as Gaussian, in less time and with no edge "ringing."
- **Gaussian:** The Gaussian filter uses a true Gaussian approximation and gives excellent results, but it is a little slower than the other filters. In some cases, it can produce an extremely slight edge "ringing" on floating-point pixels.

Blur

This blurs the edge of the matte using the method selected in the Filter menu. A value of zero results in a sharp, cutout-like hard edge. The higher the value, the more blur.

Clipping Mode

This option determines how edges are handled when performing domain of definition rendering. This is profoundly important when blurring the matte, which may require samples from portions of the image outside the current domain.

- **Frame:** The default option is Frame, which automatically sets the node's domain of definition to use the full frame of the image, effectively ignoring the current domain of definition. If the upstream DoD is smaller than the frame, the remaining area in the frame is treated as black/transparent.
- **Domain:** Setting this option to Domain respects the upstream domain of definition when applying the node's effect. This can have adverse clipping effects in situations where the node employs a large filter.
- **None:** Setting this option to None does not perform any source image clipping at all. This means that any data required to process the node's effect that is usually outside the upstream DoD is treated as black/transparent.

Contract/Expand

This slider shrinks or grows the semitransparent areas of the matte. Values above 0.0 expand the matte, while values below 0.0 contract it.

This control is usually used in conjunction with the blur to take the hard edge of a matte and reduce fringing. Since this control affects only semitransparent areas, it has no effect on a matte's hard edge.

Gamma

Matte Gamma raises or lowers the values of the matte in the semitransparent areas. Higher values cause the gray areas to be more opaque, and lower values cause the gray areas to be more transparent. Wholly black or white regions of the matte remain unaffected.

Invert

Selecting this checkbox inverts the matte, causing all transparent areas to be opaque and all opaque areas to be transparent.

Solid Matte

Solid mattes are mask nodes or images connected to the solid matte input on the node. The solid matte is applied directly to the alpha channel of the image. Generally, solid mattes are used to hold out keying in areas you want to remain opaque, such as someone with blue eyes against a blue screen.

Enabling Invert inverts the solid matte before it is combined with the source alpha.

Garbage Matte

Garbage mattes are mask nodes or images connected to the garbage matte input on the node. The garbage matte is applied directly to the alpha channel of the image. Generally, garbage mattes are used to remove unwanted elements that cannot be keyed, such as microphones and booms. They are also used to fill in areas that contain the color being keyed but that you wish to maintain.

Garbage mattes of different modes cannot be mixed within a single tool. A Matte Control node is often used after a Keyer node to add a garbage matte with the opposite effect of the matte applied to the keyer.

Enabling Invert inverts the garbage matte before it is combined with the source alpha.

Post-Multiply Image

Select this option to cause the keyer to multiply the color channels of the image against the alpha channel it creates for the image. This option is usually enabled and is on by default.

Deselect this checkbox and the image can no longer be considered premultiplied for purposes of merging it with other images. Use the Subtractive option of the Merge node instead of the Additive option.

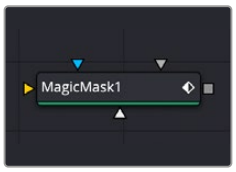
> For more information on these Merge node settings, *see Chapter 35, "Composite Nodes,"* in the Fusion Reference Manual.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Matte nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Magic Mask [MagM]



The Magic Mask node

Magic Mask Node Introduction

The Magic Mask palette uses the DaVinci Neural Engine to automagically create a mask to isolate an object in the frame, guided by user-applied paint strokes to identify the subject for isolation. Masks can be generated for either an entire object or just a part of it. This allows you to create complicated masks of complex shapes that may be difficult to isolate with other tools. For example, using the Magic Mask to modify only the wood grain section of this guitar.



(Left) Multiple strokes isolating the wood grain of the guitar, while ignoring the musician's arm; (Right) The finished shot, with a warmer wood grain

NOTE: The Magic Mask tool in Fusion is based on the Object mode of the Magic Mask tool in the Color page and is designed to be used for any subject that is not a person. Subjects like cars, pets, and food are all good candidates for Object mode. If you are trying to isolate a human being, using Magic Mask's Person mode in the Color page will likely give you better results.

Drawing Strokes in the Viewer

The Magic Mask tool works by drawing lines directly into the viewer. Once a stroke is drawn, the viewer will automatically show the mask you've created. You can draw multiple strokes to create a mask in an additive fashion and use negative strokes to remove areas from the mask. It is best to choose a frame to draw on where the object is completely in view and unobstructed.

- **To draw a positive stroke to add an object to a mask:** Load the Magic Mask node in a viewer, and left-click and drag the stroke across the object you want to choose. Alternatively you can change the Stroke Mode to Add in the Inspector and left-click and drag to draw a new positive stroke. Positive strokes are colored blue.

- **To draw a negative stroke to remove an object from a mask:** Load the Magic Mask node in a viewer, and Option-left-click and drag the stroke across the object you want to remove. Alternatively you can change the Stroke Mode to Subtract in the Inspector and left-click and drag to draw a new negative stroke. Negative strokes are colored red.
- **To delete a stroke (or group of strokes) to remove them from the mask:** Shift-drag around a stroke or group of strokes to select them (green), and press the delete key to remove them. Alternatively you can change the Stroke Mode to Select in the Inspector and left-click and drag a selection window to select one or more strokes. Clicking the delete button removes them.

Drawing positive strokes will select areas of similar contrast and color, allowing you to link complex shapes together. Generally you will need more strokes to actively define a complex object, due to the greater variety of the shapes involved. Stroke position is usually more important than stroke length.

Drawing negative strokes removes areas from the object that you don't want to isolate. This can be something simple, like removing the wheels of a car from a mask, or more complicated like removing specific books from a mask of a bookshelf. Stroke position is usually more important than stroke length.



(Left) Multiple strokes isolating the car's body, while removing the wheels and cabin; (Right) The finished shot, fed through a color corrector node to change the color of just the car's body

Inputs

The Magic Mask node includes four inputs in the Node Editor.

Input: The orange input accepts a 2D image that contains the luminance values you want to be keyed for transparency.

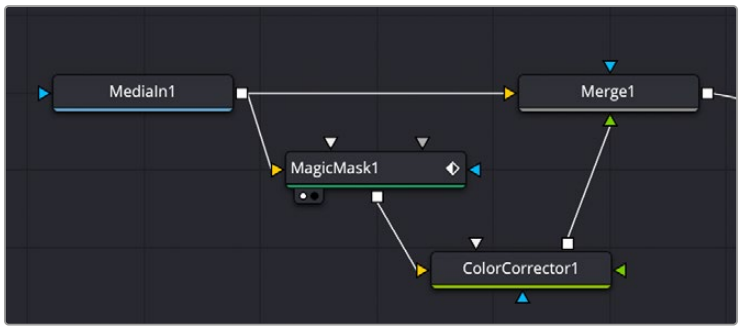
Garbage Matte: The gray garbage matte input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input causes areas of the image that fall within the matte to be made transparent. The garbage matte is applied directly to the Alpha channel of the image.

Solid Matte: The white solid matte input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input causes areas of the image that fall within the matte to be fully opaque.

Effect Mask: The optional blue input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input limits the pixels where the luminance key occurs. An effects mask is applied to the tool after the tool is processed.

Basic Node Setup

The Magic Mask takes RGB inputs from other nodes and lets you isolate specific objects in the frame to use as a mask. Tracking the mask as it changes over time is built into the tool. You can then use the output of the Magic Mask into any mask input.



A Magic Mask node, feeding a matte to a Color Corrector connected to the foreground of a Merge node

Inspector



The Magic Mask Tracking tab

Tracking Tab

You can think of each stroke you've drawn as a persistent eyedropper that samples the image that overlaps it. The mask that results from all strokes' collective analysis of the image is generated live over each frame of the clip. This means that if the camera or the subject moves, you need to motion track or otherwise adjust the position of each stroke to follow along with the motion, so the subject continues to be correctly identified. You also need to make sure that each stroke is able to analyze what it's supposed to, and turn off strokes that can't for whatever reason.

- **Tracking Controls:** These buttons control the tracking direction, from left to right:

  Track Reverse: Continuously tracks from the current frame all the way to the beginning of the clip.

  Track Reverse One Frame: Tracks one frame backwards and stops. Useful if you're tracking frame-by-frame to watch the progress of a particularly complicated bit of motion. If something goes wrong, you can back up to the last frame where the stroke was able to properly track the subject, and drag the stroke to a better location using the pointer to make it follow the subject properly. If necessary, you can go a frame at a time, dragging the stroke to a better position every time it fails to follow the feature you're using it to isolate.

Stop Tracking: Stops tracking, in cases where there's a problem with the track, and you want to make a change.

Track Forward Then Reverse: Tracks from the current frame all the way to the end of the clip, then returns to the original tracking point and tracks backwards to the beginning of the clip.

Track Forward One Frame: Tracks one frame forward and stops. Useful if you're tracking frame-by-frame to watch the progress of a particularly complicated bit of motion. If necessary, you can go a frame at a time, dragging the stroke to a better position every time it fails to follow the feature you're using it to isolate.

Track Forward: Continuously tracks from the current frame all the way to the end of the clip.

- **Go To Frame:** These buttons snap the playhead to the selected frame, from left to right:

  First Frame of Tracked Area: Moves the playhead to the first tracked frame of a range of tracked frames in preparation for tracking backwards if there are untracked frames at the beginning of the clip.

  Reference Frame: Moves the playhead to the frame on which you initially drew the strokes.

  Last Frame of Tracked Area: Moves the playhead to the last tracked frame of a range of tracked frames in preparation for tracking forwards if there are untracked frames at the end of the clip.

- **Stroke Mode:** These buttons let you change and modify the strokes drawn in the Viewer

  Add: Lets you add an additional stroke (blue) that determines what in the frame will be included in the mask.

  Subtract: Lets you add an additional stroke (red) that determines what in the frame will be excluded from the mask.

  Select: Lets you draw a selection rectangle around single or multiple strokes to select them (green).

  Delete: Deletes all selected strokes (green).

- **Clear Strokes:** These buttons remove all strokes from the selected range.

  Current: Removes strokes from the frame under the playhead.

  Range: Removes strokes from the selected render range.

  All: Removes all strokes from the entire clip.

- **Disk Cache:** These buttons allow you to control what's stored in the disk cache.

  Regenerate All: Rebuilds the disk cache.

  Clear: Deletes the frames in the current cache.

- **Reference Time:** The Reference Time determines the frame where the initial strokes are drawn. It is also the time from which tracking begins. The reference frame cannot be changed once it has been set without destroying all pre-existing tracking information.
- **Processed Frames:** Displays the range of frames that have already been tracked. Start is the earliest frame tracked and end is the last frame tracked. These fields are not user editable.
- **Mode:** Two options let you choose a tradeoff between quality and performance. Faster lets you generate a lower quality mask more quickly that's suitable for garbage matting. Better generates a higher quality mask with more detail that's more processor-intensive.

Post-Multiply Image

Select this option to cause the keyer to multiply the color channels of the image against the Alpha channel it creates for the image. This option is usually enabled and is on by default.

Deselect this checkbox and the image can no longer be considered premultiplied for purposes of merging it with other images. Use the Subtractive option of the Merge node instead of the Additive option.

For more information on these Merge node settings, *see Chapter 35, "Composite Nodes,"* in the Fusion Reference Manual.

Matte Tab

The Matte tab refines the alpha of the key, combined with any solid and garbage masks connected to the node. When using the Matte tab, set the viewer to display the Alpha channel of Magic Mask's final output.

- **Filter:** Selects the Filter that is used when blurring the matte.

  Box Blur: This option applies a Box Blur effect to the whole image. This method is faster than the Gaussian blur but produces a lower-quality result.

  Bartlett: Bartlett applies a more subtle, anti-aliased blur filter.

  Multi-Box: Multi-Box uses a box filter layered in multiple passes to approximate a Gaussian shape. With a moderate number of passes (e.g., four), a high-quality blur can be obtained, often faster than the Gaussian filter and without any ringing.

  Gaussian: Gaussian applies a smooth, symmetrical blur filter, using a sophisticated constant-time Gaussian approximation algorithm.

  Fast Gaussian: Gaussian applies a smooth, symmetrical blur filter, using a sophisticated constant-time Gaussian approximation algorithm. This mode is the default filter method.

- **Blur:** This blurs the edge of the matte using a standard, constant speed Gaussian blur. A value of zero results in a sharp, cutout-like hard edge. The higher the value, the more blur is applied to the matte.
- **Erode/Dilate:** This shrinks or grows the matte. Eroding the matte reveals more of the foreground input, while dilating the matte reveals more of the background input. Values above 0.0 dilate the matte and values below 0.0 erode it.
- **Gamma:** This raises or lowers the values of the matte in the semitransparent areas. Higher values cause the gray areas to become more opaque, and lower values cause the gray areas to become more transparent. Completely black or white regions of the matte remain unaffected.
- **Threshold:** Any value below the lower threshold becomes black or transparent in the matte. Any value above the upper threshold becomes white or opaque in the matte. All values within the range maintain their relative transparency values.
- **Restore Fringe:** This restores the edge of the matte around the keyed subject. Often when keying, the edge of the subject where you have hair is clipped out. Restore Fringe brings back that edge while keeping the matte solid.
- **Invert Matte:** When this checkbox is selected, the Alpha channel of the image is inverted, causing all transparent areas to be opaque and all opaque areas to be transparent.

Solid Matte

Solid mattes are mask nodes or images connected to the solid matte input on the node. The solid matte is applied directly to the Alpha channel of the image. Generally, solid mattes are used to hold out keying in areas you want to remain opaque, such as someone with blue eyes against a blue screen.

Enabling Invert inverts the solid matte before it is combined with the source alpha.

Garbage Matte

Garbage mattes are mask nodes or images connected to the garbage matte input on the node. The garbage matte is applied directly to the Alpha channel of the image. Generally, garbage mattes are used to remove unwanted elements that cannot be keyed, such as microphones and booms. They are also used to fill in areas that contain the color being keyed but that you wish to maintain.

Garbage mattes of different modes cannot be mixed within a single tool. A Matte Control node is often used after a Keyer node to add a garbage matte with the opposite effect of the matte applied to the keyer.

Enabling Invert inverts the garbage matte before it is combined with the source alpha.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Matte nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Matte Control [Mat]



The Matte Control node

Matte Control Node Introduction

Keyer nodes are used to create an alpha channel on an image that does not have one. The Matte Control node is used to combine and manipulate the alpha channels embedded in images as well as masks created by masking tools.

Typically, you add this node to copy a color channel or alpha channel from the foreground input to the background input, or to combine alpha channels from the two images.

Inputs

The Matte Control node includes four inputs in the Node Editor.

Background: The orange background input accepts a 2D image that receives the foreground image alpha channel (or some other channel you want to copy to the background).

Foreground: The green foreground input accepts a 2D image that contains an alpha channel (or some other channel) you want to be applied to the background image.

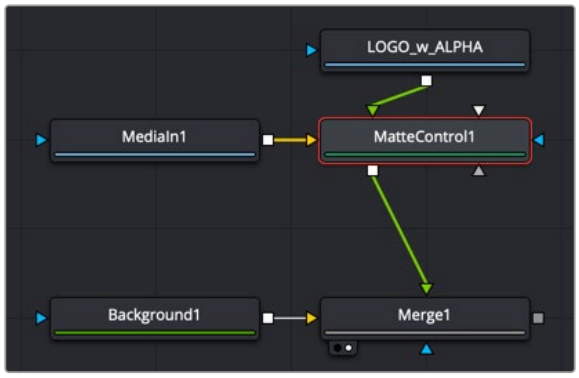
Garbage Matte: The gray garbage matte input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input causes areas of the foreground/background combination that fall within the matte to be made transparent.

Solid Matte: The white solid matte input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input causes areas of the foreground/background combination that fall within the matte to be fully opaque.

Effect Mask: The optional blue input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input limits the pixels where the matte control occurs. An effects mask is applied to the tool after the tool is processed.

Basic Node Setup

Below, the Matte Control node is set up to copy the foreground (green) input's alpha channel into the background (orange) input. The output of the Matte Control is then an image with an alpha channel used as the foreground composite in the Merge node.



A Matte Control embedded an alpha from the foreground input to the background input

Inspector



The Matte Control Matte tab

Matte Tab

The Matte tab combines and modifies alpha or color channels from an image in the foreground input with the background image.

Combine

Use this menu to select which operation is applied. The default is set to None for no operation.

- **None:** This causes the foreground image to be ignored.
- **Combine Red:** This combines the foreground red channel to the background alpha channel.
- **Combine Green:** This combines the foreground green channel to the background alpha channel.
- **Combine Blue:** This combines the foreground blue channel with the background alpha channel.
- **Combine Alpha:** This combines the foreground alpha channel with the background alpha channel.
- **Solid:** This causes the background alpha channel to become completely opaque.
- **Clear:** This causes the background alpha channel to become completely transparent.

Combine Operation

Use this menu to select the method used to combine the foreground channel with the background.

- **Copy:** This copies the foreground source over the background alpha, overwriting any existing alpha in the background.
- **Add:** This adds the foreground source to the background alpha.
- **Subtract:** This subtracts the foreground source from the background alpha.
- **Inverse Subtract:** This subtracts the background alpha from the foreground source.
- **Maximum:** This compares the foreground source and the background alpha and takes the value from the pixel with the highest value.
- **Minimum:** This compares the foreground source and the background alpha and takes the value from the pixel with the lowest value.
- **And:** This performs a logical AND on the two values.
- **Or:** This performs a logical OR on the values.
- **Merge Over:** This merges the foreground source channel over the background alpha channel.
- **Merge Under:** This merges the foreground source channel under the background alpha channel.

Filter

Selects the Filter that is used when blurring the matte.

- **Box Blur:** This option applies a Box Blur effect to the whole image. This method is faster than the Gaussian blur but produces a lower-quality result.
- **Bartlett:** Bartlett applies a more subtle, anti-aliased blur filter.
- **Multi-Box:** Multi-Box uses a box filter layered in multiple passes to approximate a Gaussian shape. With a moderate number of passes (e.g., four), a high-quality blur can be obtained, often faster than the Gaussian filter and without any ringing.
- **Gaussian:** Gaussian applies a smooth, symmetrical blur filter, using a sophisticated constant-time Gaussian approximation algorithm. In extreme cases, this algorithm may exhibit ringing; see below for a discussion of this. This mode is the default filter method.

Blur

This blurs the edge of the matte using a standard, constant speed Gaussian blur. A value of zero results in a sharp, cutout-like hard edge. The higher the value, the more blur is applied to the matte.

Clipping Mode

This option determines how edges are handled when performing domain of definition rendering. This is profoundly important when blurring the matte, which may require samples from portions of the image outside the current domain.

- **Frame:** The default option is Frame, which automatically sets the node's domain of definition to use the full frame of the image, effectively ignoring the current domain of definition. If the upstream DoD is smaller than the frame, the remaining area in the frame is treated as black/transparent.
- **Domain:** Setting this option to Domain respects the upstream domain of definition when applying the node's effect. This can have adverse clipping effects in situations where the node employs a large filter.
- **None:** Setting this option to None does not perform any source image clipping at all. This means that any data required to process the node's effect that is usually outside the upstream DoD is treated as black/transparent.

Contract/Expand

This shrinks or grows the matte similar to an Erode Dilate node. Contracting the matte reveals more of the foreground input, while expanding the matte reveals more of the background input. Values above 0.0 expand the matte, and values below 0.0 contract it.

Gamma

This raises or lowers the values of the matte in the semitransparent areas. Higher values cause the gray areas to become more opaque, and lower values cause the gray areas to become more transparent. Completely black or white regions of the matte remain unaffected.

Threshold

Any value below the lower threshold becomes black or transparent in the matte. Any value above the upper threshold becomes white or opaque in the matte. All values within the range maintain their relative transparency values.

Restore Fringe

This restores the edge of the matte around the keyed subject. Often when keying, the edge of the subject where you have hair is clipped out. Restore Fringe brings back that edge while keeping the matte solid.

Invert Matte

When this checkbox is selected, the alpha channel of the image is inverted, causing all transparent areas to be opaque and all opaque areas to be transparent.

Solid Matte

Solid mattes are mask nodes or images connected to the solid matte input on the node. The solid matte is applied directly to the alpha channel of the image. Generally, solid mattes are used to hold out areas you want to remain opaque, such as someone with blue eyes against a blue screen.

Enabling Invert inverts the solid matte before it is combined with the source alpha.

Garbage Matte

Garbage mattes are mask nodes or images connected to the garbage matte input on the node. The garbage matte is applied directly to the alpha channel of the image. Generally, garbage mattes are used to remove unwanted elements that cannot be keyed, such as microphones and booms. They are also used to fill in areas that contain the color being keyed but that you wish to maintain.

Garbage mattes of different modes cannot be mixed within a single tool. A Matte Control node is often used after a Keyer node to add a garbage matte with the opposite effect of the matte applied to the keyer.

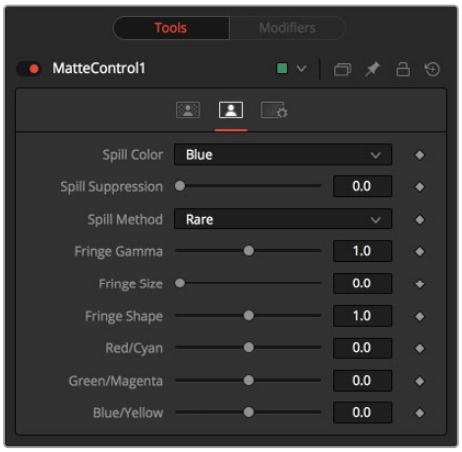
Enabling Invert inverts the garbage matte before it is combined with the source alpha.

Post-Multiply Image

Selecting this option multiplies the color channels of the image against the alpha channel it creates for the image. This option is usually enabled and is on by default.

Deselect this checkbox and the image can no longer be considered premultiplied for purposes of merging it with other images. Use the Subtractive option of the Merge node instead of the Additive option.

> For more information on these Merge node settings, *see Chapter 35, "Composite Nodes,"* in the Fusion Reference Manual.



The Matte Control Spill tab

Spill Tab

The Spill tab handles spill suppression in the Matte Control. Spill suppression is a form of color correction that attempts to remove the screen color from the fringe of the matte.

Spill is the transmission of the screen color through the semitransparent areas of the alpha channel. In the case of blue- or green-screen keying, this usually causes the color of the background to become apparent in the edges of the foreground subject.

Spill Color

This menu selects the color used as the base for all spill suppression techniques.

Spill Suppression

When this slider is set to 0, no spill suppression is applied to the image. Increasing the slider increases the strength of the spill method.

Spill Method

This selects the strength of the algorithm used to apply spill suppression to the image.

- **None:** None is selected when no spill suppression is required.
- **Rare:** This removes very little of the spill color and is the lightest of all methods.
- **Medium:** This works best for green screens.
- **Well Done:** This works best for blue screens.
- **Burnt:** This works best for blue screen. Use this mode only for very troublesome shots.

Fringe Gamma

This control can be used to adjust the brightness of the fringe or halo that surrounds the keyed image.

Fringe Size

This expands and contracts the size of the fringe or halo surrounding the keyed image.

Fringe Shape

Fringe Shape presses the fringe toward the external edge of the image or pulls it toward the inner edge of the fringe. Its effect is most noticeable while the Fringe Size value is large.

Cyan/Red, Magenta/Green, and Yellow/Blue

Use these three controls to color correct the fringe of the image.

This is useful for correcting semitransparent pixels that still contain color from the original background to match the new background.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Matte nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Primatte [Pri]



The Primatte node

NOTE: Primatte is only available in Fusion Studio.

Primatte Node Introduction

Primatte is an advanced keying tool for Fusion Studio. To use Primatte effectively, you must understand how it works. Using a series of selection buttons, Primatte assigns RGB pixels into one of the four specific zones.

- **Zone 1:** Complete background image.
- **Zone 2:** Foreground image with spill suppression and transparency.
- **Zone 3:** Foreground image with spill suppression only.
- **Zone 4:** Complete foreground image.

Depending on the type of blue- or green-screen content, you may find that the Delta Keyer or the Primatte keyer handles the specific keying task better. There is no one-solution-fits-all when it comes to keying, and in some cases, the combination of the two keyers may prove to be the best solution.

NOTE: Primatte is distributed and licensed by IMAGICA Corp. of America, Los Angeles, CA, USA. Primatte was developed by and is a trademark of IMAGICA Corp., Tokyo, Japan.

Inputs

The Primatte node includes six inputs in the Node Editor. Unlike every other tool in Fusion, the primary orange input is labeled as the Foreground input, since it accepts the green-screen or blue-screen image. The background input on the Primatte node is the green input; this is an optional input that allows Primatte to create the final merged composite.

Foreground Input: The orange input accepts a 2D image that contains blue or green screen.

Background Input: The green (optional) input accepts a 2D image layered as the background in the composite. If no image is connected, Primatte outputs the keyed foreground. Connecting an image to the background input activates Primatte's advanced edge blending options.

Replacement Image: The magenta (optional) input accepts a 2D image used as a source of Primatte's spill suppression color correction.

Garbage Matte: The gray garbage matte input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input causes areas of the image that fall within the matte to be made transparent. The garbage matte is applied directly to the alpha channel of the image.

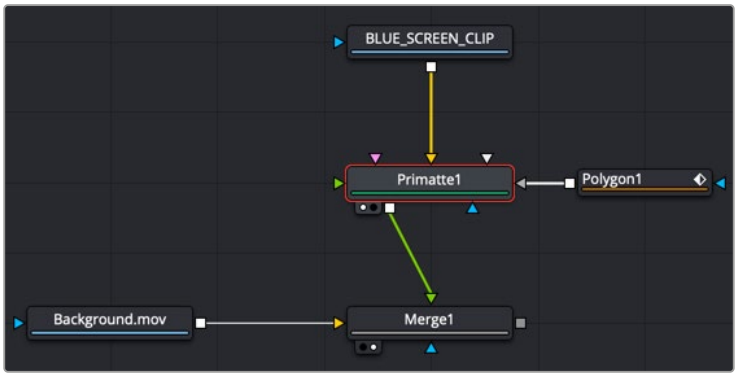
Solid Matte: The white solid matte input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input causes areas of the image that fall within the matte to be fully opaque.

Effect Mask: The optional blue input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input limits the pixels where the keying occurs. An effects mask is applied to the tool after the tool is processed.

NOTE: Connecting the background input without connecting the replacement image input uses the background image as the replacement image surf spill suppression.

Basic Node Setup

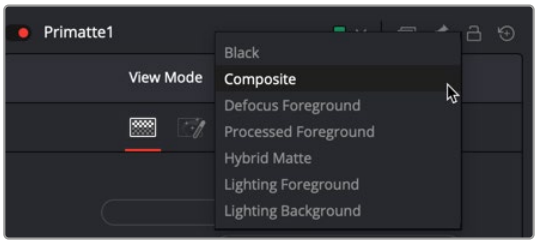
A single Primatte keyer can rarely get perfect results because most green- or blue-screen shots have problems the keyer is not made to handle. Keyers often need the help of garbage mattes or solid mattes created with a Polygon or B-Spline node. Shots can also require more than just one keyer to achieve perfect results. Below, the Primatte node has the blue-screen content connected to the orange input. Unlike other Fusion nodes, the foreground gets connected to the orange input. The result is an image with alpha that can then be connected into the foreground of a Merge node.



A Primatte node combined with a polygon matte as a garbage matte and connected to the foreground of a Merge

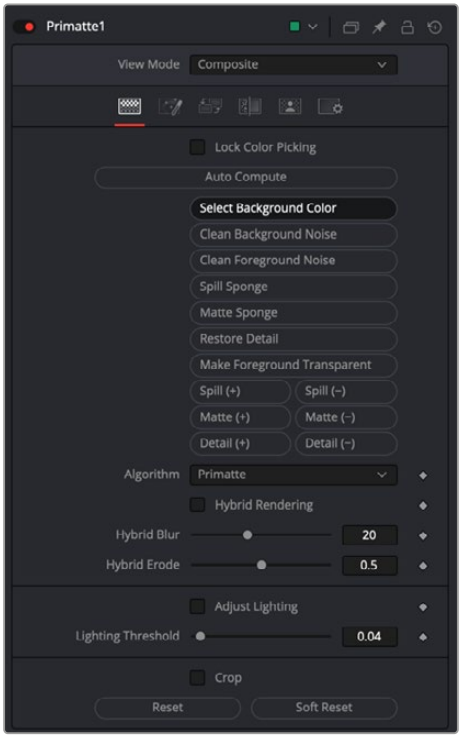
Primatte Tab View Mode

At the top of the Inspector is the View Mode menu. The default selection shows the final Composite result. You can change the view to see various intermediate stages of the keying process.



The Primatte tab View mode

- **Black:** Displays the foreground subject on a black or transparent background.
- **Composite:** The final keyed image with spill suppression, composited over the image connected to the green Background Input on the node.
- **Defocus Foreground:** Displays the output of the Pre Matte key.
- **Processed Foreground:** Displays the alpha of the key before being combined with solid and garbage masks. When displaying the matte, set the viewer to show the alpha channel.
- **Hybrid Matte:** Displays the matte generated when the Hybrid Rendering checkbox is enabled. Best viewed when adjusting the Hybrid Blur and Hybrid Erode sliders.
- **Lighting Foreground:** Displays the foreground subject over the optimized artificial backing screen that the Adjust Lighting mode creates.
- **Lighting Background:** Displays the optimized artificial backing screen that the Adjust Lighting mode creates.



The Primatte tab Inspector

Primatte Tab

The core functionality for Primatte is found in the Primatte tab. The basic workflow is based on selecting one of the operational mode buttons and then scrubbing over areas in the viewer.

Lock Color Picking

Activate this button once you finished adjusting your key to prevent making accidental changes in the viewer.

Auto Compute

The Auto Compute button is likely the first button pressed when starting to key your footage. Primatte automagically analyzes the original foreground image, determines the backing color, and sets it as the central backing color. Then, using that information, another analysis determines the foreground areas. A Clean FG Noise operation is performed using the newly determined foreground areas, and Primatte renders the composite.

NOTE: The Auto Compute button may make the next three buttons—Select Background Color, Clean Background Noise, and Clean Foreground Noise—unnecessary and make your keying operation much more straightforward. Clicking Auto Compute automatically senses the backing screen color, eliminates it, and even gets rid of some foreground and background noise. If you get good results, then jump ahead to the Spill Removal tools. If you don't get satisfactory results, continue from this point using the three buttons described below.

Select Background Color

Clicking the Select Background Color button allows you to select the screen color by scrubbing in the viewer. It uses the traditional Primatte method of taking the sampled backing screen color, projecting

a line in the opposite direction on the hue wheel, and generating artificial pixels that may represent the FG object. Then, using the artificially generated foreground pixels, it internally does the Clean FG Noise operation and creates the shape of the middle and outer polyhedrons. It then renders the composite using these generated polyhedrons. This does not automatically use the Adjust Lighting functionality, as it must be selected in a separate operation.

Clean Background Noise

Clicking this button helps to remove any white regions in the dark screen area ("noise"), or shades of the screen color that did not get picked up on the first sample. Once you click the button, scrub the mouse pointer over areas in the viewer to sample white-ish noise regions.

Clean Foreground Noise

If there are dark transparent regions in the middle of the mostly white opaque foreground object, click the Clean Foreground Noise button and scrub over the dark pixels in the foreground area until that area is as white as possible.

Spill Sponge

The Spill Sponge is the quickest method for removing color spill on your subject. Click the Spill Sponge button and scrub the mouse pointer over a screen color pixel, and the screen color disappears from the selected color region and is replaced by a complementary color, a selected color, or a color from a replacement image. These options are set in the Replace tab. Additionally, use the tools under the Fine Tuning tab or use the Spill(+) and Split(-) features to adjust the spill.

Matte Sponge

Sometimes in the Primatte operation, a 100% opaque, foreground area (all white) becomes slightly transparent (gray). To clean those transparent areas, click the Matte Sponge button and scrub over the transparent pixels. All the spill-suppression information remains intact.

Restore Detail

Clicking Restore Detail and scrubbing over background regions in the viewer turns completely transparent areas translucent. This operation is useful for restoring lost hair details, thin wisps of smoke, and the like.

Make Foreground Transparent

When this button is selected, the opaque foreground color region sampled in the viewer becomes slightly translucent. This operation is useful for the subtle tuning of foreground subjects, which are otherwise 100 percent covered with smoke or clouds. It can be used only one time on a particular color. For a more flexible way to thin out a color region, and to be able to take multiple samples, use the Matte(-) tool.

Spill(+)

Clicking the Spill(+) button returns the color spill to the sampled pixel color (and all colors like it) in the amount of one Primatte increment. This tool can be used to move the sampled color more in the direction of the color in the original foreground image. It can be used to nullify a Spill(-) step.

Spill(-)

Clicking the Spill(-) button removes from the sampled pixel color (and all colors like it) in the amount of one Primatte increment. If spill color remains, another click using this operational mode tool removes more of the color spill. Continue using this tool until all color spill has been removed from the sampled color region.

Matte(+)

Clicking the Matte(+) button makes the matte more opaque for the sampled pixel color (and all colors like it) in the amount of one Primatte increment. If the matte is still too translucent or thin, another click using this operational mode tool makes the sampled color region even more opaque. This can be used to thicken smoke or make a shadow darker to match shadows in the background imagery. It can only make these adjustments to the density of the color region on the original foreground image. It can be used to nullify a Matte(-) step.

Matte(-)

Clicking the Matte(+) button makes the matte more translucent for the sampled pixel color (and all colors like it) in the amount of one Primatte increment. If the matte is still too opaque, another click using this operational mode tool makes the sampled color region even more translucent. This can be used to thin out smoke or make a shadow thinner to match shadows in the background imagery.

Detail(+)

When this button is selected, the foreground detail becomes less visible for the sampled pixel color (and all colors like it) in the amount of one Primatte increment. If there is still too much detail, another click using this operational mode tool makes more of it disappear. This can be used to remove smoke or wisps of hair from the composite. Sample where detail is visible, and it disappears. This is for moving color regions into the 100% background region. It can be used to nullify a Detail(-) step.

Detail(-)

When this button is selected, foreground detail becomes more visible for the sampled pixel color (and all colors like it) in the amount of one Primatte increment. If detail is still missing, another click using this operational mode tool makes detail more visible. This can be used to restore lost smoke or wisps of hair. Sample where the smoke or hair just disappears and it returns to visibility. Use this for restoring color regions that were moved into the 100% background region. It may start to bring in background noise if shooting conditions were not ideal on the foreground image.

Algorithms

There are three keying algorithms available in the Primatte keyer:

- **Primatte:** The Primatte algorithm mode delivers the best results and supports both the Solid Color and the Complement Color spill suppression methods. This algorithm uses three multifaceted polyhedrons (as described later in this section) to separate the 3D RGB colorspace. It is also the default algorithm mode and, because it is computationally intensive, it may take the longest to render.
- **Primatte RT:** Primatte RT is the simplest algorithm and therefore the fastest. It uses only a single planar surface to separate the 3D RGB colorspace (as described later in this section) and, as a result, does not separate the foreground from the backing screen as carefully as the above Primatte algorithm. Another disadvantage of the Primatte RT algorithm is that it does not work well with less saturated backing screen colors, and it does not support the Complement Color spill suppression method.
- **Primatte RT+:** Primatte RT+ is in between the above two options. It uses a six planar surface color separation algorithm (as described later in this section) and delivers results in between the other two options in both quality and performance. Another disadvantage of the Primatte RT+ algorithm is that it does not work well with less saturated backing screen colors, and it does not support the Complement Color spill suppression method.

Hybrid Rendering

After sampling the backing screen color and producing acceptable edges around the foreground object, you sometimes find a transparent area within the foreground subject. This can occur when the foreground subject contains a color that is close to the backing screen color. Removing this transparency with the Clean FG Noise mode can cause the edge of the foreground subject to pick up a fringe that is close to the backing screen color. Removing the fringe is very difficult without sacrificing quality somewhere else on the image. The Hybrid Render mode internally creates two keying operations: Body and Edge. The optimized Edge operation gets the best edge around the foreground subject without any fringe effect. The Body operation deals with transparency within the foreground subject. The resultant matte is created by combining these two mattes, and then blurring and eroding the foreground subject in the Body matte and combining it with the edge matte.

To use Hybrid Rendering, start by keying the main foreground area using the Select Background Color mode (or any of the other Primatte backing screen detection methods). Activate the Hybrid Rendering checkbox. Lastly, select the Clean FG Noise button and scrub over the transparent area. The Hybrid Render mode performs the “Body/Edge” operation. The result is a final composite with perfect edges around the foreground subject with a solid foreground subject.

Hybrid Blur

Blurs the Body matte that has been automatically generated when Hybrid Rendering is activated.

Hybrid Erode

This slider dilates or erodes the Hybrid matte. You can view the results by selecting Hybrid matte in the View Mode menu.

Adjust Lighting

Before applying the Adjust Lighting operation, it is necessary to determine the backing screen color using Auto Compute or Select Background Color. After performing one of those operations, click on the Adjust Lighting button. Primatte generates an artificial clean plate and uses it to generate an evenly lit backing screen behind the foreground object. The default setting should detect all the areas that contain foreground pixels and deliver a smooth backing screen for the keying.

Lighting Threshold

Should Adjust Lighting fail to produce a smoother backing screen, adjust the Lighting Threshold slider while viewing the Lighting Background setting in the View Mode menu. This displays the optimized artificial backing screen that the Adjust Lighting mode creates.

Crop

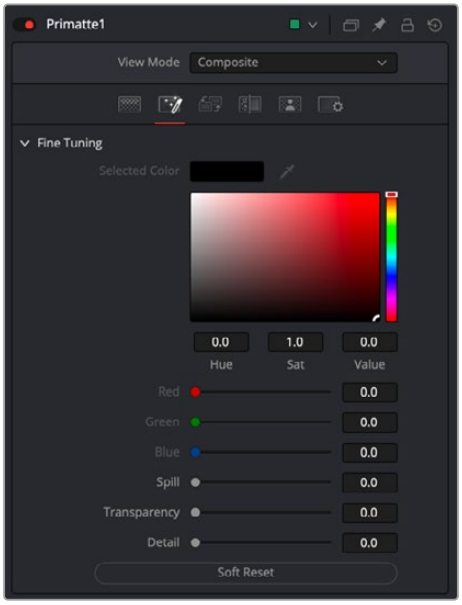
This button reveals the Crop sliders to create a rectangular garbage matte with the Primatte node. As opposed to Fusion’s Crop tool, this does not change the actual image size.

Reset

Resets all the Primatte key control data back to a blue- or green-screen.

Soft Reset

Resets just the Primatte parameters used since the Select Background Color operation was last completed.



The Primatte Fine Tuning tab

Fine Tuning Tab

The Fine Tuning tab can make refined adjustments to the spill suppression, density of the matte, and semitransparent areas. These sliders provide a bit more granularity over the Spill(+)(-), Matte(+)(-) and Detail(+)(-) buttons in the Primatte tab.

Selected Color

This shows the color selected (or registered) by the scrubbing in the viewer while the Fine Tuning tab is selected.

Fine Tuning Sliders

The color of the scrubbed pixel is registered as a reference color for fine tuning. It is displayed in the Color swatch. To perform the tuning operation, sample a color region on the image, and adjust one of the Fine Tuning sliders to achieve the desired effect.

Spill

The Spill slider can be used to remove spill from the selected color region. The more to the right the slider moves, the more spill is removed. The more to the left the slider moves, the closer the color component of the selected region is to the color in the original foreground image. If moving the slider to the right does not remove the spill, resample the color region and move the slider again.

These slider operations are additive. The result achieved by moving the slider to the right can also be achieved by clicking on the color region using the Spill(-) operational mode.

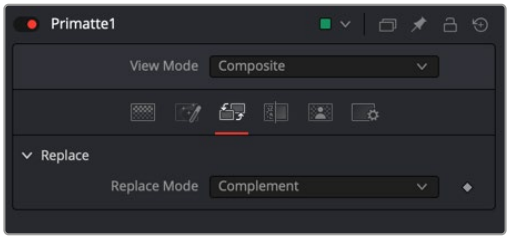
Transparency

The Transparency slider makes the matte more translucent in the selected color region. Moving this slider to the right makes the selected color region more transparent. Moving the slider to the left makes the matte more opaque. If moving the slider to the right does not make the color region translucent enough, resample the color region and again move the slider to the right. These slider operations are additive. The result achieved by moving the slider to the right can also be achieved by clicking on the color region using the Matte(-) operational mode.

Detail

The Detail slider can be used to restore lost detail. After selecting a color region, moving this slider to the left makes the selected color region more visible. Moving the slider to the right makes the color region less visible. If moving the slider to the left does not make the color region visible enough, resample the color region and again move the slider to the left.

These slider operations are additive. This result achieved by moving the slider to the left can also be achieved by clicking on the color region using the Detail(-) operational mode.



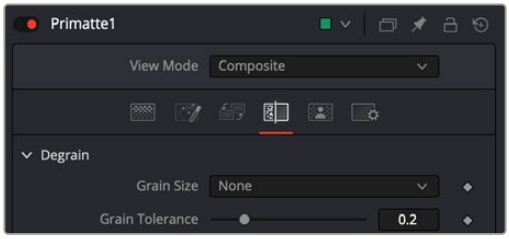
The Primatte Replace tab

Replace Tab

The Replace tab allows you to choose between the three methods of color spill replacement as covered in detail in the Spill Sponge section above. There are three options for the replacement color when removing the spill. These options are selected from the Replace mode menu.

Replace Mode

- **Complement:** Replaces the spill color with the complement of the screen color. This mode maintains fine foreground detail and delivers the best-quality results. If foreground spill is not a significant problem, this mode is the one that should be used. However, If the spill intensity on the foreground image is rather significant, this mode may often introduce serious noise in the resultant composite.
- **Image:** Replaces the spill color with colors from a defocused version of the background image or the Replace image, if one is connected to the Replace input (magenta) on the node. This mode results in a good color tone on the foreground subject even with a high-contrast background. On the negative side, the Image mode occasionally loses the fine edge detail of the foreground subjects. Another problem can occur if you later change the size of the foreground image against the background. Since the background/foreground alignment would change, the applied color tone from the defocused image might not match the new alignment.
- **Color:** Replaces the spill color with a solid color. When this option is selected, a color swatch and R,G,B sliders are displayed for selecting the color. Changing the palette color for the solid replacement, you can select a good spill replacement that matches the composite background. Its strength is that it works fine with even severe spill conditions. On the negative side, when using the Solid Color Replacement mode, fine detail on the foreground edge tends to be lost. The single palette color sometimes cannot make a good color tone if the background image has some high-contrast color areas.



The Primatte Degrain tab

Degrain Tab

The Degrain tab is used when a foreground image is highly compromised by film grain. As a result of the grain, when backing screen noise is completely removed, the edges of the foreground object often become harsh and jagged, leading to a poor key.

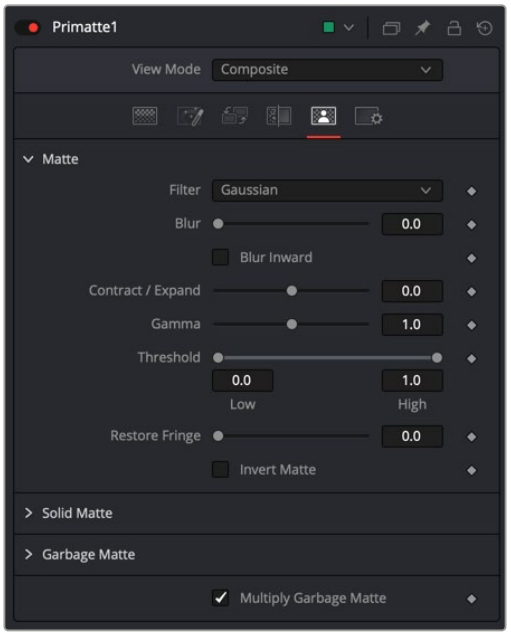
Grain Size

The Grain Size selector provides a range of grain removal from Small to Large. If the foreground image has a large amount of film grain-induced pixel noise, you may lose a good edge to the foreground object when trying to clean all the grain noise with the Clean Background Noise Operation Mode. These tools clean up the grain noise without affecting the quality of the key.

- **None:** No degraining is performed.
- **Small:** The average color of a small region of the area around the sampled pixel. This should be used when the grain is very dense.
- **Medium:** The average color of a medium-sized region of the area around the sampled pixel. This should be used when the grain is less dense.
- **Large:** The average color of a larger region of the area around the sampled pixel. This should be used when the grain is very loose.

Grain Tolerance

Adjusting this slider increases the effect of the Clean Background Noise tool without changing the edge of the foreground object.



The Primatte Matte tab

Matte Tab

The Matte tab refines the alpha of the key, combined with any solid and garbage masks connected to the node. When using the Matte tab, set the viewer to display the alpha channel of Primatte's final output.

Filter

This control selects the filtering algorithm used when applying blur to the matte.

- **Box:** This is the fastest method but at reduced quality. Box is best suited for minimal amounts of blur.
- **Bartlett:** Otherwise known as a Pyramid filter, Bartlett makes a good compromise between speed and quality.
- **Multi-Box:** When selecting this filter, the Num Passes slider appears and lets you control the quality. At 1 and 2 passes, results are identical to Box and Bartlett, respectively. At 4 passes and above, results are usually as good as Gaussian, in less time and with no edge "ringing."
- **Gaussian:** The Gaussian filter uses a true Gaussian approximation and gives excellent results, but it is a little slower than the other filters. In some cases, it can produce an extremely slight edge "ringing" on floating-point pixels.

Blur

Matte Blur blurs the edge of the matte based on the Filter menu setting. A value of zero results in a sharp, cutout-like hard edge. The higher the value, the more blur applied to the matte.

Blur Inward

Activating the Blur Inward checkbox generates the blur toward the center of the foreground subject. Conventional blurring or defocus affects the matte edges in both directions (inward and outward) and sometimes introduces a halo artifact around the edge in the composite view. Blur Inward functions only in the inward direction of the foreground subject (toward the center of the white area). The final result removes small and dark noise in the screen area without picking them up again in the Clean Background Noise mode. It can sometimes result in softer, cleaner edges on the foreground objects.

Contract/Expand

This slider shrinks or grows the semitransparent areas of the matte. Values above 0.0 expand the matte, while values below 0.0 contract it.

This control is usually used in conjunction with the Matte Blur to take the hard edge of a matte and reduce fringing. Since this control affects only semitransparent areas, it will have no effect on a matte's hard edge.

Gamma

Matte Gamma raises or lowers the values of the matte in the semitransparent areas. Higher values cause the gray areas to become more opaque, and lower values cause the gray areas to become more transparent. Completely black or white regions of the matte remain unaffected.

Since this control affects only semitransparent areas, it will have no effect on a matte's hard edge.

Threshold

This range slider sets the lower threshold using the handle on the left and sets the upper threshold using the handle on the right.

Any value below the lower threshold setting becomes black or transparent in the matte.

Any value above the upper threshold setting becomes white or opaque in the matte. All values within the range maintain their relative transparency values.

This control is often used to reject salt and pepper noise in the matte.

Restore Fringe

This restores the edge of the matte around the keyed subject. Often when keying, the edge of the subject where you have hair is clipped out. Restore Fringe brings back that edge while keeping the matte solid.

Invert Matte

When this checkbox is selected, the alpha channel created by the keyer is inverted, causing all transparent areas to be opaque and all opaque areas to be transparent.

Solid Matte

Solid mattes are mask nodes or images connected to the solid matte input on the node. The solid matte is applied directly to the alpha channel of the image. Generally, solid mattes are used to hold out keying in areas you want to remain opaque, such as someone with blue eyes against a blue screen.

Enabling Invert inverts the solid matte before it is combined with the source alpha.

Garbage Matte

Garbage mattes are mask nodes or images connected to the garbage matte input on the node. The garbage matte is applied directly to the alpha channel of the image. Generally, garbage mattes are used to remove unwanted elements that cannot be keyed, such as microphones and booms. They are also used to fill in areas that contain the color being keyed but that you wish to maintain.

Garbage mattes of different modes cannot be mixed within a single tool. A Matte Control node is often used after a Keyer node to add a garbage matte with the opposite effect of the matte applied to the keyer.

Enabling Invert inverts the garbage matte before it is combined with the source alpha.

Post-Multiply Image

Select this option to cause the keyer to multiply the color channels of the image against the alpha channel it creates for the image. This option is usually enabled and is on by default.

Deselect this checkbox and the image can no longer be considered premultiplied for purposes of merging it with other images. Use the Subtractive option of the Merge node instead of the Additive option.

> For more information on these Merge node settings, *see Chapter 35, "Composite Nodes,"* in the Fusion Reference Manual.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Matte nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

How to Key with Primatte

You begin keying with Primatte by connecting the blue- or green-screen shot to the orange foreground input on the Primatte node and the background shot for the composite into the green background input. Once the connections are made, there are four main steps to using the Primatte:

1. Select Background Color.
2. Clean the Background Noise.
3. Clean the Foreground Noise.
4. Remove Spill.

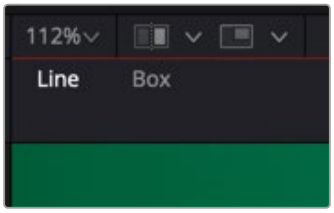
Selecting Background Color

- In the Inspector, click the Select Background Color button.
- Position the mouse pointer over the blue/green-screen area in the viewer, somewhere near the foreground subject.
- Drag over the background color.

  Primatte averages the pixels to get a single color. Sometimes Primatte works best when only a single pixel is sampled instead of a range of pixels.

  Should you have difficulties with your keying, try the Select Background Color operation again with a single dark screen pixel or single light screen pixel.

  Instead, If you want to make a rectangular selection, use the Box button in the top left-hand corner of the viewer. The Median button is the same as Line selection, except that each point sampled is the result of a 3 x 3 region based on where you click and then apply a median filter. This can potentially reduce any noisy pixels.



Primatte viewer buttons

If the foreground image has a shadow in it that you want to keep in the composite, do not select any of the dark screen pixels in the shadow. This keeps the shadow with the rest of the foreground image.

Clean Background Noise

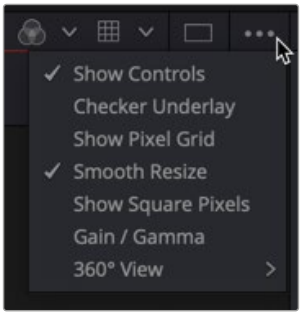
If there are any white or light gray regions in the dark screen area, this is referred to as “noise.” Technically, it is varying shades of the screen color that did not get picked up on the first sample and should be removed. You remove background noise using the Clean Background Noise button.

1. From the View Mode menu in the Inspector, select Black.
2. Above the viewer, click the Alpha Channel/RGB button.

   The image displayed changes to a black and white “matte” view of the image.
3. Click the Clean Background Noise button.
4. Drag the mouse pointer through these white or light gray regions that should be pure black.

   Primatte processes the selection and eliminates the noise.
5. Repeat this procedure as often as necessary to clear the noise from the background areas.

   Selecting Gain/Gamma from the viewer’s Options menu to increase the brightness or gamma allows you to see noise that would otherwise be invisible.



Primatte viewer Options menu

You do not need to remove every single white pixel to get good results. Most pixels displayed as a dark color close to black in a key image are considered transparent and virtually allow the background to be the final output in that area. Consequently, there is no need to eliminate all noise in the screen portions of the image. In particular, if an attempt is made to remove noise around the foreground subject meticulously, a smooth composite image is often difficult to generate.

TIP: When clearing noise from around loose, flying hair or any background/foreground transitional area, be careful not to select any of the areas near the edge of the hair. Leave a little noise around the hair as this can be cleaned up later using the Fine Tuning tools.

Clean Foreground Noise

If there are dark regions in the middle of the mostly white foreground subject, the key is not 100% in those areas. Using Clean Foreground Noise can make the matte as white as possible.

1. Keep the View Mode menu set to Black and the viewer set to the Alpha Channel.
2. Click the Clean Foreground Noise button.
3. Drag the mouse pointer through these dark pixels in the foreground that should be pure white.

   Primatte processes the selection and eliminates the noise.
4. Repeat this procedure as often as necessary to clear the noise from the foreground areas.
5. If enabled, disable Gain/Gamma from the viewer's Options menu to return to a regular viewer.

Removing Spill

The first three sections created a clean matte. At this point, the foreground can be composited onto any background image. However, if there is color spill on the foreground subject, a final operation is necessary to remove that screen spill for a more natural-looking composite.

1. From the View Mode menu, select Composite.
2. Above the viewer, click the Alpha/RGB toggle button to see the RGB image.

There are two ways in Primatte to remove the spill color:

Spill Sponge

The quickest method is to select the Spill Sponge button and then sample the spill areas away. Additional spill removal can be done using the tools under the Fine Tuning tab or by using the Spill(-) button.

Fine Tuning Tab

To use the Fine Tuning tab for spill, first scrub over the spill color in the viewer. For most images, adjusting the Spill slider is all that is required to remove any remaining spill.

NOTE: When using the slider in the Fine Tuning tab to remove spill, spill color replacement is replaced based on the setting of the Spill Replacement options.

You can use the other two sliders in the same way for different key adjustments. The Detail slider controls the matte softness for the color that is closest to the background color. For example, you can recover lost rarefied smoke in the foreground by selecting the Fine Tuning mode, clicking on the area of the image where the smoke starts to disappear and moving the Detail slider to the left. The Transparency slider controls the matte softness for the color that is closest to the foreground color. For example, if you have thick and opaque smoke in the foreground, you can make it semitransparent by moving the Transparency slider to the right after selecting the pixels in the Fine Tuning mode.

Relight [RLt]



The Relight node

Relight Node Introduction

Relight allows you to add light sources to your scene, realistically illuminating the objects within it. Unlike a simple shape node whose polygon is fixed on screen, Relight analyzes the shapes of people and objects to create a Surface Map and uses that map to reflect light appropriately. When dialed-in properly, this effect can provide the illusion of light following the curves of an object, reflecting off highlights and disappearing into shadows.

Relight is useful in situations when you wished you had another light in the scene, want to accentuate an existing light to match another shot, or want to increase the realism of day-for-night grading.

NOTE: The Relight effect creates the illusion of depth with its analysis of the surfaces in a scene, but it does not calculate 3D relationships between objects in the scene. Shapes cannot be made to cast shadows, and depth information cannot be produced.

For more information on Relight, see Chapter 162, “Resolve FX Refine,” in the *DaVinci Resolve 20 Reference Manual*.

Inputs

The Relight tool has three inputs.

Input: The orange input is used for the primary footage you wish to relight.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the filter to only those pixels within the mask. An effects mask is applied to the tool after the tool is processed.

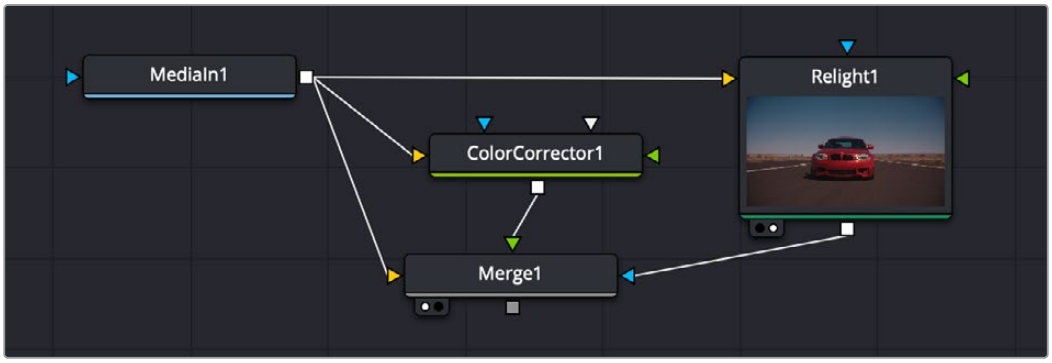
Normals: The green input is used to connect an externally created normal maps, instead of the one generated internally by the Relight tool.

Basic Node Setup

In order to use the Relight tool, you need an additional node structure to build around it. In addition to the standard MediaIn node, you will need a Merge node, and another effect like the ColorCorrector or Brightness/Contrast tools.

- **MediaIn:** The output from the MediaIn node goes to the orange inputs of the Relight, ColorCorrector (or other effect node), and the Merge node.

- **ColorCorrector (effect):** Receives the input from the MediaIn node, and its output connects to the green foreground input of the Merge node.
- **Relight:** Receives the input from the MediaIn node, and its output connects to the blue effect mask input on the Merge node. Set this node to appear in the viewer of your choice.
- **Merge:** Receives the background input from the MediaIn node, receives the green foreground input from the ColorCorrector, and the blue effect mask input from Relight. Its output goes to the rest of the node tree or media out. Set this node to appear in the viewer of your choice.



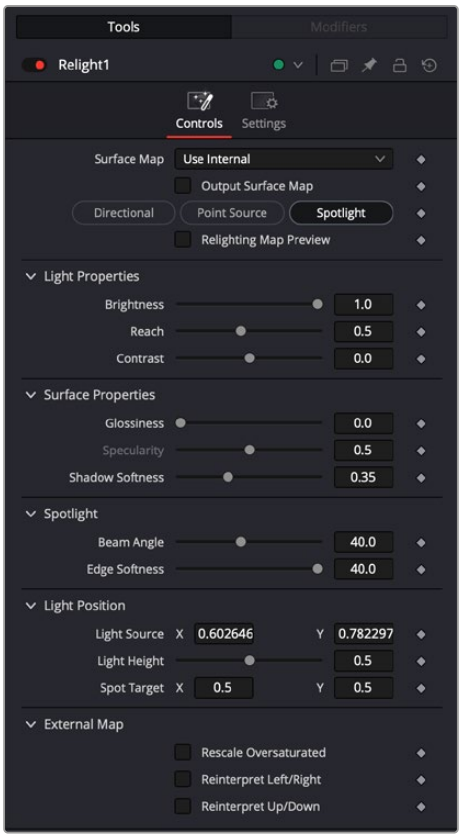
The node tree showing the surrounding node structure for the Relight tool. Note that the ColorCorrector can be replaced by any other preferred effect.

With the above node structure in place, you should see something similar to the setup below. On Viewer 1 is the final relight effect that gets passed onto the rest of the node tree. On Viewer 2 you have the relight interface that allows you to drag the lighting position around the frame. Depending on your Relighting Map Preview option, this will either be a black and white embossed image, or the RGB image with a semi-transparent background.



The Merge node in Viewer 1 (left), and the Relight node in Viewer 2 (right); both sets of on-screen controls are active.

Inspector



The Relight controls

Controls

The Controls let you adjust all the parameters of the Relight tool. There are additional on-screen controls in the viewer as well, depending on the type of Relight effect you choose.

For more information on using Relight controls, see Chapter 162, “Resolve FX Refine,” in the *DaVinci Resolve 20 Reference Manual*.

- **Surface Map:** This chooses which mode you will use to create the Surface Map.

  Use Internal: Generates the Surface Map internally to the tool.

  Use Normals Input: Receives an externally created Surface Map from another node. Connect that map to the green input of the Relight node.

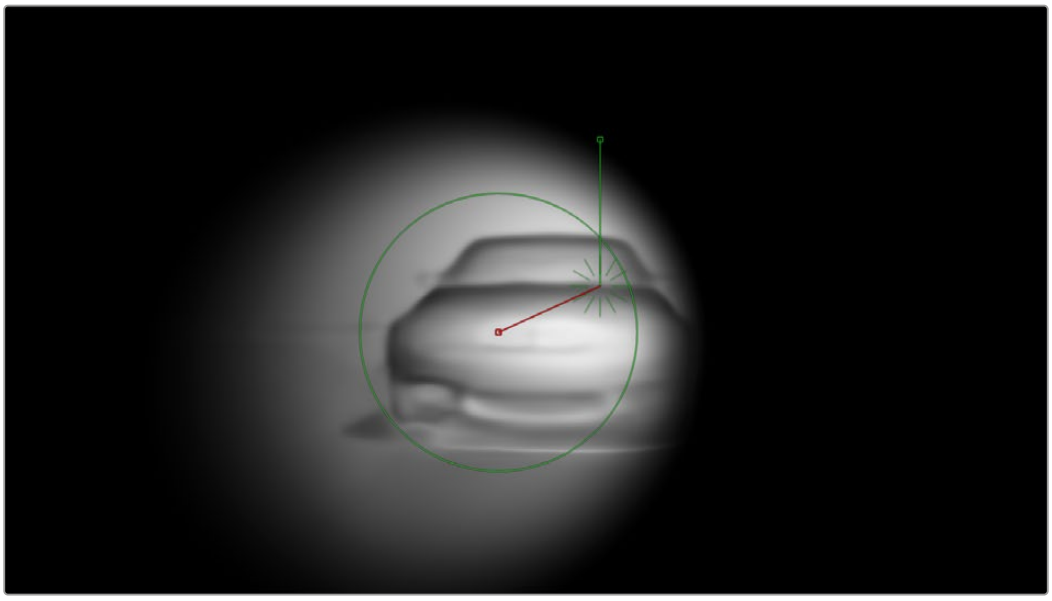
- **Output Surface Map:** Checking this box causes this copy of the effect to output a multi-colored map of the input frame. This Surface Map can then be fed into a second Relight node where the light sources are placed and adjusted.
- **Directional / Point Source / Spotlight:**

  Directional: Creates a light source that emits from a specific direction at an infinite distance.

Point Source: Creates a light source that emits from a central point.

Spotlight: Is also a point source but it is directed into a cone shape.

- **Relighting Map Preview:** Checking this box reveals the Alpha channel of the effect of this light. Unchecking this box shows the current image but with the light mask in the Alpha channel. Bright areas of the Map Preview show where the Relight is affecting the image.



The image of a car with the Map Preview box checked. Brighter areas show where and how the Relight tool will affect the image.

Light Properties

- **Brightness:** This controls the strength of any grading that is applied to the scene. It does not add brightness changes on its own but amplifies the changes you make in the grading tools.
- **Reach:** This controls how much reflections will dim with distance from the light source.
- **Contrast:** This determines how harsh or gentle the gradient is between light and dark.

Surface Properties

- **Glossiness:** Gives the illusion of a metallic sheen to reflective surfaces.
- **Specularity:** When there is any glossiness added, this controls how metallic/shiny that glossiness is.
- **Shadow Softness:** Controls how much the light penetrates into shadows due to ambient reflections, beyond where it would reach otherwise.

Light Direction (Directional only)

- **Azimuth/Elevation XY :** Controls the direction the light is coming from.

Spotlight (Spotlight only)

- **Beam Angle:** Sets how broad the cone of light is. This control is also available graphically in the viewer.
- **Edge Softness:** Controls the softness of the edges of the light cone/beam. Light Position (Point Source and Spotlight).

Light Position (Point and Spotlight only)

- These controls are for the the position of the light emitter for Point and Spotlight types. They can be adjusted here directly, but they are adjusted graphically using the Viewer controls more easily.

External Map

- **Rescale Oversaturated:** Different applications use different scalings of the Surface Map. A common format looks like the Relight plugin's internal ones but with the colors intensely saturated and can be dark or black in some areas. If your imported surface map has this appearance, check this box to re-interpret it.
- **Reinterpret Left/Right:** Check this box to swap the direction of your external surface map, if your map surfaces are appearing concave instead of convex.
- **Reinterpret Up/Down:** Check this box to swap the direction of your external surface map, if your map surfaces are appearing lit from below when the light is above them.

Ultra Keyer [UKy]



The Ultra Keyer node

Ultra Keyer Node Introduction

Like the newer Delta Keyer, the Ultra Keyer node has two keyers built in to it: a pre-matte keyer acts as a garbage matte creator and the color difference keyer that extracts fine detail and transparency. Generally, you start with the Delta Keyer as your first keyer of choice. If you do not get good results, try Primatte if you are using Fusion Studio. A good third choice is to try the Ultra Keyer.

Inputs

The Ultra Keyer node includes four inputs in the Node Editor.

Input: The orange input accepts a 2D image that contains the color you want to be keyed for transparency.

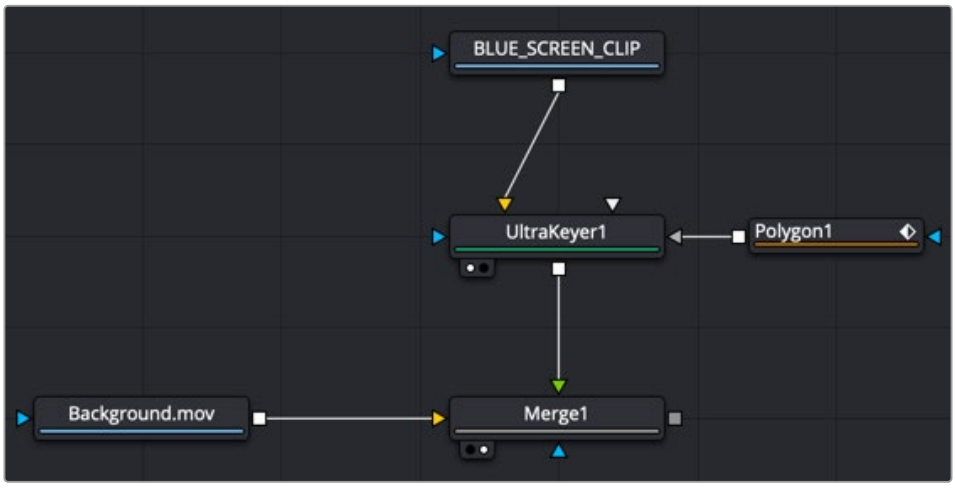
Garbage Matte: The gray garbage matte input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input causes areas of the image that fall within the matte to be made transparent. The garbage matte is applied directly to the alpha channel of the image.

Solid Matte: The white solid matte input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input causes areas of the image that fall within the matte to be fully opaque.

Effect Mask: The optional blue input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps masks. Connecting a mask to this input limits the pixels where the keying occurs. An effects mask is applied to the tool after the tool is processed.

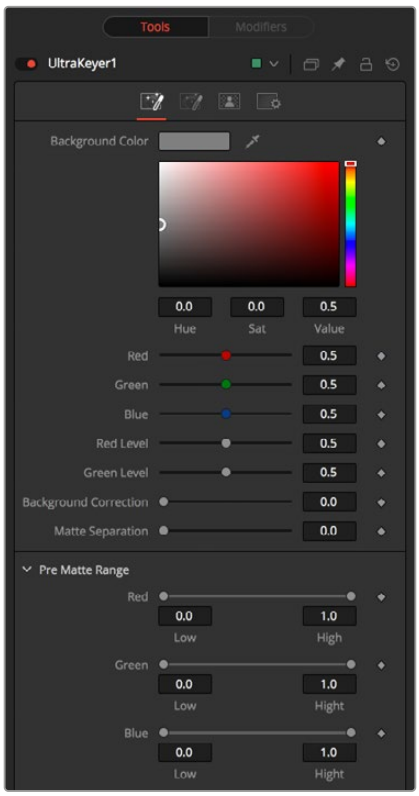
Basic Node Setup

A single keyer can rarely get perfect results because most green- or blue-screen shots have problems the keyer is not made to handle. Keyers often need the help of garbage mattes or solid mattes created with a Polygon or B-Spline node. Shots can also require more than just one keyer to achieve perfect results. Below, the Ultra Keyer node has the blue-screen content connected to the orange input. The result is an image with alpha that can then be connected into the foreground of a Merge node.



An Ultra Keyer node combined with polygon matte as a garbage matte and connected to the foreground of a Merge

Inspector



The Ultra Keyer Pre-Matte tab

Pre-Matte Tab

The Pre-Matte tab is where most keying begins. It is used to select the screen color and smooth out the color of the screen.

Background Color

The Background Color is used to select the color of the blue or green screen of the images. It is good practice to select the screen color close to the subject to be separated from the screen background.

Red Level, Green Level, Blue Level

These color sliders tune the level of the difference channels to help separate the color. When the background color is green, Red and Blue level options are provided. When the background color is blue, Red and Green level options are provided.

Background Correction

Depending on the background color selected above, the keyer iteratively merges the pre-keyed image over either a blue or green background before processing it further.

In some instances, this leads to better, more subtle edges.

Matte Separation

Matte Separation performs a pre-process on the image to help separate the foreground from the background before color selection. Generally, increase this control while viewing the alpha to eliminate the bulk of the background, but stop just before it starts cutting holes in the subject or eroding fine detail on the edges of the matte.

Pre-Matte Range

These R,G,B, and Luminance range controls update automatically to represent the current color selection. Colors are selected by selecting the Ultra Keyer node's tile in the node tree and dragging the Eyedropper into the viewer to select the colors to be used to create the matte. These range controls can be used to tweak the selection slightly, although selecting colors in the viewer is all that is required.

Lock Color Picking

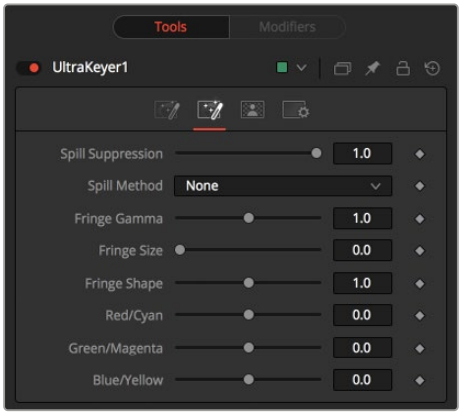
This checkbox prevents accidentally selecting more colors from the view. It is a good idea to activate this checkbox once the color selection is made for the matte. All other controls in the node remain editable.

Pre Matte Size

The Pre Matte Size control can be used to soften the general area around the keyed image. This is used to close holes in the matte often caused by spill in semitransparent areas of the subject. This can cause a small halo around the subject, which can be removed using the Matte Contract tools found later in the tool.

Reset Pre Matte Ranges

This discards all color selection by resetting the ranges but maintains all other slider and control values.



The Ultra Keyer Image tab

Image Tab

The Image tab handles the majority of spill suppression in the Ultra Keyer. Spill suppression is a form of color correction that attempts to remove the screen color from the fringe of the matte.

Spill is the transmission of the screen color through the semitransparent areas of the alpha channel. In the case of blue- or green-screen keying, this usually causes the color of the background to become apparent in the edges of the foreground subject.

Spill Suppression

When this slider is set to 0, no spill suppression is applied to the image. Increasing the slider increases the strength of the spill method.

Spill Method

This selects the strength of the algorithm used to apply spill suppression to the image.

- **None:** None is selected when no spill suppression is required.
- **Rare:** This removes very little of the spill color and is the lightest of all methods.
- **Medium:** This works best for green screens.
- **Well Done:** This works best for blue screens.
- **Burnt:** This works best for blue screens. Use this mode only for very troublesome shots.

Fringe Gamma

This control can be used to adjust the brightness of the fringe or halo that surrounds the keyed image.

Fringe Size

This expands and contracts the size of the fringe or halo surrounding the keyed image.

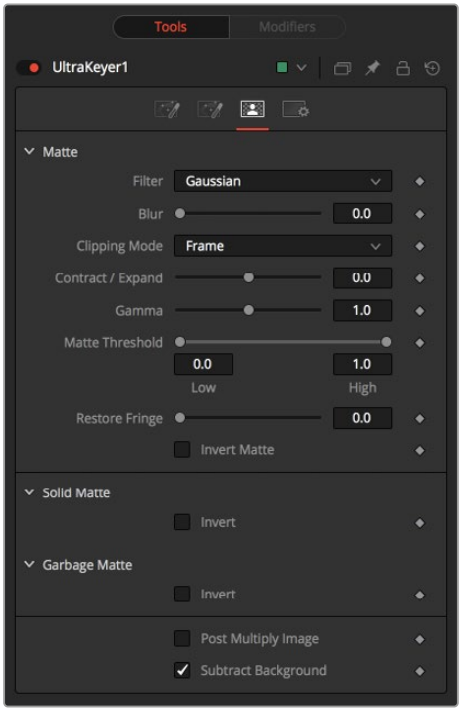
Fringe Shape

Fringe Shape presses the fringe toward the external edge of the image or pulls it toward the inner edge of the fringe. Its effect is most noticeable while the Fringe Size value is large.

Cyan/Red, Magenta/Green, and Yellow/Blue

Use these three controls to color correct the fringe of the image.

This is useful for correcting semitransparent pixels that still contain color from the original background to match the new background.



The Ultra Keyer Matte tab

Matte Tab

The Matte tab refines the alpha of the key, combined with any solid and garbage masks connected to the node. When using the Matte tab, set the viewer to display the alpha channel of the Delta Keyer's final output.

Filter

This control selects the filtering algorithm used when applying blur to the matte.

- **Box:** This is the fastest method but at reduced quality. Box is best suited for minimal amounts of blur.
- **Bartlett:** Otherwise known as a Pyramid filter, Bartlett makes a good compromise between speed and quality.
- **Multi-Box:** When selecting this filter, the Num Passes slider appears and lets you control the quality. At 1 and 2 passes, results are identical to Box and Bartlett, respectively. At 4 passes and above, results are usually as good as Gaussian, in less time and with no edge "ringing."
- **Gaussian:** The Gaussian filter uses a true Gaussian approximation and gives excellent results, but it is a little slower than the other filters. In some cases, it can produce an extremely slight edge "ringing" on floating-point pixels.

Blur

Matte Blur blurs the edge of the matte based on the Filter menu setting. A value of zero results in a sharp, cutout-like hard edge. The higher the value, the more blur applied to the matte.

Clipping Mode

This option determines how edges are handled when performing domain of definition rendering. This is profoundly important when blurring the matte, which may require samples from portions of the image outside the current domain.

- **Frame:** The default option is Frame, which automatically sets the node's domain of definition to use the full frame of the image, effectively ignoring the current domain of definition. If the upstream DoD is smaller than the frame, the remaining areas in the frame are treated as black/transparent.
- **Domain:** Setting this option to Domain respects the upstream domain of definition when applying the node's effect. This can have adverse clipping effects in situations where the node employs a large filter.
- **None:** Setting this option to None does not perform any source image clipping at all. This means that any data required to process the node's effect that would normally be outside the upstream DoD is treated as black/transparent.

Contract/Expand

This slider shrinks or grows the semitransparent areas of the matte. Values above 0.0 expand the matte, while values below 0.0 contract it.

This control is usually used in conjunction with the Matte Blur to take the hard edge of a matte and reduce fringing. Since this control affects only semitransparent areas, it has no effect on a matte's hard edge.

Gamma

Matte Gamma raises or lowers the values of the matte in the semitransparent areas. Higher values cause the gray areas to become more opaque, and lower values cause the gray areas to become more transparent. Completely black or white regions of the matte remain unaffected.

Since this control affects only semitransparent areas, it will have no effect on a matte's hard edge.

Threshold

This range slider sets the lower threshold using the handle on the left and sets the upper threshold using the handle on the right.

Any value below the lower threshold setting becomes black or transparent in the matte.

Any value above the upper threshold setting becomes white or opaque in the matte. All values within the range maintain their relative transparency values.

This control is often used to reject salt and pepper noise in the matte.

Restore Fringe

This restores the edge of the matte around the keyed subject. Often when keying, the edge of the subject where you have hair is clipped out. Restore Fringe brings back that edge while keeping the matte solid.

Invert Matte

When this checkbox is selected, the alpha channel created by the keyer is inverted, causing all transparent areas to be opaque and all opaque areas to be transparent.

Solid Matte

Solid mattes are mask nodes or images connected to the solid matte input on the node. The solid matte is applied directly to the alpha channel of the image. Generally, solid mattes are used to hold out keying in areas you want to remain opaque, such as someone with blue eyes against a blue screen.

Enabling Invert inverts the solid matte before it is combined with the source alpha.

Garbage Matte

Garbage mattes are mask nodes or images connected to the garbage matte input on the node. The garbage matte is applied directly to the alpha channel of the image. Generally, garbage mattes are used to remove unwanted elements that cannot be keyed, such as microphones and booms. They are also used to fill in areas that contain the color being keyed but that you wish to maintain.

Garbage mattes of different modes cannot be mixed within a single tool. A Matte Control node is often used after a Keyer node to add a garbage matte with the opposite effect of the matte applied to the keyer.

Enabling Invert inverts the garbage matte before it is combined with the source alpha.

Post-Multiply Image

Select this option to cause the keyer to multiply the color channels of the image against the alpha channel it creates for the image. This option is usually enabled and is on by default.

Deselect this checkbox and the image can no longer be considered premultiplied for purposes of merging it with other images. Use the Subtractive option of the Merge node instead of the Additive option.

> For more information on these Merge node settings, *see Chapter 35, "Composite Nodes,"* in the Fusion Reference Manual.

Subtract Background

This option color corrects the edges when the screen color is removed and anti-aliased to a black background. By enabling this option, the edges potentially become darker. Disabling this option allows you to pass on the color of the screen to use in other processes down the line.

Common Controls

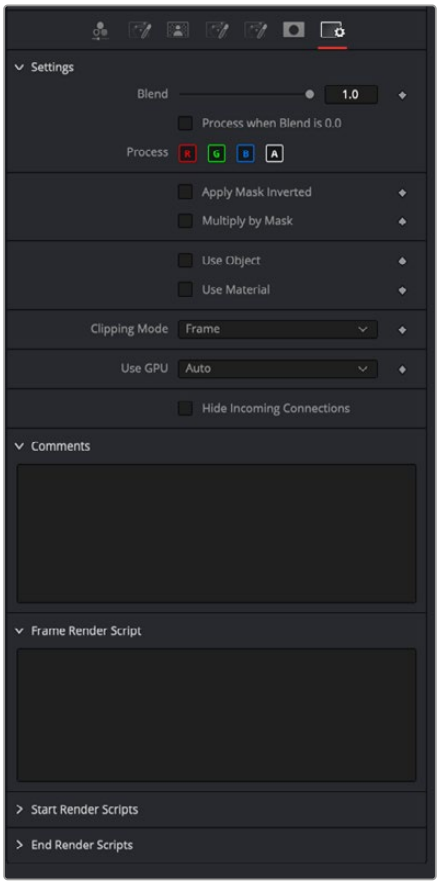
Settings Tab

The Settings tab in the Inspector is also duplicated in other matte nodes. These common controls are described in detail in the following “The Common Controls” section.

The Common Controls

Nodes that handle matte operations share a number of identical controls in the Inspector. This section describes controls that are common among matte nodes.

Inspector



Common Matte settings Inspector

Settings Tab

The Settings tab in the Inspector can be found on every tool in the Matte category. The controls are consistent and work the same way for each tool.

Blend

The Blend control is used to blend between the tool's original image input and the tool's final modified output image. When the blend value is 0.0, the outgoing image is identical to the incoming image. Usually, this causes the tool to skip processing entirely, copying the input straight to the output.

Process When Blend Is 0.0

The tool is processed even when the input value is zero. This can be useful if this node is scripted to trigger another task, but the value is set to 0.0.

Red/Green/Blue/Alpha Channel Selector

These four buttons are used to limit the effect of the tool to specified color channels. This filter is often applied after the tool has been processed.

For example, if the Red button on a Blur tool is deselected, the blur is first applied to the image, and then the red channel from the original input is copied back over the red channel of the result.

There are some exceptions, such as tools where deselecting these channels causes the tool to skip processing that channel entirely. In that case, there are a set of RGBA buttons on the Controls tab in the tool. The buttons in the Settings and the Controls tabs are identical.

Apply Mask Inverted

Enabling the Apply Mask Inverted option inverts the complete mask channel for the tool. The mask channel is the combined result of all masks connected to or generated in a node.

Multiply by Mask

Selecting this option causes the RGB values of the masked image to be multiplied by the mask channel's values. This causes all pixels of the image not included in the mask (i.e., set to 0) to become black/transparent.

Use Object/Use Material (Checkboxes)

Some 3D software can render to file formats that support additional channels. Notably, the EXR file format supports Object ID and Material ID channels, which can be used as a mask for the effect. These checkboxes determine whether the channels are used, if present. The specific Material ID or Object ID affected is chosen using the next set of controls.

Correct Edges

This checkbox appears only when the Use Object or Use Material checkboxes are selected. It toggles the method used to deal with overlapping edges of objects in a multi-object image. When enabled, the Coverage and Background Color channels are used to separate and improve the effect around the edge of the object. If this option disabled (or no Coverage or Background Color channels are available), aliasing may occur on the edge of the mask.

> For more information on the Coverage and Background Color channels, *see Chapter 18, "Understanding Image Channels,"* in the Fusion Reference Manual.

Object ID/Material ID (Sliders)

Use these sliders to select which ID is used to create a mask from the object or material channels of an image. Use the Sample button in the same way as the Color Picker: to grab IDs from the image displayed in the view. The image or sequence must have been rendered from a 3D software package with those channels included.

Clipping Mode

This option determines how edges are handled when performing domain of definition rendering. This is profoundly important when blurring the matte, which may require samples from portions of the image outside the current domain.

- **Frame:** The default option is Frame, which automatically sets the node's domain of definition to use the full frame of the image, effectively ignoring the current domain of definition. If the upstream DoD is smaller than the frame, the remaining area in the frame is treated as black/transparent.
- **Domain:** Setting this option to Domain respects the upstream domain of definition when applying the node's effect. This can have adverse clipping effects in situations where the node employs a large filter.
- **None:** Setting this option to None does not perform any source image clipping at all. This means that any data required to process the node's effect that would normally be outside the upstream DoD is treated as black/transparent.

Use GPU

The Use GPU menu has three settings. Setting the menu to Disable turns off GPU hardware-accelerated rendering. Enabled uses the GPU hardware for rendering the node. Auto uses a capable GPU if one is available, and falls back to software rendering when a capable GPU is not available

Hide Incoming Connections

Enabling this checkbox can hide connection lines from incoming nodes, making a node tree appear cleaner and easier to read. When enabled, empty fields for each input on a node are displayed in the Inspector. Dragging a connected node from the node tree into the field hides that incoming connection line as long as the node is not selected in the node tree. When the node is selected in the node tree, the line reappears.

Comments

The Comments field is used to add notes to a tool. Click in the empty field and type the text. When a note is added to a tool, a small red square appears in the lower-left corner of the node when the full tile is displayed, or a small text bubble icon appears on the right when nodes are collapsed. To see the note in the Node Editor, hold the mouse pointer over the node to display the tooltip.

Scripts

Three Scripting fields are available on every tool in Fusion from the Settings tab. They each contain edit boxes used to add scripts that process when the tool is rendering. For more details on scripting nodes, please consult the Fusion scripting documentation.

Chapter 50

Metadata Nodes

This chapter details the Metadata nodes available in Fusion.

The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Copy Metadata [Meta] 1333
Set Metadata [SMeta] 1334
Set Timecode [TcMeta] 1336
The Common Controls 1338

Copy Metadata [Meta]



The Copy Metadata node

Copy Metadata Node Introduction

Copy Metadata combines, replaces, or clears the metadata in your image. Metadata can be viewed in a subview of the viewer.

Inputs

The two inputs on the Copy Metadata node are used to connect two 2D images.

Background Input: The orange background input is used for the primary 2D image that is output from the node.

Foreground Input: The green foreground input is used for the secondary 2D image that contains metadata you want merge or overwrite onto the background image.

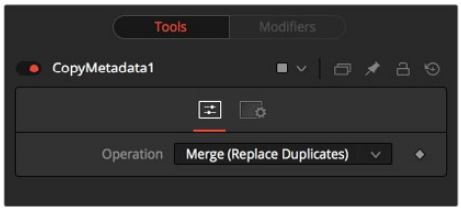
Basic Node Setup

The Copy Metadata node takes metadata from the foreground input (green) and copies it into the background input (orange). The output is the background input with modified metadata.



A Copy Metadata node copies metadata from the foreground and embeds it into the background clip

Inspector



The Copy Metadata Controls tab

Controls Tab

The Controls tab configures how metadata coming from the foreground input image gets added to the background input image.

Operation

The Operation menu determines how the metadata of the foreground and background inputs are treated.

- **Merge (Replace Duplicates):** All values are merged, but values with duplicate names are taken from the foreground input.
- **Merge (Preserve Duplicates):** All values are merged, but values with duplicate names are taken from the background input.
- **Replace:** The metadata in the foreground replaces the entire metadata in the background.
- **Clear:** All metadata is discarded.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Metadata nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Set Metadata [SMeta]



The Set Metadata node

Set Metadata Node Introduction

Set Metadata allows you to create new Name = Value pairs in the metadata. Metadata can be viewed in a subview of the viewer.

Inputs

The single input on the Set Metadata node is used to connect a 2D image that gets metadata added.

Background Input: The orange background input is used for the primary 2D image that is output from the node with the new metadata.

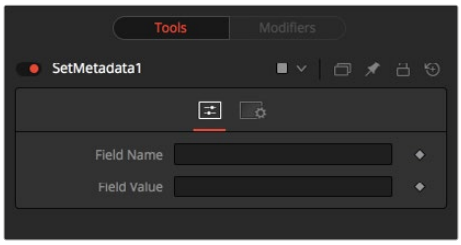
Basic Node Setup

The Set Metadata node embeds new metadata into the background input (orange). The output is the background input with new metadata.



A Set Metadata node creates new metadata and embeds it into the background clip.

Inspector



The Set Metadata Controls tab

Controls Tab

The Controls tab is where you set up the name of the metadata field and the value or information regarding the metadata.

Field Name

The name of the metadata value. Do not use spaces.

Field Value

The value assigned to the name above.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Metadata nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Set Timecode [TcMeta]



The Set Timecode node

Set Timecode Node Introduction

Set Timecode inserts dynamic timecode values into the metadata table based on the FPS settings.

Inputs

The single input on the Set Timecode node is used to connect a 2D image that gets timecode added.

Background Input: The orange background input is used for the primary 2D image that is output from the node with the new timecode.

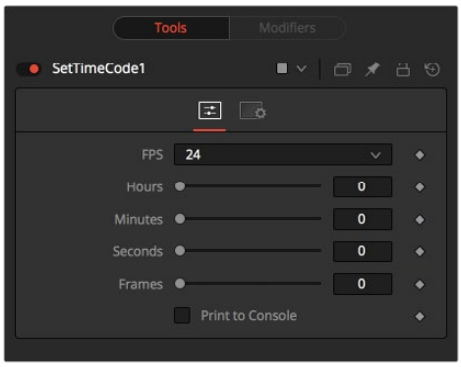
Basic Node Setup

The Set Timecode node embeds new timecode metadata into the background input (orange). The output is the background input with updated timecode.



A Set Timecode node inserts new timecode metadata into the background clip.

Inspector



The Set Timecode Controls tab

Controls Tab

The Controls tab sets the clip's starting timecode metadata based on FPS, hours, minutes, seconds, and frames.

FPS

You can choose from a variety of settings for frames per second.

Since this is a Fuse, you can easily adapt the settings to your needs by editing the appropriate piece of code for the buttons:

```
MBTNC_StretchToFit = true,
{ MBTNC_AddButton = "24" },
{ MBTNC_AddButton = "25" },
{ MBTNC_AddButton = "30" },
{ MBTNC_AddButton = "48" },
{ MBTNC_AddButton = "50" },
{ MBTNC_AddButton = "60" },
})
```

as well as for the actual values:

```
local rates = { 24, 25, 30, 48, 50, 60 }
```

Hours/Minutes/Seconds/Frames Sliders

Define an offset from the starting frame of the current comp.

Print to Console

Verbose output of the Timecode/Frame value in the Console.

The Timecode/Frames conversion is done according to the FPS settings.

The result might look like this:

```
TimeCode: 00:00:08:15
Frames:   207
```

Common Controls

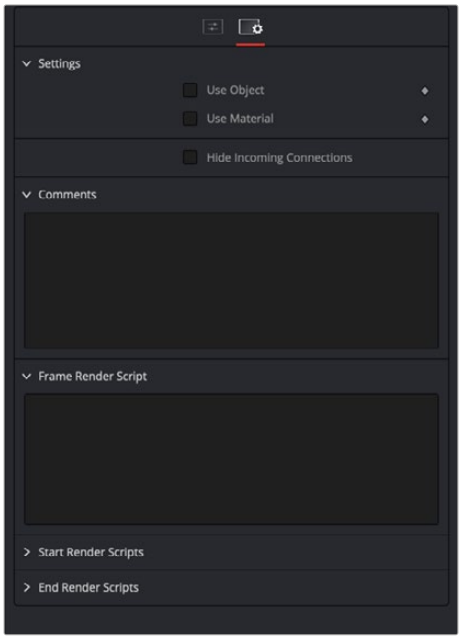
Settings Tab

The Settings tab in the Inspector is also duplicated in other Metadata nodes. These common controls are described in detail in the following “The Common Controls” section.

The Common Controls

Nodes that handle metadata operations share several identical controls in the Inspector. This section describes controls that are common among Metadata nodes.

Inspector



The Common Metadata Settings tab

Settings Tab

The Settings tab in the Inspector can be found on every tool in the Metadata category. The controls are consistent and work the same way for each tool.

Use Object/Use Material (Checkboxes)

Some 3D software can render to file formats that support additional channels. Notably, the EXR file format supports Object ID and material ID channels, which can be used as a mask for the effect. These checkboxes determine whether the channels are used, if present. The specific Material ID or Object ID affected is chosen using the next set of controls.

Correct Edges

This checkbox appears only when the Use Object or Use Material checkboxes are selected. It toggles the method used to deal with overlapping edges of objects in a multi-object image. When enabled, the Coverage and Background Color channels are used to separate and improve the effect around the edge of the object. If this option disabled (or no Coverage or Background Color channels are available), aliasing may occur on the edge of the mask.

> For more information on the Coverage and Background Color channels, *see Chapter 18, "Understanding Image Channels,"* in the Fusion Reference Manual.

Object ID/Material ID (Sliders)

Use these sliders to select which ID is used to create a mask from the object or material channels of an image. Use the Sample button in the same way as the Color Picker: to grab IDs from the image displayed in the view. The image or sequence must have been rendered from a 3D software package with those channels included.

Comments

The Comments field is used to add notes to a tool. Click in the empty field and type the text. When a note is added to a tool, a small red square appears in the lower-left corner of the node when the full tile is displayed, or a small text bubble icon appears on the right when nodes are collapsed. To see the note in the Node Editor, hold the mouse pointer over the node to display the tooltip.

Scripts

Three Scripting fields are available on every tool in Fusion from the Settings tab. They each contain edit boxes used to add scripts that process when the tool is rendering. For more details on scripting nodes, please consult the Fusion scripting documentation.

Chapter 51

Miscellaneous Nodes

This chapter details miscellaneous nodes within Fusion.

The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Auto Domain [ADoD] 1341
Change Depth [CD] 1343
Custom Tool [CT] 1344
Fields [FLDs] 1356
Frame Average [Avg] 1358
Keyframe Stretcher [KfS] 1359
Run Command [Run] 1362
Set Domain [DOD] 1365
Switch [Swi] 1367
Time Speed [TSpd] 1368
Time Stretcher [TSt] 1371
Wireless Link [Wire] 1375
The Common Controls 1376

Auto Domain [ADoD]



The Auto Domain node

Auto Domain Node Introduction

The Auto Domain node automatically sets the image's domain of definition (DoD) based on bounds of the input image's background Canvas color. It does not change the image's physical dimensions. Some EXR images come with optimized DoDs already set, but other formats do not. For formats other than EXR, this node can speed up compositions by optimizing the DoD based on the content rather than the frame's dimensions.

For example, a CG character rarely takes up the entire frame of an image. With this type of image, the Auto Domain node sets the DoD to a rectangular region by comparing image pixels with the Canvas color. The Canvas color indicates what color the pixels are outside the DoD. By default, unless a Canvas color is set using the Set Canvas Color node, the color is set to black. This default works well when an image has a premultiplied alpha channel. The result is a DoD that encompasses the portion of the clip that contains only the character. The DoD is updated on each frame to accommodate changes, such as a character walking closer to the camera. However, if a clip does not contain an alpha channel, the Set Canvas Color node can be used to define the Canvas color as solid alpha with a color that matches the solid background.

For more detail on the Set Canvas Color node, *see Chapter 34, "Color Nodes,"* in the Fusion Reference Manual.

NOTE: The Domain of Definition is a bounding box that encompasses pixels that have a nonzero value. The DoD is used to limit image-processing calculations and speeds up rendering.

Inputs

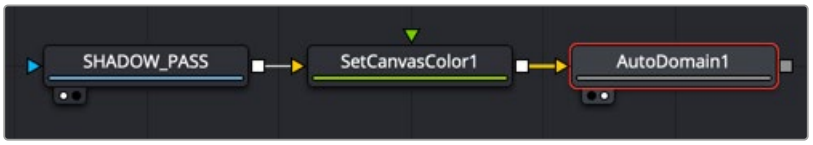
The single input on the Auto Domain node is used to connect a 2D image and an effect mask, which can be used to limit the blurred area.

Input: The orange input is used for the primary 2D image that is blurred.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the blur to only those pixels within the mask.

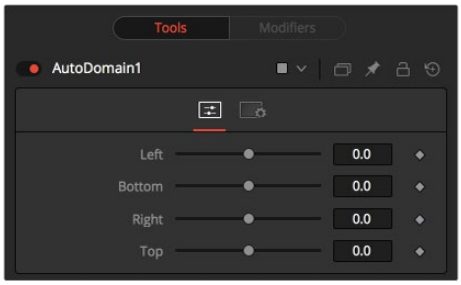
Basic Node Setup

The example below assumes that the image does not contain an alpha channel but a CG character rendered against a black background. This is common for different render passes like specular or shadows, for example. The image is connected to a Set Canvas Color node, which sets the Canvas color to black with a solid alpha. The Set Canvas Color then connects to the Auto Domain node, which detects the pixels and sets the DoD. If the original image contained a premultiplied alpha channel, the Set Canvas Color would not be needed, and the image could be connected directly into the Auto Domain node.



An Auto Domain node automatically limits the area of image processing.

Inspector



The Auto Domain Controls tab

Controls Tab

In most cases, the Auto Domain node automatically calculates the DoD bounding box; however, the rectangular shape can be modified using the Controls tab in the Inspector.

Left

Defines the left border of the search area of the ADoD. Higher values on this slider move the left border toward the right, excluding more data from the left margin.

1 represents the right border of the image; 0 represents the left border. The slider defaults to 0 (left border).

Bottom

Defines the bottom border of the search area of the ADoD. Higher values on this slider move the bottom border toward the top, excluding more data from the bottom margin.

1 represents the top border of the image; 0 represents the bottom border. The slider defaults to 0 (bottom border).

Right

Defines the right border of the search area of the ADoD. Higher values on this slider move the right border toward the left, excluding more data from the right margin.

1 represents the right border of the image; 0 represents the left border. The slider defaults to 1 (right border).

Top

Defines the top border of the search area of the ADoD. Higher values on this slider move the top border toward the bottom, excluding more data from the top margin.

1 represents the top border of the image; 0 represents the bottom border. The slider defaults to 1 (top border).

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other miscellaneous nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Change Depth [CD]



The Change Depth node

Change Depth Node Introduction

The Change Depth node has one simple use, and that is to change the bits per color channel used to process a node. This node is often used after color correcting 32-bit floating-point image files, converting them from float processing to 16-bit per channel to preserve memory and performance.

It can also be useful if, from a certain point in your node tree, you feel the need to process your images in a higher bit depth than their original one or to reduce the bit depth to save memory.

Inputs

The single input on the Change Depth node is used to connect a 2D image and an effect mask, which can be used to limit the blurred area.

Input: The orange input is used for the primary 2D image to be converted.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes or bitmaps from other tools. Connecting a mask to this input limits the blur to only those pixels within the mask.

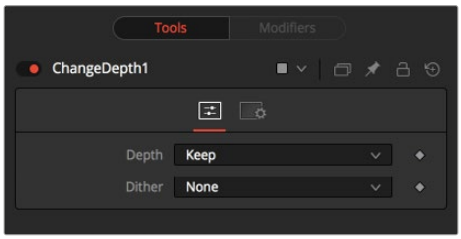
Basic Node Setup

Below, a Loader node in Fusion Studio is color corrected and then down converted from a floating-point image to a 16-bit image to save image-processing time and memory.



A Change Depth node placed after color correction is done on a floating-point image.

Inspector



The Change Depth Controls tab

Controls Tab

The two controls for this node are the Depth menu and the Dither menu. These two menus are used to convert and adjust the color depth of the image.

Depth

The Keep setting doesn't do anything to the image but instead keeps the input depth. The other options change the bit depth of the image to the respective value.

Dither

When down converting from a higher bit depth, it can be useful to add Error Diffusion or Additive Noise to camouflage artifacts that result from problematic (high-contrast) areas.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other miscellaneous nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Custom Tool [CT]



The Custom Tool node

Custom Tool Node Introduction

The Custom Tool node is the most complex and the most powerful node in Fusion. It is used to create custom expressions and filters to modify an image. In addition to providing three image inputs, the Custom Tool node allows for the connection of up to eight numeric inputs and as many as four XY position values from other controls and parameters in the node tree.

Per-pixel calculations can be performed on the Red, Green, Blue, Alpha, Z, Z-Coverage, UV texture coords, XYZ Normals, RGBA background color, and XY motion vector channels of the images.

You should be moderately experienced with scripting, or C++ programming, to understand the structure and terminology used by the Custom Tool node.

Inputs

The Custom Tool node has three image inputs, a matte input, and an effect mask input.

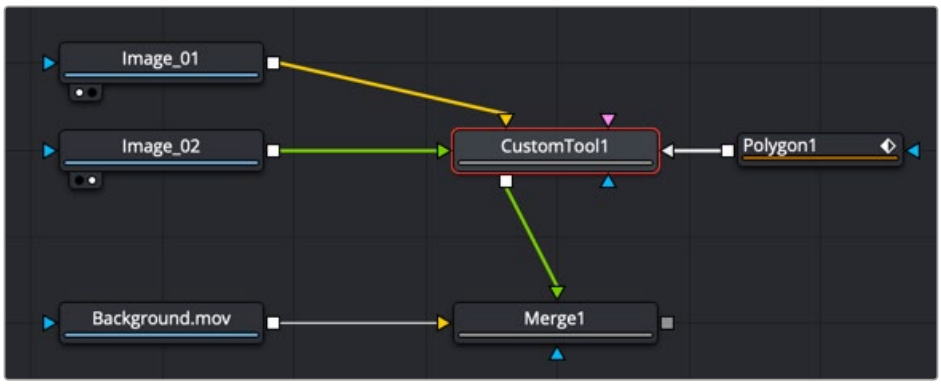
Input: The orange, green, and magenta inputs combine 2D images to make your composite. When entering them into the Custom Tool fields, they are referred to as c1, c2 and c3 (c standard for all three R, G, B channels)

Matte Input: The white input is for a matte created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a matte to this input allows a matte to be combined into any equation. When entering the matte into the Custom Tool fields, it is referred to as m1.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the Custom Tool effect to only those pixels within the mask.

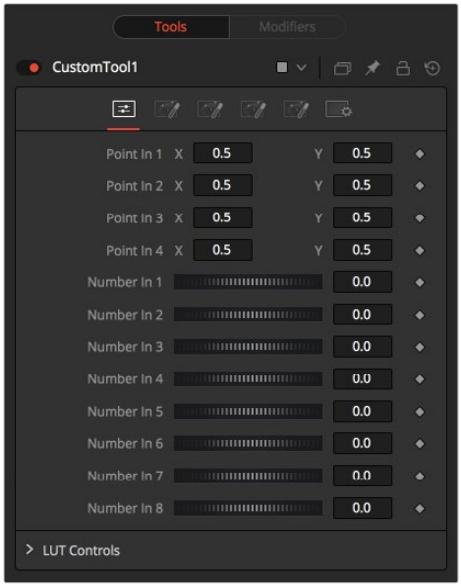
Basic Node Setup

The Custom Tool below takes two image inputs and a matte input, and then combines them using some calculation. The result can be output to a Merge or other image-processing nodes.



A Custom Tool is used to build your own effects using C++ and scripting.

Inspector



The Custom Tool Controls tab

Controls Tab

Point in 1-4, X and Y

These four controls are 2D X and Y center controls that are available to expressions entered in the Setup, Intermediate, and Channels tabs as variables p1x, p1y, ..., p4x, p4y. They are normal positional controls and can be animated or connected to modifiers as any other node might.

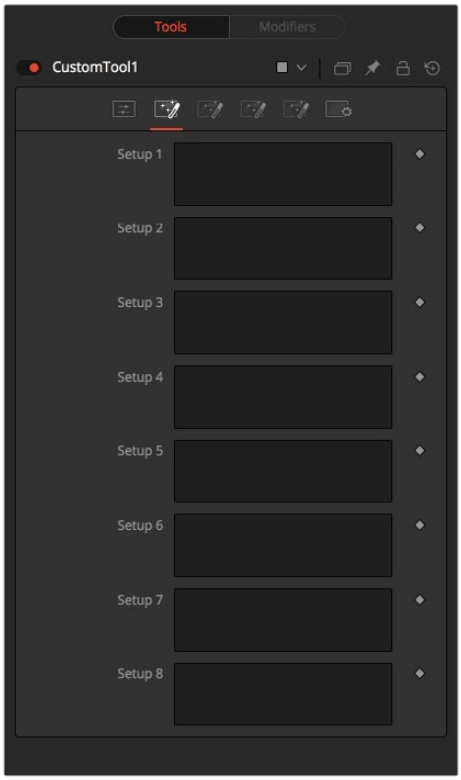
Number in 1-8

The values of these controls are available to expressions entered in the Setup, Intermediate, and Channels tabs as variables n1, n2, n3, ..., n8. They are normal slider controls and can be animated or connected to modifiers exactly as any other node might.

LUT in 1-4

The Custom Tool node provides 4 LUT splines. The values of these controls are available to expressions entered in the Setup, Intermediate, and Channels tabs using the getlut# function. For example, setting the R, G, B, and A expressions to getlut1(r1), getlut2(g1), getlut3(b1), and getlut4(a1), respectively, would cause the Custom Tool node to mimic the Color Curves node.

These controls can be renamed using the options in the Config tab to make their meanings more apparent, but expressions still see the values as n1, n2, ..., n8.



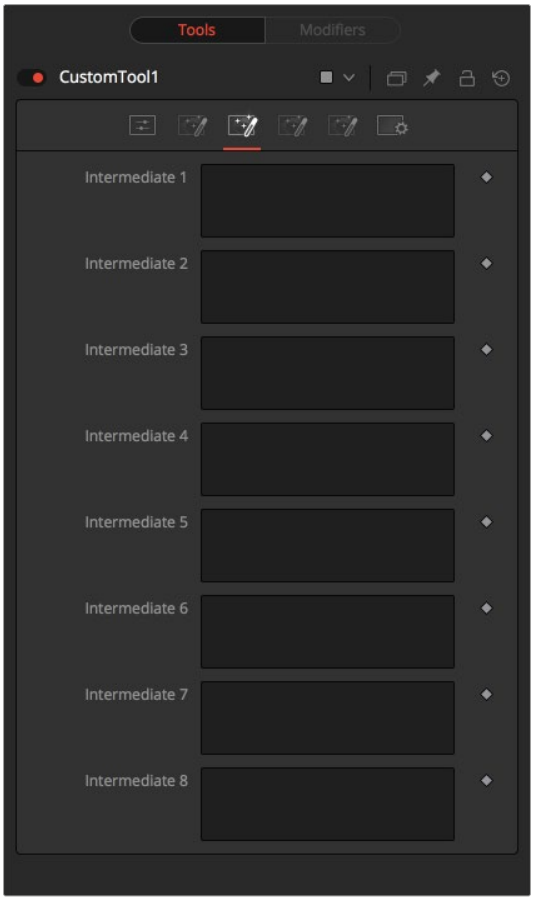
The Custom Tool Setup tab

Custom Tool Setup Tab

Setup 1-4

Up to four separate expressions can be calculated in the Setup tab of the Custom Tool node. The Setup expressions are evaluated once per frame before any other calculations are performed. The results are then made available to the other expressions in the Custom Tool node as variables s1, s2, s3, and s4.

NOTE: Because these expressions are evaluated once per frame only and not for each pixel, it makes no sense to use per-pixel variables like X and Y or channel variables like r1, g1, b1. Allowable values include constants, variables such as n1..n8, time, W and H, and functions like sin() or getr1d().

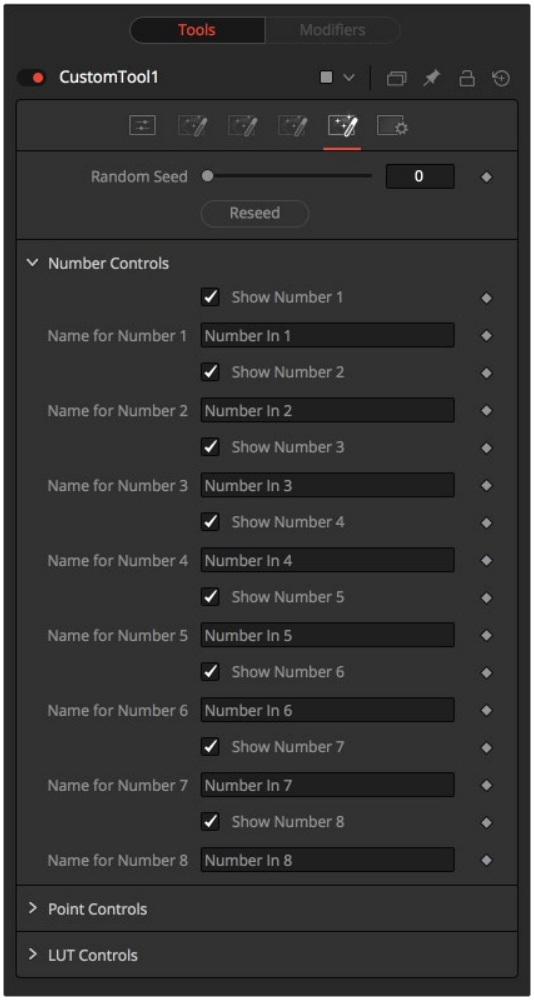


The Custom Tool Intermediate tab

Custom Tool Intermediate Tab

Intermediate 1-4

An additional four expressions can be calculated in the Inter tab. The Inter expressions are evaluated once per pixel after the Setup expressions are evaluated but before the Channel expressions are evaluated. Per-pixel channel variables like r1, g1, b1, and a1 are allowable. Results are available as variables i1, i2, i3, and i4.



The Custom Tool Config tab

Custom Tool Config Tab

Random Seed

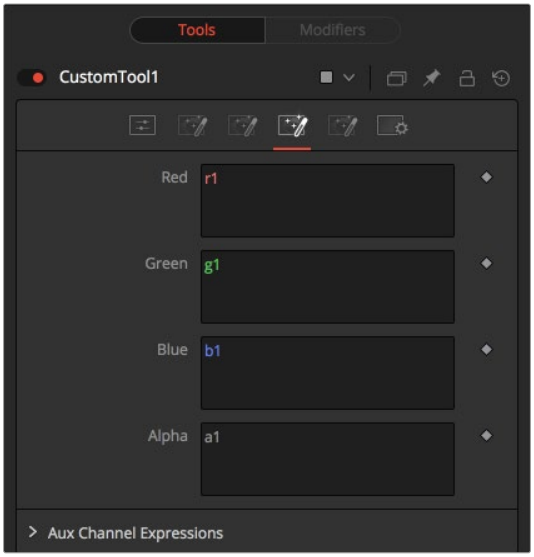
Use this to set the seed for the rand() and rands() functions. Click the Randomize button to set the seed to a random value. This control may be needed if multiple Custom Tool nodes are required with different random results for each.

Number Controls

There are eight sets of Number controls, corresponding to the eight Number In sliders in the Controls tab. Uncheck the Show Number checkbox to hide the corresponding Number In slider, or edit the Name for Number text field to change its name.

Point Controls

There are four sets of Point controls, corresponding to the four Point In controls in the Controls tab. Uncheck the Show Point checkbox to hide the corresponding Point In control and its crosshair in the viewer. Similarly, edit the Name for Point text field to change the control's name.



The Custom Tool Channel tab

Custom Tool Channels Tab

RGBA, Z, UV Expressions, and XYZ Normal Expressions

The Channel tab is used to set up one expression per each available channel of the image. Each expression is evaluated once per pixel. The result creates the value for that pixel in the output of the image.

Color Channel expressions (RGBA) should generally return floating-point values between 0.0 and 1.0. Values beyond this are clipped if the destination image is an integer. Other expression fields should produce values appropriate to their channel (e.g., between -1.0 and 1.0 for Vector and Normal fields, 0.0 to 1.0 for Coverage, or any value for Depth). The Channel expressions may use the results from both the Setup expressions (as variables s1–s4) and Inter expressions (as variables i1–i4).

Custom Tool Node Syntax

Value Variables	
`n1..n8`	Numeric Inputs
`p1x..p4x`	Position Values (X-axis)
`p1y..p4y`	Position Values (Y-axis)
`s1..s4`	Setup Expression Results
`i1..i4`	Inter Expression Results
`time`	Current Frame
`x`	Horizontal co-ordinate of the current pixel, between 0.0 and 1.0
`y`	Vertical co-ordinate of the current pixel, between 0.0 and 1.0
`w (or w1..w3)`	Width of Image (for image1..image3)
`h (or h1..h3)`	Height of Image (for image1..image3)

Value Variables	
`ax (or ax1..ax3)`	Image Aspect X (for image1..image3)
`ay (or ay1..ay3)`	Image Aspect Y (for image1..image3)

NOTE: Use w and h and ax and ay without a following number to get the dimensions and aspect of the primary image.

Channel (Pixel) Variables	
`c1..c3`	Current Channel (for image1..image3)
`r1..r3`	Red (for image1..image3)
`g1..g3`	Green (for image1..image3)
`b1..b3`	Blue (for image1..image3)
`a1..a3`	Alpha (for image1..image3)
`z1..z3`	Z-Buffer (for image1..image3)
`cv1..cv3`	Z Coverage (for image1..image3)
`u1..u3`	U Coordinate (for image1..image3)
`v1..v3 nx1..nx3`	V Coordinate (for image1..image3) X Normal (for image1..image3)
`ny1..ny3`	Y Normal (for image1..image3)
`nz1..nz3`	Z Normal (for image1..image3)
`bgr1..bgr3`	Background Red (for image1..image3)
`bgg1..bgg3`	Background Green (for image1..image3)
`bgb1..bgb3`	Background Blue (for image1..image3)
`bga1..bga3`	Background Alpha (for image1..image3)
`vx1..vx3`	X Vector (for image1..image3)
`vy1..vy3`	Y Vector (for image1..image3)
`nz1..nz3`	Z Normal (for image1..image3)

NOTE: Use c1, c2, c3 to refer to the value of a pixel in the current channel. This makes copying and pasting expressions easier. For example, if c1/2 is typed as the red expression, the result would be half the value of the red pixel from image 1, but if the expression is copied to the blue channel, now it would have the value of the pixel from the blue channel.

To refer to the red value of the current pixel in input 1, type r1. For the image in input 2, it would be r2.

- `get[ch][#]b(x, y)` Read pixel at x,y, or 0 if out of bounds—e.g., getr1b(0,0)
- `get[ch][#]d(x, y)` Read pixel at x,y or edge pixel if out of bounds—e.g., getr1d(0,0)
- `get[ch][#]w(x, y)` Read pixel at x,y or wrap if out of bounds—e.g., getr1w(0,0)

NOTE: There are a variety of methods used to refer to pixels from locations other than the current one in an image.

In the above description, [ch] is a letter representing the channel to access. The [#] is a number representing the input image. So to get the red component of the current pixel (equivalent to r), you would use getr1b(x,y). To get the alpha component of the pixel at the center of image 2, you would use geta2b(0.5, 0.5).

- `getr1b(x,y)` Output the red value of the pixel at position x, y, if there were a valid pixel present. It would output 0.0 if the position were beyond the boundaries of the image (all channels).
- `getr1d(x,y)` Output the red value of the pixel at position x, y. If the position specified were outside of the boundaries of the image, the result would be from the outer edge of the image (RGBA only).
- `getr1w(x,y)` Output the red value of the pixel at position x, y. If the position specified were outside of the boundaries of the image, the x and y coordinates would wrap around to the other side of the image and continue from there (RGBA only).

To access other channel values with these functions, substitute the r in the above examples with the correct channel variable (r, g, b, a and, for the getr1b() functions only, z, and so on), as shown above. Substitute the 1 with either 2 or 3 in the above examples to access the images from the other image inputs.

Mathematical Expressions	
`pi`	The value of pi
`e`	The value of e
`log(x)`	The base-10 log of x
`ln(x)`	The natural (base-e) log of x
`sin(x)`	The sine of x (x is degrees)
`cos(x)`	The cosine of x (x is degrees)
`tan(x)`	The tangent of x (x is degrees)
`asin(x)`	The arcsine of x, in degrees

Mathematical Expressions	
`acos(x)`	The arccosine of x, in degrees
`atan(x)`	The arctangent of x, in degrees
`atan2(x,y)`	The arctangent of x,y, in degrees
`abs(x)`	The absolute (positive) value of x
`int(x)`	The integer (whole) value of x
`frac(x)`	The fractional value of x
`sqrt(x)`	The Square Root of x
`rand(x,y)`	A random value between x and y
`rands(x,y,s)`	A random value between x and y, based on seed s
`min(x,y)`	The minimum (lowest) of x and y
`max(x,y)`	The maximum (highest) of x and y
`dist(x1,y1,x2,y2)`	The distance between point x1,y2 and x2,y2
`dist3d(x1,y1,z1,x2,y2,z2)`	The distance between 3D points x1,y2,z1 and x2,y2,z2
`noise(x)`	A smoothly varying Perlin noise value based on x
`noise2(x, y)`	A smoothly varying Perlin noise value based on x and y
`noise3(x, y, z)`	A smoothly varying Perlin noise value based on x, y and z
`if(c, x, y)`	returns x if c not 0, otherwise y

Mathematical Operators	
`!x`	1.0 if x = 0, otherwise 0.0
`-x`	(0.0 - x)
`+x`	(0.0 + x) i.e. effectively does nothing
`x ^ y`	x raised to the power of y
`x * y`	x multiplied by y
`x / y`	x divided by y
`x % y`	x modulo y, i.e. remainder of (x divided by y)
`x + y`	x plus y
`x - y`	x minus y
`x < y`	1.0 if x is less than y, otherwise 0.0
`x > y`	1.0 if x is greater than y, otherwise 0.0

Mathematical Operators	
`x <= y`	1.0 if x is less than or equal to y, otherwise 0.0
`x >= y`	1.0 if x is greater than or equal to y, otherwise 0.0
`x = y`	1.0 if x is exactly equal to y, otherwise 0.0
`x == y`	1.0 if x is exactly equal to y, otherwise 0.0, identical to above
`x <> y`	1.0 if x is not equal to y, otherwise 0.0
`x != y`	1.0 if x is not equal to y, otherwise 0.0, i.e. identical to above
`x & y`	1.0 if both x and y are not 0.0, otherwise 0.0
`x && y`	1.0 if both x and y are not 0.0, otherwise 0.0, i.e. identical to above
`x|y`	1.0 if either x or y (or both) are not 0.0, otherwise 0.0
`x||y`	1.0 if either x or y (or both) are not 0.0, otherwise 0.0

Examples

The following examples are intended to help you understand the various components of the Custom Tool node.

Rotation

To rotate an image, we need the standard equations for 2D rotation:

```
x' = x * cos(theta) - y * sin(theta) y' = x * sin(theta) + y * cos(theta)
```

Using the n1 slider for the angle theta, and a sample function, we get (for the red channel):

```
getr1b(x * cos(n1) - y * sin(n1), x * sin(n1) + y * cos(n1))
```

This calculates the current pixel's (x,y) position rotated around the origin at (0,0) (the bottom-left corner), and then fetches the red component from the source pixel at this rotated position. For centered rotation, we need to subtract 0.5 from our x and y coordinates before we rotate them, and add 0.5 back to them afterward:

```
getr1b((x-.5) * cos(n1) - (y-.5) * sin(n1) + .5, (x-.5) * sin(n1) + (y-.5) * cos(n1) + .5)
```

Which brings us to the next lesson: Setup and Intermediate Expressions. These are useful for speeding things up by minimizing the work that gets done in the Channel expressions. The Setup expressions are executed only once, and their results don't change for any pixel, so you can use these for s1 and s2, respectively.

```
cos(n1) sin(n1)
```

Intermediate expressions are executed once for each pixel, so you can use these for i1 and i2:

```
(x-.5) * s1 - (y-.5) * s2 + .5
```

```
(x-.5) * s2 + (y-.5) * s1 + .5
```

These are the x and y parameters for the getr1b() function from above, but with the Setup results, s1 and s2, substituted so that the trig functions are executed only once per frame, not every pixel. Now you can use these intermediate results in your Channel expressions:

```
getr1b(i1, i2)

getg1b(i1, i2)

getb1b(i1, i2)

geta1b(i1, i2)
```

With the Intermediate expressions substituted in, we only have to do all the additions, subtractions, and multiplications once per pixel, instead of four times. As a rule of thumb, if it doesn‘t change, do it only once.

This is a simple rotation that doesn’t take into account the image aspect at all. It is left as an exercise for you to include this (sorry). Another improvement could be to allow rotation around points other than the center.

Filtering

Our second example duplicates the functionality of a 3 x 3 Custom Filter node set to average the current pixel together with the eight pixels surrounding it. To duplicate it with a Custom Tool node, add a Custom Tool node to the node tree, and enter the following expressions into the Setup tab.

(Leave the node disconnected to prevent it from updating until we are ready.)

```
S1

1.0/w1

S2

1.0/h1
```

These two expressions are evaluated at the beginning of each frame. S1 divides 1.0 by the current width of the frame, and S2 divides 1.0 by the height. This provides a floating-point value between 0.0 and 1.0 that represents the distance from the current pixel to the next pixel along each axis.

Now enter the following expression into the first text control of the Channel tab (r).

```
(getr1w(x-s1, y-s2) + getr1w(x, y-s2) + getr1w(x+s1, y-s2) + getr1w(x+s1,
y) + getr1w(x-s1, y) + r1 +getr1w(x-s1, y+s2) + getr1w(x, y+s2) +
getr1w(x+s1, y+s2)) / 9
```

This expression adds together the nine pixels above the current pixel by calling the getr1w() function nine times and providing it with values relative to the current position. Note that we referred to the pixels by using x+s1, y+s2, rather than using x+1, y+1.

Fusion refers to pixels as floating-point values between 0.0 and 1.0, which is why we created the expressions we used in the Setup tab. If we had used x+1, y+1 instead, the expression would have sampled the same pixel over and over again. (The function we used wraps the pixel position around the image if the offset values are out of range.)

That took care of the red channel; now use the following expressions for the green, blue, and alpha channels.

```
(getg1w(x-s1, y-s2) + getg1w(x, y-s2) + getg1w(x+s1, y-s2) + getg1w(x+s1,
y) + getg1w(x-s1, y) + g1 +getg1w(x-s1, y+s2) + getg1w(x, y+s2) +
getg1w(x+s1, y+s2)) / 9

(getb1w(x-s1, y-s2) + getb1w(x, y-s2) + getb1w(x+s1, y-s2) + getb1w(x+s1,
y) + getb1w(x-s1, y) + b1 +getb1w(x-s1, y+s2) + getb1w(x, y+s2) +
getb1w(x+s1, y+s2)) / 9

(geta1w(x-s1, y-s2) + geta1w(x, y-s2) + geta1w(x+s1, y-s2) + geta1w(x+s1,
y) + geta1w(x-s1, y) + a1 + geta1w(x-s1, y+s2) + geta1w(x, y+s2) +
geta1w(x+s1, y+s2)) / 9
```

It's time to view the results. Add a Background node set to a solid color and change the color to a pure red. Add a hard-edged Rectangular effects mask and connect it to the expression just created.

For comparison, add a Custom Filter node and duplicate the settings from the image above. Connect a pipe to this node from the background to the node and view the results. Alternate between viewing the Custom Tool node and the Custom Filter while zoomed in close to the top corners of the effects mask.

Of course, the Custom Filter node renders a lot faster than the Custom Tool node we created, but the flexibility of the Custom Tool node is its primary advantage. For example, you could use an image connected to input 2 to control the median applied to input one by changing all instances of getr1w, getg1w, and getb1w in the expression to getr2w, getg2w, and getb2w, but leaving the r1, g1, and b1s as they are.

This is just one example; the possibilities of the Custom Tool node are limitless.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other miscellaneous nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Fields [FLDs]



The Fields node

Fields Node Introduction

The Fields node is a robust multipurpose utility offering several functions related to interlaced video frames. It interpolates video fields into frames or video frames into fields. Although the interlace preference and method type is defined in the MediaIn or Loader, and generators, this node can be used to assist in the standards conversion of PAL to NTSC and provides the ability to process fields and frames for specific portions of a node tree.

This node can also interlace two separate images together into a single interlace image. The background input is the dominate field 1 and the foreground is field 2.

Inputs

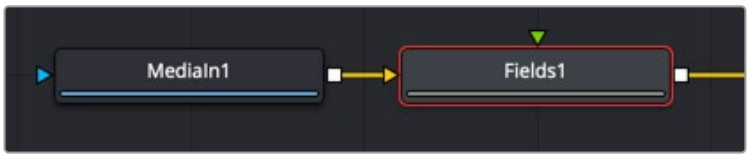
The single input on the Fields node is used to connect a 2D image and an effect mask, which can be used to limit the blurred area.

Stream1 Input: The orange background input is used for the primary 2D image that is interpolated or converted.

Stream2 Input: The optional green foreground input is only used when merging two interlaced images together.

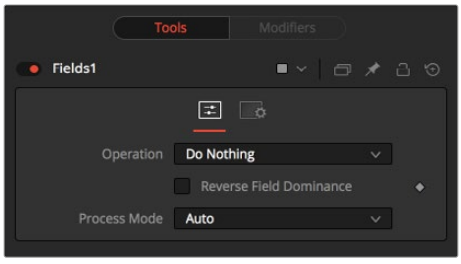
Basic Node Setup

The Fields node below is used to convert the background image from a PAL interlaced format to progressive frames.



The background image can be modified to various interlaced formats.

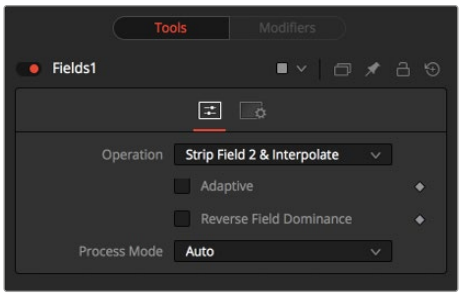
Inspector



The Fields Controls tab

Controls Tab

The Controls tab includes two menus. The Operation menu is used to select the type of field conversion performed. The Process Mode menu is used to select the field's format for the output image.



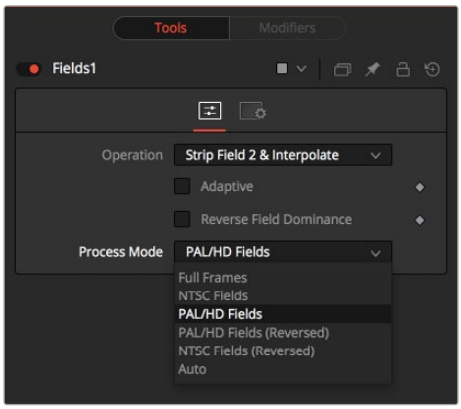
Operatiion Controls

Operation Menu

- **Do Nothing:** This causes the images to be affected by the Process Mode selection exclusively.
- **Strip Field 2:** This removes field 2 from the input image stream, which shortens the image to half of the original height.
- **Strip Field 1:** This removes field 1 from the input image stream, which shortens the image to half of the original height.
- **Strip Field 2 and Interpolate:** This removes field 2 from the input image stream and inserts a field interpolated from field 1 so that image height is maintained. Should be supplied with frames, not fields.
- **Strip Field 1 and Interpolate:** This removes field 1 from the input image stream and inserts a field interpolated from field 2 so that image height is maintained. Should be supplied with frames, not fields.
- **Interlace:** This combines fields from the input image stream(s). If supplied with one image stream, each pair of frames are combined to form half of the number of double-height frames. If supplied with two image streams, single frames from each stream are combined to form double-height images.
- **De-Interlace:** This separates fields from one input image stream. This will produce double the amount of half-height frames.

Reverse Field Dominance

When selected, the Field Order or Dominance of the image will be swapped.



Process Mode

Process Mode Menu

- **Full Frames:** This forces Frame Processing. Useful for processing frames in a part of a node tree that is otherwise field processing.
- **NTSC Fields:** This forces NTSC Field Processing. Useful for processing fields in a part of a node tree that is otherwise frame processing.
- **PAL Fields:** This forces PAL Field Processing. Useful for processing fields in a part of a node tree that is otherwise frame processing.
- **PAL Fields (Reversed):** This forces PAL-swapped Field Processing.
- **NTSC Fields (Reversed):** This forces NTSC-swapped Field Processing.
- **Auto:** This attempts to match the mode of its input images. Fields are used if the input types are mixed.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other miscellaneous nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Frame Average [Avg]



The Frame Average node

Frame Average Node Introduction

The Frame Average node averages together a series of frames to simulate clips shot with long shutter speeds. Aside from motion blur-style effects, it can be useful for time warps or noise removal.

Inputs

The single input on the Frame Average node is used to connect a 2D image that will have the averaging applied.

Input: The orange input is used for the primary 2D image that will be averaged.

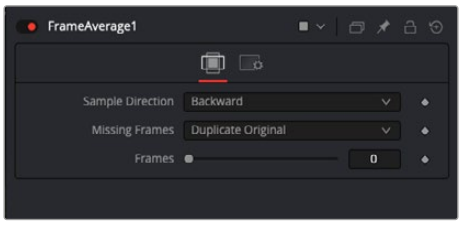
Basic Node Setup

The image connected to the orange input is frame averaged based on the settings in the Inspector.



A Frame Average node blends the input image’s frames.

Inspector



The Frame Average Controls tab

Controls Tab

The Controls tab contains the parameters for setting the duration and guidance of the averaged frames.

Sample Direction

The Sample Direction menu determines if the averaged frames are taken before the current frame, after, or a mix of the two.

- **Forward:** Averages the number of frames set by the Frames slider after the current frame.
- **Both:** Averages the number of frames set by the Frames slider, taking frames before and after the current frame.
- **Backward:** Averages the number of frames set by the Frames slider before the current frame.

Missing Frames

This control determines the behavior if a frame is missing from the clip.

- **Duplicate Original:** Uses the last original frame until a new frame is available.
- **Blank Frame:** Leaves missing frames blank.

Frames

This slider sets the number of frames that are averaged.

Keyframe Stretcher [KfS]



The Keyframe Stretcher node

Keyframe Stretcher Node Introduction

The Keyframe Stretcher node is inserted after animated nodes, so the keyframes stretch and the comp's duration is modified. It is used to scale the keyframes on the animation curve to the current duration of the clip. This is particularly useful when creating title templates in Fusion for use in DaVinci Resolve's Edit or Cut page.

TIP: The Keyframe Stretcher can be used on a single parameter by applying the Keystretcher modifier.

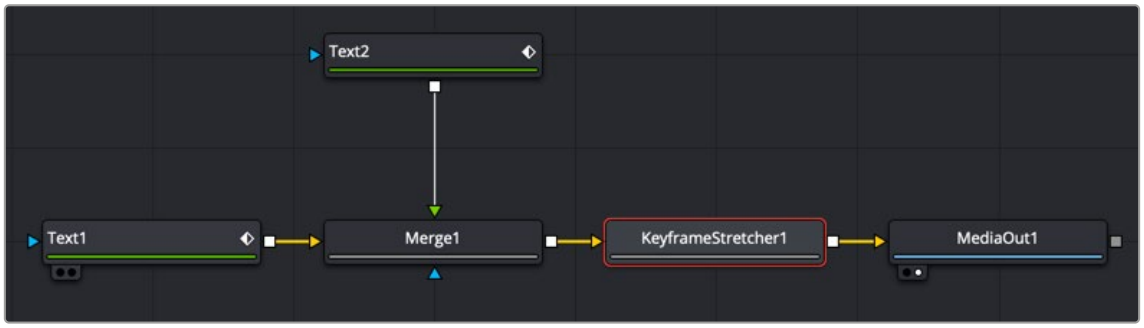
Inputs

The single input on the Keyframe Stretcher node is used to connect a 2D image that contains keyframe animation.

Input: The orange input is used for any node with keyframed animation. The input can be a Merge node that is not animated but contains foreground and background nodes that are animated.

Basic Node Setup

The Keyframe Stretcher is added just before the Media Out or Saver node. All nodes that include animation before the Keyframe Stretcher are modified if the comp changes duration.



A Keyframe Stretcher changing the animation of two text nodes.

The diagram below shows the original 50-frame animation added to a parameter. The Keyframe Stretcher Start and End would be set to 0 and 50. The second keyframe is set at frame 10, and the third keyframe is set at frame 40. Setting the Stretch Start to frame 11 and the Stretch End to frame 39 will keep the existing keyframes at the same speed (number of frames.) The middle will be stretched.



Original 50-frame animation

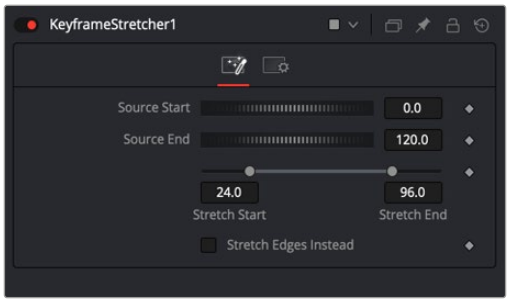
In the below example, the duration of the clip is extended to 75 frames. The first 10 frames and the last 10 frames of the animation run at the same speed as the original animation, while any animation in the middle is stretched to fill the difference.

NOTE: The actual Spline Editor will show only the original keyframe positions. The splines are not changed by the Keyframe Stretcher; only the animation is changed.



Animation modified to 75 frames but stretching only the middle of the animation

Inspector



The Keyframe Stretcher Keyframes tab

Keyframes Tab

The Keyframes tab includes Source controls for setting the source duration and Stretch controls for setting the area of the animation that gets modified.

Source Start/Source End

A source range is specified using the Source Start and Source End controls. These are typically set to match the full range of the animation spline on the Keyframes control.

Stretch Start/Stretch End

The Stretch Start and Stretch End controls let you specify a middle zone where keyframes will be stretched or squished. Handles outside the range will not get scaled. Any keyframes outside the Stretch Start and End range always remain the same number of frames from the Start and End.

Any keyframe adjustments to the original control will be correspondingly scaled back to the source curve and will match the original timing as expected.

Stretch Edges Instead

Enabling the Stretch Edges Instead checkbox overrides the Stretch Start and Stretch End controls and stretches the edges of the animation.

Run Command [Run]



The Run Command node

Run Command Node Introduction

The Run Command node is used to execute an external command or batch file at specific points during a render. You can choose to run a command at the start or the end of a render. Or you can have the command execute once for each frame.

The Run Command can be used to net render other command line applications using the Fusion Render Manager, as well as a host of other useful functions.

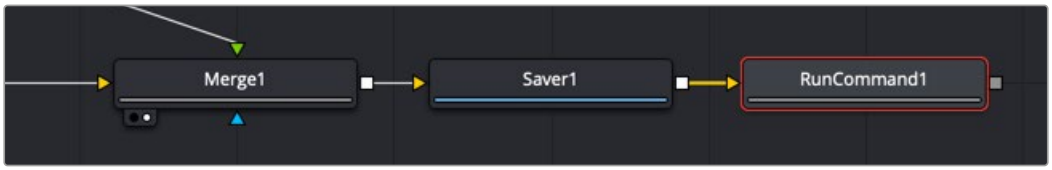
Inputs

The single input on the Run Command node is used to pass through a 2D image.

Input: The optional orange image input is not required for this node to operate. However, if it is connected to a node's output, the Run Command will only launch after the connected node has finished rendering. This is often useful when connected to a Saver, to ensure that the output frame has been fully saved to disk first. If the application launched returns a non-zero result, the node will also fail.

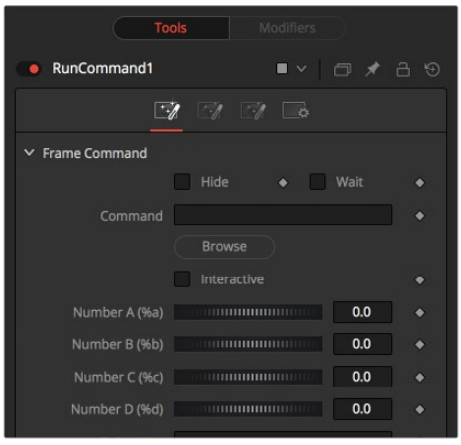
Basic Node Setup

The Run Command node can be connected after a Saver and run once the final frame is completed.



Run Command placed after a Saver node.

Inspector



The Run Command Frame tab

Frame Tab

The Frame tab is where the command to execute is selected and modified.

Hide

Enable the Hide checkbox to prevent the application or script from displaying a window when it is executed.

Wait

Enable this checkbox to cause the node to wait for a remote application or tool to exit before continuing. If this checkbox is disabled, the Fusion continues rendering without waiting for the external application.

Frame Command

This field is used to specify the path for the command to be run after each frame is rendered. The Browse button can be used to identify the path.

Interactive

This checkbox determines whether the launched application should run interactively, allowing user input.

Number A (%a) and Number B (%b)

Various wildcards can be used with the frame commands; these wildcards will be substituted at render time with the correct values.

- **%a:** Outputs the number from the Number A thumbwheel control.
- **%b:** Outputs the number from the Number B thumbwheel control.
- **%t:** Outputs the current frame number (without zero padding).
- **%s:** Substitutes using the text from the large text entry field.

If you want to add zero paddings to the numbers generated by %t, refer to the wildcard with %0x, where x is the number of characters with which to pad the value. This also works for %a and %b.

For example, test%04t.tga would return the following values at render time:

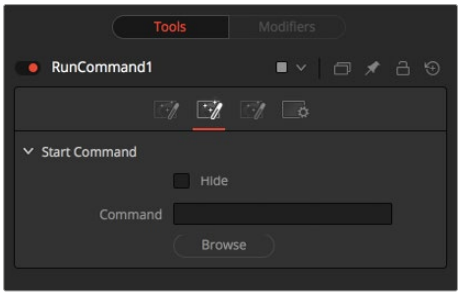
```
test0000.tga

test0001.tga

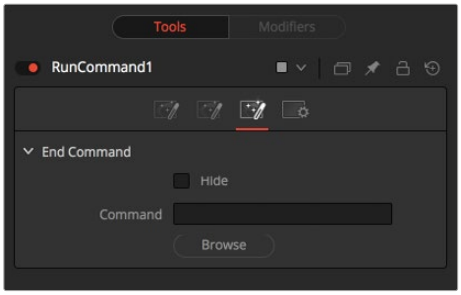
test0009.tga

test0010.tga
```

You may also pad a value with spaces by calling the wildcard as %x, where x is the number of spaces with which you would like to pad the value.



The Run Command Start tab



The Run Command End tab

Start and End Tabs

The Start and End tabs contain a file browser for a command to be run when the composition starts to render and when the composition is done rendering.

Example

To copy the saved files from a render to another directory as each frame is rendered, save the following text in a file called copyfile.bat to your C\ directory (the root folder).

```
@echo off
set parm=%1 %2
copy %1 %2 set parm=
```

Create or load any node tree that contains a Saver. The following example assumes a Saver is set to output D\ test0000.exr, test0001.exr, etc. You may have to modify the example to match.

Add a Run Command node after the Saver to ensure the Saver has finished saving first. Now enter the following text into the Run Command node's Frame Command text box:

```
C\copytest.bat D\test%04f.exr C\
```

Select the Hide Frame command checkbox to prevent the Command Prompt window from appearing briefly after every frame.

When this node tree is rendered, each file will be immediately copied to the C\ directory as it is rendered.

The Run Command node could be used to transfer the files via FTP to a remote drive on the network, to print out each frame as it is rendered, or to execute a custom image-processing tool.

The Run Command node is not restricted to executing simple batch files. FusionScript, VBScript, Jscript, CGI, and Perl files could also be used, as just a few examples.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other miscellaneous nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Set Domain [DOD]



The Set Domain node

Set Domain Node Introduction

Set Domain is used to adjust or set the active area of an image or the area of the image considered to have valid data.

It does not change the image's physical dimensions. Downstream nodes will not process anything outside the Domain of Definition (DoD), thus speeding up rendering of computation-intensive nodes.

This node provides an absolute mode, for setting the domain of definition manually, and a relative mode for adjusting the existing domain of definition.

Inputs

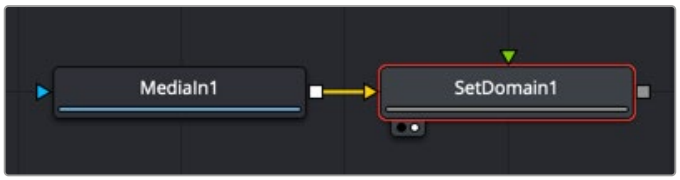
The two inputs on the Set Domain node are used to connect 2D images.

Input: The orange background input must be connected. It accepts a 2D image with the DoD you want to replace or adjust.

Foreground: The green image input is optional but also accepts a 2D image as its input. When the foreground input is connected, the Set Domain node will replace the Background input's domain of definition with the foreground's DoD.

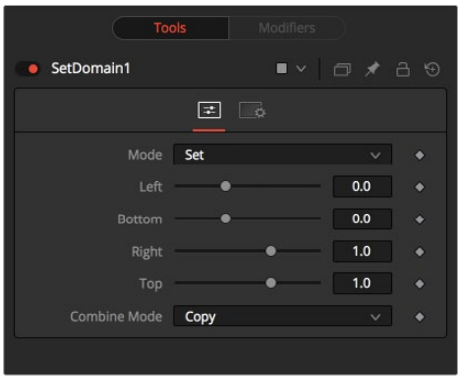
Basic Node Setup

The example below assumes an image is connected to a Set Domain node to manually configure the DoD.

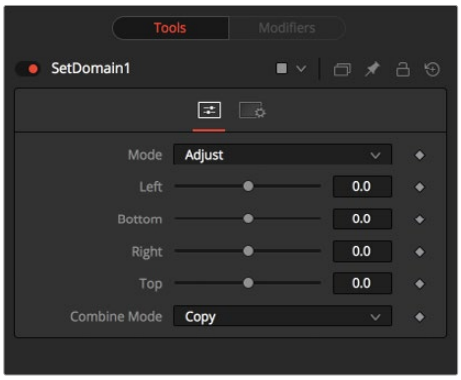


A Set Domain node manually sets the area to limit image processing.

Inspector



The Set Domain Controls tab in Set mode



The Set Domain Controls tab in Adjust mode

Controls Tab

Mode

The Mode menu has two choices depending on whether you want to adjust or offset the existing domain or set precise values for it.

The same operations can be performed in Set or in Adjust mode. In Adjust mode, the sliders default to 0, marking their respective full extent of the image. Positive values shrink the DoD while negative values expand the DoD to include more data.

Set mode defaults to the full extent of the visible image. Sliders default to a scale of 0-1 from left to right and bottom to top.

Left

Defines the left border of the DoD. Higher values on this slider move the left border toward the right, excluding more data from the left margin.

1 represents the right border of the image; 0 represents the left border. The slider defaults to 0 (left border).

Bottom

Defines the bottom border of the DoD. Higher values on this slider move the bottom border toward the top, excluding more data from the bottom margin.

1 represents the top border of the image; 0 represents the bottom border. The slider defaults to 0 (bottom border).

Right

Defines the right border of the DoD. Higher values on this slider move the right border toward the left, excluding more data from the right margin.

1 represents the right border of the image; 0 represents the left border. In Set mode, the slider defaults to 1 (right border).

Top

Defines the top border of the DoD. Higher values on this slider move the top border toward the bottom, excluding more data from the top margin.

1 represents the top border of the image; 0 represents the bottom border. In Set mode, the slider defaults to 1 (top border).

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other miscellaneous nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Switch [Swi]



The Switch node

Switch Node Introduction

Switch is a tool that enables artists to alternate between multiple input sources, allowing the selection of one output from the chosen inputs. This capability is especially useful for toggling between different visual elements within a composition. The number of input connections can vary, with the node supporting a dynamic number of inputs that can be renamed as needed. Controls in the Config tab allow users to add or remove inputs and change their names.

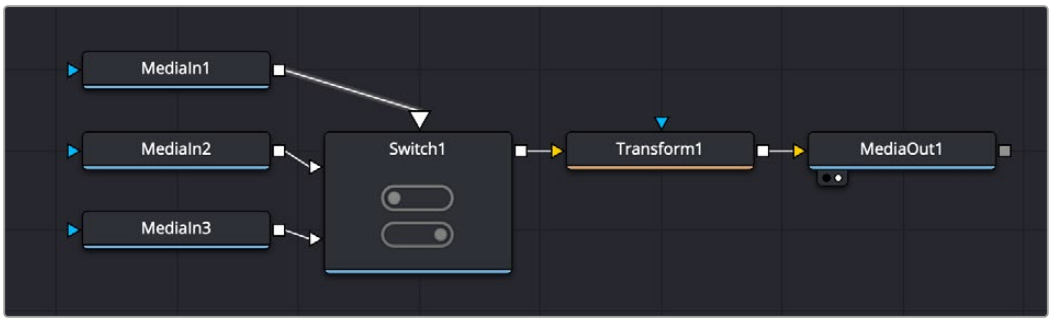
Switch works with most types of tools, including Shapes and 3D. Additionally, users can find the Switch modifier in the Modify With and Insert submenus of a control’s context menu, allowing it to be applied to any supported control.

Inputs

You can add up to nine inputs to a Switch Node using the slider in the Config tools. If you need more than nine sources, you can type that number manually into the Number of Inputs field.

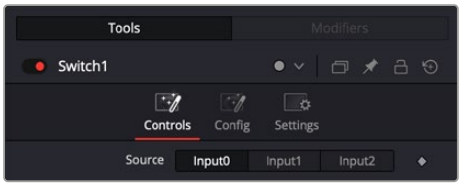
Basic Node Setup

Connect as many input sources as you wish to the Switch Node, and connect its output to the rest of the node tree.



The Switch tool receiving three media inputs, allowing you to choose one to go through the transform node

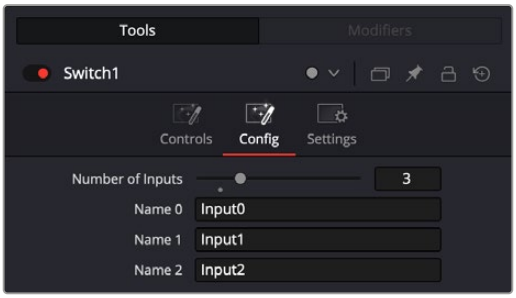
Inspector



The Switch Controls let you choose a source to pass through the switch.

Controls

The Switch Controls consist of a Source switcher that lets you choose which source to pass through the switch.



The Switch Configuration

Config

The Switch Configuration lets you choose the number of inputs and rename them.

- **Number of Inputs:** Adjust the slider to add the number of inputs you wish to use, up to nine. If you want more than nine inputs, you can click in the number field and type the number manually.
- **Name X:** You can change the name of the input from numerical to something more descriptive.

Time Speed [TSpd]



The Time Speed node

Time Speed Node Introduction

The Time Speed node allows image sequences to be sped up, slowed down, reversed, or delayed. Image Interpolation offers smooth, high-quality results. Time Speed should be used for static speed changes or to introduce delays in the footage. To apply animated changes in time, such as accelerating or decelerating time, use a Time Stretcher instead.

When operating in Flow mode, Optical Flow data is required. This node does not generate optical flow directly. You have to create it upstream using an Optical Flow node or by loading the forward/reverse vector channels from the image.

TimeSpeed does not interpolate the aux channels but instead destroys them. In particular, the Vector/BackVector channels are consumed and destroyed after computation.

Add an Optical Flow after the Time Speed node if you want to generate flow vectors for the retimed footage.

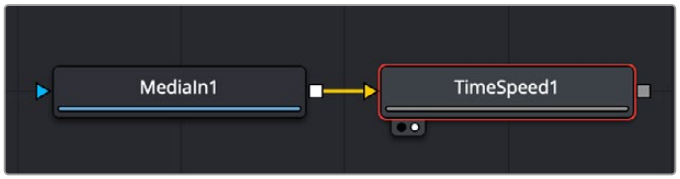
Inputs

The single input on the Time Speed node is used to connect a 2D image that will be retimed.

Input: The orange input is used for the primary 2D image that will be retimed.

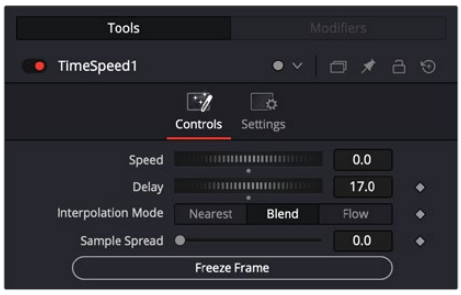
Basic Node Setup

The Time Speed node setup is as simple as connecting a 2D image into the orange background input of the node.



A MediaIn node having its speed changed in the Time Speed node.

Inspector



The Time Speed Controls tab

Speed

This control is used to adjust the Speed, in percentage values, of the outgoing image sequence. Negative values reverse the image sequence. 200% Speed is represented by a value of 2.0, 100% by 1.0, 50% by 0.5, and 10% by 0.1.

The Speed control cannot be animated.

Delay

Use this control to Delay the outgoing image sequence by the specified number of frames. Negative numbers offset time back, and positive numbers advance time.

Interpolate Mode

This menu determines the how the time speed is processed in order to improve its visual playback quality, especially in the case of clips that are slowed down. There are three choices in the menu.

- **Nearest:** The most processor efficient and least sophisticated method of processing; frames are either dropped for fast motion or duplicated for slow motion.
- **Blend:** Also processor efficient, but can produce smoother results; adjacent duplicated frames are dissolved together to smooth out slow or fast motion effects.
- **Flow:** The most processor intensive but highest quality method of speed effect processing. Using vector channels pre-generated from an Optical Flow node, new frames are generated to create slow or fast motion effects. The result can be exceptionally smooth when motion in a clip is linear. However, two moving elements crossing in different directions or unpredictable camera movement can cause unwanted artifacts.

Sample Spread

This slider is displayed only when Interpolation is set to Blend. The slider controls the strength of the interpolated frames on the current frame. A value of 0.5 blends 50% of the frame before and 50% of the frame ahead and 0% of the current frame.

Depth Ordering

This menu is displayed only when Interpolation is set to Flow. The Depth Ordering is used to determine which parts of the image should be rendered on top. This is best explained by example.

In a locked-off camera shot where a car is moving through the frame, the background does not move, so it produces small, or slow, vectors. The car produces larger, or faster, vectors.

The Depth Ordering, in this case, is Fastest on Top, since the car draws over the background.

In a shot where the camera pans to follow the car, the background has faster vectors, and the car has slower vectors, so the Depth ordering method would be Slowest on Top.

Clamp Edges

This checkbox is displayed only when Interpolation is set to Flow. Under certain circumstances, this option can remove the transparent gaps that may appear on the edges of interpolated frames. Clamp Edges can cause a stretching artifact near the edges of the frame that is especially visible with objects moving through it or when the camera is moving.

Because of these artifacts, it is a good idea to use clamp edges only to correct small gaps around the edges of an interpolated frame.

Edge Softness

This slider is only displayed when Interpolation is set to Flow and Clamp Edges is enabled. It helps to reduce the stretchy artifacts that might be introduced by Clamp Edges.

If you have more than one of the Source Frame and Warp Direction checkboxes turned on, this can lead to doubling up of the stretching effect near the edges. In this case, you'll want to keep the softness rather small at around 0.01. If you have only one checkbox enabled, you can use a larger softness at around 0.03.

Source Frame and Warp Direction

These checkboxes are displayed only when Interpolation is set to Flow. These controls determine which frames and which vectors are used to create the in-between frames. Each method ticked on will be blended into the result.

- **Prev Forward:** Takes the previous frame and uses the Forward vector to interpolate the new frame.
- **Next Forward:** Takes the next frame in the sequence and uses the Forward vector to interpolate the new frame.
- **Prev Backward:** Takes the previous frame and uses the Back Forward vector to interpolate the new frame.
- **Next Backward:** Takes the next frame in the sequence and uses the Back vector to interpolate the new frame.

Freeze Frame

Pressing this button automatically adjusts the Speed and Delay controls to freeze playback on the currently selected frame.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other miscellaneous nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Time Stretcher [TSt]



The Time Stretcher node

Time Stretcher Node Introduction

The Time Stretcher node is similar to the Time Speed node but permits the speed of the clip to be animated. Full spline control of the effect is provided, including smoothing. As a result, the Time Stretcher can be used to animate a single clip to 200, back to normal speed, pause for a second, and then play backward (like a VCR rewinding).

Image interpolation offers smooth, high-quality results, all using a spline curve to adjust time nonlinearly. To apply constant time changes such as frame rate changes, use a Time Speed instead.

When operating in Flow mode, Optical Flow data is required. This node does not generate optical flow directly; you must create it manually upstream using an Optical Flow node or by loading the forward/reverse vector channels from disk.

Flow Stretcher does not interpolate the aux channels but instead destroys them. In particular, the Vector/BackVector channels are consumed/destroyed. Add an Optical Flow after the Flow Stretcher if you want to generate flow vectors for the retimed footage.

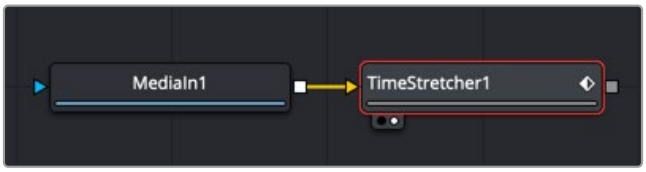
Inputs

The single input on the Time Stretcher node is used to connect a 2D image that will be time stretched.

Input: The orange input is used for the primary 2D image that will be time stretched.

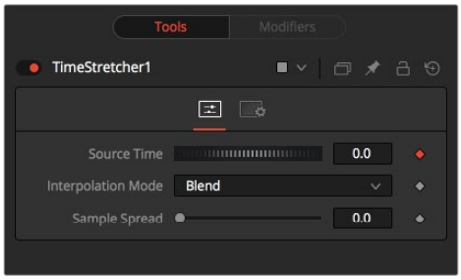
Basic Node Setup

Like the Time Speed node, the Time Stretcher setup is as simple as connecting a 2D image into the orange background input of the node.



A MediaIn node having its time ramped to various speeds in the Time Stretcher node.

Inspector



The Time Stretcher controls

Source Time

This control designates from which frame in the original sequence to begin sampling.

When a Time Stretcher node is added to the node tree, the Source Time control already contains a Bézier spline with a single keyframe set to 0.0. The keyframe position is determined by the current time when the node is added to the node tree.

> **NOTE:** The Source Time spline may not be immediately visible until Edit is selected from the Source Time's contextual menu, or Display all Splines is selected from the Spline Window's contextual menu.

Interpolate Mode

This menu determines the how the time speed is processed in order to improve its visual playback quality, especially in the case of clips that are slowed down.

There are three choices in the menu.

- **Nearest:** The most processor efficient and least sophisticated method of processing; frames are either dropped for fast motion or duplicated for slow motion.
- **Blend:** Also processor efficient but can produce smoother results; adjacent duplicated frames are dissolved together to smooth out slow or fast motion effects.

- **Flow:** The most processor intensive but highest quality method of speed effect processing. Using vector channels pre-generated from an Optical Flow node, new frames are generated to create slow or fast motion effects. The result can be exceptionally smooth when motion in a clip is linear. However, two moving elements crossing in different directions or unpredictable camera movement can cause unwanted artifacts.

Sample Spread

This slider is displayed only when Interpolation is set to Blend. The slider controls the strength of the interpolated frames on the current frame. A value of 0.5 blends 50% of the frame before and 50% of the frame ahead and 0% of the current frame.

Depth Ordering

This menu is displayed only when Interpolation is set to Flow. The Depth Ordering is used to determine which parts of the image should be rendered on top. This is best explained by example.

In a locked-off camera shot where a car is moving through the frame, the background does not move, so it produces small, or slow, vectors. The car produces larger, or faster, vectors.

The Depth Ordering in this case is Fastest on Top, since the car draws over the background.

In a shot where the camera pans to follow the car, the background has faster vectors, and the car has slower vectors, so the Depth ordering method would be Slowest on Top.

Clamp Edges

This checkbox is displayed only when Interpolation is set to Flow. Under certain circumstances, this option can remove the transparent gaps that may appear on the edges of interpolated frames. Clamp Edges can cause a stretching artifact near the edges of the frame that is especially visible with objects moving through it or when the camera is moving.

Because of these artifacts, it is a good idea to use clamp edges only to correct small gaps around the edges of an interpolated frame.

Edge Softness

This slider is displayed only when Interpolation is set to Flow and Clamp Edges is enabled. It helps to reduce the stretchy artifacts that might be introduced by Clamp Edges.

If you have more than one of the Source Frame and Warp Direction checkboxes turned on, this can lead to doubling up of the stretching effect near the edges. In this case, you'll want to keep the softness rather small at around 0.01. If you have only one checkbox enabled, you can use a larger softness at around 0.03.

Source Frame and Warp Direction

These checkboxes are displayed only when Interpolation is set to Flow. These controls determine which frames and which vectors are used to create the in-between frames. Each method ticked on will be blended into the result.

- **Prev Forward:** Takes the previous frame and uses the Forward vector to interpolate the new frame.
- **Next Forward:** Takes the next frame in the sequence and uses the Forward vector to interpolate the new frame.
- **Prev Backward:** Takes the previous frame and uses the Back Forward vector to interpolate the new frame.
- **Next Backward:** Takes the next frame in the sequence and uses the Back vector to interpolate the new frame.

Example

Make sure that the current time is either the first or last frame of the clip to be affected in the project. Add the Time Stretcher node to the node tree. This will create a single point on the Source Time spline at the current frame. The value of the Source Time will be set to zero for the entire Global Range.

Set the value of the Source Time to the frame number to be displayed from the original source, at the frame in time it will be displayed in during the project.

To shrink a 100-frame sequence to 25 frames, follow these steps:

1. Change the Current Time to frame 0.
2. Change the Source Time control to 0.0.
3. Advance to frame 24.
4. Change the Source Time to 99.
5. Check that the spline result is linear.
6. Fusion will render 25 frames by interpolating down the 100 frames to a length of 25.
7. Hold the last frame for 30 frames, and then play the clip backward at regular speed. Continue the example from above and follow the steps below.
8. Advance to frame 129.
9. Right-click on the Source Time control and select Set Key from the menu.
10. Advance to frame 229 (129 + 100).
11. Set the Source time to 0.0.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other miscellaneous nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Wireless Link [Wire]



The Wireless Link node

Wireless Link Node Introduction

The Wireless Link node helps manage the tangle of connection lines in a node tree by wirelessly connecting one 2D node to another 2D node.

Although Wireless Links can be helpful, try to keep as much of a node tree as visible as possible; otherwise, you lose one of the main benefits of a node tree.

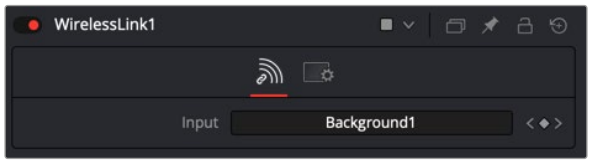
Inputs

There are no inputs on this node.

Basic Node Setup

There is no setup for this node. It is a free-standing node that connects “wirelessly” using the control in the Inspector.

Inspector



The Wireless Link Controls tab

Controls Tab

The Controls tab in the Wireless Link node contains a single Input field for the linked node.

Input

To use the Wireless Link node, in the Node Editor, drag the 2D node into the Input field of the Wireless Link node. Any change you make to the original node is wirelessly replicated in the Wireless Link node. You can use the output from the Wireless Link node to connect to a nearby node.

Common Controls

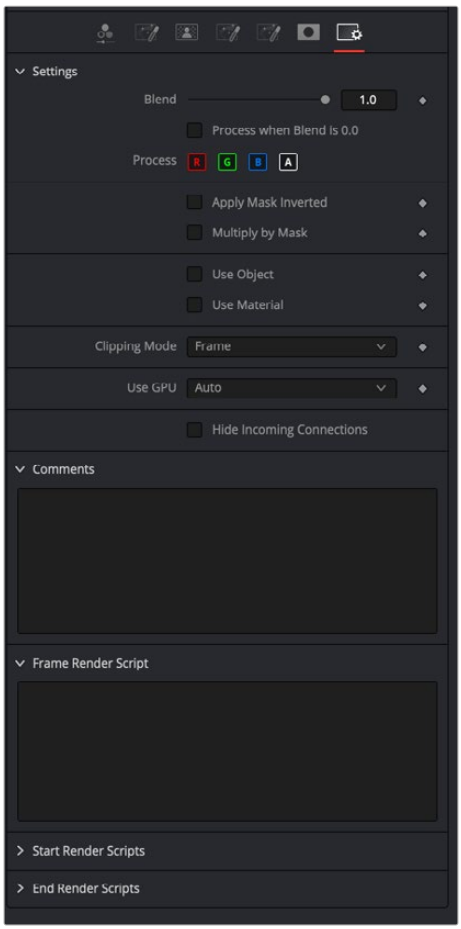
Settings Tab

The Settings tab in the Inspector is also duplicated in other miscellaneous nodes. These common controls are described in detail in the following “The Common Controls” section.

The Common Controls

Nodes that handle miscellaneous operations share a number of identical controls in the Inspector. This section describes controls that are common among miscellaneous nodes.

Inspector



The Common Miscellaneous Settings tab

Settings Tab

The Settings tab in the Inspector can be found on every tool in the miscellaneous nodes. The controls are consistent and work the same way for each tool.

Blend

The Blend control is used to blend between the tool's original image input and the tool's final modified output image. When the blend value is 0.0, the outgoing image is identical to the incoming image. Normally, this will cause the tool to skip processing entirely, copying the input straight to the output.

Process When Blend Is 0.0

The tool is processed even when the input value is zero. This can be useful if processing of this node is scripted to trigger another task, but the value of the node is set to 0.0.

Red/Green/Blue/Alpha Channel Selector

These four buttons are used to limit the effect of the tool to specified color channels. This filter is often applied after the tool has been processed.

For example, if the Red button on a Blur tool is deselected, the blur will first be applied to the image, and then the red channel from the original input will be copied back over the red channel of the result.

There are some exceptions, such as tools for which deselecting these channels causes the tool to skip processing that channel entirely. Tools that do this will generally possess a set of identical RGBA buttons on the Controls tab in the tool. In this case, the buttons in the Settings and the Controls tabs are identical.

Apply Mask Inverted

Enabling the Apply Mask Inverted option inverts the complete mask channel for the tool. The mask channel is the combined result of all masks connected to or generated in a node.

Multiply by Mask

Selecting this option will cause the RGB values of the masked image to be multiplied by the mask channel's values. This will cause all pixels of the image not included in the mask (i.e., set to 0) to become black/transparent.

Use Object/Use Material (Checkboxes)

Some 3D software can render to file formats that support additional channels. Notably, the EXR file format supports Object ID and Material ID channels, which can be used as a mask for the effect. These checkboxes determine whether the channels will be used, if present. The specific Material ID or Object ID affected is chosen using the next set of controls.

Correct Edges

This checkbox appears only when the Use Object or Use Material checkboxes are selected. It toggles the method used to deal with overlapping edges of objects in a multi-object image. When enabled, the Coverage and Background Color channels are used to separate and improve the effect around the edge of the object. If this option is disabled (or no Coverage or Background Color channels are available), aliasing may occur on the edge of the mask.

> For more information on the Coverage and Background Color channels, *see Chapter 18, "Understanding Image Channels,"* in the Fusion Reference Manual.

Object ID/Material ID (Sliders)

Use these sliders to select which ID will be used to create a mask from the object or material channels of an image. Use the Sample button in the same way as the Color Picker: to grab IDs from the image displayed in the view. The image or sequence must have been rendered from a 3D software package with those channels included.

Use GPU

The Use GPU menu has three settings. Setting the menu to Disable turns off GPU hardware-accelerated rendering. Enabled uses the GPU hardware for rendering the node. Auto uses a capable GPU if one is available and falls back to software rendering when a capable GPU is not available.

Motion Blur

- **Motion Blur:** This toggles the rendering of Motion Blur on the tool. When this control is toggled on, the tool's predicted motion is used to produce the motion blur caused by the virtual camera's shutter. When the control is toggled off, no motion blur is created.
- **Quality:** Quality determines the number of samples used to create the blur. A quality setting of 2 will cause Fusion to create two samples to either side of an object's actual motion. Larger values produce smoother results but increase the render time.
- **Shutter Angle:** Shutter Angle controls the angle of the virtual shutter used to produce the motion blur effect. Larger angles create more blur but increase the render times. A value of 360 is the equivalent of having the shutter open for one whole frame exposure. Higher values are possible and can be used to create interesting effects.
- **Center Bias:** Center Bias modifies the position of the center of the motion blur. This allows for the creation of motion trail effects.
- **Sample Spread:** Adjusting this control modifies the weighting given to each sample. This affects the brightness of the samples.

Hide Incoming Connections

Enabling this checkbox can hide connection lines from incoming nodes, making a node tree appear cleaner and easier to read. When enabled, empty fields for each input on a node will be displayed in the Inspector. Dragging a connected node from the node tree into the field will hide that incoming connection line as long as the node is not selected in the node tree. When the node is selected in the node tree, the line will reappear.

Comments

The Comments field is used to add notes to a tool. Click in the empty field and type the text. When a note is added to a tool, a small red square appears in the lower-left corner of the node when the full tile is displayed, or a small text bubble icon appears on the right when nodes are collapsed. To see the note in the Node Editor, hold the mouse pointer over the node to display the tooltip.

Scripts

Three Scripting fields are available on every tool in Fusion from the Settings tab. They each contain edit boxes used to add scripts that process when the tool is rendering. For more details on scripting nodes, please consult the Fusion scripting documentation.

Chapter 52

Optical Flow

This chapter details the Optical Flow nodes available in Fusion.

The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Optical Flow [OF] 1380
Repair Frame [REP] 1384
Smooth Motion [SM] 1386
Tween [Tw] 1388
Vector Warping Toolset (Studio Version Only) 1391
Vector Denoise [VDn] 1391
Vector Transform [VXf] 1392
Vector Warp [VWp] 1394
The Common Controls 1396

Optical Flow [OF]



The OpticalFlow node

Optical Flow Node Introduction

This node analyzes a clip connected to its input using an Optical Flow algorithm. Think of optical flow as a per-pixel motion vector that matches up features over several frames.

The computed optical flow is stored within the Vector and Back Vector aux channels of the output. These channels can be used in other nodes like the Vector Motion Blur or Vector Distort. However, Optical Flow must render twice when connecting it to a Time Stretcher or Time Speed node. These nodes require the channels A. FwdVec and B. BackVec in that order, but Optical Flow generates A. BackVec and A. FwdVec when it processes.

If you find that optical flow is too slow, consider rendering it out into OpenEXR files using a Saver node.

TIP: If the footage input flickers on a frame-by-frame basis, it is a good idea to deflicker the footage beforehand.

Inputs

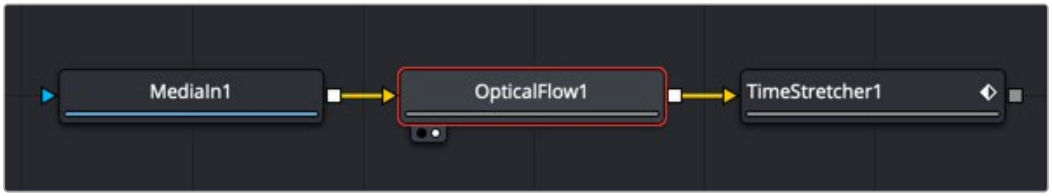
The Optical Flow node includes a single orange image input.

Input: The orange background input accepts a 2D image. This is the sequence of frames for which you want to compute optical flow. The output of the Optical Flow node includes the image and vector channels. The vector channels can be displayed by right-clicking in the viewer and choosing Channel > Vectors and then Options > Normalize Color Range.

Basic Node Setup

The Optical Flow node analyzes the frames connected to the background input. Trimming a Loader or MediaIn to only the range you need prevents analyzing unnecessary frames. The output of the node can then be connected to any node that takes advantage of vector channels, such as a Time Stretcher.

TIP: When analyzing Optical Flow vectors, consider adding a Smooth Motion node afterward with smoothing for forward/ backward vectors enabled.



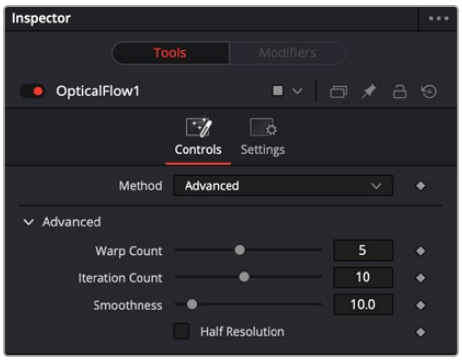
An Optical Flow node generating motion vectors on-the-fly.

Alternatively, if you find the Optical Flow node too slow to analyze the frames, consider rendering it out to an OpenEXR format using a Saver node. Then import the rendered EXR file as your new image with embedded vector channels.



An Optical Flow node rendered out through a Saver node.

Inspector



Optical Flow advanced controls

Controls Tab (Advanced)

When you add an Optical Flow, Repair Frame, or Tween node to a Comp, a Method drop-down menu in the Inspector allows you to choose between an Advanced GPU-based algorithm or a Classic CPU-based algorithm. This Advanced method is the same Optical Flow algorithm used in other DaVinci Resolve pages.

Warp Count

Decreasing this slider makes the optical flow computations faster. To understand what this option does, you must understand that the optical flow algorithm progressively warps one image until it matches with the other image. After some point, convergence is reached, and additional warps become a waste of computational time. You can tweak this value to speed up the computations, but it is good to watch what the optical flow is doing at the same time.

Iteration Count

Decreasing this slider makes the computations faster. In particular, just like adjusting the Warp Count, adjusting this option higher will eventually yield diminishing returns and not produce significantly better results. By default, this value is set to something that should converge for all possible shots and can be tweaked lower fairly often without reducing the disparity's quality.

Smoothness

This controls the smoothness of the optical flow. Higher smoothness helps deal with noise, while lower smoothness brings out more detail.

Half Resolution

The Half Resolution checkbox is used purely to speed up the calculation of the optical flow. The input images are resized down and tracked to produce the optical flow.

Controls Tab (Classic)

By choosing Classic from the Method drop-down menu in the Inspector, you can use the older CPU-based algorithm to maintain compatibility with Comps created in previous versions. This method may also be better suited for some Stereo3D processing.

When using the Classic method, a single slider at the top of the Inspector improves performance by generating proxies. The remaining Advanced section parameters tune the Optical Flow vector calculations. The default settings serve as a good standard. In most cases, tweaking of the advanced settings is not needed. Many deliver small or diminishing returns. However, depending on the settings, rendering time can easily vary by 10x. If you're interested in reducing process time, it is best to start by experimenting with the Proxy, Number of Iterations, and Number of Warps sliders and changing the filtering to Bilinear.

Proxy (for Tracking)

The Proxy slider is used purely to speed up the calculation of the optical flow. The input images are resized down by the proxy scale and tracked to produce the optical flow. The computational time is roughly proportional to the number of pixels in the image. This means a proxy scale of 2 will give a 4x speedup, and a proxy scale of 3 will give a 9x speedup.

Smoothness

This controls the smoothness of the optical flow. Higher smoothness helps deal with noise, while lower smoothness brings out more detail.

Edges

This slider is another control for smoothness but applies it based on the color channel. It tends to have the effect of determining how edges in the flow follow edges in the color images. When it is set to a low value, the optical flow becomes smoother and tends to overshoot edges. When it is set to a high value, details from the color images start to slip into the optical flow, which is not desirable. Edges in the flow end up more tightly aligning with the edges in the color images. This can result in streaked-out edges when the optical flow is used for interpolation. As a rough guideline, if you are using the disparity to produce a Z-channel for post effects like Depth of Field, then set it lower in value. If you are using the disparity to perform interpolation, you might want it to be higher in value.

Match Weight

This control sets a threshold for how neighboring groups of foreground/background pixels are matched over several frames. When set to a low value, large structural color features are matched. When set to higher values, small sharp variations in the color are matched. Typically, a good value for this slider is in the [0.7, 0.9] range. When dealing with stereo 3D, setting this option higher tends to improve the matching results in the presence of differences due to smoothly varying shadows or local lighting variations between the left and right images. The user should still perform a color match or deflickering on the initial images, if necessary, so they are as similar as possible. This option also helps with local variations like lighting differences due to light passing through a mirror rig.

Mismatch Penalty

This option controls how the penalty for mismatched regions grows as they become more dissimilar. The slider provides a choice between a balance of Quadratic and Linear penalties. Quadratic strongly penalizes large dissimilarities, while Linear is more robust to dissimilar matches. Moving this slider toward Quadratic tends to give a disparity with more small random variations in it, while Linear produces smoother, more visually pleasing results.

Warp Count

Decreasing this slider makes the optical flow computations faster. In particular, the computational time depends linearly upon this option. To understand what this option does, you must understand that the optical flow algorithm progressively warps one image until it matches with the other image. After some point, convergence is reached, and additional warps become a waste of computational time. The default value in Fusion is set high enough that convergence should always be reached. You can tweak this value to speed up the computations, but it is good to watch what the optical flow is doing at the same time.

Iteration Count

Decreasing this slider makes the computations faster. In particular, the computational time depends linearly upon this option. Just like adjusting the Warp Count, adjusting this option higher will eventually yield diminishing returns and not produce significantly better results. By default, this value is set to something that should converge for all possible shots and can be tweaked lower fairly often without reducing the disparity's quality.

Filtering

This option controls filtering operations used during flow generation. Catmull-Rom filtering will produce better results, but at the same time, turning on Catmull-Rom will increase the computation time steeply.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Optical Flow nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Repair Frame [REP]



The Repair Frame node

Repair Frame Node Introduction

Repair Frame replaces damaged or missing frames or portions of frames with scratches or other temporally transient artifacts. It requires three frames: the repair frame and two neighboring frames. An Optical Flow node is not required for generating motion vectors since the Repair Frame node computes the optical flow. However, this can make it slow to process.

Repair Frame will not pass through, but rather destroys, any aux channels after the computation is done.

See the Optical Flow node for controls and settings information.

TIP: If your footage varies in color from frame to frame, sometimes the repair can be noticeable because, to fill in the hole, Repair Frame must pull color values from adjacent frames. Consider using deflickering, color correction, or using a soft-edged mask to help reduce these kinds of artifacts.

Inputs

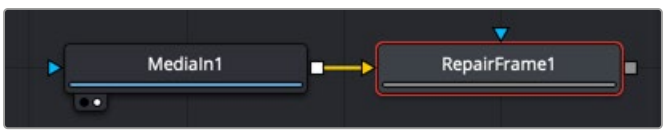
There are two inputs on the Repair Frame node. One is used to connect a 2D image that will be repaired and the other is for an effect mask.

Input: The orange input is used for the primary 2D image that will be repaired.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the repairs to certain areas.

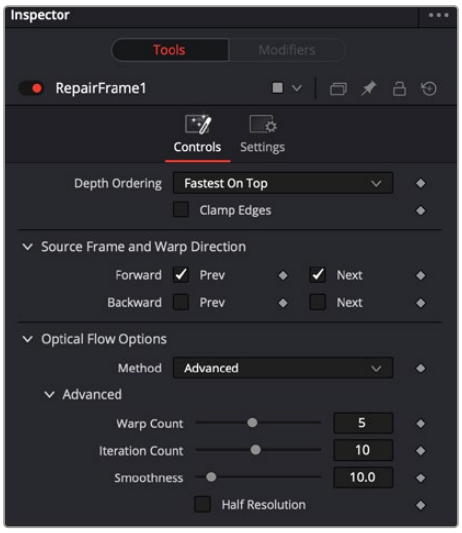
Basic Node Setup

The Repair Frame node analyzes the incoming MediaIn node and repairs single frame issues like dust or scratches.



A Repair Frame node set up to analyze a MediaIn node using internal optical flow analysis.

Inspector



The Repair Frame Controls tab

Controls Tab

The Controls tab includes options for how to repair the frames. It also includes controls for adjusting the optical flow analysis, identical to those controls in the Optical Flow node.

Depth Ordering

The Depth Ordering determines which parts of the image should be rendered on top by selecting either Fastest On Top or Slowest On Top. The examples below best explain these options.

In a locked-off camera shot where a car is moving through the frame, the background does not move, so it produces small, or slow, vectors, while the car produces larger, or faster, vectors.

The depth ordering in this case is Fastest On Top since the car draws over the background.

In a shot where the camera pans to follow the car, the background has faster vectors, and the car has slower vectors, so the Depth Ordering method is Slowest On Top.

Clamp Edges

Under certain circumstances, this option can remove the transparent gaps that may appear on the edges of interpolated frames. Clamp Edges causes a stretching artifact near the edges of the frame that is especially visible with objects moving through it or when the camera is moving.

Because of these artifacts, it is a good idea to use clamp edges only to correct small gaps around the edges of an interpolated frame.

Edge Softness

This slider is displayed only when Clamp Edges is enabled. The slider helps to reduce the stretchy artifacts that might be introduced by Clamp Edges.

If you have more than one of the Source Frame and Warp Direction checkboxes turned on, this can lead to doubling up of the stretching effect near the edges. In this case, you'll want to keep the softness rather small at around 0.01. If you have only one checkbox enabled, you can use a larger softness at around 0.03.

Source Frame and Warp Direction

These checkboxes allow you to choose which frames and vectors create the in-between frames. Each method ticked on will be blended into the result.

- **Prev Forward:** Takes the previous frame and uses the Forward vector to interpolate the new frame.
- **Next Forward:** Takes the next frame in the sequence and uses the Forward vector to interpolate the new frame.
- **Prev Backward:** Takes the previous frame and uses the Back Forward vector to interpolate the new frame.
- **Next Backward:** Takes the next frame in the sequence and uses the Back vector to interpolate the new frame.

Optical Flow Options

These settings tweak the optical flow analysis. See the Classic and Advanced Controls section for the Optical Flow node earlier in this chapter.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Optical Flow nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Smooth Motion [SM]



The Smooth Motion node

Smooth Motion Node Introduction

The Smooth Motion node smooths various AOV (Arbitrary Output Variables) channels in a clip using optical flow to look at neighboring frames. It can be used for smoothing the Disparity channel in a stereo 3D clip, where it helps reduce temporal edge/fringing artifacts, but it can also smooth a wide range of channels like vectors, normals, and Z.

It is required that the image connected to the input on the node have precomputed Vector and Back Vector channels; otherwise, this tool prints error messages in the Console window.

Check on the channels you want to temporally smooth. Be aware that if a channel selected for smoothing is not present, Smooth Motion will not fail, nor will it print any error messages.

It can also be used to smooth the Vector and Back Vector channels; however, sometimes, this can make the interpolated results worse if there are conflicting motions or objects in the shot that move around erratically, jitter, or bounce rapidly.

TIP: You can use two or more Smooth Motion nodes in sequence to get additional smoothing. With one Smooth Motion node, the previous, current, and next frames are examined for a total of 3; with two Smooth Motion nodes, 5 frames are examined; and with three Smooth Motion nodes, 7 frames are examined.

Another technique using two Smooth Motion nodes is to use the first Smooth Motion node to smooth the Vector and Back Vector channels. Use the second Smooth Motion to smooth the channels you want to smooth (e.g., Disparity). This way, you use the smoothed vector channels to smooth Disparity.

You can also try using the smoothed motion channels to smooth the motion channels.

Inputs

The Smooth Motion node includes a single orange image input.

Input: The orange image input accepts a 2D image. This is the sequence of images for which you want to compute smooth motion. This image must have precomputed Vector and Back Vector channels either generated from an Optical Flow node or saved in EXR format with vector channels.

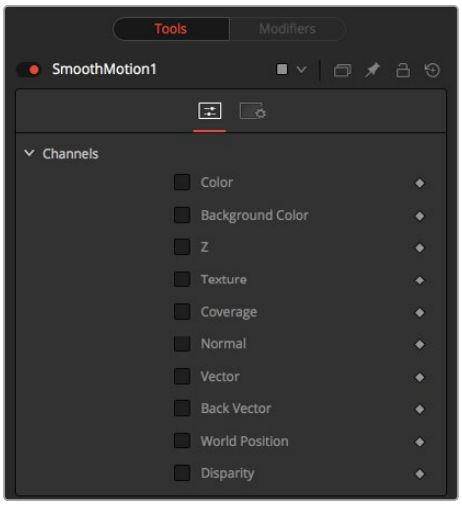
Basic Node Setup

The Smooth Motion node takes the output of the Optical Flow node for the required Vector and Back Vector channels. The Smooth Motion node can then be used to smooth those channels or AO channels.



A Smooth Motion node using Vector and Back Vector channels from the Optical Flow node.

Inspector



The Smooth Motion Controls tab

Controls Tab

The Controls tab includes checkboxes for the channels you want to smooth. If a channel selected for smoothing is not available in the input image, Smooth Motion will not fail, nor will it print any error messages to the Console.

Channel

Smooth Motion can be applied to more than just the RGBA channels. It can also be applied to the other AOV channels.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Optical Flow nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Tween [Tw]



The Tween node

Tween Node Introduction

Tween reconstructs a missing frame by interpolating between two neighboring frames using the optical flow. Tween is nearly identical in functionality to Time Speed and Time Stretcher. The major difference is that it works on two images that are not serial members of a sequence. As a consequence, it cannot use the Vector or Back Vector aux channels stored in the images. The Tween node manually generates the optical flow, so there is no need to add an Optical Flow node before the Tween node. The generated optical flow is thrown away and is not stored back into the output frames.

Since optical flow is based on color matching, it is a good idea to color correct your images to match ahead of time. Also, if you are having trouble with noisy images, it may also help to remove some of the noise ahead of time.

Tween destroys any input aux channels. See the Optical Flow node for controls and settings information.

Inputs

There are two image inputs on the Tween node and an effects mask input.

Input 0: The orange input, labeled input 0, is the previous frame to the one you are generating.

Input 1: The green input, labeled input 1, is the next frame after the one you are generating.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the Tween to certain areas.

Basic Node Setup

The Tween node receives two inputs for the two neighboring frames to the one you are generating. Below, the previous frame, frame 01, is connected to the orange input 0. The next frame, frame 03, is connected to the green input 1. The Tween node will generate frame 02 and output the sequence.



The Tween node receives two neighboring frames and generates the middle one.

Inspector



The Tween Controls tab

Controls Tab

The Controls tab includes options for how to tween frames. It also includes controls for adjusting the optical flow analysis, identical to those controls in the Optical Flow node.

Interpolation Parameter

This option determines where the frame you are interpolating is, relative to the two source frames A and B. An Interpolation Parameter of 0.0 will result in frame A, a parameter of 1.0 will result in frame B, and a parameter of 0.5 will yield a result halfway between A and B.

Depth Ordering

The Depth Ordering determines which parts of the image should be rendered on top by selecting either Fastest On Top or Slowest On Top. The examples below best explain these options.

In a locked-off camera shot where a car is moving through the frame, the background does not move, so it produces small, or slow, vectors, while the car produces larger, or faster, vectors.

The Depth Ordering in this case is Fastest On Top since the car draws over the background.

In a shot where the camera pans to follow the car, the background has faster vectors, and the car has slower vectors, so the Depth Ordering method is Slowest On Top.

Clamp Edges

Under certain circumstances, this option can remove the transparent gaps that may appear on the edges of interpolated frames. Clamp Edges causes a stretching artifact near the edges of the frame that is especially visible with objects moving through it or when the camera is moving.

Because of these artifacts, it is a good idea to use Clamp Edges only to correct small gaps around the edges of an interpolated frame.

Edge Softness

This slider is displayed only when Clamp Edges is enabled. The slider helps to reduce the stretchy artifacts that might be introduced by Clamp Edges.

If you have more than one of the Source Frame and Warp Direction checkboxes turned on, this can lead to doubling up of the stretching effect near the edges. In this case, you'll want to keep the softness rather small at around 0.01. If you have only one checkbox enabled, you can use a larger softness at around 0.03.

Source Frame and Warp Direction

These checkboxes allow you to choose which frames and vectors create the in-between frames. Each method ticked on will be blended into the result.

- **Prev Forward:** Takes the previous frame and uses the Forward vector to interpolate the new frame.
- **Next Forward:** Takes the next frame in the sequence and uses the Forward vector to interpolate the new frame.
- **Prev Backward:** Takes the previous frame and uses the Back Forward vector to interpolate the new frame.
- **Next Backward:** Takes the next frame in the sequence and uses the Back vector to interpolate the new frame.

Optical Flow Options

These settings tweak the optical flow analysis. See the Class and Advanced Controls section for the Optical Flow node earlier in this chapter.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Optical Flow nodes. These common controls are described in detail in the following "The Common Controls" section.

Vector Warping Toolset (Studio Version Only)

You can use the Vector Warping 2D imaging toolset below to map and warp a reference frame to all the other frames in the sequence. This is useful in face replacements, digital makeup, sign replacements, etc. to track and warp assets in tandem with the background. At its simplest, you:

- Apply VFX changes—make-up, face paint, scars, black eyes—on one frame,
- Ensure the areas of the reference frame are visible in the frames you are tracking,
- Add Vector Warp, set reference frame, select Generate + Texture Map and track the changes to the sequence using optical flow.

Vector Denoise [VDn]



The Vector Denoise node

Vector Denoise Node Introduction

This tool allows for a general purpose averaging operation across neighboring frames. While similar to temporal noise reduction, the source pixels to average are motion vector compensated.

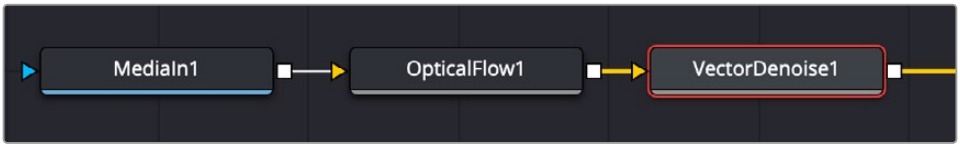
Inputs

The Vector Denoise node includes a single orange image input.

Input: The orange image input accepts a 2D image. This image must have precomputed Vector and Back Vector channels, either generated from an Optical Flow node or saved in EXR format with vector channels.

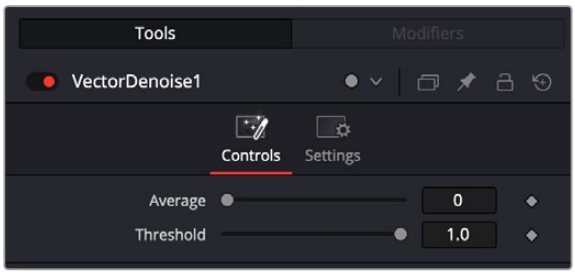
Basic Node Setup

The Vector Denoise node takes the output of the Optical Flow node for the required Vector and Back Vector channels and removes noise across multiple frames.



A Vector Denoise node using Vector and Back Vector channels from the Optical Flow node

Inspector



The Vector Denoise Controls tab

Controls Tab

The Controls tab lets you adjust the Vector Denoise settings.

- **Average:** Define the temporal window in number of frames.
- **Threshold:** Set an upper threshold to ignore flashes and brief highlights when averaging pixel values.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Optical Flow nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Vector Transform [VXf]



The Vector Transform node

Vector Transform Node Introduction

This tool allows custom transforms on the vector channel information, allowing tracking information to be resized, rotated and translated to be applied on another image or sequence. Use this to simulate transfer of kinetic energy, object responses and difference in material properties in a dynamic scene.

Inputs

The Vector Transform node includes an orange image input, and an attenuate mask input.

- **Input:** The orange image input accepts a 2D image.
- **Attenuate Mask:** The white input accepts a mask shape.

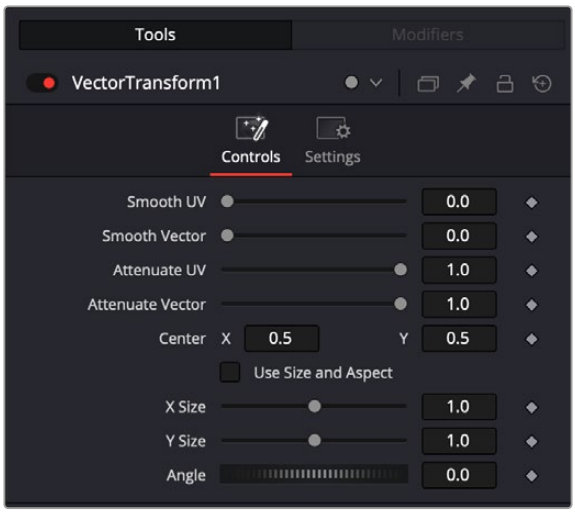
Basic Node Setup

The Vector Transform node takes the output of the Optical Flow node for the required Vector and Back Vector channels and transforms vector channel information.



A Vector Transform node using Vector and Back Vector channels from the Optical Flow node

Inspector



The Vector Transform Controls tab

Controls Tab

The Controls tab lets you adjust the Vector Transform settings.

- **Smooth UV:** Use this to average the results with the neighboring UV values. This can be used to smooth out anomalies and unwanted variations in the UV channels.
- **Smooth Vector:** Use this to smoothen vectors across neighboring frames. Useful as a pre-processing step before a Vector Warp with texture map.
- **Attenuate UV:** Amplify, decrease, or zero out the warp map effect from the vector magnitudes. This control can also decrease the strength of the UV channels.
- **Attenuate Vector:** Reduces the motion vector length, letting you control the amount of warping applied from full down to nothing.
- **Center X/Y:** Use this control to reposition the UV channels.
- **Size X/Y:** Use this control to rescale the UVs.
- **Use Size and Aspect:** Check this box to activate the parameters below.

  Size: Scales the image equally along both the X and Y axes.

  Aspect: Lets you set the aspect ratio of the image.
- **Angle:** Adjusts the angle of the UV.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Optical Flow nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Vector Warp [VWp]



The Vector Warp node

Vector Warp Node Introduction

The Vector Warp node has a main image input sequence (the background) and the map input image (a still to be warped). The output is an RGBA image and may contain extra motion vector channels (Vx, Vy. BVx, BVy) with forward and backward tracking information for use in downstream tools.

This tool distorts a nominated reference frame or image according to the motion of the source clip. The tool has two inputs: the Input Layer, which is for your source clip along with its motion vectors (generated using an Optical Flow tool) and the Texture Layer. Any clip or image connected to the Texture Layer will be warped based on the movement of the incoming vectors from the Input Layer. For warping to occur, motion vectors need to be generated for the source clip.

It is generally expected that Vector Warp follows an Optical Flow tool, which creates the required vectors.

Because Vector Warp uses motion vectors to follow movement in a clip, it’s particularly useful for adding or removing elements from non-solid surfaces like fabric or skin. This tool can be used to add or remove logos from clothing, add a scar to someone’s face, or add texture to skin

Inputs

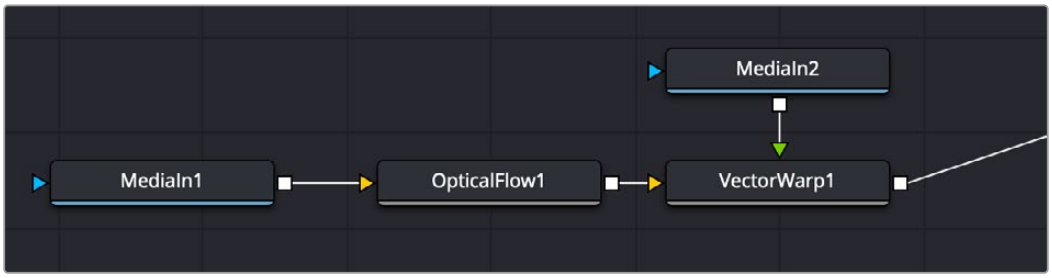
The Vector Warp node uses two inputs.

Input: The orange image input accepts a 2D image sequence. This image must have precomputed Vector and Back Vector channels, either generated from an Optical Flow node or saved in EXR format with vector channels.

Texture: The green input accepts a 2D still image to be used as the texture for warping.

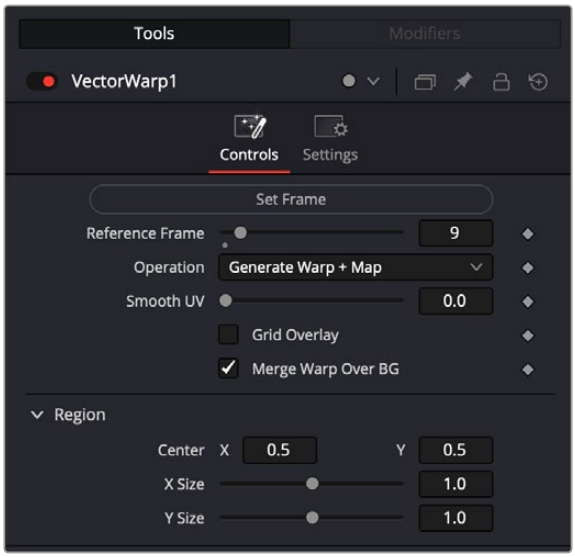
Basic Node Setup

The Vector Warp node takes the output of a clip and a 2D still image and warps them together.



A Vector Warp node

Inspector



The Vector Warp Controls tab

Controls Tab

The Controls tab lets you adjust the Vector Warp settings.

- **Set Frame:** Uses the current frame as the reference frame. Pressing this button sets the Reference Frame below.
- **Reference Frame:** Set the reference frame for the Main sequence that the Map overlay was created for.
- **Operation:** Choose the operation you wish to perform.

  Generate Warp: Creates a warp map from the reference frame to the current frame, saved in the UV channels.

  Generate Warp + Map: Creates a warp map from the reference frame to the current frame and warps the map frame.

  Apply Warp Map: Warps the map frame from a previously generated texture map from input.

  UnWarp: Freezes the map at reference frame and reverse warps the main current frame to resemble the reference frame.

- **Smooth UV:** This will average the results with the neighboring UV values. This can be used to smooth out anomalies and unwanted variations in the UV channels.
- **Grid Overlay:** Displays the warping as a viewer overlay.
- **Merge Warp over BG:** Displays the composited map image on the main image sequence.

Common Controls

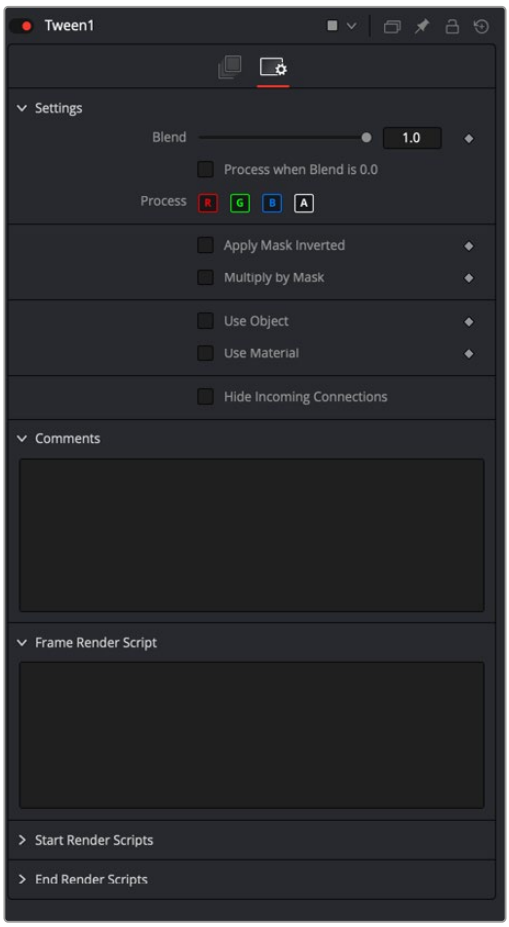
Settings Tab

The Settings tab in the Inspector is also duplicated in other Optical Flow nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

The Common Controls

Nodes that handle optical flow operations share a number of identical controls in the Inspector. This section describes controls that are common among Optical Flow nodes.

Inspector



The Common Optical Flow Settings tab

Settings Tab

The Settings tab in the Inspector can be found on every tool in the Optical Flow category. The controls are consistent and work the same way for each tool.

Blend

The Blend control is used to blend between the tool's original image input and the tool's final modified output image. When the blend value is 0.0, the outgoing image is identical to the incoming image. Normally, this causes the tool to skip processing entirely, copying the input straight to the output.

Process When Blend Is 0.0

The tool is processed even when the input value is zero. This can be useful when this node process is scripted to trigger another task, but the value of the node is set to 0.0.

Red/Green/Blue/Alpha Channel Selector

These four buttons are used to limit the effect of the tool to specified color channels. This filter is often applied after the tool has been processed.

For example, if the Red button on a Blur tool is deselected, the blur is first applied to the image, and then the red channel from the original input will be copied back over the red channel of the result.

There are some exceptions, such as tools for which deselecting these channels causes the tool to skip processing that channel entirely. Tools that do this generally possess a set of identical RGBA buttons on the Controls tab in the tool. In this case, the buttons in the Settings and the Controls tabs are identical.

Apply Mask Inverted

Enabling the Apply Mask Inverted option inverts the complete mask channel for the tool. The mask channel is the combined result of all masks connected to or generated in a node.

Multiply by Mask

Selecting this option causes the RGB values of the masked image to be multiplied by the mask channel's values. This causes all pixels of the image not included in the mask (i.e., set to 0) to become black/transparent.

Use Object/Use Material (Checkboxes)

Some 3D software can render to file formats that support additional channels. Notably, the EXR file format supports Object ID and Material ID channels, which can be used as a mask for the effect. These checkboxes determine whether the channels are used, if present. The specific Material ID or Object ID affected is chosen using the next set of controls.

Correct Edges

This checkbox appears only when the Use Object or Use Material checkboxes are selected. It toggles the method used to deal with overlapping edges of objects in a multi-object image. When enabled, the Coverage and Background Color channels are used to separate and improve the effect around the edge of the object. If this option disabled (or no Coverage or Background Color channels are available), aliasing may occur on the edge of the mask.

For more information *see Chapter 18, "Understanding Image Channels,"* in the Fusion Reference Manual.

Object ID/Material ID (Sliders)

Use these sliders to select which ID is used to create a mask from the object or material channels of an image. Use the Sample button in the same way as the Color Picker: to grab IDs from the image displayed in the view. The image or sequence must have been rendered from a 3D software package with those channels included.

Hide Incoming Connections

Enabling this checkbox can hide connection lines from incoming nodes, making a node tree appear cleaner and easier to read. When enabled, empty fields for each input on a node are displayed in the Inspector. Dragging a connected node from the node tree into the empty field hides that incoming connection line as long as the node is not selected in the node tree. When the node is selected in the node tree, the line reappears.

Comments

The Comments field is used to add notes to a tool. Click in the empty field and type the text. When a note is added to a tool, a small red square appears in the lower-left corner of the node when the full tile is displayed, or a small text bubble icon appears on the right when nodes are collapsed. To see the note in the Node Editor, hold the mouse pointer over the node to display the tooltip.

Scripts

Three Scripting fields are available on every tool in Fusion from the Settings tab. They each contain edit boxes used to add scripts that process when the tool is rendering.

> For more details on scripting nodes, see the Fusion scripting documentation.

Chapter 53

Paint Node

This chapter details the Paint node available in Fusion.

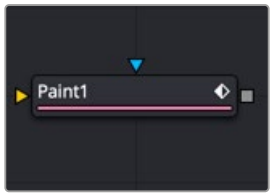
The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Paint 1400
Paint Node Introduction 1400
Inputs 1400
Basic Node Setup 1400
Types of Paint Strokes 1401
Editing Options Toolbar 1402
Inspector 1404
Paint Node Modifiers 1407
Keyboard Shortcuts 1408

Paint



The Paint node

Paint Node Introduction

Paint is an extremely flexible, stroke-based tool for wire and rig removal, image cloning, or to create custom masks and mattes rapidly. Fusion's paint can even be used to create new images and artistic elements from scratch.

Each Paint node is made up of a series of brush strokes. These strokes are vector shapes with editable brush, size, and effect. A wide range of apply modes and brush types are available.

Most Brushstrokes styles are editable polylines for fine control. They can be animated to change shape, length, and size over time. The opacity and size of a stroke can be affected by velocity and pressure when used with a supported tablet.

Unlimited undo and redo of paint provides the ability to experiment before committing changes to an image sequence. Paint strokes can be reordered, deleted, and modified with virtually infinite flexibility.

Inputs

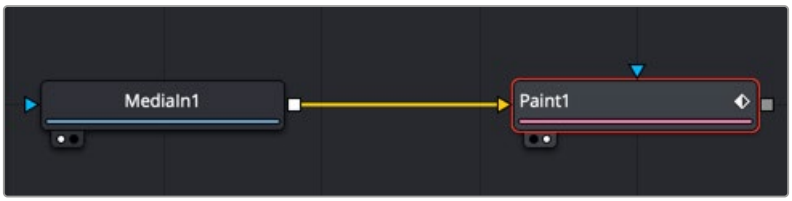
The two inputs on the Paint node are used to connect a 2D image and an effect mask which can be used to limit the painted area.

Input: It is required to connect the orange input with a 2D image that creates the size of the "canvas" on which you paint.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the Paint to only those pixels within the mask.

Basic Node Setup

The Paint node always needs an input connection. The simplest setup is to paint directly on the incoming MediaIn node.



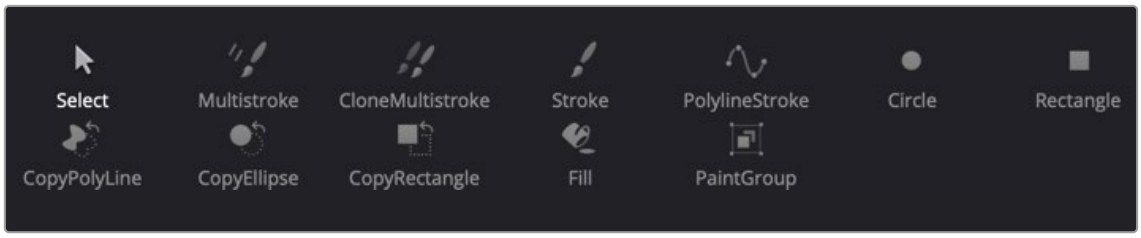
A MediaIn node connected directly to the Paint input

A more flexible setup is to use a Background node to set the size that matches the image you are painting on. In the Inspector, the background would be set to be fully transparent. Then, the Paint tool can be merged as the foreground over the actual image you want to paint on.



A Paint node merged over the top of a MediaIn for more flexibility

Types of Paint Strokes



Paint stroke types located in the viewer toolbar.

To begin working with the Paint tool, first select the paint stroke type from the Paint toolbar above the viewer. There are ten stroke types to choose from as well as two additional tools for selecting and grouping paint strokes. The stroke types and tools are described below in the order they appear in the toolbar.

- **Multistroke:** Although this is the default selection and the first actual brush type in the toolbar, Multistroke is not typically the stroke type most often used. However, it's perfect for those 100-strokes-per-frame retouching paint jobs like removing tracking markers. Multistroke is much faster than the Stroke type but is not editable after it is created. By default, Multistroke lasts for one frame and cannot be modified after it has been painted. Use the Duration setting in the Stroke controls to set the number of frames before painting. A shaded area of the Multistroke duration is visible but not editable in the Keyframes Editor. While Multistrokes aren't directly editable, they can be grouped with the PaintGroup modifier, then tracked, moved, and rotated by animating the PaintGroup instead.
- **Clone Multistroke:** Similar to Multistroke but specifically meant to clone elements from one area or image to the other. Perfect for those 100-strokes-per-frame retouching paint jobs like removing tracking markers. Clone Multistroke is faster than the Stroke type but is not editable after it is created. By default, Clone Multistroke lasts for one frame and cannot be modified after it has been painted. Use the Duration setting in the Stroke controls to set the number of frames before painting. A shaded area of the Clone Multistroke duration is visible but not editable in the Keyframes Editor.
- **Stroke:** In most cases, the Stroke tool is what people think of when they think of paint and is the tool of choice for most operations. It is a fully animatable and editable vector-based paint stroke. It can become slow if hundreds of strokes are used in an image; when creating a lot of

paint strokes, it is better to use Multistroke. The Stroke type has a duration of the entire global range. However, you can edit its duration at any time in the Keyframes Editor. When the painting is complete, choose the Select button in the Paint toolbar to avoid accidentally adding new strokes.

- **Polyline Stroke:** This provides the ability to create and manipulate a stroke in the same way that a Bézier path or polygon mask might be created. To add a Polyline Stroke, select the Polyline button and click in the viewer to add the first point. Continue clicking to add additional points to the polyline. This click append style is the default, but polyline strokes can also be created in draw append mode. Polylines can be tracked or connected to existing polylines like masks or animation paths. The Polyline Stroke has a default duration of the entire global range. However, you can edit its duration at any time in the Keyframes Editor.
- **Circle:** Creates a circular shape with animatable control over radius and center. The Circle type has a duration of the entire global range. However, you can edit its duration at any time in the Keyframes Editor.
- **Rectangle:** Creates a rectangular area. The Rectangle type has a duration of the entire global range. However, you can edit its duration at anytime in the Keyframes Editor.
- **Copy Polyline:** Allows you to create a closed Polyline area with animatable offset to clone elements from one area or image to the other. The Copy Polyline type has a duration of the entire global range. However, you can edit its duration at any time in the Keyframes Editor.
- **Copy Circle/Rectangle:** A circular or rectangular shape area with animatable offset to clone elements from one image to the other. The Copy Circle/Rectangle type has a duration of the entire global range. However, you can edit its duration at any time in the Keyframes Editor.
- **Fill:** Similar to the Wand mask tool. This tool fills similarly colored adjacent pixels with a fill color based on the selected color channel. The tool has a duration of the entire global range. However, you can edit its duration at any time in the Keyframes Editor
- **Paint Group:** Allows you to group multiple paint strokes with full control over center and size. Since Multistroke and Clone Multistroke styles cannot be tracked, grouping these paint stroke types and tracking the paint group is one common use of groups.

Editing Options Toolbar



Paint edit options are displayed in the viewer after a Polyline stroke is created.

Polyline-based paint strokes include a second toolbar in the viewer to select different editing options. The paint strokes that include this second toolbar are Polyline Stroke and Copy Polyline. The Stroke style also displays this toolbar after the stroke is selected and the Make Editable button is clicked in the Inspector.

- **Click Append:** This is the default option when creating a polyline stroke. It works more like a Bézier pen drawing tool than a paintbrush tool. Clicking sets a control point and appends the next control point when you click again in a different location.

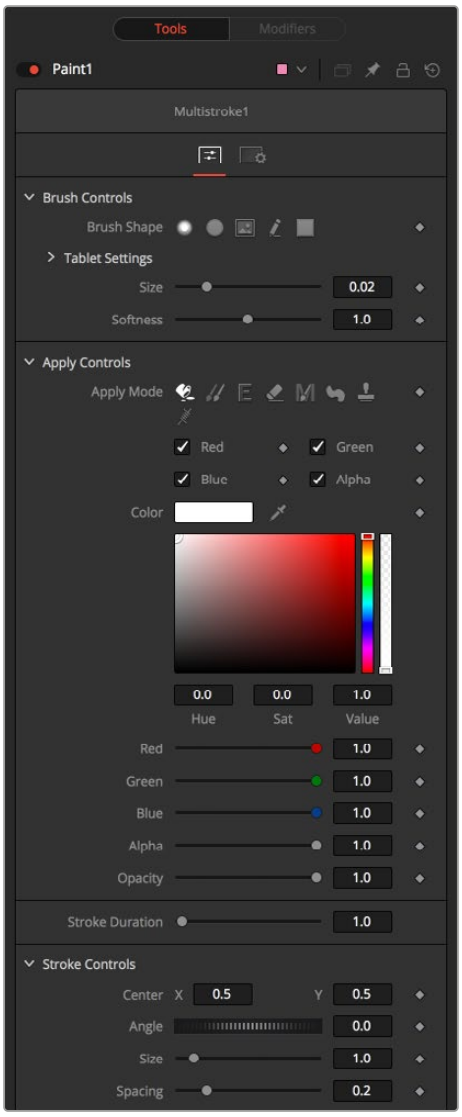
- **Draw Append:** This is a freehand drawing tool. It paints a stroke similar to drawing with a pencil on paper. You can create a new Polyline Stroke or Copy Polyline Stroke using the Draw tool, or you can extend a Stroke style after clicking the Make Editable button in the Inspector.
- **Insert:** Insert adds a new control point along the paint stroke spline.
- **Modify:** Modify allows you to safely move or smooth any exiting point along a spline without worrying about adding a new point accidentally.
- **Done:** Prevents any point along the spline from being moved or modified. Also, new points cannot be added. You can, however, move and rotate the entire spline.
- **Closed:** Closes an open polyline.
- **Smooth:** Changes the selected stroke or control point from a linear to a smooth curve.
- **Linear:** Changes the selected stroke or control point from a smooth curve to linear.
- **Select All:** Selects all the control points on the polyline.
- **Keys:** Shows or hides the control points along the polyline.
- **Handles:** Shows or hides the Bézier handles along the polyline.
- **Shape:** Places a reshape rectangle around the selected polyline control points. Using the reshape rectangle, you can deform groups of polyline control points or entire shapes much easier than modifying each point.
- **Delete:** Deletes the selected control point(s).
- **Reduce:** Opens a Freehand precision window that can be used to reduce the number of control points on a polyline. This can make the paint stroke easier to modify, especially if it has been created using the Draw tool.
- **Publish:** You can use the Publish menu to either publish control points or the path. Publishing is a form of parameter linking. It makes the selected item available for use by other controls, or to attach a control point to a tracker.
- **Follow Points:** Allows a selected point to follow the path of a published point. The point follows the published point using an offset position.
- **Roto Assist:** Enable the Roto Assist button when you begin painting with the Polyline Stroke tool. The polyline points snap to the closest edge as you click to add points to the shape. A cyan outline indicates the points that have snapped to an edge. There are three main Roto Assist options selectable through the drop-down menu:

  Multiple Points: When enabled, a single click on a high-contrast edge will add multiple points to define the entire edge, instead of having to add each point individually. This is a one time only click. The second click reverts to single point edge detection.

  Distance 8: Opens a dialog where you can set the pixel range within which searching for an edge will take place.

  Reset: Used for resetting the snap attribute of all snapped points. After resetting, the points will become unavailable for tracking.

Inspector



The Paint Controls tab

Controls Tab

Not all of the controls described here appear in all modes. Some controls are useful only in a specific Paint mode and do not appear when they are not applicable. The Controls tab is used to configure your paint settings before painting. Once a paint stroke is created, except for the Multistroke and Clone Multistroke, you can select the stroke in the viewer and update the controls.

Brush Controls

Brush Shape

The brush shape buttons select the brush tip shape. Except for the single pixel shape, you can modify the size of the brush shape in the viewer by holding down the Command or Ctrl key while dragging the mouse.

- **Soft Brush:** The Soft Brush type is a circular brush tip with soft edges.
- **Circular Brush:** A Circular Brush is a brush tip shape with hard edges.
- **Image Brush:** The Image Brush allows images from any node in the node tree, or from a file system, to be used as a brush tip.
- **Single Pixel Brush:** The Single Pixel Brush is perfect for fine detail work, creating a brush tip precisely one pixel in size. No anti-aliasing is applied to the single pixel brush.
- **Square Brush:** A Square Brush is a brush tip with hard edges.

Vary Size

Vary size settings change the stroke size based on speed or a pressure-sensitive pen and tablet.

- **Constant:** The brush tip remains a constant size over the stroke.
- **With Pressure:** The stroke size varies with the actual applied pressure.
- **With Velocity:** The stroke size varies with the speed of painting. The faster the stroke, the thinner it is.

Vary Opacity

Vary opacity settings change the stroke opacity based on speed or a pressure-sensitive pen and tablet.

- **Constant:** The brush tip remains at a constant transparency setting over the entire stroke.
- **With Pressure:** The stroke transparency varies with the applied pressure.
- **With Velocity:** The stroke transparency varies with the speed of painting. The faster the stroke, the more transparent it is.

Softness

Use this control to increase or decrease the Softness of a soft brush.

Image Source

When using the Image Source brush type, select between three possible source brush images.

- **Node:** The image source is derived from the output of a node in the node tree. Drag the node into the Inspector's Source node input field to set the source.
- **Clip:** The image source is derived from an image or sequence on disk. Any file supported by Fusion's Loader or MediaIn node can be used.
- **Brush:** Select an image to use as a brush from the menu. Images located in the Fusion > Brushes directory are used to populate the menu.

Color Space

When the Fill tool is selected, a Color Space menu selects the color space when sampling colors around the Fill tool center for inclusion in the fill range.

Channel

When the Fill tool is selected, a Channel menu selects which color channel is used in the fill paint. For example, with alpha selected, the fill occurs on contiguous pixels of the alpha channel.

Apply Controls

Apply Mode

The Apply Modes are buttons that change a brush's painting functionality.

- **Color:** The Color Apply Mode paints simple colored strokes. When used in conjunction with an image brush, it can also be used to tint the image.
- **Clone:** The Clone Apply Mode copies an area from the same image using adjustable positions and time offsets. This mode can also copy portions of one image into another image. Any image from the node tree can be used as the source image.
- **Emboss:** The Emboss Apply Mode embosses the portions of the image covered by the brush stroke.
- **Erase:** Erase reveals the underlying image through all other strokes, effectively erasing portions of the strokes beneath it without actually destroying the strokes.
- **Merge:** This Apply Mode effectively merges the brush onto the image. This mode behaves in much the same way as the Color Apply Mode but has no color controls. It is best suited for use with the image brush type.
- **Smear:** Smear the image using the direction and strength of the brushstroke as a guide.
- **Stamp:** Stamps the brush onto the image, completely ignoring any alpha channel or transparency information. This mode is best suited for applying decals to the target image.
- **Wire:** This Wire Removal Mode is used to remove wires, rigging, and other small elements in the frame by sampling adjacent pixels and draw them in toward the stroke.

Stroke Controls

The stroke controls contain parameters that adjust the entire stroke of paint as well as control it over time.

- **Size:** This control adjusts the size of the brush when the brush type is set to either Soft Brush or Circle. The diameter of the brush is drawn in the viewer as a small circle surrounding the mouse pointer. The size can also be adjusted interactively in the viewer by holding the Command or Ctrl key while dragging the mouse pointer.
- **Spacing:** The Spacing slider determines the distance between dabs (samples used to draw a continuous stroke along the underlying vector shape). Increasing this value increases the density of the stroke, whereas decreasing this value causes the stroke to assume the appearance of a dotted line.
- **Stroke Animation:** The Stroke Animation menu provides several pre-built animation effects that can be applied to a paint stroke. This menu appears only for vector strokes like Stroke and Polyline Stroke.

  All Frames: This default displays the stroke for all frames of the image connected to the orange background input of the Paint node.

  Limited Duration: This exists on the number of frames specified by the Duration slider.

  Write On: When Write On is selected, an animation spline is added to the paint stroke that precisely duplicates the timing of the paint stroke's creation. The stroke is written on the image exactly as it was drawn. To adjust the timing of the Write On effect, switch to the Spline Editor and use the Time Stretcher node to adjust the overall length of the animation spline. To smooth or manually adjust the motion, try reducing the points in the animation spline.

  Write Off: Write Off performs the reverse of Write On, drawing the stroke starting from the end and working backward to the start of the stroke.

Write On Then Off: This menu option applies a Write On and then a Write Off animation mode to the stroke.

Trail: Selecting the Trail option from the menu causes both the start and end points of the stroke to be animated simultaneously, offset from each other by the amount specified in the Duration control. This has the effect of creating a painted segment that follows the stroke path. As with the Write On and Write Off effects, this starts on the current frame when the animation mode is selected. The timing of the animation can be adjusted manually using the Spline or Keyframes Editors.

- **Duration:** Duration sets the duration of each stroke in frames. This control is present only for Multistroke and Clone Multistroke, or when the stroke animation mode is set to Limited Duration. It is most commonly employed for frame-by-frame rotoscoping through a scene.

  Each Vector stroke applied to a scene has a duration in the Keyframes Editor that can be trimmed independently from one stroke to the next. The duration can be set to 0.5, which allows each stroke to last for a single field only when the node tree is processing in Fields mode.

- **Write On and Write Off:** This range slider appears when the Stroke Animation is set to one of the Write On and Write Off methods. The range represents the beginning and end points of the stroke. Increase the Start value from 0.0 to 1.0 to erase the stroke, or increase the End value from 0.0 to 1.0 to draw the stroke on the screen. This control can be animated to good effect. It works most effectively when automatically animated through the use of the Write On and Write Off modes of the Stroke Animation menu.
- **Make Editable:** This button appears only for Vector strokes. Clicking on Make Editable turns the current stroke into a polyline spline so that the shape can be adjusted or animated.

Paint Node Modifiers

Every paint stroke created in the viewer creates an associated modifier stroke. These modifier strokes are represented as a list of paint stroke operations in the Modifiers tab of the Inspector. Each stroke you create can be modified or deleted, or applied in a different order using the modifier stack.

NOTE: The MultiStroke tools are built for speed and can contain many strokes internally without creating a huge list stack in the modifiers

Each Paint modifier stroke contains Brush controls, Apply controls, and Stroke controls identical to those found in the main Controls tab of the Inspector.

Keyboard Shortcuts

Keyboard shortcuts allow you to adjust painting styles and color without having to navigate menus.

While painting:

Hold Command or Ctrl while left-dragging to change brush size.

Hold Option or Alt while clicking to pick a color in the viewer.

While cloning:

Option-click or Alt-click to set the clone source position. Strokes start cloning from the selected location.

Hold O to temporarily enable a 50% transparent overlay of the clone source (% can be changed with pref Tweaks.CloneOverlayBlend).

Press P to toggle an opaque overlay of the clone source.

While overlay is showing:

Arrow keys change the clone source position; you can also drag crosshairs and adjust angle control or size sliders.

Option + Left/Right or Alt + Left/Right Arrow keys change the clone source angle.

Option + Up/Down or Alt + Up/Down Arrow keys change the clone source size.

Shift + Command or Shift + Ctrl can be used with the above for greater or lesser adjustments. Left and right square brackets, [and], change the clone source Time Offset (this requires a specific Clone Source node to be set in the Source Node field).

Copy Rect/Ellipse:

Shift + drag out the source to constrain the shape.

With a single stroke selected (not available on multi or polyline strokes):

Press X or Y to flip the stroke.

Paint Groups:

Command + drag or Ctrl + drag to change the position of a group's crosshair, without changing the position of the group.

Chapter 54

Particle Nodes

This chapter details the Particle nodes available in Fusion.

The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

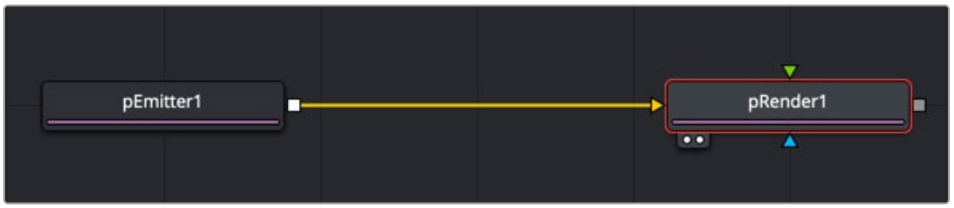
Particle Nodes 1410
pAvoid [pAv] 1410
pBounce [pBn] 1412
pChangeStyle [pCS] 1414
pCustom [pCu] 1416
pCustomForce [pCF] 1420
pDirectionalForce [pDF] 1421
pEmitter [pEm] 1423
pFlock [pFl] 1428
pFollow [pFo] 1430
pFriction [pFr] 1432
pGradientForce [pGF] 1434
pImage Emitter [pIE] 1435
pKill [pKl] 1439
pMerge [pMg] 1440
pPoint Force [pPF] 1441
pRender [pRn] 1442
About Pre-Roll 1445
Conditions 1445
Scene Tab 1446
Grid Tab 1447
Image Tab 1447
pSpawn [pSp] 1450
pTangent Force [pTF] 1452
pTurbulence [pTr] 1453
pVortex [pVt] 1455
The Common Controls 1457

Particle Nodes

The Particle nodes are used to generate a large number of duplicated objects that automatically animate. They are used to create elements like falling rain, fireworks, smoke, pixie dust, and much more. There are endless possibilities. Particles in Fusion consist of a set of nodes that are strung together in a chain for generating, modifying, and rendering particles in a 2D or 3D scene.

To begin, every particle system you create must contain two fundamental nodes:

- **pEmitter:** Used to generate the particles and control their basic look, motion. and behavior.
- **pRender:** Used to render the output of the pEmitter into a 2D or 3D scene. When creating particles, you only ever view the pRender node.



Particle Nodes

The remaining particle nodes modify the pEmitter results to simulate natural phenomena like gravity, flocking, and bounce. The names of particle nodes all begin with a lowercase p to differentiate them from non-particle nodes. They can be found in the particles category in the Effects Library.

pAvoid [pAv]



The pAvoid node

pAvoid Node Introduction

The pAvoid node is used to create a region or area within the image that affected particles attempt to avoid entering and crossing.

It has two primary controls. The first determines the distance from the region a particle should be before it begins to move away from the region. The second determines how strongly the particle moves away from the region.

A pAvoid node creates a “desire” in a particle to move away from a specific region. If the velocity of the particle is stronger than the combined distance and strength of the pAvoid region, the particle’s desire to avoid the region does not overcome its momentum and the particle crosses that region regardless.

Inputs

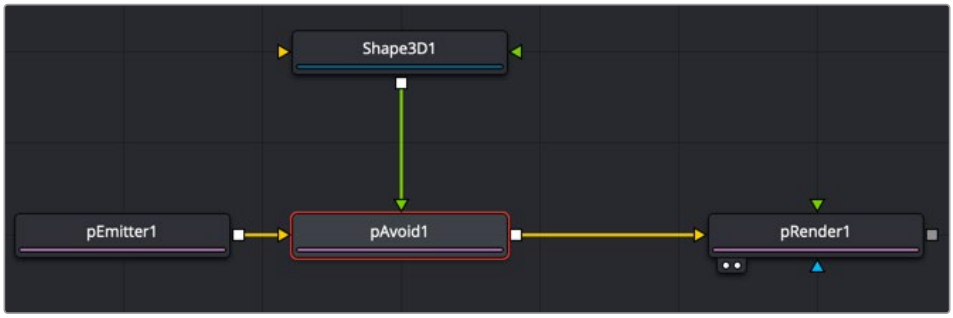
The pAvoid node has a single orange input by default. Like most particle nodes, this orange input accepts only other particle nodes. A green bitmap or mesh input appears on the node when you set the Region menu in the Region tab to either Bitmap or Mesh.

Input: The orange input takes the output of other particle nodes.

Region: The green or magenta region input takes a 2D image or a 3D mesh depending on whether you set the Region menu to Bitmap or Mesh. The color of the input is determined by whichever is selected first in the menu. The 3D mesh or a selectable channel from the bitmap defines the area particles avoid.

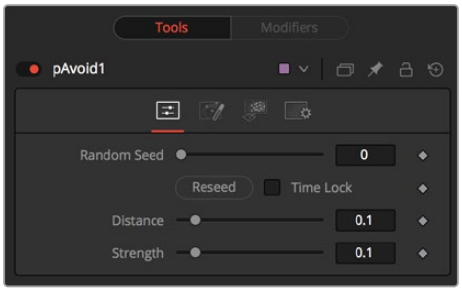
Basic Node Setup

The pAvoid node is placed in between the pEmitter and pRender. A Shape 3D node is used to create the region the particles will avoid.



A pAvoid node using a Shape 3D node as the region to avoid

Inspector



The pAvoid controls

Randomize

The Random Seed slider and Randomize button are presented whenever a Fusion node relies on a random result. Two nodes with the same seed values will produce the same random results. Click the Randomize button to randomly select a new seed value, or adjust the slider to manually select a new seed value.

Distance

Determines the distance from the region a particle should be before it begins to move away from the region.

Strength

Determines how strongly the particle moves away from the region. Negative values make the particles move toward the region instead.

Common Controls

Conditions, Style, Region, and Settings Tabs

The Conditions, Style, Region, and Settings tabs are common to all Particle nodes, so their descriptions can be found in “The Common Controls” section at the end of this chapter.

pBounce [pBn]



The pBounce node

pBounce Node Introduction

The pBounce tool is used to create a region from which affected particles will bounce off when they come into contact.

Inputs

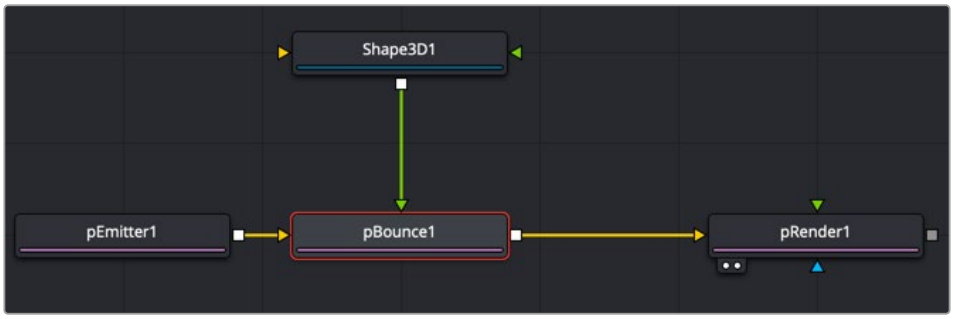
The pBounce node has a single orange input by default. Like most particle nodes, this orange input accepts only other particle nodes. A green or magenta bitmap or mesh input appears on the node when you set the Region menu in the Region tab to either Bitmap or Mesh.

Input: The orange input takes the output of other particle nodes.

Region: The green or magenta region input takes a 2D image or a 3D mesh depending on whether you set the Region menu to Bitmap or Mesh. The color of the input is determined by whichever is selected first in the menu. The 3D mesh or a selectable channel from the bitmap defines the area particles bounce off.

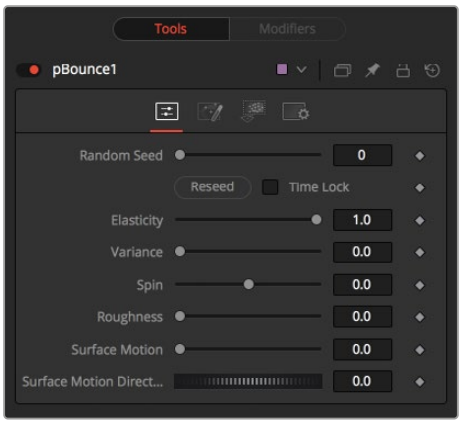
Basic Node Setup

The pBounce node is placed in between the pEmitter and pRender. A Shape 3D node is used to create the region the particles bounce off.



A pBounce node using a Shape 3D node as the region on which particles bounce off

Inspector



The pBounce controls

Randomize

The Random Seed slider and Randomize button are presented whenever a Fusion node relies on a random result.

Two nodes with the same seed values will produce the same random results. Click the Randomize button to randomly select a new seed value, or adjust the slider to manually select a new seed value.

Elasticity

Elasticity affects the strength of a bounce, or how much velocity the particle will have remaining after impacting upon the Bounce region. A value of 1.0 will cause the particle to possess the same velocity after the bounce as it had entering the bounce. A value of 0.1 will cause the particle to lose 90% of its velocity upon bouncing off of the region.

The range of this control is 0.0 to 1.0 by default, but greater values can be entered manually. This will cause the particles to gain momentum after an impact, rather than lose it. Negative values will be accepted but do not produce a useful result.

Variance

By default, particles that strike the Bounce region will reflect evenly off the edge of the Bounce region, according to the vector or angle of the region. Increasing the Variance above 0.0 will introduce a degree of variation to that angle of reflection. This can be used to simulate the effect of a rougher surface.

Spin

By default, particles that strike the region will not have their angle or orientation affected in any way. Increasing or decreasing the Spin value will cause the Bounce region to impart a spin to the particle based on the angle of collision, or to modify any existing spin on the particle. Positive values will impart a forward spin, and negative values impart a backward spin. The larger the value, the faster the spin applied to the particle will be.

Roughness

This slider varies the bounce off the surface to slightly randomize particle direction.

Surface Motion

This slider makes the bounce surface behave as if it had motion, thus affecting the particles.

Surface Motion Direction

This thumbwheel control sets the angle relative to the bounce surface.

Conditions, Style, Region, and Settings Tabs

The Conditions, Style, Region, and Settings tabs are common to all Particle nodes, so their descriptions can be found in “The Common Controls” section at the end of this chapter.

pChangeStyle [pCS]



The pChange Style node

pChange Style Node Introduction

The pChange Style node provides a mechanism for changing the appearance or style of particles that interact with a defined region. The primary controls mirror those found in the Style tab of the pEmitter node. Particles that intersect or enter the defined region change based on the parameters of this node.

Except for the pCustom node, this is the only node that modifies the particles’ appearance rather than its motion. It is often used to trigger a change in the appearance in response to some event, such as striking a barrier.

Inputs

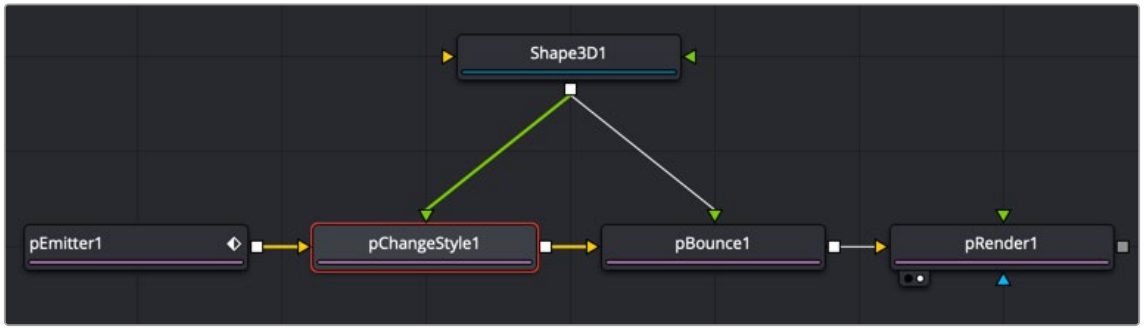
The pChange Style node has a single orange input by default. Like most particle nodes, this orange input accepts only other particle nodes. A green or magenta bitmap or mesh input appears on the node when you set the Region menu in the Region tab to either Bitmap or Mesh.

Input: The orange input takes the output of other particle nodes.

Region: The green or magenta region input takes a 2D image or a 3D mesh depending on whether you set the Region menu to Bitmap or Mesh. The color of the input is determined by whichever is selected first in the menu. The 3D mesh or a selectable channel from the bitmap defines the area where the custom particle node takes effect.

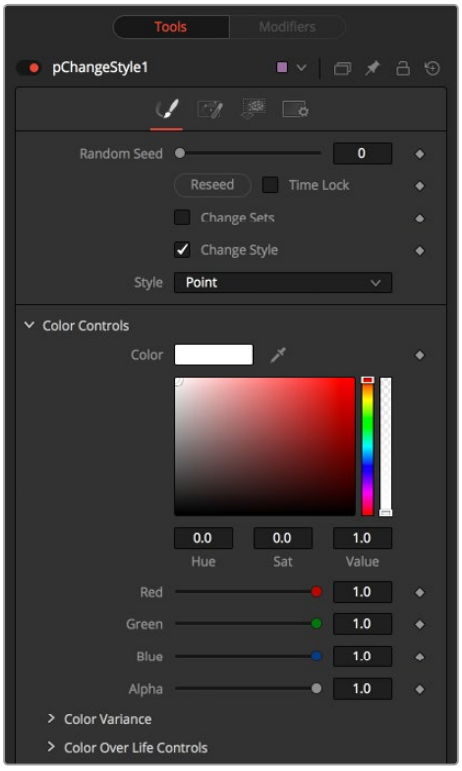
Basic Node Setup

Opposite of what you may think, to create a change in style that appears to be caused by some physical event, the pChange Style node should be placed before the node that creates the event. For example, below, the particles generated by the Emitter node change style after bouncing off a pBounce. Both the pChange Style and pBounce use the same Shape 3D node as the region. The pChange Style must be placed before the pBounce. If the pChange Style node is placed after the pBounce, the particles bounce off the region before the pChange Style calculates its effect. The particle will never get to intersect with the pChange Style node’s region, and so the style never changes.



A pChange Style node placed before the pBounce node

Inspector



The pChange Style controls

Randomize

The Random Seed slider and Randomize button are presented whenever a Fusion node relies on a random result. Two nodes with the same seed values will produce the same random results. Click the Randomize button to randomly select a new seed value, or adjust the slider to manually select a new seed value.

Change Sets

This option allows the user to change the particle's Set to become influenced by forces other than the original particle. See "The Common Controls" in this chapter to learn more about Sets.

Style

This option allows the user to change the particle's Style and thus the look. See "The Common Controls" in this chapter to learn more about Styles.

Conditions, Style, Region, and Settings Tabs

The Conditions, Style, Region, and Settings tabs are common to all Particle nodes, so their descriptions can be found in “The Common Controls” section at the end of this chapter.

pCustom [pCu]



The pCustom node

pCustom Node Introduction

The pCustom node is used to create custom expressions that affect the properties of particles. This node is similar to the Custom node, except the calculations affect particles rather than pixels.

Inputs

The pCustom node has three inputs. Like most particle nodes, this orange input accepts only other particle nodes. The green and magenta inputs are 2D image inputs for custom image calculations. Optionally, there are teal or white bitmap or mesh inputs, which appear on the node when you set the Region menu in the Region tab to either Bitmap or Mesh.

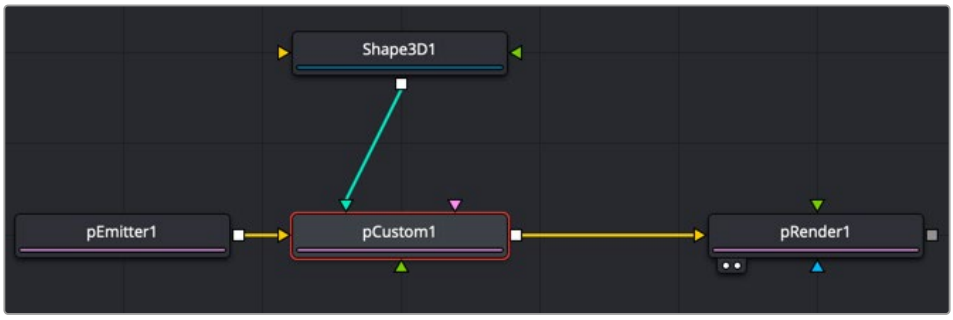
Input: The orange input takes the output of other particle nodes.

Image 1 and 2: The green and magenta image inputs accept 2D images that are used for per pixel calculations and compositing functions.

Region: The teal or white region input takes a 2D image or a 3D mesh depending on whether you set the Region menu to Bitmap or Mesh. The color of the input is determined by whichever is selected first in the menu. The 3D mesh or a selectable channel from the bitmap defines the area where the custom particle node takes effect.

Basic Node Setup

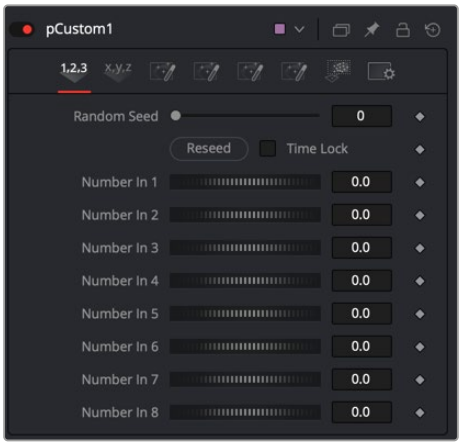
The pCustom node is placed in between the pEmitter and pRender. A Shape 3D node is used to create the region where the Custom particle event occurs.



A pCustom node using a Shape 3D node as the region where the custom event occurs

Inspector

All the same operators, functions, and conditional statements described for the Custom node apply to the pCustom node as well, including Pixel-read functions for the two image inputs (e.g., getr1w(x,y), getz2b(x,y), and so on).



The pCustom controls

Number 1-8

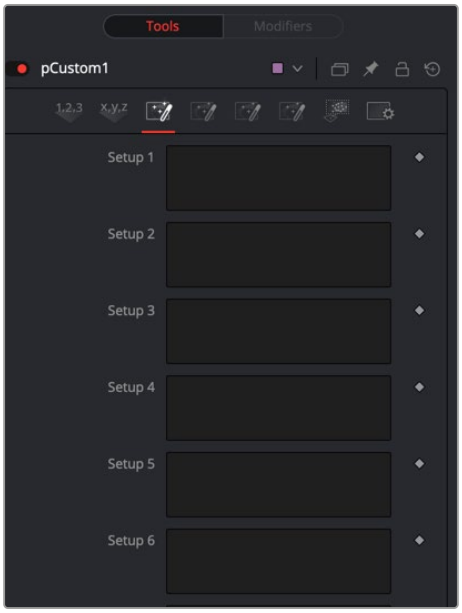
Numbers are variables with a dial control that can be animated or connected to modifiers exactly as any other control might. The numbers can be used in equations on particles at current time: n1, n2, n3, n4, ... or at any time: n1_at(float t), n2_at(float t), n3_at(float t), n4_at(float t), where t is the time you want. The values of these controls are available to expressions in the Setup and Intermediate tabs.



pCustom Position tab

Position 1-8

These eight point controls include 3D X,Y,Z position controls. They are normal positional controls and can be animated or connected to modifiers as any other node might. They are available to expressions entered in the Setup, Intermediate, and Channels tabs.

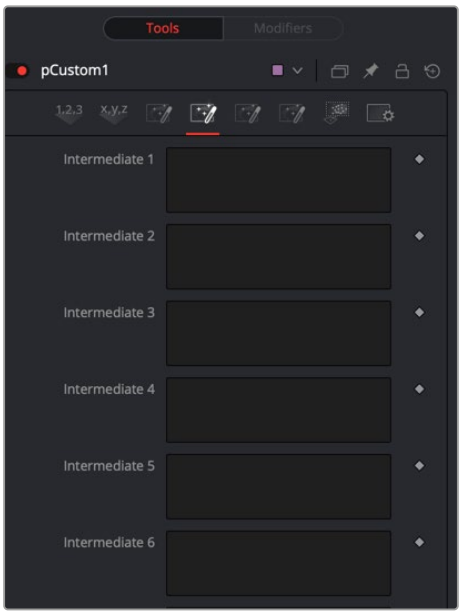


pCustom Setup tab

Setup 1-8

Up to eight separate expressions can be calculated in the Setup tab of the pCustom node. The Setup expressions are evaluated once per frame, before any other calculations are performed. The results are then made available to the other expressions in the node as variables s1, s2, s3, and s4.

Think of them as global setup scripts that can be referenced by the intermediate and channel scripts.



pCustom Intermediate tab

Inter 1-8

An additional eight expressions can be calculated in the Intermediate tab. The Intermediate expressions are evaluated once per frame, after the Setup expressions are evaluated. Results are available as variables i1, i2, i3, i4, i5, i6, i7, i8, which can be referenced by channel scripts.

Particle

Particle position, velocity, rotation, and other controls are available in the Particle tab.

The following particle properties are exposed to the pCustom control:

px, py, pz	particle position on the x, y, and z axis
vx, vy, vz	particle velocity on the x, y and z axis
rx, ry, rz	particle rotation on the x, y, and z axis
sx, sy, sz	particle spin on the x, y, and z axis
pxi1, pyi1	the 2d position of a particle, corrected for image 1's aspect
pxi2, pyi2	the 2d position of a particle, corrected for image 2's aspect
mass	not currently used by anything
size	the current size of a particle
id	the particle's identifier
r, g, b, a	the particles red, green, blue and alpha color values
rgnhit	this value is 1 if the particle hit the pCustom node's defined region
rgndist	this variable contains the particles distance from the region
condscale	the strength of the region at the particle's position
rgnix, rgniy, rgniz	values representing where on the region the particle hit
rgnnx, rgnny, rgnnz	region surface normal of the particle when it hit the region
w1, h1	image 1 width and height
w2 h2	image 2 width and height
i1, i2, i3, i4	the result of the intermediate calculations 1 through 4
s1, s2, s3, s4	the result of the setup calculations 1 through 4
n1..n8	the values of numeric inputs 1 through 8
p1x, p1y, p1z .. p4x, p4y, p4z	the values of position inputs 1 through 4
time	the current time or frame of the compositions
age	the current age of the particle
lifespan	the lifespan of the current particle

Common Controls

Conditions, Style, Region, and Settings Tabs

The Conditions, Style, Region, and Settings tabs are common to all Particle nodes, so their descriptions can be found in "The Common Controls" section at the end of this chapter.

pCustomForce [pCF]



The pCustom Force node

pCustom Force Node Introduction

The pCustom Force node allows you to change the forces applied to a particle system or subset. This node is one of the most complex and the most powerful node in Fusion. If you are experienced with scripting or C++ programming, you should find the structure and terminology used by the Custom Force node to be familiar.

The forces on a particle within a system can have their positions and rotations affected by forces. The position in XYZ and the Torque, which is the spin of the particle, are controlled by independent custom equations. The Custom Force node is used to create custom expressions and filters to modify the behavior. In addition to providing three image inputs, this node will allow for the connection of up to eight numeric inputs and as many as four XY position values from other controls and parameters in the node tree.

Inputs

The pCustom Force node has three inputs. Like most particle nodes, this orange input accepts only other particle nodes. A green and magenta are 2D image inputs for custom image calculations. Optionally there are teal or white bitmap or mesh inputs, which appear on the node when you set the Region menu in the Region tab to either Bitmap or Mesh.

Input: The orange input takes the output of other particle nodes.

Image 1 and 2: The green and magenta image inputs accept 2D images that are used for per-pixel calculations and compositing functions.

Region: The teal or white region input takes a 2D image or a 3D mesh depending on whether you set the Region menu to Bitmap or Mesh. The color of the input is determined by whichever is selected first in the menu. The 3D mesh or a selectable channel from the bitmap defines the area where the pCustom Force takes effect.

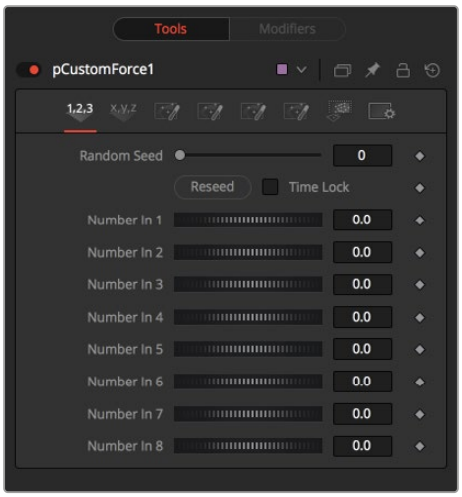
Basic Node Setup

The pCustom Force node is inserted between a pEmitter and pRender node to serve as a catalyst for particles using advanced C++ and scripting.



A pCustom Force is applied to the particles generated by the pEmitter.

Inspector



The pCustom Force controls

The tabs and controls located in the Inspector are similar to the controls found in the pCustom node. Refer to the pCustom node in this chapter for more information.

Common Controls

Conditions, Style, Region, and Settings Tabs

The Conditions, Style, Region, and Settings tabs are common to all Particle nodes, so their descriptions can be found in “The Common Controls” section at the end of this chapter.

pDirectionalForce [pDF]



The pDirectional Force node

pDirectional Force Node Introduction

This node applies a unidirectional force that pulls the affected particles in a specified direction. Its primary controls affect the strength of the force, and the angle of the force’s pull along the X, Y, and Z axis.

Since the most common use of this node is to simulate gravity, the default direction of the pull is down along the Y axis (-90 degrees), and the default behavior is to ignore regions and affect all particles.

Inputs

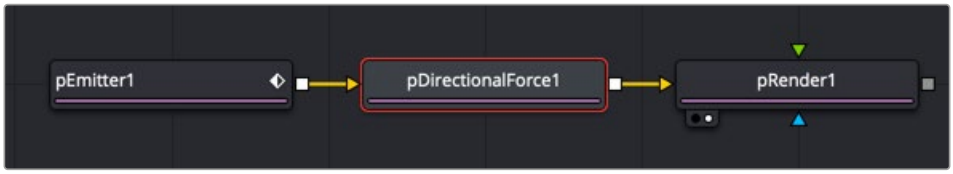
The pDirectional Force node has a single orange input by default. Like most particle nodes, this orange input accepts only other particle nodes. A green or magenta bitmap or mesh input appears on the node when you set the Region menu in the Region tab to either Bitmap or Mesh.

Input: The orange input takes the output of other particle nodes.

Region: The green or magenta region input takes a 2D image or a 3D mesh depending on whether you set the Region menu to Bitmap or Mesh. The color of the input is determined by whichever is selected first in the menu. The 3D mesh or a selectable channel from the bitmap defines the area where the directional force takes effect.

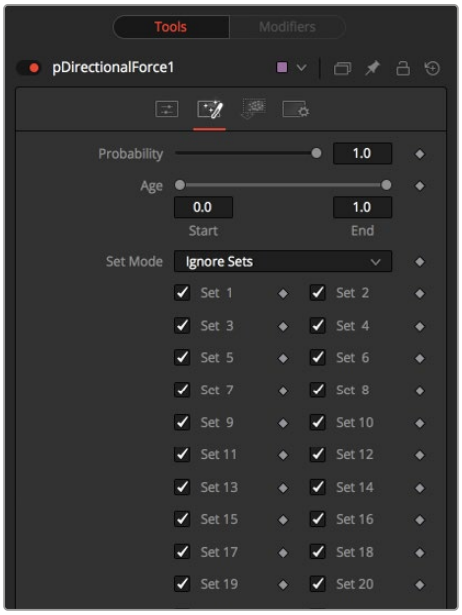
Basic Node Setup

The pDirectional Force node is placed in between the pEmitter and pRender and is often used to create gravity.



A pDirectional Force node placed between the pEmitter and pRender nodes

Inspector



The pDirectional Force controls

Randomize

The Random Seed slider and Randomize button are presented whenever a Fusion node relies on a random result. Two nodes with the same seed values will produce the same random results. Click the Randomize button to select a new seed value randomly, or adjust the slider to select a new seed value manually.

Strength

Determines the power of the force. Positive values will move the particles in the direction set by the controls; negative values will move the particles in the opposite direction.

Direction

Determines the direction in X/Y space.

Direction Z

Determines the direction in Z space.

Common Controls

Conditions, Style, Region, and Settings Tabs

The Conditions, Style, Region, and Settings tabs are common to all Particle nodes, so their descriptions can be found in “The Common Controls” section at the end of this chapter.

pEmitter [pEm]



The pEmitter node

pEmitter Node Introduction

The pEmitter node is the main source of particles (pImage Emitter is another) and will usually be the first node used in any new particle system. This node contains controls for setting the initial position, orientation, and motion of the particles, as well as controls for the visual style of each particle.

Like all other Particle nodes (with the exception of the pRender node), the pEmitter produces a particle set, not a visible image, and therefore cannot be displayed directly in a viewer. To view the output of a particle system, add a pRender node after the pEmitter.

Inputs

By default, the pEmitter node has no inputs at all. You can enable an image input by selecting Bitmap from the Style menu in the Style tab. Also, two region inputs, one for bitmap and one for mesh, appear on the node when you set the Region menu in the Region tab to either Bitmap or Mesh. The colors of these inputs change depending on the order in which they are enabled.

Style Bitmap Input: This image input accepts a 2D image to use as the particles’ image. Since this image duplicates into potentially thousands of particles, it is best to keep these images small and square—for instance, 256 x 256 pixels.

Region: The region inputs takes a 2D image or a 3D mesh depending on whether you set the Region menu to Bitmap or Mesh. The color of the input is determined by whichever is selected first in the menu. The 3D mesh or a selectable channel from the bitmap defines the area where the particles are emitted.

Basic Node Setup

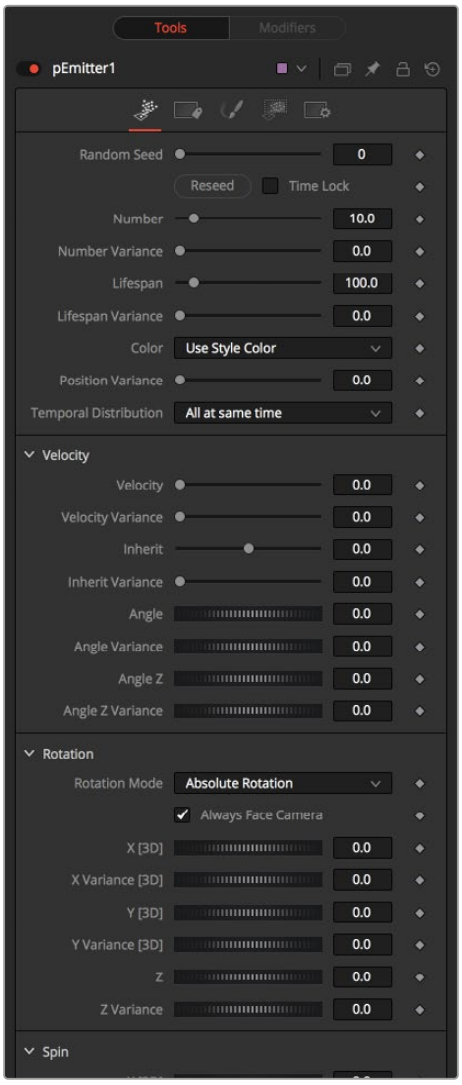
The pEmitter node starts the branch of a particle system that always ends with a pRender node. The pEmitter can feed directly into a pRender node to feed other particle nodes.



A pEmitter node connected to a pRender node is a typical setup for more particle systems.

Inspector

The pEmitter inspector is divided into four main tabs and a common settings tab. The controls tab is the first tab displayed and it contains settings that affect the general setup of the particle cells emitted by the node. These settings do not directly affect the appearance of the cells or shape of the emitter region. They modify fundamental behaviors like quantity, duration, speed, and rotation of the particle cells.



The pEmitter controls

Randomize and Random Seed

The Random Seed slider is used to seed all the variance and random number generators used by the node when creating the particle system. Two pEmitter nodes with exactly the same settings for all controls and the same random seed will generate exactly the same particle system. Changing the random seed will cause variation between the nodes. Click the Randomize button to automatically set a randomly chosen value for the Random Seed.

Number

This control is used to set the amount of new particles generated on each frame. A value of 1 would cause one new particle to be generated each frame. By frame 10, there would be a total of 10 particles in existence (unless Particle Lifespan was set to fewer than 10 frames).

Animate this parameter to specify the number of particles generated in total. For example, if only 25 particles in total are desired, animate the control to produce five particles on frame 0–4, then set a key on frame five to generate zero particles for the remainder of the project.

Number Variance

This modifies the amount of particles generated for each frame, as specified by the Number control. For example, if Number is set to 10.0 and Number Variance is set to 2.0, the emitter will produce anywhere from 9-11 particles per frame. If the value of Number Variance is more than twice as large as the value of Number, it is possible that no particles will be generated for a given frame.

Lifespan

This control determines how long a particle will exist before it disappears or 'dies.' The default value of this control is 100 frames, although this can be set to any value. The timing of many other particle controls is relative to the Lifespan of the particle. For example, the size of a particle can be set to increase over the last 80% of its life, using the Size Over Life graph in the Style tab of the pEmitter.

Lifespan Variance

Like Number Variance, the Lifespan Variance control allows the Lifespan of particles produced to be modified. If Particle Lifespan was set to 100 frames and the Lifespan Variance to 20 frames, particles generated by the emitter would have a lifespan of 90–110 frames.

Color

This provides the ability to specify from where the color of each particle is derived. The default setting is Use Style Color, which will provide the color from each particle according to the settings in the Style tab of the pEmitter node.

The alternate setting is Use Color From Region, which overrides the color settings from the Style tab and uses the color of the underlying bitmap region.

The Use Color From Region option only makes sense when the pEmitter region is set to use a bitmap produced by another node in the composition. Particles generated in a region other than a bitmap region will be rendered as white when the Use Color From Region option is selected.

Position Variance

This control determines whether or not particles can be 'born' outside the boundaries of the pEmitter region. By default, the value is set to zero, which will restrict the creation area for new particles to the exact boundaries of the defined region. Increasing this control's value above 0.0 will allow the particle to be born slightly outside the boundaries of that region. The higher the value, the 'softer' the region's edge will become.

Temporal Distribution

In general, an effect is processed per frame, based on the comp frame rate. However, processing some particles only at the exact frame boundaries can cause pulsing. To make the behavior subtly more realistic, the particles can be birthed in subframe increments.

The default, At The Same Time setting renders on frame boundaries, where as the other two settings take advantage of sub frame rendering. Randomly Distributed randomizes birth times +/- around the frame number, eg birth 10 particles at random sub times 24.1 24.85, 24.21, 24.37 etc. one particle at a time. Evenly Distributed births particles at regular sub times, eg 10 particles, birth 1 at at time at 24.0, 24.1, 24.2, 24.3, 24.4, 24.5 ... 24.8, 24.9.

These settings are influenced by the Sub Frame Accuracy setting in the pRender node. The Sub Frame Accuracy slider controls how many in-between frames are calculated between each frame. The higher the value the more accurate the particle calculation but the longer the render times.

Velocity

The controls in the Velocity section determine the speed and direction of the particle cells as the are generated from the emitter region.

Velocity and Velocity Variance

These determine the initial speed or velocity of new particles. By default, the particle has no velocity and will not move from its point of origin unless acted upon by outside forces. A velocity setting of 10.0 would cause the particle to cross the entire width of the image in one step so a velocity of 1.0 would cause the particle to cross the width of the image over 10 frames.

Velocity Variance modifies the velocity of each particle at birth, in the same manner described in Lifespan Variance and Number Variance above.

Inherit

Inherit Velocity passes the emitter region's velocity on to the particles. This slider has a wide range that includes negative and positive values. A negative value causes the particles to move in the opposite direction, a value of 1 will cause the particles to move with a velocity that matches the emitter region's velocity, and a value of 2 causes the particles to move ahead of the emitter region.

Angle and Angle Variance

This determines the angle at which particles with velocity applied will be heading at their birth.

Angle Z and Angle Z Variance

This is as above, except this control determines the angle of the particles along the Z space axis (toward or away from the camera).

Rotation

Rotation controls are used to set the orientation of particle cells and animating that orientation over time .

Rotation Mode

This menu control provides two options to help determine the orientation of the particles emitted. When the particles are spherical, the effect of this control will be unnoticeable.

- **Absolute Rotation:** The particles will be oriented as specified by the Rotation controls, regardless of velocity and heading.
- **Rotation Relative To Motion:** The particles will be oriented in the same direction as the particle is moving. The Rotation controls can now be used to rotate the particle's orientation away from its heading.

Rotation XYZ and Rotation XYZ Variance

These controls allow for Rotation of the individual particles. This can be particularly useful when dealing with a bitmap particle type, as the incoming bitmap may not be oriented in the desired direction.

Rotation XYZ Variance can be used to randomly vary the rotation by a specified amount around the center of the Rotation XYZ value to avoid having every particle oriented in the exact same direction.

Spin

Spin controls are auto animated controls that change the orientation of particle cells over time.

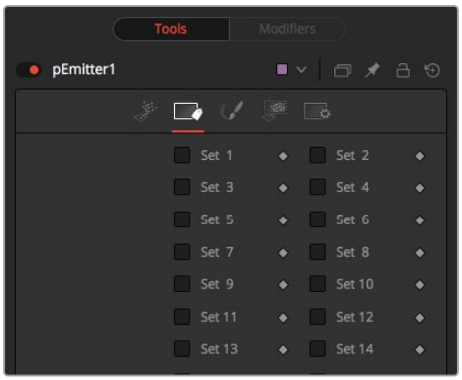
Spin XYZ and Spin Variance

These provide a spin to be applied to each particle at birth. The particles will rotate ‚x' degrees each frame, as determined by the value of Spin XYZ.

The Spin XYZ variances will vary the amount of rotation applied to each frame in the manner described by Number Variance and Lifespan Variance documented above.

Sets Tab

This tab contains settings that affect the physics of the particles emitted by the node. These settings do not directly affect the appearance of the particles. Instead, they modify behavior like velocity, spin, quantity, and lifespan.



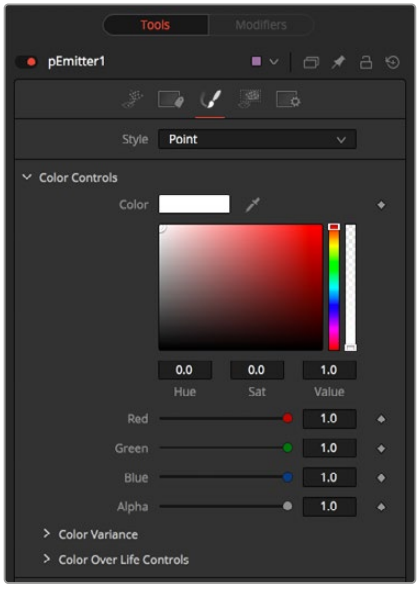
The pEmitter Sets tab

Set 1-32

To assign the particles created by a pEmitter to a given set, simply select the checkbox of the set number you want to assign. A single pEmitter node can be assigned to one or multiple sets. Once they are assigned in the pEmitter, you can enable sets in other particle nodes so they only affect particles from specific pEmitters.

Style Tab

The Style tab provides controls that affect the appearance of the particles. For detailed information about the style Tab, see the “The Common Controls” section at the end of this chapter.



The pEmitter Style tab

Region Tab

The Region tab controls the shape, size, and location of the area that emits the particle cells. This is often called the Emitter. Only one emitter region can be set for a single pEmitter node. If the pRender is set to 2D, then the emitter region will produce particles along a flat plane in Z Space. 3D emitter regions possess depth and can produce particles inside a user-defined, three-dimensional region. For more detail on the Region tab, see “The Common Controls” section at the end of this chapter.

Common Controls

Conditions, Style, Region, and Settings Tabs

The Conditions, Style, Region, and Settings tabs are common to all Particle nodes, so their descriptions can be found in “The Common Controls” section at the end of this chapter.

pFlock [pFl]



The pFlock node

pFlock Node Introduction

The pFlock node can be used to simulate the behavior of organic systems, such as a flock of birds or a colony of ants. Its use can make an otherwise mindless particle system appear to be motivated, or acting under the direction of intelligence.

The pFlock node works through two basic principles. Each particle attempts to stay close to other particles and each particle attempts to maintain a minimum distance from other particles.

The strength of these “desires” produces the seemingly motivated behavior perceived by the viewer.

Inputs

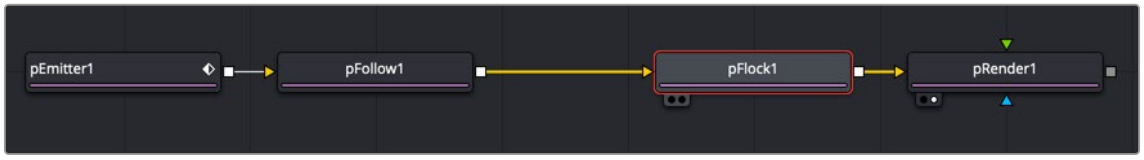
The pFlock node has a single orange input by default. Like most particle nodes, this orange input accepts only other particle nodes. A green or magenta bitmap or mesh input appears on the node when you set the Region menu in the Region tab to either Bitmap or Mesh.

Input: The orange background input takes the output of other particle nodes.

Region: The green or magenta region input takes a 2D image or a 3D mesh depending on whether you set the Region menu to Bitmap or Mesh. The color of the input is determined by whichever is selected first in the menu. The 3D mesh or a selectable channel from the bitmap defines the area where the flocking takes effect.

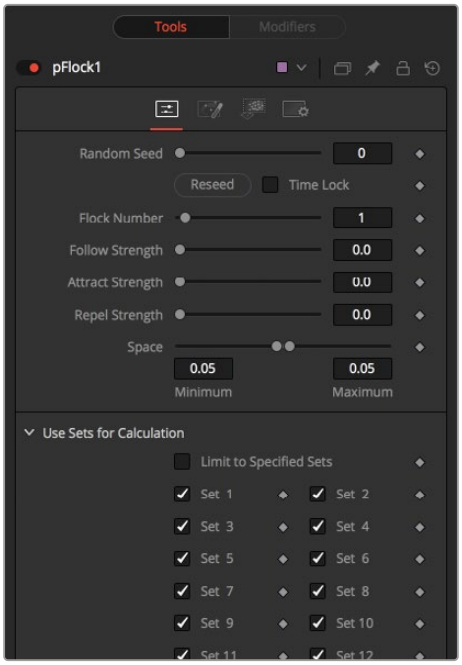
Basic Node Setup

When combined with pFollow, the pFlock node can produce natural swarming behaviors that change direction.



A pFlock node applying more herd-type mentality to particles

Inspector



The pFlock controls

Randomize

The Random Seed slider and Randomize button are presented whenever a Fusion node relies on a random result. Two nodes with the same seed values will produce the same random results. Click the Randomize button to randomly select a new seed value, or adjust the slider to manually select a new seed value.

Flock Number

The value of this control represents the number of other particles that the affected particle will attempt to follow. The higher the value, the more visible “clumping” will appear in the particle system and the larger the groups of particles will appear.

Follow Strength

This value represents the strength of each particle’s desire to follow other particles. Higher values will cause the particle to appear to expend more energy and effort to follow other particles. Lower values increase the likelihood that a given particle will break away from the pack.

Attract Strength

This value represents the strength of attraction between particles. When a particle moves farther from other particles than the Maximum Space defined in the pFlock node, it will attempt to move closer to other particles. Higher values cause the particle to maintain its spacing energetically, resolving conflicts in spacing more rapidly.

Repel Strength

This value represents the force applied to particles that get closer together than the distance defined by the Minimum Space control of the pFlock node. Higher values will cause particles to move away from neighboring particles more rapidly, shooting away from the pack.

Minimum/Maximum Space

This range control represents the distance each particle attempts to maintain between it and other particles. Particles will attempt to get no closer or farther than the space defined by the Minimum/Maximum values of this range control. Smaller ranges will give the appearance of more organized motion. Larger ranges will be perceived as disorganized and chaotic.

Common Controls

Conditions, Style, Region, and Settings Tabs

The Conditions, Style, Region, and Settings tabs are common to all Particle nodes, so their descriptions can be found in “The Common Controls” section at the end of this chapter.

pFollow [pFo]



The pFollow node

pFollow Node Introduction

Inserting the pFollow node into a particle branch causes the particles to spring back and forth toward a follow object. The follow object can be positioned in 3D or animated to create a new motion path for the particles.

Inputs

The pFollow node has a single orange input by default. Like most particle nodes, this orange background input accepts only other particle nodes. A green bitmap or mesh input appears on the node when you set the Region menu in the Region tab to either Bitmap or Mesh.

Input: The orange input takes the output of other particle nodes.

Region: The green or magenta region input takes a 2D image or a 3D mesh depending on whether you set the Region menu to Bitmap or Mesh. The color of the input is determined by whichever is selected first in the menu. The 3D mesh or a selectable channel from the bitmap defines the area where particles will follow the position point.

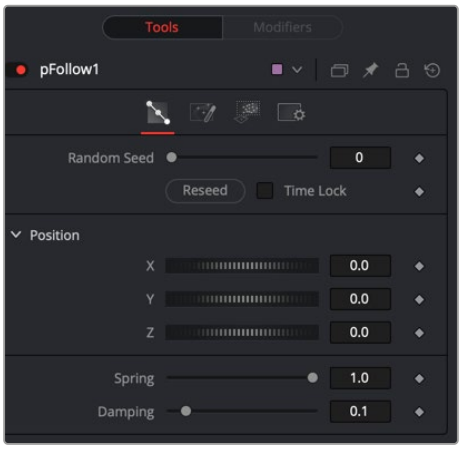
Basic Node Setup

When combined with pFlock, the pFollow node can produce natural swarming behaviors that change direction.



A pFollow node introduces a follow object that influences the particles' motion.

Inspector



The pFollow Controls tab

Random Seed

The Random Seed slider and Randomize button are presented whenever a Fusion node relies on a random result. Two nodes with the same seed values will produce the same random results. Click the Randomize button to randomly select a new seed value, or adjust the slider to manually select a new seed value.

Position XYZ

The position controls are used to create the new path by positioning the follow object. Moving the XYZ parameters displays the onscreen position of the follow object. Animating these parameters creates the new path the particles will be influenced by.

Spring

The Spring setting causes the particles to move back and forth along the path. The spread of the spring motion increases over the life of the particles depending on the distance between the particles and the follow object. Higher spring settings increase the elasticity, while lower settings decrease elasticity.

Dampen

This value attenuates the spring action. A lower setting offers less resistance to the back and forth spring action. A higher setting applies more resistance.

Common Controls

Conditions, Style, Region, and Settings Tabs

The Conditions, Style, Region, and Settings tabs are common to all Particle nodes, so their descriptions can be found in “The Common Controls” section at the end of this chapter.

pFriction [pFr]



The pFriction node

pFriction Node Introduction

The pFriction node applies resistance to the motion of a particle, slowing the particle’s motion through a defined region. This node produces two types of Friction. One type reduces the velocity of any particle intersecting/crossing the defined region, and one reduces or eliminates spin and rotation.

Inputs

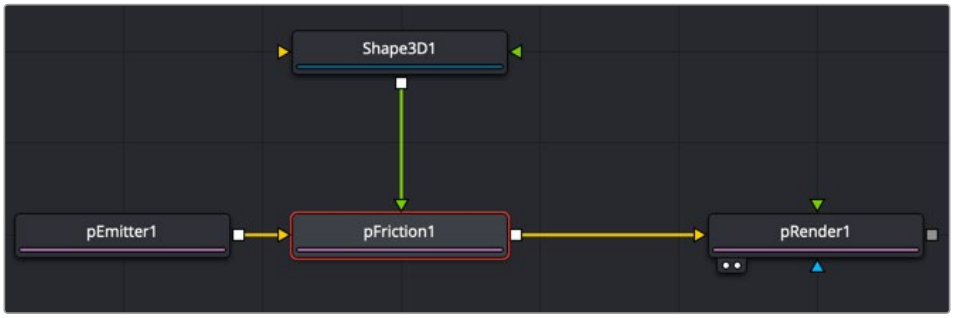
The pFriction node has a single orange input by default. Like most particle nodes, this orange input accepts only other particle nodes. A green or magenta bitmap or mesh input appears on the node when you set the Region menu in the Region tab to either Bitmap or Mesh.

Input: The orange input takes the output of other particle nodes.

Region: The green or magenta region input takes a 2D image or a 3D mesh depending on whether you set the Region menu to Bitmap or Mesh. The color of the input is determined by whichever is selected first in the menu. The 3D mesh or a selectable channel from the bitmap defines the area where the friction occurs.

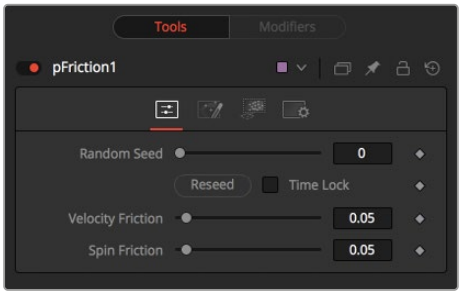
Basic Node Setup

The pFriction node is placed in between the pEmitter and pRender. A Shape 3D node is used to create the region the particles bounce off.



A pFriction node using a Shape 3D node as the region where friction is introduced to the particles.

Inspector



The pFriction controls

Random Seed

The Random Seed slider and Randomize button are presented whenever a Fusion node relies on a random result. Two nodes with the same seed values will produce the same random results. Click the Randomize button to randomly select a new seed value, or adjust the slider to manually select a new seed value.

Velocity Friction

This value represents the Friction force applied to the particle's Velocity. The larger the value, the greater the friction, thus slowing down the particle.

Spin Friction

This value represents the Friction force applied to the particle's Rotation or Spin. The larger the value, the greater the friction, thus slowing down the rotation of the particle.

Common Controls

Conditions, Style, Region, and Settings Tabs

The Conditions, Style, Region, and Settings tabs are common to all Particle nodes, so their descriptions can be found in "The Common Controls" section at the end of this chapter.

pGradientForce [pGF]



The pGradient Force node

pGradient Force Node Introduction

The particles are affected by a force generated by the gradients in the alpha values of the input image. Particles will accelerate along the gradient, moving from white to black (high values to low values).

This node can be used to give particles the appearance of moving downhill or following the contour of a provided shape.

Inputs

The pGradient Force node accepts two inputs: the default orange input from a particle node and one from a bitmap image with an alpha channel gradient. A magenta or teal bitmap or mesh input appears on the node when you set the Region menu in the Region tab to either Bitmap or Mesh.

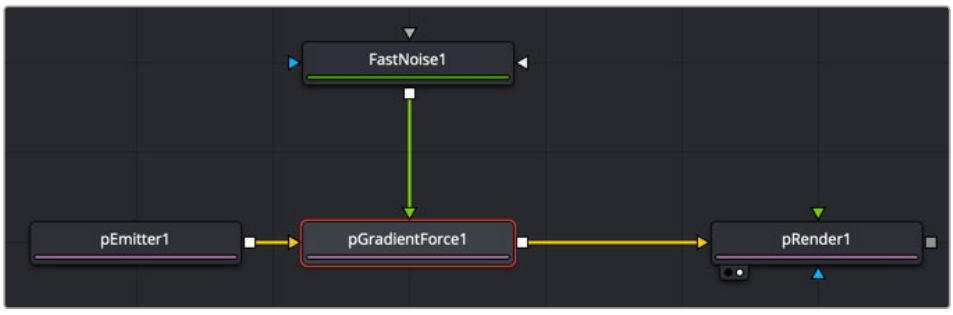
Input: The orange input takes the output of other particle nodes.

Input: The green input takes the 2D image that contains the alpha channel gradient.

Region: The magenta or teal region input takes a 2D image or a 3D mesh depending on whether you set the Region menu to Bitmap or Mesh. The color of the input is determined by whichever is selected first in the menu. The 3D mesh or a selectable channel from the bitmap defines the area where the gradient force occurs.

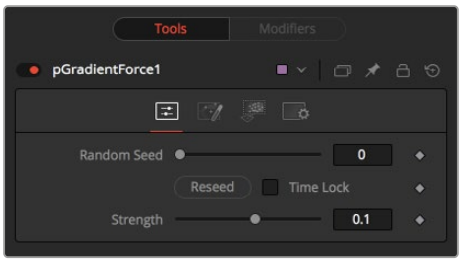
Basic Node Setup

The pGradient Force node is placed in between the pEmitter and pRender nodes. A Fast Noise node is used to create the alpha gradient used to modify the velocity of the particles.



A pGradient Force node using a Fast Noise node as the gradient to modify the particles' motion

Inspector



The pGradient Force controls

Randomize

The Random Seed slider and Randomize button are presented whenever a Fusion node relies on a random result.

Two nodes with the same seed values will produce the same random results. Click the Randomize button to randomly select a new seed value, or adjust the slider to manually select a new seed value.

Strength

Gradient Force has only one specific control, which affects the strength of the force and acceleration applied to the particles. Negative values on this control will cause the Gradient Force to be applied from black to white (low values to high values).

Common Controls

Conditions, Style, Region, and Settings Tabs

The Conditions, Style, Region, and Settings tabs are common to all Particle nodes, so their descriptions can be found in “The Common Controls” section at the end of this chapter.

pImage Emitter [pIE]



The pImage Emitter node

pImage Emitter Node Introduction

The pImage Emitter node takes an input image and treats each pixel of the image as if it were a particle. The main differences between the pImage Emitter and the normal pEmitter is that instead of emitting particles randomly within a given region, this node emits pixels in a regular 2D grid with colors based on the input image.

Inputs

The pImage Emitter node has three inputs. Like most particle nodes, the orange input accepts only other particle nodes. Green and magenta inputs are 2D image inputs for custom image calculations. Optionally, there are teal or white bitmap or mesh inputs, which appear on the node when you set the Region menu in the Region tab to either Bitmap or Mesh.

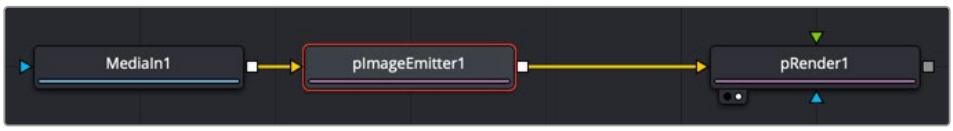
Input: Unlike most other particle nodes, the orange input on the pImage Emitter accepts a 2D image used as the emitter of the particles. If a region is defined for the emitter, this input is used to define the color of the particles.

Style Bitmap Input: This image input accepts a 2D image to use as the particles' image. Since this image duplicates into potentially thousands of particles, it is best to keep these images small and square—for instance, 256 x 256 pixels.

Region: The teal or white region input takes a 2D image or a 3D mesh depending on whether you set the Region menu to Bitmap or Mesh. The color of the input is determined by whichever is selected first in the menu. The 3D mesh or a selectable channel from the bitmap defines the area where the particles are emitted.

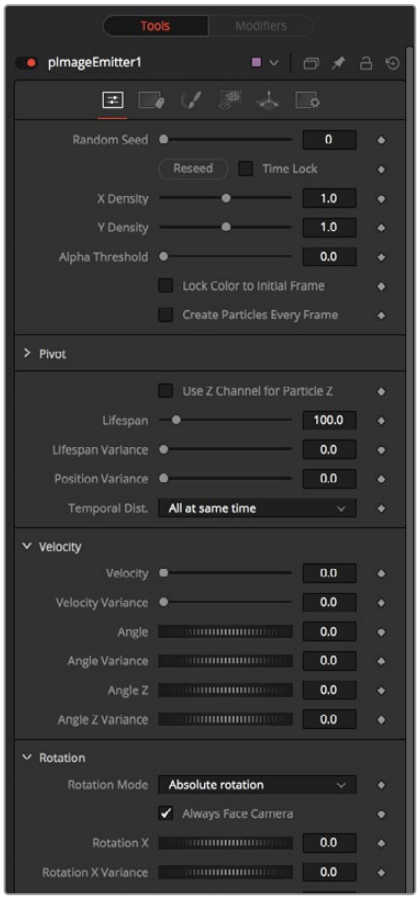
Basic Node Setup

The pImage Emitter node is placed at the start of a particle branch, replacing the location of a pEmitter node. Below, a MediaIn node is used to emit particles using the colors from the clip.



A pImage Emitter node emits particles based on an image connected to the orange input.

Inspector



The pImage Emitter controls

The great majority of controls in this node are identical to those found in the pEmitter, and those controls are documented in that previous section. Below are the descriptions of the controls unique to the pImage Emitter node.

X and Y Density

The X and Y Density sliders are used to set the mapping of particles to pixels for each axis. They control the density of the sampling grid. A value of 1.0 for either slider indicates 1 sample per pixel. Smaller values will produce a looser, more pointillistic distribution of particles, while values above 1.0 will create multiple particles per pixel in the image.

Alpha Threshold

The Alpha Threshold is used for limiting particle generation so that pixels with semitransparent alpha values will not produce particles. This can be used to harden the edges of an otherwise soft alpha channel. The higher the threshold value, the more opaque a pixel must be before it will generate a particle. Note that the default threshold of 0.0 will create particles for every pixel, regardless of alpha, although many may be transparent and invisible.

Lock Particle Color to Initial Frame

Select this checkbox to force the particles to keep the color with which they were born throughout the life of the particle. If this is off, and the input image changes on successive frames, the particles will also change color to match the image. This allows video playback on a grid of particles.

Create Particles Every Frame

Enabling this creates a whole new set of particles every frame, instead of just one set on the frame. This can lead to very large particle systems but allows some interesting effects—for example, if the particles are given some initial velocity or if emitting from an animated source. Try a small velocity, an Angle Z of -90, and a seething Fast Noise as a source to get smoothly varying clouds of particles that you could fly through. Note that if this checkbox is left off, only one set of particles is ever created, and thus animating any of the emitter's other controls will have no effect.

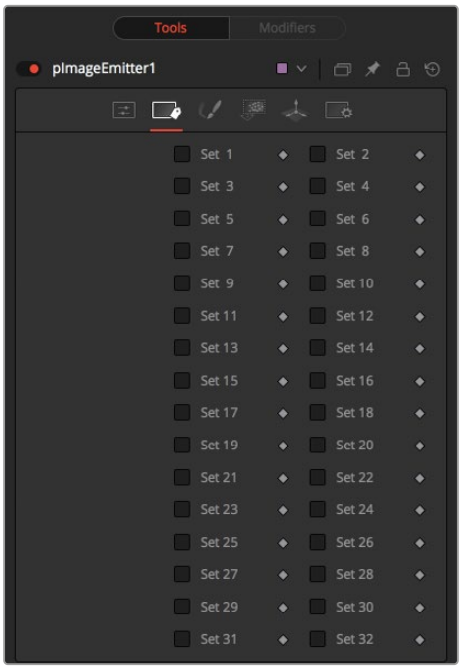
X/Y/Z Pivot

These controls allow you to position the grid of emitted particles.

Use Z Channel for Particle Z

If the input image used to generate the particles has a Z depth channel, that channel can be used to determine the initial position of the particle in Z space. This can have an interesting hollow shell effect when used in conjunction with camera rotation in the pRender node.

Sets Tab



The pImage Emitter Sets tab

NOTE: Pixels with a black (transparent) alpha channel will still generate invisible particles, unless you raise the Alpha Threshold above 0.0. This can slow down rendering significantly.

An Alpha Threshold value of 1/255 = 0.004 is good for eliminating all fully transparent pixels.

The pixels are emitted in a fixed-size 2D grid on the XY plane, centered on the Pivot position. Changing the Region from the default of All allows you to restrict particle creation to more limited areas. If you need to change the size of this grid, use a Transform 3D node after the pRender.

Remember that the various emitter controls apply only to particles when they are emitted. That is, they set the initial state of the particle and do not affect it for the rest of its lifespan. Since pImageEmitter (by default) emits particles only on the first frame, animating these controls will have no effect. However, if the Create Particles Every Frame checkbox is turned on, new particles will be emitted each frame and will use the specified initial settings for that frame.

Common Controls

Conditions, Style, Region, and Settings Tabs

The Conditions, Style, Region, and Settings tabs are common to all Particle nodes, so their descriptions can be found in “The Common Controls” section at the end of this chapter.

pKill [pKl]



The pKill node

pKill Node Introduction

The Kill node is used to destroy (kill) any particle that crosses or intersects its region. It has no specific controls, as it has only one possible affect on a particle. The controls found in the Region tab are normally used to limit this node by restricting the effect to particles that fall within a certain region, age, set, or by reducing the probability of the node applying to a given particle.

Inputs

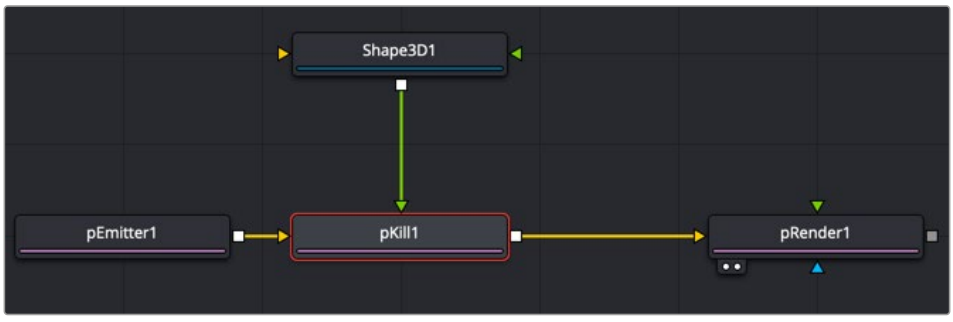
The pKill node has a single orange input by default. Like most particle nodes, this orange input accepts only other particle nodes. A green bitmap or mesh input appears on the node when you set the Region menu in the Region tab to either Bitmap or Mesh.

Input: The orange input takes the output of other particle nodes.

Region: The green or magenta region input takes a 2D image or a 3D mesh depending on whether you set the Region menu to Bitmap or Mesh. The color of the input is determined by whichever is selected first in the menu. The 3D mesh or a selectable channel from the bitmap defines the area particles are killed.

Basic Node Setup

The pKill node is placed in between the pEmitter and pRender nodes. A Shape 3D node is used to create the region where the particles die.



A pKill node using a Shape 3D node as the region where particles die

Inspector

This node only contains common controls in the Conditions and Regions tabs. The Conditions and Regions controls are used to define the location, age, and set of particles that are killed.

Conditions, Style, Region, and Settings Tabs

The Conditions, Style, Region, and Settings tabs are common to all Particle nodes, so their descriptions can be found in “The Common Controls” section at the end of this chapter.

pMerge [pMg]



The pMerge node

pMerge Node Introduction

This node has no controls whatsoever. It serves to combine particles from two streams. Any nodes downstream of the pMerge node will treat the two streams as one.

The combined particles will preserve any sets assigned to them when they were created, making it possible for nodes downstream of the pMerge to isolate specific particles when necessary.

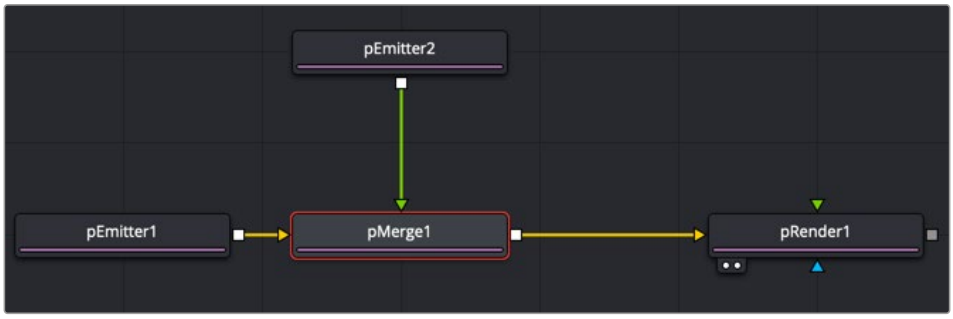
Inputs

The pMerge node has two identical inputs, one orange and one green. These two inputs accept only other particle nodes.

Particle 1 and 2 Input: The two inputs accept two streams of particles and merge them.

Basic Node Setup

The pMerge node connects two pEmitter nodes. The output of the pMerge can go on to feed other particle nodes or to a pRender.



A pMerge node combining two pEmitter nodes.

pPoint Force [pPF]



The pPoint Force node

pPoint Force Node Introduction

This node applies a force to the particles that emanate from a single point in 3D space. The pPoint Force can either attract or repel particles within its sphere of influence. There are four controls specific to the pPoint Force node.

Inputs

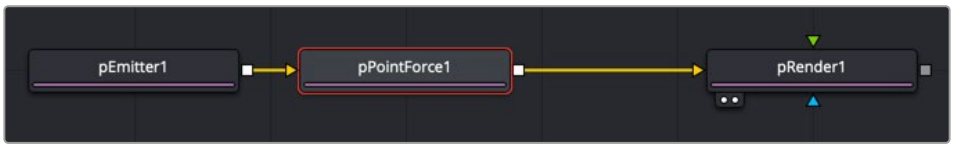
The pPoint Force node has a single orange input by default. Like most particle nodes, this orange input accepts only other particle nodes. A green bitmap or mesh input appears on the node when you set the Region menu in the Region tab to either Bitmap or Mesh.

Input: The orange input takes the output of other particle nodes.

Region: The green or magenta region input takes a 2D image or a 3D mesh depending on whether you set the Region menu to Bitmap or Mesh. The color of the input is determined by whichever is selected first in the menu. The 3D mesh or a selectable channel from the bitmap defines the area where the point force affects the particles.

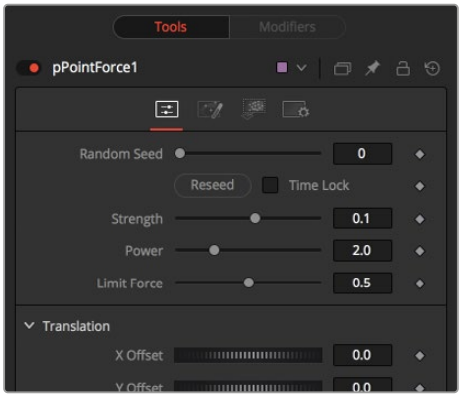
Basic Node Setup

The pPoint Force node is inserted between a pEmitter and a pRender node.



The pPoint Force node positions a tangent force that particles are attracted to or repelled from.

Inspector



The pPoint Force controls

Randomize

The Random Seed slider and Randomize button are presented whenever a Fusion node relies on a random result. Two nodes with the same seed values will produce the same random results. Click the Randomize button to randomly select a new seed value, or adjust the slider to manually select a new seed value.

Strength

This parameter sets the Strength of the force emitted by the node. Positive values represent attractive forces; negative values represent repellent forces.

Power

This determines the degree to which the Strength of the force falls off over distance. A value of zero causes no falloff of strength. Higher values will impose an ever-sharper falloff in strength of the force with distance.

Limit Force

The Limit Force control is used to counterbalance potential problems with temporal sub-sampling. Because the position of a particle is sampled only once a frame (unless sub-sampling is increased in the pRender node), it is possible that a particle can overshoot the Point Force's position and end up getting thrown off in the opposite direction. Increasing the value of this control reduces the likelihood that this will happen.

X, Y, Z Center Position

These controls are used to represent the X, Y, and Z coordinates of the point force in 3D space.

Common Controls

Conditions, Style, Region, and Settings Tabs

The Conditions, Style, Region, and Settings tabs are common to all Particle nodes, so their descriptions can be found in "The Common Controls" section at the end of this chapter.

pRender [pRn]



The pRender node

pRender Node Introduction

The pRender node converts the particle system to either an image or geometry. The default is a 3D particle system, which must be connected to a Renderer 3D to produce an image. This allows the particles to be integrated with other elements in a 3D scene before they are rendered.

Inputs

The pRender node has one orange input, a green camera input, and a blue effects mask input. Like most particle nodes, this orange input accepts only other particle nodes. A green bitmap or mesh input appears on the node when you set the Region menu in the Region tab to either Bitmap or Mesh.

Input: The orange input takes the output of other particle nodes.

Camera Input: The optional green camera input accepts a camera node directly or a 3D scene with a camera connected that is used to frame the particles during rendering.

Effect Mask: The optional blue input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input for 2D particles crops the output of the particles so they are seen only within the mask.

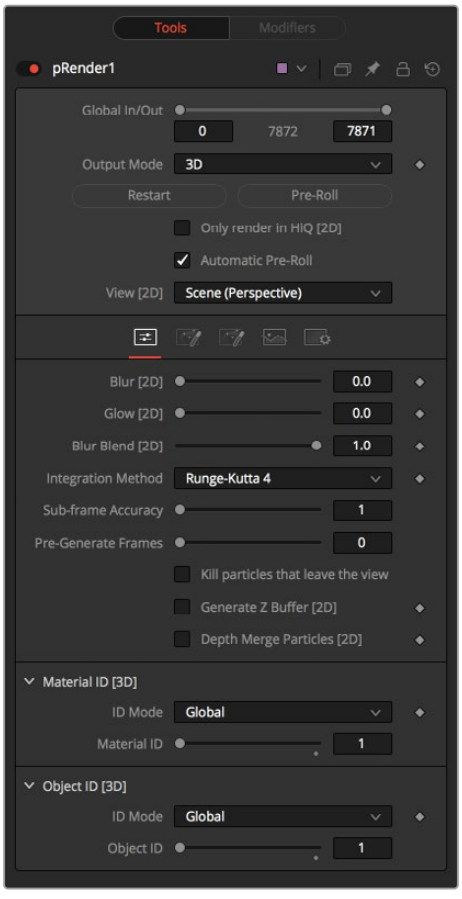
Basic Node Setup

The pRender node is always placed at the end of a particle branch. If the pRender is set to 2D, then the output connects to other 2D nodes like a Merge node. If the pRender is set to 3D, the output connects to a 3D node like a Merge 3D.



All particle branches end with a pRender node.

Inspector



The pRender controls

Output Mode (2D/3D)

While the pRender defaults to 3D output, it can be made to render a 2D image instead. This is done with the 3D and 2D buttons on the Output Mode control. If the pRender is not connected to a 3D-only or 2D-only node, you can also switch it by selecting View > 2D Viewer from the viewer's pop-up menu.

In 3D mode, the only controls in the pRender node that have any effect at all are Restart, Pre-Roll and Automatic Pre-Roll, Sub-Frame Calculation Accuracy, and Pre-Generate frames. The remaining controls affect 2D particle renders only. The pRender node also has a Camera input on the node tree that allows the connection of a Camera 3D node. This can be used in both 2D and 3D modes to allow control of the viewpoint used to render an output image.

Render and the Viewers

When the pRender node is selected in a node tree, all the onscreen controls from Particle nodes connected to it are presented in the viewers. This provides a fast, easy-to-modify overview of the forces applied to the particle system as a whole.

Pre-Roll Options

Particle nodes generally need to know the position of each particle on the last frame before they can calculate the effect of the forces applied to them on the current frame. This makes changing current time manually by anything but single frame intervals likely to produce an inaccurate image.

The controls here are used to help accommodate this by providing methods of calculating the intervening frames.

Restart

This control also works in 3D. Clicking on the Restart button will restart the particle system at the current frame, removing any particles created up to that point and starting the particle system from scratch at the current frame.

Pre-Roll

This control also works in 3D. Clicking on this button causes the particle system to recalculate, starting from the beginning of the render range up to the current frame. It does not render the image produced. It only calculates the position of each particle. This provides a relatively quick mechanism to ensure that the particles displayed in the views are correctly positioned.

If the pRender node is displayed when the Pre-Roll button is selected, the progress of the pre-roll is shown in the viewer, with each particle shown as point style only.

Automatic Pre-Roll

Selecting the Automatic Pre-Roll checkbox causes the particle system to automatically pre-roll the particles to the current frame whenever the current frame changes. This prevents the need to manually select the Pre-Roll button whenever advancing through time in jumps larger than a single frame. The progress of the particle system during an Automatic Pre-Roll is not displayed to the viewers to prevent distracting visual disruptions.

About Pre-Roll

Pre-Roll is necessary because the state of a particle system is entirely dependent on the last known position of the particles. If the current time were changed to a frame where the last frame particle state is unknown, the display of the particle is calculated on the last known position, producing inaccurate results.

To demonstrate:

1. Add a pEmitter and a pRender node to the composition.
2. View the pRender in one of the viewers.
3. Set the Velocity of the particles to 0.1.
4. Place the pEmitter on the left edge of the screen.
5. Set the Current Frame to 0.
6. Set a Render Range from 0–100 and press the Play button.
7. Observe how the particle system behaves.
8. Stop the playback and return the current time to frame 0.
9. In the pRender node, disable the Automatic Pre-Roll option.
10. Use the current time number field to jump to frame 10, and then to frames 60 and 90.

Notice how the particle system only adds to the particles it has already created and does not try to create the particles that would have been emitted in the intervening frames. Try selecting the Pre-Roll button in the pRender node. Now the particle system state is represented correctly.

For simple, fast-rendering particle systems, it is recommended to leave the Automatic Pre-Roll option on. For slower particle systems with long time ranges, it may be desirable to only Pre-Roll manually, as required.

- **Only Render in Hi-Q**
  Selecting this checkbox causes the style of the particles to be overridden when the Hi-Q checkbox is deselected, producing only fast rendering Point-style particles. This is useful when working with a large quantity of slow Image-based or Blob-style particles. To see the particles as they would appear in a final render, simply enable the Hi-Q checkbox.
- **View**
  This drop-down list provides options to determine the position of the camera view in a 3D particle system. The default option of Scene (Perspective) will render the particle system from the perspective of a virtual camera, the position of which can be modified using the controls in the Scene tab. The other options provide orthographic views of the front, top, and side of the particle system.

  It is important to realize that the position of the onscreen controls for Particle nodes is unaffected by this control. In 2D mode, the onscreen controls are always drawn as if the viewer were showing the front orthographic view. (3D mode gets the position of controls right at all times.)

  The View setting is ignored if a Camera 3D node is connected to the pRender node's Camera input on the node tree, or if the pRender is in 3D mode.

Conditions

- **Blur, Glow, and Blur Blend**
  When generating 2D particles, these sliders apply a Gaussian blur, glows, and blur blending to the image as it is rendered, which can be used to soften the particles and blend them. The result is no different than adding a Blur after the pRender node in the node tree.

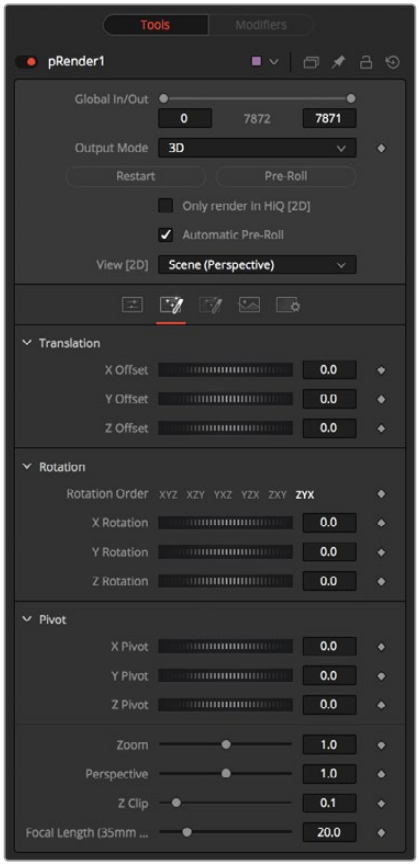
- **Sub-Frame Calculation Accuracy**
  This determines the number of sub-samples taken between frames when calculating the particle system. Higher values increase the accuracy of the calculation but also increase the amount of time to render the particle system.
- **Pre-Generate Frames**
  This control is used to cause the particle system to pre-generate a set number of frames before its first valid frame. This is used to give a particle system an initial state from which to start.

  A good example of when this might be useful is in a shot where particles are used to create the smoke rising from a chimney. Set Pre-Generate Frames to a number high enough to ensure that the smoke is already present in the scene before the render begins, rather than having it just starting to emerge from the emitter for the first few frames.
- **Kill Particles That Leave the View**
  Selecting this checkbox control automatically destroys any particles that leave the visible boundaries of the image. This can help to speed render times. Particles destroyed in this fashion never return, regardless of any external forces acting upon them.Generate Z Buffer

  Selecting this checkbox causes the pRender node to produce a Z Buffer channel in the image. The depth of each particle is represented in the Z Buffer. This channel can then be used for additional depth operations like Depth Blur, Depth Fog, and Downstream Z Merging.

  Enabling this option is likely to increase the render times for the particle system dramatically.
- **Depth Merge Particles**
  Enabling this option causes the particles to be merged using Depth Merge techniques, rather than layer-based techniques.

Scene Tab



The pRender Scene tab

Z Clip

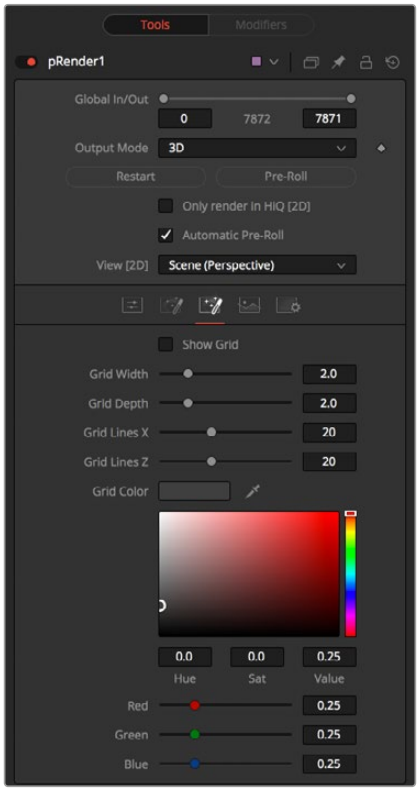
The Z Clip control is used to set a clipping plane in front of the camera. Particles that cross this plane are clipped, preventing them from impacting on the virtual lens of the camera and dominating the scene.

Grid Tab

These controls do not apply to 3D particles.

The grid is a helpful, non-rendering display guide used to orient the 2D particles in 3D space. The grid is never seen in renders, just like a center crosshair is never seen in a render. The width, depth, number of lines, and grid color can be set using the controls found in this tab.

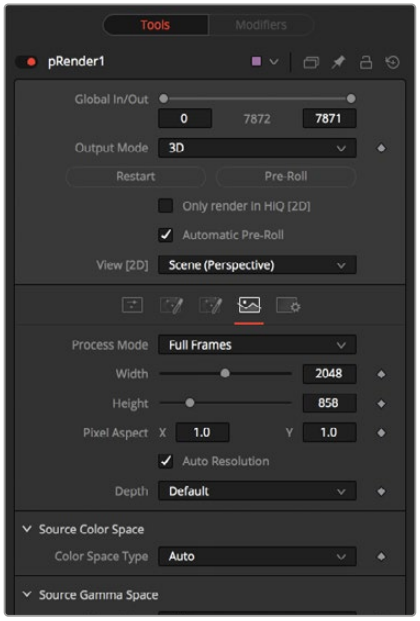
These controls cannot be animated.



The pRender Grid tab

Image Tab

The controls in this tab are used to set the resolution, color depth, and pixel aspect of the rendered image produced by the node.



The pRender Image tab

Process Mode

Use this menu control to select the Fields Processing mode used by Fusion to render changes to the image. The default option is determined by the Has Fields checkbox control in the Frame Format preferences.

Use Frame Format Settings

When this checkbox is selected, the width, height, and pixel aspect of the rendered images by the node will be locked to values defined in the composition's Frame Format preferences. If the Frame Format preferences change, the resolution of the image produced by the node will change to match. Disabling this option can be useful to build a composition at a different resolution than the eventual target resolution for the final render.

Width/Height

This pair of controls is used to set the Width and Height dimensions of the image to be rendered by the node.

Pixel Aspect

This control is used to specify the Pixel Aspect ratio of the rendered particles. An aspect ratio of 1:1 would generate a square pixel with the same dimensions on either side (like a computer display monitor), and an aspect of 0.9:1 would create a slightly rectangular pixel (like an NTSC monitor).

NOTE: Right-click on the Width, Height, or Pixel Aspect controls to display a menu listing the file formats defined in the preferences Frame Format tab. Selecting any of the listed options will set the width, height, and pixel aspect to the values for that format, accordingly.

Depth

The Depth menu is used to set the pixel color depth of the particles. 32-bit pixels require 4X the memory of 8-bit pixels but have far greater color accuracy. Float pixels allow high dynamic range values outside the normal 0...1 range, for representing colors that are brighter than white or darker than black.

Source Color Space

You can use the Source Color Space menu to set the Color Space of the footage to help achieve a linear workflow. Unlike the Gamut tool, this doesn't perform any actual color space conversion, but rather adds the source space data into the metadata, if that metadata doesn't exist. The metadata can then be used downstream by a Gamut tool with the From Image option, or in a Saver, if explicit output spaces are defined there. There are two options to choose from:

- **Auto:** Automatically reads and passes on the metadata that may be in the image.
- **Space:** Displays a Color Space Type menu where you can choose the correct color space of the image.

Source Gamma Space

Using the Curve type menu, you can set the Gamma Space of the footage and choose to remove it by way of the Remove Curve checkbox when working in a linear workflow. There are three choices in the Curve type menu:

- **Auto:** Automatically reads and passes on the metadata that may be in the image.
- **Space:** Displays a Gamma Space Type menu where you can choose the correct gamma curve of the image.
- **Log:** Brings up the Log/Lin settings, similar to the Cineon tool.

For more information, *see Chapter 39, "Film Nodes,"* in the Fusion Reference Manual.

Remove Curve

Depending on the selected Gamma Space or on the Gamma Space found in Auto mode, the Gamma Curve is removed from, or a log-lin conversion is performed on, the material, effectively converting it to a linear output space.

Motion Blur

As with other 2D nodes in Fusion, Motion Blur is enabled from within the Settings tab. You may set Quality, Shutter Angle, Sample Center, and Bias, and Blur will be applied to all moving particles.

NOTE: Motion Blur on 3D mode particles (rendered with a Renderer 3D) also requires that identical motion blur settings are applied to the Renderer 3D node.

pSpawn [pSp]



The pSpawn node

pSpawn Node Introduction

The pSpawn node makes each affected particle act as an emitter that can produce one or more particles of its own. The original particle continues until the end of its lifespan, and each of the particles it emits becomes wholly independent with a lifespan and properties of its own.

As long as a particle falls under the effect of the pSpawn node, it will continue to generate particles. It is important to restrict the effect of the node with limiters like Start and End Age, Probability, Sets and Regions, and by animating the parameters of the emitter so that the node is operative only when required.

Inputs

By default, the pSpawn node has a single orange input. Like most particle nodes, this orange input accepts only other particle nodes. You can enable an image input by selecting Bitmap from the Style menu in the Style tab. Also, two region inputs, one for bitmap and one for mesh, appear on the node when you set the Region menu in the Region tab to either Bitmap or Mesh. The colors of these inputs change depending on the order they are enabled.

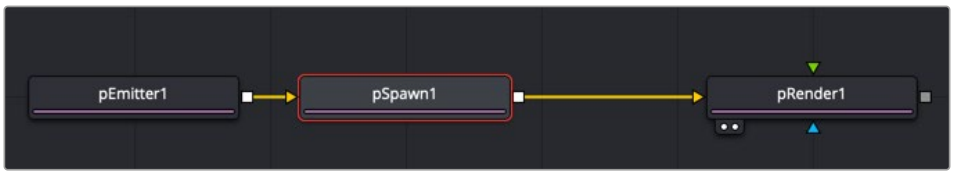
Input: The orange input accepts the output of other particle nodes.

Style Bitmap Input: This image input accepts a 2D image to use as the particles' image. Since this image duplicates into potentially thousands of particles, it is best to keep these images small and square—for instance, 256 x 256 pixels.

Region: The region inputs take a 2D image or a 3D mesh depending on whether you set the Region menu to Bitmap or Mesh. The color of the input is determined by whichever is selected first in the menu. The 3D mesh or a selectable channel from the bitmap defines the area where the particles are emitted.

Basic Node Setup

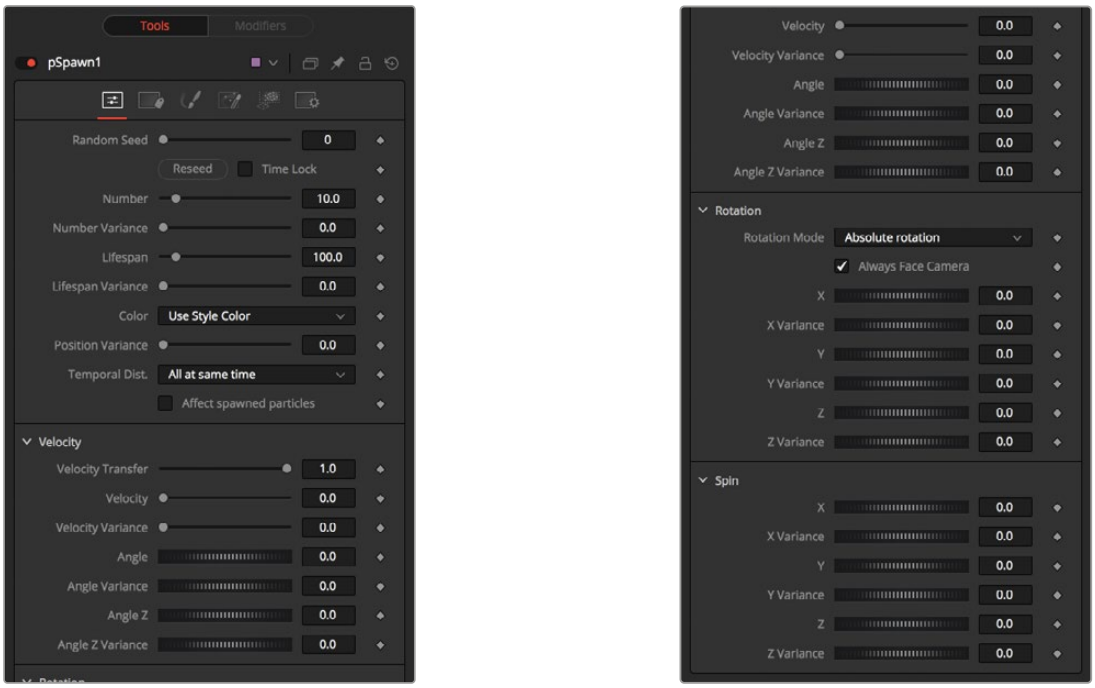
The pSpawn node is placed between the pEmitter and pRender nodes. Using the Age parameter in the pSpawn's Conditions tab, you can spawn new particles as the old ones die off. This is one way to have a trail of particles shoot up in the air like a rocket and burst into sparkling fireworks.



A pSpawn node used to generate new particles at specific points in the old particles' life

Inspector

The pSpawn node has a large number of controls, most of which exactly duplicate those found within the pEmitter node. There are a few controls that are unique to the pSpawn node, and their effects are described below.



The pSpawn controls

Affect Spawned Particles

Selecting this checkbox causes particles created by spawning to also become affected by the pSpawn node on subsequent frames. This can exponentially increase the number of particles in the system, driving render times up to an unreasonable degree. Use this checkbox cautiously.

Velocity Transfer

This control determines how much velocity of the source particle is transferred to the particles it spawns. The default value of 1.0 causes each new particle to adopt 100 percent of the velocity and direction from its source particle. Lower values will transfer less of the original motion to the new particle.

Common Controls

Conditions, Style, Region, and Settings Tabs

The Conditions, Style, Region, and Settings tabs are common to all Particle nodes, so their descriptions can be found in “The Common Controls” section at the end of this chapter.

pTangent Force [pTF]



The pTangent Force node

pTangent Force Node Introduction

This node is used to apply a tangential force to the particles—a force that is applied perpendicularly to the vector between the pTangent Force's region and the particle it is affecting.

Inputs

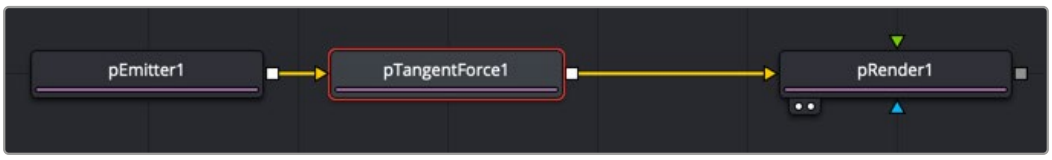
The pTangent Force node has a single orange input by default. Like most particle nodes, this orange input accepts only other particle nodes. A green bitmap or mesh input appears on the node when you set the Region menu in the Region tab to either Bitmap or Mesh.

Input: The orange input takes the output of other particle nodes.

Region: The green or magenta region input takes a 2D image or a 3D mesh depending on whether you set the Region menu to Bitmap or Mesh. The color of the input is determined by whichever is selected first in the menu. The 3D mesh or a selectable channel from the bitmap defines the area where the tangent force effects the particles.

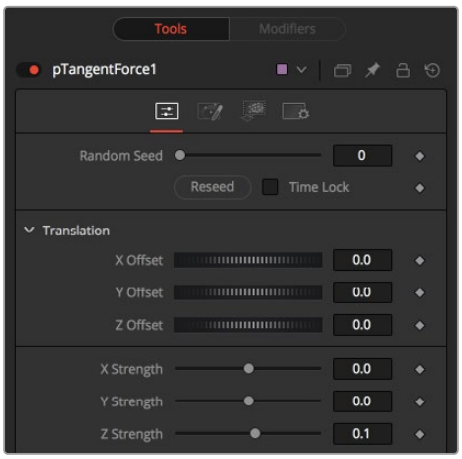
Basic Node Setup

The pTangent Force node is inserted between a pEmitter and a pRender node.



The pTangent Force node positions a tangent force that particles maneuver around.

Inspector



The pTangent Force controls

The controls for this node are used to position the offset in 3D space and to determine the strength of the tangential force along each axis independently.

Randomize

The Random Seed slider and Randomize button are presented whenever a Fusion node relies on a random result.

Two nodes with the same seed values will produce the same random results. Click the Randomize button to randomly select a new seed value, or adjust the slider to manually select a new seed value.

X, Y, Z Center Position

These controls are used to represent the X, Y, and Z coordinates of the Tangent force in 3D space.

X, Y, Z Center Strength

These controls are used to determine the Strength of the Tangent force in 3D space.

Common Controls

Conditions, Style, Region, and Settings Tabs

The Conditions, Style, Region, and Settings tabs are common to all Particle nodes, so their descriptions can be found in “The Common Controls” section at the end of this chapter.

pTurbulence [pTr]



The pTurbulence node

pTurbulence Node Introduction

The pTurbulence node imposes a frequency-based chaos on the position of each particle, causing the motion to become unpredictable and uneven. The controls for this node affect the strength and density of the Turbulence along each axis.

Inputs

The pTurbulence node has a single orange input by default. Like most particle nodes, this orange input accepts only other particle nodes. A green bitmap or mesh input appears on the node when you set the Region menu in the Region tab to either Bitmap or Mesh.

Input: The orange input takes the output of other particle nodes.

Region: The green or magenta region input takes a 2D image or a 3D mesh depending on whether you set the Region menu to Bitmap or Mesh. The color of the input is determined by whichever is selected first in the menu. The 3D mesh or a selectable channel from the bitmap defines the area of turbulence.

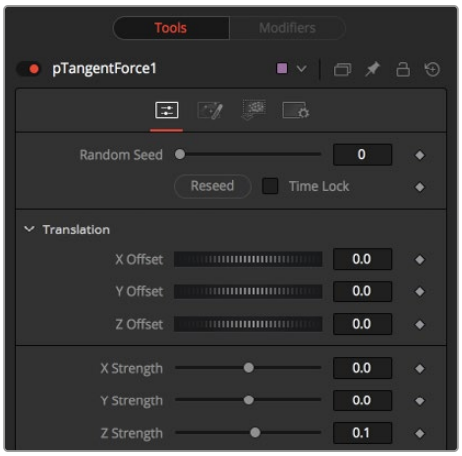
Basic Node Setup

The pTurbulence node is inserted between a pEmitter and a pRender node.



The pTurbulence node disturbs the rigid flow of particles for a more natural motion.

Inspector



The pTurbulence controls

Randomize

The Random Seed slider and Randomize button are presented whenever a Fusion node relies on a random result. Two nodes with the same seed values will produce the same random results. Click the Randomize button to randomly select a new seed value, or adjust the slider to manually select a new seed value.

X, Y, and Z Strength

The Strength control affects the amount of chaotic motion imparted to particles.

Strength Over Life

This mini Spline Editor control can be used to control the amount of turbulence applied to a particle according to its age. For example, a fire particle may originally have very little turbulence applied at the start of its life, and as it ages, the turbulence increases.

Density

Use this control to adjust the density in the turbulence field. Lower values causes more particle cells to be affected similarly, almost as if “waves” of the turbulence field run through the particles, affecting groups of cells at the same time. Higher values add finer variations to more individual particle cells causing more of a spread in the turbulence field.

Common Controls

Conditions, Style, Region, and Settings Tabs

The Conditions, Style, Region, and Settings tabs are common to all Particle nodes, so their descriptions can be found in “The Common Controls” section at the end of this chapter.

pVortex [pVt]



The pVortex node

pVortex Node Introduction

The pVortex node applies a rotational force to each particle, causing them to be drawn toward the source of the Vortex. In addition to the Common Particle Controls, the pVortex node also has the following controls.

Inputs

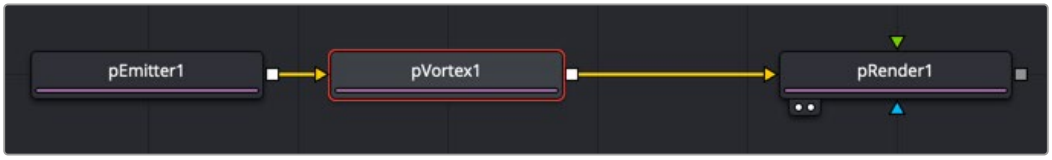
The pVortex node has a single orange input by default. Like most particle nodes, this orange input accepts only other particle nodes. A green bitmap or mesh input appears on the node when you set the Region menu in the Region tab to either Bitmap or Mesh.

Input: The orange input takes the output of other particle nodes.

Region: The green or magenta region input takes a 2D image or a 3D mesh depending on whether you set the Region menu to Bitmap or Mesh. The color of the input is determined by whichever is selected first in the menu. The 3D mesh or a selectable channel from the bitmap defines the area of the vortex.

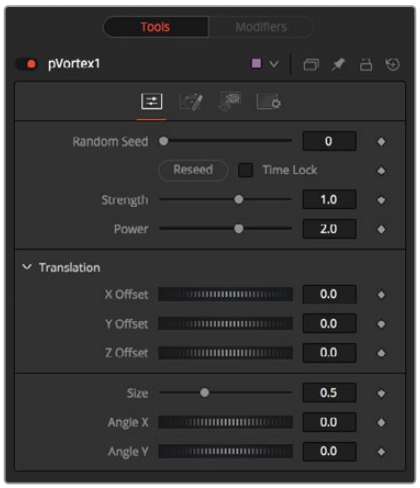
Basic Node Setup

The pVortex node is placed in between the pEmitter and pRender nodes.



A pVortex node creates a spiraling motion for particles that fall within its pull.

Inspector



The pVortex controls

Randomize

The Random Seed slider and Randomize button are presented whenever a Fusion node relies on a random result. Two nodes with the same seed values will produce the same random results. Click the Randomize button to randomly select a new seed value, or adjust the slider to manually select a new seed value.

Strength

This control determines the Strength of the Vortex Force applied to each particle.

Power

This control determines the degree to which the Strength of the Vortex Force falls off with distance.

X, Y, and Z Offset

Use these sliders to set the amount by which the vortex Offsets the affected particles.

Size

This is used to set the Size of the Vortex Force.

Angle X and Y

These sliders control the amount of rotational force applied by the Vortex along the X and Y axes.

Common Controls

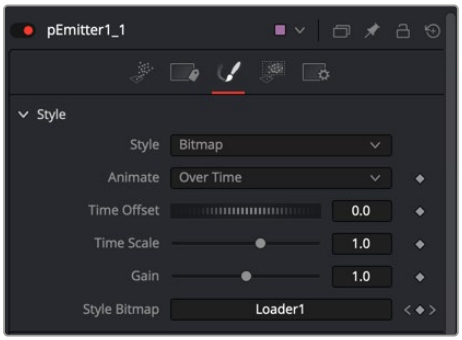
Conditions, Style, Region, and Settings Tabs

The Conditions, Region, and Settings tabs are common to all Particle nodes, so their descriptions can be found in the following “The Common Controls” section.

The Common Controls

Particle nodes share a number of identical controls in the inspector. This section describes the Style, Conditions, Region, and Settings tabs that are common among particle nodes.

Inspector



The pEmitter Style tab

Style Tab

The Style Tab is common to the pEmitter, pSpawn, pChangeStyle, and pImage Emitter. It controls the appearance of the particles using general controls like type, size, and color.

Style

The Style menu provides access to the various types of particles supported by the Particle Suite. Each style has its specific controls, as well as controls it will share with other styles.

- **Point:** This option produces particles precisely one pixel in size. Controls that are specific to Point Style are Apply Mode and Sub Pixel Rendered.

  Apply Mode: This menu applies only to 2D particles. 3D particle systems are not affected. It includes Apply modes for Add and Merge. Add combines overlapping particles by adding together the color values of each particle. Merge uses a simple over operation to combine overlapping particles.

  Sub Pixel Rendered: This checkbox determines whether the point particles are rendered with Sub Pixel precision, which provides smoother-looking motion but blurrier particles that take slightly longer to render.

- **Bitmap:** This style produces particle cells based on an image file or another node in the Node editor. When this option is selected an orange image input appears on the node in the node editor. There are several controls for affecting the appearance and animation. In addition to the controls in the Style section, a Merge section is displayed at the bottom of the inspector when Bitmap is selected as the Style. The Merge section includes controls for additive or subtractive merges when the particle cells overlap.

  Animate Over Time: This menu includes three options for determining how movie files play when they are used as particle cell bitmaps. The Over Time setting plays the movie file sequentially. For instance, when the comps is on frame 2, frame 2 of the movie file is displayed, when the comp is on frame 3, frame 3 of the movie files is displayed and so on. If a particle cell is not generated until frame 50, it begins with frame 50 of the movie file. This causes all particle cells to use the same image on any give frame of the comp. The Particle Age setting causes each particle cell to begin with the first frame of the movie file, regardless of when the particle cell is generated.

The Particle Birth Time setting causes each particle to begin with the frame that coincides with the frame of the particle cell birth time. For instance, if the particle is generated on frame 25, then it uses frame 25 of the movie file for the entire comp. Unlike the other two options, the Particle Birth Time setting holds the same frame for the duration of the comp

Time Offset: This dial is used to slip or offset the starting frame used from the movie file. For instance, setting it to 10 will cause the movie file to start at frame 10 instead of frame 1.

Time Scale: This slider is a multiplier on the frame. Instead of using an offset, it changes the starting frame by multiplying the frame by the value selected with the slider. For instance, if a value of 2 is selected then when the playhead reaches frame 2, the movie files displays frame 4 (2x2=4) and when the playhead reaches frame 8, the movie file displays frame 16 (8x2=16).

Gain: The gain slider is a multiplier of the pixel value. It is used to apply a correction to the overall Gain of the Bitmap. Let's say you have a bitmap particle cell that contains a pixel value of R0.5 G0.5 B0.4 and you add a Gain of 1.2, you end up with a pixel value of R0.6 G0.6, B0.48 (i.e., 0.4 * 1.2 = 0.48) while leaving black pixels unaffected. Higher values produce a brighter image, whereas lower values reduce both the brightness and the transparency of the image.

Style Bitmap: This control appears when the Bitmap style is selected, along with an orange Style Bitmap input on the node's icon in the Node view. Connect a 2D node to this input to provide images to be used for the particles. You can do this on the Node view, or you may drag and drop the image source node onto the Style Bitmap control from the Node Editor or Timeline, or right-click on the control and select the desired source from the Connect To menu.

- **Blob:** This option produces large, soft spherical particles, with controls for Color, Size, Fade timing, Merge method, and Noise.

  Noise: This slider only applies to 2D Blob particles. The noise slider is used to introduce a computer generated Perlin noise pattern into the blob particles in order to give the blobs more texture. A setting of 0 introduces no noise to the Blob particles and a setting of 1 introduces the maximum amount of noise.

- **Brush:** This styles produces particle cells based on any image file located in the brushes directory. There are numerous controls for affecting the appearance and animation.

  Gain: The gain slider is a multiplier of the pixel value. It is used to apply a correction to the overall Gain of the image that is used as the Brush. Let's say you have a brush particle cell that contains a pixel value of R0.5 G0.5 B0.4 and you add a Gain of 1.2, you end up with a pixel value of R0.6 G0.6, B0.48 (i.e., 0.4 * 1.2 = 0.48) while leaving black pixels unaffected. Higher values produce a brighter image, whereas lower values reduce both the brightness and the transparency of the image.

  Brush: This menu shows the names of any image files stored in the Brushes directory. The location of the Brushes directory is defined in the Preferences dialog, under Path Maps. The default is the Brushes subdirectory within Fusion's install folder.

  Use Aspect From: The Use Aspect From menu includes three settings for the aspect ratio of the brush image. You can choose image format to use the brush image's native aspect ration. Choose Frame Format to use the aspect ratio set in the Frame Format Setting in the Fusion Preferences, or choose Custom to enter your own Pixel X and Y dimensions.

- **Line:** This style produces straight line-type particles with optional "falloff." The Size to Velocity control described below (under Size Controls) is often useful with this Line type. The Fade control adjusts the amount of falloff over the length of the line.

- **Point Cluster:** This style produces small clusters of single-pixel particles. Point Clusters are similar to the Point style; however, they are more efficient when a large quantity of particles is required. This style shares parameters with the Point style. Additional controls specific to Point Cluster style are Number of Points and Number Variance.

  Sub Pixel Rendered: This checkbox determines whether the point particles are rendered with Sub Pixel precision, which provides smoother-looking motion but blurrier particles that take slightly longer to render.

  Number of Points and Variance: The value of this control determines how many points are in each Point Cluster.

Color Controls

The Color Controls select the color and Alpha values of the particles generated by the emitter.

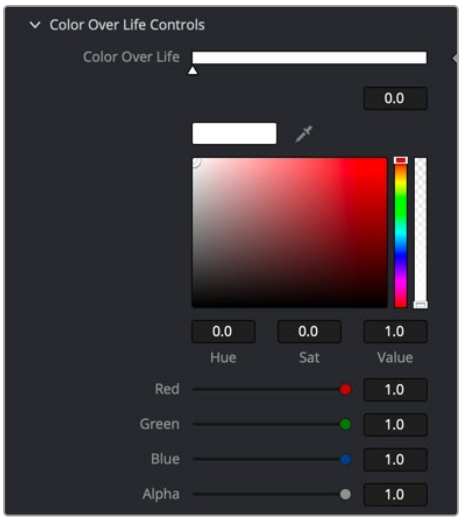
Color Variance

These range controls provide a means of expanding the colors produced by the pEmitter. Setting the Red variance range at -0.2 to +0.2 will produce colors that vary 20% on either side of the red channel, for a total variance of 40%. If the pEmitter is set to produce R0.5, G0.5, B0.5 (pure gray), the variance shown above will produce points with a color range between R0.3, G0.5, B0.5, and R0.7, G0.5, B0.5.

To visualize color space as values between 0-256 or as 0-65535, change the values used by Fusion using the Show Color As option provided in the General tab within the Preferences dialog.

Lock Color Variance

This checkbox locks the color variance of the particles. Unlocking this allows the color variance to be applied differently to each color channel, giving rise to a broader range of colors.



Particles Color Over Life controls

Color Over Life

This standard gradient control allows for the selection of a range of color values to which the particle will adhere over its lifetime.

The left point of the gradient represents the particle color at birth. The right point shows the color of the particle at the end of its lifespan.

Additional points can be added to the gradient control to cause the particle color to shift throughout its life.

This type of control can be useful for fire-type effects (for example, the flame may start blue, turn orange, and end a darker red). The gradient itself can be animated over time by right-clicking on the control and selecting Animate from the contextual menu. All points on the gradient will be controlled by a single Color Over Life spline, which controls the speed at which the gradient itself changes. You may also use the From Image modifier, which produces a gradient from the range of colors in an image along a line between two points.



Particles Size and Fade controls

Size Controls

The majority of the Size Controls are self-explanatory. The Size and Size Variance controls are used to determine the size and degree of size variation for each particle. It is worth noting that the Point style does not have size controls (each point is a single pixel in size, and there is no additional control).

When a Bitmap Particle style is used, a value of 1.0 indicates that each particle should be the same size as the input bitmap. A value of 2.0 will scale the particle up in size by 200%. For the best quality particles, always try to make the input bitmap as big, or bigger, than the largest particle produced by the system.

For the Point Cluster style, the size control adjusts the density of the cluster, or how close together each particle will get.

There are additional size controls that can be used to adjust further the size of particles based on velocity and depth.

Size to Velocity

This increases the size of each particle relative to the velocity or speed of the particle. The velocity of the particle is added to the size, scaled by the value of this control.

1.0 on this control, such as for a particle traveling at 0.1, will add another 0.1 to the size (velocity * size to velocity + size = new size). This is most useful for Line styles, but the control can be used to adjust the size of any style.

Size Z Scale

This control measures the degree to which the size of each particle changes according to its Z position. The effect is to exaggerate or reduce the impact of perspective. The default value is 1.0, which provides a relatively realistic perspective effect.

Objects on the focal plane (Z = 0.0) will be actual-sized. Objects farther along Z will become smaller. Objects closer along Z will get larger.

A value of 2.0 will exaggerate the effect dramatically, whereas a value of 0.0 will cancel the effects of perspective entirely.

Size Over Life

This spline control determines the size of a particle throughout its lifespan. The vertical scale represents a percentage of the value defined by the Size control, from 0 to 200%. The horizontal scale represents a percentage of the particle's lifespan (0 to 100%).

This graph supports all the features available to a standard spline editor. These features can be accessed by right-clicking on the graph. It is also possible to view and edit the graph spline in the larger Spline Editor.

Fade Controls

This simple range slider provides a mechanism for fading a particle at the start and end of its lifetime. Increasing the Fade In value will cause the particle to fade in at the start of its life. Decreasing the Fade Out value will cause the particle to fade out at the end of its life.

This control's values represent a percentage of the particle's overall life, therefore, setting the Fade In to 0.1 would cause the particle to fade in over the first 10% of its total lifespan. For example, a particle with a life of 100 frames would fade in from frame 0...10.

Merge Controls

This set of particle controls affects the way individual particles are merged together. The Subtractive/Additive slider works as documented in the standard Merge node. The Burn-In control will cause the particles to overexpose, or "blow out," when they are combined.

None of the Merge controls will have any effect on a 3D particle system.



Particles Blur controls

Blur Controls

This set of particle controls can be used to apply a Blur to the individual particles. Blurring can be applied globally, by age, or by Z depth position.

None of the Blur controls will have any effect on a 3D particle system.

Blur (2D) and Blur Variance (2D)

These controls apply blur to each particle. Unlike the Blur in the pRender node, this is applied to each particle independently before the particles are merged together. The Blur Variance slider modifies the amount of blur applied to each particle.

Blur Over Life

This spline graph controls the amount of blur that is applied to the particle over its life. The vertical scale represents a percentage of the value defined by the Blur control. The horizontal scale represents a percentage of the particle's lifespan.

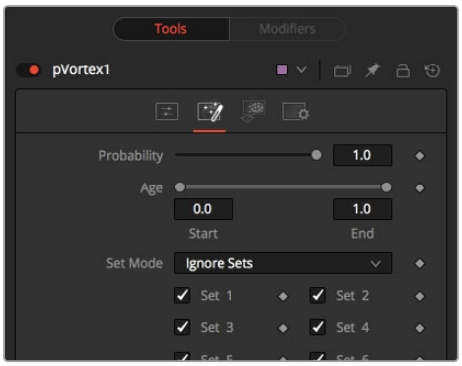
This graph supports all of the features available to a standard Spline Editor. These features can be accessed by right-clicking on the graph. It is also possible to view and edit the spline in the larger Spline editor.

Z Blur (DoF) (2D) and DoF Focus

This slider control applies blur to each particle based on its position along the Z axis.

The DoF Focus range control is used to determine what area of the image remains in focus. Lower values along Z are closer to the camera. Higher values are farther away. Particles within the range will remain in focus. Particles outside that range will have the blur defined by the Z Blur control applied to them.

Conditions



pVortex Conditions tab

Conditions Tab

The Conditions tab limits the particles that are affected by the node's behavior. You can limit the particle using probability or more specifically using sets.

Probability

The Probability slider determines the percentage of chance that the node affects any given particle.

The default value of 1.0 affects all particles. A setting of 0.6 would mean that each particle has a 60 percent chance of being affected by the control.

Probability is calculated for each particle on each frame. For example, a particle that is not affected by a force on one frame has the same chance of being affected on the next frame.

Start/End Age

This range control can be used to restrict the effect of the node to a specified percentage of the particle lifespan.

For example, to restrict the effect of a node to the last 20 percent of a particle's life, set the Start value to 0.8, and the End value remains at 1.0. The node on frames 80 through 100 only affects a particle with a lifespan of 100 frames.

Set Mode Menu

The Set Mode menu drives how the particle node influences the active particle sets. There are three options from this menu:

- **Ignore Sets:** The particle node disregards the state of the Set checkboxes and applies to all nodes.
- **Affect Specified Sets:** The particle node applies its behavior to the active Set checkboxes only.
- **Ignore Specified Sets:** The particle node applies its behavior to the inactive Set checkboxes only.

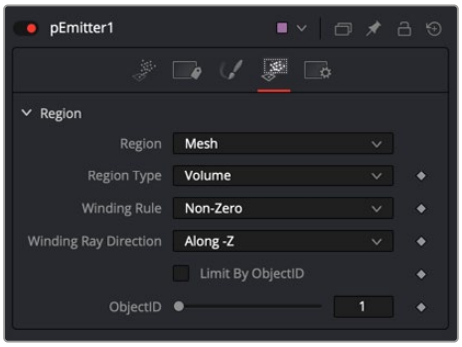
Set

The state of a Set # checkbox determines if the Particle node's effect will be applied to the particles in the set. It allows you to limit the effects of some nodes to a subset of particles.

Sets are assigned by the nodes that create particles. These include the pEmitter, pImage Emitter, pChangeStyle, and the pSpawn nodes.

Region Tab

The Region tab is used to restrict the node's effect to a geometric region or plane, and to determine the area where particles are created if it's a pEmitter node or where the behavior of a node has influence.



pEmitter Regions tab set to mesh

The Region tab is common to almost all particle nodes. In the pEmitter node Emitter Regions are used to determine the area where particles are created. In most other tools it is used to restrict the tool's effect to a geometric region or plane. There are seven types of regions, each with its own controls. Only one emitter region can be set for a single pEmitter node. If the pRender is set to 2D, then the emitter region will produce particles along a flat plane in Z Space. 3D emitter regions possess depth and can produce particles inside a user-defined, three-dimensional region.

Region Mode Menu

The Region Mode menu includes seven types of regions to define the area, each with its controls.

- **All:** In 2D, the particles will be created anywhere within the boundaries of the image. In 3D, this region describes a cube 1.0 x 1.0 x 1.0 units in size.

- **Bézier:** Bézier mode uses a user-created polyline to determine the region where particles are created. The Bézier mode works in both 2D and 3D modes; however, the Bézier polyline region can only be created in 2D.

  To animate the shape of the polyline over time or to connect it to another polyline, right-click the Shape animation label at the bottom of the inspector and select the appropriate option from the drop-down menu.
- **Bitmap:** A Bitmap source from one of the other nodes in the composition will be used as the region where particles are born.
- **Cube:** A full 3D Cube is used to determine the region within which particles are created. The height, width, depth, and XYZ positions can all be determined by the user and be animated over time.
- **Line:** A simple line control determines where particles are created. The line is composed of two end-points, which can be connected to Paths or Trackers, as necessary. This type of emitter region includes X, Y, and Z position controls for the start and end of the line
- **Mesh:** Any 3D Mesh can be used as a region. In Mesh mode, the region can also be restricted by the Object ID using the ObjectID slider. See below for a more in-depth explanation of how mesh regions work.
- **Rectangle:** The Rectangle region type is like the Cube type, except that this region has no depth in Z space. Unlike other 2D emitter regions, this region can be positioned and rotated in Z space.
- **Sphere:** This is a spherical 3D emitter region with Size and Center Z controls. Sphere (3D) is the default region type for a new pEmitter node.

Mesh Regions

Region Type

The Region Type drop-down menu allows you to choose whether the region will include the inner volume or just the surface. For example, with a pEmitter mesh region, this determines if the particles emit from the surface or the full volume.

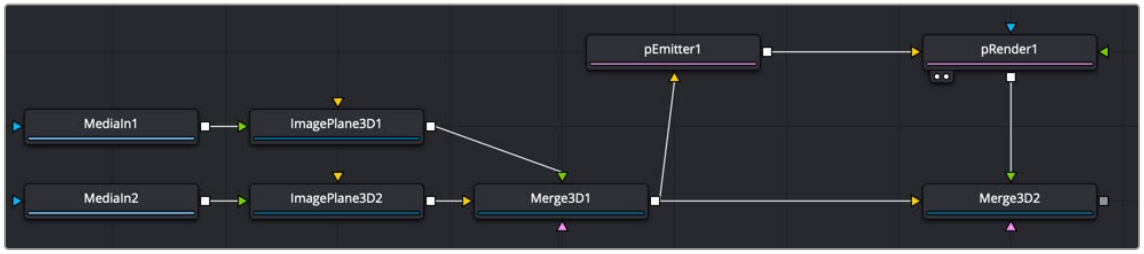
Winding Rule and Winding Ray Direction

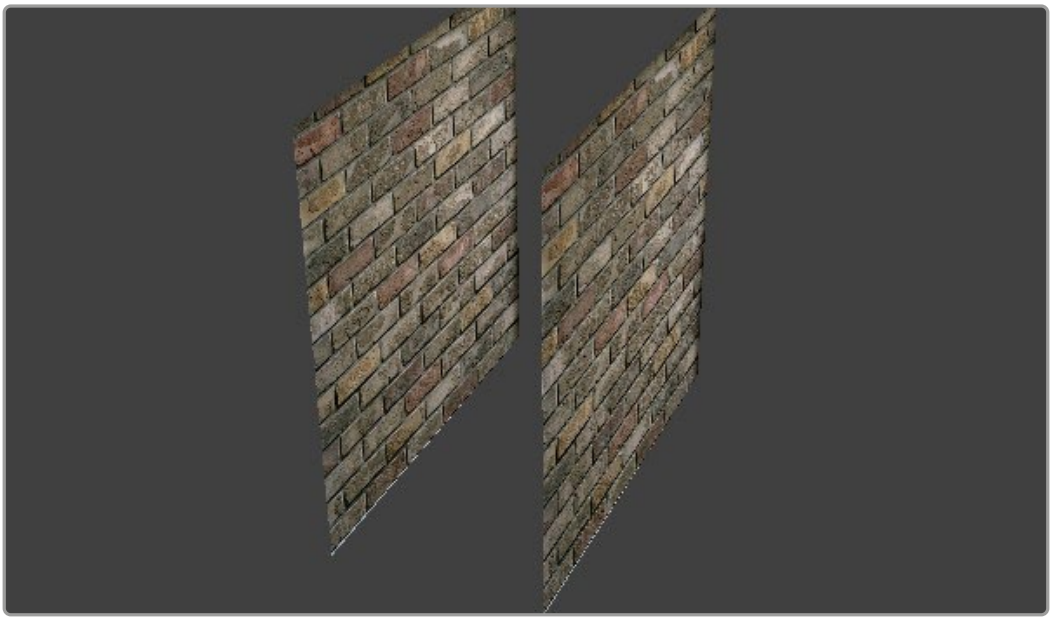
The Winding Rule and Winding Ray Direction parameters determine how the mesh region handles particle creation with meshes that are not closed, as is common in many meshes imported from external applications. This scenario is common with imported mesh geometry, and even geometry that appears closed will frequently appear to “leak” thanks to improperly welded vertices.

To determine if a particle is in the interior of an object, a ray is cast from infinity through that particle and then out to -infinity. The Winding Ray Direction determines which direction this ray is cast in. Each time a surface is pierced by the ray, it is recorded and added onto a total to generate a winding number. Going against a surfaces normal counts as +1, and going with the normal counts as -1.

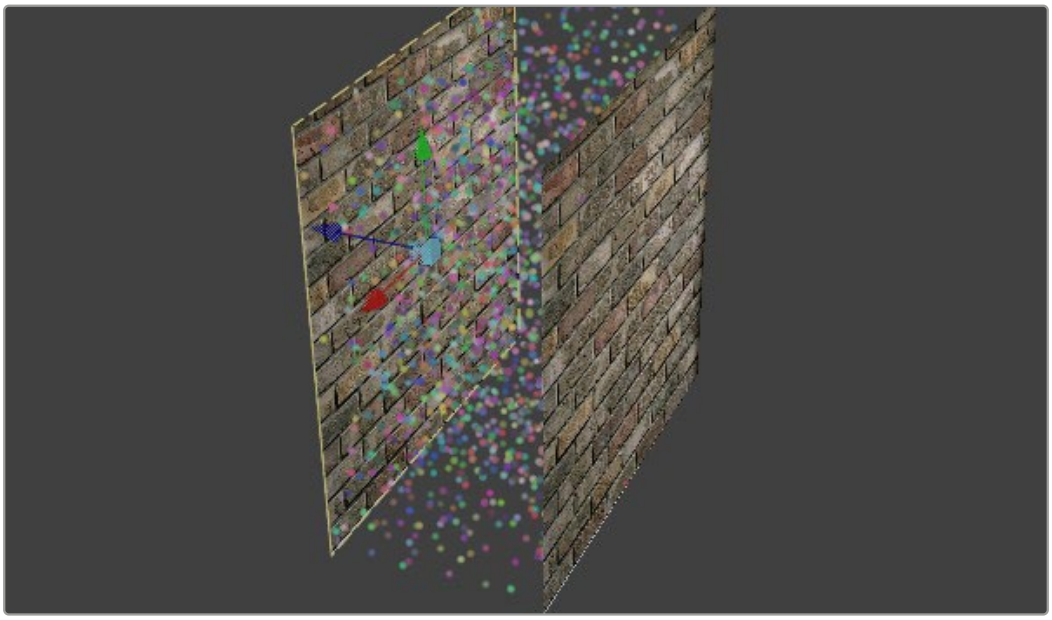
The Winding Rule is then used to determine what is inside/outside. For example, setting the Winding Rule to Odd means that only particles with odd values for the winding number are kept when creating the particles. The exact same approach is used to ensure that polylines that intersect themselves are closed properly.

For example, the following node tree and image show two image planes being used as a mesh region for particle creation.





By setting the region's Winding Ray Direction to the Z (blue) axis, this mesh can then be treated as a closed volume for purposes of particle creation, as pictured below.



Limit By ObjectID

Selecting this checkbox allows the Object ID slider to select the ObjectID used as part of the region.

Style Tab

The Style tab exists in the pEmitter, pSpawn, pChangeStyle, and pImage Emitter. It controls the appearance of the particles, allowing the look of the particles to be designed and animated over time.

Style

The Style menu provides access to the various types of particles supported by the Particle Suite. Each style has its specific controls, as well as controls it will share with other styles.

- **Point Style:** This option produces particles precisely one pixel in size. Controls that are specific to Point style are Apply Mode and Sub Pixel Rendered.
- **Bitmap Style and Brush Style:** Both the Bitmap and Brush styles produce particles based on an image file. The Bitmap style relies on the image from another node in the node tree, and the Brush style uses image files in the Brushes directory. They both have numerous controls for affecting their appearance and animation, described below.
- **Blob Style:** This option produces large, soft spherical particles, with controls for Color, Size, Fade timing, Merge method, and Noise.
- **Line Style:** This style produces straight line-type particles with optional "falloff." The Size to Velocity control described below (under Size Controls) is often useful with this Line type. The Fade control adjusts the amount of falloff over the length of the line.
- **Point Cluster Style:** This style produces small clusters of single-pixel particles. Point Clusters are similar to the Point style; however, they are more efficient when a large quantity of particles is required. This style shares parameters with the Point style. Additional controls specific to Point Cluster style are Number of Points and Number Variance.

Style Options

The following options appear only on some of the styles, as indicated below.

Apply Mode (Point and Point Cluster)

This control applies only to 2D particles; 3D particle systems are not affected.

- **Add:** Overlapping particles are combined by adding together the color values of each particle.
- **Merge:** Overlapping particles are merged.

Sub Pixel Rendered (Point and Point Cluster)

This checkbox determines whether the point particles are rendered with Sub Pixel precision, which provides smoother-looking motion but blurrier particles that take slightly longer to render.

Number of Points and Variance (Point Cluster)

The value of this control determines how many points are in each Point Cluster.

Animate (Bitmap Style)

If the Bitmap source is a movie file or image sequence, this menu determines which frame is grabbed from the source and applied to newly-created particles.

- **Over Time:** All particles use the image produced by the Style Bitmap node at the current time, and change to each successive image together in step, as time increases. A particle created at frame 1 will contain the image at frame 1 of the Style Bitmap. At frame 2, the original particle will use the image from frame 2, and so will any new particles. All created particles will share the same bitmap image from their source at all times.

- **Particle Age:** Each particle animates through the sequence of images provided by the Style Bitmap node, independently of other particles. In other words, an individual particle's appearance is taken from the Style Bitmap node at successive times, indexed by its age.
- **Particle Birth Time:** New particles take the image from the Style Bitmap node at the current time and keep it unchanged until the end of the particle's lifespan. Thus, particles generated on a given frame will all have the same appearance and will stay that way.

Time Offset (Bitmap Style)

This control allows the Bitmap source frame to be offset in time from the current frame.

Time Scale (Bitmap Style)

This control scales the time range of the source bitmap images by a specified amount. For example, a scale of 2 will cause the particle created at frame 1 to be read from the bitmap source at frame 2.

Gain (Bitmap and Brush Style)

This control applies a gain correction to the image used as the bitmap. Higher values produce a brighter image, whereas lower values reduce both the brightness and the transparency of the image.

Style Bitmap (Bitmap Style)

This control appears when the Bitmap style is selected, along with an orange Style Bitmap input on the node's icon in the Node view. Connect a 2D node to this input to provide images to be used for the particles. You can do this on the Node view, or you may drag and drop the image source node onto the Style Bitmap control from the Node Editor or Timeline, or right-click on the control and select the desired source from the Connect To menu.

Brush (Brush Style)

This menu shows the names of any image files stored in the Brushes directory. The location of the Brushes directory is defined in the Preferences dialog, under Path Maps. The default is the Brushes subdirectory within Fusion's install folder. If no images are in this directory, the only option in the menu will be None, and no particles will be rendered.

Noise (Blob Style)

Increasing this control's value will introduce grain-type noise to the blobby particle.

Fade (Line Style)

The Fade control adjusts the falloff over the line particle's length.

The default value of 1.0 causes the line to fade out completely by the end of the length.

Color Controls

The Color Controls select the color and Alpha values of the particles generated by the emitter.

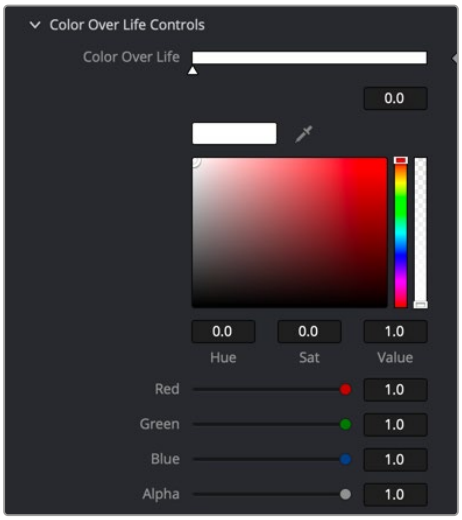
Color Variance

These range controls provide a means of expanding the colors produced by the pEmitter. Setting the Red variance range at -0.2 to +0.2 will produce colors that vary 20% on either side of the red channel, for a total variance of 40%. If the pEmitter is set to produce R0.5, G0.5, B0.5 (pure gray), the variance shown above will produce points with a color range between R0.3, G0.5, B0.5, and R0.7, G0.5, B0.5.

To visualize color space as values between 0-256 or as 0-65535, change the values used by Fusion using the Show Color As option provided in the General tab within the Preferences dialog.

Lock Color Variance

This checkbox locks the color variance of the particles. Unlocking this allows the color variance to be applied differently to each color channel, giving rise to a broader range of colors.



Particles Color Over Life controls

Color Over Life

This standard gradient control allows for the selection of a range of color values to which the particle will adhere over its lifetime.

The left point of the gradient represents the particle color at birth. The right point shows the color of the particle at the end of its lifespan.

Additional points can be added to the gradient control to cause the particle color to shift throughout its life.

This type of control can be useful for fire-type effects (for example, the flame may start blue, turn orange and end a darker red). The gradient itself can be animated over time by right-clicking on the control and selecting Animate from the contextual menu. All points on the gradient will be controlled by a single Color Over Life spline, which controls the speed at which the gradient itself changes. You may also use the From Image modifier, which produces a gradient from the range of colors in an image along a line between two points.

Chapter 55

Position Nodes

This chapter details the Position nodes available in Fusion.

The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Volume Fog [VLF] 1470
Volume Mask [VLM] 1477
Z to World Pos [Z2W] 1482
WPP Concept 1483
The Common Controls 1485

Volume Fog [VLF]



The Volume Fog node

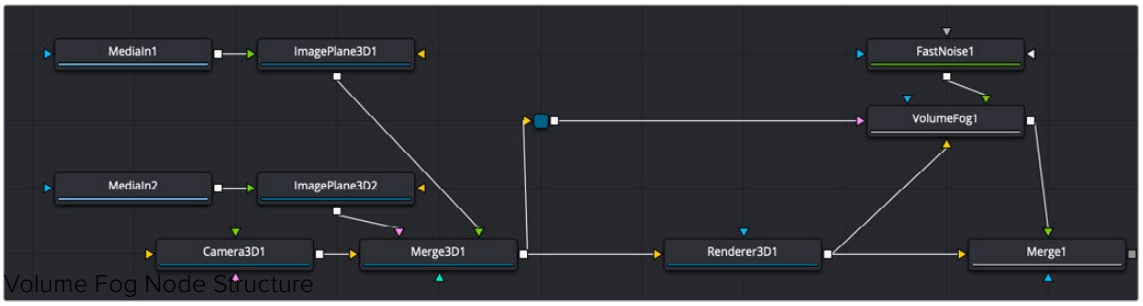
Volume Fog Node Introduction

The Volume Fog node is used to create sophisticated volumetric fog on images containing XYZ Position channels.

As opposed to 3D-rendered volumetric fog, it works on 2D images and delivers much faster results and interactive feedback when setting up the fog. See the “WPP Concept” section at the end of this chapter for further explanation of how this technology works and to learn about the required imagery.

Basic Node Setup

The Volume Fog node takes an image input; in the example below, it is a Renderer 3D with World Position enabled in the output channels. Another input is the 3D scene, which contains the camera. A Fast Noise node generates the fog texture.



Volume Fog Node Structure

A 3D scene and a rendered scene connected to Volume Fog

Inputs

The following inputs appear on the Volume Fog node in the Node Editor.

Image: The orange input accepts the primary image where the fog will be applied. This image contains a World Position Pass in the XYZ Position channels.

Fog Image: The green Fog image input is for creating volumetric fog with varying depth and extent; a 2D image can be connected here. A good starting point is to use a Fast Noise at a small resolution of 256 x 256 pixels.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the fog to certain areas.

Scene Input: The magenta scene input accepts a 3D scene containing a 3D Camera.

Inspector



The Volume Fog Shape tab

Shape Tab

The Shape tab defines the size and location of the fog volume. You can either use the Pick buttons to select the location and orientation in the viewer or use the Translation, Rotation, and Scale controls.

Shape

This menu switches between a basic spherical or rectangular volume to be placed in your image. These volumes can then be further refined using the Fog image and effect mask.

Pick

Drag the Pick button into the viewer to select the XYZ coordinates from any 3D scene or 2D image containing XYZ values, such as a rendered World Pass, to position the center of the Volume object. When picking from a 2D image, make sure it's rendered in 32-bit float to get full precision.

X, Y, Z Offset

These controls can be used to position the center of the fog volume manually or can be animated or connected to other controls in Fusion.

Rotation Pick

Drag the Pick button into the viewer to select the rotational values from any 3D Scene or 2D image containing those values, like an XYZ-Normal-Pass, to reorient the fog volume.

When picking from a 2D image, like an XYZ Normal pass, make sure it's rendered in 32-bit float to get full precision and accurate rotational values.

X, Y, Z Rotation

Use these controls to rotate the fog volume around its center.

X, Y, Z Scale

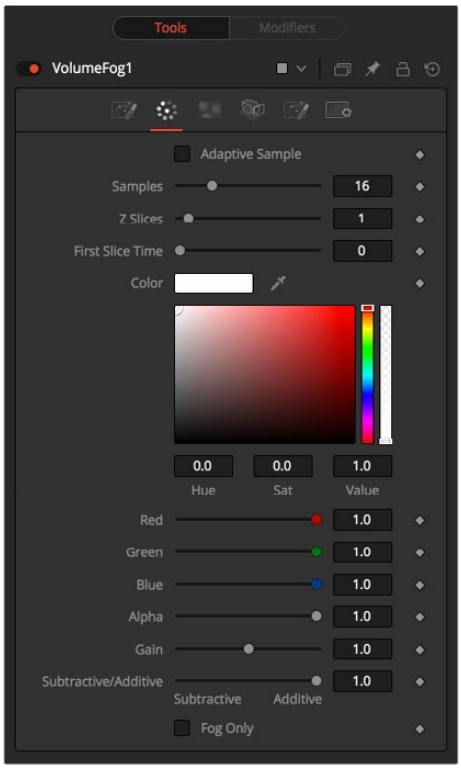
Scale the fog volume in any direction from its center to refine further the overall Size value specified below.

Size

The overall size of the fog volume created.

Soft Edge

Controls how much the fog volume is faded toward the center from its perimeter to achieve a softer look.



The Volume Fog Color tab

Color Tab

The Color tab controls the detail and color of the fog.

Adaptive Samples

Volumes images consist of multiple layers, so there may be 64 layers in a volume. This checkbox adjusts the rendering algorithm for how to best blend those layers.

Dither: Applies a form of noise to improve the blending and hide visible layer differences.

Samples

Determines how many times a “ray” shot into the volume will be evaluated before the final image is created. Not unlike raytracing, higher values lead to more detail inside the volume but also increase render times.

Z Slices

The higher the Z Slices value, the more images from the connected Fog image sequence will be used to form the depth of the volume.

You can, for example, use a Fast Noise with a high Seethe Rate to create such a sequence of images. Be careful with the resolution of the images. Higher resolutions can require a large amount of memory. As a rule of thumb, a resolution of 256 x 256 pixels with 256 Z Slices (i.e., forming a 256 x 256 x 256 cubic volume, which will use up to 256 MB for full color 32-bit float data) should give you a good starting point.

First Slice Time

Determines which frame of the Global Range is used to deliver the first slice from the connected fog image sequence.

Make sure that both Global In and Global Out, as well as the valid range of your source node, fall within the range of First Slice Time + Z Slices.

Color

Allows you to modify the color of the fog generated. This will multiply over any color provided by the connected Fog image.

Gain

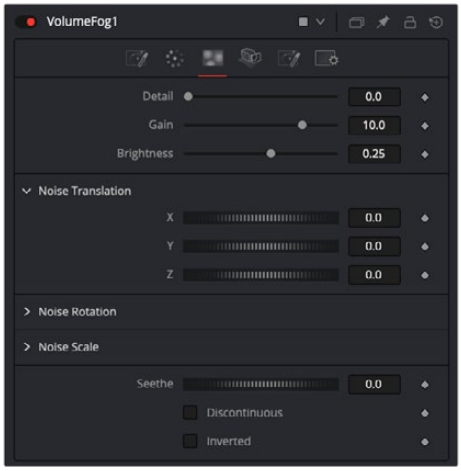
Increases or decreases the intensity of the fog. More Gain will lead to a stronger glow and less transparency in the fog. Lower values let the fog appear less dense.

Subtractive/Additive Slider

Similar to the Merge node, this value controls whether the fog is composed onto the image in Additive or Subtractive mode, leading to a brighter or dimmer appearance of the fog.

Fog Only

This option outputs the generated fog on a black background, which then can be composited manually or used as a mask on a Color Corrector for further refinement.



The Volume Fog Noise tab

Noise Tab

The Noise tab controls the shape and pattern of the noise added to the fog.

Detail

Increase the value of this slider to produce a greater level of detail in the noise result. Larger values add more layers of increasingly detailed noise without affecting the overall pattern. High values take longer to render but can produce a more natural result.

Gain

This control increases or decreases the brightest parts of the noise map.

Brightness

This control adjusts the overall brightness of the noise map, before any gradient color mapping is applied. In Gradient mode, this produces a similar effect to the Offset control.

Translation

Use the Translation coordinate control to pan and move the noise pattern.

Noise Rotation

Use the Rotation controls to orient the noise pattern in 3D.

Seethe

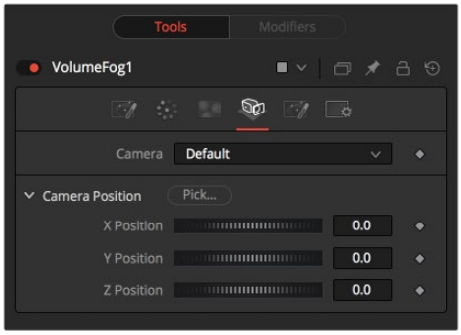
Adjust this thumbwheel control to interpolate the noise map against a different noise map. This will cause a crawling shift in the noise, like it was drifting or flowing. This control must be animated to affect the noise over time.

Discontinuous

Normally, the Noise function interpolates between values to create a smooth, continuous gradient of results. Enable this checkbox to create hard discontinuity lines along some of the noise contours. The result will be a dramatically different effect.

Inverted

Select this checkbox to invert the noise, creating a negative image of the original pattern. This is most effective when Discontinuous is also enabled.



The Volume Fog Camera tab

Camera Tab

For a perfect evaluation of a fog volume, a camera or 3D scene can be connected to the Scene input of the node.

Camera

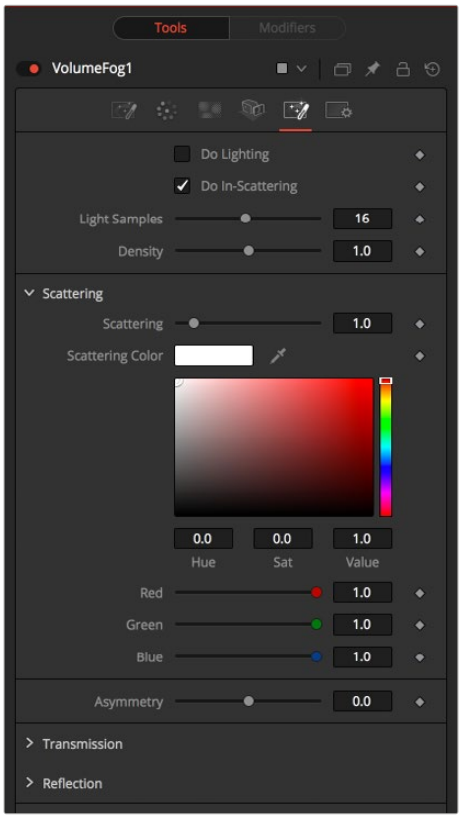
If multiple cameras are available in the connected Scene input, this menu allows the selection of the correct camera needed to evaluate the fog volume. Instead of connecting a camera, position values can be provided manually or by connecting the XYZ values to other controls.

Translation Pick

Drag the Pick button into the viewer to select XYZ coordinates from any 3D scene or 2D image containing XYZ values, like a rendered World Pass, to define the center of the camera. When picking from a 2D image, make sure it's rendered in 32-bit float to get full precision.

X, Y, Z Offset

These controls can be used to define the center of the camera manually or can be animated or connected to other controls in Fusion.



The Volume Fog Light tab

Light Tab

To utilize the controls in the Light tab, you must have actual lights in your 3D scene. Connect that scene, including Camera and Lights, to the 3D input of the node.

Do Lighting

Enables or disables lighting calculations. Keep in mind that when not using OpenCL (i.e., rendering on the CPU), these calculations may become a bit slow.

Do In-Scattering

Enables or disables light-scattering calculations. The volume will still be lit according to the state of the Do Lighting checkbox, but scattering will not be performed.

Light Samples

Determines how accurate the lighting is calculated. Higher values mean more accurate calculation at the expense of longer render times.

Density

This is similar to scattering in that it makes the fog appear thicker. With a high amount of scattering, though, the light will be scattered out of the volume before it has had much chance to travel through the fog, meaning it won't pick up a lot of the transmission color. With a high density instead, the fog still appears thicker, but the light gets a chance to be transmitted, thus picking up the transmission color before it gets scattered out. Scattering is affected by the light direction when Asymmetry is not 0.0. Density is not affected by light direction at all.

Scattering

Determines how much of the light bouncing around in the volume ends up scattering the light out of the fog. If the light scatters more, or more accurately, then there's a higher probability of the light being scattered out of the volume, hence less light is left to continue through the fog. This option can make the fog seem denser.

Asymmetry

Determines in what direction the light is scattered. A value of 0 produces uniform, or isotropic, scattering, meaning all directions have equal probability. A value greater than 0 causes "forward scattering," meaning the light is scattered more into the direction of the light rays. This is similar to what happens with water droplets in clouds. A value smaller than 0 produces "back scattering," where the light is more scattered back toward the original light source.

Transmission

Defines the color that is transmitted through the fog. The light that doesn't get scattered out will tend toward this color. It is a multiplier, though, so if you have a red light, but blue transmission, you won't see any blue.

Reflection

Changes the intensity of the light that is scattered out. Reflection can be used to modify the overall color before Emission is added. This will be combined with the color channels of the volume texture and then used to scale the values. The color options and the color channels of the volume texture are multiplied together, so if the volume texture were red, setting the Reflection color options to blue would not make the result blue. In such a case, they will multiply together to produce black.

Emission

This adds a bit of "glowing" to the fog, adding energy/light back into the calculation. If there are no lights in the scene, and the fog emission is set to be 1.0, the results are similar to no lighting, like turning off the Do Lighting option. Glowing can also be done while producing a different kind of look, by having a Transmission greater than 1. This, however, would never happen in the real world.

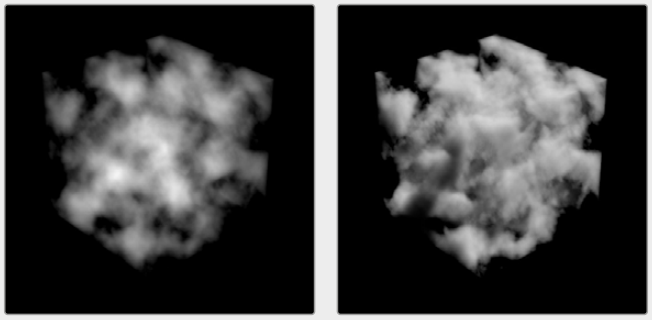
Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Position nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

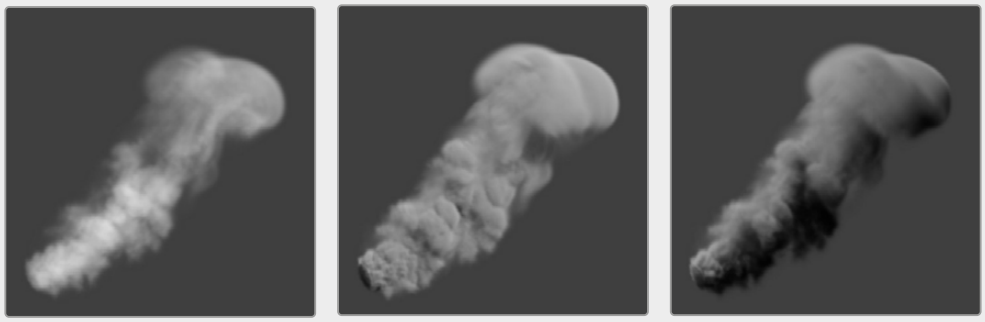
Examples

In these examples, we are looking at a volume from the outside. On the left, you see how the Volume Fog looks with straight accumulation. That means the Do Lighting option is turned off.



On the right, you see the same volume with lighting/scattering turned on, and a single point light.

Here, we have a slightly more complex Volume.



On the left with straight accumulation; in the middle with lighting, scattering, and a single point light; and on the right, the light in the scene has been moved, which also influences the look of the volume.

Volume Mask [VLM]



The Volume Mask node

Volume Mask Node Introduction

The Volume Mask node is used to create volumetric masks from images containing XYZ Position channels.

This can, for example, be used to isolate objects for color correction without the need to track or rotoscope the scene. See the “WPP Concept” section later in this chapter for further explanation on how this technology works and to learn about the required imagery.

Inputs

The following three inputs appear on the Volume Mask node in the Node Editor:

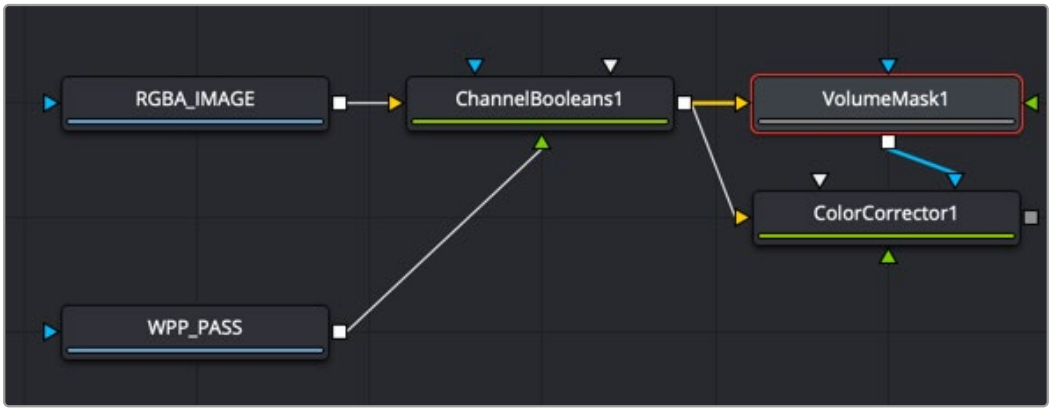
Image: The orange image input accepts a 2D image containing a World Position Pass in the XYZ Position channels.

Mask Image: An image can be connected to the green mask image input for refining the mask.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the volume mask to certain areas.

Basic Node Setup

Below, a MediaIn labeled RGBA contains the main rendered image from a 3D scene. A World Position Pass from a 3D-rendered scene is labeled WPP_PASS. The Channel Booleans node is used to map the X position pass to the red channel, Y to the green channel, and Z to the blue channel. The Volume Mask tool extracts the 3D position information from the World Position Pass to place a mask in 3D space. The mask is then used as an effect mask on a color corrector to selectively color correct objects in a 3D scene.



A Volume Mask tool takes advantage of World Position Pass for color correction in a 3D scene

Inspector



The Volume Mask Shape tab

Shape Tab

The Shape tab defines the size and location of the Volume Mask. You can either use the Pick buttons to select the location and orientation in the viewer or use the Translation, Rotation, and Scale controls.

Shape

This menu switches between a spherical or rectangular mask to be placed in your image. The mask can be further refined using the mask image input.

Translation Pick

Drag the Pick button into the viewer to select XYZ coordinates from any 3D scene or 2D image containing XYZ values, like a rendered World Pass, to position the center of the Volume Mask. When picking from a 2D image, make sure it's rendered in 32-bit float to get full precision.

X, Y, Z Offset

These controls can be used to position the center of the mask manually or can be animated or connected to other controls in Fusion.

Rotation Pick

Drag the Pick button into the viewer to select rotational values from any 3D scene or 2D image containing those values, like an XYZ Normal pass, to reorient the mask.

When picking from a 2D image, like an XYZ Normal pass, make sure it's rendered in 32-bit float, and use World Space coordinates to get full precision and the correct rotational values.

X, Y, Z Rotation

Use these controls to rotate the mask around its center.

X, Y, Z Scale

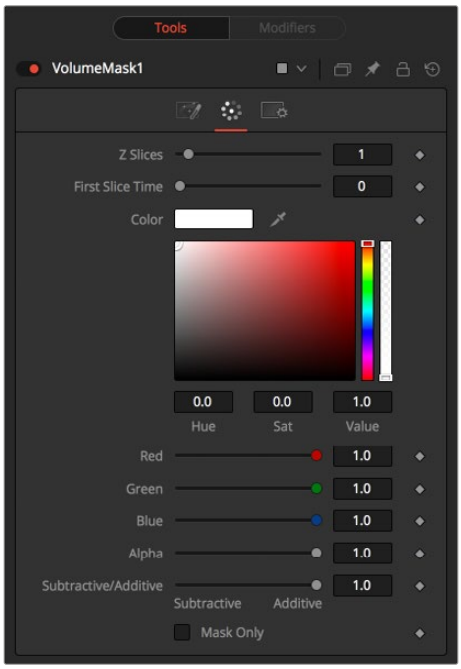
Scale the mask in any direction from its center to further refine the overall Size value specified below.

Size

The overall size, in X, Y, and Z, of the mask created.

Soft Edge

Controls how much the Volume is faded toward the center from its perimeter to achieve a softer look.



The Volume Mask Color tab

Color Tab

The Color tab controls the color and blending of the mask image.

Color

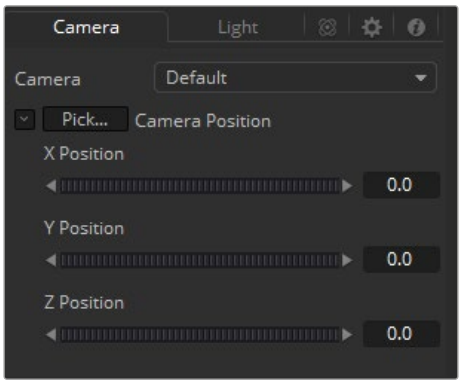
Allows you to modify the color of the generated Volume Mask. This will add to any color provided by the connected mask image.

Subtractive/Additive Slider

Similar to the Merge node, this value controls whether the mask is composed onto the image in Additive or Subtractive mode, leading to a brighter or dimmer appearance of the mask.

Mask Only

Outputs the generated mask on a black background, which then can be used as a mask on a Color Corrector for further refinement.



The Volume Mask Camera tab

Camera Tab

For a perfect evaluation of a Volume, a camera or 3D scene can be connected to the Scene input of the node.

Camera

If multiple cameras are available in the connected Scene input, this drop-down menu allows you to choose the correct camera needed to evaluate the Volume.

Instead of connecting a camera, position values can also be provided manually or by connecting the XYZ values to other controls.

Translation Pick

Drag the Pick button into the viewer to select XYZ coordinates from any 3D scene or 2D image containing XYZ values, like a rendered World Pass, to define the center of the camera.

When picking from a 2D image, make sure it's rendered in 32-bit float to get full precision.

X, Y, Z Offset

These controls can be used to define the center of the camera manually or can be animated or connected to other controls in Fusion.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Position nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Z to World Pos [Z2W]



The Z to World Position node

Z to World Pos Node Introduction

The Z to World Pos node is used to either generate a World Position Pass from a Z channel and a 3D Camera or a Z channel from a World Position Pass and a 3D Camera.

Creating a World Position Pass from Z-depth can be useful when your 3D application is not capable of creating a WPP.

It can also be used when a 3D-tracking software outputs a per-pixel Z-depth together with the 3D Camera. Thus, the Volume Mask and Volume Fog could be applied to real-world scenes. The quality of the resulting WPP depends mainly on the quality of the incoming Z channel.

See the “WPP Concept” section for further explanation on how this technology works and to learn about the required imagery.

Inputs

The following inputs appear on the node tile in the Node Editor:

Image: The orange image input accepts an image containing a World Position Pass or a Z-depth pass, depending on the desired operation.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the World Position Pass to certain areas.

Scene Input: The magenta scene input accepts a 3D scene input containing a 3D Camera.

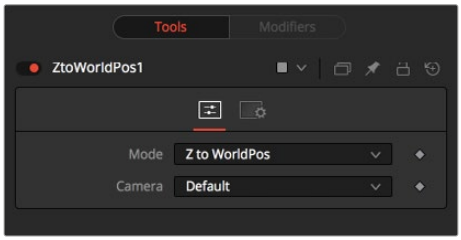
Basic Node Setup

Below, a MediaIn labeled RGBA contains the main rendered image from a 3D scene. A Z-depth pass from a 3D-rendered scene is labeled Z_PASS. The Channel Booleans node is used to map the Aux Z channel into either the red, green, or blue channel. The Z to World Position node is placed after the Channel Booleans node, and an imported 3D camera that matches the RGBA image is connected to the 3D camera input on the Z to World Position node. A Channel Booleans node is placed after the Z to World Position node, which can remap the X, Y, and Z positions for use in other nodes.



A Z to World Position node creates a World Position Pass from a Z-depth pass

Inspector



The Z to World Position Controls tab

Controls Tab

The Controls tab determines whether you are creating a World Position Pass or a Z channel. If there is more than one camera in the connected scene, this tab also selects the camera to use for the calculation.

Mode

This menu switches between creating a Z channel from a World Position Pass or vice versa.

Camera

If multiple cameras are available in the connected Scene input, this drop-down menu allows you to choose the correct camera needed to evaluate the image.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Position nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

WPP Concept

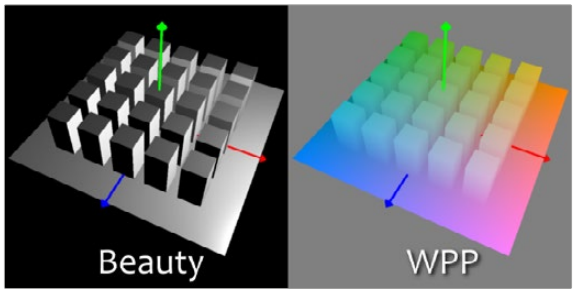
The Position nodes in Fusion offer an entirely new way of working with masks and Volumetrics for footage containing XYZ Position channels. Z to World offers the option to create those channels out of a Z channel and 3D Camera information. For this overview, we refer to the World Position Pass as WPP.

What Is a WPP?

The WPP interprets each pixel’s XYZ position in 3D space as an RGB color value.

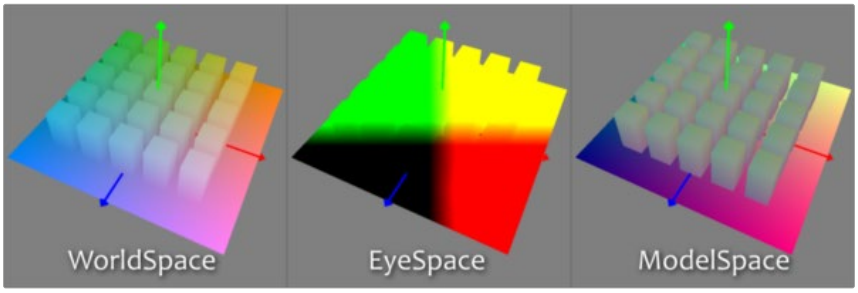
For instance, if a pixel sits at 0/0/0, the resulting pixel has an RGB value of 0/0/0 and thus will be black. If the pixel sits at 1/0/0 in the 3D scene, the resulting pixel is entirely red. Of course, if the coordinates of the pixel are something like -60/75/123, WPP interprets those values as RGB color values as well.

Due to the potentially enormous size of a 3D scene, the WPP channel should always be rendered in 32-bit floating point to provide the accuracy needed. The image below shows a 3D rendering of a scene with its center sitting at 0/0/0 in 3D Space and the related WPP. For better visibility, the WPP is normalized in this example.



Different Coordinate Spaces

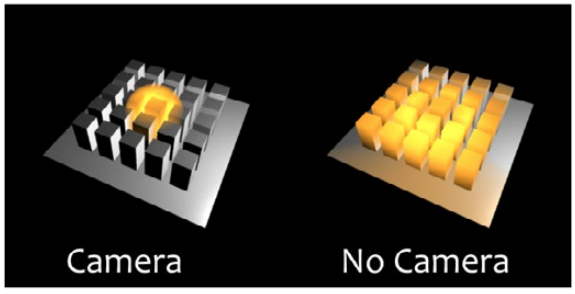
Rendering WPPs can occur in different Coordinate Spaces. These include World Space, Eye Space, and Object Space. The image below depicts how those different spaces look, although the nodes in Fusion require the WPP rendering to occur in World Space.



The Scene Input

The nodes offer a Scene input, which can either be a 3D camera or a 3D scene containing a camera. While the camera is vital for the Z to World node, Volume Mask and Volume Fog can generate their output without any camera attached or with the camera position set to 0/0/0.

However, connecting a camera that lines up with the original camera the WPP has been rendered from, or setting the camera's position manually, dramatically improves the accuracy and look of the resulting fog or mask.

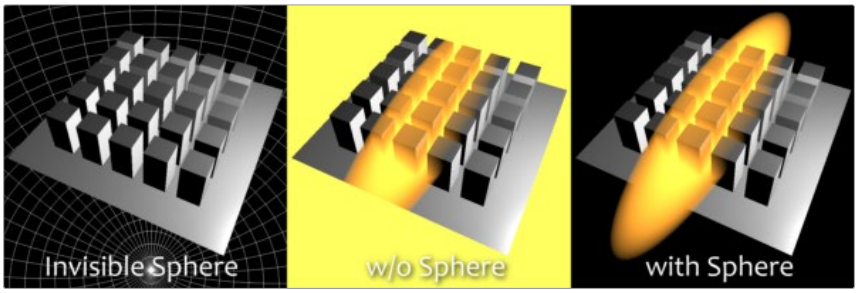


The “Invisible Sphere”

The example scene shown so far has an empty background, meaning there is nothing in the scene apart from the ground plane and the cubes.

If applying fog to a scene like that, which is larger than the ground plane, the result will look similar to the “w/o Sphere” example shown below because, with no WPP information outside the ground plane, the resulting value is 0/0/0, and the fog fills that area as well.

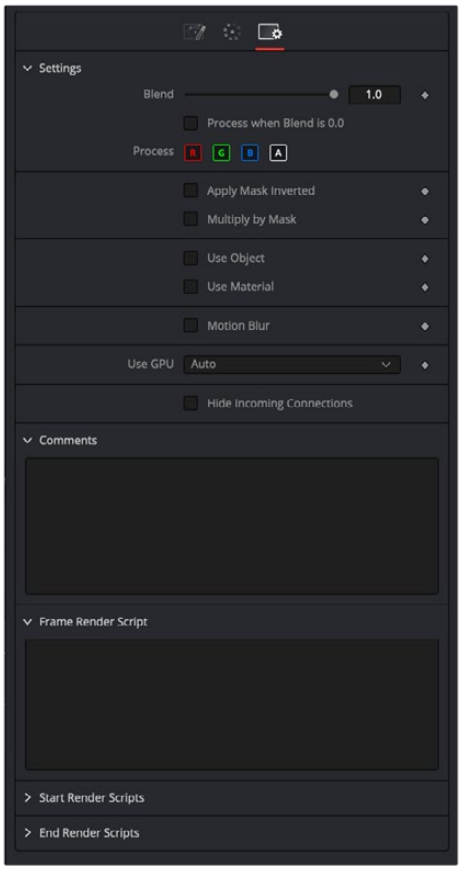
To get around that, you can add an invisible bounding sphere to your scene to create “dummy” WPP values to help the Fog node to create the correct volume as shown in the “with Sphere” example below.



The Common Controls

Nodes that handle Position operations share several identical controls in the Inspector. This section describes controls that are common among Position nodes.

Inspector



Position nodes Common Settings tab

Settings Tab

The Settings tab in the Inspector can be found on every tool in the Position category. The controls are consistent and work the same way for each tool.

Blend

The Blend control is used to blend between the tool's original image input and the tool's final modified output image. When the blend value is 0.0, the outgoing image is identical to the incoming image. Normally, this will cause the tool to skip processing entirely, copying the input straight to the output.

Process When Blend Is 0.0

The tool is processed even when the input value is zero. This can be useful if the node is scripted to trigger a task, but the node's value is set to 0.0.

Red/Green/Blue/Alpha Channel Selector

These four buttons are used to limit the effect of the tool to specified color channels. This filter is often applied after the tool has been processed.

For example, if the Red button on a Blur tool is deselected, the blur will first be applied to the image, and then the red channel from the original input will be copied back over the red channel of the result.

There are some exceptions, such as tools for which deselecting these channels causes the tool to skip processing that channel entirely. Tools that do this will generally possess a set of identical RGBA buttons on the Controls tab in the tool. In this case, the buttons in the Settings and the Controls tabs are identical.

Apply Mask Inverted

Enabling the Apply Mask Inverted option inverts the complete mask channel for the tool. The mask channel is the combined result of all masks connected to or generated in a node.

Multiply by Mask

Selecting this option will cause the RGB values of the masked image to be multiplied by the mask channel's values. This will cause all pixels of the image not included in the mask (i.e., set to 0) to become black/transparent.

Use Object/Use Material (Checkboxes)

Some 3D software can render to file formats that support additional channels. Notably, the EXR file format supports Object and Material ID channels, which can be used as a mask for the effect. These checkboxes determine whether the channels will be used if present. The specific Material ID or Object ID affected is chosen using the next set of controls.

Correct Edges

This checkbox appears only when the Use Object or Use Material checkboxes are selected. It toggles the method used to deal with overlapping edges of objects in a multi-object image. When enabled, the Coverage and Background Color channels are used to separate and improve the effect around the edge of the object. If this option is disabled (or no Coverage or Background Color channels are available), aliasing may occur on the edge of the mask.

> For more information on the Coverage and Background Color channels, *see Chapter 18, "Understanding Image Channels,"* in the Fusion Reference Manual

Object ID/Material ID (Sliders)

Use these sliders to select which ID will be used to create a mask from the object or material channels of an image. Use the Sample button in the same way as the Color Picker: to grab IDs from the image displayed in the viewer. The image or sequence must have been rendered from a 3D software package with those channels included.

Motion Blur

- **Motion Blur:** This toggles the rendering of Motion Blur on the tool. When this control is toggled on, the tool's predicted motion is used to produce the motion blur caused by the virtual camera's shutter. When the control is toggled off, no motion blur is created.
- **Quality:** Quality determines the number of samples used to create the blur. A quality setting of 2 will cause Fusion to create two samples to either side of an object's actual motion. Larger values produce smoother results but increase the render time.
- **Shutter Angle:** Shutter Angle controls the angle of the virtual shutter used to produce the motion blur effect. Larger angles create more blur but increase the render times. A value of 360 is the equivalent of having the shutter open for one full frame exposure. Higher values are possible and can be used to create interesting effects.
- **Center Bias:** Center Bias modifies the position of the center of the motion blur. This allows for the creation of motion trail effects.
- **Sample Spread:** Adjusting this control modifies the weighting given to each sample. This affects the brightness of the samples.

Use GPU

The Use GPU menu has three settings. Setting the menu to Disable turns off GPU hardware-accelerated rendering. Enabled uses the GPU hardware for rendering the node. Auto uses a capable GPU if one is available and falls back to software rendering when a capable GPU is not available.

Hide Incoming Connections

Enabling this checkbox can hide connection lines from incoming nodes, making a node tree appear cleaner and easier to read. When enabled, empty fields for each input on a node will be displayed in the Inspector. Dragging a connected node from the node tree into the field will hide that incoming connection line as long as the node is not selected in the node tree. When the node is selected in the node tree, the line will reappear.

Comments

The Comments field is used to add notes to a tool. Click in the empty field and type the text. When a note is added to a tool, a small red square appears in the lower-left corner of the node when the full tile is displayed, or a small text bubble icon appears on the right when nodes are collapsed. To see the note in the Node Editor, hold the mouse pointer over the node to display the tooltip.

Scripts

Three Scripting fields are available on every tool in Fusion from the Settings tab. They each contain edit boxes used to add scripts that process when the tool is rendering. For more details on scripting nodes, please consult the Fusion scripting documentation.

Chapter 56

Resolve Connect

This chapter details the single node found in the Resolve Connect category, available only in standalone Fusion Studio.

> The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.
>
> For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

External Matte Saver [EMS] 1489
External Matte Saver Node Introduction 1489
Inputs 1489
Basic Node Setup 1489
Inspector 1490

External Matte Saver [EMS]



The External Matte Saver node

NOTE: The Resolve Connect category and External Matte Saver node are available only in Fusion Studio.

External Matte Saver Node Introduction

The External Matte Saver node renders multiple mattes into multiple channels of an EXR file. This file is intended to import into DaVinci Resolve's Color page as an efficient way to deliver multiple mattes for color grading. To perform the same operation with a traditional Saver node, a Channel Boolean is needed to place each matte into a channel, and then name that channel in the Saver. It requires a bit more of a set up. This node streamlines the process by providing multiple inputs and naming the channel based on the node.

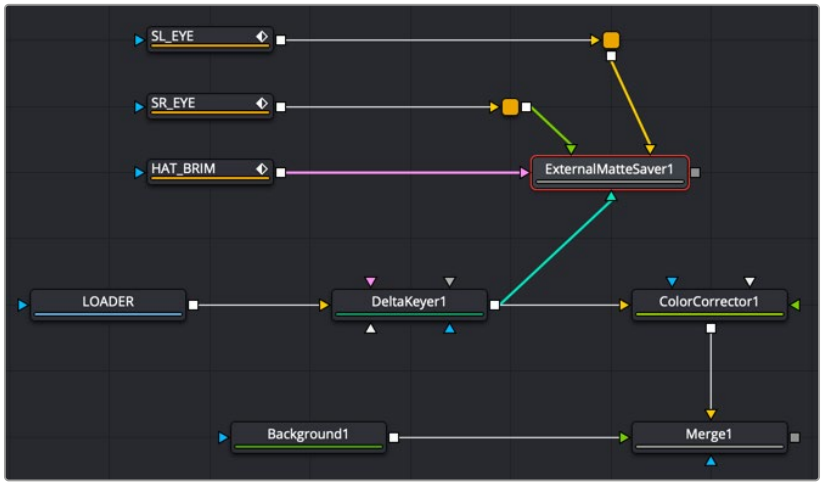
Inputs

By default, the node provides a single input for a 2D image you want to save as a matte.

Input: Although initially there is only a single orange input for a matte to connect, the Inspector provides an Add button for adding additional inputs. Each input uses a new color, but all accept 2D RGBA images.

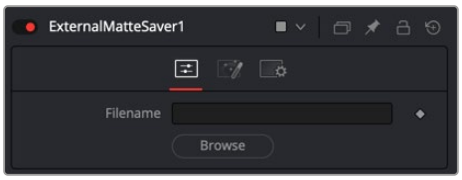
Basic Node Setup

The External Matte Saver node can be placed at the end of any node tree branch for saving mattes. Below, it is connected to the Delta Keyer as well as three other mattes. Each matte beyond the initial one is connected by clicking the Add button in the Inspector first, and then a new input is provided on the node.



An External Matte Saver node added as a separate branch in a node tree to render the mattes

Inspector



The External Matte Saver Controls tab

Controls Tab

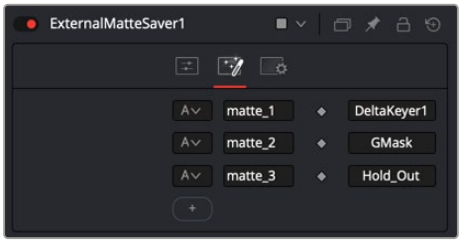
The Controls tab is used to name the saved file and determine where on your hard drive the file is stored.

Filename

Enter the name you want to use for the EXR file in the Filename field. At the end of the name, append the .exr extension to ensure that the file is saved as an EXR file.

Browse

Clicking the Browse button opens a standard file browser window where you can select the location to save the file.



The External Matte Saver Mattes tab

Mattes Tab

The Mattes tab is where you set up the number of mattes saved in the file, the name for each channel, and the RGBA channels saved from each input.

Channels menu

The Channels menu allows you to select which channels are saved in the matte. You can choose the alpha channel, the RGB channels, or the RGBA channels.

Channels Name

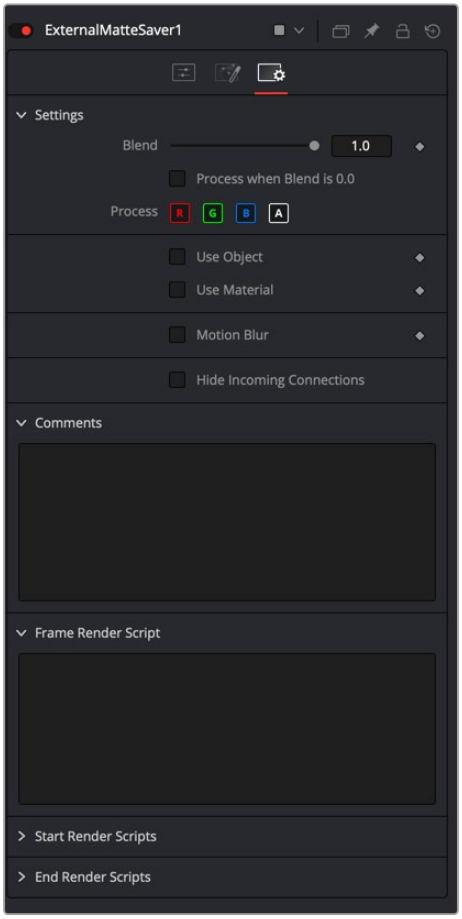
The Channels Name field allows you to customize the name of the matte channel you are saving. This name is displayed in DaVinci Resolve's Color page.

Node Name

The Node Name field displays the source of the matte. This is automatically populated when you connect a node to the input.

Add

Clicking the Add button adds an input on the node and another set of fields for you to configure and name the new matte channel.



The External Matte Saver Settings tab

Settings Tab

The Settings Tab in the Inspector is similar to settings found in the Saver tool. The controls are consistent and work the same way as the Settings in other tools.

Blend

The Blend control is used to blend between the tool's original image input and the tool's final modified output image. When the blend value is 0.0, the outgoing image is identical to the incoming image. Normally, this causes the tool to skip processing entirely, copying the input straight to the output.

Process When Blend Is 0.0

The tool is processed even when the input value is zero. This can be useful if processing of this node is scripted to trigger another task, but the value of the node is set to 0.0.

Red/Green/Blue/Alpha Channel Selector

These four buttons are used to limit the effect of the tool to specified color channels. This filter is often applied after the tool has been processed.

For example, if the Red button on a Blur tool is deselected, the blur is first applied to the image, and then the red channel from the original input is copied back over the red channel of the result.

There are some exceptions, such as tools for which deselecting these channels causes the tool to skip processing that channel entirely. Tools that do this generally possess a set of identical RGBA buttons on the Controls tab in the tool. In this case, the buttons in the Settings and the Controls tabs are identical.

Apply Mask Inverted

Enabling the Apply Mask Inverted option inverts the complete mask channel for the tool. The mask channel is the combined result of all masks connected to or generated in a node.

Multiply by Mask

Selecting this option causes the RGB values of the masked image to be multiplied by the mask channel's values. This causes all pixels of the image not included in the mask (i.e., set to 0) to become black/transparent.

Use Object/Use Material (Checkboxes)

Some 3D software can render to file formats that support additional channels. Notably, the EXR file format supports Object ID and Material ID channels, which can be used as a mask for the effect. These checkboxes determine whether the channels are used, if present. The specific Material ID or Object ID affected is chosen using the next set of controls.

Correct Edges

This checkbox appears only when the Use Object or Use Material checkboxes are selected. It toggles the method used to deal with overlapping edges of objects in a multi-object image. When enabled, the Coverage and Background Color channels are used to separate and improve the effect around the edge of the object. If this option is disabled (or no Coverage or Background Color channels are available), aliasing may occur on the edge of the mask.

> For more information on the Coverage and Background Color channels, *see Chapter 18, "Understanding Image Channels,"* in the Fusion Reference Manual.

Object ID/Material ID (Sliders)

Use these sliders to select which ID is used to create a mask from the object or material channels of an image. Use the Sample button in the same way as the Color Picker: to grab IDs from the image displayed in the viewer. The image or sequence must have been rendered from a 3D software package with those channels included.

Motion Blur

- **Motion Blur:** This toggles the rendering of Motion Blur on the tool. When this control is toggled on, the tool's predicted motion is used to produce the motion blur caused by the virtual camera's shutter. When the control is toggled off, no motion blur is created.
- **Quality:** Quality determines the number of samples used to create the blur. A quality setting of 2 causes Fusion to create two samples to either side of an object's actual motion. Larger values produce smoother results but increase the render time.
- **Shutter Angle:** Shutter Angle controls the angle of the virtual shutter used to produce the motion blur effect. Larger angles create more blur but increase the render times. A value of 360 is the equivalent of having the shutter open for one full frame exposure. Higher values are possible and can be used to create interesting effects.
- **Center Bias:** Center Bias modifies the position of the center of the motion blur. This allows for the creation of motion trail effects.
- **Sample Spread:** Adjusting this control modifies the weighting given to each sample. This affects the brightness of the samples.

Hide Incoming Connections

Enabling this checkbox can hide connection lines from incoming nodes, making a node tree appear cleaner and easier to read. When enabled, empty fields for each input on a node will be displayed in the Inspector. Dragging a connected node from the node tree into the field will hide that incoming connection line as long as the node is not selected in the node tree. When the node is selected in the node tree, the line will reappear.

Comments

The Comments field is used to add notes to a tool. Click in the empty field and type the text. When a note is added to a tool, a small red square appears in the lower-left corner of the node when the full tile is displayed, or a small text bubble icon appears on the right when nodes are collapsed. To see the note in the Node Editor, hold the mouse pointer over the node to display the tooltip.

Scripts

Three Scripting fields are available on every tool in Fusion from the Settings tab. They each contain edit boxes used to add scripts that process when the tool is rendering. For more details on scripting nodes, please consult the Fusion scripting documentation.

Chapter 57

Shape Nodes

This chapter details the Shape nodes available in Fusion.

Contents

sBoolean 1495
sBSpline [sBSp] 1498
sChangeStyle [sCs] 1499
sDuplicate 1501
sEllipse 1503
sExpand 1506
sGrid 1507
sJitter 1509
sMerge 1510
sNGon 1512
sOutline 1515
sPolygon [sPly] 1516
Adding Points 1520
sRectangle 1521
sRender 1524
sStar 1527
sText [sTxt] 1530
sTransform 1531
Common Controls 1532

sBoolean



The sBoolean node

The sBoolean node combines or excludes overlapping areas of two shapes based on a menu of boolean operations.

Like almost all shape nodes, you can only view the sBoolean node's results through a sRender node.

External Inputs

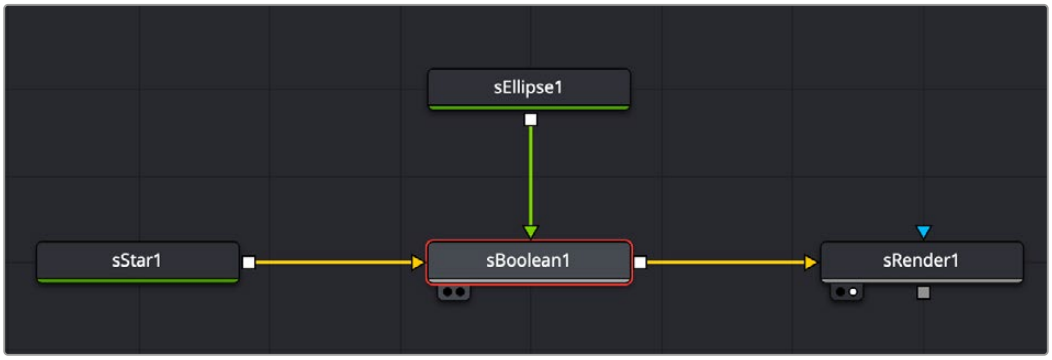
The following inputs appear on the node's tile in the Node Editor. Except when using the subtract boolean operation, which shape you connect into which input does not matter.

Input1: [orange, required] This input accepts the output of another shape node. This input is used as the base shape when the subtract boolean operation is chosen.

Input2: [green, optional] This input accepts the output of another shape node. This input is used to cut the base shape hole when the subtract boolean operation is chosen.

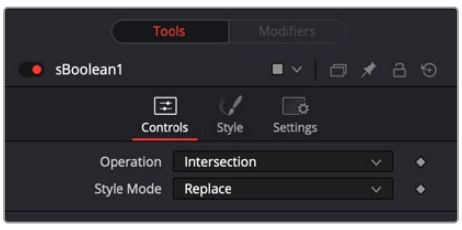
Basic Node Setup

The sBoolean node is used to combine two shape nodes. The output of the sBoolean can then be output to another shape node or a sRender node for viewing.



Star and ellipse shapes combined in an sBoolean node, then output to an sRender for viewing and further processing

Inspector



The sBoolean Controls tab

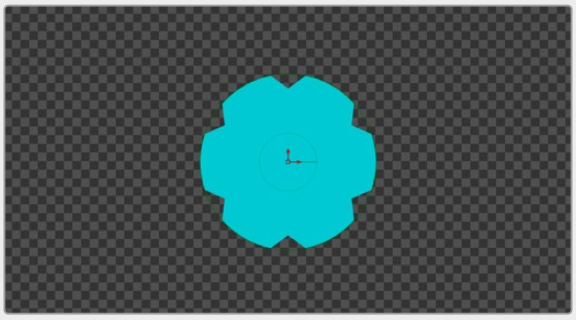
Controls

The Controls tab is primarily used to select the boolean operation that determines how the two shapes are combined.

Operation

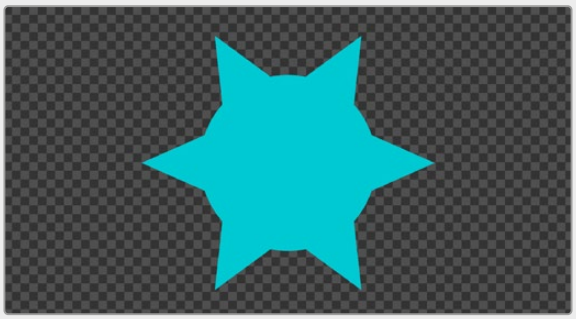
The operation menu includes four boolean operations:

Intersection: Sometimes called an AND operation, this setting will only show areas where the two shapes overlap. The result is only where input 1 AND input 2 overlap.



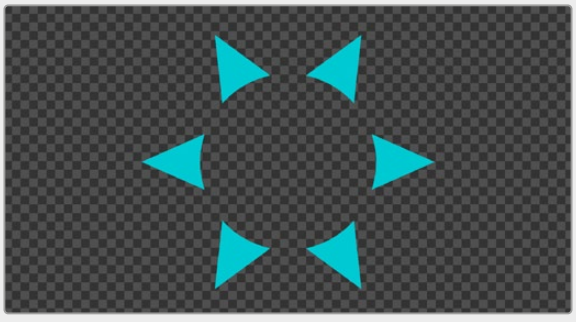
Star and ellipse shapes with an sBoolean node set to intersection

Union: Sometimes called an OR operation, this setting will only show areas where either of the two shapes exists. The result is where either input 1 OR input 2 exists. The Union setting is similar to the result of the sMerge node.



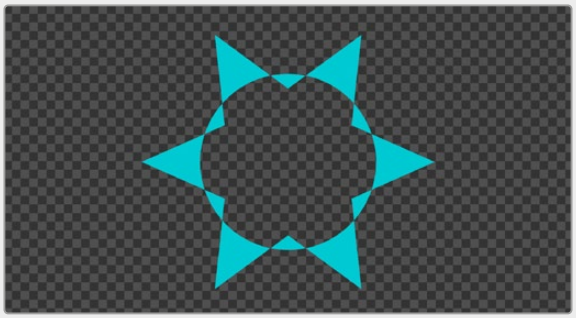
Star and ellipse shapes with an sBoolean node set to union

Subtract: Sometimes called a NOT operation, this setting outputs the shape of input 1 but eliminates the areas where input 2 overlaps. The result is input 1 minus input 2.



Star and ellipse shapes with an sBoolean node set to subtract

Xor: Sometimes called an AND NOT operation, this setting outputs the shape of input 1 or input 2 but eliminates the areas where they overlap. The result is (input 1 minus input 2) + (input 2 minus input 1).

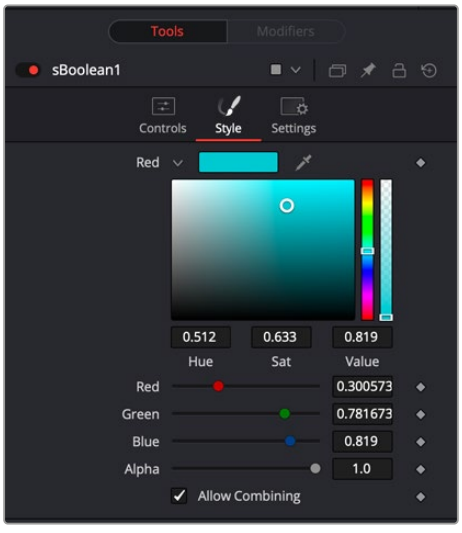


Star and ellipse shapes with an sBoolean node set to xor

Style Mode

The Style mode menu only includes one option. The Replace setting replaces the color and alpha level of the incoming shapes with the color set in the Style tab.

Style Tab



The sBoolean Style tab

Style

Any color assigned to the individual shape nodes is replaced by the color set using the Style tab controls.

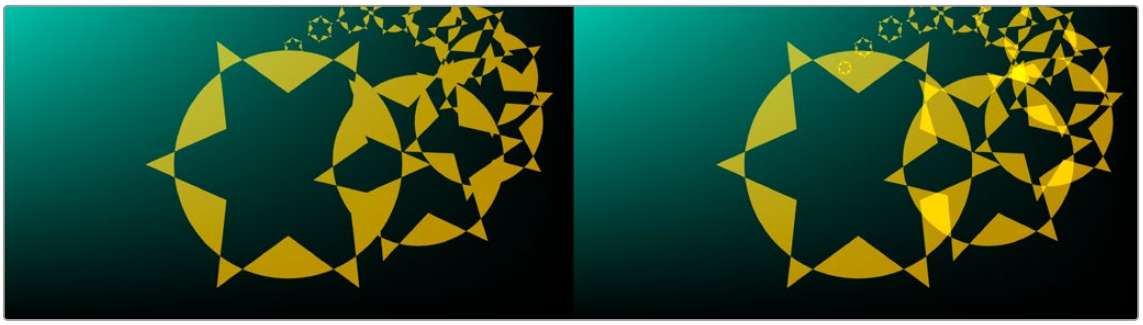
Color

The color controls determine the color of the output shape from the sBoolean node. To choose a shape color, you can click the color disclosure arrow, use the color swatch, or drag the eyedropper into the viewer to select a color from an image. The RGBA sliders or number fields can be used to enter each color channel's value or the strength of the Alpha channel.

Allow Combining

When this checkbox is enabled, the Alpha channel value is maintained even when passing through other nodes downstream that may cause the shape to overlap with copies of itself. When disabled, the Alpha channel value may increase when the shape overlaps itself.

For instance, if an ellipse's Alpha channel is set to .5, enabling the Allow Combining checkbox maintains that value even if the shape passes through a duplicate or grid node that causes the shape to overlap. Disabling the checkbox causes the Alpha channel values to be compounded at each overlapping area. When using the sBoolean node, the individual shape node checkboxes are ignored, and the sBoolean node's checkbox determines the Alpha channel's behavior.



Allow Combining Enabled (left), Allow Combining Disabled (right)

Common Controls

Settings tab

The Settings tab in the Inspector is common to all Shape nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

sBSpline [sBSp]



The sBSpline node

Introduction

sBSpline is identical to a sPolygon in all respects except one. Where sPolygon use Bézier splines, this node uses B-Splines. Where Bézier splines employ a central point and two handles to manage the smoothing of the spline segment, a B-Spline requires only a single point. This means that a sBSpline shape requires far fewer control points to create a nicely smoothed shape.

Usage



The sBSpline controls

When first added to a Node Editor, the sBSpline shape consists of only Center control, which is visible onscreen.

Points are added to the sBSpline by clicking in the viewer. Each new point is connected to the last one created, but instead of the spline going directly through each control point, sBSpline control points only influence the spline shape. The control point pulls the spline in its direction to create a smooth curve.

Like the sPolygon shape tool, the sBSpline shape auto-animates. Adding this node to the Node Editor adds a keyframe to the current frame. Moving to a new frame and changing the shape creates a new keyframe and interpolate between the two defined shapes.

sChangeStyle [sCs]



The sChangeStyle node

sChangeStyle Node Introduction

The sChangeStyle node assigns a new style to the combined shape tree, overriding the style (color and Allow Combining settings) set in the upstream shape tools.

Inputs

The single input on the sChangeStyle tool is used to connect the Shape node you want to change the style of.

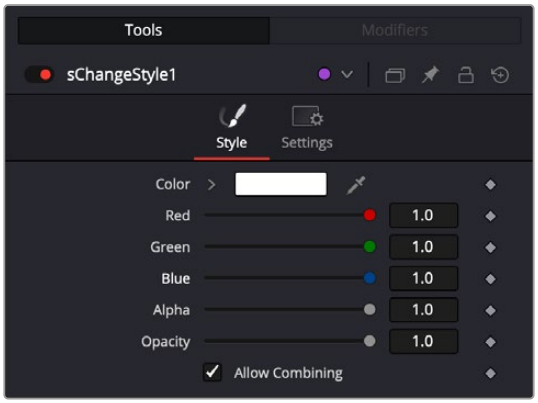
Basic Node Setup

Connect the shape output that you wish change the style of to the input. The output of the node then gets passed to the remainder of the shape node tree.



A simple node tree showing the sChangeShape node taking the output of an sEllipse shape, performing the style change, and passing it to the sRender node.

Inspector



The sChangeStyle controls

Controls

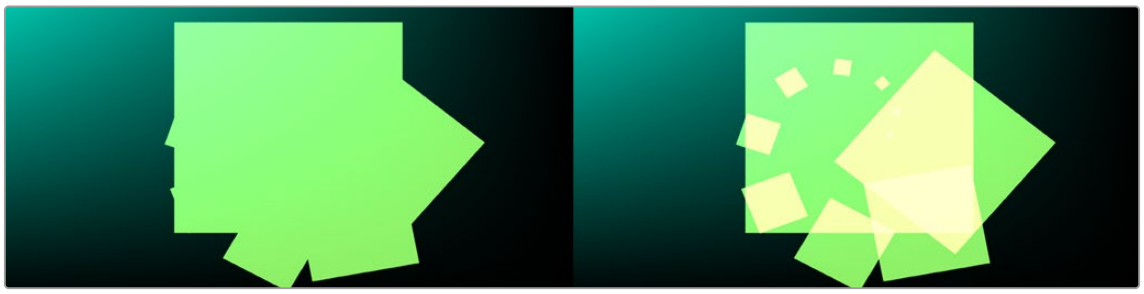
The Style tab is used to assign color to the shape and control its transparency.

Color

The Color parameter controls determine the color of the fill and border from the sPolygon node. To choose a shape color, you can click the color disclosure arrow and use the color swatch, or drag the eye dropper into the viewer to select a color from an image. The RGBA sliders or number fields can be used to enter the value of each color channel or the strength of the Alpha channel.

Allow Combining

When this checkbox is enabled, the Alpha channel value is maintained, even when passing through other nodes downstream that may cause the shape to overlap with copies of itself. When disabled, the Alpha channel value may increase when the shape overlaps itself. For instance, if a rectangle Alpha channel is set to .5, enabling the Allow Combining checkbox maintains that value, even if the shape passes through a duplicate or grid node that causes the shape and Alpha channel to overlap. Disabling the checkbox causes the Alpha channel values to be compounded at each overlapping area.



Allow Combining enabled (left), Allow Combining disabled (right)

sDuplicate



The sDuplicate node

The sDuplicate node creates copies of the input shape, offsetting each copy's position, size, and rotation. Like almost all shape nodes, you can only view the sDuplicate node's results through a sRender node.

External Inputs

The following input appears on the node's tile in the Node Editor:

Input1: [orange, required] This input accepts the output of another Shape node. The shape connected to this input is copied and offset based on the controls in the Inspector.

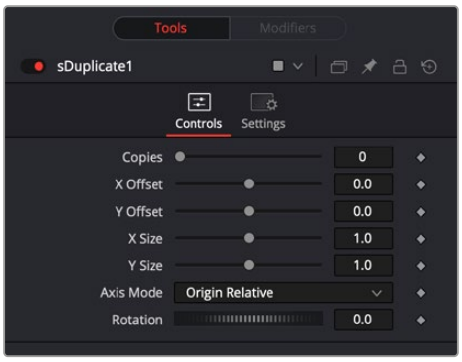
Basic Node Setup

The sDuplicate node takes a single input that is most often from a single or compound shape from a sMerge or sBoolean node. The sDuplicate node creates copies of the incoming shape and offsets them to create a pattern. The output of the sDuplicate can then be output to another shape node or to a sRender node for viewing or compositing into the greater node tree.



An sEllipse shape connected to an sDuplicate and then output to an sRender for viewing and combining with other elements

Inspector



The sDuplicate Controls tab

Controls

The Controls tab is used to determine the number of copies and set their position, size, and rotation offset.

Copies

This slider determines the number of copies created by the node. The number does not include the original shape, so entering a value of five will produce five copies plus the original.

X and Y Offset

These sliders set the X and Y distance between each of the copies. Each copy is offset from the previous copy by the value entered in the X and Y number fields. The copies all start at 0, the center of the original shape, and are offset from there. Using Fusion's normalized coordinate system, entering X Offset at 0.5 would move each copy half the frame's width to the right. Entering -1.0 would move each copy to the left by the width of the frame.

X and Y Size

Sets the X and Y size offset based on the previous shape size. For instance, an X and Y value of 1.0 creates copies identical in size to the original but. Entering a value of X and Y of 0.5 will cause each copy to be half the size of the copy before it.

Axis Mode

The Axis mode menu provides four options for determining how each copy determines its rotational pivot point.

- **Absolute:** Allows you to set an X and Y position for the axis of rotation based on the original shape's location. The axis of rotation is then copied and offset with each duplicated shape.
- **Origin Relative:** Each copy uses its center point as its axis of rotation.
- **Origin Absolute:** Each copy uses the center of the original shape as its axis of rotation.
- **Progressive:** Compounds each shape copy by progressively transforming each copy based on the previous shape's position, rotation, and scale.

X and Y Pivot

The X and Y pivot controls are displayed when the Axis mode is set to Absolute. You can use these position controls to place the axis of rotation.

Rotation

Determines an offset rotation applied to each copy. The rotation is calculated from the offset rotation of the previous copy. To rotate all copies identically, use the Angle parameter on the original shape or use a sTransform node.

Common Controls

Settings tab

The Settings tab in the Inspector is common to all Shape nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

sEllipse



The sEllipse node

The sEllipse node is used to create circular shapes. Like almost all shape nodes, you can only view the sEllipse node’s results through a sRender node.

External Inputs

This node generates shapes and does not have any inputs.

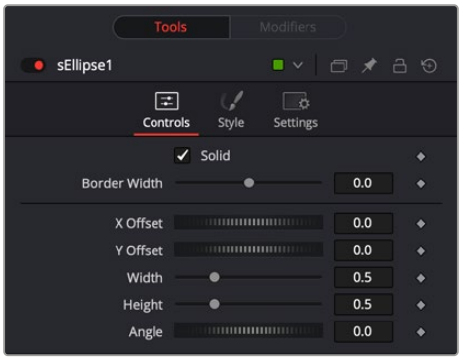
Basic Node Setup

The sEllipse node is a shape generator, meaning it generates a shape and therefore has no input. The output of the sEllipse can go into a sRender for viewing and further compositing or, more likely, connect to another shape node like sGrid or sDuplicate.



An sEllipse node connecting to an sGrid node, and then viewed using an sRender node

Inspector



The sEllipse Controls tab

Controls

The Controls tab is used to define the elliptical shape characteristics, including fill, border, size, and position.

Solid

When enabled, the Solid checkbox fills the elliptical shape with the color defined in the Style tab. When disabled, an outline created by the Border Width control is displayed, and the center is made transparent.

Border Width

This parameter expands or contracts the border around the shape. Although it can be used when the Solid checkbox is enabled, it is primarily used to determine the outline thickness when the checkbox is disabled.

Cap style

When the Solid checkbox is disabled, three cap style options are displayed. The cap styles can create lines with flat, rounded, or squared ends. Flat caps have flat, squared ends, while rounded caps have semi-circular ends. Squared caps have projecting ends that extend half the line width beyond the end of the line.

The caps are not visible unless the length is below 1.0.

Position

The position parameter is only displayed when the Solid checkbox is disabled. It allows you to position the starting point of the shape. When used in conjunction with the length parameter, it positions the gap in the ellipse outline.

Length

The length parameter is only displayed when the Solid checkbox is disabled. A length of 1.0 is a closed shape. Setting the length below 1.0 creates an opening or gap in the outline. Keyframing the length parameters allows you to create write-on style animations.

X and Y Offset

These parameters are used to position the shape left, right, up, and down in the frame. The shape starts in the center of the frame, and the parameters are used to offset the position. The offset coordinates are normalized based on the width of the frame. An X offset of 0.0 is centered, and a value of 0.5 places the center of the shape directly on the right edge of the frame.

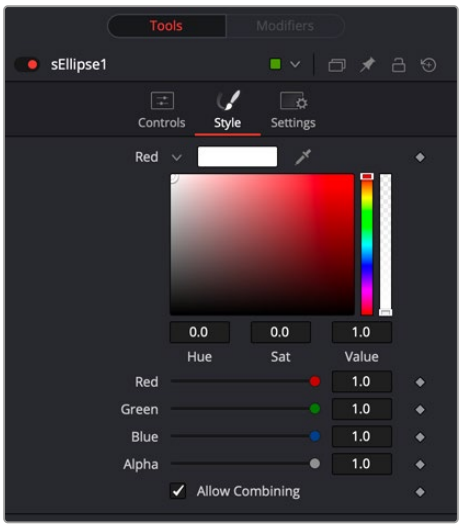
Width/Height

The width and height determine the vertical and horizontal size of the ellipse. If the values are identical, then you have a perfect circle.

Angle

The angle rotates the shape, which on a perfect circle doesn't change the image all that much, but if you create an oval or an outline with a short length, you can rotate the shape based on the center axis.

Style Tab



The sEllipse Style tab

Style

The Style tab is used to assign a color to the shape and control its transparency.

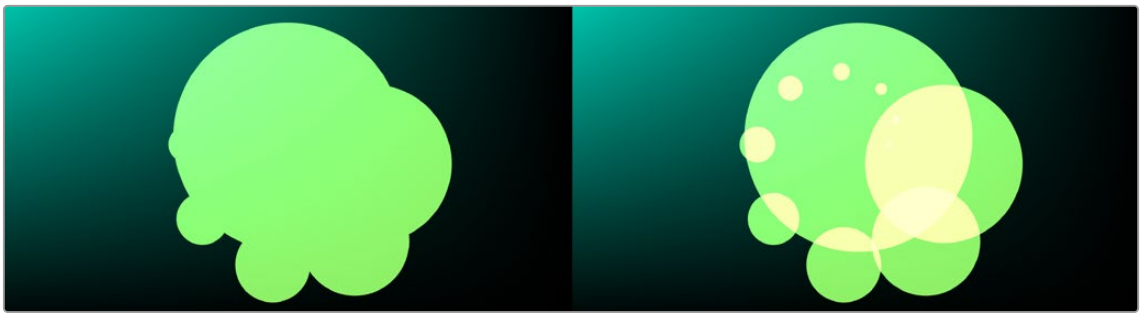
Color

The color controls determine the color of the fill and border. To choose a shape color, you can click the color disclosure arrow, use the color swatch, or drag the eyedropper into the viewer to select a color from an image. The RGBA sliders or number fields can be used to enter each color channel's value or the strength of the Alpha channel.

Allow Combining

When this checkbox is enabled, the Alpha channel value is maintained even when passing through other nodes downstream that may cause the shape to overlap with copies of itself. When disabled, the Alpha channel value may increase when the shape overlaps itself.

For instance, if an ellipse's Alpha channel is set to .5, enabling the Allow Combining checkbox maintains that value even if the shape passes through a Duplicate or Grid node that causes the shape to overlap. Disabling the checkbox causes the Alpha channel values to be compounded at each overlapping area.



Allow Combining Enabled (left), Allow Combining Disabled (right)

Common Controls

Settings tab

The Settings tab in the Inspector is common to all Shape nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

sExpand



The sExpand node

The sExpand node is used to dilate or erode shapes. Like almost all Shape nodes, you can only view the sExpand node’s results through a sRender node.

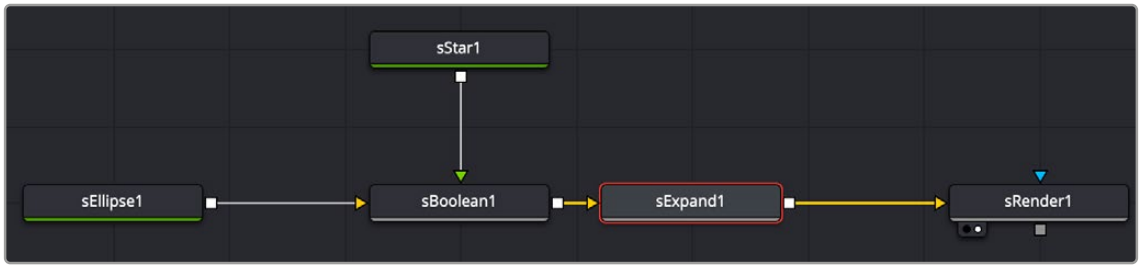
External Inputs

The following input appears on the node’s tile in the Node Editor.

Input1: [orange, required] This input accepts the output of another shape node. This shape or compound shape connected to this input is either eroded or dilated.

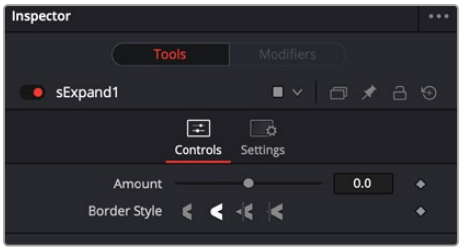
Basic Node Setup

The sExpand node takes a single input that is most often from a compound shape. However, it can be used on single shapes like sStars and sNgons. The output of the sExpand can then be output to another shape node or to a sRender node for viewing or compositing into the greater node tree.



Star and ellipse shapes combined in an sBoolean node, then output to an sExpand for dilating or eroding

Inspector



The sExpand Controls tab

Controls

The Controls tab includes all of the parameters for the sExpand node.

Amount

A positive value dilates the shape while a negative value erodes it.

Border Style

The border style controls how the expanded or contracted shapes join at the corners. There are four styles provided as options. Bevel squares off the corners. Round creates rounded corners. Miter and Miter Clip maintain pointed edges, until a certain threshold. The Threshold is set by the Miter limit slider.

Miter Limit

The Miter parameter is only displayed when the Miter or Miter Clip border style is selected. The miter limit determines when the pointed edges become beveled based on the shape's thickness.

Common Controls

Settings tab

The Settings tab in the Inspector is common to all Shape nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

sGrid



The sGrid node

The sGrid node replicates the shape on an X and Y grid and adds the ability to offset the rows and columns. Like almost all Shape nodes, you can only view the sGrid node's results through a sRender node.

External Inputs

The following input appears on the node's tile in the Node Editor.

Input1: [orange, required] This input accepts the output of another Shape node. The shape connected to this input is replicated on a custom grid.

Basic Node Setup

The sGrid node takes a single input from a single or compound shape from am sMerge or sBoolean node. The sGrid node places the incoming shape on a grid of rows and columns. The output of the sGrid can then be output to another Shape node or a sRender node for viewing or compositing into the greater node tree.



An sEllipse shape connected to an sGrid and then output to an sRender for viewing and combining with other elements

Inspector



The sGrid Controls tab

Controls

The Controls tab is used to determine the number of grid cells and their offset position.

Grid Cells X and Y

These parameters set the number of cells on the grid, both horizontally and vertically. For instance, entering 5 in the X and Y number field creates five rows of the shape and five columns.

X and Y Offset

Sets the X and Y distance between the rows and columns. An offset value of 0.0 will have all the rows and columns on top of each other. Entering X Offset at 1.0 would spread the columns the width to the frame.

Common Controls

Settings tab

The Settings tab in the Inspector is common to all Shape nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

sJitter



The sJitter node

The sJitter node is most often used to randomly position an array of shapes generated from a sGrid or sDuplicate node. However, it includes an auto-animating random mode that can be used to distort and randomly jitter single shapes.

Like almost all Shape nodes, you can only view the sJitter node's results through a sRender node.

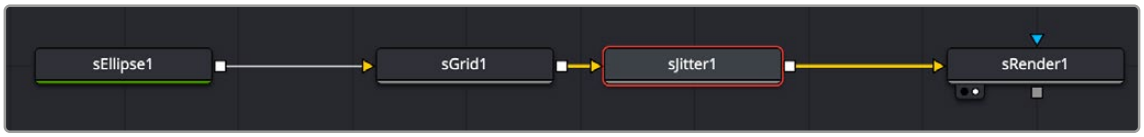
External Inputs

The following input appears on the node's tile in the Node Editor.

Input1: [orange, required] This input accepts the output of another Shape node. The shape connected to this input is offset, distorted, and animated based on the sJitter node settings.

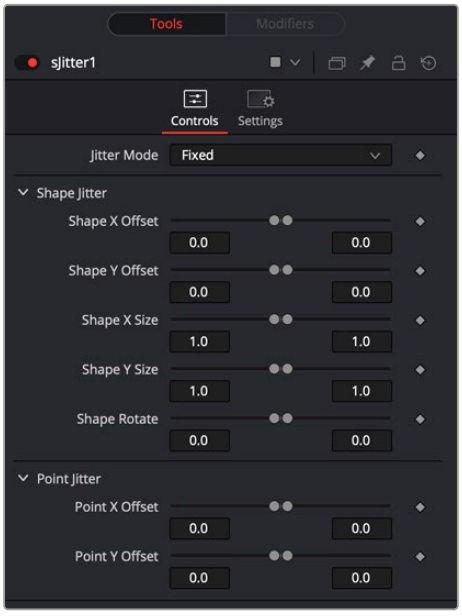
Basic Node Setup

The sJitter node takes an array of shapes from a sGrid or sDuplicate node and randomly changes their position, size, and rotation. The output of the sExpand can then be output to another Shape node or to a sRender node for viewing or compositing into the greater node tree.



An array of shapes created by the sGrid node input into an sJitter node to randomly offset or scale the shapes

Inspector



The sJitter Controls tab

Controls

The Controls tab offers range sliders that determine the variation amount for offset, size, and rotation. The Point Jitter parameters are used to offset the invisible points that create the vector shapes.

Jitter Mode

The Jitter Mode menu allows you to choose between static position and size offsets or enabling an auto-animation mode. Leaving the default Fixed selection allows you to offset a grid of shapes, animating with keyframes or modifiers if needed. The Random menu selection auto-animates the parameters based on the range you define using the range sliders. If all the range sliders are left in the default position, no random animation is created. Increasing the range on any given parameter will randomly animate that parameter between the range slider values.

Shape X and Y Offset

These parameters set the horizontal and vertical offset from the shape array's original position. This is done randomly, so not all shapes in the array will offset by the same amount.

Shape X and Y Size

These parameters set the horizontal and vertical scaling for each shape in an array. The left range value decreases the scale, and the right range value increases the scale. This is done randomly, so not all shapes in the array will scale by the same amount.

Shape Rotate

This parameter rotates each shape in an array.

Point Jitter

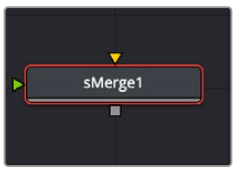
The X and Y Point Jitter parameters use the vector control points to distort the shape. This can be used to give a distressed appearance to ellipses or wobbly animation to other shapes.

Common Controls

Settings tab

The Settings tab in the Inspector is common to all Shape nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

sMerge



The sMerge node

The sMerge node combines shapes similar to a standard Merge node, except the sMerge node can accept more than two shape inputs.

Like almost all Shape nodes, you can only view the sMerge node's results through a sRender node.

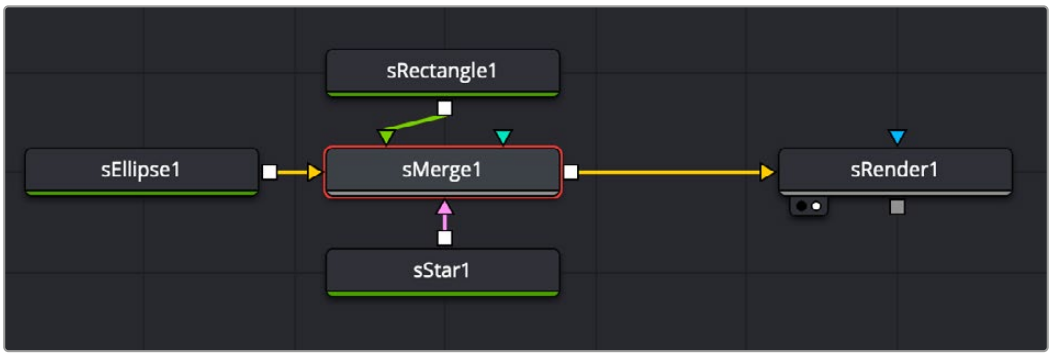
External Inputs

The node displays only two inputs first, but as each shape node is connected, a new input appears on the node, assuring there is always one free to add a new shape into the composite.

Input[#]: These multi-colored inputs are used to connect multiple Shape node. There is no limit to the number of inputs this node can accept. The node dynamically adds more inputs as needed, ensuring that there is always at least one input available.

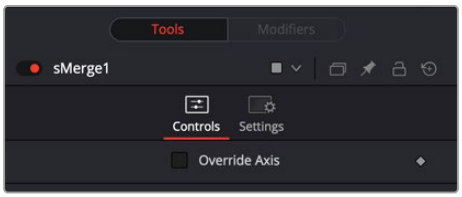
Basic Node Setup

The sMerge node is used to combine two Shape nodes. In terms of layering, each subsequent input is placed over the previous input. For instance, the first shape connected to the orange input is the bottom-most shape, the green input is layered over it, and if a third shape is connected to the pink input, that is the topmost layer.



Three shapes are combined in an sMerge node, then output to an sRender for viewing and further processing.

Inspector



The sMerge Controls tab

Controls

The only control for the sMerge node is the Override Axis checkbox, which overrides the shape's axis.

Common Controls

Settings tab

The Settings tab in the Inspector is common to all Shape nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

sNGon



The sNGon node

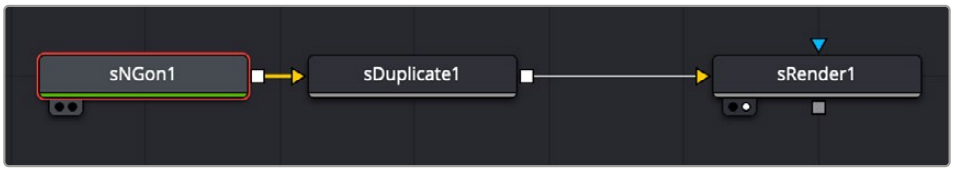
The sNGon node is used to create multi-sided shapes like triangles, pentagons, and octagons. Like almost all Shape nodes, you can only view the sNGon node's results through a sRender node.

External Inputs

This node generates shapes and does not have any inputs.

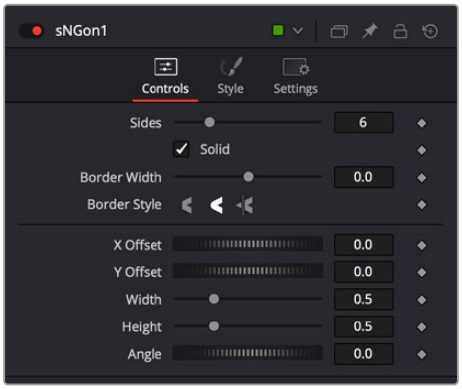
Basic Node Setup

The sNGon node is a shape generator, meaning it generates a shape and therefore has no inputs. The output of the sNGon can go into a sRender for viewing and further compositing or, more likely, connect to another Shape node like sGrid or sDuplicate.



An sNGon node connecting to an sDuplicate node, and then viewed using an sRender node.

Inspector



The sNGon Controls tab

Controls

The Controls tab is used to define multi-sided shape characteristics, including fill, border, size, and position.

Solid

When enabled, the Solid checkbox fills the NGon shape with the color defined in the Style tab. When disabled, an outline created by the Border Width control is displayed, and the center is made transparent.

Border Width

This parameter expands or contracts the border around the shape. Although it can be used when the Solid checkbox is enabled, it is primarily used to determine the outline thickness when the checkbox is disabled.

Border Style

The Border Style parameter controls how the sides of the NGon join at the corners. There are three styles provided as options. Bevel squares off the corners. Round creates rounded corners. Miter maintains pointed corners.

Cap style

When the Solid checkbox is disabled, three cap style options are displayed. The cap styles can create lines with flat, rounded, or squared ends. Flat caps have flat, squared ends, while rounded caps have semi-circular ends. Squared caps have projecting ends that extend half the line width beyond the end of the line.

The caps are not visible unless the length is below 1.0.

Position

The Position parameter is only displayed when the Solid checkbox is disabled. It allows you to position the starting point of the shape. When used in conjunction with the Length parameter, it positions the gap in the outline.

Length

The Length parameter is only displayed when the Solid checkbox is disabled. A length of 1.0 is a closed shape. Setting the length below 1.0 creates an opening or gap in the outline. Keyframing the Length parameters allows you to create write-on style animations.

X and Y Offset

These parameters are used to position the shape left, right, up, and down in the frame. The shape starts in the center of the frame, and the parameters are used to offset the position. The offset coordinates are normalized based on the width of the frame. So, an X offset of 0.0 is centered and a value of 0.5 places the center of the shape directly on the right edge of the frame.

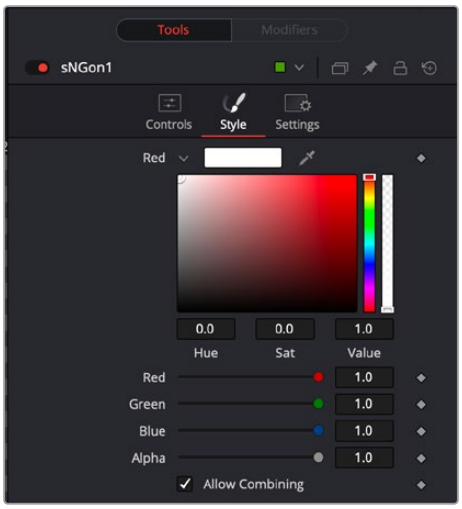
Width/Height

The Width and Height parameters determine the vertical and horizontal size of the ellipse. If the values are identical, then all sides are of equal length.

Angle

The Angle parameter rotates the shape based on the center axis.

Style Tab



The sNGon Style tab

Style

The Style tab is used to assign a color to the shape and control its transparency.

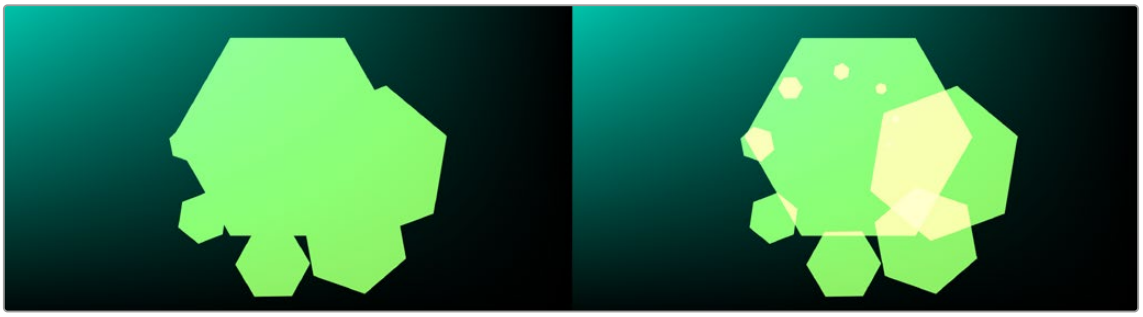
Color

The Color controls determine the color of the fill and border. To choose a shape color, you can click the color disclosure arrow, use the color swatch, or drag the eyedropper into the viewer to select a color from an image. The RGBA sliders or number fields can be used to enter each color channel's value or the strength of the Alpha channel.

Allow Combining

When this checkbox is enabled, the Alpha channel value is maintained even when passing through other nodes downstream that may cause the shape to overlap with copies of itself. When disabled, the Alpha channel value may increase when the shape overlaps itself.

For instance, if a NGon Alpha channel is set to .5, enabling the Allow Combining checkbox maintains that value even if the shape passes through a duplicate or grid node that causes the shape to overlap. Disabling the checkbox causes the Alpha channel values to be compounded at each overlapping area.



Allow Combining Enabled (left), Allow Combining Disabled (right)

Settings tab

The Settings tab in the Inspector is common to all Shape nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

sOutline



The sOutline node

The sOutline node is used to create outlines from merged or boolean compound shapes. The individual shapes retain their own style, color, size, position, and other characteristics. The only difference is the border thickness, border style, position, and length are applied to all incoming shapes uniformly in the sOutline node.

Like almost all shape nodes, you can only view the sOutline node's results through a sRender node.

External Inputs

The following input appears on the node's tile in the Node Editor:

Input1: [orange, required] This input accepts the another shape node's output, but more likely a compound shape from a sMerge or sBoolean. An outline is created from the compound shape connected to this input.

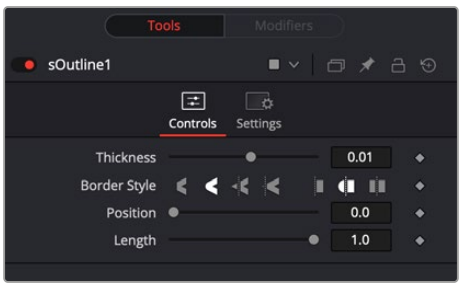
Basic Node Setup

The sOutline node takes a single input that is most often from a compound shape, however, it can sometimes be useful on single shapes to create double outlines. The output of the sOutline can then be output to another Shape node or to a sRender node for viewing or compositing into the greater node tree.



A compound shape from an sBoolean node is connected to an sOutline for creating a complex outlined shape

Inspector



The sOutline Controls tab

Controls

The Controls tab is used to define the outline thickness, border and cap style, position, and length that is applied to the compound shape connected to the input.

Thickness

This parameter controls the width of the outline.

Border Style

The Border Style parameter controls how the outline joins at the corners. There are three styles provided as options. Bevel squares off the corners. Round creates rounded corners. Miter maintains pointed corners.

Cap style

Three Cap Style options are used to create lines with flat, rounded, or squared ends. Flat caps have flat, squared ends, while rounded caps have semi-circular ends. Squared caps have projecting ends that extend half the line width beyond the end of the line.

The caps are not visible unless the length is below 1.0.

Position

The Position parameter allows you to position the starting point of the shape. When used in conjunction with the Length parameter, it positions the gap in the outline.

Length

The Length parameter controls the end position of the outline. A length of 1.0 is a closed shape. Setting the length below 1.0 creates an opening or gap in the outline. Keyframing the Length parameters allows you to create write-on style animations.

Common Controls

Settings tab

The Settings tab in the Inspector is common to all Shape nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

sPolygon [sPly]



The sPolygon node

The sPolygon tool allows users to draw custom shapes. Through intuitive point manipulation, this tool allows users to effortlessly generate intricate forms, curves, and lines, which tie in with the existing shape tools in Fusion. It provides the freedom to design custom elements, making it a valuable tool in motion graphics design.

Inputs

This node generates shapes and does not have any inputs.

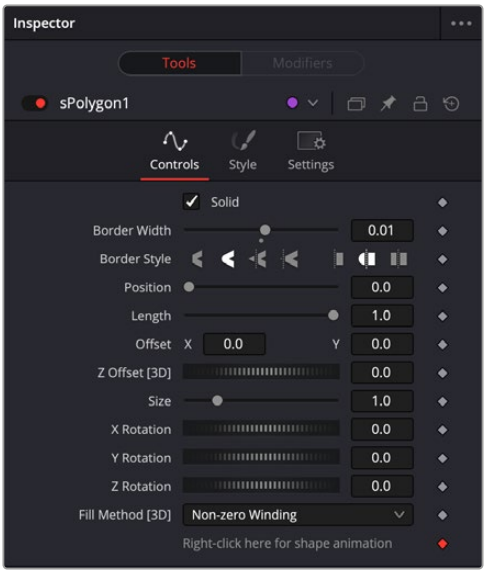
Basic Node Setup

The sPolygon node is a shape generator, meaning it generates a shape and therefore has no inputs. The output of the sPolygon can go into a sRender for viewing and further compositing or, more likely, connect to another Shape node like sGrid or sDuplicate.



An sPolygon node connecting to an sDuplicate node, and then viewed using an sRender node

Inspector



The sPolygon Controls tab

Controls

The Controls tab is used to define the polygon characteristics, including fill, border, size, and position.

Solid

When enabled, the Solid checkbox fills the rectangle shape with the color defined in the Style tab. When disabled, an outline created by the Border Width control is displayed, and the center is made transparent.

Border Width

This parameter expands or contracts the border around the shape. Although it can be used when the Solid checkbox is enabled, it is primarily used to determine the outline thickness when the checkbox is disabled.

Border Style

The Border Style parameter controls how the sides of the rectangle join at the corners. There are three styles provided as options. Bevel squares off the corners. Round creates rounded corners. Miter maintains pointed corners.

Cap style

The cap styles can create lines with flat, rounded, or squared ends. Flat caps have flat, squared ends, while rounded caps have semi-circular ends. Squared caps have projecting ends that extend half the line width beyond the end of the line.

The caps are not visible unless the length is below 1.0.

Position

The Position parameter is only displayed when the Solid checkbox is disabled. It allows you to position the starting point of the shape. When used in conjunction with the Length parameter, it positions the gap in the outline.

Length

The Length parameter is only displayed when the Solid checkbox is disabled. A length of 1.0 is a closed shape. Setting the length below 1.0 creates an opening or gap in the outline. Keyframing the Length parameters allows you to create write-on style animations.

X and Y Offset

These parameters are used to position the shape left, right, up, and down in the frame. The shape starts in the center of the frame, and the parameters are used to offset the position. The offset coordinates are normalized based on the width of the frame. So an X offset of 0.0 is centered and a value of 0.5 places the center of the shape directly on the right edge of the frame.

Z Offset [3D]

If used with the 3D toolset, you can offset the polygon in back and forth along the Z-axis by adjusting these controls.

Size

Use the Size control to adjust the scale of the polygon shape, without affecting the relative behavior of the points that compose the shape or setting a keyframe in the shape animation.

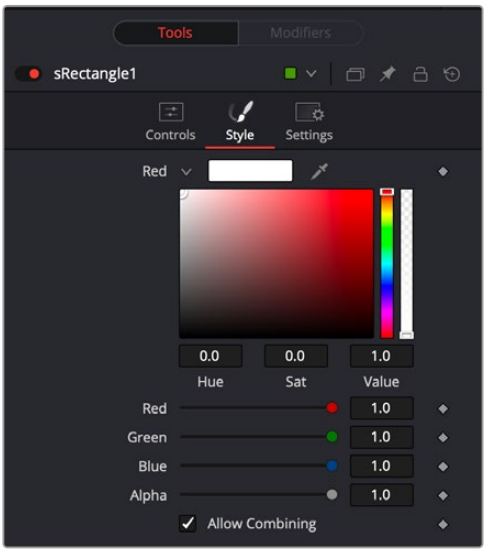
X,Y,Z Rotation

Use these three controls to adjust the rotation angle of the shape along any axis.

Fill Method [3D]

The Fill Method menu offers two different techniques for dealing with overlapping regions of a polyline. If overlapping segments in a shape are causing undesirable holes to appear, try switching the setting of this control from Alternate to Non Zero Winding.

Style Tab



The sPolygon Style tab

Style

The Style tab is used to assign color to the shape and control its transparency.

Color

The Color parameter controls determine the color of the fill and border from the sPolygon node. To choose a shape color, you can click the color disclosure arrow and use the color swatch, or drag the eye dropper into the viewer to select a color from an image. The RGBA sliders or number fields can be used to enter the value of each color channel or the strength of the Alpha channel.

Allow Combining

When this checkbox is enabled, the Alpha channel value is maintained, even when passing through other nodes downstream that may cause the shape to overlap with copies of itself. When disabled, the Alpha channel value may increase when the shape overlaps itself. For instance, if a rectangle Alpha channel is set to .5, enabling the Allow Combining checkbox maintains that value even if the shape passes through a duplicate or grid node that causes the shape and Alpha channel to overlap. Disabling the checkbox causes the Alpha channel values to be compounded at each overlapping area.

Common Controls

Settings tab

The Settings tab in the Inspector is common to all Shape nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Adding Points

Adding Points to a polygonal shape is relatively simple. Immediately after adding the node to the Node Editor, there are no points, but the tool will be in Click Append mode. Click once in the viewer wherever a point is required for the shape. Continue clicking to draw the shape. When the shape is complete, click on the initial point again to close the shape.

When the shape is closed, the mode of the polyline will change to Insert and Modify. This allows for the adjusting and adding of additional points to the shape by clicking on segments of the polyline. To lock down the shape and prevent accidental changes, switch the Polyline mode to Done using the Polyline toolbar or contextual menu.



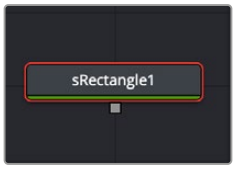
The sPolygon toolbar

When a Polygon (or B-Spline) shape is added to a node, a toolbar appears above the viewer, offering easy access to modes. Hold the pointer over any button in the toolbar to display a tooltip that describes that button's function.

- **Click:** Click is the default option when creating a polyline (or B-Spline) shape. It is a Bézier style drawing tool. Clicking sets a control point and appends the next control point when you click again in a different location.
- **Draw:** Draw is a freehand drawing tool. It creates a shape similar to drawing with a pencil on paper. You can create a new shape using the Draw tool, or you can extend an existing open spline by clicking the Draw tool and starting to draw from the last control point.
- **Insert and Modify:** Insert adds a new control point along the spline and lets you modify it.
- **Modify Only:** Modify allows you to safely move or smooth any exiting point along a spline without worrying about adding new points accidentally.
- **Done:** Prevents any point along the spline from being moved or modified. Also, new points cannot be added. You can, however, move and rotate the entire spline.
- **Close:** Closes an open spline.
- **Smooth:** Changes the selected control point from a linear to a smooth curve.
- **Linear:** Changes the selected control point from a smooth curve to linear.
- **Select All:** Selects all the control points on the spline.
- **Show Key Points:** Shows or hides the control points along the spline.
- **Show All Handles:** Shows or hides the Bézier handles along the polyline.
- **Shape Box:** Places a reshape rectangle around the selected spline shape. Using the reshape rectangle, you can deform groups of control points or entire shapes much easier than modifying each point.
- **Delete Points:** Deletes the selected control point(s).
- **Reduce Points:** Opens a Freehand precision window that can be used to reduce the number of controls points on a spline. This can make the paint stroke easier to modify, especially if it has been created using the Draw tool.

Change the way the toolbar is displayed by right-clicking on the toolbar and selecting from the options displayed in the toolbar's contextual menu.

sRectangle



The sRectangle node

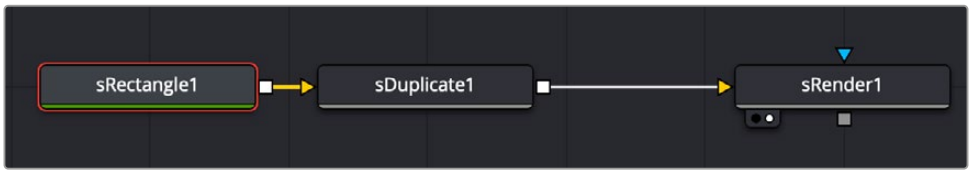
The sRectangle node is used to create rectangular shapes. Like almost all shape nodes, you can only view the sRectangle node's results through a sRender node.

External Inputs

This node generates shapes and does not have any inputs.

Basic Node Setup

The sRectangle node is a shape generator, meaning it generates a shape and therefore has no inputs. The output of the sRectangle can go into a sRender for viewing and further compositing or, more likely, connect to another Shape node like sGrid or sDuplicate.



An sRectangle node connecting to an sDuplicate node, and then viewed using an sRender node

Inspector



The sRectangle Controls tab

Controls

The Controls tab is used to define the rectangle characteristics, including fill, border, size, and position.

Solid

When enabled, the Solid checkbox fills the rectangle shape with the color defined in the Style tab. When disabled, an outline created by the Border Width control is displayed, and the center is made transparent.

Border Width

This parameter expands or contracts the border around the shape. Although it can be used when the Solid checkbox is enabled, it is primarily used to determine the outline thickness when the checkbox is disabled.

Border Style

The Border Style parameter controls how the sides of the rectangle join at the corners. There are three styles provided as options. Bevel squares off the corners. Round creates rounded corners. Miter maintains pointed corners.

Cap style

When the Solid checkbox is disabled, three Cap Style options are displayed. The cap styles can create lines with flat, rounded or squared ends. Flat caps have flat, squared ends, while rounded caps have semi-circular ends. Squared caps have projecting ends that extend half the line width beyond the end of the line.

The caps are not visible unless the length is below 1.0.

Position

The Position parameter is only displayed when the Solid checkbox is disabled. It allows you to position the starting point of the shape. When used in conjunction with the Length parameter, it positions the gap in the outline.

Length

The Length parameter is only displayed when the Solid checkbox is disabled. A length of 1.0 is a closed shape. Setting the length below 1.0 creates an opening or gap in the outline. Keyframing the Length parameters allows you to create write-on style animations.

X and Y Offset

These parameters are used to position the shape left, right, up, and down in the frame. The shape starts in the center of the frame, and the parameters are used to offset the position. The offset coordinates are normalized based on the width of the frame. So an X offset of 0.0 is centered and a value of 0.5 places the center of the shape directly on the right edge of the frame.

Width/Height

The Width and Height parameters determine the vertical and horizontal size of the rectangle. If the values are identical, then you have a square.

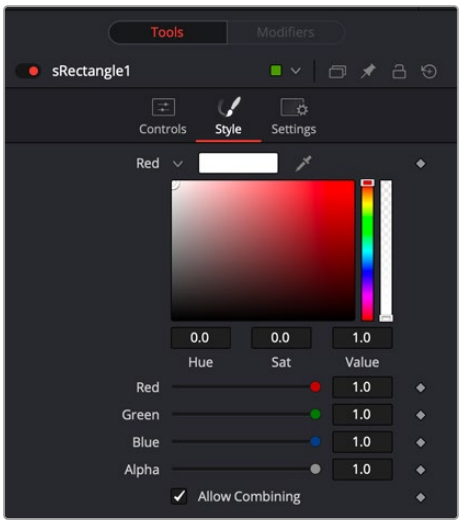
Corner Radius

This parameter determines if the corners of the rectangle are sharp or rounded. A value of 0.0 produces sharp corners, while a value of 1.0 will create a circle from a staring square shape or a pill shape from a rectangle.

Angle

The Angle parameter rotates the shape based on the center axis.

Style Tab



The sRectangle Style tab

Style

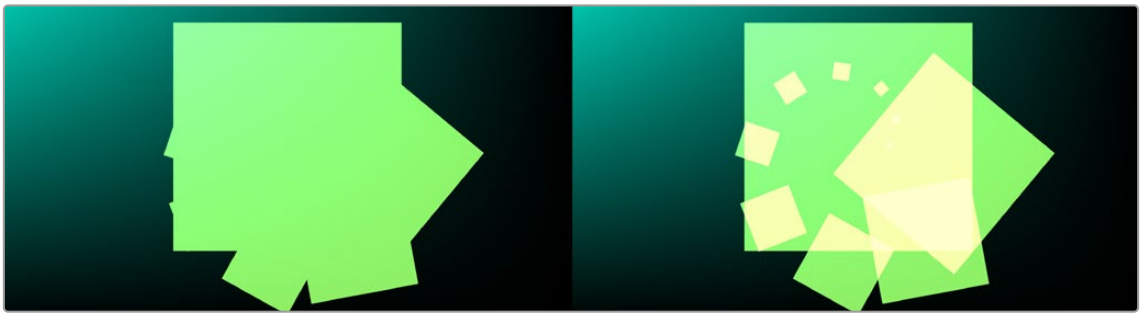
The Style tab is used to assign color to the shape and control its transparency.

Color

The Color parameter controls determine the color of the fill and border from the sRectangle node. To choose a shape color, you can click the color disclosure arrow and use the color swatch, or drag the eye dropper into the viewer to select a color from an image. The RGBA sliders or number fields can be used to enter the value of each color channel or the strength of the Alpha channel.

Allow Combining

When this checkbox is enabled, the Alpha channel value is maintained even when passing through other nodes downstream that may cause the shape to overlap with copies of itself. When disabled, the Alpha channel value may increase when the shape overlaps itself. For instance, if a rectangle Alpha channel is set to .5, enabling the Allow Combining checkbox maintains that value even if the shape passes through a duplicate or grid node that causes the shape and Alpha channel to overlap. Disabling the checkbox causes the Alpha channel values to be compounded at each overlapping area.



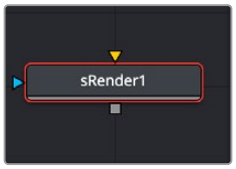
Allow Combining Enabled (left), Allow Combining Disabled (right)

Common Controls

Settings tab

The Settings tab in the Inspector is common to all Shape nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

sRender



The sRender node

The sRender node converts the vector shapes to an image. The output of the sRender allows the vector shapes to be integrated with other elements in a composite.

Inputs

There is one input on the Background node for an Effect Mask input.

Input1: [orange, required] This input accepts the output of your final shape node. A rendered bitmap image is created from the sRender node for compositing into the rest of your comp.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the displayed area to only those pixels within the mask.

Basic Node Setup

The sRender node is always placed at the end of a string of Shape nodes. The output connects to other 2D nodes like a Soft Glow node.



Multiple Shape nodes connected to the sRender node and then processed and composited with a title

Inspector



The sRender Image tab

Image Tab

The controls in this tab are used to set the resolution, color depth, and pixel aspect of the image produced by the sRender node.

Process Mode

Use this menu control to select the Fields Processing mode used by Fusion to render the resulting image. The default Full Frames option is appropriate for progressive formats.

Width/Height

This pair of controls are used to set the Width and Height dimensions of the image to be created by the sRender node.

Pixel Aspect

This control is used to specify the Pixel Aspect ratio of the created images. An aspect ratio of 1:1 would generate a square pixel with the same dimensions on either side (like a computer display monitor), and an aspect of 0.9:1 would create a slightly rectangular pixel (like an NTSC monitor).

> **NOTE:** Right-click on the Width, Height, or Pixel Aspect controls to display a menu listing the file formats defined in the preferences Frame Format tab. Selecting any of the listed options will set the width, height, and pixel aspect to the values for that format, accordingly.

Auto Resolution

When this checkbox is selected, the width, height, and pixel aspect of the image created by the node will be locked to values defined in the composition's Frame Format preferences. If the Frame Format preferences change, the resolution of the image produced by the node will change to match. Disabling this option can be useful to build a composition at a different resolution than the eventual target resolution for the final render.

Depth

The Depth button array is used to set the pixel color depth of the image created by the Creator node. 32-bit pixels require 4X the memory of 8-bit pixels but have far greater color accuracy. Float pixels allow high dynamic range values outside the normal 0..1 range, for representing colors that are brighter than white or darker than black.

Source Color Space

You can use the Source Color Space menu to set the Color Space of the footage to help achieve a linear workflow. Unlike the Gamut tool, this doesn't perform any actual color space conversion, but rather adds the source space data into the metadata, if that metadata doesn't exist. The metadata can then be used downstream by a Gamut tool with the From Image option, or in a Saver, if explicit output spaces are defined there. There are two options to choose from:

- **Auto:** Automatically reads and passes on the metadata that may be in the image.
- **Space:** Displays a Color Space Type menu where you can choose the correct color space of the image.

Source Gamma Space

Using the Curve Type menu, you can set the Gamma Space of the footage and choose to remove it by way of the Remove Curve check box when working in a linear workflow. There are three choices in the Curve Type menu

- **Auto:** Automatically reads and passes on the metadata that may be in the image.
- **Space:** Displays a Gamma Space Type menu where you can choose the correct gamma curve of the image.
- **Log:** Brings up the Log/Lin settings, similar to the Cineon tool.

For more information, *see Chapter 39, "Film Nodes,"* in the Fusion Reference Manual.

Remove Curve

Depending on the selected Gamma Space or on the Gamma Space found in Auto mode, the Gamma Curve is removed from, or a log-lin conversion is performed on, the material, effectively converting it to a linear output space.

Common Controls

Settings tab

The Settings tab in the Inspector is common to all Shape nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

sStar



The sStar node

The sStar node is used to create multi-point star shapes. Like almost all Shape nodes, you can only view the sStar node's results through a sRender node.

External Inputs

This node generates shapes and does not have any inputs.

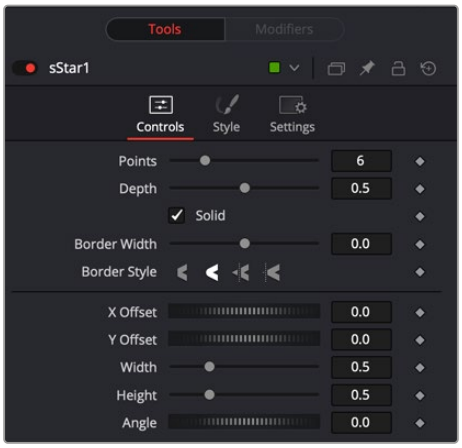
Basic Node Setup

The sStar node is a shape generator, meaning it generates a shape and therefore has no inputs. The output of the sStar can go into a sRender for viewing and further compositing or, more likely, connect to another shape node like sGrid or sDuplicate.



An sStar node connecting to an sDuplicate node, and then viewed using an sRender node

Inspector



The sStar Controls tab

Controls

The Controls tab is used to define the star shape's characteristics, including number of points, depth, fill, border, size, and position.

Points

This slider determines the number of points or arms on the star.

Depth

The depth slider controls the inner radius or width of the arms. A depth of 0.001 makes hair-thin arms, while a depth of 1.0 makes a faceted circle.

Solid

When enabled, the Solid checkbox fills the star shape with the color defined in the Style tab. When disabled, an outline created by the Border Width control is displayed, and the center is made transparent.

Border Width

This parameter expands or contracts the border around the shape. Although it can be used when the Solid checkbox is enabled, it is primarily used to determine the outline thickness when the checkbox is disabled.

Border Style

The Border Style parameter controls how the sides of the star join at the corners. There are three styles provided as options. Bevel squares off the corners. Round creates rounded corners. Miter maintains pointed corners.

Cap style

When the Solid checkbox is disabled, three cap style options are displayed. The cap styles can create lines with flat, rounded or squared ends. Flat caps have flat, squared ends while rounded caps have semi-circular ends. Squared caps have projecting ends that extend half the line width beyond the end of the line.

The caps are not visible unless the length is below 1.0.

Position

The Position parameter is only displayed when the Solid checkbox is disabled. It allows you to position the starting point of the shape. When used in conjunction with the length parameter, it positions the gap in the outline.

Length

The Length parameter is only displayed when the Solid checkbox is disabled. A length of 1.0 is a closed shape. Setting the length below 1.0 creates an opening or gap in the outline. Keyframing the Length parameters allows you to create write-on style animations.

X and Y Offset

These parameters are used to position the shape left, right, up, and down in the frame. The shape starts in the center of the frame, and the parameters are used to offset the position. The offset coordinates are normalized based on the width of the frame. So an X offset of 0.0 is centered and a value of 0.5 places the center of the shape directly on the right edge of the frame.

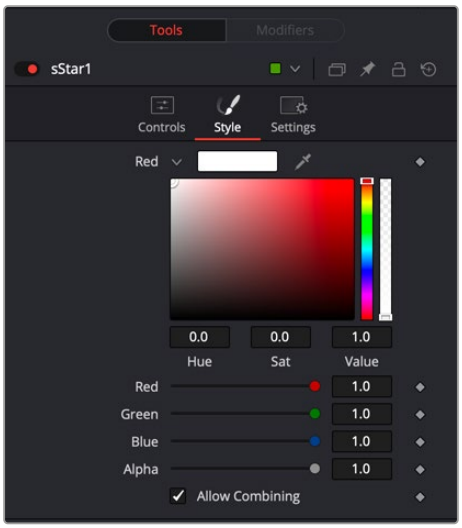
Width/Height

The Width and Height parameters determine the vertical and horizontal size of the star. If the values are identical, then all arms of the star are of equal length.

Angle

The Angle parameter rotates the shape based on the center axis.

Style Tab



The sStar Style tab

Style

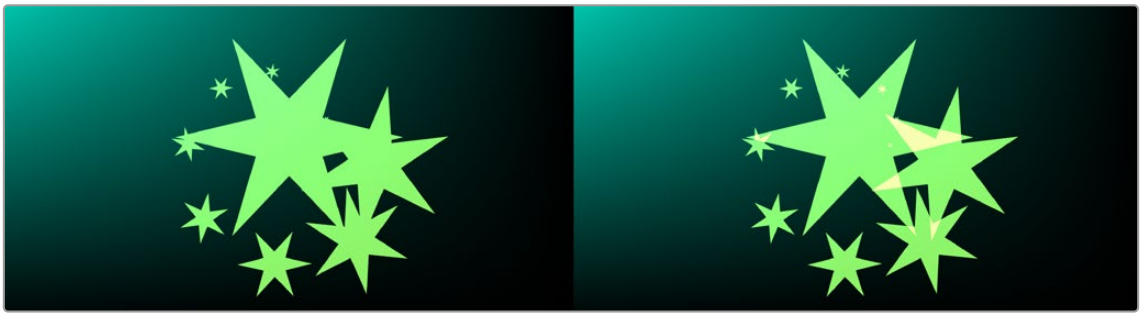
The Style tab is used to assign color to the shape and control its transparency.

Color

The Color controls determine the color of the fill and border from the sStar node. To choose a shape color, you can click the color disclosure arrow and use the color swatch, or drag the eye dropper into the viewer to select a color from an image. The RGBA sliders or number fields can be used to enter the value of each color channel or the strength of the Alpha channel.

Allow Combining

When this checkbox is enabled, the Alpha channel value is maintained even when passing through other nodes downstream that may cause the shape to overlap with copies of itself. When disabled, the Alpha channel value may increase when the shape overlaps itself. For instance, if a star Alpha channel is set to .5, enabling the Allow Combining checkbox maintains that value even if the shape passes through a duplicate or grid node that causes the shape and Alpha channel to overlap. Disabling the checkbox causes the Alpha channel values to be compounded at each overlapping area.



Allow Combining Enabled (left), Allow Combining Disabled (right)

Common Controls

Settings tab

The Settings tab in the Inspector is common to all Shape nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

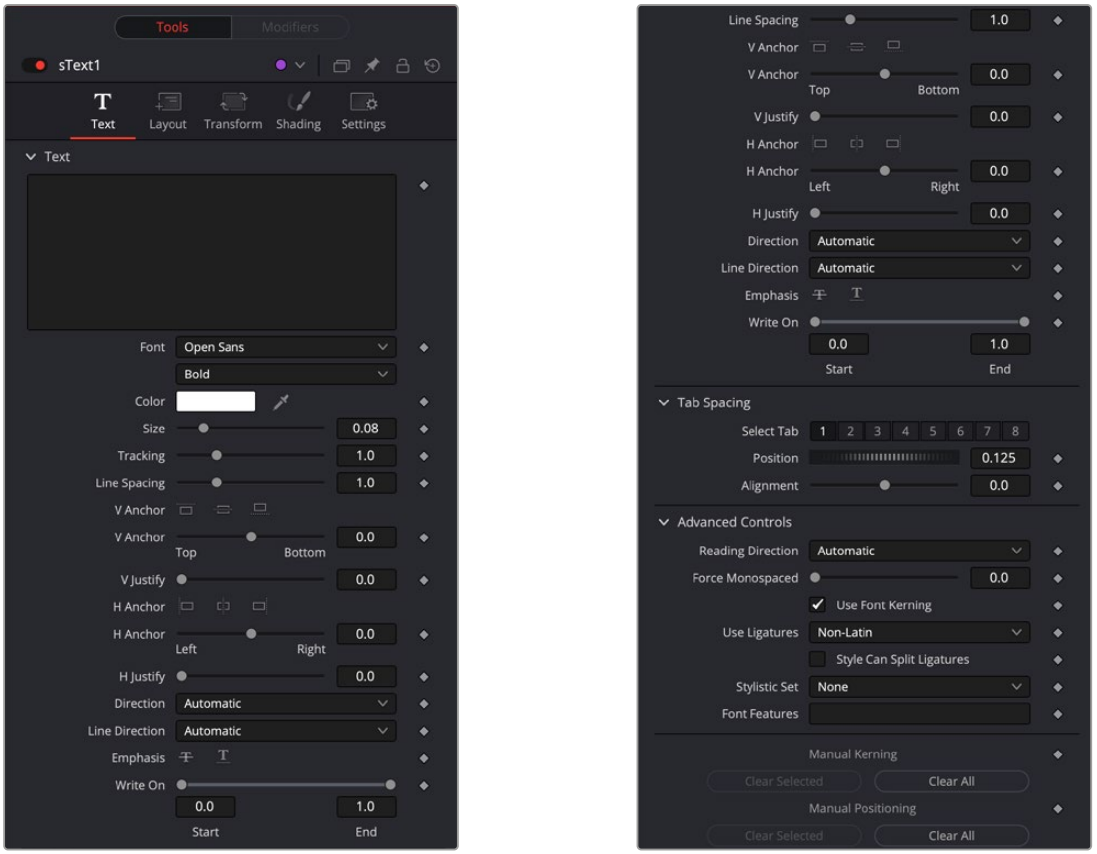
sText [sTxt]



The sText node

Introduction

Designed for the Fusion Shape environment, the sText node is an advanced character generator capable of multiple styles, 3D transformations, and several layers of shading. The sText node is a “Shape” version of the 2D Text+ node. The controls for this node are mostly identical to the controls for the 2D version in almost all respects, except sText doesn’t support gradient colors.



The sText controls

sTransform



The sTransform node

The sTransform node is used to add an additional set of transform controls to the existing controls that are contained in Shape nodes. These additional transforms can be used to create hierarchical animations. For instance, you can use a sStar's built-in Angle control to spin the star around. The star can then be output to an sTransform node. The rotation control in the sTransform can be used to orbit the star around the frame.

Like almost all Shape nodes, you can only view the sStar node's results through a sRender node.

External Inputs

The following input appears on the node's tile in the Node Editor:

Input1: [orange, required] This input accepts the output of another Shape node. The shape connected to this input is moved, scaled, and rotated based on the sTransform settings.

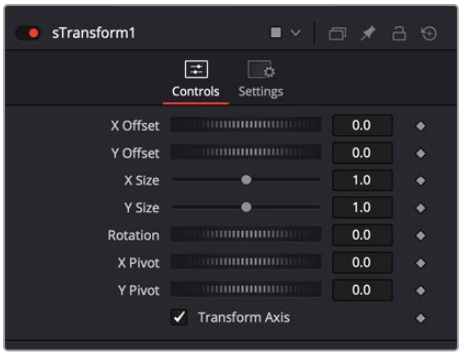
Basic Node Setup

The sTransform node takes the input from another Shape node and adds another set of transforms or hierarchical animation. The output of the sTransform can go into a sRender for viewing and further compositing.



An sStar node connecting to an sTransform node, and then viewed using an sRender node

Inspector



The sTransform Controls tab

Controls

The Controls tab is used to define the add a set of transform controls to the incoming shape.

X and Y Offset

These parameters are used to position the shape left, right, up, and down in the frame. The shape starts in the center of the frame, and the parameters are used to offset the position. The offset coordinates are normalized based on the width of the frame. So an X offset of 0.0 is centered and a value of 0.5 places the center of the shape directly on the right edge of the frame.

X and Y Size

The X and Y Size determine the vertical and horizontal scaling of the incoming shape. If the values are different then the shape will be skewed from its original design.

Rotation

The dial rotates the shape based on the pivot controls.

X and Y Pivot

These parameters position the axis of rotation for the incoming shape. The pivot point is visible in the viewer as a red X. The X can be dragged in the viewer for positioning.

Transform Axis

Check this box to apply the transform to the shape's axis.

Common Controls

Settings tab

The Settings tab in the Inspector is common to all Shape nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Common Controls

Nodes that handle Shape operations share a number of identical controls in the Inspector. This section describes controls that are common amongst Shape nodes.

Settings Tab

The Settings tab in the Inspector can be found on every Shape node. Most of the controls listed here are only found in the sRender node but a few are common to all Shape nodes.

Blend (sRender only)

The Blend control is used to blend between the tool's original image input and the tool's final modified output image. When the blend value is 0.0, the outgoing image is identical to the incoming. Normally, this causes the tool to skip processing entirely, copying the input straight to the output.

Process When Blend Is 0.0 (sRender only)

The tool is processed even when the input value is zero. This can be useful if processing of this node is scripted to trigger another task but the value of the node is set to 0.0.

Red/Green/Blue/Alpha Channel Selector (sRender only)

These four buttons are used to limit the effect of the tool to specified color channels. This filter is often applied after the tool has been processed.

For example, if the red button on a Blur tool is deselected, the blur is first applied to the image, then the red channel from the original input is copied back over the red channel of the result.

There are some exceptions, such as tools where deselecting these channels causes the tool to skip processing that channel entirely. Tools that do this generally possess a set of identical RGBA buttons on the Controls tab in the tool. In this case, the buttons in the Settings and the Control tabs are identical.

Apply Mask Inverted (sRender only)

Enabling the Apply Mask Inverted option inverts the complete mask channel for the tool. The mask channel is the combined result of all masks connected to or generated in a node.

Multiply By Mask (sRender only)

Selecting this option causes the RGB values of the masked image to be multiplied by the mask channel's values. This causes all pixels of the image not included in the mask (i.e., set to 0) to become black/transparent.

Motion Blur (sRender only)

- **Motion Blur:** This toggles the rendering of motion blur on the tool. When this control is toggled on, the tool's predicted motion is used to produce the motion blur caused by the virtual camera's shutter. When the control is toggled off, no motion blur is created.
- **Quality:** Quality determines the number of samples used to create the blur. A quality setting of 2 causes Fusion to create two samples to either side of an object's actual motion. Larger values produce smoother results but increase the render time.
- **Shutter Angle:** Shutter Angle controls the angle of the virtual shutter used to produce the motion blur effect. Larger angles create more blur but increase the render times. A value of 360 is the equivalent of having the shutter open for one whole frame exposure. Higher values are possible and can be used to create interesting effects.
- **Center Bias:** Center Bias modifies the position of the center of the motion blur. This allows for the creation of motion trail effects.
- **Sample Spread:** Adjusting this control modifies the weighting given to each sample. This affects the brightness of the samples.

Use GPU (sRender only)

The user GPU menu has three settings. Setting the menu to Disable turns off hard accelerated rendering using the graphics card in your computer. Auto uses a capable GPU if one is available and falls back to software rendering when a capable GPU is not available.

Hide Incoming Connections

Enabling this checkbox can hide connection lines from incoming nodes, making a node tree appear cleaner and easier to read. When enabled, empty fields for each input on a node will be displayed in the Inspector. Dragging a connected node from the node tree into the empty field will hide that incoming connection line as long as the node is not selected in the node tree.. When the node is selected in the node tree the line will reappear.

Comments

The Comments field is used to add notes to a tool. Click in the empty field and type the text. When a note is added to a tool, a small red square appears in the lower left corner of the node when the full tile is displayed or a small text bubble icon appears on the right when nodes are collapsed. To see the note in the Node Editor, hold the mouse pointer over the node to display the tooltip.

Scripts

Three Scripting fields are available on every tool in Fusion from the Settings tab. They each contain edit boxes used to add scripts that process when the tool is rendering. For more information on scripting nodes, see the Fusion scripting documentation.

Chapter 58

Stereo Nodes

This chapter details the Stereo nodes available in Fusion. Stereoscopic nodes are available only in Fusion Studio and DaVinci Resolve Studio.

The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Anaglyph [Ana] 1535
Combiner [Com] 1539
Disparity [Dis] 1540
Disparity To Z [D2Z] 1544
Global Align [GA] 1548
New Eye [NE] 1550
Splitter [Spl] 1553
Stereo Align [SA] 1555
Z To Disparity [Z2D] 1560
The Common Controls 1563

Anaglyph [Ana]



The Anaglyph node

NOTE: The Anaglyph node is available only in Fusion Studio and DaVinci Resolve Studio.

Anaglyph Node Introduction

The Anaglyph node is used to create stereoscopic images by combining separate left eye and right eye images. It is most commonly used at the end of a stereoscopic workflow to display or deliver the final result.

Inputs

The three inputs on the Anaglyph node are the left eye input, right eye input, and effect mask.

Left Eye Input: The orange input is used to connect the 2D image representing the left eye in the stereo comp.

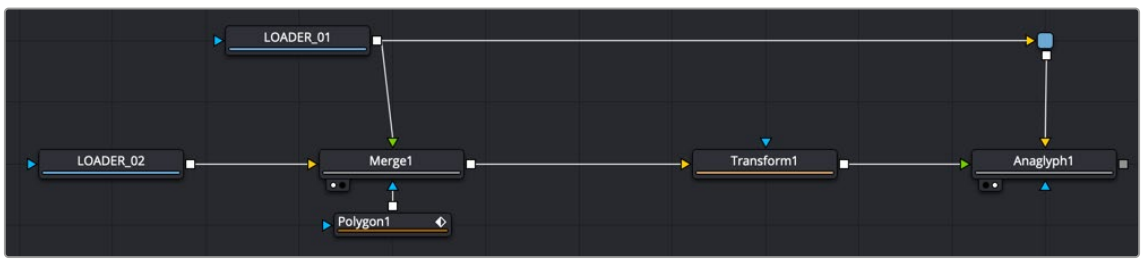
Right Eye Input: The green input is used to connect the 2D image representing the right eye in the stereo comp.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the stereoscopic creation to only those pixels within the mask.

Basic Node Setup

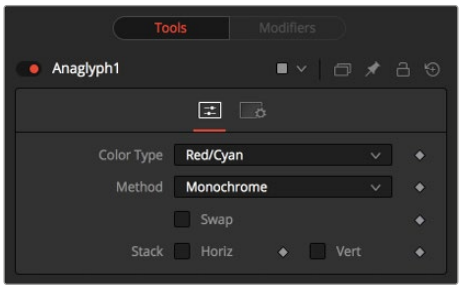
The Anaglyph node is usually placed at the end of a stereoscopic node tree to display the final result.

When using separate images for the left and right eye, the left eye image is connected to the orange input, and the right eye image is connected to the green input of the node. When using either horizontally or vertically stacked images containing both left eye and right eye information, these only connect to the orange input.



An Anaglyph node using separate left and right eye inputs

Inspector



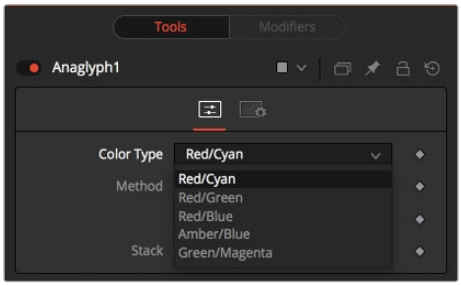
The Anaglyph Controls tab

Controls Tab

Using the parameters in the Controls tab, the separate images are combined to create a stereoscopic output.

Color Type Menu

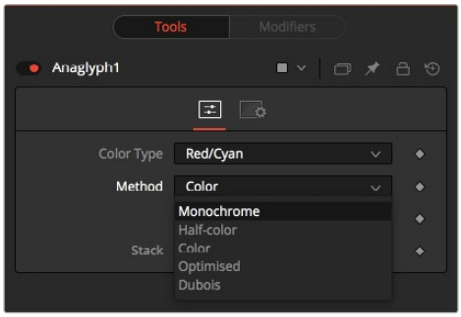
The Color Type menu allows you to choose between different color encodings to fit your preferred display device. To match your stereo glasses, you can choose between Red/Cyan, Red/Green, Red/Blue, Amber/Blue, and Green/Magenta encoding; Red/Cyan is the most commonly used.



The Anaglyph Color Type menu

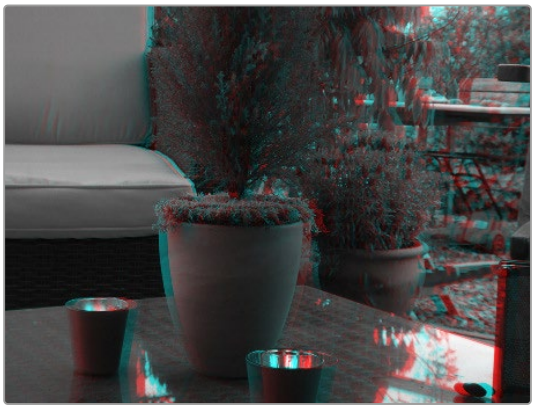
Method

In addition to the color used for encoding the image, you can also choose five different methods from the Method menu: Monochrome, Half-color, Color, Optimized, and Dubois. These methods are described below.



The Anaglyph Method menu

- **Monochrome:** Assuming you are using a Red/Cyan Color Type, the left eye contains the luminance of the left image and is placed in the output of the red channel. The right eye contains the luminance of the right image and is placed in the output green and blue channels.
- **Half-Color:** Assuming you are using a Red/Cyan Color Type, the left eye contains the luminance of the left image and is placed in the output of the red channel. The right eye contains the color channels from the right image that match the glasses' color for that eye.



Monochrome



Half-Color

- **Color:** The left eye contains the color channels from the left image that match the glasses' color for that eye. The right eye contains the color channels from the right image that match the glasses' color for that eye.
- **Optimized:** Used with red/cyan glasses, for example, the resulting brightness of what shows through the left eye is substantially less than the brightness of the right eye. Using typical ITU-R 601 ratios for luminance as a guide, the red eye would give 0.299 brightness, while the cyan eye would give 0.587+0.114=0.701 brightness—over twice as bright. The difference in brightness between the eyes can produce what are referred to as retinal rivalry or binocular rivalry, which can destroy the stereo effect. The Optimized method generates the right eye in the same fashion as the Color method. The left eye also uses the green and blue channels but in combination with increased brightness that reduces retinal rivalry. Since it uses the same two channels from each of the source images, it doesn't reproduce the remaining one. For example, 1.05× the green and 0.45× the blue channels of the left image is placed in the red output channel, and the green and blue channels of the right image are placed in the output green and blue channels. Red from both the left and right images is not used.



Color



Optimized

- **Dubois:** Images with fairly saturated colors can produce retinal rivalry with the Half-color, Color, and Optimized methods because the color is visible in only one eye. For example, with red/cyan glasses, a saturated green object looks black in the red eye, and green in the cyan eye. The Dubois method uses the spectral characteristics of (specifically) red/cyan glasses and CRT (Trinitron) phosphors to produce a better anaglyph and in the end, tends to reduce retinal rivalry caused by such color differences in each eye. It also tends to reduce ghosting produced when one eye 'leaks' into the other eye. The particular calculated matrix in Fusion is designed for red/cyan glasses and isn't available for other glasses types. Since it is also derived from CRT color primaries, it may not give the best results with a common LCD (though it'll still likely produce less retinal rivalry and ghosting than the other methods).



Dubois

Swap Eyes

Allows you to swap the left and right eye inputs easily.

Horiz Stack

Takes an image that contains both left and right eye information stacked horizontally. These images are often referred to as "crosseyed" or "straight stereo" images. You only need to connect that one image to the orange input of the node. It then creates an image half the width of the original input, using the left half of the original image for the left eye and the right half of the original image for the right eye. Color encoding takes place using the specified color type and method.

Vert Stack

Takes an image that contains both left and right eye information stacked vertically. You only need to connect that one image to the orange input of the node. It then creates an image half the height of the original input, using the bottom half of the original image for the left eye and the top half of the original image for the right eye. Color encoding takes place using the specified color type and method.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Stereo nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Combiner [Com]



The Combiner node

NOTE: The Combiner node is available only in Fusion Studio and DaVinci Resolve Studio.

Combiner Node Introduction

The Combiner node takes two stereoscopic inputs and creates so-called “stacked images” with the left and right eye, either side by side or on top of each other. Stereoscopic nodes are available only in Fusion Studio and DaVinci Resolve Studio.

Inputs

The two inputs on the Combiner node are used to connect the two images that get combined in a stacked stereo image.

Image 1 Input: The orange input is used to connect the 2D image representing the left eye in the stereo comp.

Image 2 Input: The green input is used to connect the 2D image representing the right eye in the stereo comp.

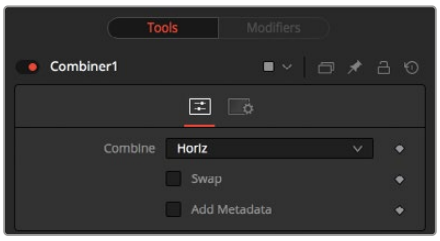
Basic Node Setup

Below, a left eye image and right eye image are connected to the Combiner node to create a single stacked stereo image. It can be more efficient to render out the stacked stereo images as EXR files than to generate the disparity on-the-fly.



Left and right eye images are connected into a Combiner node to generate a stacked stereo image.

Inspector



The Combiner Controls tab

Controls Tab

To stack the images, the left eye image is connected to the orange input, and the right eye image is connected to the green input of the node.

Combine

The Combine menu provides three options for how the two images are made into a stacked stereo image.

- **None:** No operation will take place. The output image is identical to the left eye input.
- **Horiz:** Both images will be stacked horizontally, or side-by-side, with the image connected to the left eye input on the left. This will result in an output image double the width of the input image.
- **Vert:** Both images will be stacked vertically, or on top of each other, with the image connected to the left eye input on the bottom. This will result in an output image double the height of the input image.
- **Layers:** This exposes layer name fields, allowing you to rename the layers.

Swap Eyes

Allows you to easily swap the left and right eye input.

Add Metadata

Metadata is carried along with the images and can be added to the existing metadata using this checkbox. To view Metadata, use the viewer's SubView menu set to Metadata.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Stereo nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Disparity [Dis]



The Disparity node

NOTE: The Disparity node is available only in Fusion Studio and DaVinci Resolve Studio.

Disparity Node Introduction

Disparity generates the left/right shift between the frames in a stereo pair. It also generates the vertical disparity between the left/right images, which is usually a lot smaller than the horizontal disparity and ideally should be 0 to minimize viewing discomfort. When viewing the output of the Disparity node in the views, the human eye can distinguish quality/detail in the Disparity map better by looking at either the grayscale X disparity or Y disparity rather than looking at the combined XY disparity as a Red/Green color image.

The generated disparity is stored in the output image's Disparity aux channel, where the left image contains the left > right disparity, and the right image contains the right > left disparity. Because disparity works based on matching regions in the left eye to regions in the right eye by comparing colors and gradients of colors, colors in the two eyes must be as similar as possible. Thus, it is a good idea to color correct ahead of time. It is also a good idea to crop away any black borders around the frames, as this confuses the disparity tracking (and also causes problems if you are using the Color Corrector's histogram match ability to do the color matching).

In Stack mode, left and right outputs deliver the same image. If the left and right images have a global vertical offset larger than a few pixels, it can help the disparity tracking algorithm if you vertically align features in the left/right eyes ahead of time using a Transform node. Small details tend to get lost in the tracking process when you have a large vertical offset between left/right eyes.

Consider using a SmoothMotion node to smooth your disparity channel. This can help reduce time-dependent flickering when warping an eye. Also, think about whether you want to remove lens distortion before computing disparity. If you do not, your Disparity map becomes a combined Disparity and Lens Distortion map. This can have advantages and disadvantages.

One disadvantage is that if you then do a vertical alignment, you are also removing lens distortion effects. When trying to reduce the computation time, start first with adjusting the Proxy and Number of Iterations sliders.

The Disparity node does not support RoI or DoD.

Inputs

The two inputs on the Disparity node are used to connect the left and right images.

Left Input: The orange input is used to connect either the left eye image or the stacked image.

Right Input: The green input is used to connect the right eye image. This input is available only when the Stack Mode menu is set to Separate.

Outputs

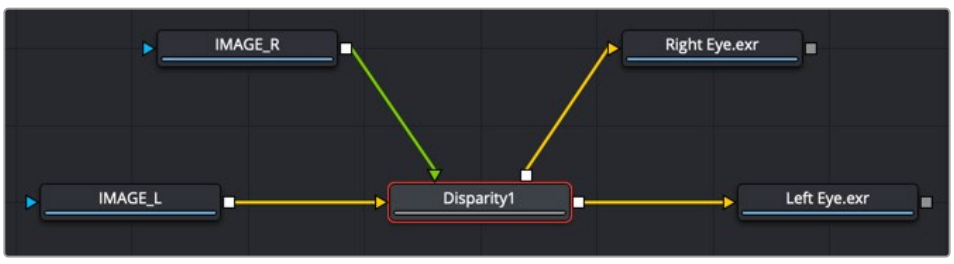
Unlike most nodes in Fusion, Disparity has two outputs for the left and right eye.

Left Output: This holds the left eye image with a new disparity channel, or a Stacked Mode image with a new disparity channel.

Right Output: This holds the right eye image with a new disparity channel. This output is visible only if Stack Mode is set to Separate.

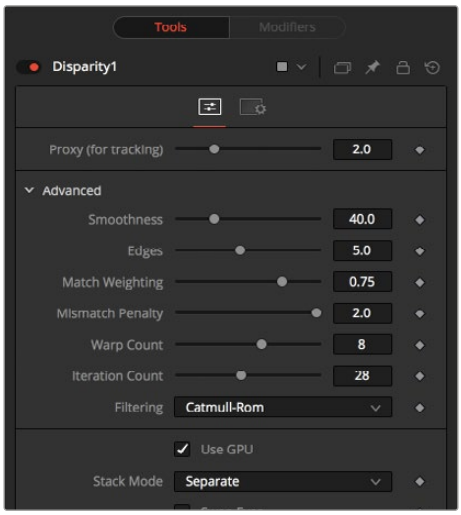
Basic Node Setup

Below, a left eye image and right eye image are connected to the Disparity node. The Disparity node then outputs each eye to Saver nodes. It can be more efficient to render out the stereo images as EXR files than to generate the disparity on-the-fly.



Left and right eye images are connected into a Disparity node to generate and render out a stereo image

Inspector



The Disparity Controls tab

Proxy (for Tracking)

The input images are resized down by the proxy scale, tracked to produce the disparity, and then the resulting disparities are scaled back up. This option is purely to speed up the calculation of the disparity, which can be slow. The computational time is roughly proportional to the number of pixels in the image. This means a proxy scale of 2 gives a 4x speedup, and a proxy scale of 3 gives a 9x speedup. In general, 1:1 proxy will give the most detailed flow, but keep in mind that this is highly dependent on the amount of noise and film grain. If noise is present in large quantities, it can completely obliterate any gains moving from 2:1 to 1:1 proxy. In some situations, it can even make things worse. You can think of the Proxy setting as acting as a simplistic low-pass filter for removing noise/grain.

Advanced

The Advanced settings section has parameter controls to tune the Disparity map calculations. The default settings have been chosen to be the best default values from experimentation with many different shots and should serve as a good standard. In most cases, tweaking of the Advanced settings is not needed.

Smoothness

This controls the smoothness of the disparity. Higher values help deal with noise, while lower values bring out more detail.

Edges

This slider is another control for smoothness but applies it based on the color channel. It tends to have the effect of determining how edges in the disparity follow edges in the color images. When it is set to a lower value, the disparity becomes smoother and tends to overshoot edges. When it is set to a higher value, edges in the disparity align more tightly with the edges in the color images, and details from the color channels start to slip into the disparity, which is not usually desirable.

As a rough guideline, if you are using the disparity to produce a Z channel for post effects like depth of field, experiment with higher values, but if you are using the disparity to do interpolation, you might want to keep the values lower.

In general, if the Edges slider is set is too high, there can be problems with streaked out edges when the disparity is used for interpolation.

Match Weight

This controls how matching is done between neighboring pixels in the left image and neighboring pixels in the right image. When a lower value is used, large structural color features are matched. When higher values are used, small sharp variations in the color are matched. Typically, a good value for this slider is in the [0.7, 0.9] range. Setting this option higher tends to improve the matching results in the presence of differences due to smoothly varying shadows or local lighting variations between the left and right images. You should still color match the initial images so they are as similar as possible; this option tends to help with local variations (e.g., lighting differences due to light passing through a mirror rig).

Mismatch Penalty

This controls how the penalty for mismatched regions grows as they become more dissimilar. The slider provides a choice between a balance of Quadratic (lower values) and Linear (higher values) penalties. Lower value Quadratic settings strongly penalize large dissimilarities, while higher value Linear settings are more robust to dissimilar matches. Moving this slider toward lower tends to give a disparity with more small random variations in it, while higher values produce smoother, more visually pleasing results.

Warp Count

Turning down the Warp Count makes the disparity computations faster. In particular, the computational time depends linearly upon this option. To understand what this option does, you need to understand that the Disparity algorithm progressively warps the left image until it matches with the right image. After some point, convergence is reached, and additional warps are just a waste of computational time. The default value in Fusion is set high enough that convergence should always be reached. You can tweak this value to speed up the computations, but it is good to watch how the disparity is degrading in quality at the same time.

Iteration Count

Turning down the Iteration Count makes the disparity computations faster. In particular, the computational time depends linearly upon this option. Just like adjusting Warp Count, at some point adjusting this option higher will yield diminishing returns and will not produce significantly better results. By default, this value is set to something that should converge for all possible shots and can be tweaked lower fairly often without reducing the disparity's quality.

Filtering

This menu determines the filtering operations used during flow generation. Catmull-Rom filtering will produce better results, but at the same time, it increases the computation time steeply.

Stack Mode

This menu determines how the input images are stacked.

When set to Separate, the Right Input and Output will appear, and separate left and right images must be connected.

Swap Eyes

Enabling this checkbox causes the left and right images to swap.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Stereo nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Disparity To Z [D2Z]



The DisparityToZ node

NOTE: The Disparity to Z node is available only in Fusion Studio and DaVinci Resolve Studio.

Disparity to Z Node Introduction

Disparity To Z takes a 3D camera and an image containing a disparity channel as inputs, and outputs the same image but with a newly computed Z channel.

Optionally, this node can output Z into the RGB channels. Ideally, either a stereo Camera 3D or a tracked stereo camera is connected into Disparity To Z. However, if no camera is connected, the node provides artistic controls for determining a Z channel. The depth created by this node can be used for post effects like fogging or depth of field (DoF).

The Z values produced become more incorrect the larger (negative) they get. The reason is that disparity approaches a constant value as Z approaches -infinity. So Z = -1000 and Z = -10000 and Z = -100000 may map to D=142. 4563 and D=142. 4712 and D=142. 4713. As you can see, there is only 0.0001 in D to distinguish between 10,000 and 100,000 in Z. The maps produced by disparity are not accurate enough to make distinctions like this.

Inputs

The three inputs on the Disparity To Z node are used to connect the left and right images and a camera node.

Left Input: The orange input is used to connect either the left eye image or the stack image.

Right Input: The green input is used to connect the right eye image. This input is available only when the Stack Mode menu is set to Separate.

Stereo Camera: The magenta input is used to connect a stereo camera node.

Outputs

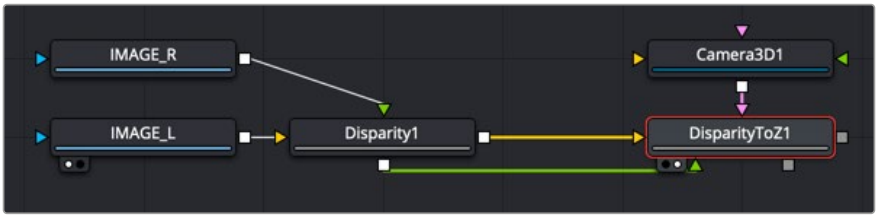
Unlike most nodes in Fusion, Disparity To Z has two outputs for the left and right eye.

Left Output: This holds the left eye image with a new Z channel, or a Stacked Mode image with a new disparity channel.

Right Output: This holds the right eye image with a new Z channel. This output is visible only if Stack Mode is set to Separate.

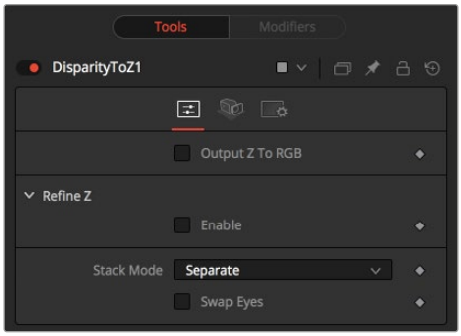
Basic Node Setup

Disparity To Z takes a 3D camera and stereo images containing a disparity channel as inputs. The output is an image with a newly computed Z channel.



A Disparity to Z node creates an image with a Z channel.

Inspector



The Disparity To Z Controls tab

Controls Tab

In addition to outputting Z values in the Z channel, this tab promotes the color channels to float32 and outputs the Z values into the color channels as {z, z, z, 1}. This option is useful to get a quick look at the Z channel.

NOTE: Z values are negative, becoming more negative the further you are from the camera. The viewers only show 0.0 to 1.0 color, so to visualize other data it has to be converted via a normalization method to fit in a display 0-1 range. To do this, right-click in the viewer and choose Options > Show Full Color Range.

Output Z to RGB

Rather than keeping the Z values within the associated aux channel only, they will be copied into the RGB channels for further modification with any of Fusion's nodes.

Refine Z

The Enable checkbox refines the depth map based upon the RGB channels. The refinement causes edges in the flow to align more closely with edges in the color channels. The downside is that unwanted details in the color channels start to show up in the flow. You may want to experiment with using this option to soften out harsh edges for Z-channel post effects like depth of field or fogging.

HiQ Only

Activating this checkbox causes the Refine Z option to process only when rendering is set to High Quality. You can ensure High Quality is enabled by right-clicking to the left or right of the transport controls in the main toolbar.

Strength

Increasing this slider does two things. It smooths out the depth in constant color regions and moves edges in the Z channel to correlate with edges in the RGB channels.

Increasing the refinement has the undesirable effect of causing texture in the color channel to show up in the Z channel. You will want to find a balance between the two.

Radius

This is the radius of the smoothing algorithm.

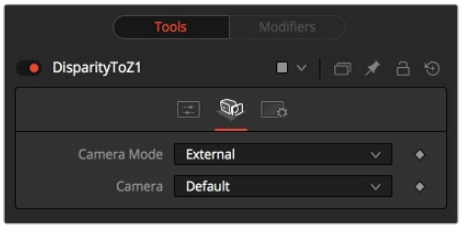
Stack Mode

This menu determines how the input images are stacked.

When set to Separate, the Right Input and Output will appear, and separate left and right images must be connected.

Swap Eyes

Enabling this checkbox causes left and right images to be swapped.



The Disparity To Z Camera tab

Camera Tab

If you need correct real-world Z values because you are trying to match some effect to an existing scene, you should use the External Camera options to get precise Z values back. If any Z-buffer will suffice and you are not that particular about the exact details of how it is offset and scaled, or if there is no camera available, the Artistic option might be helpful.

- **External Mode:** An input is available on the node to connect an existing stereo Camera 3D. This can either be a single stereo Camera 3D (i.e., its eye separation is set to non-zero), or a pair of (tracked) Camera 3Ds connected via the Camera 3D > Stereo > Right Camera input.
- **Artistic Mode:** If you do not have a camera, you can adjust these controls to produce an “artistic” Z channel whose values will not be physically correct but will still be useful. To reconstruct the Disparity > Z Curve, pick (D, Z) values for a point in the foreground and a point in the background.

TIP: If artistic mode is a little too “artistic” for you and you want more physically-based parameters to adjust (e.g., convergence and eye separation), create a dummy Camera 3D, connect it into the Disparity To Z > Camera input, and then fiddle with the Camera 3D’s controls.

Foreground Disparity (Pick from Left Eye)

When the camera Mode is set to Artistic, a Foreground Disparity slider is available. This is the disparity for the closest foreground object. It will get mapped to the depth value specified by the Foreground Depth control. Any objects with disparity outside of the range [ForegroundDisparity, BackgroundDisparity] will have their disparity values clipped to this range leading to flat areas in the Z channel, so make sure that you pick values that enclose the actual disparity range.

Background Disparity (Pick from Left Eye)

When the camera Mode is set to Artistic, a Background Disparity is available. This is the disparity for the furthest background object. It will get mapped to the depth value specified by the Background Depth control. One way to think of this input is as the upper limit to disparity values for objects at -infinity. This value should be for the left eye. The corresponding value in the right eye will be the same in magnitude but negative.

Foreground Depth

This is the depth to which Foreground Disparity will be mapped. Think of this as the depth of the nearest object. Note that values here are positive depth.

Background Depth

This is the depth to which Background Disparity will be mapped. Think of this as the depth of the most distant object.

Falloff

Falloff controls the shape of the depth curve between the requested foreground and background depths. When set to Hyperbolic, the disparity-depth curve behaves roughly like depth = constant/disparity. When set to Linear, the curve behaves like depth = constant * disparity. Hyperbolic tends to emphasize Z features in the foreground, while linear gives foreground/background features in the Z channel equal weighting.

Unless there's a specific reason, choose Hyperbolic, as it is more physically accurate, while Linear does not correspond to nature and is purely for artistic effect.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Stereo nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Global Align [GA]



The GlobalAlign node

NOTE: The Global Align node is available only in Fusion Studio and DaVinci Resolve Studio.

Global Align Node Introduction

As opposed to Stereo Align, this node does not utilize optical flow at all. It's meant as a fast and convenient way to do simple stereo alignment for both X and Y as well as rotation.

Global Align comes in handy at the beginning of the node chain to visually correct major differences between the left and right eye before calculating Disparity.

Manual correction of large discrepancies between left and right, as well as applying an initial color matching, helps Disparity generate more accurate results.

Inputs

The two inputs on the Global Align node are used to connect the left and right images.

Left Input: The orange input is used to connect either the left eye image or the stack image.

Right Input: The green input is used to connect the right eye image. This input is available only when the Stack Mode menu is set to Separate.

Outputs

Unlike most nodes in Fusion, Global Align has two outputs for the left and right eye.

Left Output: This outputs the newly aligned left eye image.

Right Output: This outputs the newly aligned right eye image.

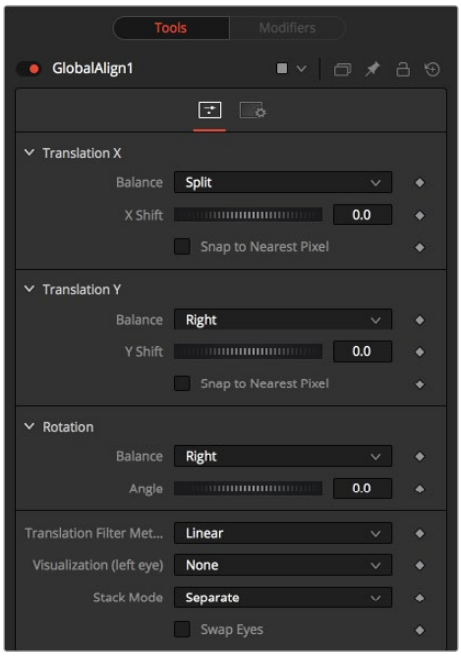
Basic Node Setup

Global Align is typically placed at the beginning of the node tree. Below it is inserted between the left and right eye images and the Disparity node to visually correct major differences.



A Global Align node used to manually correct left and right eye discrepancies

Inspector



The Global Align Controls tab

Controls Tab

The Controls tab includes translation and rotation controls to align the stereo images manually.

Translation X and Y

- **Balance:** Determines how the global offset is applied to the stereo footage.
- **None:** No translation is applied.
- **Left Only:** The left eye is shifted, while the right eye remains unaltered.
- **Right Only:** The right eye is shifted, while the left eye remains unaltered.
- **Split Both:** Left and right eyes are shifted in opposite directions.

Snap to Nearest Pixel

While adjusting the X or Y shift dial, this option ensures that the image is shifted in full pixel amounts only to maintain optimum quality. This avoids sub-pixel rendering of the image, which could result in subtle blurring.

Rotation

- **Balance:** Determines how the global rotation is applied to the stereo footage.
- **None:** No rotation is applied.
- **Left Only:** The left eye is rotated, while the right eye remains unaltered.
- **Right Only:** The right eye is rotated, while the left eye remains unaltered.
- **Split Both:** Left and right eyes are rotated in opposite directions.

Angle

This dial adjusts the angle of the rotation. Keep in mind that the result depends on the Balance settings. If only rotating one eye by, for example, 10 degrees, a full 10-degree rotation will be applied to that eye.

When applying rotation in Split mode, one eye will receive a -5 degree and the other eye a +5 degree rotation.

Translation Filter Method

This menu chooses the filter method that delivers the best results depending on the content of your footage.

Visualization

This control allows for different color encodings of the left and right eye to conveniently examine the results of the above controls without needing to add an extra Anaglyph or Combiner node.

Set this to None for final output.

Stack Mode

Determines how the input images are stacked.

When set to Separate, the right input and output will appear, and separate left and right images must be connected.

Swap Eyes

With Stacked Mode, image stereo pairs' left and right images can be swapped.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Stereo nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

New Eye [NE]



The New Eye node

NOTE: The New Eye node is available only in Fusion Studio and DaVinci Resolve Studio.

New Eye Node Introduction

The New Eye node constructs a new image by interpolating between two existing stereo images using the embedded disparity channels. This node can also be used to replace one view with a warped version of the other. In Stack Mode, L and R outputs will output the same image.

You can map the left eye onto the right eye and replace it. This can be helpful when removing errors from certain areas of the frame.

New Eye does not interpolate the aux channels but instead destroys them. In particular, the disparity channels are consumed/destroyed. Add another Disparity node after the New Eye if you want to generate Disparity for the realigned footage.

Inputs

The two inputs on the New Eye node are used to connect the left and right images.

Left Input: The orange input is used to connect either the left eye image or the stack image.

Right Input: The green input is used to connect the right eye image. This input is available only when the Stack Mode menu is set to Separate.

Outputs

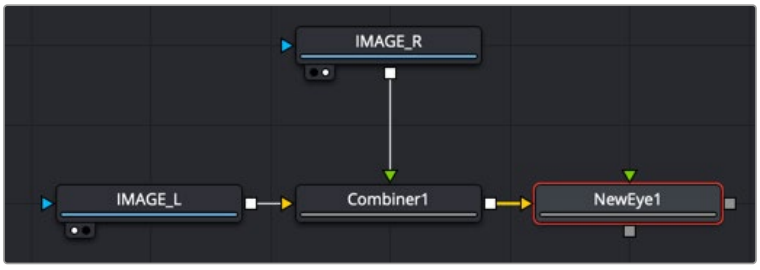
Unlike most nodes in Fusion, New Eye has two outputs for the left and right eye.

Left Output: This outputs the left eye image with a new disparity channel, or a Stacked Mode image with a new disparity channel.

Right Output: This outputs the right eye image with a new disparity channel. This output is visible only if Stack Mode is set to Separate.

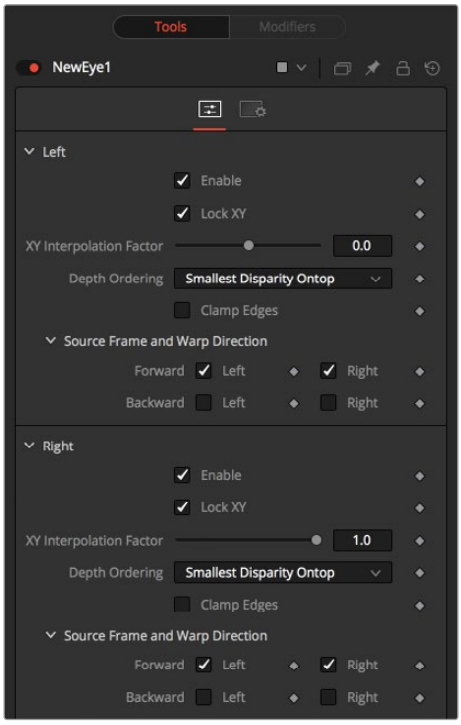
Basic Node Setup

The New Eye node in the example below constructs a new image by interpolating between the two stereo images using the embedded disparity channels.



A New Eye node creates a new stereo image using embedded disparity

Inspector



The New Eye Controls tab

Controls Tab

The Controls tab is divided into identical parameters for the left eye and right eye. The parameters are used to select which eye to recreate and the methods used for the interpolation.

Enable

The Enable checkbox allows you to activate the left or right eye independently. The New Eye will replace enabled eye with an interpolated eye. For example, if the left eye is your “master” eye and you are recreating the right eye, you would disable the left eye and enable the right eye.

Lock XY

Locks the X and Y interpolation parameters. When they are unlocked, you can provide separate interpolation factors for using the X and Y disparity. For example, if you are working with the right eye and you have the X Interpolation slider set to 1.0 and the Y Interpolation slider set to -1.0, you will be effectively interpolating the left eye onto the right eye but vertically aligned to the left eye.

XY Interpolation Factor

Interpolation determines where the interpolated frame is positioned, relative to the two source frames: A slider position of -1.0 outputs the frame Left and a slider position of 1.0 outputs the frame Right. A slider position of 0.0 outputs a result that is halfway between Left and Right.

Depth Ordering

The Depth Ordering is used to determine which parts of the image should be rendered on top. When warping images, there is often overlap. When the image overlaps itself, there are two options for which values should be drawn on top.

- **Largest Disparity On Top:** The larger disparity values will be drawn on top in the overlapping image sections.
- **Smallest Disparity On Top:** The smaller disparity values will be drawn on top in the overlapping image sections.

Clamp Edges

Under certain circumstances, this option can remove the transparent gaps that may appear on the edges of interpolated frames. Clamp Edges will cause a stretching artifact near the edges of the frame that is especially visible with objects moving through it or when the camera is moving.

Because of these artifacts, it is a good idea to use Clamp Edges only to correct small gaps around the edges of an interpolated frame.

Softness

Helps to reduce the stretchy artifacts that might be introduced by Clamp Edges.

If you have more than one of the Source Frame and Warp Direction checkboxes turned on, this can lead to doubling up of the stretching effect near the edges. In this case, you’ll want to keep the softness rather small at around 0.01. If you have only one checkbox enabled, you can use a larger softness at around 0.03.

Source Frame and Warp Direction

The output of this node is generated by combining up to four different warps. You can choose to use either the color values from the left or right frame in combination with the Forward (left > right) Disparity or the Backward (right > left) Disparity. Sometimes you will want to replace an existing eye. For example, if you want to regenerate the right eye, you would only use left eye warps.

It's good to experiment with various options to see which gives the best effect. Using both the left and right eyes can help fill in gaps on the left/right side of images. Using both the Forward/Backward Disparity can give a doubling-up effect in places where the disparities disagree with each other.

- **Left Forward:** Takes the Left frame and uses the Forward Disparity to interpolate the new frame.
- **Right Forward:** Takes the Right frame and uses the Forward Disparity to interpolate the new frame.
- **Left Backward:** Takes the Left frame and uses the Back Disparity to interpolate the new frame.
- **Right Backward:** Takes the Right frame and uses the Back Disparity to interpolate the new frame.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Stereo nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Splitter [Spl]



The Splitter node

NOTE: The Splitter node is available only in Fusion Studio and DaVinci Resolve Studio.

Splitter Node Introduction

The Splitter takes a stacked input image—for example, created with the Combiner—and provides two output images: a left eye and a right eye.

Inputs

The two inputs on the Splitter node are used to connect the left and right images.

Left Input: The orange input is used to connect a stacked stereo image.

Outputs

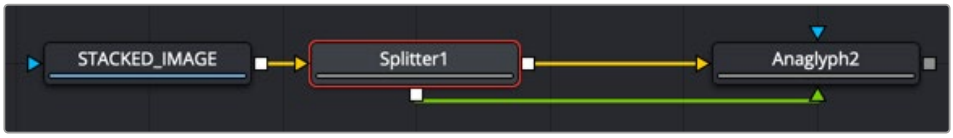
Unlike most nodes in Fusion, the Splitter node has two outputs for the left and right eye.

Left Output: This outputs the left eye image.

Right Output: This outputs the right eye image.

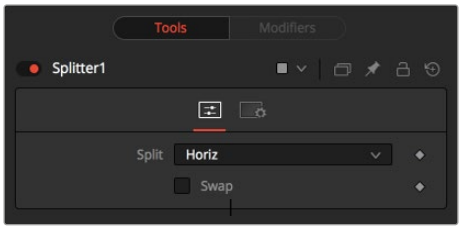
Basic Node Setup

Below, a stacked stereo image is connected to the input on a Splitter where a left and right eye is output.



A Splitter node creates a left and right image from a stacked stereo image

Inspector



The Splitter Controls tab

Controls Tab

The Controls tab is used to define the type of stacked image connected to the node's input.

Split

The Split menu contains three options for determining the orientation of the stacked input image.

- **None:** No operation takes place. The output image on both outputs is identical to the input image.
- **Horiz:** The node expects a horizontally stacked image. This will result in two output images, each being half the width of the input image.
- **Vert:** The node expects a vertically stacked image. This will result in two output images, each being half the height of the input image.

Swap Eyes

Allows you to easily swap the left and right eye outputs.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Stereo nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Stereo Align [SA]



The StereoAlign node

NOTE: The Stereo Align node is available only in Fusion Studio and DaVinci Resolve Studio.

Stereo Align Node Introduction

This extremely versatile node for fixing Stereo issues can be used for performing any of the following actions or combinations thereof:

- Vertical alignment of one eye to the other
- Changing the convergence
- Changing the eye separation

By combining these operations in one node, you can execute them using only a single image resampling. In essence, this node can be thought of as applying scales and translation to the disparities and then using the modified disparities to interpolate between the views.

NOTE: Changing the eye separation can cause holes to appear, and it may not be possible to fill them since the information needed may not be in either image. Even if the information is there, the disparity may have mismatched the holes. You may need to fill the holes manually. This node modifies only the RGBA channels.

TIP: Stereo Align does not interpolate the aux channels but instead destroys them. In particular, the disparity channels are consumed/destroyed. Add another Disparity node after the StereoAlign if you want to generate Disparity for the realigned footage.

Inputs

The two inputs on the Stereo Align node are used to connect the left and right images.

Left Input: The orange input is used to connect either the left eye image or the stack image.

Right Input: The green input is used to connect the right eye image. This input is available only when the Stack Mode menu is set to Separate.

Outputs

Unlike most nodes in Fusion, Stereo Align has two outputs for the left and right eye.

Left Output: This outputs the left eye image with a new disparity channel, or a Stacked Mode image with a new disparity channel.

Right Output: This outputs the right eye image with a new disparity channel. This output is visible only if Stack Mode is set to Separate.

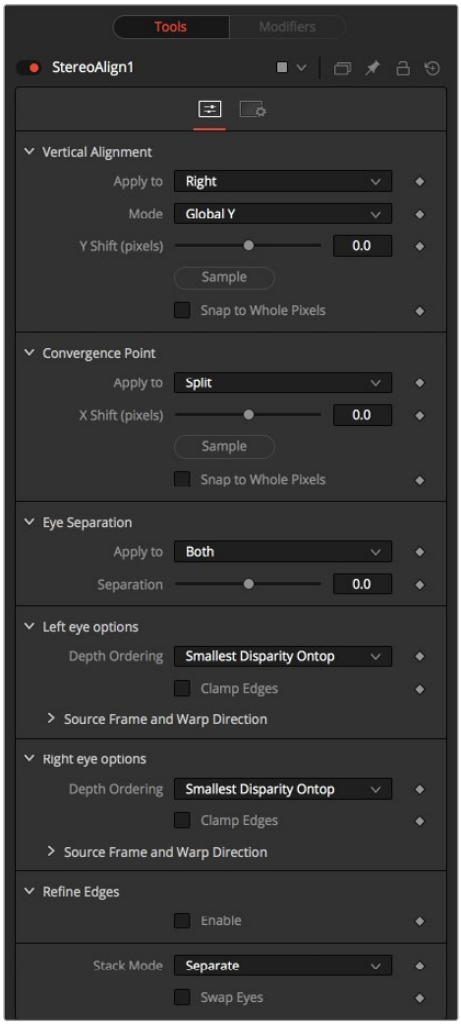
Basic Node Setup

Below, the Stereo Align receives the left and right eye images with disparity. Once the adjustments are made, another Disparity node is added after it to generate disparity for the realigned footage.



A Stereo Align node destroys the disparity channel, so another Disparity node is placed after it

Inspector



The Stereo Align Controls tab

Controls Tab

Vertical Alignment

This option determines how the vertical alignment is split between two eyes. Usually, the left eye is declared inviolate, and the right eye is aligned to it to avoid resampling artifacts.

When doing per pixel vertical alignment, it may be helpful to roughly pre-align the images by a global Y-shift before disparity computation because the disparity generation algorithm can have problems resolving small objects that move large distances.

Also, be aware that you must be careful about lens distortion because even if two cameras are perfectly vertically aligned, they will still have vertical disparities due to lens distortion. As a best practice, remove the lens distortion before computing the disparity. When a vertical alignment of the right eye is done, you have essentially removed the Y-component of the lens distortion in the right eye, and it will look wrong later when you try to distort it again.

Apply to

- **Right:** Only the right eye is adjusted.
- **Left:** Only the left eye is adjusted.
- **Both:** The vertical alignment is split evenly between the left and right eyes.

Mode

- **Global:** The eyes are simply translated up or down by the Y-shift to match up.
- **Per Pixel:** The eyes are warped pixel-by-pixel using the disparity to vertically align.

> Keep in mind that this can introduce sampling artifacts and edge artifacts.

Y-shift

Y-shift is available only when the Mode menu is set to Global. You can either adjust the Y-shift manually to get a match or drag the Sample button into the viewer, which picks from the disparity channel of the left eye. Also remember, if you use this node to modify disparity, you can't use the Sample button while viewing the node's output.

Snap to Whole Pixels

You can snap the global shift to whole pixels by enabling this option. When enabled, there is no resampling of the image, but rather a simple shift is done so there will be no softening or image degradation.

Convergence Point

The Convergence Point section is used as a global X-translation of L/R images.

Apply to

This menu determines which eyes are affected by convergence. You can choose to apply the convergence to the left eye, right eye, or split between the two. In most cases, this will be set to Split. If you set the eyes to Split, then the convergence will be shared 50-50 between both eyes. Sharing the convergence between both eyes means you get half the shift in each eye, which in turn means smaller holes and artifacts that need to be fixed later. The tradeoff is that you've resampled both eyes rather than keeping one eye as a pure reference master.

X-shift

You can either use the slider to adjust the X-shift manually to get a match or use the Sample button to pick from the disparity channels for easy point-to-feature alignment.

Snap

You can snap the global shift to whole pixels using this option. In this mode, there is no resampling of the image, but rather a simple shift is done so there will be no softening or image degradation.

Eye Separation

Eye Separation changes the distance between the left/right eyes, causing objects in the left/right eyes to converge/diverge further depending on their distance from the camera.

This has the same effect as the Eye Separation option in the Camera 3D node.

Separation

This is a scale factor for eye separation.

- When set to 0.0, this leaves the eyes unchanged.
- Setting it to 0.1 increases the shifts of all objects in the scene by a factor of 10% in each eye.
- Setting it to 0.1 will scale the shifts of all objects 10% smaller.

Unlike the Split option for vertical alignment, which splits the alignment effect 50-50 between both eyes, the Both option will apply 100-100 eye separation to both eyes. If you are changing eye separation, it can be a good idea to enable per-pixel vertical alignment, or the results of interpolating from both frames can double up.

Left/Right Eye Options

The left and right eye options contain depth ordering and warp direction controls independently for the left and right eye.

Depth Ordering

The Depth Ordering is used to determine which parts of the image should be rendered on top. When warping images, there is often overlap. When the image overlaps itself, there are two options for which values should be drawn on top.

- **Largest Disparity On Top:** The larger disparity values will be drawn on top in the overlapping image sections.
- **Smallest Disparity On Top:** The smaller disparity values will be drawn on top in the overlapping image sections.

Clamp Edges

Under certain circumstances, this option can remove the transparent gaps that may appear on the edges of interpolated frames. Clamp Edges will cause a stretching artifact near the edges of the frame that is especially visible with objects moving through it or when the camera is moving.

Because of these artifacts, it is a good idea to use Clamp Edges only to correct small gaps around the edges of an interpolated frame.

Edge Softness

Helps to reduce the stretchy artifacts that might be introduced by Clamp Edges.

If you have more than one of the Source Frame and Warp Direction checkboxes turned on, this can lead to doubling up of the stretching effect near the edges. In this case, you'll want to keep the softness rather small at around 0.01. If you have only one checkbox enabled, you can use a larger softness at around 0.03.

Source Frame and Warp Direction

The output of this node is generated by combining up to four different warps. You can choose to use either the color values from the left or right frame in combination with the Forward (left > right)

Disparity or the Backward (right > left) Disparity. Sometimes you will want to replace an existing eye. For example, if you want to regenerate the right eye, you would use only left eye warps.

It's good to experiment with various options to see which gives the best effect. Using both the left and right eyes can help fill in gaps on the left/right side of images. Using both the Forward/Backward Disparity can give a doubling-up effect in places where the disparities disagree with each other.

- **Left Forward:** Takes the Left frame and uses the Forward Disparity to interpolate the new frame.
- **Right Forward:** Takes the Right frame and uses the Forward Disparity to interpolate the new frame.
- **Left Backward:** Takes the Left frame and uses the Back Disparity to interpolate the new frame.
- **Right Backward:** Takes the Right frame and uses the Back Disparity to interpolate the new frame.

Stack Mode

In Stack Mode, L and R outputs will output the same image.

If High Quality is off, the interpolations are done using nearest-neighbor sampling, leading to a more "noisy" result. To ensure High Quality is enabled, right-click under the viewers, near the transport controls, and choose High Quality from the pop-up menu.

Swap Eyes

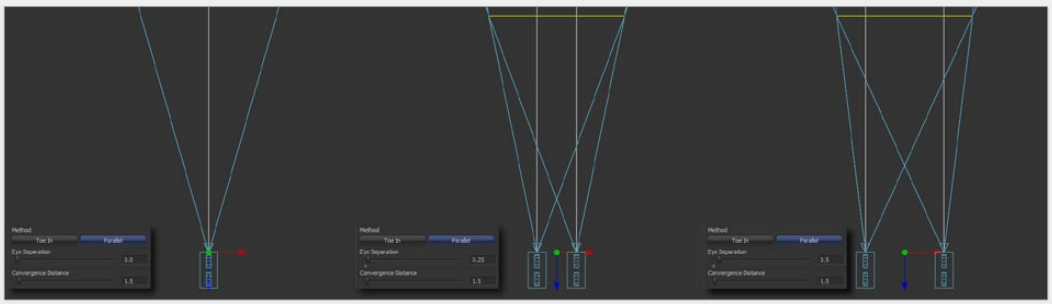
Allows you to easily swap the left and right eye outputs.

Common Controls

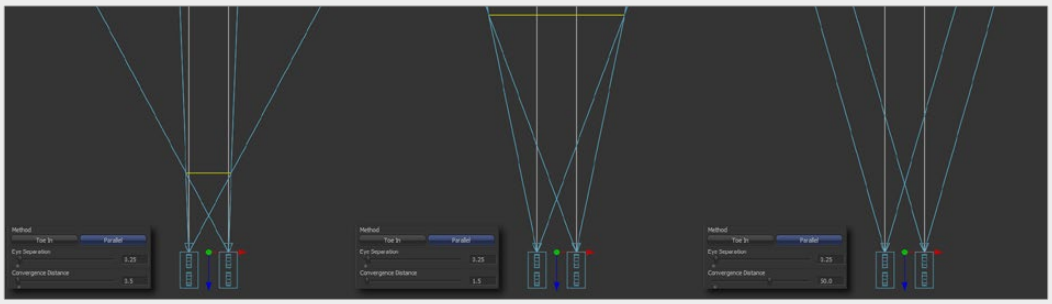
Settings Tab

The Settings tab in the Inspector is also duplicated in other Stereo nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Example



Different settings for Eye Separation...



...and example settings for Convergence

Z To Disparity [Z2D]



The Z To Disparity node

NOTE: The Z To Disparity node is available only in Fusion Studio and DaVinci Resolve Studio.

Z To Disparity Node Introduction

Z To Disparity takes a stereo camera and an image containing a Z channel and outputs the same image but with disparity channels in it. This is useful for constructing a Disparity map from CG renders, which will be more accurate than the Disparity map created from the Disparity node.

Inputs

The three inputs on the Z To Disparity node are used to connect the left and right images and a camera node.

Left Input: The orange input is used to connect either the left eye image or the stack image.

Right Input: The green input is used to connect the right eye image. This input is available only when the Stack Mode menu is set to Separate.

Stereo Camera: The magenta input is used to connect a stereo perspective camera, which may be either a Camera 3D with eye separation, or a tracked L/R Camera 3D.

Outputs

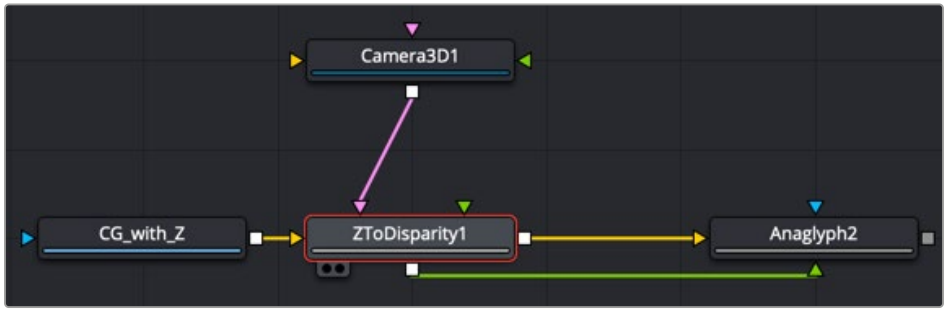
Unlike most nodes in Fusion, Z To Disparity has two outputs for the left and right eye.

Left Output: This outputs the left eye image containing a new disparity channel, or a Stacked Mode image with a new disparity channel.

Right Output: This outputs the right eye image with a new disparity channel. This output is visible only if Stack Mode is set to Separate.

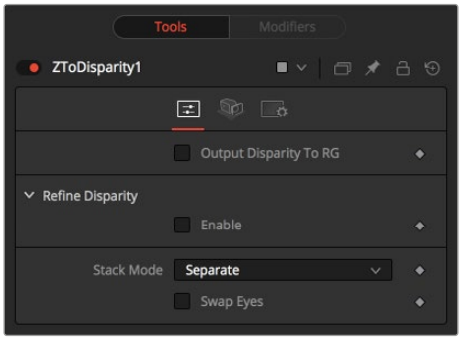
Basic Node Setup

Below, a stereo camera and an image containing a Z channel is connected to a Z To Disparity node. The same image is output with disparity channels.



A Z To Disparity node takes an image with a Z channel and creates a disparity channel

Inspector



The Z To Disparity Controls tab

Controls Tab

The Controls tab includes settings that refine the conversion algorithm.

Output Disparity To RGB

In addition to outputting disparity values into the disparity channel, activating this checkbox causes Z To Disparity to also output the disparity values into the color channels as {x, y, 0, 1}.

When activated, this option will automatically promote the RGBA color channels to float32. This option is useful for a quick look to see what the disparity channel looks like.

Refine Disparity

This refines the Disparity map based on the RGB channels.

Strength

Increasing this slider does two things. It smooths out the depth in constant color regions and moves edges in the Z channel to correlate with edges in the RGB channels. Increasing the refinement has the undesirable effect of causing texture in the color channel to show up in the Z channel. You will want to find a balance between the two.

Radius

This is the pixel-radius of the smoothing algorithm.

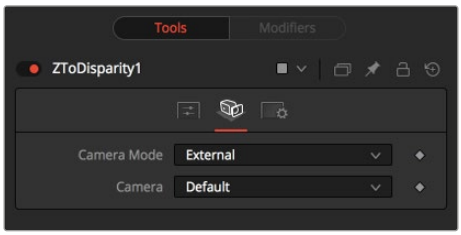
Stack Mode

In Stack Mode, L and R outputs will output the same image.

If HiQ is off, the interpolations are done using nearest-neighbor sampling, leading to a more “noisy” result.

Swap Eyes

This allows you to easily swap the left and right eye outputs.



The Z To Disparity Camera tab

Camera Tab

The Camera tab includes settings for selecting a camera and setting its conversion point if necessary.

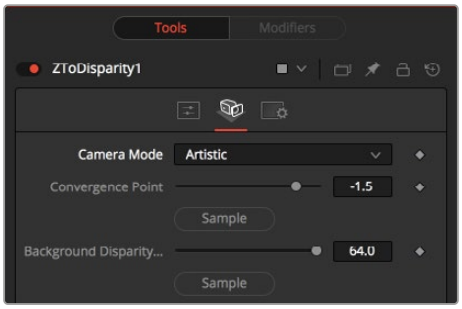
Camera Mode

If you need correct real-world disparity values because you are trying to match some effect to an existing scene, you should use the External setting to get precise disparity values back. When External is selected, a magenta camera input is available on the node to connect an existing stereo Camera 3D node, and use the Camera settings to determine the Disparity settings.

If you just want any disparity and do not particularly care about the exact details of how it is offset and scaled, or if there is no camera available, then the Artistic setting might be helpful.

Camera

If you connect a Merge 3D node that contains multiple cameras to the camera input, the Camera menu allows you to select the camera to use.



Artistic Camera mode

If you do not have a camera, you can adjust the artistic controls to produce a custom disparity channel whose values will not be physically correct but will be good enough for compositing hacks. There are two controls to adjust:

Convergence Point

This is the Z value of the convergence plane. This corresponds to the negative of the Convergence Distance control that appears in Camera 3D. At this distance, objects in the left and right eyes are at exactly the same position (i.e., have zero disparity).

Objects that are closer appear to pop out of the screen, and objects that are further appear behind the screen.

Background Disparity (Sample from Left Eye)

This is the disparity of objects in the distant background. You can think of this as the upper limit to disparity values for objects at infinity. This value should be for the left eye. The corresponding value in the right eye will be the same in magnitude but negative.

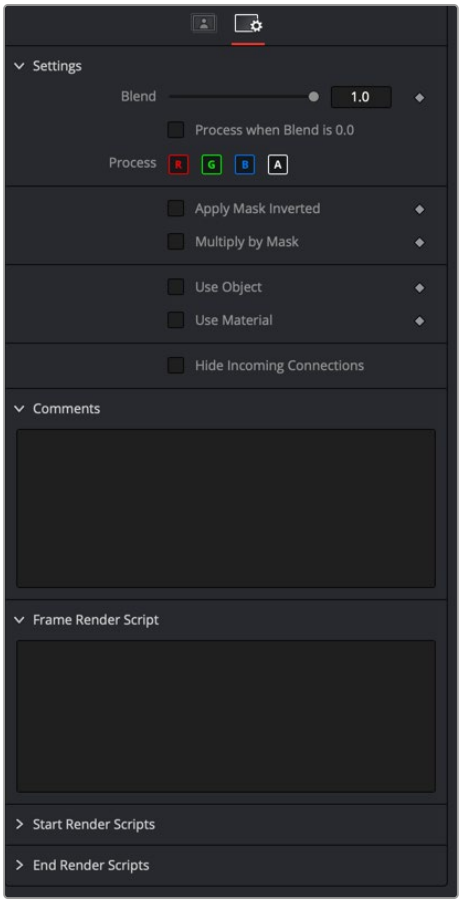
Settings Tab

The Settings tab in the Inspector is also duplicated in other Stereo nodes. These common controls are described in detail in the following “The Common Controls” section.

The Common Controls

The Common Settings tab can be found on virtually every tool found in Fusion. The following controls are specific settings for Stereo nodes.

Settings Tab



The Common Settings Stereo 3D Settings tab

Blend

The Blend control is used to blend between the tool’s original image input and the tool’s final modified output image. When the blend value is 0.0, the outgoing image is identical to the incoming image. Normally, this will cause the tool to skip processing entirely, copying the input straight to the output.

Process When Blend Is 0.0

The tool is processed even when the input value is zero. This can be useful if the node is scripted to trigger a task, but the node's value is set to 0.0.

Red/Green/Blue/Alpha Channel Selector

These four buttons are used to limit the effect of the tool to specified color channels. This filter is often applied after the tool has been processed.

For example, if the Red button on a Blur tool is deselected, the blur will first be applied to the image, and then the red channel from the original input will be copied back over the red channel of the result.

There are some exceptions, such as tools for which deselecting these channels causes the tool to skip processing that channel entirely. Tools that do this will generally possess a set of identical RGBA buttons on the Controls tab in the tool. In this case, the buttons in the Settings and the Controls tabs are identical.

Apply Mask Inverted

Enabling the Apply Mask Inverted option inverts the complete mask channel for the tool. The mask channel is the combined result of all masks connected to or generated in a node.

Multiply by Mask

Selecting this option will cause the RGB values of the masked image to be multiplied by the mask channel's values. This will cause all pixels of the image not included in the mask (i.e., set to 0) to become black/transparent.

Use Object/Use Material (Checkboxes)

Some 3D software can render to file formats that support additional channels. Notably, the EXR file format supports Object ID and Material ID channels, which can be used as a mask for the effect. These checkboxes determine whether the channels will be used, if present. The specific Material ID or Object ID affected is chosen using the next set of controls.

Correct Edges

This checkbox appears only when the Use Object or Use Material checkboxes are selected. It toggles the method used to deal with overlapping edges of objects in a multi-object image. When enabled, the Coverage and Background Color channels are used to separate and improve the effect around the edge of the object. If this option is disabled (or no Coverage or Background Color channels are available), aliasing may occur on the edge of the mask.

> For more information on the Coverage and Background Color channels, *see Chapter 18, "Understanding Image Channels,"* in the Fusion Reference Manual.

Object ID/Material ID (Sliders)

Use these sliders to select which ID will be used to create a mask from the object or material channels of an image. Use the Sample button in the same way as the Color Picker: to grab IDs from the image displayed in the viewer. The image or sequence must have been rendered from a 3D software package with those channels included.

Hide Incoming Connections

Enabling this checkbox can hide connection lines from incoming nodes, making a node tree appear cleaner and easier to read. When enabled, empty fields for each input on a node will be displayed in the Inspector. Dragging a connected node from the node tree into the field will hide that incoming connection line as long as the node is not selected in the node tree. When the node is selected in the node tree, the line will reappear.

Comments

The Comments field is used to add notes to a tool. Click in the empty field and type the text. When a note is added to a tool, a small red square appears in the lower-left corner of the node when the full tile is displayed, or a small text bubble icon appears on the right when nodes are collapsed. To see the note in the Node Editor, hold the mouse pointer over the node to display the tooltip.

Scripts

Three Scripting fields are available on every tool in Fusion from the Settings tab. They each contain edit boxes used to add scripts that process when the tool is rendering. For more details on scripting nodes, please consult the Fusion scripting documentation.

Chapter 59

Tracking Nodes

This chapter details the Tracking nodes available in Fusion.

The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Camera Tracker [CTra] 1567
Camera Tab 1570
Solve Tab 1572
Track Filtering 1575
Operations on Selected Tracks 1576
Export 1577
3D Scene Transform 1578
Options 1580
Understanding Camera Tracking 1581
Planar Tracker Node [PTRA] 1583
What the Planar Tracker Saves 1583
A Typical Planar Tracker Workflow 1584
Track Mode 1587
Steady Mode 1589
Corner Pin Mode 1589
Merge Mode 1590
Options Tab 1592
Surface Tracker [SFt] 1593
Tracker [TRA] 1602
IntelliTrack AI Point Tracker 1604
Legacy Point Tracker Onscreen Controls 1604
The Common Controls 1619

Camera Tracker [CTra]



The Camera Tracker node

Camera tracking is match moving, and a vital link between 2D and 3D, allowing compositors to integrate 3D renders into live-action scenes. The Camera Tracker node is used to calculate the path of a live-action camera and generate a virtual camera in 3D space. This virtual camera's motion is intended to be identical to the motion of the actual camera that shot the scene. Using the calculated position and movement of the virtual camera provides the flexibility to add 3D elements to a live-action scene. Also, the Camera Tracker creates a point cloud in 3D space that can be used to align objects and other 3D models to the live-action scene.

For more information about how to use the Camera Tracker, *see Chapter 26, "3D Camera Tracking,"* in the Fusion Reference Manual.



The original video clip



3D objects composited on to the video clip that use the camera tracker to remain aligned with the objects in the frame as the image moves

Inputs

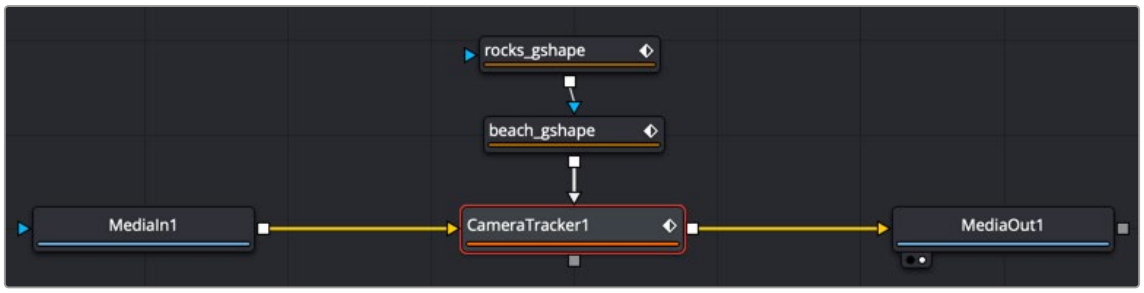
The Camera Tracker has two inputs:

Background: The orange image input accepts a 2D image you want tracked.

Occlusion Mask: The white occlusion mask input is used to mask out regions that do not need to be tracked. Regions where this mask is white will not be tracked. For example, a person moving in front of and occluding bits of the scene may be confusing to the tracker, and a quickly-created rough mask around the person can be used to tell the tracker to ignore the masked-out bits.

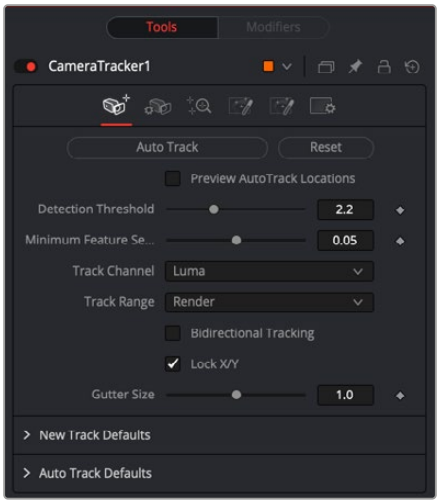
Basic Node Setup

The Camera Tracker background input is used to connect the image you want tracked. Polygon masks can be connected into the occlusion mask input to identify areas the tracker should ignore.



The Camera Tracker with occlusion masks

Inspector



The Camera Tracker tab

Track Tab

The Track tab contains the controls you need to set up an initial analysis of the scene.

Auto Track

Automatically detects trackable features and tracks them through the source footage. Tracks will be automatically terminated when the track error becomes too high, and new tracks are created as needed. The values of the Detection Threshold and Minimum Feature Separation sliders can be used to control the number and distribution of auto tracks.

Reset

Deletes all the data internal to the Camera Tracker node, including the tracking data and the solve data (camera motion path and point cloud). To delete only the solve data, use the Delete button on the Solve tab.

Preview AutoTrack Locations

Turning this checkbox on will show where the auto tracks will be distributed within the shot. This is helpful for determining if the Detection Threshold and Minimum Feature Separation need to be adjusted to get an even spread of trackers.

Detection Threshold

Determines the sensitivity to detect features. Automatically generated tracks will be assigned to the shot, and the Detection Threshold will force them to be either in locations of high contrast or low contrast.

Minimum Feature Separation

Determines the spacing between the automatically generated tracking points. Decreasing this slider causes more auto tracks to be generated. Keep in mind that a large number of tracking points will also result in a lengthier solve.

Track Channel

Used to nominate a color channel to track: red, green, blue, or luminance. When nominating a channel, choose one that has a high level of contrast and detail.

Track Range

Used to determine which frames are tracked:

- **Global:** The global range, which is the full duration of the Timeline.
- **Render:** The render duration set on the Timeline.
- **Valid:** The valid range is the duration of the source media.
- **Custom:** A user-determined range. When this is selected, a separate range slider appears to set the start and end of the track range.

Bidirectional Tracking

Enabling this will force the tracker to track backward after the initial forward tracking. When tracking backward, new tracks are not started but rather existing tracks are extended backward in time. It is recommended to leave this option on, as long tracks help give better solved cameras and point clouds.

Gutter Size

Trackers can become unstable when they get close to the edge of the image and either drift or jitter or completely lose their pattern. The Camera Tracker will automatically terminate any tracks that enter the gutter region. Gutter size is given as a percentage of pattern size. By default, it's 100% of pattern size, so a 0.04 pattern means a 0.04 gutter.

New Track Defaults

There are three methods in which the Camera Tracker node can analyze the scene, and each has its own strengths when dealing with certain types of camera movement.

- **Tracker:** Internally, all the Trackers use the Optical Flow Tracker to follow features over time and then further refine the tracks with the trusted Fusion Tracker or Planar Tracker. The Planar Tracker method allows the pattern to warp over time by various types of transforms to find the best fit. These transforms are:
  - Translation
  - Translation and Rotation Translation, Rotation, and Scale Affine
  - Perspective

It is recommended to use the default TRS setting when using the Planar Tracker. The Affine and Perspective settings need large patterns in order to track accurately.

- **Close Tracks When Track Error Exceeds:** Tracks will be automatically terminated when the tracking error gets too high. When tracking a feature, a snapshot of the pixels around a feature are taken at the reference time of the track. This is called a pattern, and that same pattern of pixels is searched for at future times. The difference between the current time pattern and the reference time pattern is called the track error. Setting this option higher produces longer but increasingly less accurate tracks.
- **Solve Weight:** By default, each track is weighted evenly in the solve process. Increasing a track's weight means it has a stronger effect on the solved camera path. This is an advanced option that should be rarely changed.

Auto Track Defaults

Set a custom prefix name and/or color for the automatically generated tracks. This custom color will be visible when Track Colors in the Options tab is set to User Assigned.

Camera Tab



The Camera Tracker Camera tab

The controls of the Camera tab let you specify the physical aspects of the live-action camera, which will be used as a starting point when searching for solve parameters that match the real-world camera. The more accurate the information provided in this section, the more accurate the camera solve.

The Camera tab includes controls relating to the lens and gate aspects of the camera being solved for.

Focal Length

Specify the known constant focal length used to shoot the scene or provide a guess if the Refine Focal Length option is activated in the Solve tab.

Film Gate

Choose a film gate preset from the drop-down menu or manually enter the film back size in the Aperture Width and Aperture Height inputs. Note that these values are in inches.

Aperture Width

In the event that the camera used to shoot the scene is not in the preset drop-down menu, manually enter the aperture width (inches).

Aperture Height

In the event that the camera used to shoot the scene is not in the preset drop-down menu, manually enter the aperture height (inches).

Resolution Gate Fit

This defines how the image fits the sensor size. Often, film sensors are sized to cover a number of formats, and only a portion of the sensor area is recorded into an image.

For example, a 16:9 image is saved out of a full aperture-sized sensor.

Typically, fit to Width or Height is the best setting. The other fit modes are Inside, Outside, or Stretched.

Center Point

This is where the camera lens is aligned to the camera. The default is (0.5, 0.5), which is the middle of the sensor.

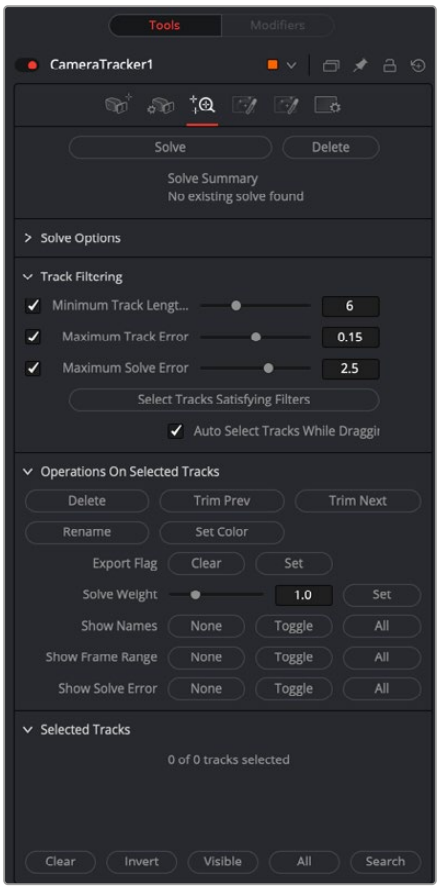
Use Source Pixel Aspect

This will use the squeeze aspect of the pixels that is loaded in the image. HD is square pixels, but NTSC has a pixel aspect ratio of 0.9:1, and Anamorphic CinemaScope is 2:1 aspect. Disabling this option exposes Pixel X and Y number fields where you can customize the source pixel aspect.

Auto Camera Planes

When enabled, the camera's image plane and far plane are automatically moved to enclose the point cloud whenever a solve completes. Sometimes, though, the solver can atypically fling points off really deep into the scene, consequently pushing the image plane very far out. This makes the resulting scene unwieldy to work with in the 3D views. In these cases, disable this option to override this default behavior (or delete the offending tracks).

Solve Tab



The Camera Tracker Solve tab

The Solve tab is where the tracking data is used to reconstruct the camera's motion path along with the point cloud. It is also where clean up of bad or false tracks is done, and other operations on the tracks can be performed, such as defining which marks are exported in the Point Cloud 3D. The markers can also have their weight set to affect the solve calculations.

> For example, a good camera solve may have already been generated, but there are not enough locators in the point cloud in an area where an object needs to be placed, so adding more tracks and setting their Solve Weight to zero will not affect the solved camera but will give more points in the point cloud.

Solve

Pressing Solve will launch the solver, which uses the tracking information and the camera specifications to generate a virtual camera path and point cloud, approximating the motion of the physical camera in the live-action footage. The console will automatically open, displaying the progress of the solver.

Delete

Delete will remove any solved information, such as the camera and the point cloud but will keep all the tracking data.

Average Solve Error

Once the camera has been solved, a summary of the solve calculation is displayed at the top of the Inspector. Chief among those details is the Average Solve Error. This number is a good indicator of whether the camera solve was successful. It can be thought of as the difference (measured in pixels) between tracks in the 2D image and the reconstructed 3D locators reprojected back onto the image through the reconstructed camera. Ultimately, in trying to achieve a low solve error, any value less than 1.0 pixels will generally result in a good track. A value between 0.6 and 0.8 pixels is considered excellent.

Clean Tracks by Filter

Clicking this button selects tracks based on the Track Filtering options. If the Auto Delete Tracks By Filter checkbox is activated, the selected tracks will be deleted as well.

Clean Foreground Tracks

Clicking this button makes a selection of the tracks on fast-moving objects that would otherwise cause a high solve error. The selection is determined by the Foreground Threshold slider.

Foreground Threshold

This slider sets the detection threshold for finding the tracks on fast-moving objects. The higher the value, the more forgiving.

Auto Delete Tracks by Filter

With this checkbox enabled, tracks that are selected by the Clean Tracks By Filter button will be deleted. Enable the checkbox, and then press Clean Tracks By Filter. Any track that meets the filtering options is then selected and deleted.

Auto Delete Foreground Tracks

With this checkbox enabled, tracks that are selected by the Clean Foreground Tracks button will be deleted. Enable the checkbox, and then press Clean Foreground Tracks. Any track that meets the foreground threshold criteria is deleted.

Accept Solve Error

This slider sets an acceptable maximum threshold level for the solve error. If the solve error is greater than this value, the Camera Tracker will sweep the focal length setting in an attempt to bring the solve error under the Accept Solve Error value. If the solver cannot find a solution, the Camera Tracker will display a message in the console that the solver failed. If a solution cannot be found, ideally you should try to input the correct focal length or alternatively manually clean some noisy tracks then re-solve.

Auto Select Seed Frames

With this enabled, the Camera Tracker nominates two frames that will be used as a reference for initiating the solve. These two frames are initially solved for and a camera is reconstructed, and then gradually more frames are added in, and the solution is "grown" outward from the seed frames. The choice of seed frames strongly affects the entire solve and can easily cause the solve to fail. Seed frames can be found automatically or defined manually.

Disabling this will allow the user to select their own two frames. Manual choice of seed frames is an option for advanced users. When choosing seed frames, it is important to satisfy two conflicting desires: the seed frames should have lots of tracks in common yet be far apart in perspective (i.e., the baseline distance between the two associated cameras is long).

Refine Focal Length

Enabling this will allow the solver to adjust the focal length of the lens to match the tracking points. You can prevent the focal length being adjusted by setting the Focal Length parameter in the Camera tab.

Enable Lens Parameter

When enabled, lens distortion parameters are exposed to help in correcting lens distortion when solving.

- **Refine Center Point:** Normally disabled, camera lenses are typically centered in the middle of the film gate, but this may differ on some cameras. For example, a cine camera may be set up for Academy 1.85, which has a sound stripe on the left, and shooting super35, the lens is offset to the right.
- **Refine Lens Parameters:** This will refine the lens distortion or curvature of the lens. There tends to be larger distortion on wide angle cameras.

NOTE: When solving for the camera's motion path, a simulated lens is internally created to model lens distortion in the source footage. This simulated lens model is much simpler than real-world lenses but captures the lens distortion characteristics important for getting an accurate camera solve. Two types of distortion are modeled by Camera Tracker:

Radial Distortion: The strength of this type of distortion varies depending on the distance from the center of the lens. Examples of this include pincushion, barrel, and mustache distortion. Larger values correspond to larger lens curvatures. Modeling radial distortion is especially important for wide angle lenses and fisheye lenses (which will have a lot of distortion because they capture 180 degrees of an environment and then optically squeeze it onto a flat rectangular sensor).

Tangential Distortion: This kind of distortion is produced when the camera's imaging sensor and physical lens are not parallel to each other. It tends to produce skew distortions in the footage similar to distortions that can be produced by dragging the corners of a corner pin within Fusion. This kind of distortion occurs in very cheap consumer cameras and is practically non-existent in film cameras, DSLRs, and pretty much any kind of camera used in film or broadcast. It is recommended that it be left disabled.

Enable Lens Parameters

When disabled, the Camera Tracker does not do any lens curvature simulations. This is the default setting and should remain disabled if there is a very low distortion lens or the lens distortion has already been removed from the source clip. Activating the Enable Lens Parameters checkbox determines which lens parameters will be modeled and solved for. Parameters that are not enabled will be left at their default values. The following options are available:

- **Radial Quadratic:** Model only Quadratic radial lens curvature, which is either barrel or pincushion distortion. This is the most common type of distortion. Selecting this option causes the low and high order distortion values to be solved for.
- **Radial Quartic:** Model only Quartic radial lens curvature, which combines barrel and pincushion distortion. This causes the low and high order distortion values to be solved for.
- **Radial & Tangential:** Model and solve for both radial and tangential distortion. Tangential relates to misaligned elements in a lens.
- **Division Quadratic:** Provides a more accurate simulation of Quadratic radial lens curvature. This causes the low and high order distortion values to be solved for.

- **Division Quartic:** Provides a more accurate simulation of Quartic radial lens curvature. This causes the low and high order distortion values to be solved for.

  Lower Order Radial Distortion: This slider is available for all simulations. It determines the quadratic lens curvature.

  Higher Order Radial Distortion: This slider is available only for Quartic simulations. Determines the quartic lens curvature.

  Tangential Distortion X/Y: These sliders are available only for Tangential simulations. Determines skew distortion.

Track Filtering

The Camera Tracker can produce a large number of automatically generated tracks. Rather than spending a lot of time individually examining the quality of each track, it is useful to have some less time-intensive ways to filter out large swaths of potentially bad tracks. The following input sliders are useful for selecting large amounts of tracks based on certain quality metrics, and then a number of different possible operations can be made on them. For example, weaker tracks can be selected and deleted, yielding a stronger set of tracks to solve from. Each filter can be individually enabled or disabled.

Minimum Track Length (Number of Markers)

Selects tracks that have a duration shorter than the slider's value. Short tracks usually don't get a chance to move very far and thus provide less perspective information to the solver than a longer track, yet both short and long tracks are weighted evenly in the solve process, making long tracks more valuable to the solver. Locators corresponding to shorter tracks are also less accurately positioned in 3D space than those corresponding to longer tracks. If the shot has a lot of long tracks, it can be helpful to delete the short tracks. For typical shots, using a value in the range of 5 to 10 is suggested. If there are not a lot of long tracks (e.g., the camera is quickly rotating, causing tracks to start and move off frame quickly), using a value closer to 3 is recommended.

Maximum Track Error

Selects tracks that have an average track error greater than the slider's value. When tracking, tracks are automatically terminated when their track error exceeds some threshold. This auto termination controls the maximum track error, while this slider controls the average track error. For example, tracks following the foliage in a tree tend to be inaccurate and sometimes may be detected by their high average error.

Maximum Solve Error

Selects tracks that have a solve error greater than the slider's value. One of the easiest ways to increase the accuracy of a camera solve is to select the 20% of the tracks with the highest solve error and simply delete them (although this can sometimes make things worse).

Select Tracks Satisfying Filters

Selects the tracks within the scene that meet the above Track Filtering values. Note that when this button is pressed, the tracks that satisfy the filter values are displayed in the Selected Tracks area of the Solve tab and are colored in the viewer. This button is useful when Auto Select Tracks While Dragging Sliders is turned off or if the selection, for example, was accidentally lost by mis-clicking in the viewer.

Auto Select Tracks While Dragging Sliders

When this is ticked on, dragging the above sliders (Minimum Track Length, Maximum Track Error, Maximum Solve Error) will cause the corresponding tracks to be interactively selected in the viewer.

Operations on Selected Tracks

Tracks selected directly in the viewer with the mouse or selected via track filtering can have the following operations applied:

Delete	Removes the tracks from the set. When there are bad tracks, the simplest and easiest option is to simply delete them.
Trim Previous	Cuts the tracked frames from the current frame to the start of the track. Sometimes it can be more useful to trim a track than to delete it. For example, high-quality long tracks that become inaccurate when the feature they are tracking starts to become occluded or when the tracked feature moves too close to the edge of the image.
Trim Next	Cuts the tracked frames from the current frame to the end of the track.
Rename	Replaces the current auto generated name with a new name.
Set Color	Allows for user-assigned color of the tracking points.
Export Flag	This controls whether the locators corresponding to the selected tracks will be exported in the point cloud. By default, all locators are flagged as exportable.
Solve Weight	By default, all the tracks are used and equally weighted when solving for the camera's motion path. The most common use of this option is to set a track's weight to zero, so it does not influence the camera's motion path but still has a reconstructed 3D locator. Setting a track's weight to values other than 1.0 or 0.0 should only be done by advanced users. Onscreen display of track names and values are controlled by these functions:
None	Clears/hides the selected tracks.
Toggle	Swaps the selected tracks and unselected sets.
All	Selects all tracks.
Show Names	Displays the track name. By default, these are numbers.
Show Frame Range	Displays the start and end frame of a track.
Show Solve Error	Displays the amount of solve error of each selected track.

Selected Tracks

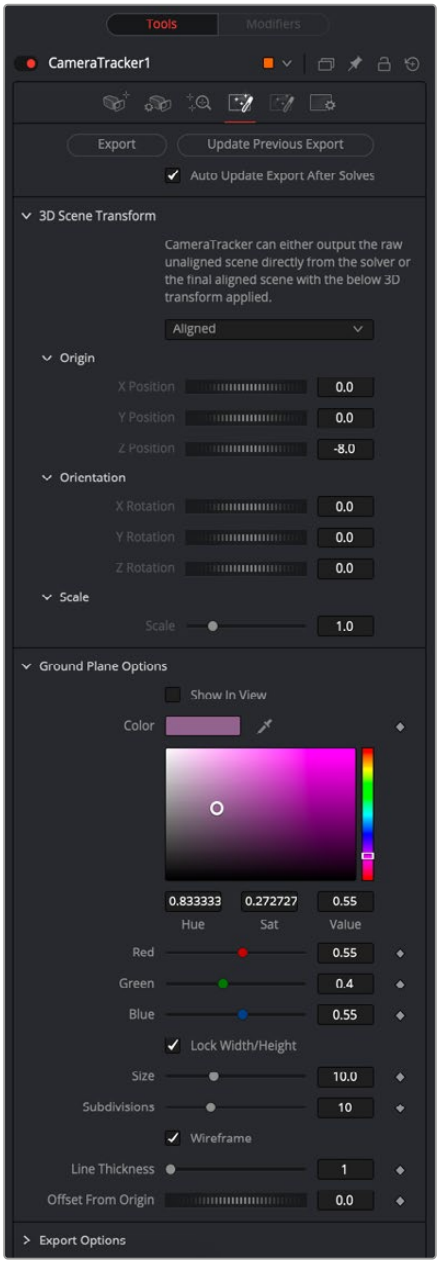
This area displays the properties of a track point or group of points. It has options to:

- **Clear:** Deselects all tracks and clears this area.
- **Invert:** Deselects the current selected tracks and selects the other tracks.
- **Visible:** Selects all the trackers at the current frame.
- **All:** Selects all trackers on all frames.
- **Search:** Selects tracks whose names contain a substring.

TIP: Also select tracks directly in the 2D viewer using the mouse or in the 3D viewer by selecting their corresponding locators in the point cloud.

Export

The Export tab lets you turn the tracked and solved data this node has generated into a form that can be used for compositing.



The Camera Tracker Export tab

Export

The Export button will create a basic setup that can be used for 3D match moving:

- A Camera 3D with animated translation and rotation that matches the motion of the live-action camera and an attached image plane.
- A Point Cloud 3D containing the reconstructed 3D positions of the tracks.
- A Shape 3D set to generate a ground plane.

- A Merge 3D merging together the camera, point cloud, and ground plane. When the Merge 3D is viewed through the camera in a 3D viewer, the 3D locators should follow the tracked footage.
- A Renderer 3D set to match the input footage.

The export of individual nodes can be enabled/disabled in the Export Options tab.

Update Previous Export

When this button is clicked, the previously exported nodes are updated with any new data generated. These previously exported nodes are remembered in the Previous Export section at the bottom of this section. Here's an example of how this is handy:

1. Solve the camera and export.
2. Construct a complex Node Editor based around the exported nodes for use in set extension.
3. The camera is not as accurate as preferred or perhaps the solver is rerun to add additional tracks to generate a denser point cloud. Rather than re-exporting the Camera 3D and Point Cloud 3D nodes and connecting them back in, just press the Update Previous Export button to "overwrite" the existing nodes in place.

Automatically Update Previous Export After Solves

This will cause the already exported nodes (Camera 3D, Point Cloud 3D, Lens Distort, Renderer 3D, and the ground plane) to auto update on each solve.

3D Scene Transform

Although the camera is solved, it has no idea where the ground plane or center of the scene is located. By default, the solver will always place the camera in Fusion's 3D virtual environment so that on the first frame it is located at the origin (0, 0, 0) and is looking down the Z-axis. You have the choice to export this raw scene without giving the Camera Tracker any more information, or you can set the ground plane and origin to simplify your job when you begin working in the 3D scene. The 3D Scene Transform controls provide a mechanism to correctly orient the physical ground plane in the footage with the virtual ground plane in the 3D viewer. Adjusting the 3D Scene Transform does not modify the camera solve but simply repositions the 3D scene to best represent the position of the live-action camera.

NOTE: If you export the scene and then make changes in the 3D Scene Transform, it is important to manually click Update Previous Export to see the results in the exported nodes.

Aligned/Unaligned

The Aligned/Unaligned menu locks or unlocks the origin and ground plane settings. When set to Unaligned, you can select the ground plane and origin either manually or by selecting locators in the viewer. When in unaligned mode, a 3D Transform control in the 3D viewer can be manually manipulated to adjust the origin.

Once alignment of the ground plane and origin has been completed, the section is locked by switching the menu to Aligned.

Set from Selection

When set to unaligned, buttons labeled Set from Selection are displayed under the Origin, Orientation, and Scale sections. Clicking these buttons takes the selecting locators in the viewer and aligns the ground plane or origin based on the selection.

For instance, to set the ground plane, do the following:

1. After solving, set the 3D Scene Transform menu to Unaligned.
2. Find a frame where the ground plane is at its largest and clearest point.
3. In the viewer, drag a selection rectangle around all the ground plane locators.
4. Hold Shift and drag again to add to the selection.
5. In the Orientation section, make sure the Selection Is menu correctly matches the orientation of the selected locators.
6. Click the Set from Selection button located under the Orientation parameters.
7. Set the 3D Scene Transform menu back to Aligned.

To get the best result when setting the ground plane, try to select as many points as possible belonging to the ground and having a wide separation.

> **TIP:** When selecting points for the ground plane, it is helpful to have the Camera Tracker node viewed in side-by-side 2D and 3D views. It may be easier to select tracks belonging to the ground by selecting tracks from multiple frames in the 2D viewer rather than trying to box select locators in the 3D viewer.

Setting the origin can help you place 3D objects in the scene with more precision. To set the origin, you can follow similar steps, but only one locator is required for the origin to be set. When selecting a locator for the origin, select one that has a very low solve error.

Ground Plane Options

These controls let you adjust the ground plane for the scene, which is a crucial step in making sure the composite looks correct.

Color	Sets the color of the ground plane.
Size	Controls how big the ground plane can be set.
Subdivision Level	Shows how many polygons are in the ground plane.
Wireframe	Sets whether the ground plane is set as wireframe or solid surface when displayed in 3D.
Line Thickness	Adjusts how wide the lines will draw in the viewer.
Offset	By default, the center of the ground plane is placed at the origin (0, 0, 0). This can be used to shift the ground plane up and down along the Y-axis.

Export Options

Provides a checkbox list of what will be exported as nodes when the Export button is pressed. These options are: Camera, Point Cloud, Ground Plane, Renderer, Lens Distortion, and Enable Image Plane in the camera.

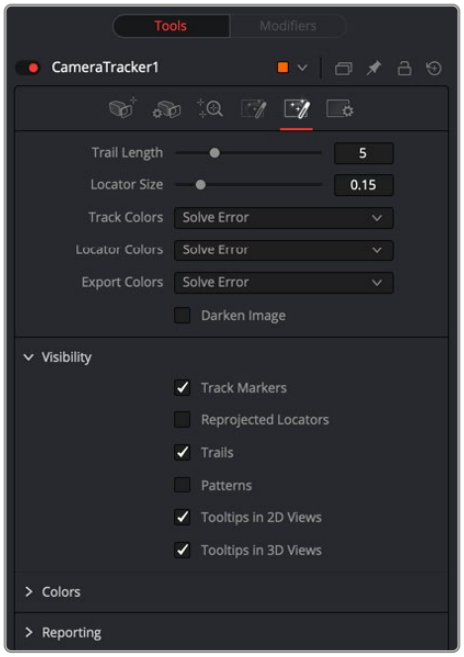
The Animation menu allows you to choose between animating the camera and animating the point cloud. Animating the camera leaves the point cloud in a locked position while the camera is keyframed to match the live-action shot. Animating the point cloud does the opposite. The camera is locked in position while the entire point cloud is keyframed to match the live-action shot.

Previous Export

When the Update Previous Export button is clicked, the previously exported nodes listed here are updated with any new data generated (this includes the camera path and attributes, the point cloud, and the renderer).

Options

The Options tab lets you customize the Camera Tracker's onscreen controls so you can work most effectively with the scene material you have.



The Camera Tracker Options tab

Trail Length

Displays trail lines of the tracks overlaid on the viewer. The amount of frames forward and back from the current frame is set by length.

Location Size

In the 3D viewer, the point cloud locators can be sized by this control.

Track Colors, Locator Colors, and Export Colors each have options for setting their color to one of the following: User Assigned, Solve Error, Take From Image, and White.

Track Colors	The onscreen tracks in the 2D viewer.
Locator Colors	The point cloud locators in the 3D viewer.
Export Colors	The colors of the locators that get exported within the Point Cloud node.

Darken Image

Dims the brightness of the image in viewers to better see the overlaid tracks. This affects both the 2D and 3D viewers.

Visibility

Toggles which overlays will be displayed in the 2D and 3D viewers. The options are: Tracker Markers, Trails, Tooltips in the 2D viewer, Tooltips in the 3D viewer, Reprojected Locators, and Tracker Patterns.

Colors

Sets the color of the overlays.

- **Selection Color:** Controls the color of selected tracks/locators.
- **Preview New Tracks Color:** Controls the color of the points displayed in the viewer when the Preview AutoTrack Locations option is enabled.
- **Solve Error Gradient:** By default, tracks and locators are colored by a green-yellow-red gradient to indicate their solve error. This gradient is completely user adjustable.

Reporting

Outputs various parameters and information to the Console.

Understanding Camera Tracking

On large productions, camera tracking or 3D match moving is often handed over to experts who have experience with the process of tracking and solving difficult shots. There is rarely a shot where you can press a couple of buttons and have it work perfectly. It does take an understanding of the whole process and what is essential to get a good solved track.

The Camera Tracker must solve for hundreds of thousands of unknown variables, which is a complex task. For the process to work, it is essential to get good tracking data that exists in the shot for a long time. False or bad tracks will skew the result. This section explains how to clean up false tracks and other techniques to get a good solve.

Workflow

Getting a good solve is a repeated process.

Track > Solve > Refine Filters > Solve > Cleanup tracks > Solve > Cleanup from point cloud > Solve > Repeat.

Initially, there are numerous tracks, and not all are good, so a process of filtering and cleaning up unwanted tracks to get to the best set is required. At the end of each cleanup stage, pressing Solve ideally gives you a progressively lower solve error. This needs to be below 1.0 for it to be good for use with full HD content, and even lower for higher resolutions. Refining the tracks often but not always results in a better solve.

False Tracks

False tracks are caused by a number of conditions, such as moving objects in a shot, or reflections and highlights from a car. There are other types of false tracks like parallax errors where two objects are at different depths, and the intersection gets tracked. These moiré effects can cause the track to creep. Recognizing these False tracks and eliminating them is the most important step in the solve process.

Track Lengths

Getting a good set of long tracks is essential; the longer the tracks are, the better the solve. The Bi-Directional tracking option in the Tracker tab is used to extend the beginning of tracks in time. The longer in time a track exists and the more tracks that overlap in time of a shot, the more consistent and accurate the solve.

Seed Frames

Two seed frames are used in the solve process. The algorithm chooses two frames that are as far apart in time yet share the same tracks. That is why longer tracks make a more significant difference in the selection of seed frames.

The two Seed frames are used as the reference frames, which should be from different angles of the same scene. The solve process will use these as a master starting point to fit the rest of the tracks in the sequence.

There is an option in the Solve tab to Auto Detect Seed Frames, which is the default setting and most often a good idea. However, auto detecting seed frames can make for a longer solve. When refining the Trackers and re-solving, disable the checkbox and use the Seed 1 and Seed 2 sliders to enter the previous solve's seed frames. These seed frames can be found in the Solve Summary at the top of the Inspector after the initial solve.

Refine Filters

After the first solve, all the Trackers will have extra data generated. These are solve errors and tracking errors.

Use the refine filters to reduce unwanted tracks, like setting minimum tracker length to eight frames. As the value for each filter is adjusted, the Solve dialog will indicate how many tracks are affected by the filter. Then Solve again.

Onscreen Culling

Under the Options tab, set the track to 20. This will display each track on footage with +-20 frame trail. When scrubbing/playing through the footage, false tracks can be seen and selected onscreen, and deleted by pressing the Delete key. This process takes an experienced eye to spot tracks that go bad. Then Solve again.

You can view the exported scene in a 3D perspective viewer. The point cloud will be visible. Move and pan around the point cloud, select and delete points that seem to have no inline with the image and the scene space. Then Solve again.

Repeat the process until the solve error is below 1.0 before exporting.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Tracking nodes. These common controls are described in detail in the following "The Common Controls" section.

Planar Tracker Node [PTRA]



The Planar Tracker node

The Planar Tracker node is designed to solve a match-moving problem that commonly comes up during post-production. As an example, live-action footage can often contain a planar surface such as a license plate or a road sign that needs new numbers in the license plate or a new city's name on the road sign. Often, the problem is that the camera is moving in the shot, so the license plate or road sign is continuously changing perspective. You cannot just merge a new license plate over the existing one without accounting for the perspective distortions. A time-intensive way to solve this problem would be to use a Corner Pin node and manually keyframe the four corners. The Planar Tracker automates this keyframing process and tracks the perspective distortions of a planar surface over time. That tracking data is then applied with those same perspective distortions to a different foreground.

For more information on using the Planar Tracker, *see Chapter 23, "Planar Tracking,"* in the Fusion Reference Manual.

TIP: Part of using a Planar Tracker is also knowing when to give up and fall back to using Fusion's Tracker node or to manual keyframing. Some shots are simply not trackable, or the resulting track suffers from too much jitter or drift. The Planar Tracker is a useful time-saving node in the artist's toolbox, but while it may track most shots, it is not a 100% solution.

What the Planar Tracker Saves

While the Planar Tracker does save the resulting final track in the composition on disk, it does not save temporary tracking information, such as the individual point trackers (compared with the Camera Tracker, which does save the individual point trackers). Some consequences of this include:

- The point trackers no longer appear in the viewer when a comp containing a Planar Tracker node is saved and reloaded.
- Tracking may not be resumed after a comp containing a Planar Tracker node has been saved and reloaded. In particular, this also applies to auto saves. For this reason, it is good to complete all planar tracking within one session.
- The size of composition files is kept reasonable (in some situations, a Planar Tracker can produce hundreds of megabytes of temporary tracking data).
- Saving and loading of compositions is faster and more interactive.

Inputs

The Planar Tracker has four inputs:

Background: The orange background image input accepts a 2D image with the planar surface to be tracked.

Corner Pin 1: The green corner pin 1 input accepts a 2D image to be pinned on top of the background. There may be multiple corner pin inputs, named Corner Pin 1, Corner Pin 2, etc.

Occlusion Mask: The white occlusion mask input is used to mask out regions that do not need to be tracked. Regions where this mask is white will not be tracked. For example, a person moving in front of and occluding bits of the pattern may be confusing the tracker, and a quickly-created rough rotomask around the person can be used to tell the tracker to ignore the masked-out bits.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the output of the Planar Tracker to certain areas.

Basic Node Setup

A basic Planar Tracker setup consists of just two nodes: a MediaIn connects to the background input and the Planar Tracker can be used as a separate branch from the rest of the node tree. Once the tracking is completed, a Planar Transform node should be generated to use the planar tracking data.



A Planar Tracker can be isolated on its own branch of a node tree.

A Typical Planar Tracker Workflow

The following steps outline the workflow with the Planar Tracker:

1. Remove lens distortion: The more lens distortion in the footage, the more the resulting track will slide and wobble.
2. Connect footage: Connect a Loader or MediaIn node that contains a planar surface to the orange background input and view the Planar Tracker node in a viewer.
3. Select a reference frame: Move to a frame where the planar surface to be tracked is not occluded and click the Set button to set this as a reference frame.
4. Choose the pattern: In the viewer, make sure the onscreen controls are visible, and draw a polygon around the planar surface you want to track. This is called the “pattern.” In most cases, this will probably be a rectangle, but an arbitrary closed polygon can be used. The pixels enclosed by this region will serve as the pattern that will be searched for on other frames. Note that it is important that the pattern is drawn on the reference frame. Do not confuse the pattern with the region to corner pin (which always has four corners and is separately specified in Corner Pin mode).
5. Adjust render range: In the Keyframes Editor, adjust the render range to match the range of frames where the planar surface is visible.

6 Adjust track options: Frequently changed options include Tracker, Motion Type, and Track Channel.

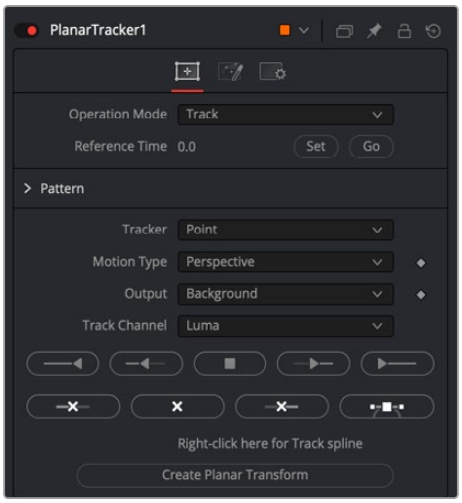
7 Mask out occluders: If moving objects partially cover up the planar surface, you may wish to connect an occlusion mask to the Planar Tracker. When using the Hybrid tracker, providing a mask to deal with occluding objects is strongly recommended, while with the Point tracker it is recommended to try tracking without a mask.

8 Track: Click the Go button to return to the reference frame. Press the Track To End button and wait for the track to complete. Click the Go button to return to the reference frame again. Press the Track To Start button and wait for the track to complete. Note that the tracks in the viewer are not selectable or deletable as they are in a Camera Tracker.

9 Check track quality: Visually inspect the track to see how accurate it is. Does it stick to the surface? Switching to Steady mode can help here.

10 Use the track: At this point, in most cases, you will create a Planar Transform node from the Inspector and use it to apply the tracked perspective distortion onto masked images. If the image you are applying the tracking data to is a full frame unmasked clip, you can use the Steady, Corner Pin, and Stabilize operation modes in the Planar Tracker.

Inspector



The Planar Tracker Controls tab

Controls tab

The Controls tab contains controls for determining how the Planar Tracker will be used, setting the reference frame, and initiating the track.

Operation Mode

The Operation Mode menu selects the purpose of the Planar Tracker node. The Planar Tracker has four modes of operation:

- **Track:** Used to isolate a planar surface and track its movement over time. Then, you can create a Planar Transform node that uses this data to match move another clip in various ways.
- **Steady:** After analyzing a planar surface, this mode removes all motion and distortions from the planar surface, usually in preparation for some kind of paint or roto task, prior to “unsteadying” the clip to restore the motion.

- **Corner Pin:** After analyzing a planar surface, this mode computes and applies a matching perspective distortion to a foreground image you connect to the foreground input of the Planar Tracker node and merges it on top of the tracked footage.
- **Stabilize:** After analyzing a planar surface, this mode allows smoothing of a clip's translation, rotation, and scale over time. This is good for getting unwanted vibrations out of a clip while retaining the overall camera motion that was intended.

The last three modes (Steady, Corner Pin, and Stabilize) use the tracking data produced in Track mode.

> **NOTE:** None of the operations can be combined together. For example, both Corner Pin and Stabilize cannot be done at the same time, nor can a track be done while in corner pinning mode.

Reference Time

The Reference Time determines the frame where the pattern is outlined. It is also the time from which tracking begins. The reference frame cannot be changed once it has been set without destroying all pre-existing tracking information, so scrub through the footage to be tracked and choose carefully. The reference frame must be chosen carefully to give the best possible quality track.

You choose a reference frame by moving the playhead to an appropriate frame and then clicking the Set button to choose that frame.

Pattern Polygon

You specify which region of the image you want to track by drawing a polygon on the reference frame. Typically, when you first add a Planar Tracker node, you are immediately ready to start drawing a polygon in the viewer, so it's best to do this right away. When choosing where to draw a polygon, make sure the region selected belongs to a physically planar surface in the shot. In a pinch, a region that is only approximately planar can be used, but the less planar the surface, the poorer the quality of the resulting track.

As a rule of thumb, the more pixels in the pattern, the better the quality of the track. In particular, this means the reference frame pattern should be:

- As large as possible.
- As much in frame as possible.
- As unoccluded as possible by any moving foreground objects.
- At its maximal size (e.g., when tracking an approaching road sign, it is good to pick a later frame where it is 400 x 200 pixels big rather than 80 x 40 pixels).
- Relatively undistorted (e.g., when the camera orbits around a flat stop sign, it is better to pick a frame where the sign is face-on parallel to the camera rather than a frame where it is at a highly oblique angle).

If the pattern contains too few pixels or not enough trackable features, this can cause problems with the resulting track, such as jitter, wobble, and slippage. Sometimes dropping down to a simpler motion type can help in this situation.

After you've drawn a pattern, a set of Pattern parameters lets you transform and invert the resulting polygon, if necessary.

Track Mode

Track mode is unlike the other three options in the Operation menu in that is the only option that initiates the planar tracking. The other modes use the tracking data generated by the Track mode.

Tracker

There are two available trackers to pick from:

- **Point:** Tracks points from frame to frame. Internally, this tracker does not actually track points-per-se but rather small patterns like Fusion's Tracker node. The point tracker possesses the ability to automatically create its internal occlusion mask to detect and reject outlier tracks that do not belong to the dominant motion. Tracks are colored green or red in the viewer, depending on whether the point tracker thinks they belong to the dominant motion or they have been rejected. The user can optionally supply an external occlusion mask to further guide the Point tracker.
- **Hybrid Point/Area:** Uses an Area tracker to track all the pixels in the pattern. Unlike the Point tracker, the Area tracker does not possess the ability to automatically reject parts of the pattern that do not belong to the dominant motion, so you must manually provide it with an occlusion mask. Note that for performance reasons, the Hybrid tracker internally first runs the Point tracker, which is why the point tracks can still be seen in the viewer.

  There is no best tracker. They each have their advantages and disadvantages:

  Artist Effort (occlusion masks): The Point tracker will automatically create its internal occlusion mask. However, with the Hybrid tracker, you need to spend more time manually creating occlusion masks.

  Accuracy: The Hybrid tracker is more accurate and less prone to wobble, jitter, and drift since it tracks all the pixels in the pattern rather than a few salient feature points.

  Speed: The Hybrid tracker is slower than the Point tracker.

In general, it is recommended to first quickly track the shot with the Point tracker and examine the results. If the results are not good enough, then try the Hybrid tracker.

Motion Type

Determines how the Planar Tracker internally models the distortion of the planar surface being tracked. The five distortion models are:

- Translation.
- Translation, Rotation (rigid motions).
- Translation, Rotation, Scale (takes squares to squares, scale is uniform in x and y).
- Affine - includes translation, rotation, scale, skew (maps squares to parallelograms).
- Perspective (maps squares to generic quadrilaterals).

Each successive model is more general and includes all previous models as a special case.

When in doubt, choose Perspective for the initial track attempt. If the footage being tracked has perspective distortions in it, and the Planar Tracker is forced to work with a simpler motion type, this can end up causing the track to slide and wobble.

Sometimes with troublesome shots, it can help to drop down to a simpler motion model—for example, when many track points are clustered on one side of the tracked region or when tracking a small region where there are not many trackable pixels.

Output

Controls what is output from the Planar Tracker node while in the Track operation mode.

- **Background:** Outputs the input image unchanged.
- **Background - Preprocessed:** The Planar Tracker does various types of preprocessing on the input image (e.g., converting it to luma) before tracking. It can be useful to see this when deciding which track channel to choose.
- **Mask:** Outputs the pattern as a black and white mask.
- **Mask Over Background:** Outputs the pattern mask merged over the background.

Track Channel

Determines which image channel in the background image is tracked. It is good to pick a channel with high contrast, lots of trackable features, and low noise. Allowed values are red, green, blue, and luminance.

Tracking Controls

These controls are used to control the Tracker. Note that while tracking, only track to a new frame if the current frame is already tracked or it is the reference frame.

- **Track to start:** Tracks from the current frame backward in time to the start (as determined by the current render range).
- **Step tracker to previous frame:** Tracks from the current frame to the previous frame.
- **Stop tracking:** Stops any ongoing tracking operations.
- **Step tracker to next frame:** Tracks from the current frame to the next frame.
- **Track to end:** Tracks from the current frame forward in time to the end (as determined by the current render range).
- **Trim to start:** Removes all tracking data before the current frame.
- **Delete:** Deletes all tracking data at all times. Use this to destroy all current results and start tracking from scratch.
- **Trim to end:** Removes all tracking data after the current frame. This can be useful, for example, to trim the end of a track that has become inaccurate when the pattern starts to move off frame.

Show Splines

This button to the right of the “Trim to end” button opens the Spline Editor and shows the splines associated with the Planar Tracker node. This can be useful for manually deleting points from the Track and Stable Track splines.

Right-Click Here for Track Spline

While tracking, a spline containing 4 x 4 matrices at each keypoint is created. This is known as the “Track spline” or just “Track” for short. These matrices completely describe the distortions of the tracked pattern.

Create Planar Transform

After tracking footage, this button can be pressed to create a Planar Transform node on the Node Editor. The current tracking data is embedded in the Planar Transform node so that it can replicate the planar distortions tracked by the Planar Tracker node. Unless you are compositing a full frame foreground that matches the same dimensions as the raster, it is best to create a Planar Transform and use it to apply motion to the foreground.

Steady Mode

In Steady mode, the Planar Tracker transforms the background plate to keep the pattern as motionless as possible. Any leftover motion is because the Planar Tracker failed to follow the pattern accurately or because the pattern did not belong to a physically planar surface.

Steady mode is not very useful for actual stabilization, but is useful for checking the quality of a track. If the track is good, during playback the pattern should not move at all while the rest of the background plate distorts around it. It can be helpful to zoom in on parts of the pattern and place the mouse cursor over a feature and see how far that feature drifts away from the mouse cursor over time.

Steady Time

This is the time where the pattern's position is snapshotted and frozen in place. It is most common to set this to the reference frame.

Invert Steady Transform

This causes the Planar Tracker node to reverse the effects of the steady transform. This means two Planar Tracker nodes connected back to back with the second set to invert should give back the original image. If you place an effects node in between the two, then the effect will be locked in place. This should only be used to accomplish effects that cannot be done through corner pinning, since it involves two resamplings, causing softening of the background image.

Clipping Mode

Determines what happens to the parts of the background image that get moved off frame by the steady transform:

- **Domain:** The off frame parts are kept.
- **Frame:** The off frames parts are thrown away.

> Domain mode is useful when Steady mode is being used to "lock" an effect to the pattern. As an example, consider painting on the license plate of a moving car. One way to do this is to use a Planar Tracker node to steady the license plate, then a Paint node to paint on the license plate, and then a second Planar Tracker to undo the steady transform. If the Clipping mode is set to Domain, the off frame parts generated by the first Planar Tracker are preserved so that the second Planar Tracker can, in turn, map them back into the frame.

Corner Pin Mode

In Corner Pin mode, one or more textures can be attached to a previously tracked planar surface and undergo the same perspective distortions as the surface.

The corner pin workflow with Planar Tracker is:

1. Track: select a planar surface in the shot that you wish to attach a texture to or replace the texture on. Track the shot (see the tracking workflow in the Track section).
2. Switch the Operation Mode to Corner Pin: When Corner Pin mode is entered from Track mode, the pattern polygon is hidden and a corner pin control is shown in the viewer.
3. Connect in the texture: In the Node Editor, connect the output of the MediaIn node containing the texture to the Corner Pin 1 input on the Planar Tracker node.
4. Adjust corner pin: Drag the corners of the corner pin in the viewer until the texture is positioned correctly. Sometimes the Show Grid option is useful when positioning the texture. Additionally, if it helps to position it more accurately, scrub to other times and make adjustments to the corner pin.
5. Review: Play back the footage and make sure the texture "sticks" to the planar surface.

Merge Mode

Controls how the foreground (the corner pinned texture) is merged over the background (the tracked footage). If there are multiple corner pins, this option is shared by all of them. There are four options to pick from:

- BG only
- FG only
- FG over BG
- BG over FG

Number of Corner Pins

Use the + and - buttons to increase or decrease the number of corner pins. Each time an additional corner pin is created, a corresponding input appears on the node in the Node Editor.

Corner Pin 1 Input Group

Each corner pin has a group of related inputs:

- **Enable:** Controls the visibility of the corner pin in the viewer.
- **Show Grid:** Shows a grid over the corner pin. This can be useful when positioning the corners.
- **Merge Options:** Controls merging of corner pin texture over the background. (See the documentation for the Merge node.)
- **Reference Time Positions:** The positions of the four corners at the reference time. If the track was not perfect, these positions can be animated to make adjustments on top of the track.

Stabilize Mode

Stabilize mode is used to smooth out shakiness in the camera by applying a transform that partially counteracts the camera shake. This stabilizing transform (contained in the Stable Track spline) is computed by comparing neighboring frames.

> **NOTE:** Stabilize mode only smooths out motions, while Steady mode tries to completely “lock off” all motion.

Be aware that the Planar Tracker stabilizes based on the motion of the pattern, so it is essential to choose the pattern carefully. If the motion of the pattern does not represent the motion of the camera, then there may be unexpected results. For example, if tracking the side of a moving truck and the camera is moving alongside it, the Planar Tracker smooths the combined motion of both the truck and the mounted camera. In some cases, this is not the desired effect. It may be better to choose the pattern to be on some fixed object like the road or the side of a building, which would result in smoothing only the motion of the camera.

One unavoidable side effect of the stabilization process is that transparent edges appear along the edges of the image. These edges appear because the stabilizer does not have any information about what lies off frame, so it cannot fill in the missing bits. The Planar Tracker node offers the option to either crop or zoom away these edges. When filming, if the need for post-production stabilization is anticipated, it can sometimes be useful to film at a higher resolution (or lower zoom).

The Planar Tracker Stabilization Workflow

1. Track: Select a roughly planar region that represents the motion that you want to stabilize. Track the shot (see the tracking workflow in the Track section).
2. Switch the Operation Mode to Stabilize: Until a stabilization is computed, the Planar Tracker will just output the input footage.
3. Adjust stabilization options: Frequently changed options are Parameters to Smooth and Smoothing Radius.
4. Compute stabilization: Press the Compute Stabilization button and wait for the stabilization computations to finish. Play back the output of the Planar Tracker node to see the effects of the stabilization. Notice that transparent edges have been introduced around the edges of the image by the stabilization transform.
5. Refine: Adjust the stabilization options and recompute the stabilization as many times as desired.
6. Handle transparent edges (optional): Set the Frame Mode to either Zoom or Crop as desired and then click the Auto Zoom or Auto Crop button. Play back the footage to observe the effects. If there is too much zoom or the image has been cropped too small, try reducing the amount of smoothing.

Parameters to Smooth

Specify which of the following parameters to smooth:

- X Translation
- Y Translation
- Rotation
- Scale

Smoothing Window

When stabilizing a particular frame, this determines how the contributions of neighboring frames are weighted. Available choices are Box and Gaussian.

Smoothing Radius (Frames)

Determines the number of frames whose transforms are averaged together to compute the stabilization. A larger Smoothing Radius results in more smoothing but introduces more transparent edges.

Compute Stabilization

Clicking this button runs the stabilizer, overwriting the results of any previous stabilization. As soon as the stabilization is finished, the output of the Planar Tracker node will be immediately updated with the stabilization applied.

> **NOTE:** The stabilizer uses the Track spline (created by the tracker) to produce the Stable Track spline. Both of these splines' keyframes contain 4 x 4 matrices, and the keyframes are editable in the Spline Editor.

Clipping Mode

Determines what happens to the parts of the background image that get moved off frame by the stabilization:

- **Domain:** The off frame parts are kept.
- **Frame:** The off frames parts are thrown away.

Frame Mode

This controls how transparent edges are handled. The available options include:

- **Full:** Do nothing. Leaves the transparent edges as is.
- **Crop:** Crops away the transparent edges. When this option is selected, the size of Planar Tracker's output image is smaller than the input image. No image resamplings occur. In Crop mode, use the Auto Crop button or manually adjust the crop window by changing the X Offset, Y Offset, and Scale sliders.

  Auto Crop Button: When this button is clicked, the Planar Tracker will examine all the frames and pick the largest possible crop window that removes all the transparent edges. The computed crop window will always be centered in frame and pixel aligned. When clicked, Auto Crop updates the X/Y Offset and Scale sliders.

- **Zoom:** Scales the image bigger until the transparent edges are off frame. Choosing this option causes an image resampling to occur. The downside of this approach is that it reduces the quality (slightly softens) of the output image. In Zoom mode, use the Auto Zoom button or manually adjust the zoom window by changing the X Offset, Y Offset, and Scale sliders.

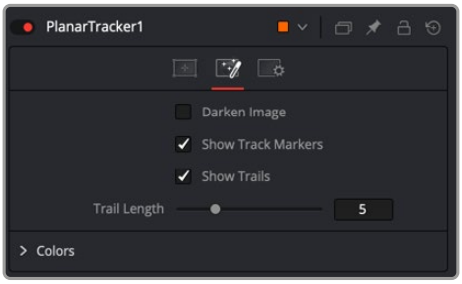
  Auto Zoom: When this button is clicked, the Planar Tracker will examine all the frames and pick the smallest possible zoom factor that removes all the transparent edges. The computed zoom window will always be centered in frame. When clicked, Auto Zoom updates the X/Y Offset and Scale sliders.

Right-Click Here for Stable Track Spline

Right-clicking over this label provides access to a spline whose keyframes contain 4 x 4 matrices which in turn represent the stabilization transforms. This is mostly here for completeness and for advanced users.

Options Tab

These controls affect the look of onscreen controls in the viewer.



The Planar Tracker Options tab

Darken Image

Darkens the image while in Track mode in order to better see the controls and tracks in the viewer. The Shift+D keyboard shortcut toggles this.

Show Track Markers

Toggles the display of the dots marking the location of trackers at the current time.

Show Trails

Toggles the display of the trails following the location of trackers.

Trail Length

Allows changing the length of tracker trails. If the pattern is moving very slowly, increasing the length can sometimes make the trails easier to follow in the viewer. If the pattern is moving very fast, the tracks can look like spaghetti in the viewer. Decreasing the length can help.

Inlier/Outlier Colors

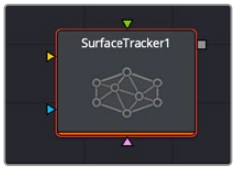
When tracking, the tracker analyzes the frame and detects which of the multitudinous tracks belong to the dominant motion and which ones represent anomalous, unexplainable motion. By default, tracks belonging to the dominant motion are colored green (and are called inliers) and those that do not belong are colored red (and are called outliers). Only the inlier tracks are used when computing the final resulting track.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Tracking nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Surface Tracker [SFt]



The Surface Tracker node

Surface Tracker Node Introduction

The Surface Tracker is used for applying textures and effects to flexible or deformable moving surfaces, like cloth and skin. The Surface Tracker ensures that the folding and deforming of the texture or effect is realistically matched to the folding and deforming of the surface that it is attached to in the image.

The Surface tracker is a complex effect, with many different options and controls to create the most realistic composites possible. But in all cases you simply move from left to right, using the buttons in the effect’s settings.

1. Click the Bounds button, and click on the screen to define the area on the surface where you wish to apply the new texture. The bounds should select one surface at a time and not go over/on object edges.
2. Click on the Mesh button to define the control points along the natural folds and curves of the surface within the bounds.
3. Click on the Track button to analyze the motion of the surface over time, using the mesh you just set.
4. Click on the Result button to composite the texture on top of the moving mesh.

Across all the tabs in the Surface Tracker, the Tracked Range shows the start and end frame numbers of your clip that contain tracking data. Any existing tracking data will be reset if you modify the bounds, or mesh layout, and you will need to re-track it.



The Surface Tracker in action, overlaying the logo (left), and warping and conforming it realistically to the folds of the shirt (right).

For more information on using the Surface Tracker, see Chapter 168, "Resolve FX Transform," in the *DaVinci Resolve 20 Reference Manual*.

Inputs

There are four inputs to the Surface Tracker node.

Input: The yellow input accepts the media with the surface that you want to track and composite onto.

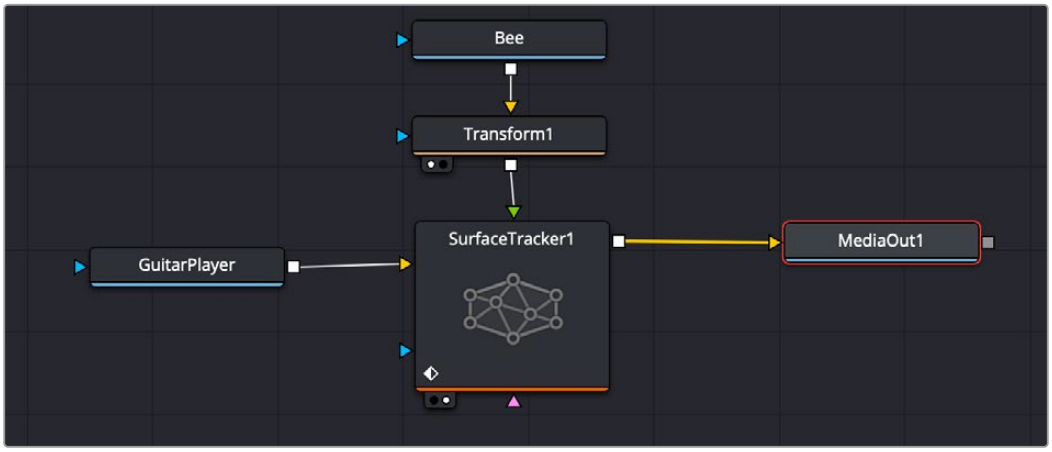
Overlay: The green input accepts still graphics, video, or the output of a node tree. This is the image that you want to overlay onto the surface.

Occlusion Mask: The pink occlusion mask input is used to mask out regions that do not need to be tracked. Regions where this mask is white will not be tracked. For example, a person moving in front of and occluding bits of the scene may be confusing to the tracker, and a quickly-created rough mask around the person can be used to tell the tracker to ignore the masked-out bits.

Effect Mask: The optional blue effect mask input accepts a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the filter to only those pixels within the mask. An effects mask is applied to the tool after the tool is processed.

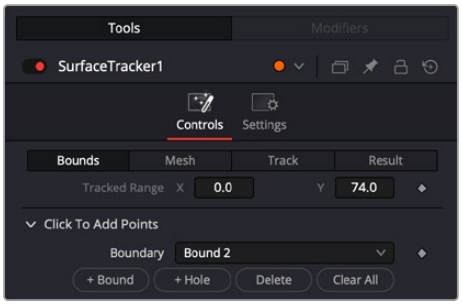
Basic Node Setup

Connect the media with the surface you want to composite onto to the main yellow input. Connect the media with the graphic or video that you want to overlay onto the surface to the green overlay input. Optionally, connect effect or occlusion masks to their respective inputs, if needed.



A simple node tree showing video clip of a guitar player, using the surface tracker to overlay a small graphic of a bee on his shirt.

Inspector



The Bounds tab in the Surface Tracker controls

Bounds

This is where you are prompted to draw an area on the image that you want to apply the texture to. Ideally, you want to choose a frame in the clip that has the surface you want to track as close and as flat-on to the camera as possible. This frame becomes the default reference frame for the effect. Simply click on the image to make multiple points around the area you want to define (a minimum of three points are required), and the effect will connect them automatically as a polygon. Boundary points should be on the surface for tracking, not around it. If points are accidentally placed on object edges or the background instead, it will cause errors in the tracking.

You may add multiple boundaries and multiple holes to the same image, and all boundaries will join together into one shape before subtracting the holes. The purpose of multiple boundaries and holes are to allow you to mask out complex surfaces with regions not for tracking. For example, you can

make a boundary around someone's face, but use multiple holes to exclude the eyes and mouth as those change appearance quickly as the subject blinks and talks, making them exceptionally difficult to track.

Multiple boundaries are designed to work on the same surface; using multiple boundaries on different surfaces within the image will cause poor results. Separate surfaces are best tracked by using multiple instances of the Surface Tracker instead.

The boundary is not useful for rotoscoping objects, only for analyzing the motion of a surface. Rotoscoping of foreground objects can be done as described in the Tracking Using a Mask section below.

- **Click to Add Points:** These are the controls used to make boundary areas on the image.

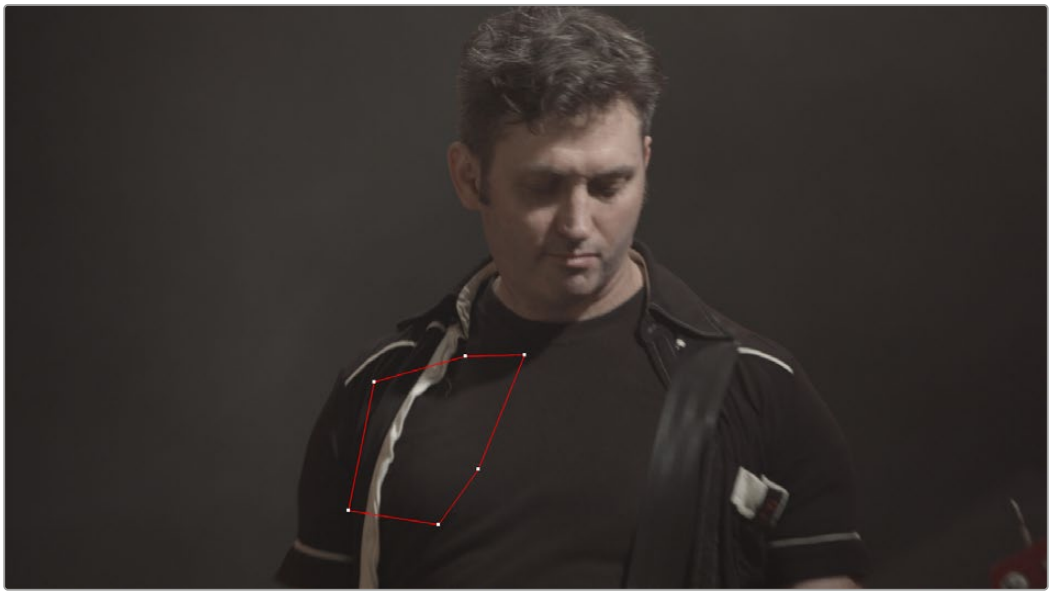
  Boundary: The currently selected boundary area on the image. You can change the current boundary or hole by using the drop-down menu. If there are issues with the boundary (i.e., crossed edges, too few points, etc.), an error message will appear next to the boundary name.

  + Bound: Adds a boundary polygon to the image. You can have multiple boundaries defined for the same surface, and they will be joined together into one shape.

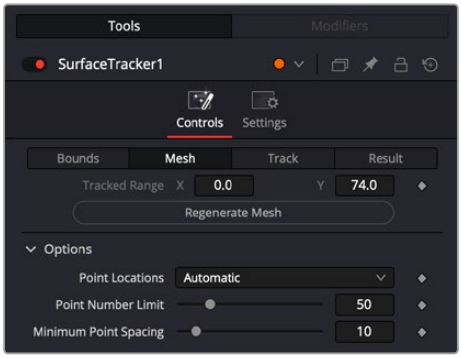
  + Hole: Adds an area to exclude tracking from the surface. Used to remove overly complex and changing parts of a surface from the effect's calculations.

  Delete: Lets you delete the boundary or hole selected in the Boundary field. Currently, this operation can not be undone, so be careful around this button.

  Clear All: Lets you delete all boundaries and holes in the effect.



A simple boundary drawn on his shirt in the viewer



The Mesh tab in the Surface Tracker controls

Mesh

With the boundaries now defined, you can establish the mesh over the surface to be tracked. The initial mesh serves as the starting point for tracking, so for best results, the mesh should be placed when your surface is as flat-on and closest to the camera as possible.

The mesh can be manually adjusted by dragging current points (multiple points can be moved by holding command and dragging an area around the points) or clicking on an area to place new ones. Option-clicking a point will remove it from the mesh. Manually adding mesh points can be useful if there are flexible portions of the surface known to be a point where the texture should fold but which is not initially identified when placing the mesh. For example, an extended arm may not have a mesh point active at the elbow, but with a mouse click (and a basic knowledge of human anatomy), you can manually add one there. Another use case would be to add points to an area with little detail, but you know will develop detail as it rotates into view. Placing mesh points in blank areas that show little detail will not improve results and may actually degrade them, so very few points on smooth areas is desirable. A good rule of thumb is that if more than five points are created inside the boundaries, the mesh may be good enough for tracking. Points cannot be added or deleted after tracking.

A good mesh is a sign of good tracking as they go hand in hand. If your mesh has few points detected and the results don't track tightly, you can artificially improve your results. For example, make sure that the content is clearly visible. Tracking and point detection can fail if the surface is too bright (HDR), too dark, or low contrast due to working in log. To improve results, you can temporarily add a brightness/contrast node before the Surface Tracker to adjust contrast manually, then run the Surface Tracker. Simply disable the brightness/contrast node again after the tracking has finished.

- **Regenerate Mesh:** This button recreates the mesh at the current frame and settings and discards any manual changes that have been made.

  Point Locations: These controls determine which mode is used to generate the mesh points.

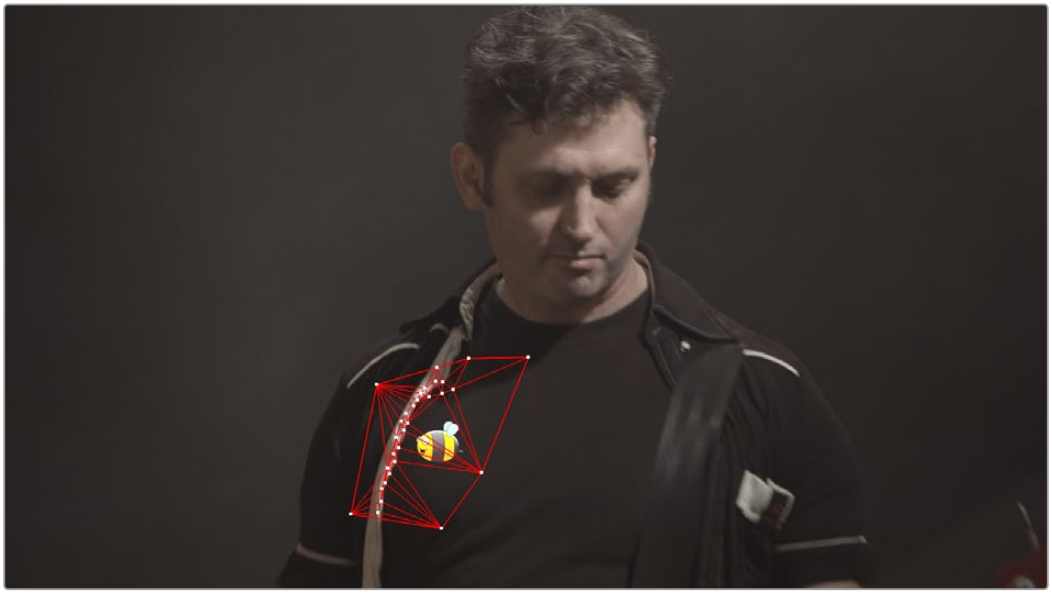
  Automatic: The Surface Tracker automatically locates the best details for tracking, and places mesh points on those details. In most cases, Automatic will give you the best results.

  Point Number Limit: Lets you set the maximum amount of mesh points to place. The number of mesh points affects how the overlay will flex and fold, and should be set by the expected flexibility of the surface. They do not directly affect tracking speed.

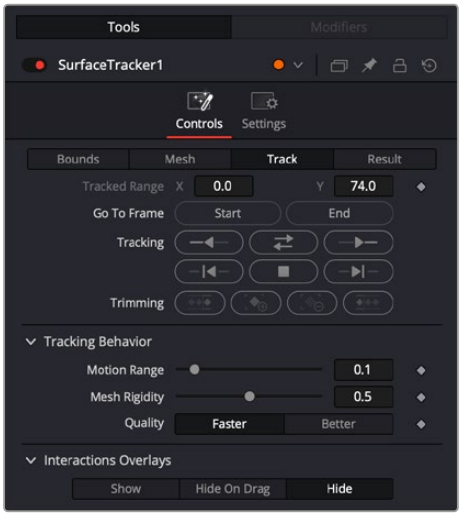
  Minimum Point Spacing: Lets you set the minimum distance between mesh points. This prevents mesh points from bunching together too tightly in highly detailed areas and interfering with each other.

  Uniform: Creates a regular grid of points in the boundary and relies on internal shape models to deform the grid during tracking. This may give better results than Automatic mode, if your surface is highly detailed.

  - **Horizontal Spacing:** Adjusts the horizontal distance between adjacent points in the uniform mesh.
  - **Vertical Spacing:** Adjusts the vertical distance between adjacent points in the uniform mesh.



The automatic mesh applied to the boundary. The bee is now composited onto the shirt, but it does not look very realistic yet.



The Track tab in the Surface Tracker controls

Track

This section of the effect tracks the motion of the mesh points over time. Use the Tracking Direction controls to complete the tracking procedure across the clip. After the tracking is complete, the effect does no more analysis, and any additional contrast assistance from previous nodes can be discarded after this step. For simple cases, the tracking may run completely without user intervention, but for more complex cases, a number of additional controls are used to refine the track.

Fixing a broken track will require some manual adjustments to both the mesh and the tracking data. At any time while tracking, mesh points can be adjusted either one by one, or with the lasso selection tool (Command - Drag) around the entire mesh or just parts of it. The lasso selection tool also has the ability to stretch, scale, and rotate the mesh points, making it easier to manually deform and place the mesh exactly where you want it to be, if for some reason the automatic tracking fails. However, additional mesh points cannot be added or deleted at this stage, without first resetting the tracking.

Manually editing the mesh points at a frame only affects that frame directly; it does not affect neighboring tracking data. However, it does affect neighboring interpolated positions, which makes it good for tracking from that point onwards. For example:

- Tracking is good until frame 150 when a few points separate from the rest, disturbed by a lighting change. At frame 150 they can be manually moved back into place, then if tracking is resumed they will continue with their corrected positions.
- Tracking exists until frame 150, after which there is no data for a few frames until it resumes at 160. Editing the points at either frames 150 or 160 will cause the intermediate frames to be interpolated with these changes applied.
- Tracking exists up to frame 150 and at frame 160 onwards, frames 151-159 are interpolated. At frame 155 some points are adjusted. This creates tracking data at this frame, and now frames 151-154 interpolate from the last tracked frame to this current one, and then frames 156-159 from the current to the next tracked frame.

Go to Start: This button will snap the playhead to the first frame that has been tracked, or the frame where the boundary was created.

Go To End: This button will snap the playhead to the last frame that has been tracked.

Tracking Direction Controls: These controls perform the tracking analysis in the direction indicated. They are from left to right: Track Backwards, Track Backwards One Frame, Pause Tracking, Track Forwards Then Reverse, Track Forwards One Frame, Track Forwards. Using these controls, you can set the track across the whole clip or limit the tracking to just a section of it.

Trimming: These controls allow you to manipulate the track data after it's been run.

- **Remove Tracking Data Before This Frame:** Clears all the tracking data before the current frame.
- **Track From Last Known Match:** This button will match the tracking data to the current frame and connect it to the last frame where tracking data exists. If there has been too much motion between those times, you can move the points approximately into place and try matching again. For example, the tracking occurs without issue until frame 150, but something moving across the scene causes an error in the track for six frames (151-157). The tracking data can be paused before it hits these frames, or these frame's tracking data can be deleted after the fact. Then at frame 158 you can press this button to continue the track from the last known good track at frame 150.
- **Remove Tracking Data At This Frame:** Clears the tracking data at the current frame.
- **Remove Tracking Data After This Frame:** Clears all the tracking data after the current frame.

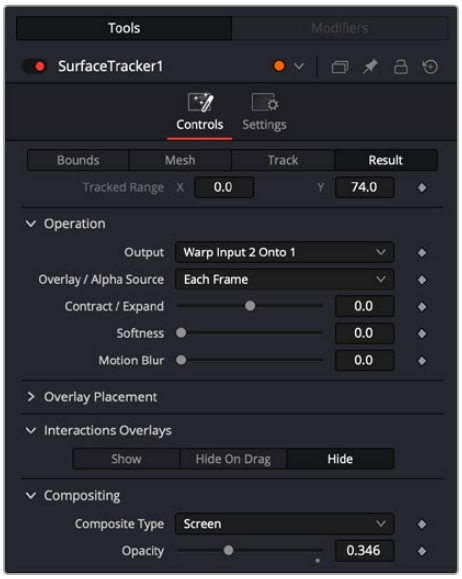
Tracking Behavior: These options allow you adjust the way the tracking operates.

- **Motion Range:** Controls how far to search for matches. Increase this slider if the surface is moving too quickly and the tracking fails to find it.
- **Mesh Rigidity:** Controls how flexible the mesh is. Reduce this slider to make the mesh points more flexible. This gives you best results if the clip is noise-free and warping smoothly. Increase this slider to get a "stiffer" tracked surface with less internal warping.
- **Quality:** The tracker can be very computationally intensive; it is recommended to use the Faster setting while making adjustments to the mesh, then use Better for the final track once all the parameters have been set correctly.

Interaction Overlays: These buttons control the visibility of the mesh overlays on the Viewer.

- **Show:** Shows the mesh at all times.
- **Hide on Drag:** Shows the mesh until a point is dragged, then it is hidden, allowing you to see the surface underneath for fine adjustments.
- **Hide:** Hides the mesh at all times.

NOTE: The Surface Tracker creates data at each frame of the source clip. If the source clip is not at the same frame rate as the timeline, tracking may seem to stutter or skip frames. You can "bake" retiming into a clip by right-clicking on it on the Edit page and choosing "Render In Place" to create a result with exactly one frame for every timeline frame.



The Result tab in the Surface Tracker controls

Result

Once the tracking data is complete, the Result section lets you determine what and how to composite a video layer on top of the moving surface.

Output: There are six options for what to do with the resulting tracking data.

- **Warp Input 2 Onto 1:** Warps the Surface Tracker's overlay input to the shape and position of the tracked mesh and composites it onto the clip.
- **Warp Input 2 Onto Blank:** Warps the Surface Tracker's overlay input to the shape and position of the tracked mesh and passes the Alpha channel information, allowing you to perform the actual composite in another node or effect.
- **Warp Input 1's Alpha:** Applies the results of the Surface Tracker to the Effect Mask input.
- **Create Mesh Alpha Mask:** Produces an Alpha mask, which follows the position and shape of the outline of the mesh. All inputs and alphas are ignored.
- **Stabilize-Warp Input 1:** Performs an inverse warp to the Reference Frame (again, this should be a frame that ideally is as close to and flat-on to the camera as possible). The result is a stationary image within the boundary with all the lighting changes and details of the surface's motion shown on it. This now stabilized image can be modified with any of DaVinci Resolve's existing tools, bypassing any issues caused by this area's motion in the original video clip. Once modified, you warp the result back onto the original footage by using a copy of the original Surface Tracker node but with the Rewarp Stabilized Clip setting here instead.

- **Rewarp Stabilized Clip:** Composites the output of the Stabilize-Warp Input 1 selection back onto the original video clip, after it's been modified. This option is designed to be used in conjunction with a copy of a Surface node using Stabilize-Warp Input 1. It seamlessly blends whatever additional modifications that were made between the two Surface Trackers back together.

Overlay/Alpha Source: Allows you to select which frames of the overlay or alpha are used to create the result.

- **First Frame:** Useful for freezing content from an overlay clip to use a still frame as a static texture.
- **Reference Frame:** Starts at the currently defined reference frame (see below), instead of at frame 0. The reference frame is the frame which the mesh overlay is initially positioned in its "un-warped" state, and the positioning controls are active.
- **Each Frame:** Treats the overlay or alpha as a clip, warping each frame in turn. This is necessary for animated overlays or re-warping stabilized content.

Contract/Expand: Allows you to add or remove pixels from the edges of the overlay or alpha. It is important to note that the expanded area is not tracked and will not be stabilized to the content.

Softness: Blurs the edge of the Expand/Contract tool to blend more realistically into the scene.

Motion Blur: This slider lets you add motion blur to the overlay to preserve the illusion that it was actually captured with the scene.

Overlay Placement: These controls let you set the initial transform options for the image you want to overlay on the mesh.

- **Go To Reference:** Moves the playhead to the selected reference frame.
- **Set to Current:** Sets the current frame under the playhead as the reference frame.
- **Reference Frame:** Shows the frame number at which the overlay is originally positioned and un-warped (the point from which the warp to every other frame is calculated). The positioning controls appear on screen when the clip is at the reference frame. A frame where the surface's view is largest and least warped is the best starting point for rendering onto other frames.
- **Positioning:** Lets you choose how to adjust the overlay position.

  Sliders: Choosing this option exposes a set of transform sliders (X/Y position, zoom, rotate, etc.) that let you manipulate the overlay's initial position.

  Interactive Canvas: Choosing this option exposes draggable frame controls in the Viewer, allowing you to transform the overlay using mouse controls. The Viewers shows a large grid with 9 regions; you can drag in any of them to move, skew, or warp the overlay.

  Interactive Pins: Adjusting the image in this mode is done by manually placing control points, called pins, in the Timeline Viewer. Adding one pin only gives you position control. At least two points are required for scaling and rotation. Dragging on one of the pins scales or rotates the image around the other pin. Using three pins, you can create perspective distortions by dragging any one of the pins. You can add up to four pins for unique corner pinning distortions.
- **Reset Position:** Resets the on-screen overlay back to its default state.

Interaction Overlays: These buttons control the visibility of the mesh overlays on the Viewer.

- **Show:** Shows the mesh at all times.
- **Hide on Drag:** Shows the mesh until a point is dragged, then it is hidden, allowing you to see the surface underneath for fine adjustments.
- **Hide:** Hides the mesh at all times.

Compositing: Lets you set the composite blend mode for the overlay operation. For more detailed information on each blend mode, see Chapter 50, “Compositing and Transforms in the Timeline.”

- **Composite Type:** A drop-down menu of all the possible composite operations.
- **Opacity:** This slider controls the transparency of the overlay.



The completed surface tracking composites the bee graphic onto his shirt. The composite type was screen, and the opacity reduced to blend into the shirt. The bee moves realistically as the shirt folds and deforms as he plays.

Tracker [TRA]



The Tracker node

Tracker Node Introduction

The Tracker is used to detect and follow one or more pixel patterns across frames in moving video. The tracking data can then be used to control the position or values of other nodes in the composition (for example, the center of a Light Rays node). Additionally, trackers can be used to stabilize an image or to apply destabilization to one image based on the motion of another.

> For more information, *see Chapter 22, "Using the Tracker Node,"* in the Fusion Reference Manual.

Inputs

The Tracker has three inputs:

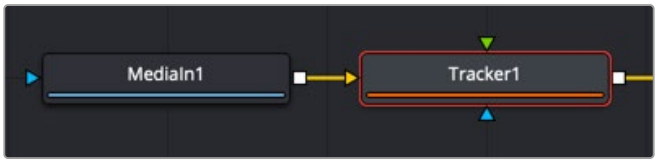
Background: The orange image input accepts the main 2D image to be tracked.

Foreground: The optional green foreground accepts a 2D image to be merged on top of the background as a corner pin or match move.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the tracking to certain areas.

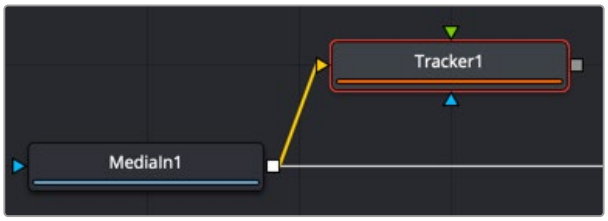
Basic Node Setup

Tracker nodes can be applied inline with other nodes or as a branch from the clip you want to track. When used inline, an image can be stabilized by connecting it to the orange background input. After the image is tracked, setting the Tracker's Operation menu to Match Move will apply a stabilization to the connected image.



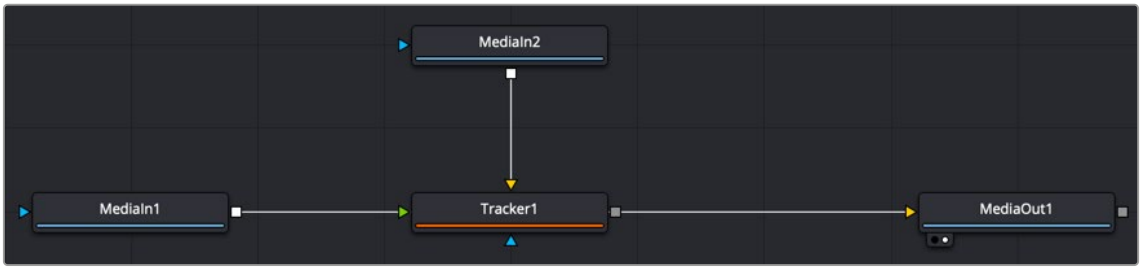
A Tracker node connected inline to an image for stabilization

When used as an offshoot from the node tree, an image can be tracked and then that tracking data is published for use on another node somewhere else in the node tree. The output of the Tracker does not need to connect to another node. The tracking data is published and can be used via the Connect To contextual menu.



A Tracker node branched from the node tree

The Tracker can also work as a replacement for a Merge tool in match-moving setups. Below, the Tracker tracks the image connected to the orange background input and applies the tracking data to the image connected to the foreground input. The same foreground-over-background merge capabilities are available in the Tracker node.



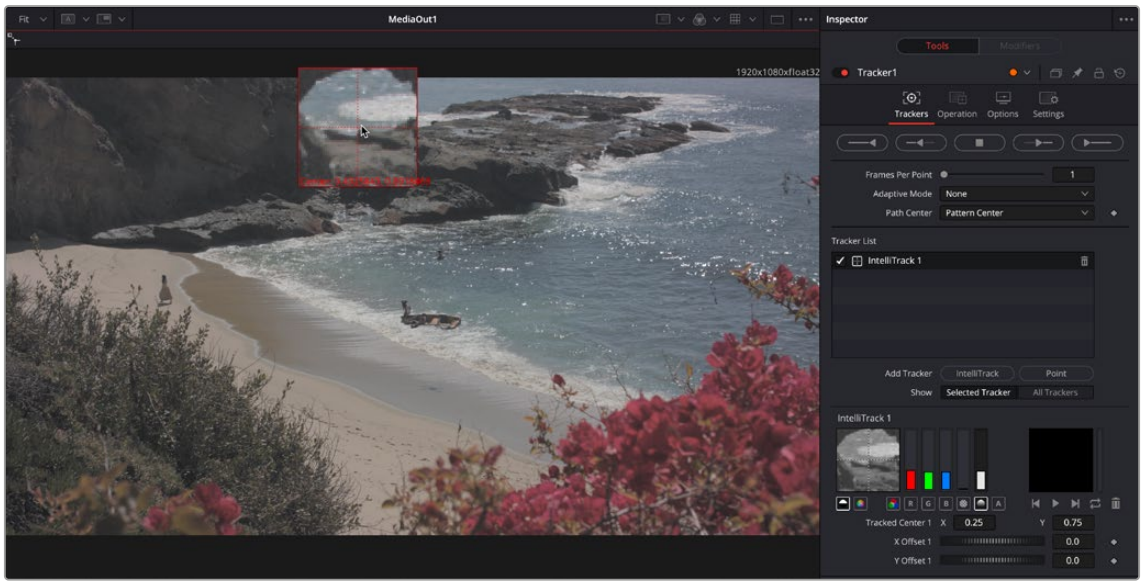
A Tracker node set up to apply a match move to the foreground input

IntelliTrack AI Point Tracker

The Fusion Tracker node now defaults to using IntelliTrack, the new DaVinci Neural Engine tracker. Using AI makes tracking easier and more precise, especially when tracking many points.

IntelliTrack, like other AI tools, doesn't use predetermined rulesets or heuristics. Instead, it is trained on real world examples. This makes it less likely to fail in scenarios like tracking a subject behind brief occlusions. For most cases it will be more precise and more robust than the normal point tracker.

All the IntelliTrack analysis happens behind the scenes, and there is no additional user input required, other than to choose the IntelliTrack setting, set the tracking crosshairs over high contrast areas on the same moving subject, and then press one of the tracker direction controls.



Adding a Tracking Point in the IntelliTrack mode in the Fusion Tracker

NOTE: The original Point tracker is still available to use by choosing the Point button.

Legacy Point Tracker Onscreen Controls

Though Fusion now defaults to IntelliTrack, if you wish to continue to use the old Point tracker, it is still available by clicking the Point button. Each pattern in the tracker has its own set of onscreen controls used to select the pixels in the image to be tracked. These controls are visible in the viewers whenever you select a tracker in the node tree.



Tracker onscreen controls

Pattern Rectangle

In the viewer, the tracker displays a solid-line red rectangle called the pattern rectangle. Every pixel within the rectangle makes up the pattern used for tracking. You can resize the pattern if necessary by dragging on the rectangle's borders.



A pattern rectangle identifies the area to track

Search Rectangle

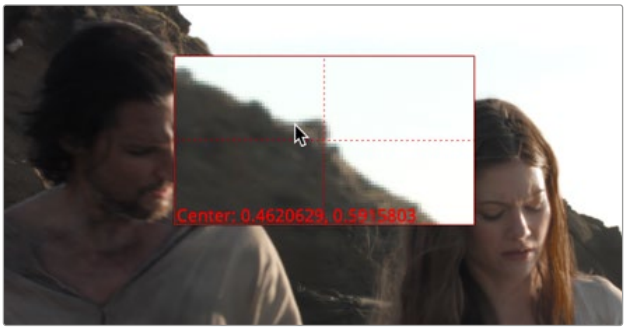
Whenever the mouse moves over the pattern rectangle, a second rectangle with a dashed outline appears. The dashed outline represents the search area, which determines how far away from the current pattern the Tracker looks in the next frame. The search area should always be larger than the pattern, and it should be large enough to encompass the largest frame-to-frame movement in the scene. Faster moving objects require larger search areas, and slower moving objects can get away with smaller search areas. The larger the search area, the longer it takes to track, so try not to make the search area larger than necessary. If the selected Tracker has a custom name, the name of that Tracker is displayed as a label at the bottom right of the search area rectangle.



The search rectangle is the area searched from frame to frame to locate the pattern.

Repositioning the Tracker

The pattern rectangle has a small handle in the upper-left corner. Dragging on the handle repositions the pattern. An enlarged view of the pattern is displayed under your mouse pointer to assist with the precise positioning of the pattern. This thumbnail disappears when the mouse button is released. You can adjust the magnification ratio in the Inspector's Options tab.

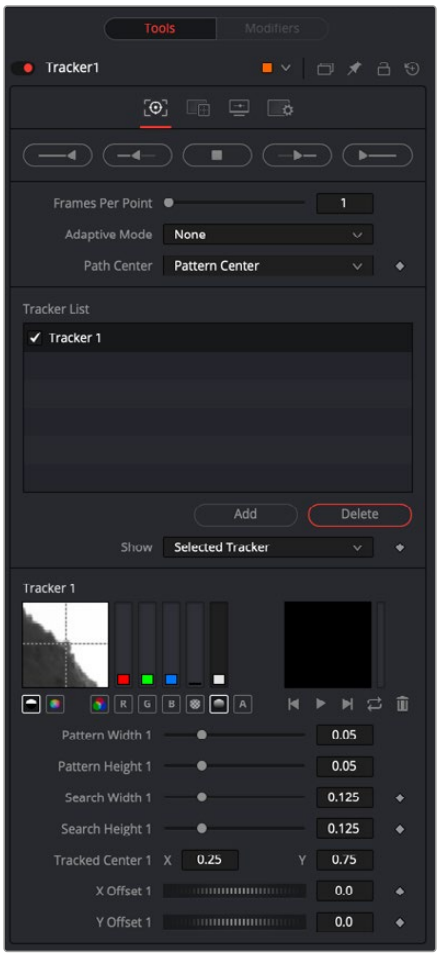


Dragging the handle magnifies the pattern rectangle for precise placement.

TIP: There is no limit to the number of trackers that can be used in one composition, or the number of objects that use the tracking data. There is also no limit to the number of patterns that can be tracked by a single Tracker node. This chapter serves as a reference for the various controls in the Tracker, but we strongly suggest you read the more general information *see Chapter 22, "Using the Tracker Node,"* in the Fusion Reference Manual.

The Tracker can be employed in two forms: as a node in the Node Editor or as a modifier attached to a parameter. When used as a node in the Node Editor, the image tracked comes from the input to the Tracker node. When used as a modifier, controls appear in the Modifiers tab for the node with the connected control. Tracker Modifiers can track only one pattern, but the image source can come from anywhere in the composition. Use this technique when tracking a quick position for an element.

Inspector



The Tracker node Trackers tab

Trackers Tab

The Trackers tab contains controls for creating, positioning, and initiating tracking operations. After tracking, offset controls are used to improve the alignment of the image following the track.

Track Buttons

There are four buttons to initiate tracking, and one button in the middle used to stop a track in progress. These buttons can track the current pattern forward or backward in time. Holding the pointer over each button displays a tooltip with the name of each button.

The buttons operate as follows:

- **Track Reverse:** Clicking this button causes all active trackers to begin tracking, starting at the end of the render range and moving backward through time until the beginning of the render range.
- **Track Reverse From Current Time:** Clicking this button causes all active trackers to begin tracking, starting at the current frame and moving backward through time until the beginning of the render range.
- **Stop Tracking:** Clicking this button or pressing ESC stops the tracking process immediately. This button is active only when tracking is in process.

- **Track Forward From Current Time:** Clicking this button causes all active trackers to begin tracking, starting at the current frame and moving forward through time until the end of the render range.
- **Track Forward:** Clicking this button causes all active trackers to begin tracking, starting at the first frame in the render range and moving forward through time until the end of the render range.

Tracking Behavior Controls

The following controls all affect how trackers adapt to changing patterns, how the resulting track path is defined, and how many keyframes are generated.

Frames Per Path Point

This slider determines how often the Tracker sets a keyframe on the path. The default is 1, which sets a keyframe on the tracked path at every frame.

Increasing the value causes the tracked path to be less accurate. This may be desirable if the track is returning fluctuating results, but under normal circumstances, leave this control at its default value.

> **TIP:** If the project is field rendered, a value of 1 sets a keyframe on every field. Since the Tracker is extremely accurate, this will result in a slight up-and-down jittering due to the position of the fields. For better results when tracking interlaced footage in Field mode, set the Frames Per Path Point slider to a value of 2, which results in one keyframe per frame of your footage.

Adaptive Mode

Fusion is capable of reacquiring the tracked pattern, as needed, to help with complex tracks. This menu determines the Adaptive tracking method.

- **None:** When set to None, the tracker searches for the original pattern in each frame.
- **Every Frame:** When set to Every Frame, the tracker reacquires the pattern every frame. This helps the Tracker compensate for gradual changes in profile and lighting over time.
- **Best Match:** When set to Best Match, the tracker compares the original selected pattern to the pattern acquired at each frame. If the variation between the two patterns exceeds the threshold amount defined by the Match Tolerance control, the tracker does not reacquire the pattern on that frame. This helps to avoid Tracker drift caused by transient artifacts that cross the pattern's path (such as a shadow).

Path Center

This menu determines how the Tracker behaves when repositioning a pattern. This menu is particularly useful when a pattern leaves the frame or changes so significantly that it can no longer be tracked.

- **Pattern Center:** When Pattern Center is selected in the menu, the tracked path continues from the center of the new path. This is appropriate when replacing an existing path entirely.
- **Track Center (append):** When Track Center (append) is selected in the menu, the path tracked by a new pattern will be appended to the existing path. The path created is automatically offset by the required amount. This setting is used to set a new tracking pattern when the original pattern moves out of the frame or gets obscured by other objects. This technique work bests if the new pattern is located close to the position of the original pattern to avoid any problems with parallax or lens distortion.

Tracker List

A Tracker node can track multiple patterns. Each tracker pattern created in the current Tracker node is managed in the Tracker List.

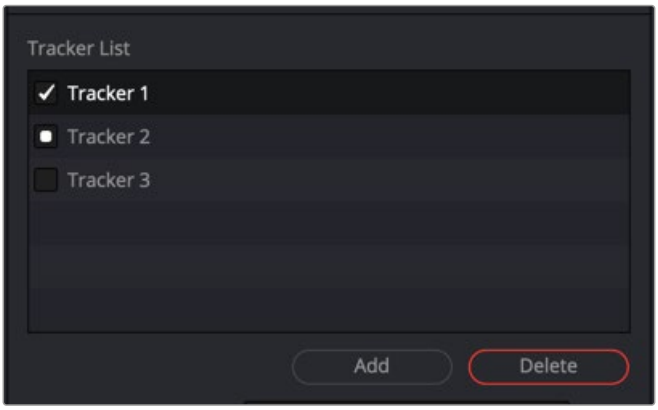
Tracker List

The Tracker List shows the names of all trackers created.

- Each tracker pattern appears in the list by name, next to a small checkbox. Clicking the name of the tracker pattern will select that tracker pattern.
- The controls below the list will change to affect that tracker pattern only. Click a selected tracker pattern once to rename the tracker pattern to something more descriptive.
- Clicking the checkbox changes the state of the tracker.

Tracker States

- **Enabled (black checkbox):** An enabled pattern will re-track each time the track is initiated. Its path data is available for use by other nodes, and the data is available for Stabilization and Corner Positioning.
- **Suspended (white circle):** A Suspended pattern does not re-track when the track is initiated. The data is locked to prevent additional changes. The data from the path is still available for other nodes, and the data is available for advanced Tracking modes like Stabilization and Corner Positioning.
- **Disabled (clear):** A Disabled pattern does not create a path when tracking is initialized, and its data is not available to other nodes or for advanced Tracking operations like Stabilization and Corner Positioning.



Three tracking patterns from top to bottom: enabled, suspended, and disabled

Add/Delete Tracker

Use these buttons to add or delete IntelliTrack or Point trackers from your Tracker List.

Show

This menu selects what controls are displayed in the Tracker node controls. They do not affect the operation of the tracker; they only affect the lower half of the Inspector interface.

- **Selected Tracker Details:** When Selected Tracker Details is chosen, the controls displayed pertain only to the currently selected tracker. You will have access to the Pattern window and the Offset sliders.
- **All Trackers:** When All Trackers is selected, the pattern window for each of the added tracking patterns is displayed simultaneously below the Tracker List.

Left Pattern Display

The pattern display has two side-by-side image windows and a series of status bars. The window on the left shows the pattern initially selected, while the one on the right shows a real-time display of the current pattern as tracking progresses.

As the onscreen controls move while tracking, the display in the leftmost window updates to show the pattern. As the pattern moves, the vertical bars immediately to the right of the image indicate the clarity and contrast of the image channels.

The best channel or channels get selected for tracking based on clarity. These channels have a gray background in the vertical bar representing that channel. You can use the automatic tracking or override the selection and choose the channel by selecting the button beneath the channel to track.



Tracker pattern display

Under normal circumstances, the channel selected shows in the pattern display. If the selected channel is blue, then a grayscale representation of the blue channel for the pattern appears. The image is represented in color only when you activate the Full Color button.

Override this behavior by selecting the Show Full Color button beneath the pattern display instead of the Show Selected Channel button.

> **TIP:** Because Fusion looks for the channel with the highest contrast automatically, you might end up tracking a noisy but high-contrast channel. Before tracking, it's always a good idea to zoom in to your footage and check the RGB channels individually.

Right Pattern Display

The pattern display on the right indicates the actual pattern acquired for tracking. This display is black until tracking the selected pattern for the first time. The pattern display becomes active during tracking, displaying the pattern that Fusion acquires from frame to frame.

As the tracking occurs, the pattern from each frame accumulates into a Flipbook, which can be played back in the pattern window after tracking by using the transport controls at the bottom of the window.

While the track is progressing, the vertical bar immediately to the right of the pattern shows how confident Fusion is that the current pattern matches the initially selected pattern. A green bar indicates a high degree of confidence that the current pattern matches the original, a yellow bar indicates less certainty, and a red bar indicates extreme uncertainty.

After tracking, the pattern display shows a small Flipbook of the track for that pattern to help identify problem frames for the track.

Tracker Sizes

In addition to onscreen controls, each tracker has a set of sizing parameters that let you adjust the pattern and search box.

- **Pattern Width and Height:** Use these controls to adjust the width and height of the selected tracker pattern manually. The size of the tracker pattern can also be adjusted in the viewer, which is the normal method, but small adjustments are often easier to accomplish with the precision of manual controls.
- **Search Width and Height:** The search area defines how far Fusion will look in the image from frame to frame to reacquire the pattern during tracking. As with the Pattern Width and Height, the search area can be adjusted in the viewer, but you may want to make small adjustments manually using these controls.

Tracked Center

This positional control indicates the position of the tracker's center. To remove a previously tracked path from a tracker pattern, right-click this parameter and select Remove Path from the contextual menu.

X and Y Offset

The Offset controls help to create a track for objects that may not provide very well defined or reliable patterns. The Offset controls permit the tracking of something close to the intended object instead. Use these Offsets to adjust the desired position of the path, while the tracker pattern rectangle is positioned over the actual tracking location.

The Offset can also be adjusted directly in the viewer by activating the Offsets button in the viewer toolbar.

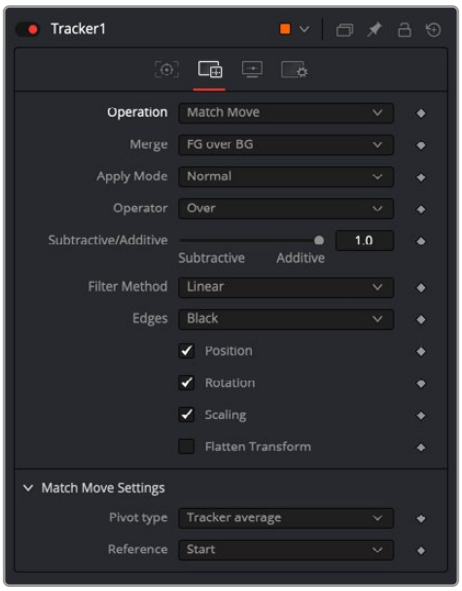


The tracker offset icon in the upper left of the viewer is used to offset the tracking pattern from the intended object.

Operation Tab

While the Trackers tab controls let you customize how the Tracker node analyzes motion to create motion paths, the Operation tab puts the analyzed motion data to use, performing image transforms of various kinds.

The Tracker node is capable of performing a wide variety of functions, from match moving an object into a moving scene, smoothing out a shaky camera movement, or replacing the content of a sign. Use the options and buttons in the Operation tab to select the function performed by the Tracker node.



The Tracker Operation tab

Operation Menu

The Operation menu contains four functions performed by the Tracker. The remaining controls in this tab fine-tune the result of this selection.

- **None:** The Tracker performs no additional operation on the image beyond merely locating and tracking the chosen pattern. This is the default mode, used to create a path that will then drive another parameter on another node.
- **Match Move:** When only the orange background input is connected, this mode stabilizes the image. When a foreground image is connected to the green foreground input, the foreground image matches the position, rotation, and scaling based on the tracking patterns. Stabilizing and match move require a minimum of one tracking pattern to determine position, and two or more to determine scaling and rotation.
- **Corner Positioning:** The Corner Positioner mode tracks the four corners of a rectangular object and replaces the contents with a new image. This function requires a minimum of four tracking patterns. If there are not enough tracking patterns, new tracking patterns are added until the total equals four.
- **Perspective Positioning:** This mode is the inverse of the Corner Positioning mode. Rather than replacing the contents of the rectangle, the four trackers are mapped to the four corners of the image. This is generally used to remove perspective from an image. Like the Corner Positioning mode, this mode requires four tracking patterns, which automatically get added if there are fewer patterns.

Additional Layering Controls

When you choose any operation other than None, a series of additional controls appear.

Merge

The Merge control determines what is done (if anything) with the image provided to the green Foreground input of the Tracker. This menu appears when the operation is set to anything other than None.

- **BG Only:** The foreground input is ignored; only the background is affected. This is used primarily when stabilizing the background image.
- **FG Only:** The foreground input is transformed to match the movement in the background, and this transformed image is passed through the Tracker's output. This Merge technique is used when match moving one layer's motion to another layer's motion.
- **FG Over BG:** The foreground image is merged over the background image, using the Merge method described by the Apply Mode control that appears.
- **BG Over FG:** The background is merged over the foreground. This technique is often used when tracking a layer with an Alpha channel so that a more static background can be applied behind it.

Apply Mode and Operator Menus

This menu provides a variety of options that determine how the two layers should be combined. The options in this menu are identical to those found in the Merge node.

- **Apply Modes:** The Apply Mode setting determines the math used when blending or combining the foreground and background pixels.

  Normal: The Default merge mode uses the foreground's alpha channel as a mask to determine which pixels are transparent and which are not. When this is active, another menu shows possible operations, including Over, In, Held Out, Atop, and XOr.

  Screen: Screen merges the images based on a multiplication of their color values. The alpha channel is ignored, and layer order becomes irrelevant. The resulting color is always lighter. Screening with black leaves the color unchanged, whereas screening with white will always produce white. This effect creates a similar look to projecting several film frames onto the same surface. When this is active, another menu shows possible operations, including Over, In, Held Out, Atop, and XOr.

  Dissolve: Dissolve mixes two image sequences together. It uses a calculated average of the two images to perform the mixture.

  Multiply: Multiplies the values of a color channel. This will give the appearance of darkening the image as the values are scaled from 0 to 1. White has a value of 1, so the result would be the same. Gray has a value of 0.5, so the result would be a darker image or, in other words, an image half as bright.

  Overlay: Overlay multiplies or screens the color values of the foreground image, depending on the color values of the background image. Patterns or colors overlay the existing pixels while preserving the highlights and shadows of the color values of the background image. The background image is not replaced but is mixed with the foreground image to reflect the original lightness or darkness of the background image.

  Soft Light: Soft Light darkens or lightens the foreground image, depending on the color values of the background image. The effect is similar to shining a diffused spotlight on the image.

  Hard Light: Hard Light multiplies or screens the color values of the foreground image, depending on the color values of the background image. The effect is similar to shining a harsh spotlight on the image.

  Color Dodge: Color Dodge uses the foreground's color values to brighten the background image. This is similar to the photographic practice of dodging by reducing the exposure of an area of a print.

  Color Burn: Color Burn uses the foreground's color values to darken the background image. This is similar to the photographic practice of burning by increasing the exposure of an area of a print.

  Darken: Darken looks at the color information in each channel and selects the background or foreground image's color value, whichever is darker, as the result color. Pixels lighter than the merged colors are replaced, and pixels darker than the merged color do not change.

Lighten: Lighten looks at the color information in each channel and selects the background or foreground image's color values, whichever is lighter, as the result color value. Pixels darker than the merged color are replaced, and pixels lighter than the merged color do not change.

Difference: Difference looks at the color information in each channel and subtracts the foreground color values from the background color values or the background from the foreground, depending on which has the greater brightness value. Merging with white inverts the color. Merging with black produces no change.

Exclusion: Exclusion creates an effect similar to but lower in contrast than the Difference mode. Merging with white inverts the base color values. Merging with black produces no change.

Hue: Hue creates a result color with the luminance and saturation of the background color values and the hue of the foreground color values.

Saturation: Saturation creates a result color with the luminance and hue of the base color and the saturation of the blend color.

Color: Color creates a result color with the luminance of the background color value and the hue and saturation of the foreground. This preserves the gray levels in the image and is useful for coloring monochrome images.

Luminosity: Luminosity creates a color with the hue and saturation of the background color and the luminance of the foreground color. This mode creates an inverse effect from that of the Color mode.

- **Operator Modes:** This menu is used to select the Operation Mode of the merge. It determines how the foreground and background are combined to produce a result. This drop-down menu is visible only when the Merge node's Apply Mode is set to either Normal or Screen.

> **NOTE:** For an excellent description of the math underlying the Operation modes, read "Compositing Digital Images," Porter, T., and T. Duff, SIGGRAPH 84 proceedings, pages 253-259. Essentially, the math is as described below.

> **TIP:** Some modes not listed in the Operator drop-down menu (Under, In, Held In, Below) are easily obtained by swapping the foreground and background inputs and choosing a corresponding mode.

The formula used to combine pixels in the merge is always fg * x + bg * y. The different operations determine exactly what x and y are, as shown in the description for each mode.

The Operator modes are as follows:

- **Over:** The Over mode adds the foreground layer to the background layer by replacing the pixels in the background with the pixels from the Z wherever the foreground's alpha channel is greater than 1.

```
x = 1, y = 1-[foreground Alpha]
```

- **In:** The In mode multiplies the alpha channel of the background input against the pixels in the foreground. The color channels of the foreground input are ignored. Only pixels from the foreground are seen in the final output. This essentially clips the foreground using the mask from the background.

```
x = [background Alpha], y = 0
```

- **Held Out:** Held Out is essentially the opposite of the In operation. The pixels in the foreground image are multiplied against the inverted alpha channel of the background image. You can accomplish exactly the same result using the In operation and a Matte Control node to invert the matte channel of the background image.

  ```
  x = 1-[background Alpha], y = 0
  ```

- **ATop:** ATop places the foreground over the background only where the background has a matte.

  ```
  x = [background Alpha], y = 1-[foreground Alpha]
  ```

- **XOr:** XOr combines the foreground with the background wherever either the foreground or the background have a matte, but never where both have a matte.

  ```
  x = 1-[background Alpha], y = 1-[foreground Alpha]
  ```

- **Subtractive/Additive:** This slider controls whether Fusion performs an Additive merge, a Subtractive merge, or a blend of both. This slider defaults to Additive merging for most operations, assuming the input images are premultiplied (which is usually the case). If you don't understand the difference between Additive and Subtractive merging, here's a quick explanation:
  - An Additive merge is necessary when the foreground image is premultiplied, meaning that the pixels in the color channels have been multiplied by the pixels in the alpha channel. The result is that transparent pixels are always black since any number multiplied by 0 always equals 0. This obscures the background (by multiplying with the inverse of the foreground alpha), and then simply adds the pixels from the foreground.
  - A Subtractive merge is necessary if the foreground image is not premultiplied. The compositing method is similar to an Additive merge, but the foreground image is first multiplied by its alpha to eliminate any background pixels outside the alpha area.

  In most software applications, you will find the Additive/Subtractive option displayed as a simple checkbox. Fusion lets you blend between the Additive and Subtractive versions of the merge operation, which is occasionally useful for dealing with problem composites with edges that are calling attention to themselves as too bright or too dark.

  > For example, using a Subtractive setting on a premultiplied image may result in darker edges. Using an Additive setting with a non-premultiplied image may result in lightening the edges. By blending between Additive and Subtractive, you can tweak the edge brightness to be just right for your situation.

Filter Method (Match Move)

Determines which filter to use to handle image transforms made using the Tracker node. This menu appears only when the Operation Mode is set to Match Move.

- **Box:** This is a simple interpolation resize of the image.
- **Linear:** This uses a simplistic filter, which produces relatively clean and fast results.
- **Quadratic:** This filter produces a nominal result. It offers a good compromise between speed and quality.
- **Cubic:** This produces better results with continuous-tone images but is slower than Bi-Cubic. If the images have fine detail in them, the results may be blurrier than desired.
- **Catmull-Rom:** This produces good results with continuous-tone images that are resized down. It produces sharp results with finely detailed images.
- **Gaussian:** This is very similar in speed and quality to Bi-Cubic.

- **Mitchell:** This is similar to Catmull-Rom but produces better results with finely detailed images. It is slower than Catmull-Rom.
- **Lanczos:** This is very similar to Mitchell and Catmull-Rom but is a little cleaner and also slower.
- **Sinc:** This is an advanced filter that produces very sharp, detailed results; however, it may produce visible “ringing” in some situations.
- **Bessel:** This is similar to the Sinc filter but may be slightly faster.

Edges

This menu selects how the revealed edges are handled when the image is moved to match position and scaling.

- **Black Edges:** Out-of-frame edges revealed by Stabilization are left black.
- **Wrap:** Portions of the image moved off frame to one side are used to fill edges that are revealed on the opposite side.
- **Duplicate:** The last valid pixel on an edge is repeated to the edge of the frame.
- **Mirror:** Image pixels are mirrored to fill to the edge of the frame.

Position, Rotation, and Scaling Checkboxes (Match Move)

The Position, Rotation, and Scaling checkboxes appear only when the mode is set to Match Move. They determine what components of motion that Stabilization will attempt to correct in the image. For example, if only the Position checkbox is selected, no attempt will be made to correct for Rotation and Scaling in the image.

Flatten Transformation (Match Move)

This checkbox appears only when the mode is set to Match Move. Like most transformations in Fusion, Stabilization is concatenated with other sequential transformations by default. Selecting this checkbox will flatten the transform, breaking any concatenation taking place and applying the transform immediately.

Mapping Type

The Mapping Type control appears only in the Corner Positioning mode. There are two options in the menu:

- **Bi_Linear:** The first method is Bi-Linear, where the foreground image is mapped into the background without any attempt to correct for perspective distortion. This is identical to how previous versions of Fusion operated.
- **Perspective:** The foreground image is mapped into the background taking perspective distortion into account. This is the preferred setting since it maps better to the real world than the older Bi-Linear setting.

Corner Selector (Corner or Perspective Positioning)

When the operation of the Tracker is set to either Corner or Perspective Positioning modes, four drop-down menus appear. These options choose which trackers map to each of the four corners of the rectangle. This is useful when a Tracker has more than four patterns selected, and you must choose which patterns the positioners use.

Rotate Clockwise and Counter-Clockwise Buttons (Corner or Perspective Positioning)

These controls appear only when the operation of the Tracker is set to either Corner or Perspective Positioning modes. They are used to rotate the foreground image by 90 degrees before it is applied to the background.

Stabilize Settings

The Tracker node automatically outputs several steady and unsteady position outputs to which other controls in the Node Editor can be connected. The Stable Position output provides X and Y coordinates to match or reverse motion in a sequence. These controls are available even when the operation is not set to Match Move, since the Stable Position output is always available for connection to other nodes.

Match Move Settings

These settings determine how tracking data is correlated with the reference pattern for making transforms.

Pivot Type

The Pivot type menu determines how the anchor point for rotation is selected.

- **Tracker Average:** Averages the location based on the tracking points.
- **Selected Tracker:** Provides a menu where one of the current trackers can be selected as the pivot point.
- **Manual:** Displays X and Y position number fields where you can manually position the pivot points.

Reference

The Reference mode determines the “snapshot frame” based on the frame where the pattern is first selected. All Stabilization is intended to return the image back to that reference.

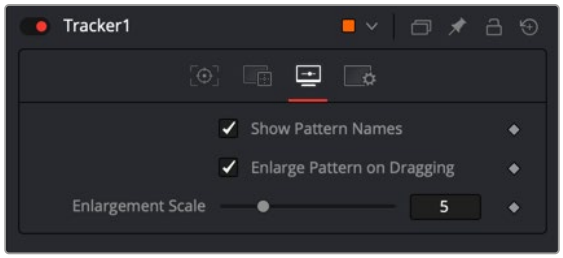
- **Select Time:** Lets you select the current frame.
- **Start:** The Snapshot Frame is determined to be the first frame in the tracked path. All Stabilization is intended to return the image back to that reference.
- **Start and End:** The Start and End Reference mode is somewhat different from all other Reference modes. Where the others are intended to take a snapshot frame to which all stabilization returns, immobilizing the image, the Start and End mode is intended to smooth existing motion, without removing it. This mode averages the motion between the Start and End of the path, drawing a straight line between those points.

  When this mode is active, it reveals the Reference Intermediate Points control. Increasing the value of this control increases the number of points in the path used by the Reference, smoothing the motion from a straight line between Start and End without making it wholly linear.
- **End:** The Snapshot Frame is determined to be the last frame in the tracked path. All Stabilization is intended to return the image back to that reference.

TIP: By default, the Tracker displays a single displacement path of the tracked data in the Spline Editor. To view X and Y paths of the tracked points in the Spline Editor, go to Preferences > Globals > Splines.

Display Options Tab

The Display Options tab lets you customize the look of onscreen controls in the viewer.



The Tracker Display Options tab

Show Pattern Names

This option defines whether the Tracker's pattern names will be displayed in the viewer. Switch it off to see the pattern rectangle instead.

Enlarged Pattern on Dragging

This option defines whether there is a magnified thumbnail view when positioning the pattern rectangle.

Enlargement Scale

The zoom factor that is used when positioning the pattern rectangle when the above option is activated. TIP: The outputs of a tracker (seen in the Connect to... menu) can also be used by scripts. They are:

- **SteadyPosition:** Steady Position
- **UnsteadyPosition:** Unsteady Position
- **SteadyAxis:** Steady Axis
- **SteadySize:** Steady Size
- **UnsteadySize:** Unsteady Size
- **SteadyAngle:** Steady Angle
- **UnsteadyAngle:** Unsteady Angle
- **Position1:** Tracker 1 Offset position
- **PerspectivePosition1:** Tracker 1 Perspective Offset position
- **PositionX1:** Tracker 1 Offset X position (3D Space)
- **PositionY1:** Tracker 1 Offset Y position (3D Space)
- **PerspectivePositionX1:** Tracker 1 Perspective Offset X position (3D Space)
- **PerspectivePositionY1:** Tracker 1 Perspective Offset Y position (3D Space)
- **SteadyPosition1:** Tracker 1 Steady Position
- **UnsteadyPosition1:** Tracker 1 Unsteady Position (likewise for the 2nd, 3rd, and so on)

Common Controls

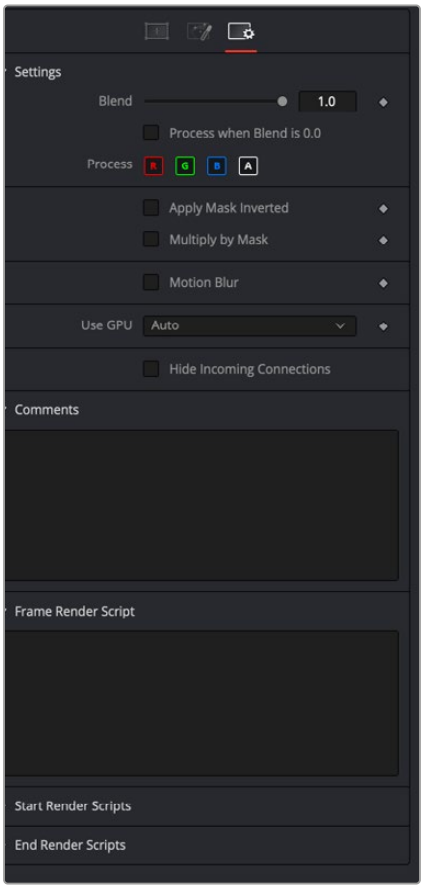
Settings Tab

The Settings tab in the Inspector is also duplicated in other Tracking nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

The Common Controls

Nodes that handle tracking operations share several identical controls in the Inspector. This section describes controls that are common among Tracking nodes.

Inspector



The Common Tracking Controls Settings tab

Settings Tab

The Settings tab in the Inspector can be found on every tool in the Tracking category. The controls are consistent and work the same way for each tool.

Blend

The Blend control is used to blend between the tool's original image input and the tool's final modified output image. When the blend value is 0.0, the outgoing image is identical to the incoming image. Normally, this will cause the tool to skip processing entirely, copying the input straight to the output.

Process When Blend Is 0.0

The tool is processed even when the input value is zero. This can be useful if the node is scripted to trigger a task, but the node's value is set to 0.0.

Red/Green/Blue/Alpha Channel Selector

These four buttons are used to limit the effect of the tool to specified color channels. This filter is often applied after the tool has been processed.

For example, if the Red button on a Blur tool is deselected, the blur will first be applied to the image, and then the red channel from the original input will be copied back over the red channel of the result.

There are some exceptions, such as tools for which deselecting these channels causes the tool to skip processing that channel entirely. Tools that do this will generally possess a set of identical RGBA buttons on the Controls tab in the tool. In this case, the buttons in the Settings and the Controls tabs are identical.

Apply Mask Inverted

Enabling the Apply Mask Inverted option inverts the complete mask channel for the tool. The mask channel is the combined result of all masks connected to or generated in a node.

Multiply by Mask

Selecting this option will cause the RGB values of the masked image to be multiplied by the mask channel's values. This will cause all pixels of the image not included in the mask (i.e., set to 0) to become black/transparent.

Use Object/Use Material (Checkboxes)

Some 3D software can render to file formats that support additional channels. Notably, the EXR file format supports Object ID and Material ID channels, which can be used as a mask for the effect. These checkboxes determine whether the channels will be used, if present. The specific Material ID or Object ID affected is chosen using the next set of controls.

Correct Edges

This checkbox appears only when the Use Object or Use Material checkboxes are selected. It toggles the method used to deal with overlapping edges of objects in a multi-object image. When enabled, the Coverage and Background Color channels are used to separate and improve the effect around the edge of the object. If this option is disabled (or no Coverage or Background Color channels are available), aliasing may occur on the edge of the mask.

> For more information on the Coverage and Background Color channels, *see Chapter 18, "Understanding Image Channels,"* in the Fusion Reference Manual.

Object ID/Material ID (Sliders)

Use these sliders to select which ID will be used to create a mask from the object or material channels of an image. Use the Sample button in the same way as the Color Picker: to grab IDs from the image displayed in the viewer. The image or sequence must have been rendered from a 3D software package with those channels included.

Motion Blur

- **Motion Blur:** This toggles the rendering of Motion Blur on the tool. When this control is toggled on, the tool's predicted motion is used to produce the motion blur caused by the virtual camera's shutter. When the control is toggled off, no motion blur is created.
- **Quality:** Quality determines the number of samples used to create the blur. A quality setting of 2 will cause Fusion to create two samples to either side of an object's actual motion. Larger values produce smoother results but increase the render time.

- **Shutter Angle:** Shutter Angle controls the angle of the virtual shutter used to produce the motion blur effect. Larger angles create more blur but increase the render times. A value of 360 is the equivalent of having the shutter open for one full frame exposure. Higher values are possible and can be used to create interesting effects.
- **Center Bias:** Center Bias modifies the position of the center of the motion blur. This allows for the creation of motion trail effects.
- **Sample Spread:** Adjusting this control modifies the weighting given to each sample. This affects the brightness of the samples.

Use GPU

The Use GPU menu has three settings. Setting the menu to Disable turns off GPU hardware-accelerated rendering. Enabled uses the GPU hardware for rendering the node. Auto uses a capable GPU if one is available and falls back to software rendering when a capable GPU is not available.

Hide Incoming Connections

Enabling this checkbox can hide connection lines from incoming nodes, making a node tree appear cleaner and easier to read. When enabled, empty fields for each input on a node will be displayed in the Inspector. Dragging a connected node from the node tree into the field will hide that incoming connection line as long as the node is not selected in the node tree. When the node is selected in the node tree, the line will reappear.

Comments

The Comments field is used to add notes to a tool. Click in the empty field and type the text. When a note is added to a tool, a small red square appears in the lower-left corner of the node when the full tile is displayed, or a small text bubble icon appears on the right when nodes are collapsed. To see the note in the Node Editor, hold the mouse pointer over the node to display the tooltip.

Scripts

Three Scripting fields are available on every tool in Fusion from the Settings tab. They each contain edit boxes used to add scripts that process when the tool is rendering. For more details on scripting nodes, please consult the Fusion scripting documentation.

Chapter 60

Transform Nodes

This chapter details the Transform nodes available in Fusion.

The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Camera Shake [CSh] 1623
Crop [CRP] 1626
DVE [DVE] 1629
Letterbox [LBX] 1632
Planar Transform [Pxf] 1635
Resize [RSZ] 1637
Scale [SCL] 1640
Transform [XF] 1643
The Common Controls 1647

Camera Shake [CSh]



The Camera Shake node

Camera Shake Node Introduction

This node can simulate a variety of camera shake-style motions from organic to mechanical. It is not the same as the Shake Modifier, which generates random number values for parameters.

> For more information on the Shake modifier, *see Chapter 64, "Modifiers,"* in the Fusion Reference Manual.

The Camera Shake node concatenates its result with adjacent transformation nodes for higher-quality processing.

Inputs

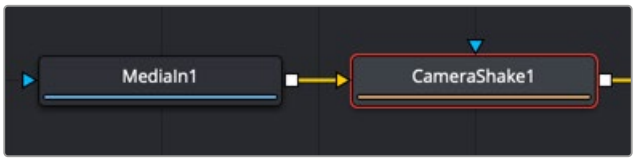
The two inputs on the Camera Shake node are used to connect a 2D image and an effect mask, which can be used to limit the camera shake area.

Input: The orange input is used for the primary 2D image that shakes.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the camera shake area to only those pixels within the mask. An effects mask is applied to the tool after the tool is processed.

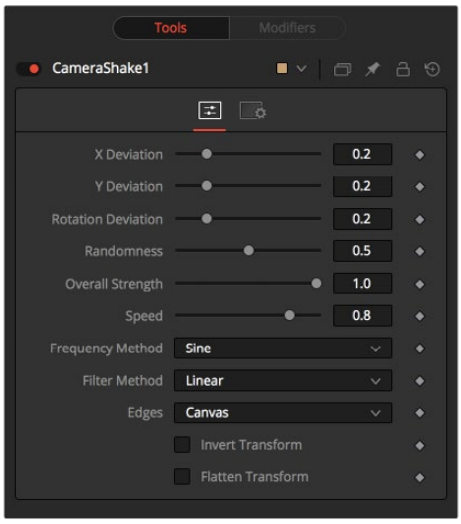
Basic Node Setup

The Camera Shake background input is used to connect the image you want to transform. Polygon masks can be connected into the occlusion mask input to identify areas the camera shake should ignore.



The Camera Shake node can be connected directly after a MediaIn node or any node providing a 2D output.

Inspector



The Camera Shake Controls tab

Controls Tab

The Controls tab includes parameters for adjusting the offsets, strength, speed, and frequency of the simulated camera shake movement.

Deviation X and Y

These controls determine the amount of shake applied to the image along the horizontal (X) and vertical (Y) axes. Values between 0.0 and 1.0 are permitted. A value of 1.0 generates shake positions anywhere within the boundaries of the image.

Rotation Deviation

This determines the amount of shake that is applied to the rotational axis. Values between 0.0 and 1.0 are permitted.

Randomness

Higher values in this control cause the movement of the shake to be more irregular or random. Smaller values cause the movement to be more predictable.

Overall Strength

This adjusts the general amplitude of all the parameters and blends that affect in and out. A value of 1.0 applies the effect as described by the remainder of the controls.

Speed

Speed controls the frequency, or rate, of the shake.

Frequency Method

This selects the overall shape of the shake. Available frequencies are Sine, Rectified Sine, and Square Wave. A Square Wave generates a much more mechanical-looking motion than a Sine.

Filter Method

When rescaling a pixel, surrounding pixels are often used to give a more realistic result. There are various algorithms for combining these pixels, called filters. More complex filters can give better results but are usually slower to calculate.

The best filter for the job often depends on the amount of scaling and on the contents of the image itself.

- **Box:** This is a simple interpolation resize of the image.
- **Linear:** This uses a simplistic filter, which produces relatively clean and fast results.
- **Quadratic:** This filter produces a nominal result. It offers a good compromise between speed and quality.
- **Cubic:** This produces better results with continuous-tone images. If the images have fine detail in them, the results may be blurrier than desired.
- **Catmull-Rom:** This produces good results with continuous-tone images that are resized down. This produces sharp results with finely detailed images.
- **Gaussian:** This is very similar in speed and quality to Bi-Cubic.
- **Mitchell:** This is similar to Catmull-Rom but produces better results with finely detailed images. It is slower than Catmull-Rom.
- **Lanczos:** This is very similar to Mitchell and Catmull-Rom but is a little cleaner and also slower.
- **Sinc:** This is an advanced filter that produces very sharp, detailed results; however, it may produce visible “ringing” in some situations.
- **Bessel:** This is similar to the Sinc filter but may be slightly faster.

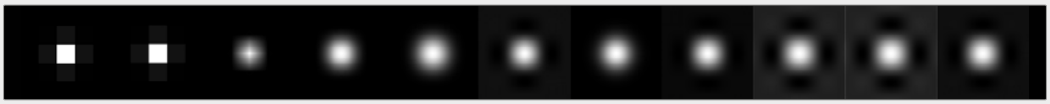
Window Method (Sinc and Bessel Only)

Some filters, such as Sinc and Bessel, require an infinite number of pixels to calculate exactly. To speed up this operation, a windowing function is used to approximate the filter and limit the number of pixels required. This control appears when a filter that requires windowing is selected.

- **Hanning:** This is a simple tapered window.
- **Hamming:** Hamming is a slightly tweaked version of Hanning that does not taper all the way down to zero.
- **Blackman:** A window with a more sharply tapered falloff.
- **Kaiser:** A more complex window with results between Hamming and Blackman.

Most of these filters are useful only when making an image larger. When shrinking images, it is common to use the Bi-Linear filter, however, the Catmull-Rom filter will apply some sharpening to the results and may be useful for preserving detail when scaling down an image.

Example



Resize filters. From left to right: Nearest Neighbor, Box, Linear, Quadratic, Cubic, Catmull-Rom, Gaussian, Mitchell, Lanczos, Sinc, and Bessel.

Edges

This menu determines how the Edges of the image are treated.

- **Canvas:** This causes the edges that are revealed by the shake to be the canvas color—usually transparent or black.
- **Wrap:** This causes the edges to wrap around (the top is wrapped to the bottom, the left is wrapped to the right, and so on).
- **Duplicate:** This causes the Edges to be duplicated, causing a slight smearing effect at the edges.
- **Mirror:** Image pixels are mirrored to fill to the edge of the frame.

Invert Transform

Select this control to Invert any position, rotation, or scaling transformation. This option might be useful for exactly removing the motion produced in an upstream Camera Shake.

Flatten Transform

The Flatten Transform option prevents this node from concatenating its transformation with adjacent nodes. The node may still concatenate transforms from its input, but it will not concatenate its transformation with the node at its output.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Transform nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Crop [CRP]



The Crop node

Crop Node Introduction

The Crop node can be used to cut out a portion of an image or to offset the image into a larger image area. However, unlike using a mask, this node actually changes the resolution of the image.

TIP: You can crop an image in the viewer by activating the Allow Box Selection in the upper-left corner of the viewer while the Crop node is selected and viewed. Then, drag a crop rectangle around the area of interest to perform the operation.

NOTE: Because this node changes the physical resolution of the image, animating the parameters is not advised.

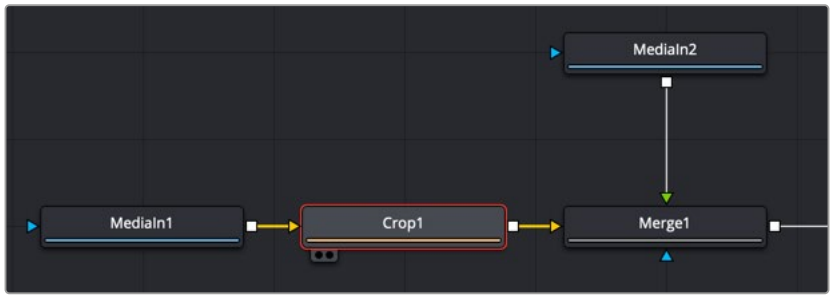
Inputs

The single input on the Crop node is used to connect a 2D image for cropping.

Input: The orange input is used for the primary 2D image you want to crop.

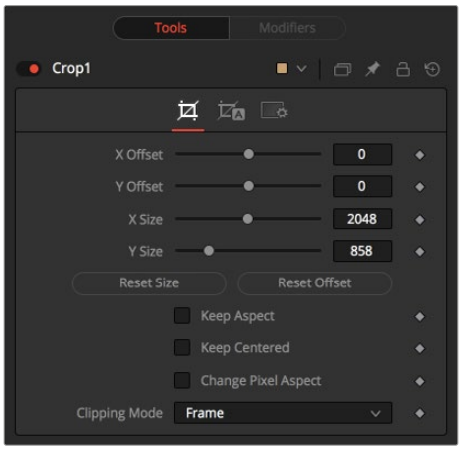
Basic Node Setup

Below, the Crop node is inserted between the MediaIn1 node and the background input of the Merge. Unlike using a mask tool, cropping the MediaIn1 changes the resolution of the clip. The cropped MediaIn1 node connected to the orange background input also sets the resolution of the Merge output.



The Crop node can be used to cut out a portion of an image.

Inspector



The Crop Controls tab

Controls Tab

The Controls tab provides XY Offset and XY Size methods for cropping the image.

Offset X and Y

These controls position the image off the screen by pushing it left/right or up/down. The cropped image disappears off the edges of the output image. The values of these controls are measured in pixels.

Size X and Y

Use these controls to set the vertical and horizontal resolution of the image output by the Crop node. The values of these controls are measured in pixels.

Keep Aspect

When toggled on, the Crop node maintains the aspect of the input image.

Keep Centered

When toggled on, the Crop node automatically adjusts the X and Y Offset controls to keep the image centered. The XY Offset sliders are automatically adjusted, and control over the cropping is done with the Size sliders or the Allow Box Selection button in the viewer.

Reset Size

This resets the image dimensions to the size of the input image.

Reset Offset

This resets the X and Y Offsets to their defaults.

Change Pixel Aspect

Enable this checkbox to reveal a Pixel Aspect control that can be used to change the image's pixel aspect.

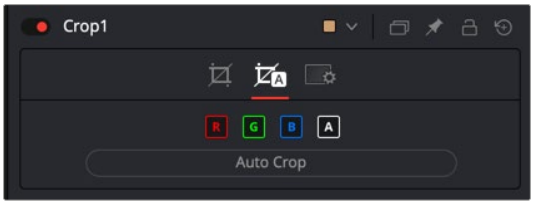
Clipping Mode

This option sets the mode used to handle the edges of the image when performing domain of definition (DoD) rendering. This is profoundly important for nodes like Blur, which may require samples from portions of the image outside the current domain.

- **Frame:** The default option is Frame, which automatically sets the node's domain of definition to use the full frame of the image, effectively ignoring the current domain of definition. If the upstream DoD is smaller than the frame, the remaining area in the frame will be treated as black/transparent.
- **Domain:** Setting this option to Domain will respect the upstream DoD when applying the node's effect. This can have adverse clipping effects in situations where the node employs a large filter.
- **None:** Setting this option to None does not perform any source image clipping at all. This means that any data required to process the node's effect that would normally be outside the upstream DoD is treated as black/transparent.

Auto Crop Tab

Auto Crop tab analyzes the selected channel and crops the image based on that channel's boundaries. The adjustments from auto crop are seen in the Crop tab parameters.



The Auto Crop tab

RGBA Color Channels

Select which channels are examined for an Auto Crop. This is useful for auto cropping images with non-solid backgrounds in a specific color channel, like a blue color gradient. Toggling the channel off causes Auto Crop to ignore it when evaluating the image.

Auto Crop

This evaluates the image and attempts to determine the background color. It then crops each side of the image to the first pixel that is not that color.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Transform nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

DVE [DVE]



The DVE node

DVE Node Introduction

The DVE (Digital Video Effects) node is a 3D-image transformation similar to nodes found in old, tape-based online editing suites. The node encompasses image rotations, perspective changes, and Z moves. The axis can be defined for all transformations.

Inputs

The three inputs on the DVE node are used to connect a 2D image, DVE mask, and an effect mask, which can be used to limit the DVE area.

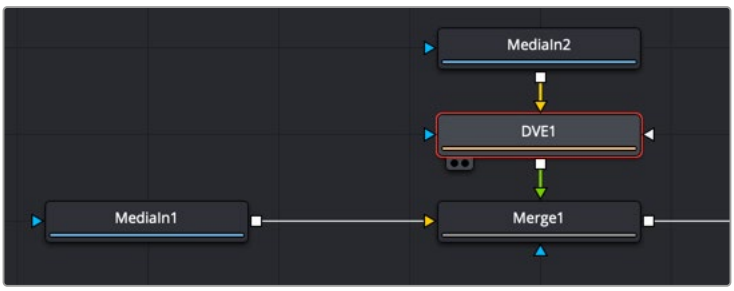
Input: The orange input is used for the primary 2D image that is transformed by the DVE.

DVE Mask: The white DVE mask input is used to mask the image prior to the DVE transform being applied. This has the effect of modifying both the image and the mask.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input causes the DVE to modify only the image within the mask. An effects mask is applied to the tool after the tool is processed.

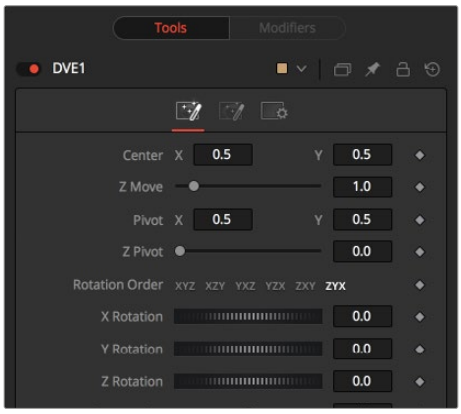
Basic Node Setup

In the example below, the DVE node is inserted between the MediaIn2 node and the foreground input of the Merge. The MediaIn1 node is manipulated in the DVE node and composited over the top of the MediaIn1 node.



The DVE node modifying the foreground input of a Merge node

Inspector



The DVE Controls tab

Controls Tab

The Controls tab includes all the transform parameters for the DVE.

Pivot X, Y, and Z

Positions the axis of rotation and scaling. The default is 0.5, 0.5 for X and Y, which is in the center of the image, and 0 for Z, which is at the center of Z space.

Rotation Order

Use these buttons to determine in what order rotations are applied to the image.

XYZ Rotation

These controls are used to rotate the image around the pivot along the X-, Y- and Z-axis.

Center X and Y

This positions the center of the DVE image onscreen. The default is 0.5, 0.5, which positions the DVE in the center of the image.

Z Move

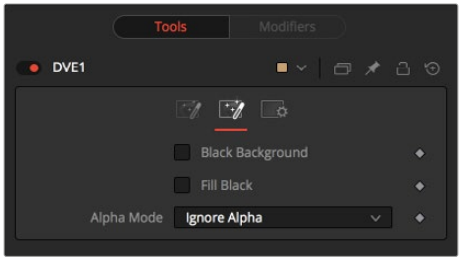
This zooms the image in and out along the Z-axis. Visually, when this control is animated, the effect is similar to watching an object approach from a distance.

Perspective

This adds additional perspective to an image rotated along the X- or Y-axis, similar to changing the Field of View and zoom of a camera.

Masking Tab

The DVE node allows pre-masking of its input image. This offers the ability to create transformations from the masked area of the image while leaving the remainder of the image unaffected.



The DVE Masking tab

Unlike regular effect masks, the masking process occurs before the transformation. All the usual mask types can be applied to the DVE mask.

Black Background

Toggle this on to erase the area outside the mask from the transformed image.

Fill Black

Toggle this on to erase the area within the mask (before transformation) from the DVE's input, effectively cutting the masked area out of the image. Enabling both Black Background and Fill Black will show only the masked, transformed area.

Alpha Mode

This determines how the DVE will handle the alpha channel of the image when merging the transformed image areas over the untransformed image.

- **Ignore Alpha:** This causes the input image's alpha channel to be ignored, so all masked areas will be opaque.
- **Subtractive/Additive:** These cause the internal merge of the pre-masked DVE image over the input image to be either Subtractive or Additive.

  An Additive setting is necessary when the foreground DVE image is premultiplied, meaning that the pixels in the color channels have been multiplied by the pixels in the alpha channel. The result is that transparent pixels are always black, since any number multiplied by 0 always equals 0. This obscures the background (by multiplying with the inverse of the foreground alpha), and then simply adds the pixels from the foreground.

  A Subtractive setting is necessary if the foreground DVE image is not premultiplied. The compositing method is similar to an Additive merge, but the foreground DVE image is first multiplied by its own alpha, to eliminate any background pixels outside the alpha area.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Transform nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Letterbox [LBX]



The Letterbox node

Letterbox Node Introduction

Use the Letterbox node to adapt existing images to the frame size and aspect ratios of any other format. The most common use of this node is to convert film resolution images to HD-sized frames for viewing on an external television monitor. Horizontal or vertical black edges are automatically added where necessary to compensate for aspect ratio differences. This node actually changes the resolution of the image.

NOTE: Because this node changes the physical resolution of the image, animating the controls is not recommended.

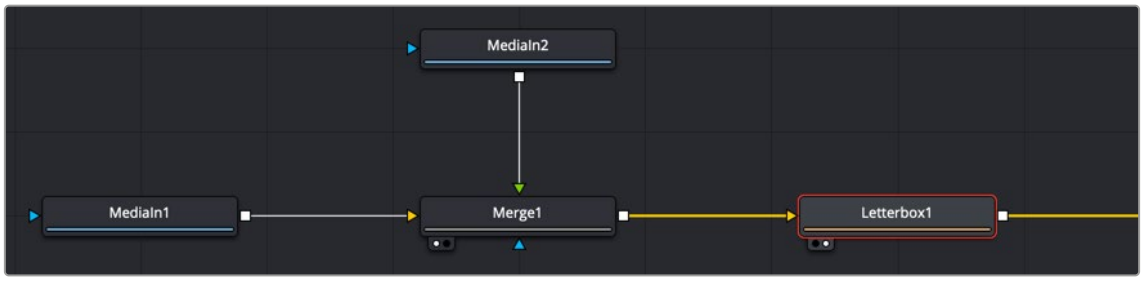
Inputs

The single input on the Letterbox node is used to connect a 2D image for letterbox/cropping.

Input: The orange input is used for the primary 2D image you want to letterbox/crop.

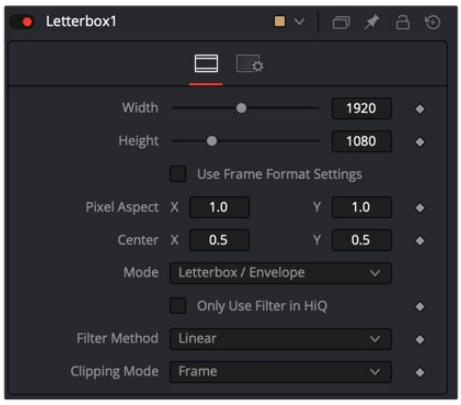
Basic Node Setup

The Letterbox node is used in the example below to change the resolution of the Merge node's output. Depending on how the resolution is modified, side pillars, or a horizontal letterbox mask, is applied to "fill in" the frame area, which the Merge node output does not cover.



The Letterbox node converting the Merge output resolution and adding letterbox masking where needed.

Inspector



The Letterbox Controls tab

Controls Tab

The Controls tab includes parameters for adjusting the resolution and pixel aspect of the image. It also has the option of letterboxing or pan-and-scan formatting.

Width and Height

The values of these controls determine the size of the output image as measured in pixels.

> **TIP:** You can use the formatting contextual menu to quickly select a resolution from a list. Place the pointer over the Width or Height controls, and then right-click to display the contextual menu. The bottom of the menu displays a Select Frame Format submenu with available frame formats. Select any one of the choices from the menu to set the Height, Width, and Aspect controls automatically.

Auto Resolution

Activating this checkbox automatically sets the Width and Height sliders to the Frame Format settings found in the Preferences window for Fusion Studio or to the resolution of the DaVinci Resolve Timeline.

Pixel Aspect X and Y

These controls determine the pixel aspect ratio of the output image.

Center X and Y

This Center control repositions the image window when used in conjunction with Pan-and-Scan mode. It has no effect on the image when the node is set to Letterbox mode.

Mode

This control is used to determine the Letterbox node's mode of operation.

- **Letterbox/Envelope:** This corrects the aspect of the input image and resizes it to match the specified width.
- **Pan-and-Scan:** This corrects the aspect of the input image and resizes it to match the specified height. If the resized input image is wider than the specified width, the Center control can be used to animate the visible portion of the resized input.

Filter Method

When rescaling a pixel, surrounding pixels are often used to give a more realistic result. There are various algorithms for combining these pixels, called filters. More complex filters can give better results but are usually slower to calculate. The best filter for the job often depends on the amount of scaling and on the contents of the image itself.

- **Box:** This is a simple interpolation resize of the image.
- **Linear:** This uses a simplistic filter, which produces relatively clean and fast results.
- **Quadratic:** This filter produces a nominal result. It offers a good compromise between speed and quality.
- **Cubic:** This produces better results with continuous-tone images. If the images have fine detail in them, the results may be blurrier than desired.
- **Catmull-Rom:** This produces good results with continuous-tone images that are resized down. This produces sharp results with finely detailed images.
- **Gaussian:** This is very similar in speed and quality to Bi-Cubic.
- **Mitchell:** This is similar to Catmull-Rom but produces better results with finely detailed images. It is slower than Catmull-Rom.
- **Lanczos:** This is very similar to Mitchell and Catmull-Rom but is a little cleaner and also slower.
- **Sinc:** This is an advanced filter that produces very sharp, detailed results; however, it may produce visible “ringing” in some situations.
- **Bessel:** This is similar to the Sinc filter but may be slightly faster.

Window Method (Sinc and Bessel Only)

Some filters, such as Sinc and Bessel, require an infinite number of pixels to calculate exactly. To speed up this operation, a windowing function is used to approximate the filter and limit the number of pixels required. This control appears when a filter that requires windowing is selected.

- **Hanning:** This is a simple tapered window.
- **Hamming:** Hamming is a slightly tweaked version of Hanning that does not taper all the way down to zero.
- **Blackman:** A window with a more sharply tapered falloff.
- **Kaiser:** A more complex window with results between Hamming and Blackman.

Most of these filters are useful only when making an image larger. When shrinking images, it is common to use the Bi-Linear filter; however, the Catmull-Rom filter will apply some sharpening to the results and may be useful for preserving detail when scaling down an image.

Example



Different resize filters. From left to right: Nearest Neighbor, Box, Linear, Quadratic, Cubic, Catmull-Rom, Gaussian, Mitchell, Lanczos, Sinc, and Bessel.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Transform nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Planar Transform [Pxf]



The Planar Transform node

The Planar Transform node applies perspective distortions generated by a Planar Tracker node onto any input mask or masked image. The Planar Transform node can be used to reduce the amount of time spent on rotoscoping objects. The workflow here centers around the notion that the Planar Tracker node can be used to track objects that are only roughly planar. After an object is tracked, a Planar Transform node can then be used to warp a rotospline, making it approximately follow the object over time. Fine-level cleanup work on the rotospline then must be done.

Depending on how well the Planar Tracker followed the object, this can result in substantial time savings in the amount of tedious rotoscoping. The key to using this technique is recognizing situations where the Planar Tracker performs well on an object that needs to be rotoscoped.

A rough outline of the workflow involved is:

1. Track: Using a Planar Tracker node, select a pattern that represents the object to be rotoscoped. Track the shot (see the tracking workflow in the Track section for the Planar Tracker node).
2. Create a Planar Transform node: Press the Create Planar Transform button on the Planar Tracker node to do this. The newly created Planar Transform node can be freely cut and pasted into another composition as desired.
3. Rotoscope the object: Move to any frame that was tracked by the Planar Tracker. When unsure if a frame was tracked, look in the Spline Editor for a tracking keyframe on the Planar Transform node. Connect a Polygon node into the Planar Transform node. While viewing the Planar Transform node, rotoscope the object.
4. Refine: Scrub the timeline to see how well the polygon follows the object. Adjust the polyline on frames where it is off. It is possible to add new points to further refine the polygon.

Inputs

The Planar Transform has only two inputs:

Image Input: The orange image input accepts a 2D image on which the transform will be applied.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the output of the Planar Transform to certain areas.

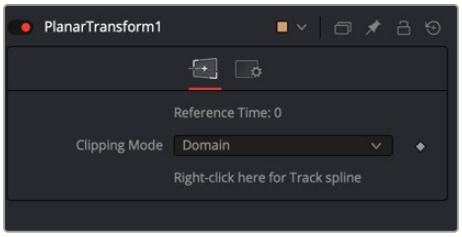
Basic Node Setup

The example below uses a Planar Transform node between a masked MediaIn2 node and the foreground input to a Merge node. The background MediaIn1 node connects to the Planar Tracker, which was used to generate the Planar Transform. Once the Planar Transform is created, the Planar Tracker is no longer needed in the node tree.



A Planar Transform creating a match move

Inspector



The Planar Transform Controls tab

Controls Tab

The Planar Transform node has very few controls, and they are all located in the Controls tab. It's designed to apply the analyzed Planar Tracking data as a match move,

Reference Time

This is the reference time that the pattern was taken from in the Planar Tracker node used to produce the Planar Transform.

Right-Click Here for Track Spline

The Track spline contains information about the perspective distortions stored in 4 x 4 matrices. When a Planar Transform node is exported from a Planar Tracker node, the Track spline produced by the Planar Tracker is shared by connecting it with the Planar Transform node. A consequence of this sharing of the Track spline is that if the track is changed in the Planar Tracker node, the Planar Transform will be automatically updated. Note that this spline can be examined in the Spline Editor, which is useful for seeing the extent of tracked frames.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Tracking nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Resize [RSZ]



The Resize node

Resize Node Introduction

Use the Resize node to increase or decrease the resolution of an input image. This is useful for converting images from one format to another (for example, from film to video resolution). This node actually changes the resolution of the image.

NOTE: Because this node changes the physical resolution of the image, animating the controls is not advised.

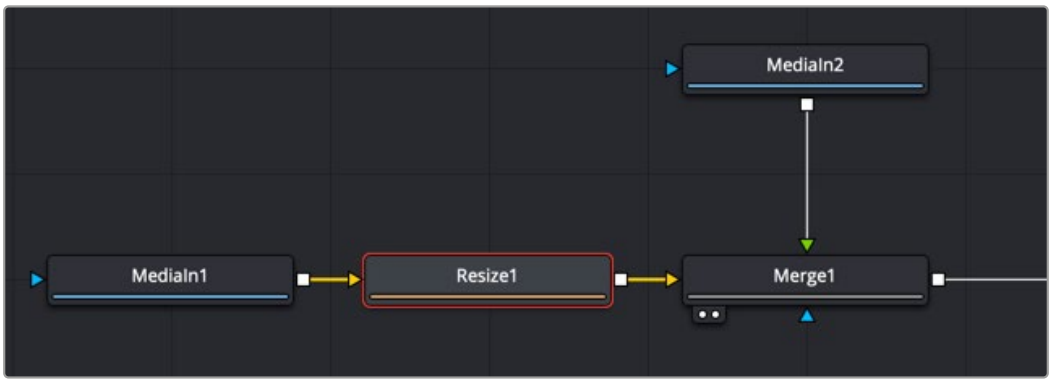
Inputs

The single input on the Resize node is used to connect a 2D image for resizing.

Input: The orange input is used for the primary 2D image you want to resize.

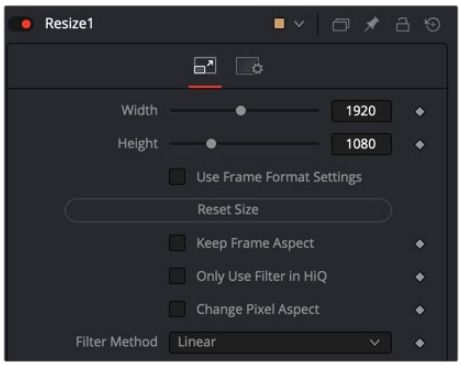
Basic Node Setup

Below, the Resize node is inserted between the MediaIn1 node and the background input of the Merge. Unlike using a Transform tool, resizing the MediaIn1 changes the resolution of the clip. The resized MediaIn1 node connected to the orange background input also sets the resolution of the Merge output.



The Resize node can be used to scale an image and change its resolution.

Inspector



The Resize Controls tab

Controls Tab

The Controls tab includes parameters for changing the resolution of the image. It uses pixel values in the Width and Height controls.

Width

This controls the new resolution for the image along the X-axis.

Height

This controls the new resolution for the image along the Y-axis.

> **TIP:** You can use the formatting contextual menu to quickly select a resolution from a list. Place the mouse pointer over the Width or Height controls, and then right-click to display the contextual menu. The bottom of the menu displays a Select Frame Format submenu with available frame formats. Select any one of the choices from the menu to set the Height and Width controls automatically.

Auto Resolution

Activating this checkbox automatically sets the Width and Height sliders to the Frame Format settings found in the Preferences window for Fusion Studio or the resolution in the DaVinci Resolve Timeline.

Reset Size

Resets the image dimensions to the original size of the image.

Keep Frame Aspect

When toggled on, the Resize node maintains the aspect of the original image, preserving the original ratio between width and height.

Only Use Filter in HiQ

The Resize node will normally use the fast Nearest Neighbor filter for any non-HiQ renders, where speed is more important than full accuracy. Disable this checkbox to force Resize to always use the selected filter for all renders.

Change Pixel Aspect

Enable this checkbox to reveal a Pixel Aspect control that can be used to change the image's pixel aspect.

Filter Method

When rescaling a pixel, surrounding pixels are often used to give a more realistic result. There are various algorithms for combining these pixels, called filters. More complex filters can give better results but are usually slower to calculate. The best filter for the job often depends on the amount of scaling and on the contents of the image itself.

- **Box:** This is a simple interpolation resize of the image.
- **Linear:** This uses a simplistic filter, which produces relatively clean and fast results.
- **Quadratic:** This filter produces a nominal result. It offers a good compromise between speed and quality.
- **Cubic:** This produces better results with continuous-tone images. If the images have fine detail in them, the results may be blurrier than desired.
- **Catmull-Rom:** This produces good results with continuous-tone images that are resized down. This produces sharp results with finely detailed images.
- **Gaussian:** This is very similar in speed and quality to Bi-Cubic.
- **Mitchell:** This is similar to Catmull-Rom but produces better results with finely detailed images. It is slower than Catmull-Rom.
- **Lanczos:** This is very similar to Mitchell and Catmull-Rom but is a little cleaner and also slower.
- **Sinc:** This is an advanced filter that produces very sharp, detailed results; however, it may produce visible “ringing” in some situations.
- **Bessel:** This is similar to the Sinc filter but may be slightly faster.

Window Method (Sinc and Bessel Only)

Some filters, such as Sinc and Bessel, require an infinite number of pixels to calculate exactly. To speed up this operation, a windowing function is used to approximate the filter and limit the number of pixels required. This control appears when a filter that requires windowing is selected.

- **Hanning:** This is a simple tapered window.
- **Hamming:** Hamming is a slightly tweaked version of Hanning that does not taper all the way down to zero.
- **Blackman:** A window with a more sharply tapered falloff.
- **Kaiser:** A more complex window with results between Hamming and Blackman.

Most of these filters are useful only when making an image larger. When shrinking images, it is common to use the Bi-Linear filter; however, the Catmull-Rom filter will apply some sharpening to the results and may be useful for preserving detail when scaling down an image.

Example



Different resize filters. From left to right: Nearest Neighbor, Box, Linear, Quadratic, Cubic, Catmull-Rom, Gaussian, Mitchell, Lanczos, Sinc, and Bessel.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Transform nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Scale [SCL]



The Scale node

Scale Node Introduction

The Scale node is almost identical to the Resize node, except that Resize uses exact dimensions, whereas the Scale node uses relative dimensions to describe the change to the source image's resolution. This node actually changes the resolution of the image.

NOTE: Because this node changes the physical resolution of the image, animating the controls is not advised.

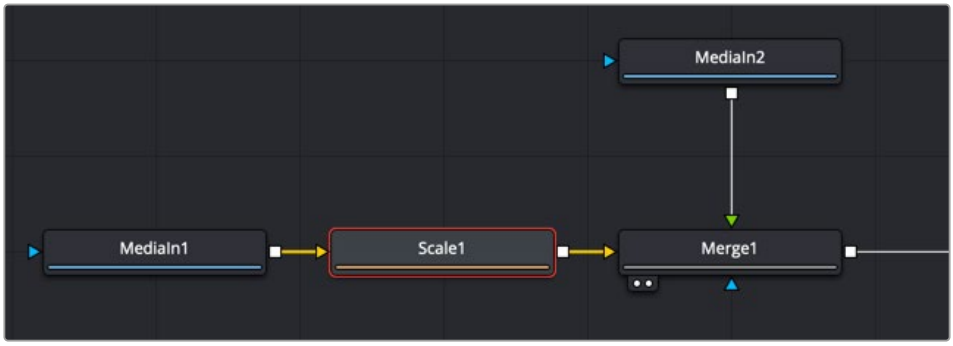
Inputs

The single input on the Scale node is used to connect a 2D image for scaling.

Input: The orange input is used for the primary 2D image you want to scale.

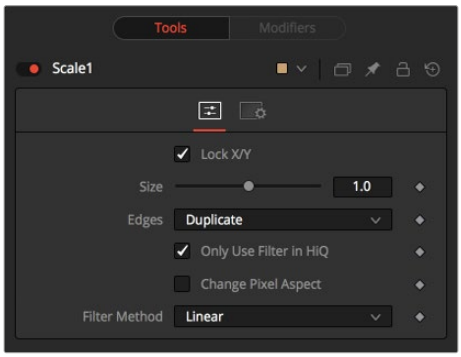
Basic Node Setup

Below, the Scale node is inserted between the MediaIn1 node and the background input of the Merge. Unlike using a Transform tool, scaling the MediaIn1 changes the resolution of the clip. The resized MediaIn1 node connected to the orange background input also sets the resolution of the Merge output.



The Scale node can be used to scale an image and change its resolution.

Inspector



The Scale Controls tab

Controls Tab

The Controls tab includes parameters for changing the resolution of the image. It uses a multiplier of size to set the new resolution. An Edges menu allows you to determine how the edges of the frame are handled if the scaling decreases.

Lock X/Y

When selected, only a Size control is shown, and changes to the image's scale are applied to both axes equally. If the checkbox is cleared, individual Size controls appear for both X and Y Size.

Size

The Size control is used to set the scale used to adjust the resolution of the source image. A value of 1.0 would have no affect on the image, while 2.0 would scale the image to twice its current resolution. A value of 0.5 would halve the image's resolution.

Only Use Filter in HiQ

The Scale node will normally use the fast Nearest Neighbor filter for any non-HiQ renders, where speed is more important than full accuracy. Disable this checkbox to force Scale to always use the selected filter for all renders.

Change Pixel Aspect

Enable this checkbox to reveal a Pixel Aspect control that can be used to change the image's pixel aspect.

Filter Method

When rescaling a pixel, surrounding pixels are often used to give a more realistic result. There are various algorithms for combining these pixels, called filters. More complex filters can give better results but are usually slower to calculate. The best filter for the job often depends on the amount of scaling and on the contents of the image itself.

- **Box:** This is a simple interpolation resize of the image.
- **Linear:** This uses a simplistic filter, which produces relatively clean and fast results.
- **Quadratic:** This filter produces a nominal result. It offers a good compromise between speed and quality.

- **Cubic:** This produces better results with continuous-tone images. If the images have fine detail in them, the results may be blurrier than desired.
- **Catmull-Rom:** This produces good results with continuous-tone images that are resized down. This produces sharp results with finely detailed images.
- **Gaussian:** This is very similar in speed and quality to Bi-Cubic.
- **Mitchell:** This is similar to Catmull-Rom but produces better results with finely detailed images. It is slower than Catmull-Rom.
- **Lanczos:** This is very similar to Mitchell and Catmull-Rom but is a little cleaner and also slower.
- **Sinc:** This is an advanced filter that produces very sharp, detailed results; however, it may produce visible "ringing" in some situations.
- **Bessel:** This is similar to the Sinc filter but may be slightly faster.

Window Method (Sinc and Bessel Only)

Some filters, such as Sinc and Bessel, require an infinite number of pixels to calculate exactly. To speed up this operation, a windowing function is used to approximate the filter and limit the number of pixels required. This control appears when a filter that requires windowing is selected.

- **Hanning:** This is a simple tapered window.
- **Hamming:** Hamming is a slightly tweaked version of Hanning that does not taper all the way down to zero.
- **Blackman:** A window with a more sharply tapered falloff.
- **Kaiser:** A more complex window with results between Hamming and Blackman.

Most of these filters are useful only when making an image larger. When shrinking images, it is common to use the Bi-Linear filter; however, the Catmull-Rom filter will apply some sharpening to the results and may be useful for preserving detail when scaling down an image.

Example



Different resize filters. From left to right: Nearest Neighbor, Box, Linear, Quadratic, Cubic, Catmull-Rom, Gaussian, Mitchell, Lanczos, Sinc, and Bessel.

NOTE: Because this node changes the physical resolution of the image, animating the controls is not advised.

Transform [XF]



The Transform node

Transform Node Introduction

The Transform node can be used for simple 2D transformations of the image, such as moving, rotating, and scaling. The image's aspect can also be modified using the Transform node.

The Transform node concatenates its result with adjacent Transformation nodes. The Transform node does not change the image's resolution.

Inputs

The two inputs on the Transform node are used to connect a 2D image and an effect mask, which can be used to limit the transformed area.

Input: The orange input is used for the primary 2D image that gets transformed.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the transform area to only those pixels within the mask. An effects mask is applied to the tool after the tool is processed.

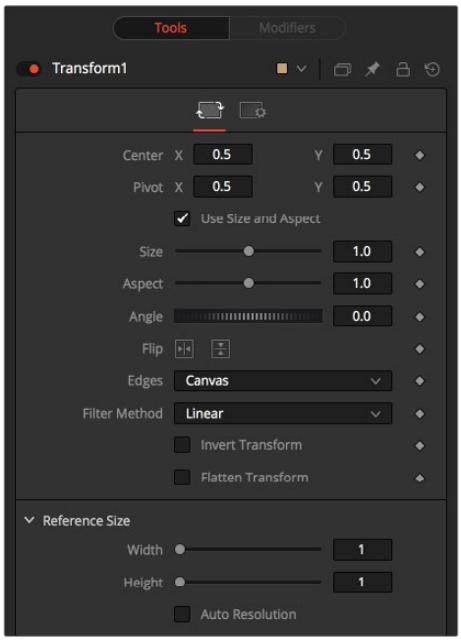
Basic Node Setup

The Transform node in the example below is inserted between the MediaIn2 node and the foreground input of the Merge. Unlike using a Scale or Resize tool, transforming the MediaIn2 does not change the resolution of the clip. For that reason, it is the tool most often used to scale, move, and rotate a clip.



The Transform node can be used to scale an image without changing its resolution.

Inspector



The Transform Controls tab

Controls Tab

The Controls tab presents multiple ways to transform, flip (vertical), flop (horizontal), scale, and rotate an image. It also includes reference size controls that can reinterpret the coordinates used for width and height from relative values of 0-1 into pixel values based on the image's resolution.

Center X and Y

This sets the position of the image on the screen. The default is 0.5, 0.5, which places the image in the center of the screen. The value shown is always the actual position multiplied by the reference size. See below for a description of the reference size.

Pivot X and Y

This positions the axis of rotation and scaling. The default is 0.5, 0.5, which is the center of the image.

Use Size and Aspect

This checkbox determines whether the Transform node provides independent Size controls for the X and Y scale or if Size and Aspect controls are used instead.

Size

This modifies the scale of the image. Values range from 0 to 5, but any value greater than zero can be entered into the edit box. If the Use Size and Aspect checkbox is selected, this control will scale the image equally along both axes. If the Use Size and Aspect option is off, independent control is provided for X and Y.

Aspect

This control changes the aspect ratio of an image. Setting the value above 1.0 stretches the image along the X-axis. Values between 0.0 and 1.0 stretch the image along the Y-axis. This control is available only when the Use Size and Aspect checkbox is enabled.

Angle

This control rotates the image around the axis. Increasing the Angle rotates the image in a counterclockwise direction. Decreasing the Angle rotates the image in a clockwise direction.

Flip Horizontally and Vertically

Toggle this control on to flip the image along the X- or Y-axis.

Edges

This menu determines how the edges of the image are treated when the edge of the raster is exposed.

- **Canvas:** This causes the edges of the image that are revealed to show the current Canvas Color. This defaults to black with no Alpha and can be set using the Set Canvas Color node.
- **Wrap:** This wraps the edges of the image around the borders of the image. This is useful for seamless images to be panned, creating an endless moving background image.
- **Duplicate:** This causes the edges of the image to be duplicated as best as possible, continuing the image beyond its original size.
- **Mirror:** Image pixels are mirrored to fill to the edge of the frame.

Filter Method

When rescaling a pixel, surrounding pixels are often used to give a more realistic result. There are various algorithms for combining these pixels, called filters. More complex filters can give better results but are usually slower to calculate. The best filter for the job often depends on the amount of scaling and on the contents of the image itself.

- **Box:** This is a simple interpolation resize of the image.
- **Linear:** This uses a simplistic filter, which produces relatively clean and fast results.
- **Quadratic:** This filter produces a nominal result. It offers a good compromise between speed and quality.
- **Cubic:** This produces better results with continuous-tone images. If the images have fine detail in them, the results may be blurrier than desired.
- **Catmull-Rom:** This produces good results with continuous-tone images that are resized down. This produces sharp results with finely detailed images.
- **Gaussian:** This is very similar in speed and quality to Bi-Cubic.
- **Mitchell:** This is similar to Catmull-Rom but produces better results with finely detailed images. It is slower than Catmull-Rom.
- **Lanczos:** This is very similar to Mitchell and Catmull-Rom but is a little cleaner and also slower.
- **Sinc:** This is an advanced filter that produces very sharp, detailed results; however, it may produce visible “ringing” in some situations.
- **Bessel:** This is similar to the Sinc filter but may be slightly faster.

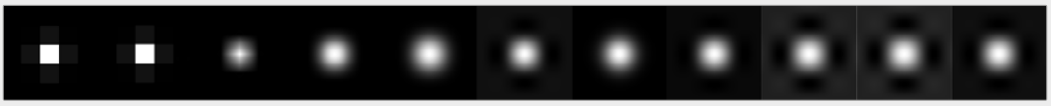
Window Method (Sinc and Bessel Only)

Some filters, such as Sinc and Bessel, require an infinite number of pixels to calculate exactly. To speed up this operation, a windowing function is used to approximate the filter and limit the number of pixels required. This control appears when a filter that requires windowing is selected.

- **Hanning:** This is a simple tapered window.
- **Hamming:** Hamming is a slightly tweaked version of Hanning that does not taper all the way down to zero.
- **Blackman:** A window with a more sharply tapered falloff.
- **Kaiser:** A more complex window with results between Hamming and Blackman.

Most of these filters are useful only when making an image larger. When shrinking images, it is common to use the Bi-Linear filter; however, the Catmull-Rom filter will apply some sharpening to the results and may be useful for preserving detail when scaling down an image.

Example



Different resize filters. From left to right: Nearest Neighbor, Box, Linear, Quadratic, Cubic, Catmull-Rom, Gaussian, Mitchell, Lanczos, Sinc, and Bessel.

Invert Transform

Select this control to invert any position, rotation, or scaling transformation. This option is useful when connecting the Transform to the position of a tracker for the purpose of reintroducing motion back into a stabilized image.

Flatten Transform

The Flatten Transform option prevents this node from concatenating its transformation with adjacent nodes. The node may still concatenate transforms from its input, but it will not concatenate its transformation with the node at its output.

Reference Size

The controls under the Reference Size menu do not directly affect the image. Instead they allow you to control how Fusion represents the position of the Transform node's center.

Normally, coordinates are represented as values between 0 and 1, where 1 is a distance equal to the full width or height of the image. This allows for resolution independence, because you can change the size of the image without having to change the value of the center.

One disadvantage to this approach is that it complicates making pixel-accurate adjustments to an image. To demonstrate, imagine an image that is 100 x 100 pixels in size. To move the center of the image to the right by 5 pixels, we would change the X value of the transform center from 0.5, 0.5 to 0.55, 0.5. We know the change must be 0.05 because 5/100 = 0.05.

The Reference Size controls allow you to specify the dimensions of the image. This changes the way the control values are displayed, so that the Center shows the actual pixel positions in the X and Y number fields of the Center control. Extending our example, if you set the Width and Height to 100 each, the Center would now be shown as 50, 50, and we would move it 5 pixels toward the right by entering 55, 50.

Internally, the Transform node still stores this value as a number between 0 to 1, and if you were to query the Center controls value via scripting, or publish the Center control for use by other nodes, then you would retrieve the original normalized value. The change is visible only in the value shown for Transform Center in the node control.

Reference Width and Height Sliders

Set these to the width and height of the image to change the way that Fusion displays the values of the Transform node's Center control.

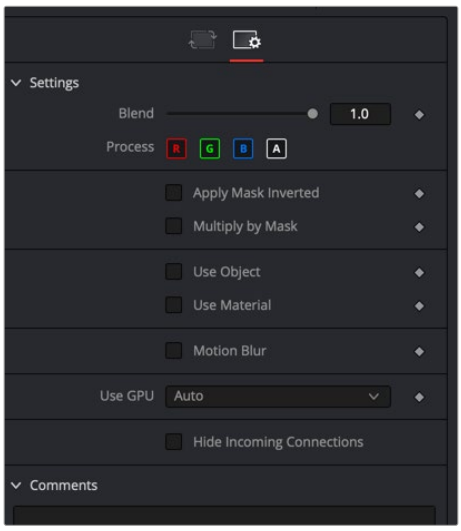
Auto Resolution

Enable this checkbox to use the current frame format settings in Fusion Studio or the timeline resolution in DaVinci Resolve to set the Reference Width and Reference Height values.

The Common Controls

Nodes that handle Transform operations share several identical controls in the Inspector. This section describes controls that are common among Transform nodes.

Inspector



The Transform Common controls

Settings Tab

The Settings tab in the Inspector can be found on every tool in the Transform category. The Settings controls are even found on third-party Transform-type plugin tools. The controls are consistent and work the same way for each tool.

Blend

The Blend control is used to blend between the tool's original image input and the tool's final modified output image. When the blend value is 0.0, the outgoing image is identical to the incoming image. Normally, this will cause the tool to skip processing entirely, copying the input straight to the output.

Process When Blend Is 0.0

The tool is processed even when the input value is zero. This can be useful if the node is scripted to trigger a task, but the node's value is set to 0.0.

Red/Green/Blue/Alpha Channel Selector

These four buttons are used to limit the effect of the tool to specified color channels. This filter is often applied after the tool has been processed.

For example, if the Red button on a Blur tool is deselected, the blur will first be applied to the image, and then the red channel from the original input will be copied back over the red channel of the result.

There are some exceptions, such as tools for which deselecting these channels causes the tool to skip processing that channel entirely. Tools that do this will generally possess a set of identical RGBA buttons on the Controls tab in the tool. In this case, the buttons in the Settings and the Controls tabs are identical.

Apply Mask Inverted

Enabling the Apply Mask Inverted option inverts the complete mask channel for the tool. The mask channel is the combined result of all masks connected to or generated in a node.

Multiply by Mask

Selecting this option will cause the RGB values of the masked image to be multiplied by the mask channel's values. This will cause all pixels of the image not included in the mask (i.e., set to 0) to become black/transparent.

Use Object/Use Material (Checkboxes)

Some 3D software can render to file formats that support additional channels. Notably, the EXR file format supports Object ID and Material ID channels, which can be used as a mask for the effect. These checkboxes determine whether the channels will be used, if present. The specific Material ID or Object ID affected is chosen using the next set of controls.

Correct Edges

This checkbox appears only when the Use Object or Use Material checkboxes are selected. It toggles the method used to deal with overlapping edges of objects in a multi-object image. When enabled, the Coverage and Background Color channels are used to separate and improve the effect around the edge of the object. If this option is disabled (or no Coverage or Background Color channels are available), aliasing may occur on the edge of the mask.

> For more information on the Coverage and Background Color channels, *see Chapter 18, "Understanding Image Channels,"* in the Fusion Reference Manual.

Object ID/Material ID (Sliders)

Use these sliders to select which ID will be used to create a mask from the object or material channels of an image. Use the Sample button in the same way as the Color Picker: to grab IDs from the image displayed in the viewer. The image or sequence must have been rendered from a 3D software package with those channels included.

Motion Blur

- **Motion Blur:** This toggles the rendering of Motion Blur on the tool. When this control is toggled on, the tool's predicted motion is used to produce the motion blur caused by the virtual camera's shutter. When the control is toggled off, no motion blur is created.
- **Quality:** Quality determines the number of samples used to create the blur. A quality setting of 2 will cause Fusion to create two samples to either side of an object's actual motion. Larger values produce smoother results but increase the render time.
- **Shutter Angle:** Shutter Angle controls the angle of the virtual shutter used to produce the motion blur effect. Larger angles create more blur but increase the render times. A value of 360 is the equivalent of having the shutter open for one full frame exposure. Higher values are possible and can be used to create interesting effects.
- **Center Bias:** Center Bias modifies the position of the center of the motion blur. This allows for the creation of motion trail effects.
- **Sample Spread:** Adjusting this control modifies the weighting given to each sample. This affects the brightness of the samples.

Use GPU

The Use GPU menu has three settings. Setting the menu to Disable turns off GPU hardware-accelerated rendering. Enabled uses the GPU hardware for rendering the node. Auto uses a capable GPU if one is available and falls back to software rendering when a capable GPU is not available.

Hide Incoming Connections

Enabling this checkbox can hide connection lines from incoming nodes, making a node tree appear cleaner and easier to read. When enabled, empty fields for each input on a node will be displayed in the Inspector. Dragging a connected node from the node tree into the field will hide that incoming connection line as long as the node is not selected in the node tree. When the node is selected in the node tree, the line will reappear.

Comments

The Comments field is used to add notes to a tool. Click in the empty field and type the text. When a note is added to a tool, a small red square appears in the lower-left corner of the node when the full tile is displayed, or a small text bubble icon appears on the right when nodes are collapsed. To see the note in the Node Editor, hold the mouse pointer over the node to display the tooltip.

Scripts

Three Scripting fields are available on every tool in Fusion from the Settings tab. They each contain edit boxes used to add scripts that process when the tool is rendering. For more details on scripting nodes, please consult the Fusion scripting documentation.

Chapter 61

USD Nodes

This chapter details the Universal Scene Descriptor (USD) nodes available in Fusion.

The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Universal Scene Descriptor (USD) **Files** 1651

USD Scene Tree Dialog for Object Selection in USD Files 1652

uCamera [uCa] 1653

uDuplicate [uDp] 1656

uExport [uEx] 1661

uImage Plane [uIm] 1662

uLoader [uLd] 1664

uMaterialX [uMX] 1665

uMerge [uMg] 1667

uRenderer [uRn] 1668

uReplaceMaterial [uRM] 1671

uShape [uSh] 1674

uSwitch [uSw] 1677

uTransform [uXf] 1679

uVariant [uVa] 1681

uVisibility [uVis] 1682

uVolume [uVo] 1684

USD Lights

uCylinder Light [uCL] 1686

uDisk Light [uDi] 1687

uDistant Light [uDL] 1689

uDome Light [uDo] 1690

uRectangle Light [uRL] 1692

uSphere Light [uSL] 1694

USD Material

uNormalMap [uNm] 1696

uShader [uSd] 1697

uTexture [uTx] 1700

uTextureTransform [uTXf] 1701

The Common Controls 1702

Universal Scene Descriptor (USD) Files

The Universal Scene Descriptor USD framework is a set of open standards for describing, composing, and simulating a 3D environment. The USD file format is more than just a simple format - it's an open source 3D scene description used for creating and sharing 3D content across various applications.

DaVinci Resolve and Fusion can import USD (.usdc, .usdz, and .usda) 3D information including geometry, lighting, cameras, materials, and animation. A new collection of USD tools has been added to Fusion, allowing users to manipulate, re-light, and render these USD files.

The USD Loader

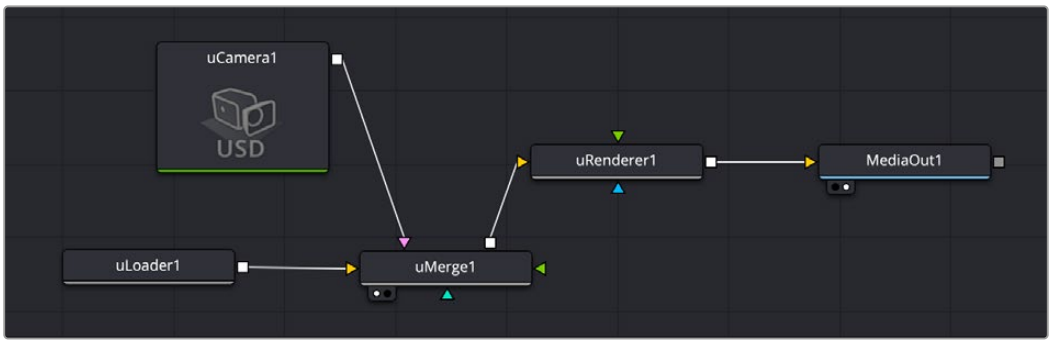
The uLoader allows you to import USD files and adjust its options.



The uLoader node

To Import a USD file:

1. Add a uLoader tool to your Fusion Composition.
2. Select the USD file you want to load in the File Browser. If you wish to change this file in the future, you can do so in the uLoader Inspector without disrupting your node tree.
3. Connect the output of the uLoader to a uMerge tool.
4. Connect any other USD tools, like uCamera or any of the uLights, to manipulate the scene.
5. Connect the output of the uMerge tool to a uRenderer.
6. Connect the output of the uRenderer to the MediaOut tool.

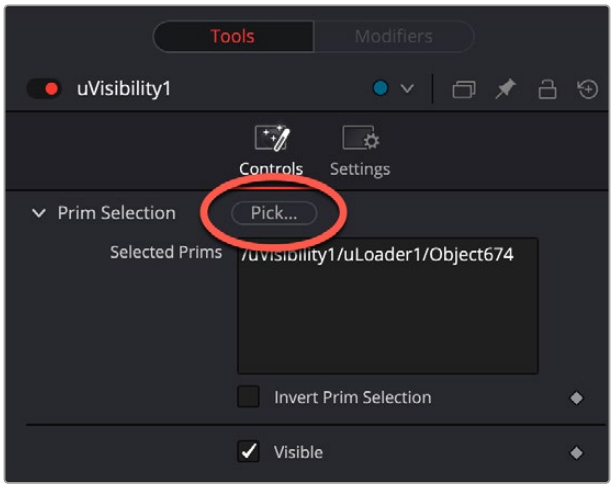


A very simple USD import node tree

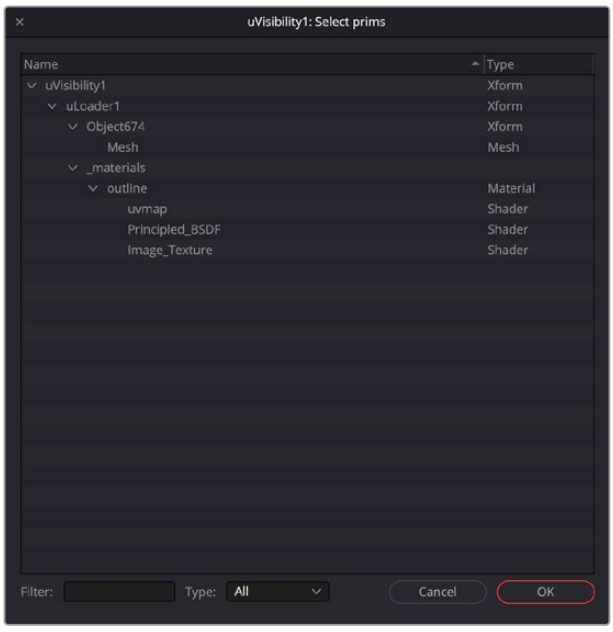
USD Scene Tree Dialog for Object Selection in USD Files

The Scene Tree window is an intuitive visual hierarchy to represent USD scenes. It lets you choose and modify individual elements inside a USD scene. The USD Scene Tree is important for selective effects using the uVisibility, uVariant, uReplace Material, and uTransform nodes, but also used for overriding imported scene objects in uCamera and the various USD light tools.

The USD Scene Tree is activated by pressing the Pick button in the Prim Selection box in certain USD tools.



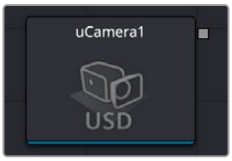
The Prim Selection Pick button



The USD Scene Tree dialog

The tree shows the USD hierarchy, and the type column shows what kind of asset it is. Users can select, filter, isolate, and choose specific elements in a USD scene to focus on logical components and make targeted adjustments. Holding Command allows the toggle selection of multiple objects, and Shift allows selecting a range of objects with a click.

uCamera [uCa]



The USD Camera Node

uCamera Node Introduction

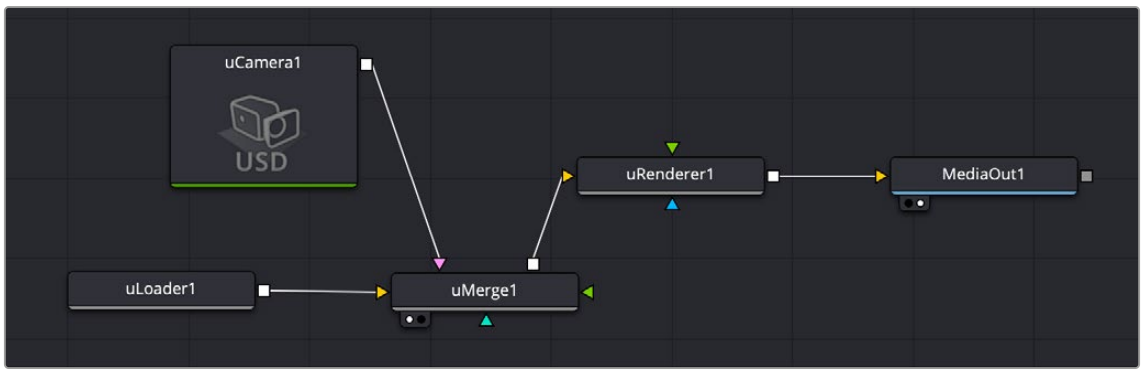
The uCamera node generates a virtual camera for viewing the USD scene in its 3D environment. It closely emulates the settings used in real cameras to make matching live-action or 3D-rendered elements as seamless as possible. Adding any cameras to a 3D composite allows you to frame the elements in a composite how you want and animate the camera during a scene to create moving camera shots.

Inputs

There is one Scene Input in yellow to attach a USD scene. This allows you to adjust any existing cameras within the scene.

Basic Node Setup

The output of a uCamera node should be connected to a uMerge node to build it into a USD 3D scene with other assets and nodes.



The uCamera node is connected to a uMerge node to give standard camera controls to a USD scene.

Displaying a uCamera node directly in the viewer shows only an empty scene; there is nothing for the camera to see. To view the scene through the camera, view the uMerge node where the uCamera is connected, or any node downstream of that uMerge. Then right-click on the viewer and select uCamera > [Camera name] from the contextual menu. Right-clicking on the axis label found in the lower corner of each USD viewer also displays the Camera submenu.

The aspect of the viewer may be different from the aspect of the camera, so the camera view may not match the actual boundaries of the image rendered by the uRenderer node. Guides can be enabled to represent the portion of the view that the camera sees and assist you in framing the shot. Right-click on the viewer and select an option from the Guides > Frame Aspect submenu. The default option uses the format enabled in the Composition > Frame Format preferences. To toggle the guides on or off, select Guides > Show Guides from the viewers' contextual menu, or use the Command-G (macOS) or Ctrl-G (Windows) keyboard shortcut when the viewer is active.

Inspector



The uCamera Controls tab

Controls Tab

The Controls tab contains some of the most fundamental camera settings, including the camera's clipping plains, focal length, and film back properties.

Override Selection

Pressing the Pick button lets you choose an existing camera in an attached USD file to allow adjusting the cameras in an imported scene.

Projection Type

The Projection Type menu is used to select between Perspective and Orthographic cameras. Generally, real-world cameras are perspective cameras. An orthographic camera uses parallel orthographic projection, a technique where the view plane is perpendicular to the viewing direction. This produces a parallel camera output that is undistorted by perspective.

Orthographic cameras present controls only for the near and far clipping planes and a control to set the viewing scale.

Near/Far Clip

The clipping planes are used to limit what geometry in a scene is rendered based on an object's distance from the camera's focal point. Clipping planes ensure objects that are extremely close to the camera, as well as objects that are too far away to be useful, are excluded from the final rendering.

The default perspective camera ignores this setting unless the Adaptive Near/Far Clip checkbox located under the Near/Far Clip control is disabled.

The clip values use units, so a far clipping plane of 20 means that any object more than 20 units from the camera is invisible to the camera. A near clipping plane of 0.1 means that any object closer than 0.1 units is also invisible.

NOTE: A smaller range between the near and far clipping planes allows greater accuracy in all depth calculations. If a scene begins to render strange artifacts on distant objects, try increasing the distance for the Near Clip plane.

Exposure

When selected, the renderer automatically adjusts the camera's Near/Far Clipping plane to match the extents of the scene. This setting overrides the values of the Near and Far clip range controls described above. This option is not available for orthographic cameras.

Focal Length

In the real world, a lens' focal length is the distance from the center of the lens to the film plane. The shorter the focal length, the closer the focal plane is to the back of the lens. The focal length is measured in millimeters. The angle of view and focal length controls are directly related. Smaller focal lengths produce a wider angle of view, so changing one control automatically changes the other to match.

The relationship between focal length and angle of view is angle = 2 * arctan[aperture / 2 / focal_ length]. Use the vertical aperture size to get the vertical angle of view and the horizontal aperture size to get the horizontal angle of view.

Focal Distance

Like a focal point on a real-world camera, this setting defines the distance from the camera to an object and is used to calculate depth of field.

F Stop

This is used to define the aperture size of the synthetic lens; it will affect exposure and is used to calculate depth of field.

Film Back

This section allows you to control the technical parameters of the non-lens part of the camera.

Horizontal/Vertical Aperture

The Horizontal Aperture Width and Vertical Aperture Height sliders control the dimensions of the camera's aperture or the portion of the camera that lets light in on a real-world camera. In video and film cameras, the aperture is the size of the opening that defines the area of each frame exposed, also known as sensor size.

Lens Shift X/Y

Provides controls to manually offset the camera lens width and height.

Shutter Close/Open

These define when the shutter of the camera will open and close; this will affect Exposure and Motion Blur.

Stereo Role

If the camera is stereo, this defines the center of role for the two cameras, left, center, or right.

Control Visibility

This section allows you to selectively activate the onscreen controls that are displayed along with the camera.

- **Show View Controls:** Displays or hides all camera onscreen controls in the viewers.
- **Frustum:** Displays the actual viewing cone of the camera.
- **View Vector:** Displays a white line inside the viewing cone, which can be used to determine the center of the frame.
- **Near Clip:** The Near Clipping plane. This plane can be subdivided for better visibility.
- **Far Clip:** The Far Clipping plane. This plane can be subdivided for better visibility.
- **Focal Plane:** The plane set by the Focus Distance. This plane can be subdivided for better visibility.
- **Convergence Distance:** The point of convergence when using Stereo mode. This plane can be subdivided for better visibility.

Common Controls

Transform and Settings Tab

The Transform and Settings tab in the Inspector are also duplicated in other USD nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

uDuplicate [uDp]



The USD Duplicate node

Introduction

Similar to the 2D and 3D Duplicate nodes, the uDuplicate node can be used to duplicate any USD asset in a scene, applying a successive transformation to each, and creating repeating models and complex arrays of objects. The options in the Jitter tab allow non-uniform transformations, such as random positioning or sizes.

Inputs

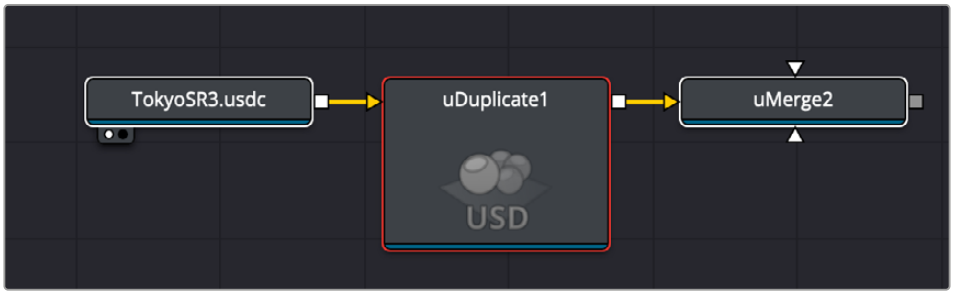
The uDuplicate node has a single input by default where you connect a USD scene. An optional Mesh input appears based on the settings of the node.

SceneInput: The orange Scene Input is a required input. The scene or object you connect to this input is duplicated based on the settings in the Control tab of the Inspector.

MeshInput: A green optional mesh input appears when the Region's tab Region menu is set to mesh. The mesh can be any 3D model, either generated in Fusion or imported.

Basic Node Setup

The output of a uDuplicate node typically connects to a uMerge node, integrating it into a larger scene. The USD scene you want duplicated, in this case a USD scene of a Tokyo shophouse, is connected to the orange input.



A USD scene of a Tokyo shophouse is duplicated.

Usage



The uDuplicate controls

Controls Tab

The Controls tab includes all the parameters you can use to create, offset, and scale copies of the USD object connected to the scene input on the node.

USD Instancing

With this setting enabled, Fusion uses the same USD data for each copy, to increase efficiency. However, in some occasions, when the copies differ too much from each other, it will reduce efficiency instead, so you can disable this setting if that's the case.

Copies

Use this range control to set the number of copies made. Each copy is a copy of the last copy, so if this control is set to [0,3], the parent is copied, then the copy is copied, then the copy of the copy is copied, and so on. This allows some interesting effects when transformations are applied to each copy using the controls below.

Setting the First Copy to a value greater than 0 excludes the original object and shows only the copies.

Time Offset

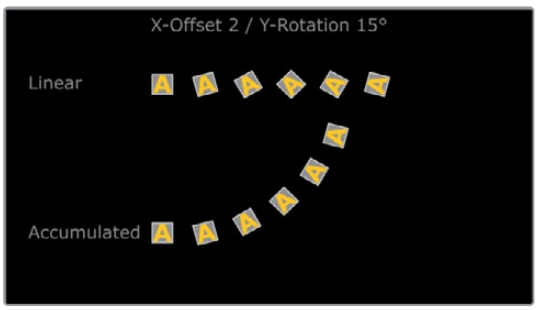
Use the Time Offset slider to offset any animations that are applied to the source geometry by a set amount per copy. For example, set the value to -1.0 and use a cube set to rotate on the Y-axis as the source. The first copy shows the animation from a frame earlier; the second copy shows animation from a frame before that, etc. This can be used with great effect on textured planes—for example, where successive frames of a clip can be shown.

Transform Method



- **Linear:** When set to Linear, transforms are multiplied by the number of the copy, and the total scale, rotation, and translation are applied in turn, independent of the other copies.
- **Accumulated:** When set to Accumulated, each object copy starts at the position of the previous object and is transformed from there. The result is transformed again for the next copy

Transform Order



With this menu, the order in which the transforms are calculated can be set. It defaults to Scale-Rotation-Transform (SRT).

Using different orders results in different positions of your final objects.

Translation

The X, Y, and Z Offset sliders set the offset position applied to each copy. An X offset of 1 would offset each copy 1 unit along the X-axis from the last copy.

Rotation

The buttons along the top of this group of rotation controls set the order in which rotations are applied to the geometry. Setting the rotation order to XYZ would apply the rotation on the X-axis first, followed by the Y-axis rotation, then the Z-axis rotation.

The three Rotation sliders set the amount of rotation applied to each copy.

Pivot

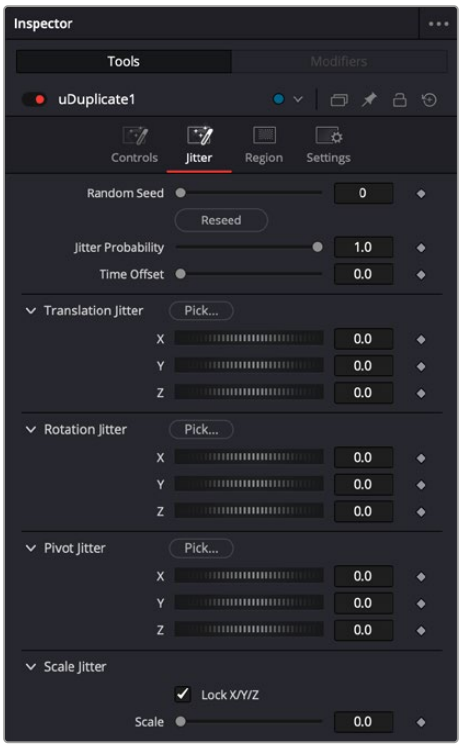
The pivot controls determine the position of the pivot point used when rotating each copy.

Scale

- **Lock:** When the Lock XYZ checkbox is selected, any adjustment to the duplicate scale is applied to all three axes simultaneously. If this checkbox is disabled, the Scale slider is replaced with individual sliders for the X, Y, and Z scales.
- **Scale:** The Scale controls tell Duplicate how much scaling to apply to each copy.

Jitter

The options in the Jitter tab allow you to randomize the position, rotation, and size of all the copies created in the Controls tab.



uDuplicate Jitter tab

Random Seed

The Random Seed slider is used to generate a random starting point for the amount of jitter applied to the duplicated objects. Two Duplicate nodes with identical settings but different random seeds produce two completely different results.

Randomize

Click the Randomize button to auto generate a random seed value.

Jitter Probability

Adjusting this slider determines the percentage of copies that are affected by the jitter. A value of 1.0 means 100% of the copies are affected, while a value of 0.5 means 50% are affected.

Time Offset

Use the Time Offset slider to offset any animations that are applied to the source geometry by a set amount per copy. For example, set the value to –1.0 and use a cube set to rotate on the Y-axis as the source. The first copy shows the animation from a frame earlier; the second copy shows animation from a frame before that, etc. This can be used with great effect on textured planes—for example, where successive frames of a clip can be shown.

Translation Jitter

Use these three controls to adjust the amount of variation in the X, Y, and Z translation of the duplicated objects.

Rotation Jitter

Use these three controls to adjust the amount of variation in the X, Y, and Z rotation of the duplicated objects.

Pivot Jitter

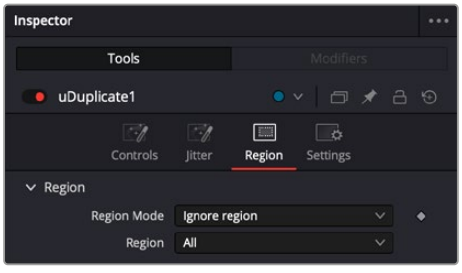
Use these three controls to adjust the amount of variation in the rotational pivot center of the duplicated objects. This affects only the additional jitter rotation, not the rotation produced by the Rotation settings in the Controls tab.

Scale Jitter

Use this control to adjust the amount of variation in the scale of the duplicated objects. Disable the Lock XYZ checkbox to adjust the scale variation independently on all three axes.

Region Tab

The options in the Region tab allow you to define an area in the viewer where the copies can appear or are prevented from appearing. Like most parameters in Fusion, this area can be animated to cause the copied object to pop on and off the screen based on the region's shape and setting.



uDuplicate Region tab

Region Tab

The Region section includes two settings for controlling the shape of the region and the affect the region has on the duplicate objects.

- **Region Mode:** There are three options in the Region Mode menu. The default, labeled “Ignore region” bypasses the node entirely and causes no change to the copies of objects from how they are set in the Controls and Jitter tabs. The menu option labeled “When inside region” causes the copied objects to appear only when their position falls inside the region defined in this tab. The last menu option, “When not Inside region” causes the copied objects to appear only when their position falls outside the region defined in this tab.
- **Region:** The Region menu determines the shape of the region. The five options include cube, sphere, and rectangle primitive shapes. The mesh option allows you to connect a 3D model into the green mesh input on the node. The green input appears only after the Region menu is set to Mesh. The All setting refers to the entire scene. This allows the copies to pop on and off if the Region mode is animated. When the Region menu is set to Mesh, two other options are displayed. These are described below.

  Limit by Object ID: When a scene with multiple meshes is connected to the green Mesh input on the node, all the meshes are used as the region. Enabling this checkbox allows you to use the Object ID slider to select the ID for the mesh you want to use as the Region.

  Object ID: When the Limit by Object ID checkbox is enabled, this slider selects the number ID for the mesh object you want to use for the Region.

uExport [uEx]



The uExport node

uExport Node Introduction

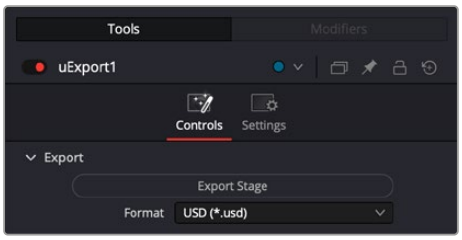
The uExport node is designed to export scenes in the Universal Scene Description (USD) format. The node allows users to export their USD scene, including geometry, materials, animations, and lighting, into a USD file. This upgrade marks an important advancement in Fusion’s functionality, enabling users to seamlessly share complex USD assets, and facilitating smoother workflows in collaborative environments.

The node offers controls to choose between different USD formats such as usd, .usda, .usdc, and .usdz. It is also possible to export specific nodes in the source tree as a hierarchy of linked USD files.

Inputs

There is one input on the Background node for a USD scene input.

Inspector



The uExport controls

Export

The Export controls let you set the format of the exported USD scene.

- **Export Stage:** Press this button to open a file browser to select where to store your USD scene.
- **Format:** Choose a format for your USD scene.

  USD (*.usd)

  USD UTF-8 (*.usda)

  USD binary (*.usdc)

  USD packaged (*.usdz)

Settings Tab

The Settings tab in the Inspector is also duplicated in other USD nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

uImage Plane [uIm]



The USD Image Plane node

uImage Plane Node Introduction

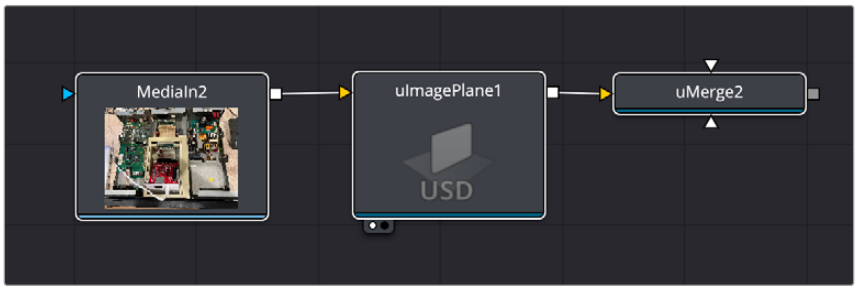
The uImage Plane node produces 2D planar geometry in 3D space. The node is most commonly used to represent 2D images on “cards” in the 3D space. The aspect of the image plane is determined by the aspect of the image that is input.

Inputs

The input on this node is for the image that will be texture mapped to the plane.

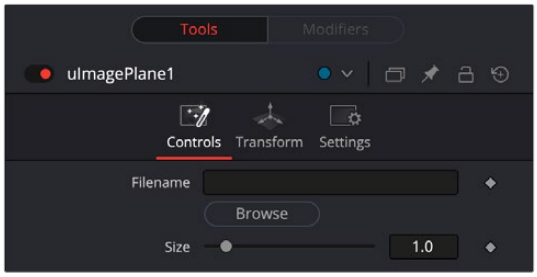
Input: The yellow input is used for the connection of a MediaIn image.

Basic Node Setup



The uImage Plane Node accepts a 2D image and places it into a 3D USD scene.

Inspector



The uImage Plane controls

Controls Tab

Most of the Controls tab is taken up by common controls. The Image Plane-specific controls at the top of the Inspector allow minor adjustments.

Filename

The uImagePlane has an input so an image can be piped directly into the node. This node can also directly load an image by browsing, and it will show the name and path in the filename. Use the Browse button to open the file browser and select an image.

Size

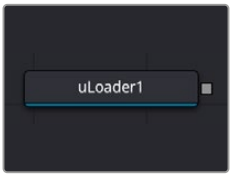
Size will set the size of the image plane in the USD 3D scene.

Common Controls

Transform and Settings Tab

The Transform and Settings tab in the Inspector are also duplicated in other USD nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

uLoader [uLd]



The USD Loader node

uLoader Node Introduction

The uLoader node is used for importing USD files into your Fusion Composition. It will load .usd, .usda, .usdc, and .usdz formats.

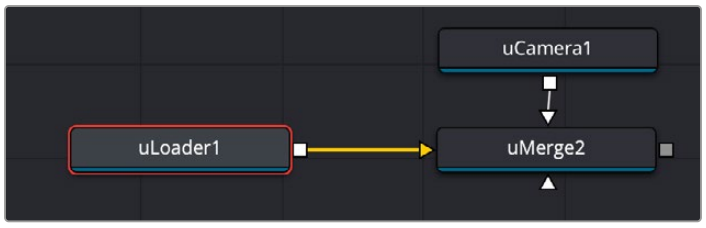
Replacing Materials with Imported MaterialX files

Users can load MaterialX files (.mtlx) from the the USD Loader and use the uReplaceMaterial tool to assign imported materials to a nominated USD object to achieve a desired look.

Inputs

None.

Basic Node Setup



The uLoader node imports a USD scene and then merges it with a USD Camera for view control.

Inspector



The USD Loader controls

Controls Tab

Filename

This control allows you to source a USD file to bring into the Fusion project. Use the Browse button to open the file browser and select a USD file (.usd, .usda, .usdc, or .usdz).

Trim

Lets you select an In and Out point of a USD file to bring only that range into the Fusion project.

Time Scale

Animated .usd scenes can have their animation speed adjusted to be faster and slower using time scale.

Frame Offset

Will define when an animation will start, so timing can be moved along the timeline.

Reverse

When checked, this will reverse the animation of the USD file.

Loop

When checked, this will loop the animation of the USD file.

Reload

This button will reload the USD file set in the Filename field above. This lets you make changes to the USD file in another application and see the results in Fusion by reloading it.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other USD nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

uMaterialX [uMX]



The uMaterialX node

uMaterialX Node Introduction

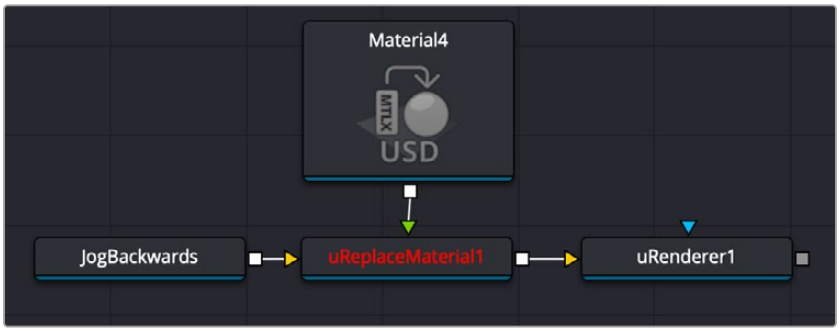
The uMaterialX node simply loads a MaterialX (.mtlx) definition file and creates a USD scene file creating only those materials.

Inputs

There are no inputs to the uMaterialX node. Only a file requester to load the .mtlx file.

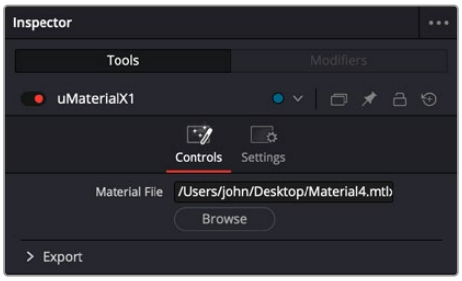
Basic Node Setup

Connect the material output to the rest of the USD tools.



A node tree showing a MaterialX definition (.mtlx) file called Material4 loaded into a uMaterialX node, then passing that material to a uReplaceMaterial node to change the material in a USD scene called JogBackwards.

Inspector



The uMaterialX controls

Controls

The single control is a a Material File selector with a browse button to load the .mtlx file from your computer.

Settings Tab

The Settings tab in the Inspector is also duplicated in other USD nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

uMerge [uMg]



The USD Merge node

uMerge Node Introduction

The uMerge node is the primary node in Fusion that you use to combine separate 3D elements into the same 3D environment.

For example, in a scene created with an image plane, a camera, and a light, the camera would not be able to see the image plane and the light would not affect the image plane until all three objects are introduced into the same environment using the uMerge node.

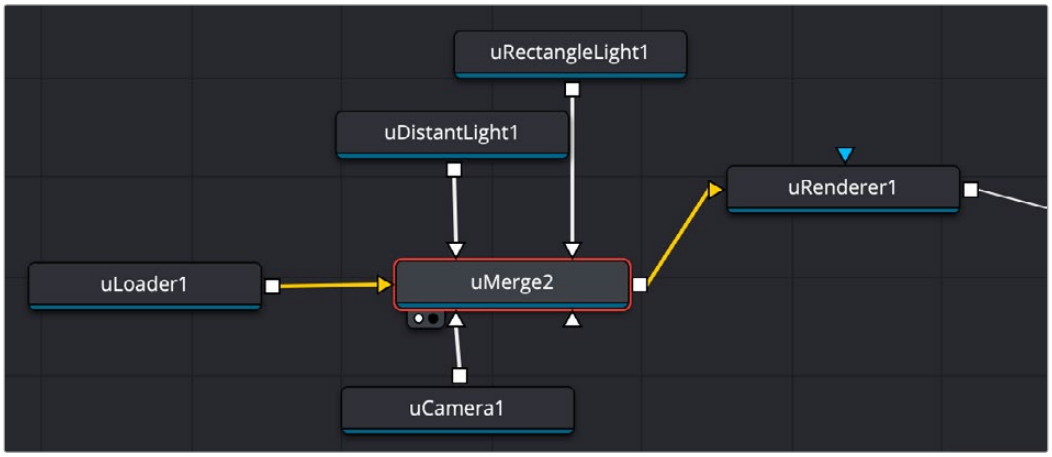
The uMerge provides the standard transformation controls found on most nodes in Fusion's USD suite. Unlike those nodes, changes made to the translation, rotation, or scale of the uMerge affect all the objects connected to the uMerge. This behavior forms the basis for all parenting in Fusion's 3D environment.

Inputs

The uMerge node displays three inputs initially, but as each input is connected a new input appears on the node, assuring there is always one free to add a new element into the scene.

SceneInput[#]: These inputs are used to connect USD image planes, 3D cameras, lights, entire USD scenes, as well as other uMerge nodes. There is no limit to the number of inputs this node can accept. The node dynamically adds more inputs as needed, ensuring that there is always at least one input available for connection.

Basic Node Setup

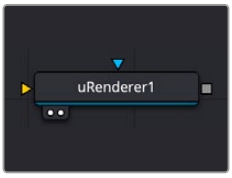


The uMerge node combines the 3D USD scene from the uLoader, two light sources, and the camera, then passes that view out to the uRenderer node to flatten into a 2D image.

Transform and Settings Tab

The Transform and Settings tab in the Inspector is also duplicated in other USD nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

uRenderer [uRn]



The uRenderer node

uRenderer Node Introduction

The uRenderer node converts the USD 3D scene into a 2D image, using either a default perspective camera or one of the cameras found in the scene. Every USD scene in a composition terminates with at least one uRenderer node.

Inputs

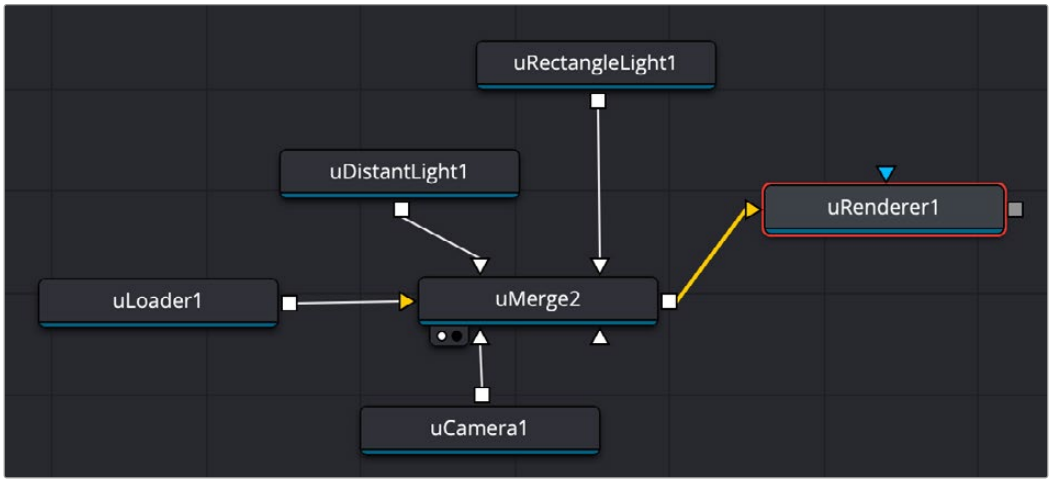
The uRenderer node has two inputs. The main scene input takes in the uMerge or other USD nodes that need to be converted to 2D. The effect mask limits the uRenderer output.

SceneInput: The orange scene input is a required input that accepts a USD scene that you want to convert to 2D.

EffectMask: The blue effects mask input uses a 2D image to mask the output of the node.

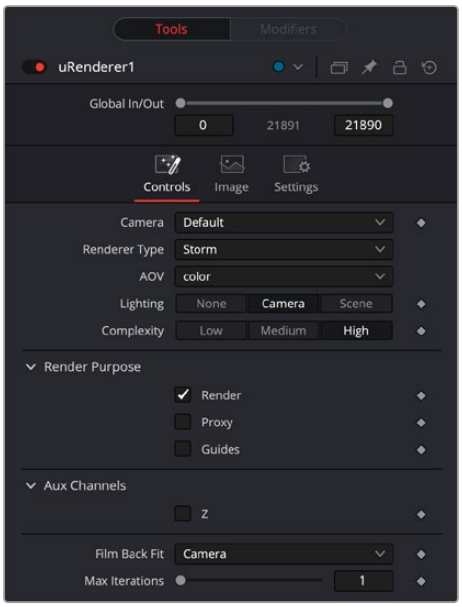
Basic Node Setup

All USD scenes must end with a uRenderer node. The uRenderer node is used to convert a 3D scene into a 2D image. Below, the uRenderer node takes the output of a uMerge node and renders the 3D scene into a 2D image.



The uRenderer node takes the input from the uMerge node and flattens it into a 2D image for display.

Inspector



USD Renderer controls

Controls Tab

Camera

The Camera menu is used to select which camera from the scene is used when rendering. The Default setting uses the first camera found in the scene. If no camera is located, the default perspective view is used instead.

Renderer Type

This menu lists the available render engines. Storm is currently the only renderer available.

AOV

AOV or Arbitrary Output Variable, are additional passes that the render can output; these are also known as Deep Channels or Auxiliary channels. Color is one; others include Depth and Camera Depth.

- **Color:** This outputs the main full-color render of the USD scene. This will include textures and lighting
- **Depth:** This outputs a black-and-white depth map. This particular depth map is normalized for every frame, showing dynamic black to white values.
- **PrimID:** This outputs a numeric value for each different prim in the USD scene.
- **Camera Depth:** This outputs an accurate camera depth map, which calculates the camera's position relative to the USD scene. The image is rendered as a float 32bit image; at first glance this will appear black.

Lighting

Controls how the USD scene is to be lit.

- **None:** Disables all lighting.
- **Camera:** Enabled objects are lit by the default camera light
- **Scene:** Enabled objects are lit by any lights in the USD scene.
- **Enable Sky Dome:** When checked, the sky dome texture is displayed and rendered from the uRenderer.

Complexity

This defines how much detail is rendered: Low, Medium, or High detail.

Aux Channel Z

This option enables rendering of the Z channel, the depth channel that represents the distance of each pixel from the camera. When enabled, the camera depth map is rendered to the Z channel.

Film Back Fit

Defines how the rendered image fits the camera's Film Back. If the resolution of the images being rendered is not the same relative dimensions as the film back then these options set how to fit the image.

Max Iterations

This defines the amount of iterations of the rendering algorithm that will be processed.



USD Renderer Image controls

Image Tab

Process Mode: Lets you choose whether the output of the node will be processed as Full Frames or via one of the specified interlaced methods.

Image: Lets you control the output image size and aspect ratio.

- **Width/Height:** Lets you set the pixel dimensions for the output of the node.
- **Pixel Aspect XY:** Lets you set a custom pixel aspect ratio for non-square video formats.
- **Auto Resolution:** Sets the resolution to the Timeline resolution. Uncheck to manually set the width/height above.
- **Depth:** Lets you manually set the bit depth of the output from the uRenderer.

Render Color Space

Allows you to render the output to Linear or sRGB color spaces.

Source Color Space

The Color Space Type menu sets the color space of the media to help achieve a linear workflow. Unlike the Gamut tool, this doesn‘t perform any actual color space conversion but rather adds the source space data into the metadata, if that metadata doesn‘t already exist. The metadata can then be used downstream by a Gamut tool with the From Image option, or in a Saver if explicit output spaces are defined there.

- **Auto:** Passes along any metadata that might be in the incoming image.
- **Space:** Allows the user to set the color space from a variety of options.

Source Gamma Space

The Curve Type menu automatically determines or allows you to choose the Gamma setting for the image and allows the option to remove the Gamma curve to help achieve a linear workflow.

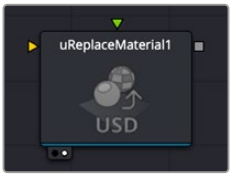
- **Auto:** Passes along any metadata that might be in the incoming image.
- **Space:** Allows you to choose a specific setting from a Gamma Space drop-down menu, while a visual graph lets you see a representation of the gamma setting you’ve selected.
- **Log:** Similar to the Log-Lin node, this option reveals specific log-encoded gamma profiles so that you can select the one that matches your content. A visual graph shows a representation of the log setting you’ve selected. When Cineon is selected from the Log Type menu, additional Lock RGB, Level, Soft Clip, Film Stock Gamma, Conversion Gamma, and Conversion table options are presented to finesse the gamma output.
- **Remove Curve:** Depending on the selected gamma space or on the gamma space found in Auto mode, the associated gamma curve is removed from the material, effectively converting it to output in a linear color space.
- **Pre-Divide/Post-Mujltiply:** Lets you convert “straight” Alpha channels into pre-multiplied Alpha channels, when necessary.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other USD nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

uReplaceMaterial [uRM]



The uReplaceMaterial node

uReplaceMaterial Node Introduction

The uReplaceMaterial node lets you override an object’s material with a new material. This can be a standard color material, texture file, or MaterialX file. The material can be applied to all objects in the scene, or limited to specific objects by using the Prim Selection controls.

Inputs

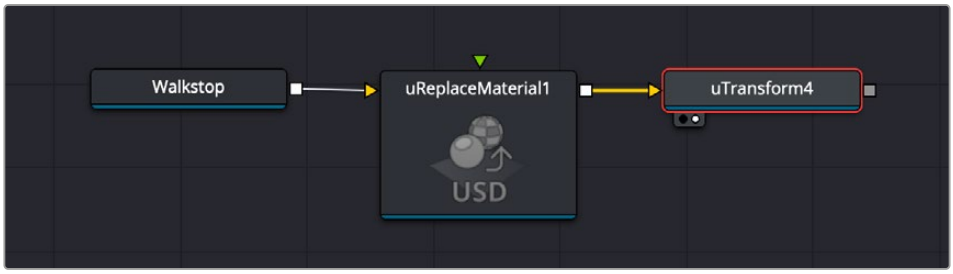
There are two inputs to the uReplaceMaterial node.

Scene Input: The yellow input accepts a USD scene file from a uLoader tool.

Material Input: The green input accepts material files from various tools like uMaterialX, uShader, or even picked from another USD scene using a uLoader tool.

Basic Node Setup

Connect the material output to the rest of the USD tools.



A simple node tree showing a uReplaceMaterial tool accepting input from a USD scene called Walkstop, replacing a material in the scene, then passing that scene to a uTransform node to adjust its transform parameters.

Inspector



The uReplaceMaterial controls

Controls

The uReplaceMaterial controls let you change the materials used in a USD scene.

- **Prim Selection:** Pick opens the Scene Tree window that lets you choose specific materials or objects in a USD scene.

  Invert Prim Selection: Checking this box inverts the Prim Selection to exclude the materials or objects in the Selected Prims box.

- **Type:** Select the type of material you want to replace the original with.

  Color: Chooses a standard Color material and opens the color controls.

  Texture: Presents a file requester for loading a texture file.

  MaterialX: Presents a file requester for loading a MaterialX file, or choosing a Material from a file attached to the green Material input on the node. Choosing this option closes all other controls.

- **Diffuse:** Lets you change the color of the material, or load a texture.

  Color: Opens a standard color picker to choose a new color for the material.

  Red: Lets you adjust the amount of red in the Material.

  Green: Lets you adjust the amount of green in the Material.

  Blue: Lets you adjust the amount of blue in the Material.

- **Emissive:** Emissive adds a global lighting effect to a shape, creating a layer of color over the existing texture. It can be adjusted to different levels of intensity and can be used to simulate the emission of light.

  Color: Opens a standard color picker to choose a new color for the lighting.

  Red: Lets you adjust the amount of red in the lighting.

  Green: Lets you adjust the amount of green in the lighting.

  Blue: Lets you adjust the amount of blue in the lighting.

- **Workflow Mode:** Lets you set the mode between Metallic and Specular.
- **Metallic:** When enabled, this control shows the Metallic slider. The Metallic slider adds a reflective quality to objects, creating the appearance of metal. It enhances the reflective properties of a shape.
- **Specular:** When enabled, this shows the Specular Color controls. Specular Color determines the color of light that reflects from a shiny surface. The more specular a material is, the glossier it appears.
- **Roughness:** Affects the amount of light reflection and scattering on a surface, giving it a smoother or more textured appearance, depending on the setting. It is a key element in creating realistic reflections.
- **Clearcoat:** Affects the appearance of a surface by adding a glossy layer on top of it, mimicking the effect of a protective coating. It is commonly used in creating materials like car paint or polished metal.
- **Clearcoat Roughness:** Adjusts the level of imperfection or smoothness of the glossy layer added on top of a surface. This can affect the realism and reflection of materials like car paint or polished metal.
- **Opacity:** Reduces the material and object opacity, impacting the color and Alpha values of the shape.
- **Opacity Threshold:** Adjusts where the opacity starts to appear.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other USD nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

uShape [uSh]



The USD Shape node

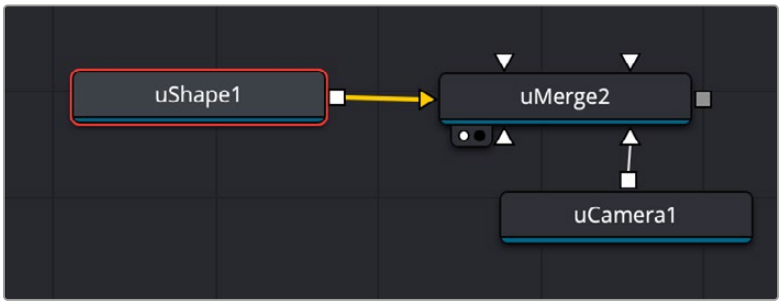
uShape Node Introduction

The uShape node is used to produce several basic primitive 3D shapes, including capsule, torus, planes, cubes, spheres, and cylinders.

Inputs

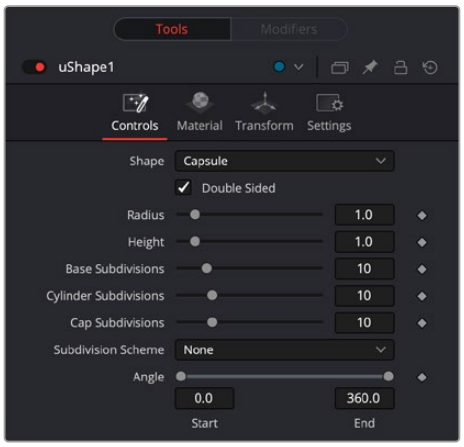
None.

Basic Node Setup



The uShape node generates a 3D shape that is merged with a uCamera node to adjust the view.

Inspector



The USD Shape Controls

Controls Tab

Shape

This menu allows you to select the primitive geometry produced by the Shape 3D node. The remaining controls in the Inspector change to match the selected shape.

The Shape primitives are Capsule, Cone, Cube, Cylinder, Ico sphere, Plane, Sphere and Torus.

Double Sided

This will make the polygons have two faces and receive lighting on both sides.

Lock Width/Height/Depth

Only for Plane and Cube shapes. If this checkbox is selected, the width, height, and depth controls are locked together as a single size slider. Otherwise, individual controls over the size of the shape along each axis are provided.

Size Width/Height/Depth

Only for Plane and Cube shapes. Used to control the size of the shape.

Radius

When a Capsule, Sphere, Cylinder, Cone, or Torus is selected in the shape menu, this control sets the radius of the selected shape.

Height

When a Capsule, Cone, or Cylinder is selected in the shape menu, this control sets the height of the selected shape.

Top Radius

When a cone is selected in the Shape menu, this control is used to define a radius for the top of a cone, making it possible to create truncated cones.

Subdivision Level/Base/Height/Cap/Cylinder

The Subdivision controls are used to determine the tessellation of the mesh on all shapes. The higher the subdivision, the more vertices and polygons in each shape.

Subdivision Scheme

The Subdivision Scheme defines which algorithm is used to tessellate the polygons in the shapes. The methods are None, Bilinear, Loop and Catmull-Clark.

Angle

When the Capsule, Cone, Cylinder, Sphere, or Torus shape is selected in the Shape menu, this range control determines how much of the shape is drawn. A start angle of 180° and end angle of 360° would only draw half of the shape.

Latitude

When a Sphere or Torus is selected in the Shape menu, this range control is used to crop or slice the object by defining a latitudinal subsection of the object.

Cap Bottom/Top

When Cylinder or Cone is selected in the Shape menu, the Bottom Cap and Top Cap checkboxes are used to determine if the end caps of these shapes are created or if the shape is left open

Section

When the Torus is selected in the Shape menu, Section controls the thickness of the tube making up the torus.

Material Tab



The USD Shape Material controls

Diffuse Mode

Diffuse describes the base surface characteristics without any additional effects like reflections or specular highlights.

- **Color:** Provides controls to modify the shape's diffuse color.
- **Texture:** This control allows you to use a source image as the diffuse surface color. Use the Browse button to open the file browser and select an image file.

Emissive

Emissive adds a global lighting effect to a shape, creating a layer of color over the existing texture. It can be adjusted to different levels of intensity and can be used to simulate the emission of light.

Workflow Mode

Lets you set the mode between Metallic and Specular.

- **Metallic:** When enabled, this control shows the Metallic slider. The Metallic slider adds a reflective quality to objects, creating the appearance of metal. It enhances the reflective properties of a shape.
- **Specular:** When enabled, this shows the Specular Color controls. Specular Color determines the color of light that reflects from a shiny surface. The more specular a material is, the glossier it appears.

Roughness

Affects the amount of light reflection and scattering on a surface, giving it a smoother or more textured appearance depending on the setting. It is a key element in creating realistic reflections.

Clearcoat

Affects the appearance of a surface by adding a glossy layer on top of it, mimicking the effect of a protective coating. It is commonly used in creating materials like car paint or polished metal.

Clearcoat Roughness

Adjusts the level of imperfection or smoothness of the glossy layer added on top of a surface. This can affect the realism and reflection of materials like car paint or polished metal.

Opacity

Reduces the material and object opacity, impacting the color and Alpha values of the shape.

Common Controls

Transform and Settings Tab

The Transform and Settings tab in the Inspector is also duplicated in other USD nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

uSwitch [uSw]



The uSwitch node

uSwitch Node Introduction

uSwitch functions similarly to Switch but is specifically designed for the USD environment.

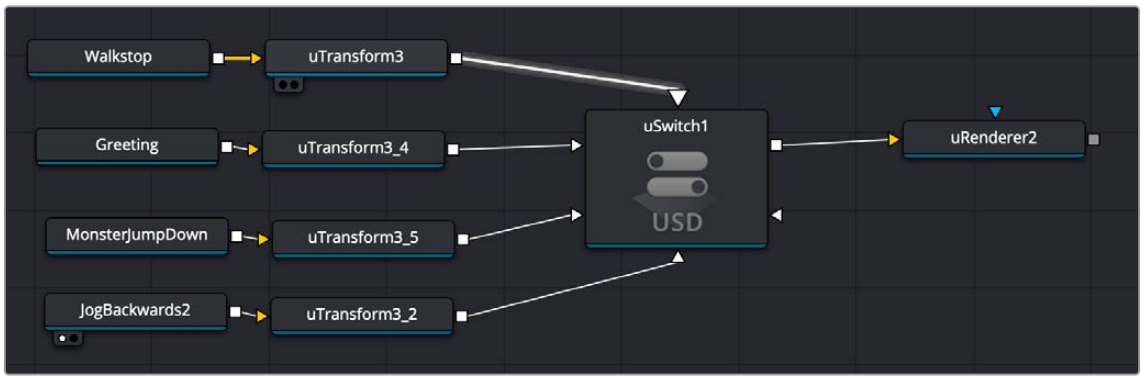
uSwitch is a tool that enables artists to alternate between multiple USD input sources, allowing the selection of one output from the chosen inputs. This capability is especially useful for toggling between different visual elements within a composition. The number of input connections can vary, with the node supporting a dynamic number of inputs that can be renamed as needed. Controls in the Config tab allow users to add or remove inputs and change their names.

Inputs

You can add up to nine inputs to a uSwitch node, using the slider in the Config tools. If you need more than nine sources, you can type that number manually into the Number of Inputs field.

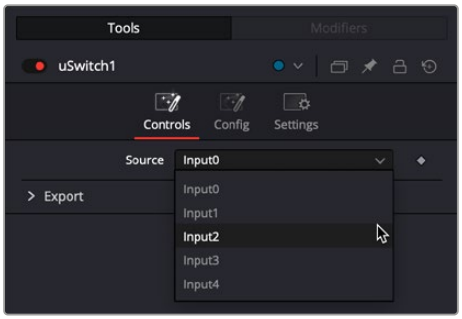
Basic Node Setup

Connect as many input sources as you wish to the uSwitch Node and connect its output to the rest of the node tree.



The uSwitch tool receiving five USD media inputs, allowing you to choose one to go out to uRender.

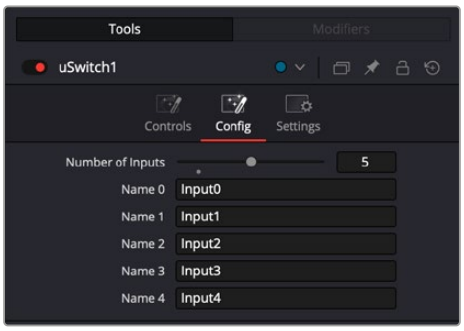
Inspector



The uSwitch controls let you choose a source to pass through the switch.

Controls

The Switch Controls consist of a Source switcher that lets you choose which source to pass through the switch.



The uSwitch configuration

Config

The uSwitch configuration lets you choose the number of inputs and rename them.

- **Number of Inputs:** Adjust the slider to add the number of inputs you wish to use, up to nine. If you want more than nine inputs, you can click in the number field and type the number manually.
- **Name X:** You can change the name of the input from numerical to something more descriptive.

Settings Tab

The Settings tab in the Inspector is also duplicated in other USD nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

uTransform [uXf]



The USD Transform node

uTransform Node Introduction

If you have multiple objects or prims in your USD scene, the uTransform node allows you to choose individual groups of objects and adjust their transform independently from the rest of the scene.

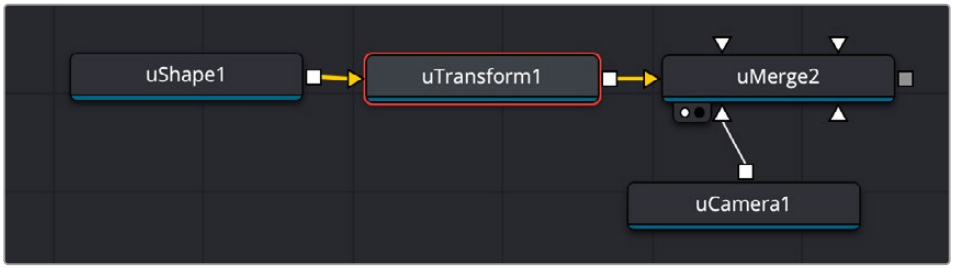
Inputs

The uTransform node has a single required input for a USD scene or USD object.

Scene Input: The orange scene input is connected to a USD scene or USD object to apply a second set of transformation controls.

Basic Node Setup

The uTransform node adds 3D position, rotation, and pivot control onto any existing transforms in the USD node prior to it. You can combine multiple uTransform nodes together to build parenting or hierarchical movement.



The uTransform node gives you another set of 3D transform controls that can be used to modify those in the Transform tab on the uShape node.

Inspector



The USD Transform controls

Controls Tab

Pick

Pick opens the Scene Tree window that lets you choose specific materials or objects in a USD scene. For more information on using the *USD Scene Tree*, see its section at the beginning of this chapter.

Translation

Transform controls are used to position the USD shape in 3D space.

Rotation Order

Use these buttons to select the order used to apply the rotation along each axis of the object. For example, XYZ would apply the rotation to the X-axis first, followed by the Y-axis, and then the Z-axis.

Rotation

Use these controls to rotate the shape around its pivot point. If the Use Target checkbox is selected, then the rotation is relative to the position of the target; otherwise, the global axis is used.

Pivot

A pivot point is the point around which an object rotates. Normally, an object rotates around its own center, which is considered to be a pivot of 0,0,0. These controls can be used to offset the pivot from the center.

Scale

If the Lock X/Y/Z checkbox is checked, a single scale slider is shown. This adjusts the overall size of the object. If the Lock checkbox is unchecked, individual X, Y, and Z sliders are displayed to allow scaling in any dimension.

Scale

Selecting the Use Target checkbox enables a set of controls for positioning an XYZ target. When the target is enabled, the object always rotates to face the target. The rotation of the object becomes relative to the target.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other USD nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

uVariant [uVa]



The uVariant node

uVariant Node Introduction

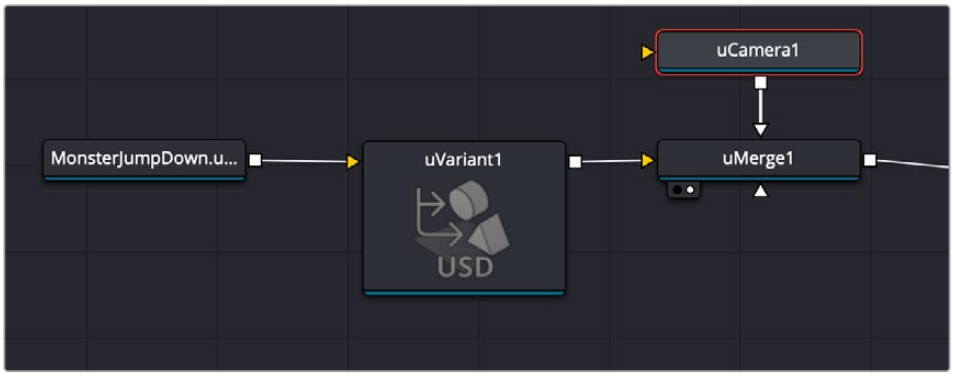
Connecting the uVariant node to the USD loader lets you switch between different alternatives or “variants” for assets that have been set in the USD scene.

Inputs

The uVariant node has a single required input for a USD scene or USD object.

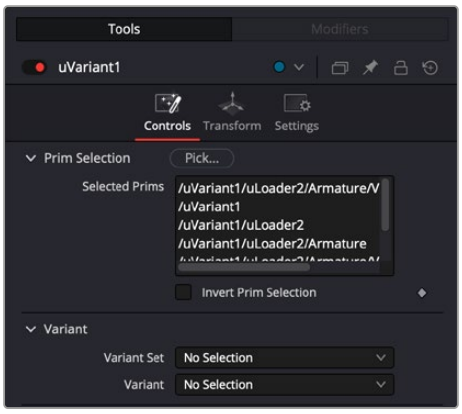
Basic Node Setup

The uVariant node lets you choose between any of the set variants created inside a USD scene. It accepts a USD scene input.



The uVariant node accepts the input from the MonsterJumpDown USD scene.

Inspector



The uVariant controls

Controls Tab

Pick

Pick opens the Scene Tree window that lets you choose specific materials or objects in a USD scene. For more information on using the *USD Scene Tree*, see its section at the beginning of this chapter.

Selected Prims

Shows a list of the prims that have been picked.

Invert Prim Selection

Check this box to swap the selected prims with the unselected ones. This action is animatable using keyframes.

Variant Set

This drop-down menu will show which of the scene elements have variants to choose from. If no variants are available it will show No Selection.

Variant

Once you have selected a variant set above, the variant options to choose between will appear in this dropdown.

uVisibility [uVis]



The uVisibility node

uVisibility Node Introduction

Connecting the uVisibility node to the USD loader lets you control the visibility of an object or branch of objects in a scene.

Inputs

The uVisibility node has a single required input for a USD scene or USD object.

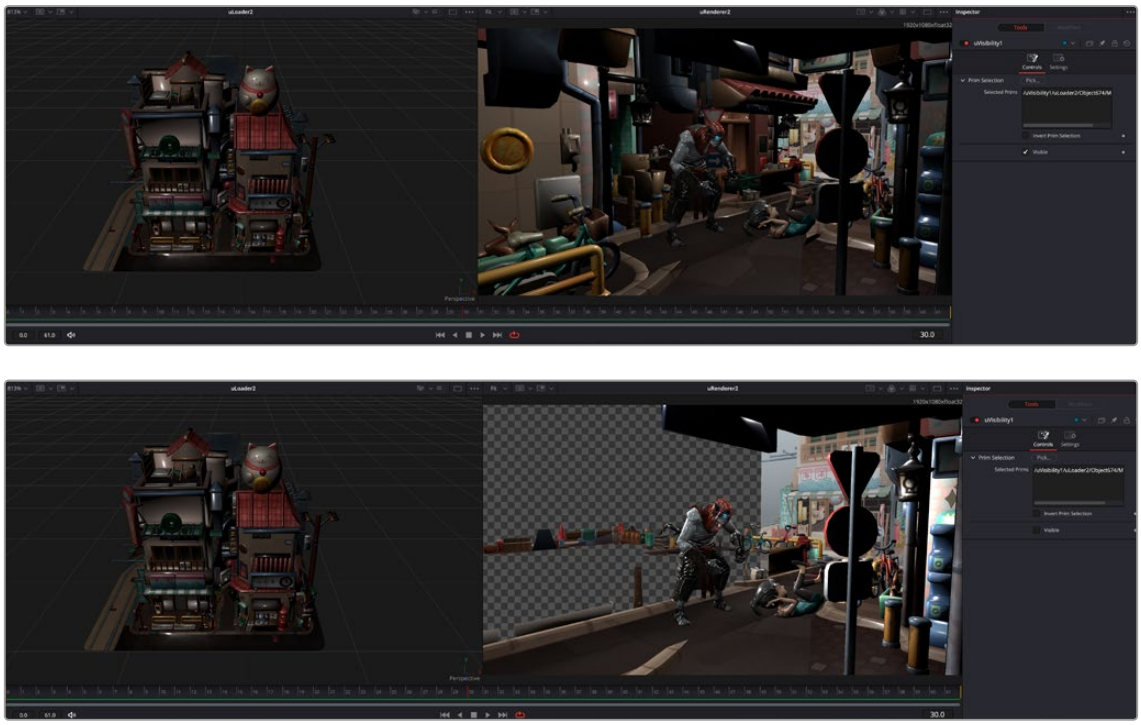
Scene Input: The orange scene input is connected to a USD scene or USD object to apply a second set of transformation controls.

Basic Node Setup

The uVisibility node accepts a USD scene and then outputs the visible elements to the rest of the node tree.

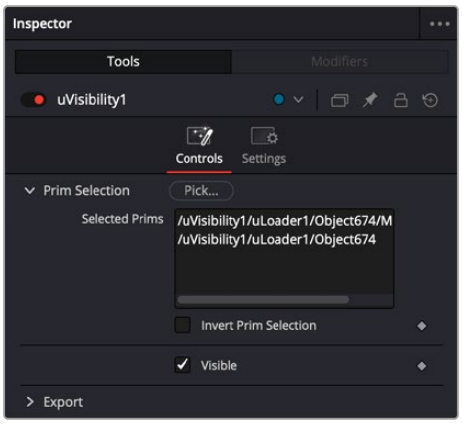


The uVisibility node



Using the uVisibility tool, you can show or hide individual 3D assets. In the above, the cityscape model (left) has been marked visible, and below it is hidden with the visible box unchecked.

Inspector



The uVisibility controls

Controls Tab

Pick

Pick opens the Scene Tree Dialog that lets you choose specific materials or objects in a USD scene. As you select the prims, they will highlight in the viewer. See the beginning of this chapter for information on using the Scene Tree Dialog.

Selected Prims

Shows a list of the prims that have been picked.

Invert Prim Selection

Check this box to swap the selected prims with the unselected ones. This action is animatable using keyframes.

Visible

Toggling this checkbox will add or remove the selected prims from the view. This action is animatable using keyframes.

uVolume [uVo]



The uVolume tool

Introduction

uVolume is a versatile and powerful tool that simplifies the process of importing volumetric VDB files into a USD environment. This node allows you to adjust VDB parameters such as density, temperature, and color, giving you complete control over the appearance and behavior of the imported VDB file.

Usage

When uVolume is added to the Node Editor, a File dialog is displayed automatically allowing the selection of a VDB file from your media drive.

You can add either a single VDB file or a VDB sequence (animated VDB).

Once a VDB is brought in using the uVolume node, the uVolume can be used for trimming, looping, and changing the speed of the imported asset.



The uVolume controls for manipulating VDB files

You can adjust the density with the Scale control, allowing for precise and fine-tuned adjustments.

The Emission section provides color controls for the VDB. Using the Mode selection pulldown will give you the option to select a method for how to change the color of the VDB. Emission modes include:

- **Color:** Applies an overall ‘diffuse’ color to import VDB. Note, USD lights will not affect a VDB with color set to white.
- **Field:** Scatters an overall color to a nominated Emission Field.
- **Blackbody:** Applies color and intensity based on temperature; this is suitable for rendering fire and explosions.
- **Gradient:** Allows you to map a range of colors over a specified field. The Gradient option provides you with a number of controls to target a field, then apply nominated colors to that field.

The uVolume provides the standard transformation controls found on most nodes in Fusion’s USD suite.

USD Lights

USD light tools are a powerful addition to Fusion's USD toolset. They allow for greater control and flexibility when creating and manipulating lighting in your USD scenes. There are six lights available in the USD toolset, they are:

uCylinder Light [uCL]



The USD Cylinder light

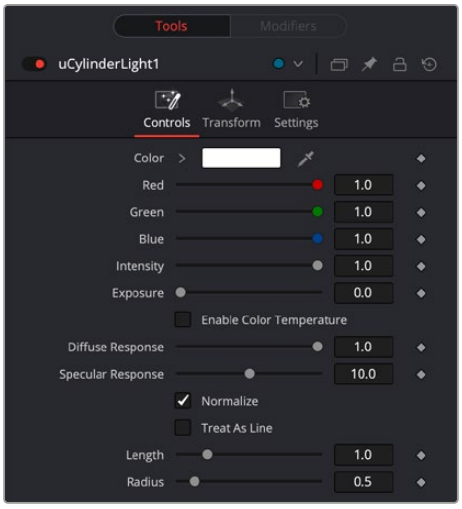
uCylinder Light Node Introduction

A uCylinder light is geometric light that has length and diameter, similar to fluorescent light. This light shows an onscreen control. Translation and rotation of the control place this in the scene.

Inputs

There is one Scene Input in yellow to attach a USD scene; this allows you to adjust any existing lights within the scene.

Inspector



The USD Cylinder Light controls

Controls Tab

The Controls tab is used to set the color and brightness of the uCylinder light. The position, direction, and scale of the light source is controlled in the Transform tab.

Override Selection

Pressing the Pick button lets you choose an existing light in an attached USD file to allow adjusting the lights in an imported scene.

Color

Use this standard Color control to set the color of the light.

Intensity

Use this slider to set the intensity of the light. A value of 0.2 indicates 20% percent light.

Exposure

This will change how much light will expose a scene; this is similar to Intensity.

Color Temperature/Enable

This sets the color temperature of a light source. Its default is 6500K, which is daylight temperature.

Diffuse Response

Controls the amount the light will contribute to the Diffuse color of a material.

Specular Response

Controls the amount the light will contribute to the Specular color of a material.

Normalize

This will normalize the light's contribution in the scene.

Treat As Line

This will make the light a simpler line light source.

Length

Defines the length of the light.

Radius

Defines the diameter of the light.

uDisk Light [uDi]



The USD Disk light

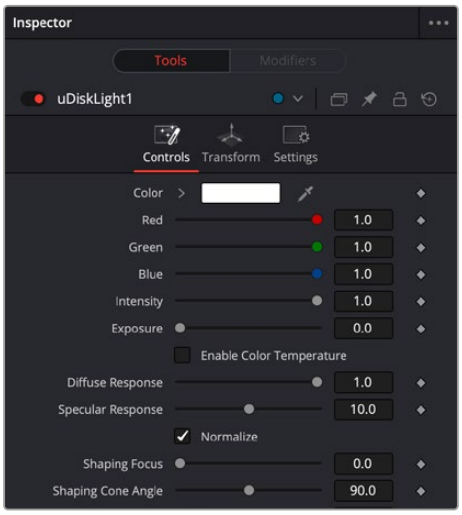
uDisk Light Node Introduction

A uDisk light is a round flat light, similar to an umbrella or softlight. The Controls tab is used to set the color and brightness of the uDisk light. The position, direction, and scale of the light source is controlled in the Transform tab.

Inputs

There is one Scene Input in yellow to attach a USD scene; this allows you to adjust any existing lights within the scene.

Inspector



The USD Disk Light controls

Controls Tab

The Controls tab is used to set the color and brightness of the uDisk light. The position, direction, and scale of the light source is controlled in the Transform tab.

Override Selection

Pressing the Pick button lets you choose an existing light in an attached USD file to allow adjusting the lights in an imported scene.

Color

Use this standard Color control to set the color of the light.

Intensity

Use this slider to set the intensity of the light. A value of 0.2 indicates 20% percent light.

Exposure

This will change how much light will expose a scene; this is similar to Intensity.

Color Temperature / Enable

This sets the color temperature of a light source. Its default is 6500K, which is daylight temperature.

Diffuse Response

Controls the amount the light will contribute to the Diffuse color of a material.

Specular Response

Controls the amount the light will contribute to the Specular color of a material.

Normalize

This will normalize the light's contribution in the scene.

Shaping Focus

This defines focus point of the lens effect of this light.

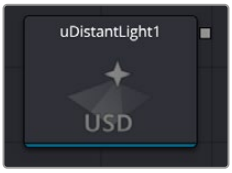
Shaping Cone Angle

Defines the spread of the light.

Radius

Defines the diameter of the light.

uDistant Light [uDL]



The USD Distant light

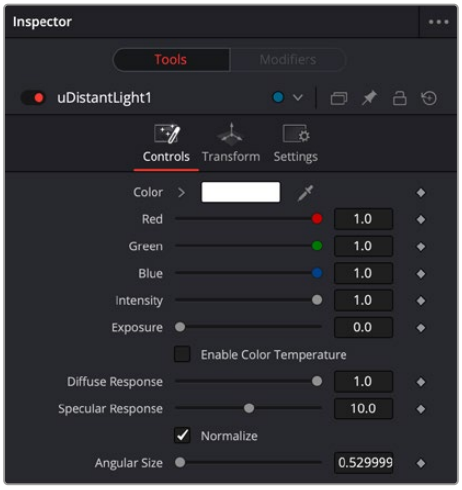
uDistant Light Node Introduction

A uDistant light is similar to a directional light and simulates distance source, similar to sunlight. This light shows an onscreen control. The rotation of the control is used to determine from where in the scene the light direction is pointing.

Inputs

There is one Scene Input in yellow to attach a USD scene; this allows you to adjust any existing lights within the scene.

Inspector



The USD Distant Light controls

Controls Tab

The Controls tab is used to set the color and brightness of the uDistant light. The position, direction, and scale of the light source is controlled in the Transform tab.

Override Selection

Pressing the Pick button lets you choose an existing light in an attached USD file to allow adjusting the lights in an imported scene.

Color

Use this standard Color control to set the color of the light.

Intensity

Use this slider to set the intensity of the light. A value of 0.2 indicates 20% percent light.

Exposure

This will change how much light will expose a scene; this is similar to Intensity.

Color Temperature/Enable

This sets the color temperature of a light source. Its default is 6500K, which is daylight temperature.

Diffuse Response

Controls the amount the light will contribute to the Diffuse color of a material.

Specular Response

Controls the amount the light will contribute to the Specular color of a material.

Normalize

This will normalize the light's contribution in the scene.

Angular Size

Defines the spread of the light.

uDome Light [uDo]



The USD Dome light

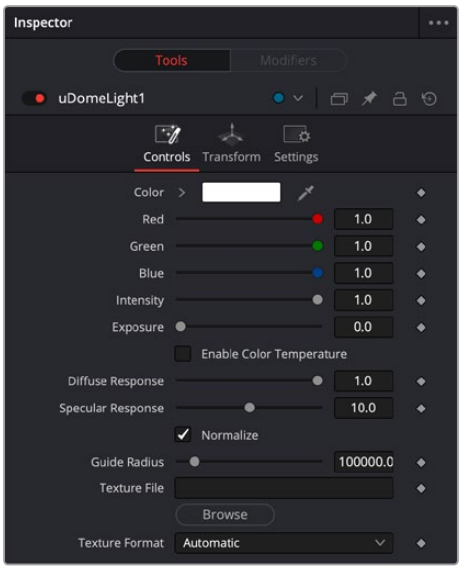
uDome Light Node Introduction

A uDome light uses SDR and HDR images and surrounds the entire scene in a sphere light mapped with the image, similar to environment light. This light shows an onscreen control. The rotation of the control is used to determine from where in the scene the light direction is pointing.

Inputs

There is one Scene Input in yellow to attach a USD scene; this allows you to adjust any existing lights within the scene.

Inspector



The USD Dome Light controls

Controls Tab

The Controls tab is used to set the color and brightness of the uDome light. The position, direction, and scale of the light source is controlled in the Transform tab.

Override Selection

Pressing the Pick button lets you choose an existing light in an attached USD file to allow adjusting the lights in an imported scene.

Color

Use this standard Color control to set the color of the light.

Intensity

Use this slider to set the intensity of the light. A value of 0.2 indicates 20% percent light.

Exposure

This will change how much light will expose a scene; this is similar to Intensity.

Color Temperature/Enable

This sets the color temperature of a light source. Its default is 6500K, which is daylight temperature.

Diffuse Response

Controls the amount the light will contribute to the Diffuse color of a material.

Specular Response

Controls the amount the light will contribute to the Specular color of a material.

Normalize

This will normalize the light's contribution in the scene.

Guide Radius

This sets how large the dome is around the scene. The default is set to large and far away, similar to the outdoors. This radius can be set smaller to simulate a room size environment.

Texture File/Browse

This will load an image to map on the Domelight.

Texture Format

Sphere mapped images can be stored in different schemes; the common schemes are Lat-Long, MirrorBall, Angular, and Cube Mapped Vertical Cross.

uRectangle Light [uRL]



The USD Rectangle light

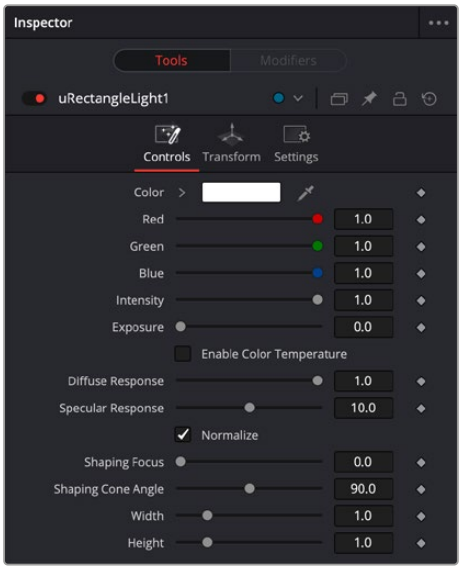
uRectangle Light Node Introduction

A uRectangle light is a flat light, similar to softlight or window . The Controls tab are used to set the color and brightness of the uRectangle light. The position, direction, and scale of the light source is controlled in the Transform tab.

Inputs

There is one Scene Input in yellow to attach a USD scene; this allows you to adjust any existing lights within the scene.

Inspector



The USD Rectangle Light controls

Controls Tab

The Controls tab is used to set the color and brightness of the uDome light. The position, direction, and scale of the light source is controlled in the Transform tab.

Override Selection

Pressing the Pick button lets you choose an existing light in an attached USD file to allow adjusting the lights in an imported scene.

Color

Use this standard Color control to set the color of the light.

Intensity

Use this slider to set the intensity of the light. A value of 0.2 indicates 20% percent light.

Exposure

This will change how much light will expose a scene; this is similar to Intensity.

Color Temperature/Enable

This sets the color temperature of a light source. Its default is 6500K, which is daylight temperature.

Diffuse Response

Controls the amount the light will contribute to the Diffuse color of a material.

Specular Response

Controls the amount the light will contribute to the Specular color of a material.

Normalize

This will normalize the light's contribution in the scene.

Shaping Focus

This defines focus point of the lens effect in this light.

Shaping Cone Angle

Defines the angle of spread of the light.

Width

Defines how wide the light appears.

Height

Defines how high the light appears.

uSphere Light [uSL]



The USD Sphere light

uSphere Light Node Introduction

A uSphere light is local light that is similar to point light but also has radius. This light shows an onscreen control. The rotation of the control is used to determine from where in the scene the light direction is pointing.

Inputs

There is one Scene Input in yellow to attach a USD scene; this allows you to adjust any existing lights within the scene.

Inspector



The USD Sphere Light controls

Controls Tab

The Controls tab is used to set the color and brightness of the uSphere light. The position, direction, and scale of the light source is controlled in the Transform tab.

Override Selection

Pressing the Pick button lets you choose an existing light in an attached USD file to allow adjusting the lights in an imported scene.

Color

Use this standard Color control to set the color of the light.

Intensity

Use this slider to set the intensity of the light. A value of 0.2 indicates 20% percent light.

Exposure

This will change how much light will expose a scene; this is similar to Intensity.

Color Temperature/Enable

This sets the color temperature of a light source. Its default is 6500K, which is daylight temperature.

Diffuse Response

Controls the amount the light will contribute to the Diffuse color of a material.

Specular Response

Controls the amount the light will contribute to the Specular color of a material.

Normalize

This will normalize the light's contribution in the scene.

Treat as Point

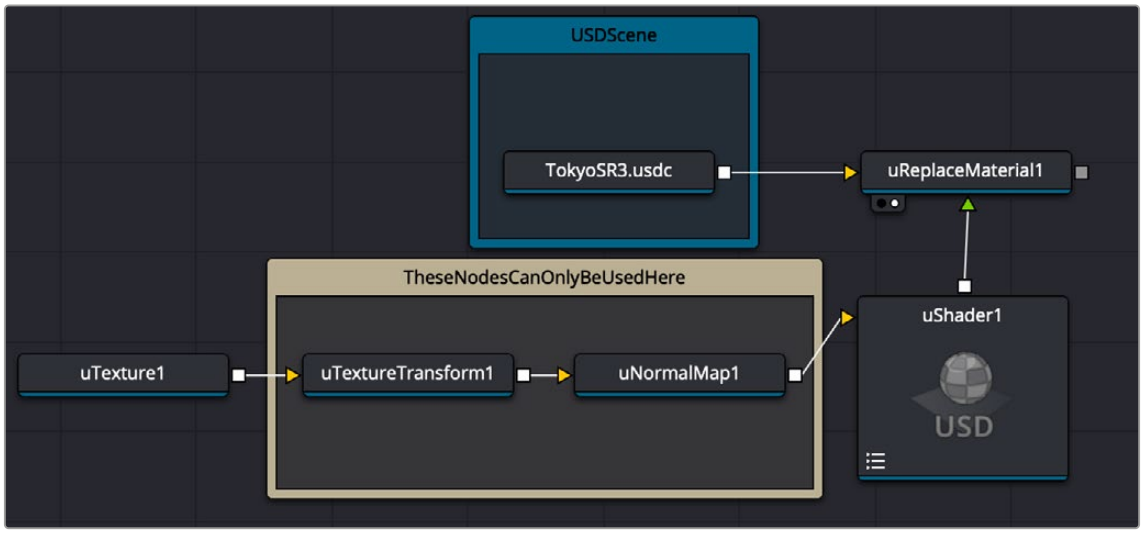
This will make the light a simple classic point source.

Radius

This sets the size of the sphere light.

USD Material

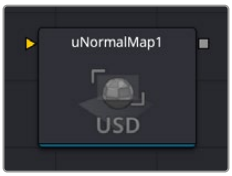
The uTexture, uTextureTransform, uNormalMap, and uShader nodes give you the ability to import and manipulate textures directly in your USD scene without having to wrap them in a USD file first. These tools add flexibility and control to your USD workflow within the Fusion page.



A sample texture replacement node tree using the USD Materials tools

The uTexture, uTextureFormat, and uNormalMap tools are designed to work around and feed a texture to the uShader tool, where it is manipulated using either the USD or MaterialX models to change its appearance. The uShader is commonly connected to a uReplaceMaterial node, which selects the specific parts of the entire USD scene to apply the new texture to. This workflow lets you quickly iterate texture changes internally in Fusion, without requiring the modification of the original USD scene in external 3D modeling software.

uNormalMap [uNm]

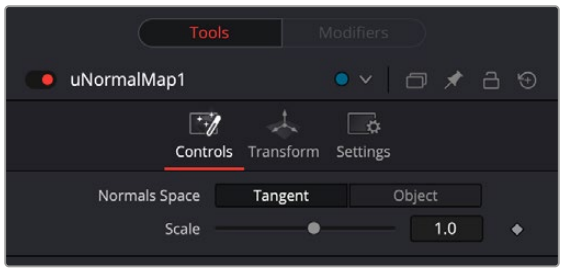


The USD Normal Map

Introduction

The uNormalMap tool provides several controls to modify a “normals” texture before it's passed on to the uShader. This allows you to manipulate the texture in various ways and ensure that it is optimized for your specific needs. Once any necessary adjustments have been made, the modified image is then passed on to the uShader.

Usage



The uNormal Map controls

To use uNormalMap, place it after a uTexture node and before a uShader node. This node node only works between the the loaded texture and uShader.

uShader [uSd]



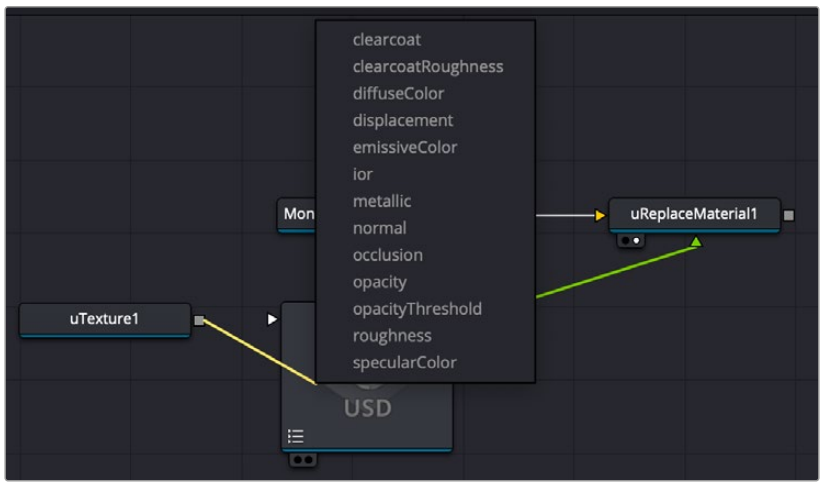
The USD Shader

uShader Node Introduction

The uShader node is used to create materials and looks using either the USD or MaterialX shader tools. Used in conjunction with uReplaceMaterial, this tool gives an easy and intuitive way to iterate material changes to a USD source.

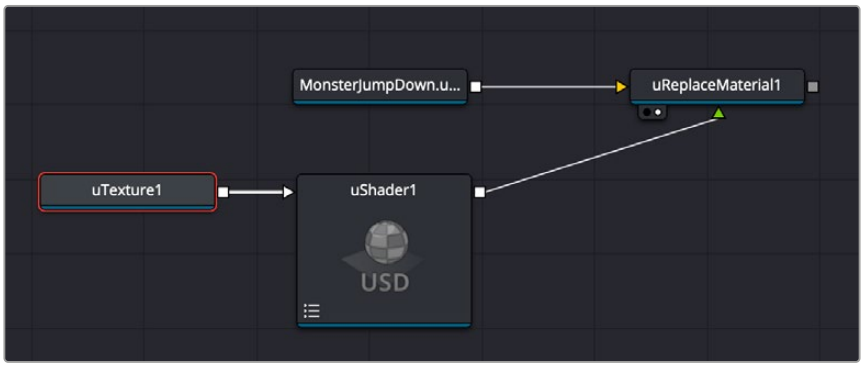
Inputs

The uShader node has variable inputs that allow you to import different assets into the tool, such as a texture, by using a uTexture node. When connecting the tool, you can choose from a drop list of all possible connections. Connecting a texture to any input will disable its corresponding slider and expose new Channel, Scale, and Bias controls to let you choose which channel the property uses, with their strength and base values.



Attaching a uTexture node to the uShader node presents a drop list, showing all possible inputs to connect to.

Basic Node Setup



The uShader is accepting an input from the uTexture node, modifying that texture, then passing it to the uReplace Material node. uReplace Material is selecting specific materials from the MonsterJumpDown USD scene to apply the newly modified texture to.

Inspector



The uShader controls in USD mode

Shader Model

Here you can choose what toolset to work with, either USD or MaterialX tools.

Shader Model: USD

Diffuse Mode

Diffuse describes the base surface characteristics without any additional effects like reflections or specular highlights.

- **Color:** Provides controls to modify the shape's diffuse color.

Emissive

Emissive adds a global lighting effect to a shape, creating a layer of color over the existing texture. It can be adjusted to different levels of intensity and can be used to simulate the emission of light.

Workflow Mode

Lets you set the mode between Metallic and Specular.

- **Metallic:** When enabled, this control shows the Metallic slider. The Metallic slider adds a reflective quality to objects, creating the appearance of metal. It enhances the reflective properties of a shape.
- **Specular:** When enabled, this shows the Specular Color controls. Specular Color determines the color of light that reflects from a shiny surface. The more specular a material is, the glossier it appears.

Roughness

Affects the amount of light reflection and scattering on a surface, giving it a smoother or more textured appearance, depending on the setting. It is a key element in creating realistic reflections.

Clearcoat

Affects the appearance of a surface by adding a glossy layer on top of it, mimicking the effect of a protective coating. It is commonly used in creating materials like car paint or polished metal.

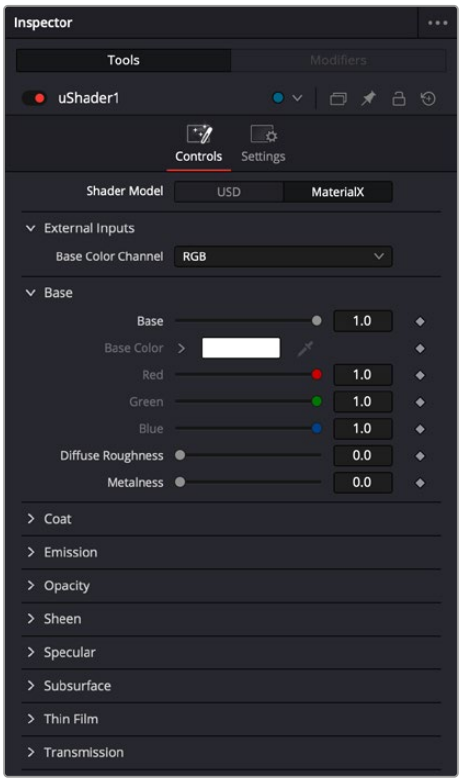
Clearcoat Roughness

Adjusts the level of imperfection or smoothness of the glossy layer added on top of a surface. This can affect the realism and reflection of materials like car paint or polished metal.

Opacity

Reduces the material and object opacity, impacting the color and Alpha values of the shape.

Shader Model: MaterialX



The uShader controls in MaterialX mode

Controls

MaterialX is an open standard maintained by the Academy Software Foundation designed for creating material and look development content that is platform independent. MaterialX is under constant development, and you can find the latest specifications at the MaterialX website: materialx.org

uTexture [uTx]

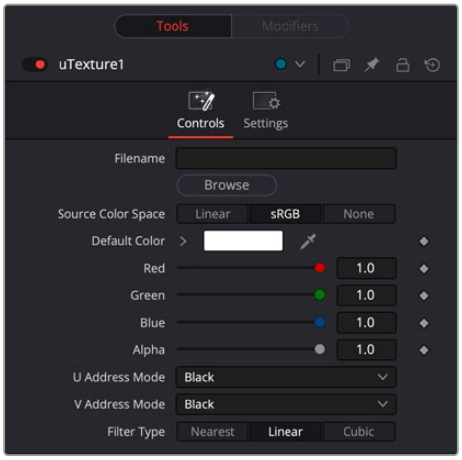


The USD Texture

Introduction

To create a material/surface in the uShader, you can import a texture file using this tool. The uTexture sources image files and uses them to build the material/surface in the uShader.

Usage

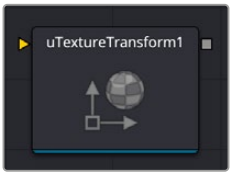


The uTexture controls

After adding the node to the Node Editor, use the Browse button to select an image file from your hard drive. Alternatively, you can click and drag an image from the Media Pool directly onto the “File Name” control in the Inspector to load an image which has already been imported into DaVinci Resolve.
Use the “Source Color Space” control to choose how the image will be read by the uTexture tool. It’s important select Linear when loading a Normals image to get the best results.

U/V Address Mode controls offer a useful way to determine how the edges of the image are treated when exposed due to scaling.

uTextureTransform [uTXf]

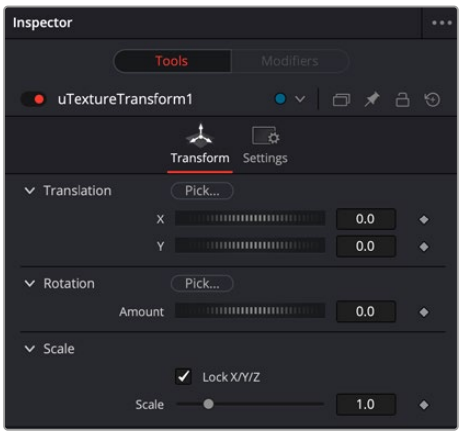


The USD Texture Transform

Introduction

uTextureTransform provides you with granular control over the scale, rotation, and position of each loaded texture file. With uTexture Transform, you can quickly and easily tweak the size, orientation, and placement of any texture file sourced from the uTexture node.

Usage



The uTexture Transform controls

To use uTextureTransform, place it after a uTexture node and before a uShader node. This node only works between the the loaded texture and uShader.

Adjust the Scale, Rotation, and Position settings to your liking.

The Scale setting allows you to increase or decrease the size of the texture file, while the Position setting lets you move the file around within your project. You can also use the Rotation setting to rotate the texture file to any angle you desire.

The Common Controls

Nodes that handle USD operations share several identical controls in the Inspector. This section describes controls that are common among USD nodes.

Transform Tab

The Transform tab can be found in the Inspector in most USD nodes. The controls are the same as the uTransform node, which can be applied separately.

Translation

Transform controls are used to position the USD shape in 3D space.

Rotation Order

Use these buttons to select the order used to apply the rotation along each axis of the object. For example, XYZ would apply the rotation to the X-axis first, followed by the Y-axis, and then the Z-axis.

Rotation

Use these controls to rotate the shape around its pivot point. If the Use Target checkbox is selected, then the rotation is relative to the position of the target; otherwise, the global axis is used.

Pivot

A pivot point is the point around which an object rotates. Normally, an object rotates around its own center, which is considered to be a pivot of 0,0,0. These controls can be used to offset the pivot from the center.

Scale

If the Lock X/Y/Z checkbox is checked, a single scale slider is shown. This adjusts the overall size of the object. If the Lock checkbox is unchecked, individual X, Y, and Z sliders are displayed to allow scaling in any dimension.

Settings Tab

The Settings tab in the Inspector can be found on every tool in the USD category. The controls are consistent and work the same way for each tool.

Hide Incoming Connections

Enabling this checkbox can hide connection lines from incoming nodes, making a node tree appear cleaner and easier to read. When enabled, empty fields for each input on a node will be displayed in the Inspector. Dragging a connected node from the node tree into the field will hide that incoming connection line as long as the node is not selected in the node tree. When the node is selected in the node tree, the line will reappear.

Comments

The Comments field is used to add notes to a tool. Click in the empty field and type the text. When a note is added to a tool, a small red square appears in the lower-left corner of the node when the full tile is displayed, or a small text bubble icon appears on the right when nodes are collapsed. To see the note in the Node Editor, hold the mouse pointer over the node to display the tooltip.

Scripts

Three Scripting fields are available on every tool in Fusion from the Settings tab. They each contain edit boxes used to add scripts that process when the tool is rendering. For more details on scripting nodes, please consult the Fusion scripting documentation.

Chapter 62

VR Nodes

This chapter details the Virtual Reality (VR) nodes available in Fusion.

The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

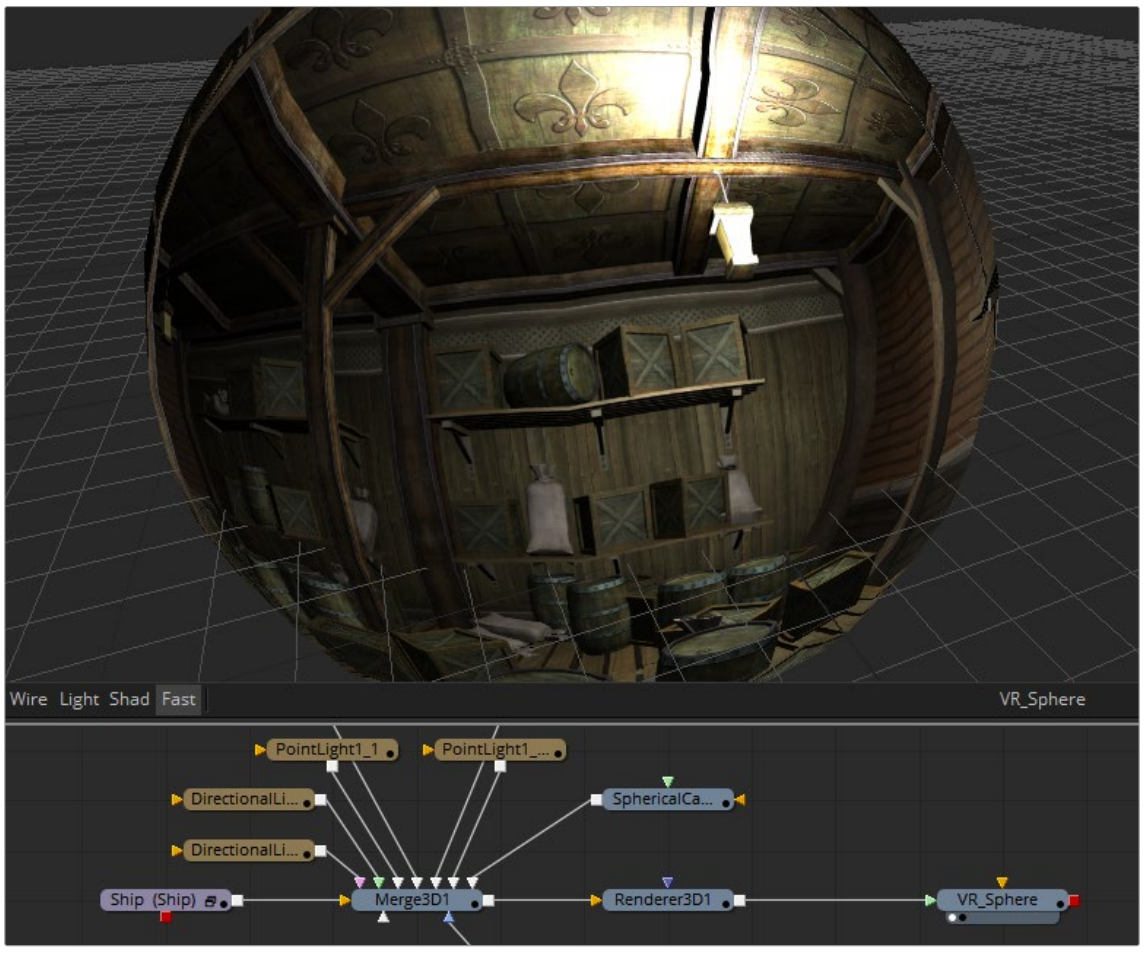
VR Nodes 1704
Fusion VR 180 Support (Studio Version Only) 1705
Lat Long Patcher [LLP] 1705
Pano Map [PaM] 1707
Spherical Camera [3SC] 1709
Spherical Stabilizer 1709
The Common Controls 1712

VR Nodes

You can create and fix spherical (360°) video, often described as Virtual Reality, or VR, using Fusion's set of VR nodes. Dome productions, planetariums, and other special-venue theaters have utilized the flexibility of Fusion and its 3D system to produce and deliver special content for years.

The equirectangular (lat-long) format often used for 360° video is similar to how a globe is represented by a flat world map, with the poles at the top and bottom edges of the image and the forward viewpoint at the center.

TIP: You can create stereo VR using two stacked Lat Long images, one for each eye.



A VR Sphere displayed in the Node Editor and Fusion's 3D viewer

Fusion supports several common spherical image formats and can easily convert between them.

VCross and HCross: VCross and HCross are the six square faces of a cube laid out in a cross, vertically or horizontally, with the forward view in the center of the cross in a 3:4 or 4:3 image.

VStrip and HStrip: VStrip and HStrip are the six square faces of a cube laid vertically or horizontally in a line, ordered as Left, Right, Up, Down, Back, Front (+X, -X, +Y, -Y, +Z, -Z) in a 1:6 or 6:1 image.

LatLong: LatLong is a single 2:1 image in an equirectangular mapping.

You can display both spherical video and live 3D scenes from the comp directly to headsets, including those from Oculus Rift and HTC Vive.

Fusion's "Fix it in post" tools for VR make it easy to perform several important tasks that are common in these types of productions.

Fusion VR 180 Support (Studio Version Only)

A number of Fusion's existing VR tools now support 180-degree angle of view. The 180-degree VR format allows you to produce immersive content and simplifies the production process for VR experiences by focusing on a forward-facing perspective.

Lat Long Patcher [LLP]



The Lat Long Patcher node

NOTE: The VR category and Lat Long node are available only in Fusion Studio and DaVinci Resolve Studio.

Lat Long Patcher Node Introduction

Equirectangular stitched images often need patches, paintwork, or other VFX applied. The Lat Long Patcher extracts and de-warps a section of a lat-long (equirectangular) image to be treated, and can warp and merge fixes back over the original. You can quickly pick a section of the spherical image to patch or paint, and then apply it back to the original image. Note that matching rotations are used in both Extract and Apply modes, allowing a node's operation to be easily reversed by a copy or instance with the same rotation settings.

Input

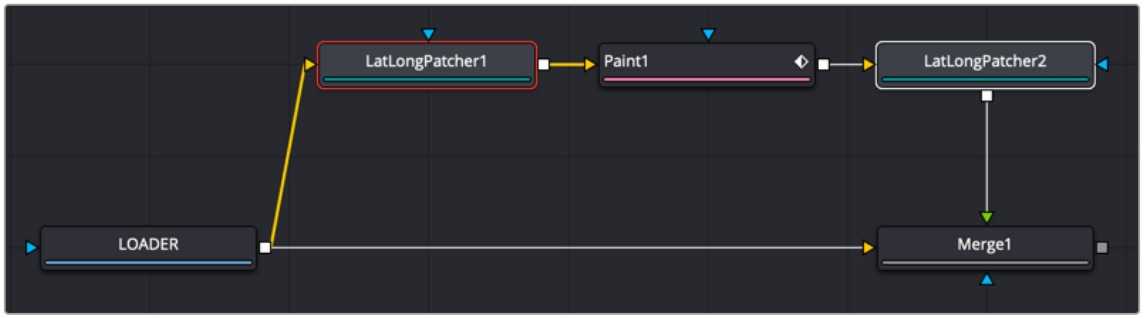
The Lat Long Patcher node includes two inputs. The orange input accepts a 2D image in an equirectangular format, where the X-axis represents 0–360 degrees longitude, and the Y-axis represents –90 to +90 degrees latitude. The effect mask input is provided, although rarely used on VR nodes.

Image Input: The orange image input accepts a equirectangular (lat-long) 2D RGBA image.

Effect Mask: The effect mask input is provided, although rarely used on VR nodes.

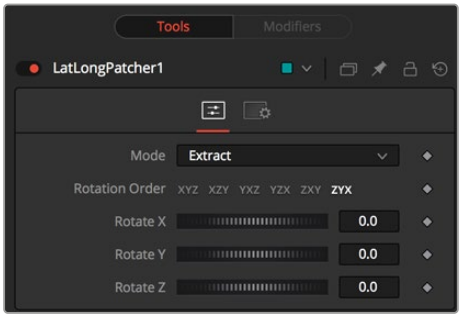
Basic Node Setup

The Loader node connects to the input on a Lat Long Patcher node. The output of a Lat Long Patcher node is set to Extract. It is then connected to whatever image-processing operation is required. A second Lat Long Patcher node set to Apply takes an input from the processed extraction and merges it over the top of the original source.



Two Lat Long Patchers used to repair a section

Inspector



The Lat Long Patcher Controls tab

Controls Tab

The Controls tab is used to extract and later reapply a section from an equirectangular image. Rotation controls allow you to select the exact portion you need to repair.

Mode

- **Extract:** Pulls a de-warped 90-degree square image from the equirectangular image.
- **Apply:** Warps and merges a 90-degree square image over the equirectangular image. Because the square image's alpha is used, this allows, for example, paint strokes or text drawn over a transparent black background to be applied to the original equirectangular image, avoiding any double-filtering from de-warping and re-warping the original.
- **Apply180:** Use to extract and apply a patch to a VR180 image. This gives control to apply traditional 2D effects, like paint, to a 180 distorted clip.

Rotation Order

These buttons choose the ordering of the rotations around each axis. For example, XYZ rotates first around the X axis (pitch/tilt), then around the Y axis (pan/yaw), and then around the Z axis (roll). Any of the six possible orderings can be chosen.

Rotation

These dials rotate the spherical image around each of the X, Y, and Z axes, offering independent control over pitch/tilt, pan/yaw, and roll, respectively.

Common Controls

Settings Tab

The Settings tab in the Inspector is duplicated in other VR nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Pano Map [PaM]



The Pano Map node

NOTE: The VR category and Pano Map node are available only in Fusion Studio and DaVinci Resolve Studio.

Pano Map Node Introduction

The Pano Map node converts images from one spherical layout to another, such as from a cube map to an equirectangular format. The node can also perform rotations of the spherical images when converting.

Input

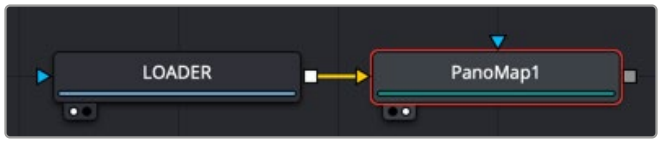
The Pano Map node includes two inputs. The orange input accepts a 2D image in an equirectangular, cube map or other spherical formats. The effect mask input is provided, although rarely used on VR nodes.

Image Input: The orange Image input accepts a spherical formatted 2D RGBA image that gets converted to another spherical format.

Effect Mask: The effect mask input is provided, although rarely used on VR nodes.

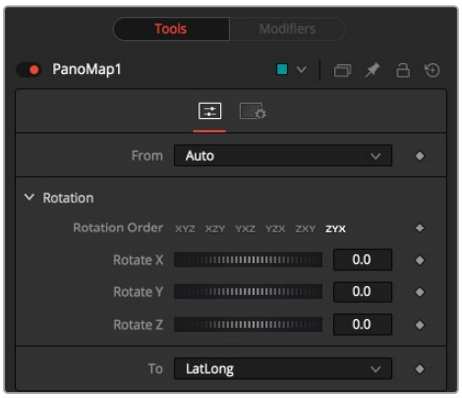
Basic Node Setup

In the example below, a Loader node containing a Lat Long image connects to the input on a Pano Map node. The Pano Map node is used to convert the image to an H-Cross format. It is then connected to whatever image-processing operation is required.



A Pano Map used to convert from one VR format to another

Inspector



The Pano Map Controls tab

Controls Tab

The Controls tab is used to determine the format of the input image and the desired output format.

From/To

- **Auto:** Auto detects the incoming image layout from the metadata and image frame aspect.
- **VCross and HCross:** VCross and HCross are the six square faces of a cube laid out in a cross, vertically or horizontally, with the forward view in the center of the cross in a 3:4 or 4:3 image.
- **VStrip and HStrip:** VStrip and HStrip are the six square faces of a cube laid vertically or horizontally in a line, ordered as Left, Right, Up, Down, Back, Front (+X, -X, +Y, -Y, +Z, -Z) in a 1:6 or 6:1 image.
- **LatLong:** LatLong is a single 2:1 image in equirectangular mapping.
- **VR 180** is a 1:1 180 degree stereoscopic image.

Rotation Order

These buttons choose the ordering of the rotations around each axis. For example, XYZ rotates first around the X axis (pitch/tilt), then around the Y axis (pan/yaw), and then around the Z axis (roll). Any of the six possible orderings can be chosen.

Rotation

These dials rotate the spherical image around each of the X, Y, and Z axes, offering independent control over pitch/tilt, pan/yaw, and roll, respectively.

Common Controls

Settings Tab

The Settings tab in the Inspector is duplicated in other VR nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Spherical Camera [3SC]



The Spherical Camera node

Spherical Camera Node Introduction

The Spherical Camera is not located in the VR category of the Effects Library but in the 3D category. However, it is commonly used in creating and fixing VR content, so it is referenced here. The Spherical Camera allows the 3D Renderer node to output an image covering all viewing angles, laid out in several different formats. This image may be used, for example, as a skybox texture or reflection map or viewed in a VR headset. The Image Width setting in the 3D Renderer sets the size of each square cube face so that the resulting image may be a multiple of this size horizontally and vertically.

For more detail on the Spherical Camera node, *see Chapter 29, "3D Nodes,"* in the Fusion Reference Manual.

Spherical Stabilizer



The Spherical Stabilizer node

NOTE: The VR category and Spherical Stabilizer node are available only in Fusion Studio and DaVinci Resolve Studio.

Spherical Stabilizer Node Introduction

VR live action often uses handheld cameras, causing shaky footage to be a common problem. The Spherical Stabilizer node automatically identifies and tracks visible features in the footage, and then analyzes their movement to identify pan, tilt, and roll rotations. After tracking, it is then possible to smooth out or stabilize the rotation of the footage.

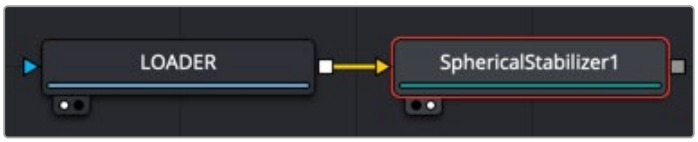
Inputs

The Spherical Stabilizer node has a single orange input.

Image: This orange image input node requires an image in a spherical layout, which can be any of Lat Long (2:1 equirectangular), VR180, Horizontal/Vertical Cross, or Horizontal/Vertical Strip.

Basic Node Setup

In the example below, a 2:1 Lat Long image is connected to the input of the Spherical Stabilizer node. Once the image is stabilized, the output of the Spherical Stabilizer node is a steadied clip.



Spherical Stabilizer set up to steady a 2:1 Lat Long clip

Inspector



The Spherical Stabilizer Controls tab

Controls Tab

The Controls tab contains parameters to initiate the tracking and modify the results for stabilization or smoothing.

Reject Dominant Motion Outliers While Tracking

With this control activated (the default setting), features that move contrary to the majority of other features are ignored. This helps ignore the movement of subjects in the shot, preferring stable and consistent markers from the surrounding environment.

Track Controls

These buttons initiate tracking and analysis of the shot. Be aware that the reference frame used for stabilization is set to the first frame tracked.

- Track Backward from End Frame starts tracking backward from the end of the current render range.
- Track Backward from Current Time starts tracking backward from the current frame.
- Stop ceases tracking, preserving all results so far.
- Track Forward from Current Time starts tracking forward from the start of the current render range.
- Track Forward from Start Frame starts tracking forward from the current time.

Append to Track

- Replace causes the Track Controls to discard any previous tracking results and replace them with the newly-created track.
- Append adds the new tracking results to any earlier tracks.

Stabilization Strength

This control varies the amount of smoothing or stabilization applied, from 0.0 (no change) to 1.0 (maximum).

Smoothing

The Spherical Stabilizer node can eliminate all rotation from a shot, fixing the forward viewpoint (Still mode, 0.0) or gently smooth out any panning, rolling, or tilting to increase viewer comfort (Smooth mode, 1.0). This slider allows either option or anything in between.

Offset Rotation

Often a shot is not entirely level and needs the horizon to be realigned, or perhaps a desired pan should be reintroduced after fully stabilizing the shot. The Offset Rotation controls allow additional manual control of the Spherical Stabilizer's rotation of the footage, for pitch/tilt (X), pan/yaw (Y), and roll (Z), respectively. Rotation is always performed in the order X, Y, Z.

Common Controls

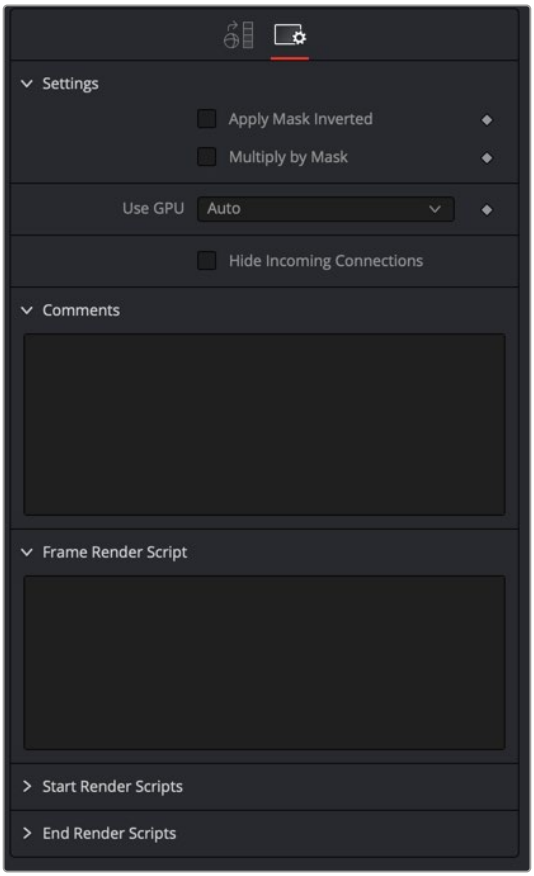
Settings Tab

The Settings tab in the Inspector is duplicated in other VR nodes. These common controls are described in detail in the following "The Common Controls" section.

The Common Controls

Nodes that handle VR operations share a number of identical controls in the Inspector. This section describes controls that are common among VR nodes.

Inspector



The VR Common controls

Settings Tab

The Settings tab in the Inspector can be found on every tool in the VR category. The controls are consistent and work the same way for each tool.

Apply Mask Inverted

Enabling the Apply Mask Inverted option inverts the complete mask channel for the tool. The mask channel is the combined result of all masks connected to or generated in a node.

Use GPU

The Use GPU menu has three settings. Setting the menu to Disable turns off hardware-accelerated rendering using the graphics card in your computer. Auto uses a capable GPU if one is available and falls back to software rendering when a capable GPU is not available.

Hide Incoming Connections

Enabling this checkbox can hide connection lines from incoming nodes, making a node tree appear cleaner and easier to read. When enabled, empty fields for each input on a node are displayed in the Inspector. Dragging a connected node from the node tree into the empty field hides that incoming connection line as long as the node is not selected in the node tree. When the node is selected in the node tree, the line reappears.

Comments

The Comments field is used to add notes to a tool. Click in the empty field and type the text. When a note is added to a tool, a small red square appears in the lower-left corner of the node when the full tile is displayed, or a small text bubble icon appears on the right when nodes are collapsed. To see the note in the Node Editor, hold the mouse pointer over the node to display the tooltip.

Scripts

Three Scripting fields are available on every tool in Fusion from the Settings tab. They each contain edit boxes used to add scripts that process when the tool is rendering. For more details on scripting nodes, please consult the Fusion scripting documentation.

Chapter 63

Warp Nodes

This chapter details the Warp nodes available in Fusion.

The abbreviations next to each node name can be used in the Select Tool dialog when searching for tools and in scripting references.

For purposes of this document, node trees showing MediaIn nodes in DaVinci Resolve are interchangeable with Loader nodes in Fusion Studio, unless otherwise noted.

Contents

Coordinate Space [CdS] 1715
Corner Positioner [CPn] 1717
Dent [Dnt] 1719
Displace [Dsp] 1720
Drip [DRP] 1723
Grid Warp [Grd] 1725
Lens Distort [Lens] 1732
Perspective Positioner [PPn] 1735
Vector Distortion [Dst] 1736
Vortex [Vtx] 1739
The Common Controls 1741

Coordinate Space [CdS]



The Coordinate Space node

Coordinate Space Node Introduction

The Coordinate Space node changes the coordinate space of the image from rectangular to polar or from polar to rectangular.

Inputs

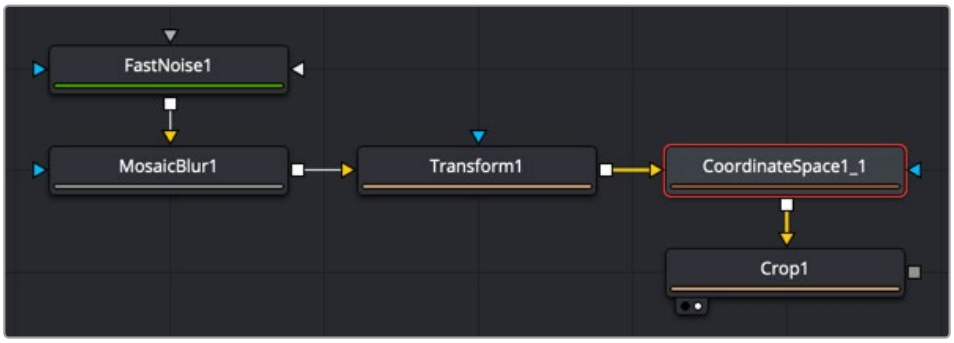
The two inputs on the Coordinate Space node are used to connect a 2D image and an effect mask, which can be used to limit the distorted area.

Input: The orange input is used for the primary 2D image that is distorted.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the distortion to only those pixels within the mask. An effects mask is applied to the tool after the tool is processed.

Basic Node Setup

The Coordinate Space node is used below to make a circular pattern based on a Fast Noise, Mosaic Blur (DaVinci Resolve Resolve FX only), and a Transform node. The Crop node at the end is used to set the desired resolution.



The Coordinate Space node can help create motion graphics backgrounds

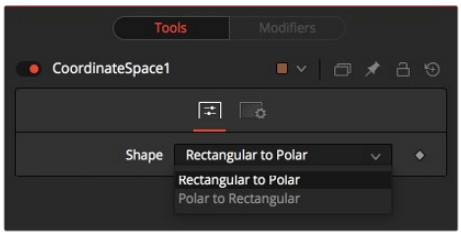
Example

To demonstrate a basic tunnel effect that can be achieved with this node:

1. Add a Text+ node with some text, and then animate it to move along a path from the top of the frame to the bottom.
2. Connect the output of the Text+ node to a Coordinate Space node.
3. Select Polar to Rectangular from the Shape menu.

As the text moves from top to bottom along the original path, it appears to move from an infinite distance in the Coordinate Space node. It may be necessary to flip the text using a Transform node to make it appear the correct way in the Coordinate Space node. Another common use for the Coordinate Space node is to use it in pairs: two of them set to different Shape settings with a Drip or Transform node in between. When used in this way, the effect gets modified while the image remains the same.

Inspector



The Coordinate Space Controls tab

Controls Tab

The Controls tab Shape menu switches between Rectangular to Polar and Polar to Rectangular. Consider the following example to demonstrate the two coordinate spaces.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Warp nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Corner Positioner [CPn]



The Corner Positioner node

Corner Positioner Node Introduction

The Corner Positioner can be used to position the four corners of an image interactively. This would typically be used to replace a sign or other rectangular portion of a scene. Connect all corners to Paths or Trackers for animation purposes.

Inputs

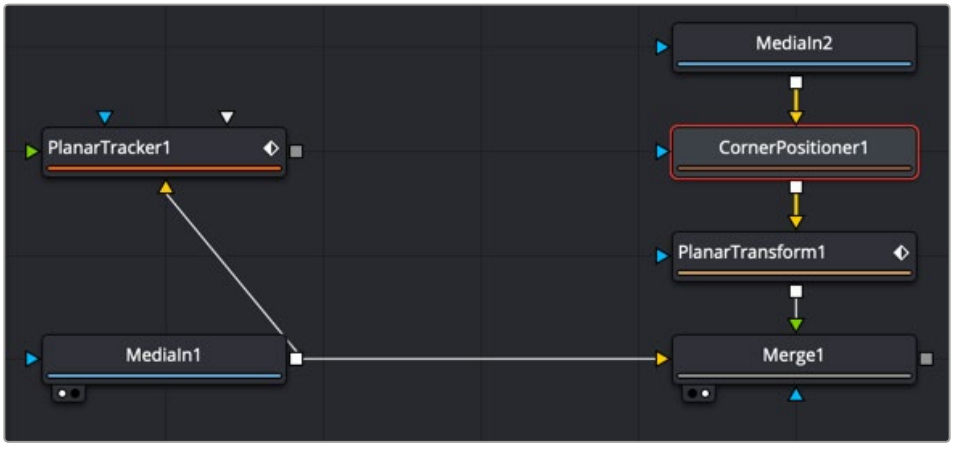
The two inputs on the Corner Positioner node are used to connect a 2D image and an effect mask, which can be used to limit the warped area.

Input: The orange input is used for the primary 2D image that is warped.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the Corner Positioner to only those pixels within the mask. An effects mask is applied to the tool after the tool is processed.

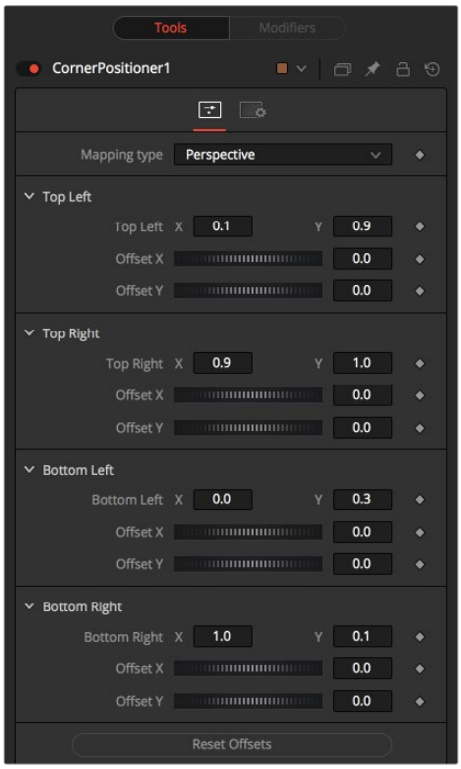
Basic Node Setup

Below, the Corner Positioner is used to position the rectangular corners of the MediaIn2 to fit within a specific area of the MediaIn1 node. A Planar Tracker tracks the background, and then a Planar Transform is used to keep the Corner Positioner in place as the background clip moves. Once the planar tracking is completed, and the Planar Transform is created, the Planar Tracker node is no longer needed, and you can delete it.



The Corner Positioner corner pins a clip in place and uses a Planar Transform to match move it to the background.

Inspector



The Corner Positioner Controls tab

Controls Tab

The Controls tab includes transform and offset adjustments for the four corners of the image

Mapping Type

This determines the method used to project the image caused by the Corner Positioner. In Bi-Linear mode, a straight 2D warping takes place. In Perspective mode, the image is calculated with the offsets in 2D space and then mapped into a 3D perspective.

Corners X and Y

There are four points in the Corner Positioner. Drag these around to position each corner of the image interactively. Attach these control points to any of the usual modifiers.

The image input is deformed and perspective corrected to match the position of the four corners.

Offset X and Y

These controls can be used to offset the position of the corners slightly. This is useful when the corners are attached to Trackers with patterns that may not be positioned exactly where they are needed.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Warp nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Dent [Dnt]



The Dent node

Dent Node Introduction

The Dent node creates a circular deformation of an image similar to a Fish Eye Lens effect, with the choice of six different Dent filters.

Inputs

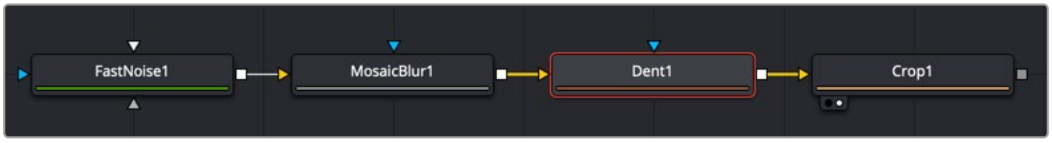
The two inputs on the Dent node are used to connect a 2D image and an effect mask, which can be used to limit the warped area.

Input: The orange input is used for the primary 2D image that is warped.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the Dent to only those pixels within the mask. An effects mask is applied to the tool after the tool is processed.

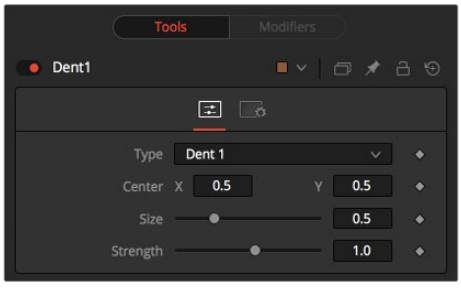
Basic Node Setup

The Dent node is used below to make a circular pattern based on a Fast Noise and a Mosaic Blur (DaVinci Resolve Resolve FX only). The Crop node at the end is used to set the desired resolution.



The Dent node can help create lens distortion effects or a motion graphics background.

Inspector



The Dent Controls tab

Controls Tab

The adjustments in the Controls tab are used to change the Dent style, position, size, and strength.

Type

Select the type of Dent filter to use from this menu. All parameters for the Dent can be keyframed.

Dent 1

This creates a bulge dent.

Kaleidoscope

This creates a dent, mirrors it, and inverts it.

Dent 2

This creates a displacement dent.

Dent 3

This creates a deform dent.

Cosine Dent

This creates a fracture to a center point.

Sine Dent

This creates a smooth rounded dent.

Center X and Y

This positions the Center of the Dent effect on the image. The default values are 0.5, 0.5, which center the effect in the image.

Size

This changes the size of the area affected by the dent. Animate this slider to make the dent grow.

Strength

This changes the overall strength of the dent.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Warp nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Displace [Dsp]



The Displace node

Displace Node Introduction

The Displace node uses a map image to displace or refract another image. This is useful for creating a vast range of effects from bevels and heat distortion to glass and water effects.

Inputs

There are three inputs on the Displace node: The primary image, the displacement map foreground image, and an effect mask.

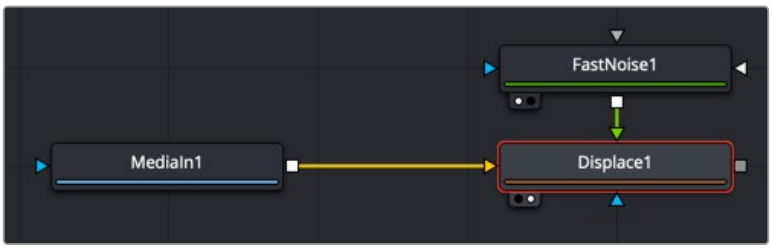
Input: The orange image input is a required connection for the primary image you wish to displace.

Foreground Image: The green input is also required as the image used to displace the background. Once connected, you can choose red, green, blue, alpha, or luminance channel to create the displacement.

Effect Mask: The optional blue effect mask input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the displacement to only those pixels within the mask. An effects mask is applied to the tool after it is processed.

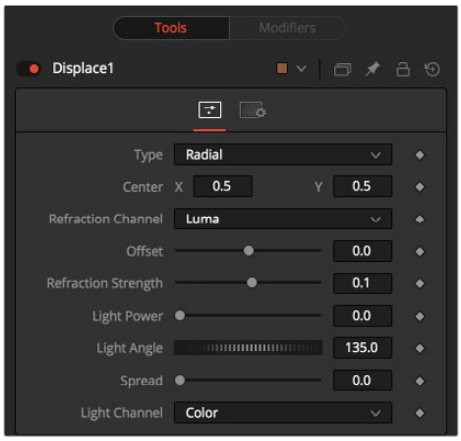
Basic Node Setup

Below, the Displace node uses a Fast Noise to generate the Displace map. Increasing the seethe rate can produce heat distortion or flag waving effects.



The Displace node using a Fast Noise node for the Displace map

Inspector



The Displace Controls tab

Controls Tab

The Controls tab is used to change the style, position, size , strength, and lighting (embossing) of the displacement.

Type

The Type menu is used to choose in what mode the Displace node operates. The Radial mode uses the map image that refracts each pixel out from the center, while X/Y mode provides control over the amount of displacement along each axis individually.

NOTE: There is one set of Refraction controls while in Radial mode, and two sets in XY mode-one for each of the X and Y channels.

Center (Radial Only)

The Center control defines the point from which pixels are displaced toward or away.

Refraction Channel

This drop-down menu controls which channel from the foreground image is used to displace the image. Select from Red, Green, Blue, Alpha, or Luminance channels. In XY mode, this control appears twice, once for the X displacement and once for the Y displacement.

Refraction Strength (Radial)

Controls the strength of the refraction. Higher values cause stronger or more pronounced refraction.

X and Y Refraction (X/Y)

Two separate sliders appear to control the Refraction strength along the X- and Y-axis separately. Otherwise, this is exactly like Refraction Strength.

Light Power

This controls the intensity, or strength, of the simulated light, causing bright and dim areas to form according to the contour of the refraction image. Higher values cause the bright and dim areas to be more pronounced.

Light Angle

This sets the angle of the simulated light source.

Spread

This widens the Displacement effect and takes the edge off the Refraction map. Higher values cause the ridges or edges to spread out.

Light Channel

Select the channel from the refraction image to use as the simulated light source. Select from Color, Red, Green, Blue, Alpha, or Luminance channels.

NOTE: The Radial mode pushes pixels inward or outward from a center point, based on pixel values from the Displacement map. The XY mode uses two different channels from the map to displace pixels horizontally and vertically, allowing more precise results. Using the XY mode, the Displace node can even accomplish simple morphing effects. The Light controls allow directional highlighting of refracted pixels for simulating a beveled look.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Warp nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Drip [DRP]



The Drip node

Drip Node Introduction

The Drip node creates a ripple effect over the entire image, which has the potential to animate outward from a central source. There are a variety of different Drip effects from which to choose.

Inputs

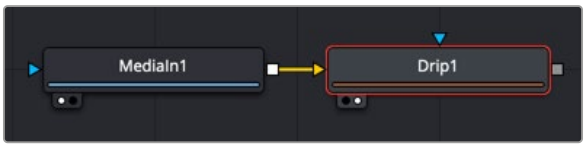
The two inputs on the Drip node are used to connect a 2D image and an effect mask, which can be used to limit the warped area.

Input: The orange input is used for the primary 2D image that is warped.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the warping to only those pixels within the mask. An effects mask is applied to the tool after the tool is processed.

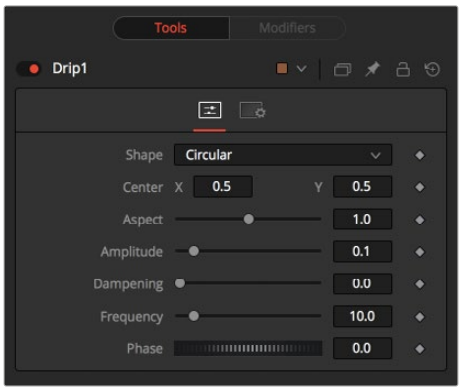
Basic Node Setup

Below, the Drip node is used to make rippling water-style effects using a MediaIn node.



The Drip node can be connected directly after a MediaIn node or any node providing a 2D output.

Inspector



The Drip Controls tab

Controls Tab

The Controls tab is used to change the style, position, size , strength, and phase for animating the “ripples” of the Drip.

Shape

Use this control to select the shape of the Drip.

Circular

This creates circular ripples. This is the default Drip mode.

Square

This creates even-sided quadrilateral drips.

Random

This creates a randomly dispersed noise that distorts your image and is similar to a particle effect.

Horizontal

This creates horizontal waves that move in one direction.

Vertical

This creates vertical waves that move in one direction.

Exponential

This creates a Drip effect that looks like a diamond shape with inverted, curved sides (an exponential curve flipped and mirrored).

Star

This creates an eight-way symmetrical star-shaped ripple that acts as a kaleidoscope when the phase is animated.

Radial

This creates a star-shaped ripple that emits from a fixed pattern.

Center X and Y

Use this control to position the center of the Drip effect in the image. The default is 0.5, 0.5, which centers the effect in the image.

Aspect

Control the aspect ratio of the various Drip shapes. A value of 1.0 causes the shapes to be symmetrical. Smaller values cause the shape to be taller and narrower, while larger values cause shorter and wider shapes.

Amplitude

The Amplitude of the Drip effect refers to the peak height of each ripple. Use the slider to change the amount of distortion the Drip applies to the image. A value of 0.0 gives all ripples no height and therefore makes the effect transparent. A maximum Amplitude of 10 makes each ripple extremely visible and completely distorts the image. Higher numbers can be entered via the text entry boxes.

Dampening

Controls the Dampening, or falloff, of the amplitude as it moves away from the center of the effect. It can be used to limit the size or area affected by the Drip.

Frequency

This changes the number of ripples emanating from the center of the Drip effect. A value of 0.0 indicates no ripples. Move the slider up to a value of 100 to correspond with the density of desired ripples.

Phase

This controls the offset of the frequencies from the center. Animate the Phase value to make the ripple emanate from the center of the effect.

Common Controls

Settings Tab

The Settings Tab in the Inspector is also duplicated in other Warp nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Grid Warp [Grd]



The Grid Warp node

Grid Warp Node Introduction

The Grid Warp node is a 2D deformation grid with flexible vertices. The image is deformed so that the source grid matches the destination grid.

Inputs

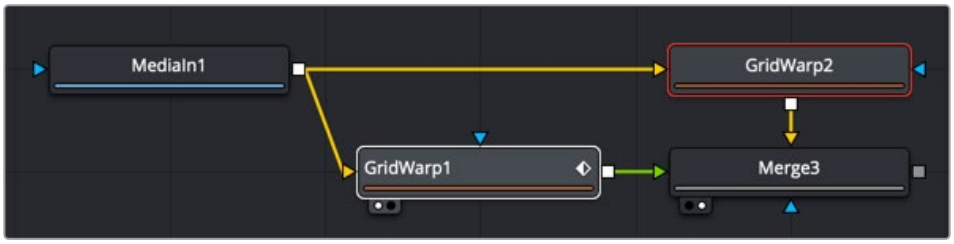
The two inputs on the Grid Warp node are used to connect a 2D image and an effect mask, which can be used to limit the warped area.

Input: The orange input is used for the primary 2D image that is warped.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the warping to only those pixels within the mask. An effects mask is applied to the tool after the tool is processed.

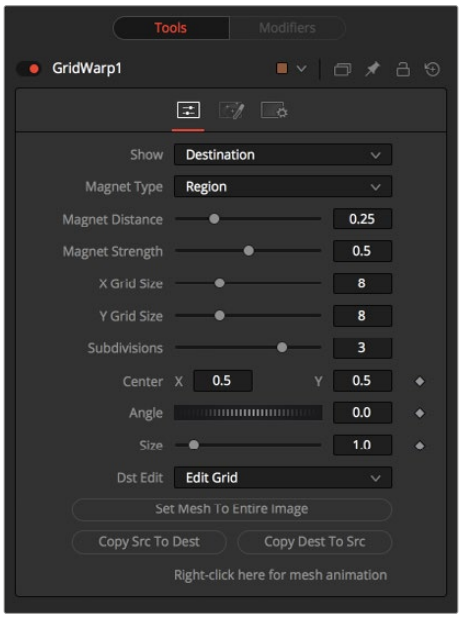
Basic Node Setup

Below, two Grid Warp nodes are used to warp different areas of a frame. This can be useful for chaining the framing in a shot or adding slight movement to a still. Using the Copy Src to Dest button, only the area modified by Grid Warp 2 is pasted as the foreground in the Merge.



The Grid Warp can be used to shift areas within a shot for reframing or adding animation to a still.

Inspector



The Grid Warp Controls tab

Controls Tab

The Controls tab contains parameters that configure the onscreen grid as well the type of distortion applied when a control point on the grid is moved.

Source and Destination

The Source and Destination buttons determine whether the source grid or destination grid is currently active. Only one grid can be displayed or manipulated at a time. The selected button is highlighted to indicate that it is the currently active grid.

All other controls in this tab affect the grid selected by this control.

Selection Type

These three buttons determine the selection types used for manipulating the points. There are three options available.

Selected

When in Selected mode, adjustments to the grid are applied only to the currently selected points. This mode is identical to normal polyline operation.

Region

In Region mode, all points within the area around the mouse pointer move when the mouse button is clicked. New points that enter the region during the move are ignored. Choosing this option exposes Magnet Distance and Magnet Strength controls to determine the size and falloff of the area.

Magnetic

In Magnetic mode, all points within the area around the mouse pointer move when the mouse button is clicked. New points that enter the region during the move are affected as well. Choosing this option exposes Magnet Distance and Magnet Strength controls to determine the size and falloff of the area.

Magnet Distance

The default node for selecting and manipulating the grid is a Magnet node. The magnet is represented in the viewer by a circle around the mouse pointer. The Magnet Distance slider controls how large the region of effect for the magnet is, as in the size of the circle. Drag on the grid and any vertex within the range of the slider moves.

To increase the size of the magnet, increase the value of this slider. Alternately, adjust the size of the magnet by holding down the D key while dragging the mouse.

Magnet Strength

The Magnet Strength slider increases or decreases the falloff of the magnet cursor's effect. At a setting of 0.0, the magnetic cursor has no effect, and vertices do not move at all. As the values increase, the magnet causes a greater range of motion in the selected vertices. Use smaller values for a more sensitive adjustment and larger values for broad-sweeping changes to the grid.

X and Y Grid Size

The X and Y Grid Size sliders control the number of divisions in the grid. Where the X and Y divisions intersect, a control vertex is created.

Be aware that changing either of these controls after applying changes in the grid resets the entire grid. Set the X and Y grid sizes to the appropriate resolution before making detailed adjustments to the grid.

Subdivision Level

The Subdivision Level determines how many subdivisions there are between each set of divisions. Subdivisions do not generate vertices at intersections. The more subdivisions, the smoother the deformation is likely to be, but the slower it is to render.

Center

The Center coordinates determine the exact center of the grid. The onscreen Center control is invisible while editing the grid. Select the Edit Rect mode, and the grid center becomes visible and available for editing.

Use the Center control to move the grid through a scene without affecting the animation applied to the individual vertices. For example, while deforming lips, track the motion of the face with a Tracker, and connect the grid center to the Tracker. This matches the grid with slight movements of the head while focusing on the deformation of the lips.

Angle

This Angle control rotates the entire grid.

Size

The Size control increases or decreases the scale of the grid.

Edit Buttons

There are four edit modes available, each of which can be selected by clicking on the appropriate button.

Edit None

Set the grid to Edit None mode to disable the display of all onscreen controls.

Edit Grid

The Edit Grid mode is the default mode. While this mode is enabled, the grid is drawn in the viewer, and the control vertices of the grid can be manipulated directly.

Edit Rectangle

When the grid is in Edit Rectangle mode, the onscreen controls display a rectangle that determines the dimensions of the grid. The sides of the rectangle can be adjusted to increase or decrease the grid's dimension. This mode also reveals the onscreen Center control for the grid.

Edit Line

The Edit Line mode is beneficial for creating grids around organic shapes. When this mode is enabled, all onscreen controls disappear, and a spline can be drawn around the shape or object to be deformed. While drawing the spline, a grid is automatically created that best represents that object. Additional controls for Tolerance, Over Size, and Snap Distance appear when this mode is enabled. These controls are documented below.

Set Mesh to Entire Image

The Set Mesh to Entire Image button automatically resets the size of the grid to the exact dimensions of the image. Any adjustments to vertices within the grid are reset.

Copy Buttons

These two buttons provide a technique for copying the exact shape and dimensions of the source grid to the destination, or the destination grid to the source. This is particularly useful after setting the source grid to ensure that the destination grid's initial state matches the source grid before beginning a deformation.

Point Tolerance

This control is visible only when the Edit Line mode is enabled. The Point Tolerance slider determines how much tessellation the grid applies to match the density of points in the spline closely. The lower this value, the fewer vertices there are in the resulting grid, and the more uniform the grid appears. Higher values start applying denser grids with variations to account for regions in the spline that require more detail.

Oversize Amount

This control is visible only when the Edit Line mode is enabled. The Oversize Amount slider is used to set how large an area around the spline should be included in the grid. Higher values create a larger border, which can be useful when blending a deformation back into the source image.

Snap Distance

This control is visible only when the Edit Line mode is enabled. The Snap Distance slider dictates how strongly the drawn spline attracts surrounding vertices. If a vertex is close enough to a spline's edge, the vertex moves to line up with the spline. The higher the value, the farther the reach of the spline.

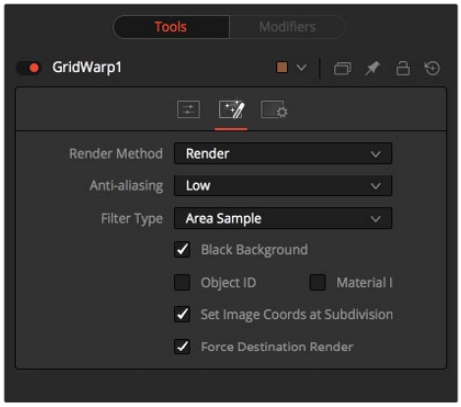
Right-Click Here for Mesh Animation

The grids are static by default. Right-clicking on the Right-Click Here for Mesh Animation label provides a contextual menu with options for animating the grid or connecting it to another grid in the composition.

The grid uses a Polychange spline. Any adjustment to the control points adds or modifies the keyframe for all points on that spline.

Right-Click Here for Shape Animation

This label appears only in the Edit Line mode. Right-clicking on the Right-Click Here for Shape Animation label reveals a pop-up menu used to animate the shaping polyline or to connect it to other polylines.



The Grid Warp Render tab

Render Tab

The Render tab controls the final rendered quality and appearance of the warping.

Render Method

The Render Method drop-down menu is used to select the rendering technique and quality applied to the mesh. The three settings are arranged in order of quality, with the first, Wireframe, as the fastest and lowest of quality. The default mode is Render, which produces final resolution, full-quality results.

Anti-Aliasing

The Anti-Aliasing control appears only as a checkbox when in Wireframe Render mode.

In other modes, it is a drop-down menu with three levels of quality. Higher degrees of anti-aliasing improve image quality dramatically but vastly increase render times. The Low setting may be an appropriate option while setting up a large dense grid or previewing a node tree, but rarely for a final render.

Filter Type

When the Render Method is set to something other than Wireframe mode, the Filter Type menu is visible and set to Area Sample. This setting prevents the grid from calculating area samples for each vertex in the grid, providing good render quality. Super Sample can provide even better results but requires much greater render times.

Wireframe Width

This slider appears only when the Render Method is set to Wireframe. It determines the width of the lines that make up the wireframe.

Anti-Aliased

This checkbox appears only when the Render Method is set to Wireframe. Use this checkbox to enable/disable anti-aliasing for the lines that make up the wireframe.

Black Background

The Black Background checkbox determines whether pixels outside of the grid in the source image are set to black or if they are preserved.

Object ID and Material ID

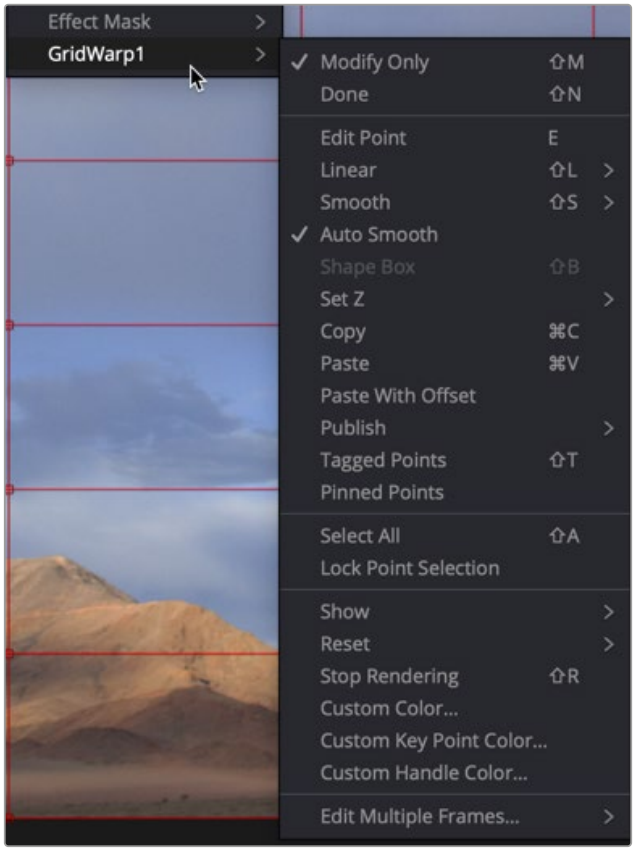
Enable the Object ID or Material ID checkboxes to have the grid output the proper ID channel in the final render.

Set Image Coordinates at Subdivision Level

This checkbox defaults to enabled and sets the image coordinates at the subdivision level.

Force Destination Render

This checkbox defaults to enabled and forces a destination render.



Grid Warp Contextual Menu options

Contextual Menu Options

The Grid Warp node places a submenu for both source and destination grids in the viewer's contextual menu. Both menus have the exact same name, where only the menu for the active grid is populated with options. The other menu is empty. The contextual menu options are also available from the toolbar that appears in the viewer.

Modify Only/Done

These two options set the mesh to Modify Only and Done modes, respectively. Select Modify Only to edit the mesh or Modify Done to prevent any further changes to a mesh.

Smooth/Linear

Use Smooth and Linear options to apply or remove smoothing from selected vertices.

Auto Smooth Points

When Auto Smooth Points is enabled, the vertices in the grid are automatically smoothed whenever they are moved. This is generally enabled by default.

Z Under/Z Same/Z Over

When two vertices in a grid overlap, one ends up getting clipped by the other. Z Under, Z Same, and Z Over are used to select which vertices are rendered on top and which are rendered behind.

Select All

This option selects all points in the mesh.

Show Key Points, Handles, Grid, and Subdivisions

Use these four options to enable or disable the display of the grid, key points (vertices), Bézier handles, and subdivisions in the viewers.

Reset Selected Points

This resets Selected Points (vertices) to their default positions.

Reset All Points

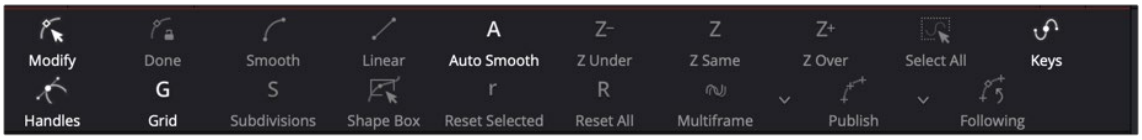
This resets all points (vertices) in the mesh to their default positions.

Stop Rendering

This option stops rendering, which disables all rendering of the Grid Warp node until the mode is turned off. This is frequently useful when making a series of fine adjustments to a complex grid.

Grid Warp Toolbar

Whenever the Grid Warp node is selected and is in Edit Grid mode, the Grid Warp toolbar is displayed in the views. This toolbar provides a variety of options for manipulating and adjusting the grid. The toolbar buttons in this toolbar are described in the preceding “Contextual Menu Options” section.



The Grid Warp viewer toolbar

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Warp nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Lens Distort [Lens]



The Lens Distort node

Lens Distort Node Introduction

This node can be used to remove or add lens distortion in an image. The lens distortion in an image depends on the lens, quality, number of elements, and many other factors.

One reason to remove lens distortion is to composite with an undistorted layer. For example, compositing a 3D element over a distorted live-action layer will cause unwanted effects like straight lines not matching up on the foreground and background. The resulting composite will not look believable.

Inputs

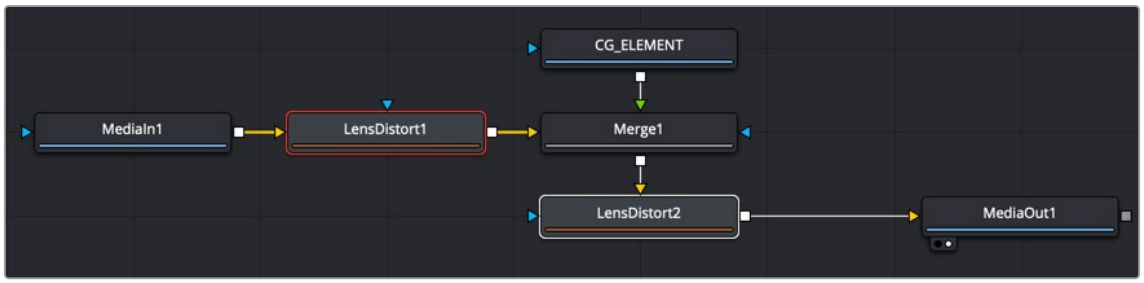
The two inputs on the Lens Distort node are used to connect a 2D image and an effect mask, which can be used to limit the distorted area.

Input: The orange input is used for the primary 2D image that is distorted.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the distortion to only those pixels within the mask. An effects mask is applied to the tool after the tool is processed.

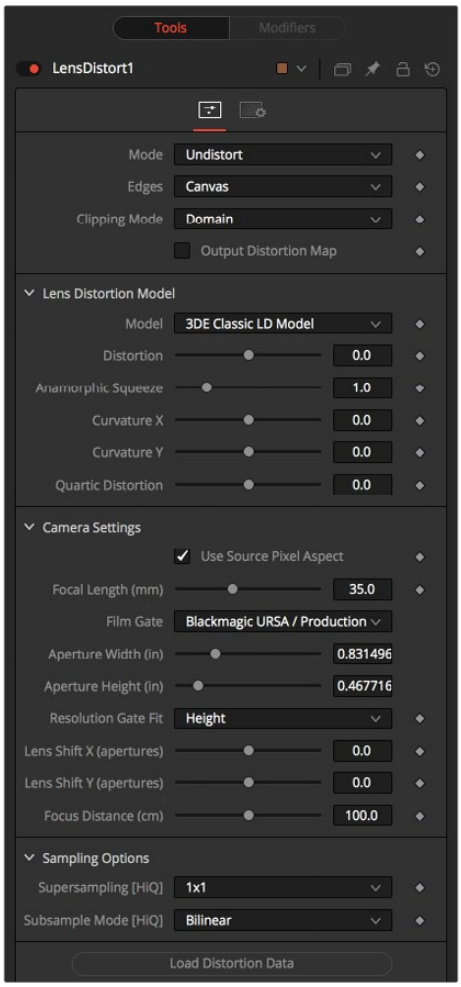
Basic Node Setup

A simplified example below applies the Lens Distort in Undistort mode to the MediaIn1 live-action layer, composites the 3D elements, and finally applies the Lens Distort at the end with the same settings, but this time in Distort mode to get the original look and distortion back into the image.



Lens Distort applied on the live-action media at the beginning of the node tree, and once again at the end

Inspector



The Lens Distort Controls tab

Controls Tab

The Controls tab presents various ways to customize or build the lens distortion model you want. Camera Settings allow you to specify the camera used to capture the content.

Mode

Undistort removes the lens distortion to create a flattened image. Distort brings the original lens distortion back into the image.

Edges

Determines how samples that fall outside the frame are treated.

- **Canvas:** Pixels outside the frame are set to the default canvas color. In most cases, this is black with no alpha.
- **Duplicate:** Pixels outside the frame are duplicated. This results in “smeared” edges but is useful when, for example, applying a blur because in that case black pixels would result in the unwanted blurring between the actual image and the black canvas.

Clipping Mode

- **Domain:** Retains all pixels that might be moved out of the frame for later re-distorting.
- **Frame:** Pixels moved outside the frame are discarded.

Output Distortion Map

Outputs the location of pixels as a warped screen-coordinate map.

Camera Settings

The options known from the Camera 3D are duplicated here. They can either be set manually or connected to an already existing Camera 3D.

Lens Distortion Model

Select the appropriate 3D Equalizer Lens Distortion model here: 3DE Classic Model, 3DE4 Anamorphic, 3DE4 Radial Fisheye, or 3DE4 Radial. Please consult the 3D Equalizer manual for further explanation. The sliders in the 3DE Classic LD Model are most likely best suited for manually applying (un)distortion, without having imported lens data.

Supersampling [HiQ]

Sets the number of samples used to determine each destination pixel. As always, higher supersampling leads to higher render times. 1×1 bilinear is usually of sufficient quality, but with high lens distortion near the edges of the lens, there are noticeable differences to higher settings.

Supersampling Mode [HiQ]

The type of sample done for each supersample. Nearest leads to a crisper but more aliased image. Bi-Linear gives a blurrier result.

Load Distortion Data

Allows the user to load a Lens Distortion profile created, for example, by the 3D Equalizer.

How to Manually Determine Lens Distortion

In an ideal world, one would have exact lens parameters from each lens that was used during the shoot, and one could use those values to undistort the image. However, in the real world, those parameters have not been taken on set or don't match. Another approach is to use software like 3D Equalizer, which analyzes the footage and delivers a dataset that can be imported into the Lens Distort node right away.

And finally, one could try to manually eyeball the amount of lens distortion using the control sliders. To do that, one could either look for horizontal or vertical lines in the footage that are supposed to be straight and straighten them out using the controls, or shoot a full-frame checkerboard pattern on set as a reference.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Warp nodes. These common controls are described in detail at the end of this chapter in "The Common Controls" section.

Perspective Positioner [PPn]



The Perspective Positioner node

Perspective Positioner Node Introduction

The Perspective Positioner is the complementary node to the Corner Positioner node. It “unpins” an image by positioning corner points on a perspective distorted area, thereby removing the perspective from the image. This function can also be used to wobble and warp the image by animating the points over time.

Inputs

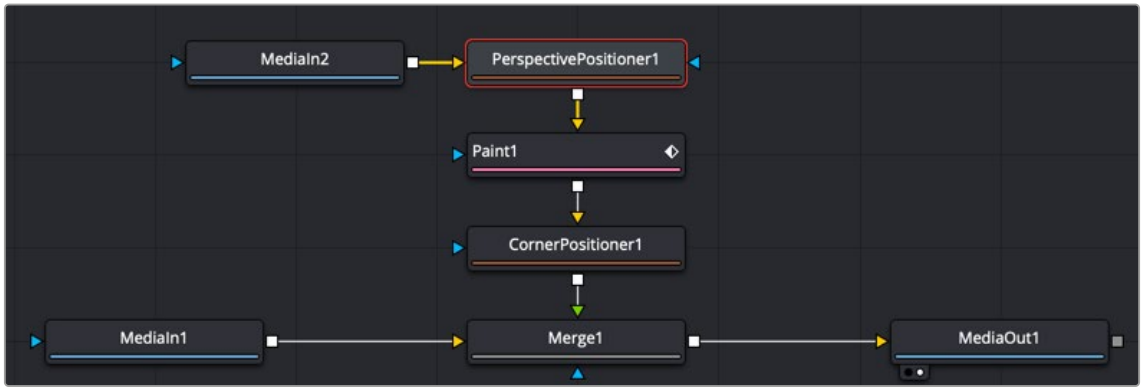
The two inputs on the Perspective Positioner node are used to connect a 2D image and an effect mask, which can be used to limit the transformed area.

Input: The orange input is used for the primary 2D image that is transformed.

Effect Mask: The blue input is for a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the transform to only those pixels within the mask. An effects mask is applied to the tool after the tool is processed.

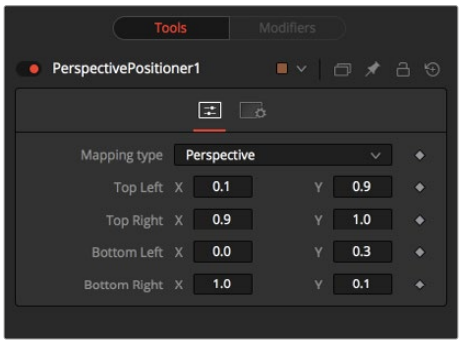
Basic Node Setup

In the example below, the Perspective Positioner is used to unpin a perspective distorted area of the MediaIn2 in order to paint on the flat texture. The MediaIn2 is then corner pinned back into place. The Perspective Positioner and Corner Positioner do not concatenate, so some softness is introduced with these nodes.



The Perspective Positioner unpins an image to paint on a texture map.

Inspector



The Perspective Positioner Controls tab

Controls Tab

The Controls tab contains parameters for selecting vector channels and controlling how much distortion they apply to an image.

Mapping Type

The Mapping Type menu is used to select the type of transform used to distort the image. Bi-Linear is available for support of older projects. Leaving this on Perspective is strongly suggested since the Perspective setting maps the real world more accurately.

Corners X and Y

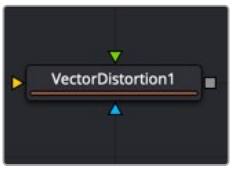
There are the four control points of the Perspective Positioner. Interactively drag these in the viewers to position each corner of the image. You can refine their position using the Top, Bottom, Left, and Right controls in the Inspector.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Warp nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Vector Distortion [Dst]



The Vector Distortion node

Vector Distortion Node Introduction

The Vector Distortion node distorts the main source image along the X- and Y-axis separately, based on the vector channel data in the source image or vector channels from a second reference image.

Inputs

There are three inputs on the Vector Distort node for the primary 2D image, the distort image with vector channels and an effect mask.

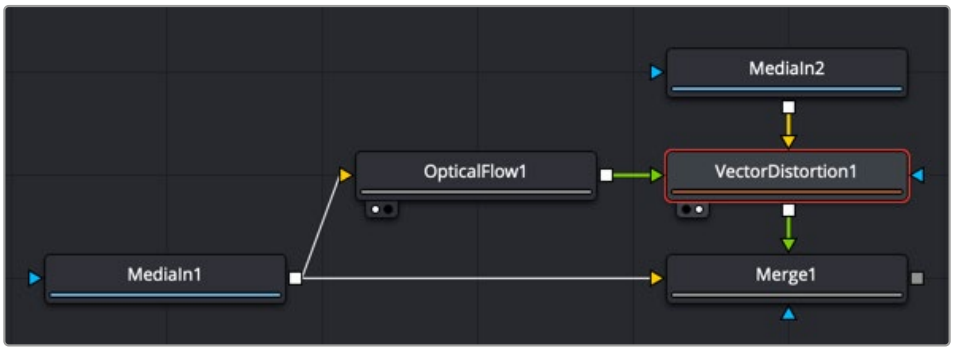
Input: The orange image input is a required connection for the primary image you wish to distort. If this image has vector channels, they are used in the distortion.

Distort: The green input is an optional distort image input used to distort the background image based on vector channels. Once connected, it overrides vector channels in the input image.

Effect Mask: The optional blue effect mask input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the displacement to only those pixels within the mask. An effects mask is applied to the tool after it is processed.

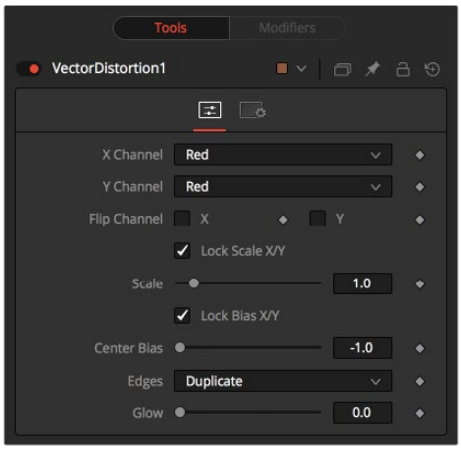
Basic Node Setup

In the example below, the MediaIn1 node uses an Optical Flow node to generate the vector channels, which are then passed on to the Vector Distort node. MediaIn2 is distorted and composited over the background.



The Vector Distort node relies on the image having vector channels or an Optical Flow node to generate them.

Inspector



The Vector Distortion Controls tab

Controls Tab

The Controls tab contains parameters for selecting vector channels and controlling how much distortion they apply to an image.

X Channel and Y Channel

These two menus are used to select which channel of the (green) distort image input is used to distort the X and Y channels. If no distort reference image is connected, then channels from the main orange input are used instead.

Flip Channel X and Flip Channel Y

Use these checkboxes to flip the direction of the distortion along the specified axis.

Lock Scale X/Y

Select this checkbox to separate the Scale slider into separate Scale X and Scale Y sliders.

Scale

Use the Scale slider to apply a multiplier to the values of the distortion reference image.

Lock Bias X/Y

Select this checkbox to separate the Bias slider into separate Bias X and Bias Y sliders.

Center Bias

Use the Center Bias slider to shift or nudge the distortion along a given axis.

Edges

This menu determines how the edges of the image are treated.

- **Canvas:** This causes the edges that are revealed by the shake to be the canvas color—usually transparent or black.
- **Duplicate:** This causes the edges to be duplicated, causing a slight smearing effect at the edges.

Glow

Use this slider to add a glow to the result of the vector distortion.

Common Controls

Settings Tab

The Settings tab in the Inspector is also duplicated in other Warp nodes. These common controls are described in detail at the end of this chapter in “The Common Controls” section.

Vortex [Vtx]



The Vortex node

Vortex Node Introduction

The Vortex effect appears as a swirling whirlpool in specified regions of the image. The Vortex can be made to move and grow by animating the various controls.

Inputs

There are two inputs on the Vortex node for the primary 2D image and the effect mask.

Input: The orange image input is a required connection for the primary image you wish to swirl.

Effect Mask: The optional blue effect mask input expects a mask shape created by polylines, basic primitive shapes, paint strokes, or bitmaps from other tools. Connecting a mask to this input limits the swirling vortex to only those pixels within the mask. An effects mask is applied to the tool after it is processed.

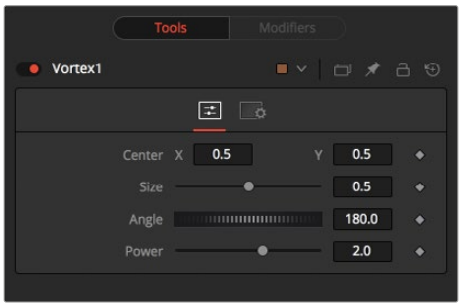
Basic Node Setup

Below, the Vortex is applied to text for creating motion graphics. Since the Vortex will cause the text to swirl outside the text boundary, a Set Domain node is used to expand the text boundary, ensuring that the text is not cropped when the Vortex is applied.



The Vortex is used to create swirling whirlpool effects.

Inspector



The Vector Distortion Controls tab

Controls Tab

The Controls tab contains parameters for adjusting the position, size, and strength of the Vortex effect.

Center X and Y

This control is used to position the center of the Vortex effect on the image. The default is 0.5, 0.5, which positions the effect in the center of the image.

Size

Size changes the area affected by the Vortex. You can drag the circumference of the effect in the viewer or use the Size slider.

Angle

Drag the rotation handle in the viewer or use the thumbwheel control to change the amount of rotation in the Vortex. The higher the angle value, the greater the swirling effect.

Power

Increasing the Power slider makes the Vortex smaller but tighter. It effectively concentrates it inside the given image area.

Common Controls

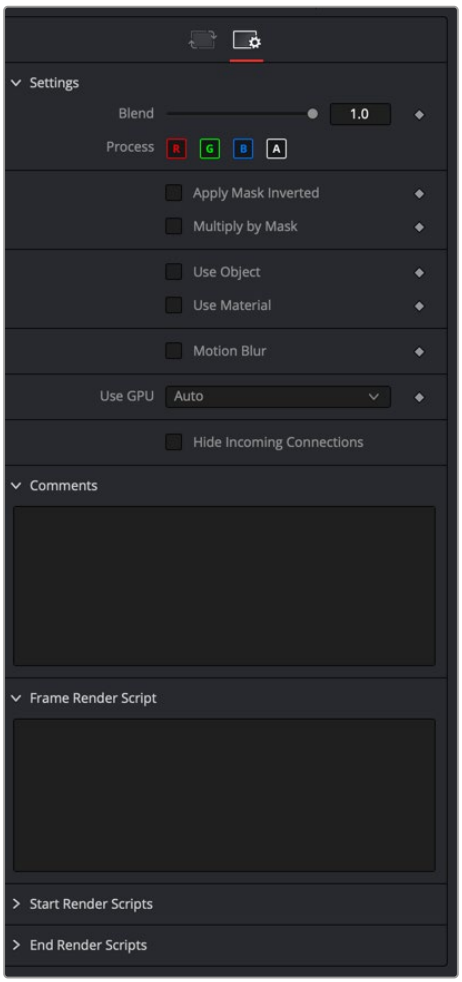
Settings Tab

The Settings tab in the Inspector is also duplicated in other Warp nodes. These common controls are described in detail in the following “The Common Controls” section.

The Common Controls

Nodes that handle Warp operations share a number of identical controls in the Inspector. This section describes controls that are common among Warp nodes.

Inspector



Warp Common Controls tab

Settings Tab

The Settings tab in the Inspector can be found on every tool in the Warp category. The Settings controls are even found on third-party Warp-type plugin tools. The controls are consistent and work the same way for each tool.

Blend

The Blend control is used to blend between the tool's original image input and the tool's final modified output image. When the blend value is 0.0, the outgoing image is identical to the incoming image. Normally, this will cause the tool to skip processing entirely, copying the input straight to the output.

Process When Blend Is 0.0

The tool is processed even when the input value is zero. This can be useful if the node is scripted to trigger a task, but the node's value is set to 0.0.

Red/Green/Blue/Alpha Channel Selector

These four buttons are used to limit the effect of the tool to specified color channels. This filter is often applied after the tool has been processed.

For example, if the Red button on a Blur tool is deselected, the blur will first be applied to the image, and then the red channel from the original input will be copied back over the red channel of the result.

There are some exceptions, such as tools for which deselecting these channels causes the tool to skip processing that channel entirely. Tools that do this will generally possess a set of identical RGBA buttons on the Controls tab in the tool. In this case, the buttons in the Settings and the Controls tabs are identical.

Apply Mask Inverted

Enabling the Apply Mask Inverted option inverts the complete mask channel for the tool. The mask channel is the combined result of all masks connected to or generated in a node.

Multiply by Mask

Selecting this option will cause the RGB values of the masked image to be multiplied by the mask channel's values. This will cause all pixels of the image not included in the mask (i.e., set to 0) to become black/transparent.

Use Object/Use Material (Checkboxes)

Some 3D software can render to file formats that support additional channels. Notably, the EXR file format supports Object ID and Material ID channels, which can be used as a mask for the effect. These checkboxes determine whether the channels will be used, if present. The specific Material ID or Object ID affected is chosen using the Object and Material ID sliders explained below.

Correct Edges

This checkbox appears only when the Use Object or Use Material checkboxes are selected. It toggles the method used to deal with overlapping edges of objects in a multi-object image. When enabled, the Coverage and Background Color channels are used to separate and improve the effect around the edge of the object. If this option is disabled (or no Coverage or Background Color channels are available), aliasing may occur on the edge of the mask.

> For more information on the Coverage and Background Color channels, *see Chapter 18, "Understanding Image Channels,"* in the Fusion Reference Manual.

Object ID/Material ID (Sliders)

Use these sliders to select which ID will be used to create a mask from the object or material channels of an image. Use the Sample button in the same way as the Color Picker: to grab IDs from the image displayed in the viewer. The image or sequence must have been rendered from a 3D software package with those channels included.

Motion Blur

- **Motion Blur:** This toggles the rendering of Motion Blur on the tool. When this control is toggled on, the tool's predicted motion is used to produce the motion blur caused by the virtual camera's shutter. When the control is toggled off, no motion blur is created.
- **Quality:** Quality determines the number of samples used to create the blur. A quality setting of 2 will cause Fusion to create two samples to either side of an object's actual motion. Larger values produce smoother results but increase the render time.
- **Shutter Angle:**L Shutter Angle controls the angle of the virtual shutter used to produce the motion blur effect. Larger angles create more blur but increase the render times. A value of 360 is the equivalent of having the shutter open for one full frame exposure. Higher values are possible and can be used to create interesting effects.
- **Center Bias:** Center Bias modifies the position of the center of the motion blur. This allows for the creation of motion trail effects.
- **Sample Spread:** Adjusting this control modifies the weighting given to each sample. This affects the brightness of the samples.

Use GPU

The Use GPU menu has three settings. Setting the menu to Disable turns off GPU hardware-accelerated rendering. Enabled uses the GPU hardware for rendering the node. Auto uses a capable GPU if one is available and falls back to software rendering when a capable GPU is not available.

Hide Incoming Connections

Enabling this checkbox can hide connection lines from incoming nodes, making a node tree appear cleaner and easier to read. When enabled, empty fields for each input on a node will be displayed in the Inspector. Dragging a connected node from the node tree into the field will hide that incoming connection line as long as the node is not selected in the node tree. When the node is selected in the node tree, the line will reappear.

Comments

The Comments field is used to add notes to a tool. Click in the empty field and type the text. When a note is added to a tool, a small red square appears in the lower-left corner of the node when the full tile is displayed, or a small text bubble icon appears on the right when nodes are collapsed. To see the note in the Node Editor, hold the mouse pointer over the node to display the tooltip.

Scripts

Three Scripting fields are available on every tool in Fusion from the Settings tab. They each contain edit boxes used to add scripts that process when the tool is rendering. For more details on scripting nodes, please consult the Fusion scripting documentation.

Chapter 64

Modifiers

This chapter details the modifiers available in Fusion.

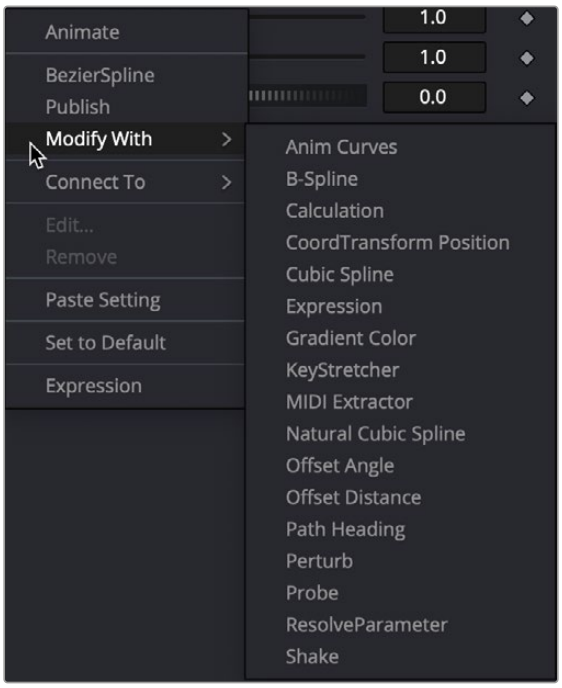
Contents

Modifiers 1745
Anim Curves 1746
Bézier Spline 1748
B-Spline 1750
Calculation 1750
CoordTransform Position 1753
Cubic Spline 1754
Custom Poly 1755
Expression 1757
From Image 1762
Gradient Color 1764
Key Stretcher Modifier 1765
MIDI Extractor 1766
Natural Cubic Spline 1769
Offset (Angle, Distance, Position) 1770
Path 1773
Perturb 1774
Probe 1776
Publish 1778
Resolve Parameter 1779
Shake 1779
Track 1781
Vector Result 1782
XY Path 1784

Modifiers

Modifiers are extensions to a node's standard set of parameters found in the Inspector; in fact, modifiers are designed to control other parameters. They can be as simple as a motion path or linking two parameters. However, they can also be elaborate expressions, procedural functions, external data, third-party plugins, or scripted Fuses. You can add modifiers by right-clicking over a parameter in the Inspector and choosing a modifier from the menu. Alternatively, you can right-click a control in the viewer. Not all modifiers are displayed in the right-click menu for all parameters. Some modifiers work only on specific parameter types.

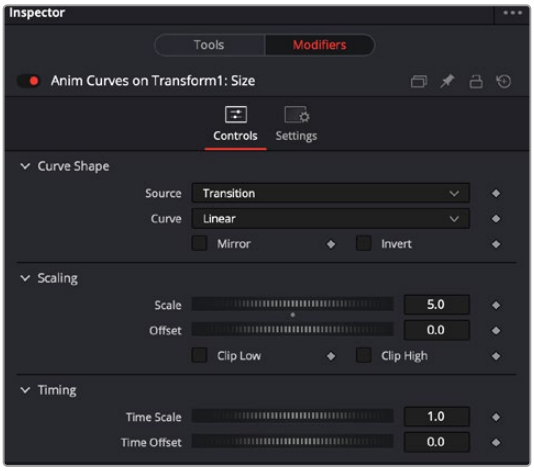
NOTE: Text3D and Text+ have additional text-specific modifiers, which are covered in their nodes' sections.



Modifiers right-click menu

Anim Curves

The Animation Curves modifier (Anim Curves) is used to dynamically adjust the timing, values, and acceleration of an animation, even if you decide to change the duration of a comp. Using this modifier makes it infinitely easier to stretch or squish animations, create smooth motion, add bouncing properties, or mirror animations curves without the complexity of manually adjusting splines.



The Inspector for the Anim Curves modifier

When creating Fusion templates for the Edit and Cut page in DaVinci Resolve, the Anim Curves modifier allows the keyframed animation you've created in Fusion to stretch and squish appropriately as the transition, title, or effect's duration changes on the Edit and Cut page Timelines.

Curve Shape Controls

The controls for the Anim Curves modifier appear in the modifier's tab of the Inspector. The Curve Shape controls determine the acceleration or shape of the animation curve.

- **Source:** This drop-down menu has three options based on how the comp is created from DaVinci Resolve's Edit page.

  Transition: This setting is automatically selected when the comp is created from an Edit page transition effect. If the duration of the transition is updated in the Edit page, the timing of the animation updates as well.

  Duration: Use this setting when the comp is created from a clip on the Edit page. The animation timing will update if the clip's duration changes by trimming.

  Custom: Displays an Input dial to manually control the timing.

- **Input:** This dial is only visible when Source is set to Custom. It is used to change the input keyframe value.
- **Curve:** The Curve drop-down menu selects the interpolation method used between keyframes. The three choices are: linear, easing, or custom.

  Linear: The default Linear interpolation method maintains a fixed, consistent acceleration between keyframes.

  Easing: Displays interpolation menus for both the start of the curve (In) and the end of the curve (Out).

  Custom: Opens a mini Spline Editor to customize the interpolation from the start of the animation to the end.

- **Mirror:** Plays the animation forward, and after reaching the end, it returns to the starting value. This causes the initial animation to be twice as fast, since the second half of the comp is used for the reverse animation.
- **Invert:** Flips the animation curve upside-down so that the values start high and end low.

Scaling

The Scale parameters modify the animation values using relative adjustments.

- **Scale:** This number is a multiplier applied to the value of the keyframes. If the Scale value is 2 and a keyframe has a value of 0, it remains 0. If the Scale value is 2 and a keyframe has a value of 10, the result is as if the keyframe is set to 20. This can be thought of as the ending value for the animation. It is best to set this while viewing the last frame in the comp.
- **Offset:** The offset is added to the keyframe values and can be thought of as the starting value for the animation. It is best to set this while viewing the first frame in the comp.
- **Clip Low:** Ensures the output value never dips below 0.0.
- **Clip High:** Ensures the output value never exceeds 1.0.

Timing

The Timing parameters adjust the animation timing using relative values.

- **Time Scale:** Stretches or squishes the animation, causing it to run faster or slower. A value of 1.0 keeps the animation running for the comp's duration (unless you have customized the animation using other controls in the Modifier).
- **Time Offset:** This value delays the animation as a fraction of its total duration. A value of 0.0 applies no delay. A value of 0.5 delays the animation starting point midway into the comp's duration.

Using the Anim Curves Modifier to Create a Custom Transition

To understand how to use the Anim Curves modifier for a transition, let's create a simple scaling dissolve.

1. Add a standard cross dissolve in the Edit page Timeline.
2. Right-click over the transition and choose Convert to Fusion Cross Dissolve.
3. Right-click over the transition again and choose Open in Fusion page.
4. Add a Transform node to MediaIn1 and to MediaIn2.
5. Select the Transform node attached to the MediaIn2.
6. In the Inspector, right-click the Size control, then choose Modify With > Anim Curves from the contextual menu. Adding this modifier to the Size control will cause the slider to animate from 0 to 1 for the Cross Dissolve's duration.
7. Select the Transform node attached to the MediaIn1.
8. In the Inspector, right-click the Size control, then choose Modify With > Anim Curves from the contextual menu.
9. At the top of the Inspector, click the Modifier tab and click the Invert button. Inverting the animation curves causes MediaIn1 to scale opposite of MediaIn2.
10. In the Modifiers tab, set the Curve drop-down menu to Easing, and experiment with the different ease-in ease-out curve types from the In/Out drop-down menus.

Once you create a Macro from this node tree and save it as a Transition template, you can apply it in the Edit page Timeline. If you change the transition duration in the Edit page, the animation timing will update appropriately.

Using the Anim Curves Modifier with Paths

To understand how to use the Anim Curves modifier's with a Path modifier, let's use the premise that you want to create text that falls from the top of the frame and bounces as it reaches the bottom of the frame.

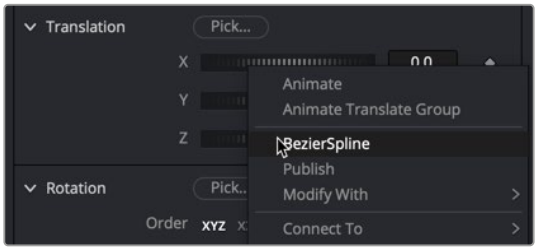
1. In Fusion, create two keyframes that cause text to start at the top of the frame and drop to the bottom. This automatically creates a Path modifier.
2. In the Inspector's Modifier tab, right-click over the Displacement parameter and choose Insert > Anim Curves. The animation is normalized to the duration of the comp.
3. Set the Source menu to Duration, since this is not a transition and we are not customizing the duration.
4. From the Curve menu, choose Easing, then for the Out menu, choose Bounce.
5. Play the animation to see the Bounce animation.
6. To make the bounce occur halfway down the frame, change the Scale to .05.
7. To make the animation run twice as fast, enter 2.0 in the Time Scale parameter.

Once you create a macro from this node tree and save it as a Title template, you can apply it in the Edit page Timeline. If you change the title's duration in the Edit page, the animation timing will update appropriately.

TIP: To view the resulting animation curve in the Spline Editor, select the parameter name in the Spline Editor's header. The spline is updated as you change the controls.

Bézier Spline

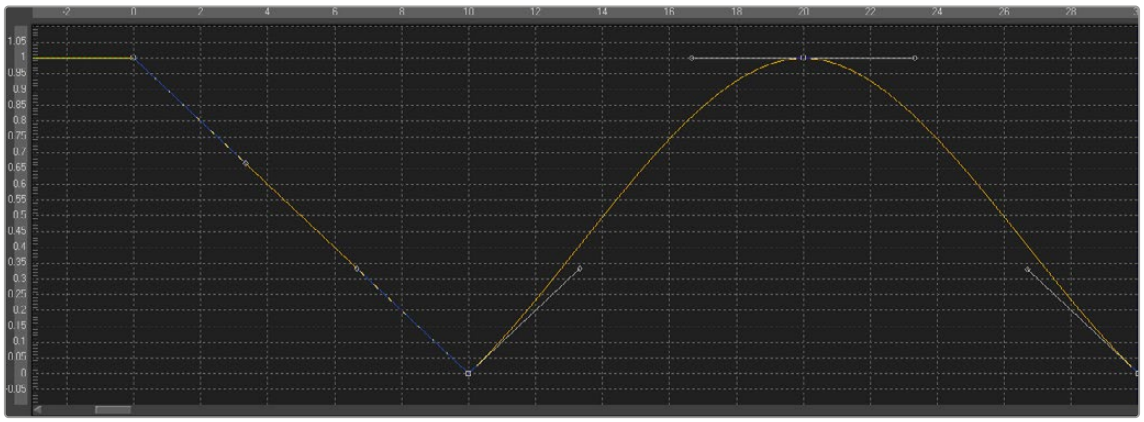
The Bézier Spline is one of the animation modifiers in Fusion and is typically applied to numerical values rather than point values. It is automatically applied when you keyframe a parameter or each time you right-click a number field and select Animate.



Bézier Spline modifier menu

Usage

You can add the Bézier Spline to the Spline Editor by right-clicking a number field and selecting BezierSpline. Since this is the most common choice for animation splines, it is separated from the Modify With menu for quicker access. Selecting BezierSpline from the menu adds a keyframe at the current location and displays a Bézier Spline in the Spline Editor.



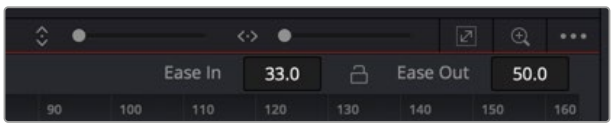
Bézier Spline Editor

Unlike most modifiers, this modifier has no actual Controls tab in the Inspector. However, the Spline Editor displays the Bézier Spline, and it can be controlled there. The Bézier Spline offers individual control over each control point's smoothness using Bézier handles. The smoothness is applied in multiple ways:

- To make the control points smooth, select them, and press Shift-S. The handles can be used to modify the smoothness further.
- To make the control points linear, select them, and press Shift-L. These operations can also be performed using the contextual menu.
- Select the control point(s), right-click, and select Smooth or Linear. The menu also allows the user to smooth a spline using a convolution analysis called a Savitzky-Golay filter. Select the control point(s), right-click, and select Smooth Points -Y Dialog.

Ease In/Out

Traditional Ease In/Out can also be modified by using the number field virtual sliders in the Spline Editor. Select the control points you want to modify, right-click, and select Ease In/Out... from the contextual menu. Then use the number field virtual sliders to control the Ease In/Out numerically.

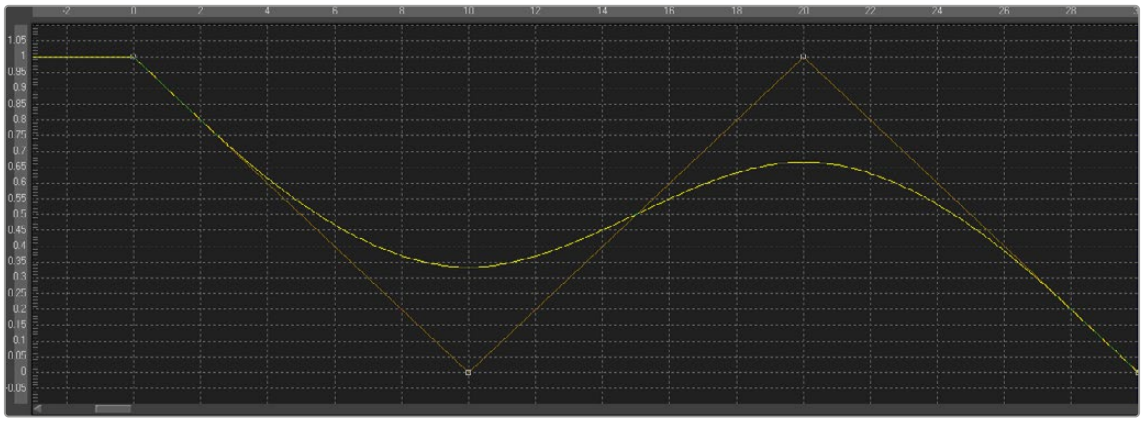


Spline Ease In/Out modifier

B-Spline

An alternative to the Bézier Spline, B-spline is another animation modifier in Fusion and is typically applied to numerical values rather than point values. It is applied by right-clicking a parameter and selecting Modify With > B-Spline.

Usage



B-Spline Editor

- This animation spline modifier has no actual Controls tab. However, the Spline Editor displays the B-spline, and it can be controlled there. Notice that, though the actual value of the second keyframe is 0, the value of the resulting spline is 0.33 due to the unique smoothing and weighing algorithms of a B-spline.
- The weight can be modified by clicking the control point to select it, holding the W key, and moving the mouse left and right to lower or increase the tension. This is also done with multiple selected control points simultaneously.

Calculation

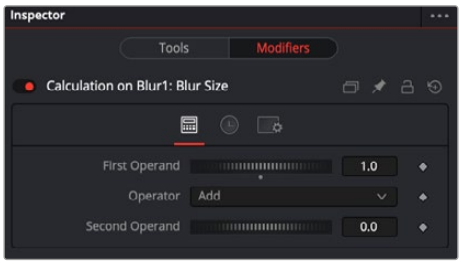
Calculations are used to create indirect connections between parameters. A Calculation can perform a mathematical expression based on two operands, where each operand can be connected to another parameter or set manually.

Additionally, using Time offsets and Scale controls in the Time tab, the Calculation control can access values of a parameter at times other than the current time.

The Calculation's most common use is for connecting two parameters when one value range or scope is inappropriate for the other parameter.

NOTE: The Expression modifier is essentially a more flexible version of the Calculation modifier, with a single exception. It is far easier to manipulate the timing of the operands provided in the Calculation modifier than it is to do so with an Expression.

Inspector



Calculation Calc tab

Calc Tab

The Calc tab includes two dials used for the connected parameter and value that gets mathematically combined. The Operator menu selects how the Second Operand value combines with the parameter's value.

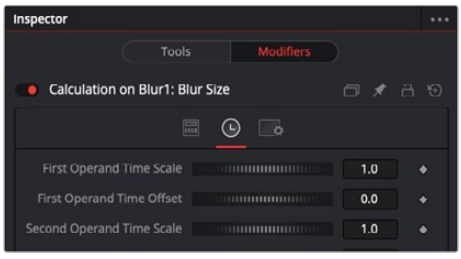
First and Second Operand

These sliders are connected to published or animated parameters or manually set to the desired values for the calculation.

Operator

Select from the mathematical operations listed in this menu to determine how the two operands are combined. Clicking the drop-down arrow opens the menu with the following options:

- Add
- Subtract (First - Second)
- Multiply
- Divide (First / Second)
- Divide (Second / First)
- Subtract (Second - First)
- Minimum
- Maximum
- Average
- First only



Calculation Time tab

Time Tab

The Time tab is used to modify the time of the Calculation modifier. The controls here retime the speed of the effect or offset it in time.

First and Second Operand Time Scale

These sliders multiply the frame number and return the value of the operands at the multiplied frame number. A value of 1 returns the value of the operand at frame x when the composition is on frame x. For example, if the first operand is animated with a value of 1 to 50 from frame 0 to 10, then a scale of 0.5 would cause the calculation to return a value of 25 at frame 10 (effectively slowing the animation by half for the purposes of the calculation).

First Operand and Second Operand Time Offset

These sliders return the value of the operand at the Time Offset specified. A value of 10 would return the value of the operand 10 frames forward in time, and -10 would return the value of the operand 10 frames back in time. See the example below for a practical example.

Example

The following example uses a calculation to apply blur to a Text node in inverse proportion to the size of the text.

- Open a new composition that starts on frame 0 and ends on frame 100.
- At frame 0, add a Text+ node to the composition.
- Enter a small amount of text and set the size to 0.05
- Click the Keyframe button to the right of the Size slider to add a keyframe.
- Move to frame 100 and set the Size value to 0.50.
- Connect a Blur node after the Text+ node.
- View the Blur node in one of the viewers.

  To have the blur decrease in strength as the text gets bigger, a simple “pick whip”-style parameter linking does not work. The controls cannot be directly connected together because the values of the Text Size control are getting bigger instead of smaller.

- Right-click the Blur Size and select Modify With > Calculation from the contextual menu.

  This adds a Calculation modifier to the Blur node. At the top of the Inspector, a new set of controls appears in the Modifiers tab while the Blur node is selected.

- At the top of the Inspector, select the Modifiers tab (F11).
- Right-click the First Operand slider and select Connect To > Text 1 > Size from the contextual menu.

  Although the Blur Size is now connected to the Text Size parameter, this connection isn’t very useful. The maximum value of the Blur Size control is 0.5, which is hardly noticeable as a blur.

- Set the Operator drop-down menu to Multiply.
- Set the Second Operand slider to 100.
- Now the first operand is multiplied by 100, and adjusting the dial gives you a much blurrier blur.
- Switch to the Time tab of the modifier and set the First Operand Time Scale to -1.0.

  Normally, the first operand gets the value of the control it is connected to from the same frame as the current time. So at frame 10, the first operand is set to the same value as the Text size at frame 10. By setting this value to -1, the value is read from one frame back in time whenever the current time of the composition advances by 1 frame.

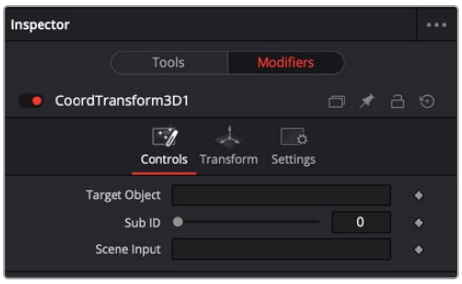
  However, this means that the Calculation would be reading the value of the Text size at frame -10 when we are at frame 10 in the composition.

- To correct for this, set the First Operand Time Offset slider to 100.
- Return to the Tools tab at the top of the Inspector and press Play (Spacebar) to watch how the value of the Blur Size relates to the value of the Text Size.

CoordTransform Position

Because of the hierarchical nature of 3D in Fusion, the original position of an object in a 3D scene often fails to indicate the current position of the object. For example, an image plane might initially have a position at 1, 2, 1, but then be scaled, offset, and rotated by other tools further downstream in the node tree, ending up with an absolute location of 10, 20, 5. This can complicate connecting an object further downstream in the composition directly to the position of an upstream object. The Coordinate Transform modifier can be added to any set of the XYZ coordinate controls to calculate the current position of a given object at any point in the scene hierarchy. To add a Coordinate Transform modifier, right-click the numeric input on any node, and select Modify With > CoordTransform Position from the contextual menu.

Inspector



The Coordinate Transform modifier Controls tab

Controls Tab

The Controls tab has two fields for the target and scene input. The target is for the node continuing the original coordinates, while the scene input is used for the scene with the new coordinates.

Target Object

This control is connected to the 3D tool that produces the original coordinates to be transformed. To connect a tool, drag the node from the Node Editor into the text edit control, or right-click the control and select the tool from the contextual menu. It is also possible to type the tool's name directly into the control.

SubID

The SubID slider can be used to target an individual sub-element of certain types of geometry, such as an individual character produced by a Text 3D tool or a specific copy created by a Duplicate 3D tool.

Scene Input

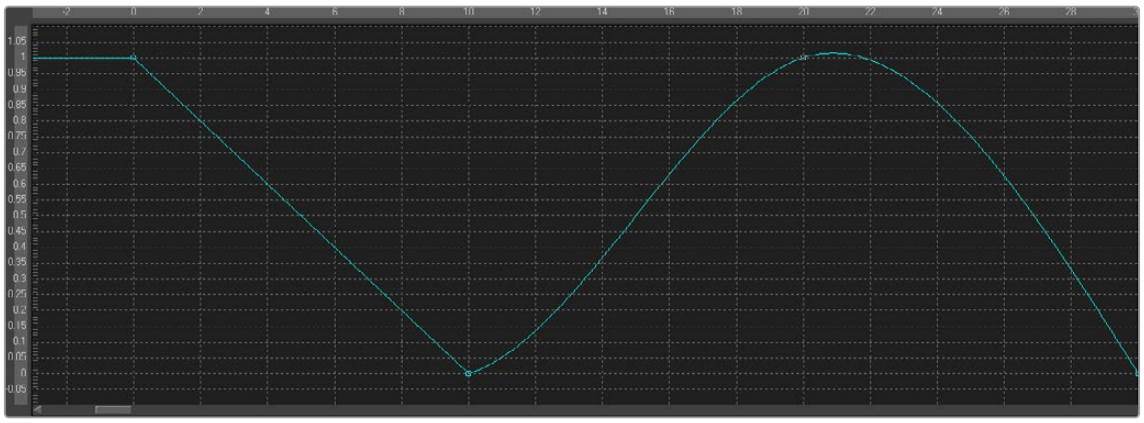
This control should be connected to the 3D tool, which outputs the scene containing the object at the new location. To connect a tool, drag and drop a tool tile from the Node Editor into the text edit control, or right-click the control and select an object from the Connect To pop-up menu.

Cubic Spline

The Cubic Spline is another animation modifier in Fusion that is normally applied to numerical values rather than point values. It can be applied by right-clicking a numerical control and selecting Modify With > Natural Cubic Spline.

Usage

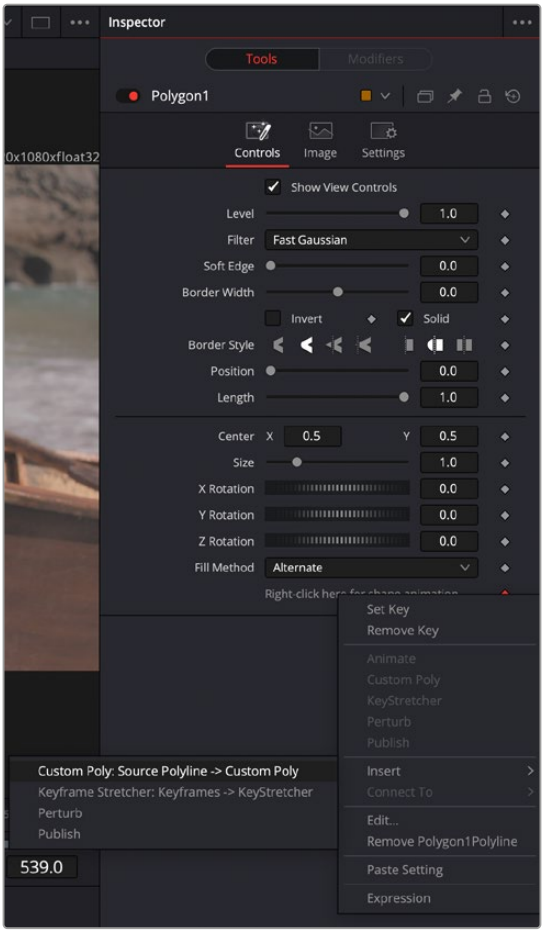
Being an animation spline, this modifier has no actual Controls tab. However, its effect can be seen and influenced in the Spline Editor.



Cubic Spline Editor

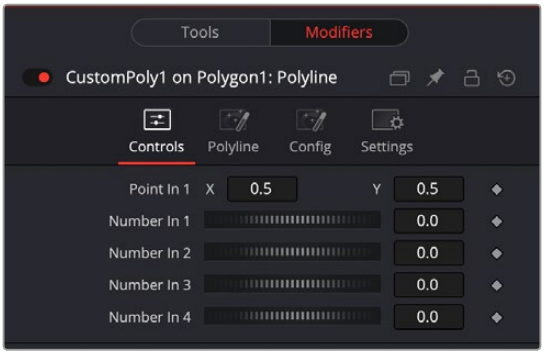
Custom Poly

The Custom Poly modifier can be added to Polygon masks or paths. Similar in function to the Custom and pCustom tools, existing points can be repositioned in the polyline, or replaced completely with a new set of points. The expressions are evaluated for each point on the output polygon. The modifier can be applied by right-clicking on the “Right-click here for shape animation” text at the bottom of the Polygon controls, and selecting Insert > Custom Poly from the contextual menu.



Selecting the Custom Poly modifier

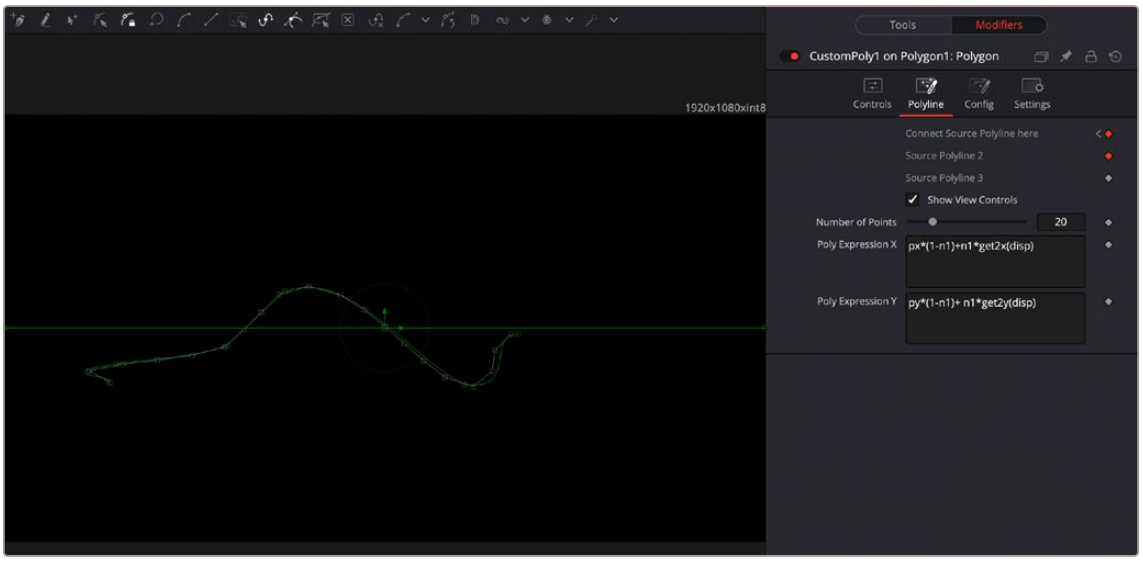
Inspector



The Custom Poly Controls tab

Controls Tab

The Custom Poly Controls tab are available once you select the modifiers tab on the Polygon's inspector. It has a single point input, and number variables to use in animating the expressions. The default is one point and four numbers, but they can be expanded to a maximum of nine. Additional points and numbers can be added from the Config tab.



The Custom Poly Polyline tab

Polyline Tab

The polyline tab exposes the controls to connect and modify the polyline.

Connect Source Polyline here

Allows you to connect other polylines to the modifier.

Show View Controls

This checkbox toggles the polyline controls visibility in the viewer.

Number of Points

The number of output points can be set, controlling the amount of custom subdivision of the polyline. A value of zero uses the number of original source points.

Poly Expression X/Y

Fields for entering the mathematical expressions. For example, px*(1-n1)+n1*get2x(disp) and py*(1-n1)+ n1*get2y(disp).

It uses most of the same expression variables as the Expression modifier below (i.e., n1..n9, p1x..p9x, p1y..p9y, math functions, etc.), and it adds:

- px,py for the current point on the source polygon
- disp for the point's displacement on the polyline (0.0 is start, 1.0 is end)
- index for the current point's index (zero-based)
- num for the number of output points
- getx(disp), gety(disp) to get values from anywhere on the polyline
- getx_at(disp, time), gety_at(disp, time) to get values from polylines at other times
- Similarly, get2x/y(), get3x/y(), get2x/y_at(), get3x/y_at() for the second and third source polys

Config Tab

The config tab allows you to select the amount of Points and Numbers shown in the modifier (to a maximum of 9), and to give them custom names for easier organization.

Expression

An expression is a variable or a mathematical calculation added to a parameter, rather than a straight numeric value. You can add an expression to any parameter in Fusion, or you can add the Expression modifier, which adds several tabs to the modifier Inspector. Adding this modifier to a parameter adds the ability to manipulate that parameter based on any number of controls, either positional or value-based. This modifier offers exceptional flexibility compared to the more limited Calculation or Offset modifiers, but it is unable to access values from frames other than the current time.

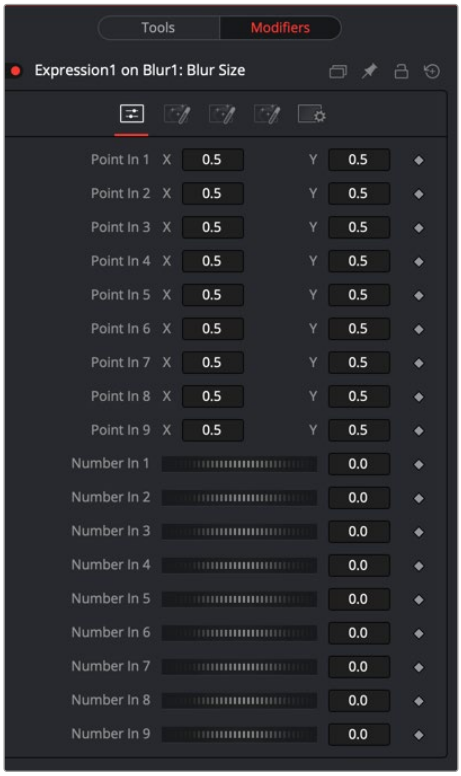
The Expression modifier accepts nine value inputs and nine position inputs that are used as part of a user-defined mathematical expression to output a value.

To add the Expression modifier to a parameter, right-click the parameter in the Inspector and choose Modify With > Expression from the contextual menu. The type of value returned by the Expression depends entirely on the type of control it is modifying.

When used with a value control (like a slider), the Expression in the Number Out tab is evaluated to create the result. When used to modify a positional control (like Center), the Point Out tab controls the result.

The Inspector's Modifiers tab contains the controls for the Expression modifier, described below.

Inspector

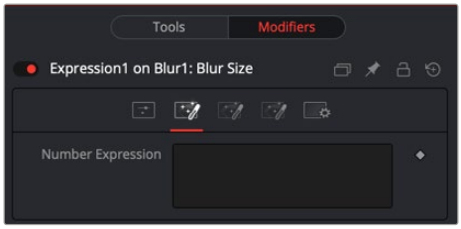


Expressions modifier controls

Controls Tab

This tab provides nine number controls and nine point controls. The values of the number controls can be referred to in an expression as n1 through n9. The X-coordinate of each point control can be referred to as p1x through p9x, while the Y-coordinate is p1y through p9y.

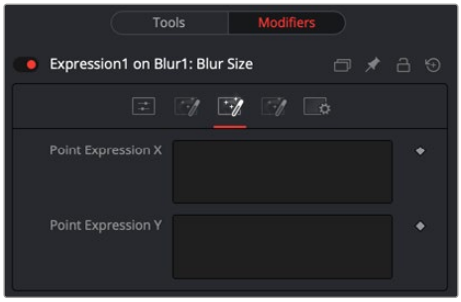
These values can be set manually, connected to other parameters, animated, and even connected to other Expressions or Calculations.



The Number Out tab

Number Out Tab

Mathematical formulas are entered using the Number In and Point In values from the Controls tab. The output modifies the control in which the Expression was applied. See below for the syntax to use in this field.



The Point Out tab

Point Out Tab

The two text boxes in this tab use mathematical formulas, accessing the Number In and Point In values from the Controls tab. The output value modifies the control in which the Expression was applied. The Expression in the top text box control is used to calculate the X-axis value, and the bottom text box is used to calculate the Y-axis control. See below for the syntax to use in this field.



The Expressions modifier Config tab

Config Tab

A good expression is reused over and over again. As a result, it can be useful to provide more descriptive names for each parameter or control and to hide the unused ones. The Config Tab of the Expressions modifier can customize the visibility and name for each of the nine point and number controls.

Random Seed

The Random Seed control sets the starting number for the Rand() function. The rand(x, y) function produces a random value between X and Y, producing a new value for every frame. As long as the setting of this Random Seed slider remains the same, the values produced at frame x are always the same. Adjust the Seed slider to a new value to get a different value for that frame.

Show Number or Point X

There are eighteen of these checkbox controls, one for each of the nine Number and nine Point inputs. Enable this checkbox to display the control for Number x or Point x in the Controls tab.

Name for Number or Point X

There are eighteen of these edit controls, one for each of the nine Number and nine Point inputs. Type a new name for the input into this edit control to assign a new name for the Input's label in the Controls tab.

Expression Syntax Formulas

Formulas are entered into the Number Out or Point Out tabs as part of an expression. They can be made up of the following functions:

n1...n9	The value of Number Input 1..9.
p1x...p9x	The X of Positional Control 1..9.
p1y...p9y	The Y of Positional Control 1..9.
time	The current time (frame number).
pi	The value of pi.
e	The value of e.
log(x)	The base-10 log of x.
ln(x)	The natural (base-e) log of x.
sin(x)	The sine of x (x is degrees).
cos(x)	The cosine of x (x is degrees).
tan(x)	The tangent of x (x is degrees).
asin(x)	The arcsine of x, in degrees.
acos(x)	The arccosine of x, in degrees.
atan(x)	The arctangent of x, in degrees.
atan2(x, y)	The arctangent of x,y, in degrees.
abs(x)	The absolute (positive) value of x.
int(x)	The integer (whole) value of x.
frac(x)	The fractional value of x.
sqrt(x)	The Square Root of x.
rand(x, y)	A random value between x and y.
rands(x, y, s)	A random value between x and y, based on seed s.
min(x, y)	The minimum (lowest) of x and y.
max(x, y)	The maximum (highest) of x and y.
dist(x1, y1, x2, y2)	The distance between point x1,y2 and x2,y2.
dist3d(x1,y1,z1,x2,y2,z2)	The distance between 3D points x1,y2,z1 and x2,y2,z2
noise(x)	A smoothly varying Perlin noise value based on x
noise2(x, y)	A smoothly varying Perlin noise value based on x and y
noise3(x, y, z)	A smoothly varying Perlin noise value based on x, y and z
if(c, x, y)	Returns x if c not 0, otherwise y.

Expression Syntax Operators

Operators are used to evaluate statements. They are combined with functions to perform logical and mathematical calculations in the Number Out and Point Out tabs.

x + y	x plus y.
x - y	x minus y.
x > = y	1.0 if x is less than y, otherwise 0.0.
x > y	1.0 if x is greater than y, otherwise 0.0.
!x	1.0 if x = 0, otherwise 0.0.
-x	(0.0 - x).
+x	(0.0 + x) (effectively does nothing).
x ^ y	x raised to the power of y.
x y	x multiplied by y.
x y	x divided by y.
x % y	x modulo y, (remainder of (x divided by y)).
x > = y	1.0 if x is less than or equal to y, otherwise 0.0.
x > = y	1.0 if x is greater than or equal to y, otherwise 0.0.
x = y	1.0 if x is exactly equal to y, otherwise 0.0.
x == y	1.0 if x is exactly equal to y, otherwise 0.0 (identical to above).
x < > y	1.0 if x is not equal to y, otherwise 0.0.
x != y	1.0 if x is not equal to y, otherwise 0.0 (identical to above).
x & y	1.0 if both x and y are not 0.0, otherwise 0.0.
x && y	1.0 if both x and y are not 0.0, otherwise 0.0 (identical to above).
x | y	1.0 if either x or y (or both) are not 0.0, otherwise 0.0.
x || y	1.0 if either x or y (or both) are not 0.0, otherwise 0.0 (identical to above).

Example 1

To make a numeric control equal to the Y value of a motion path, add an expression to the desired target control and connect the Path to Point In 1. Enter the formula:

```
p1y
```

into the Number Out field.

Example 2

To make the result of the Expression's Number Out be the largest of Number In 1 and Number In 2, multiplied by the cosine of Number In 3, plus the X coordinate of Point In 1, enter the formula:

```
max(n1, n2) * cos(n3) + p1x
```

into the Number Out field.

Example 3

Add a Background node set to solid black and a Hotspot node. Set the Hotspot size to 0.08 and set the Strength to maximum. Modify the Hotspot center with an expression. Change the current frame to 0.

Set n1 to 0.0 and add a Bézier Spline. At frame 29, set the value of n1 to 1.0. Select both points and loop the spline using the Spline Editor. Now enter the following equations into the Point Out tab of the expression.

```
X-Axis Expression

n1

Y-Axis Expression

0.5 + sin(time*50) 4
```

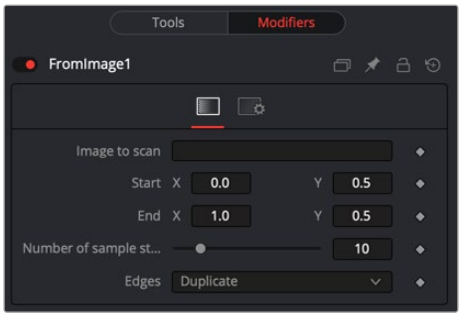
Render out a preview and look at the results. (Try this one with motion blur.)

From Image

The From Image modifier only works on gradients, like the gradient on a Background node. It takes samples of an image along a user-definable line and creates a gradient from those samples.

Unlike other modifiers, From Image is not located in the Modify With menu. This modifier can be applied by right-clicking a Gradient bar in the Inspector and selecting From Image.

Inspector



From Image modifier

Controls Tab

The From Image controls tab in the Inspector is used to select the node that contains the image you want to sample. It allows you to define the starting and ending point in the image as well as how many color samples to use in creating the gradient.

Image to Scan

Drop into this box the node from the Node Editor that you want to be color sampled.

Start X/Y, End X/Y

These two point controls define the Start and End points of the line along which the samples are taken from the image defined in the Image to Scan box.

The points can also be moved directly in the viewer.

Number of Sample Steps

This defines how many individual color samples are taken along the line. You can also see the result of this setting when you look at the actual node's Gradient control. The more sample steps you define here, the more individual color stops appear on the Gradient control. It is also possible to first create a gradient using the From Image modifier and then remove that modifier from the Gradient control again. The created gradient stays intact and can then be fine tuned by hand.

Edges

Edges determines how the edges of the image are treated when the sample line extends over the actual frame of the image to be sampled.

Black

This outputs black for every point on the sample line outside of the image bounds.

Wrap

This wraps the edges of the line around the borders of the image.

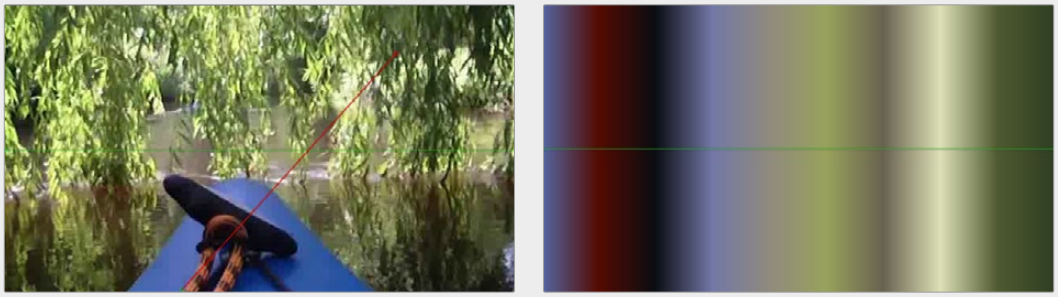
Duplicate

This causes the edges of the image to be duplicated as best as possible, continuing the image beyond its original size.

Color

This outputs a user-definable color instead of black for every point on the sample line outside of the image bounds.

Example



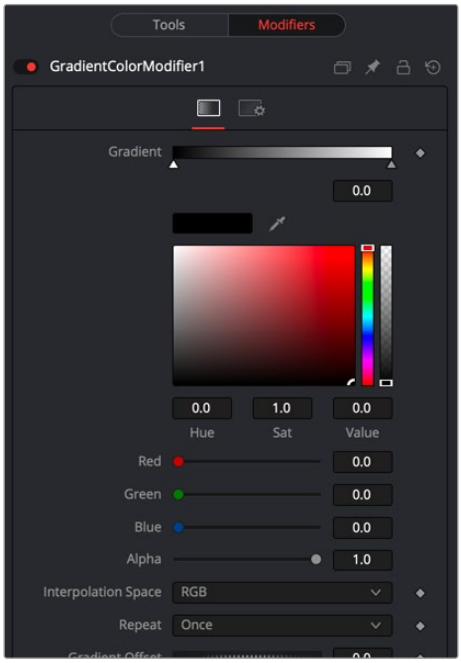
The source image on the left shows the color selection line in red. The image on the right shows the resulting gradient from that selection.

Gradient Color

The Gradient Color modifier allows you to create a customized gradient and map it into a specific time range to control a value. Use the Start and End time controls to set the frames for the animation. If the Start and End time values are set to 0, then the modifier returns the value at the starting point of the gradient. You can use the Offset control to animate the gradient manually.

It can be applied by right-clicking a parameter and selecting Modify With > Gradient Color.

Inspector



The Gradient modifier Controls tab

Controls Tab

The Controls tab consists of a Gradient bar where you add and adjust points of the gradient. Start Time and End Time thumbwheels at the bottom of the Inspector determine the time range the gradient is mapped into.

Gradient

The Gradient control consists of a bar where it is possible to add, modify, and remove points of the gradient. Each point has its color. It is possible to animate the color as well as the position of the point. Furthermore, a From Image modifier can be applied to the gradient to evaluate it from an image.

Gradient Interpolation Method

The gradient is linearly interpolated from point to point in RGB color space by default. This can sometimes result in unwanted colors. Choosing another color space may provide a better result.

Repeat

Defines how the left and right borders of the gradient are treated.



Gradients set to Once, Repeat, and Ping Pong from top to bottom, respectiμvely, and shifting the gradient to the left.

- **Once:** When using the Gradient Offset control to shift the gradient, the border colors keep their values. Shifting the default gradient to the left results in a white border on the left; shifting it to the right results in a black border on the right.
- **Repeat:** When using the Gradient Offset control to shift the gradient, the border colors are wrapped around. Shifting the default gradient to the left results in a sharp jump from white to black; shifting it to the right results in a sharp jump from black to white.
- **Ping Pong:** When using the Gradient Offset control to shift the gradient, the border colors ping-pong back and forth. Shifting the default gradient to the left results in the edge fading from white back to black; shifting it to the right results in the edge fading from black back to white.

Gradient Offset

Allows you to pan through the gradient.

Time Controls

The Start Time and End Time thumbwheels determine the time range the gradient is mapped into. This is set in frames. The same effect can be achieved by setting the Gradient to Once and animating the offset thumbwheel.

Key Stretcher Modifier

The Keyframe Stretcher modifier is primarily used when creating title templates in Fusion for use in DaVinci Resolve's Edit page or Cut page. The Keyframe Stretcher modifier is added to an animated parameter, so the keyframes in the animated parameter stretch when the template is trimmed in the Timeline. It can be applied by right-clicking a parameter and selecting Modify with > KeyStretcher.

For more information on the Keyframe Stretcher Modifier controls, refer to the "Keyframe Stretcher", *see Chapter 51, "Miscellaneous Nodes,"* in the Fusion Reference Manual.

MIDI Extractor

The MIDI Extractor modifier provides the ability to modify the value of a control using the values stored in a MIDI file. This modifier relies on some knowledge of MIDI, which is beyond the scope of this manual.

The value produced by the modifier is extracted from the MIDI event selected in the Mode menu. Each mode can be trimmed so that only specific messages for that event are processed—for example, only some notes are processed, while others are ignored. The value of the event can be further scaled or modified by additional factors, such as Scale, Velocity, Attack, and Decay.

It can be applied by right-clicking a parameter and selecting Modify With > MIDI Extractor.

Inspector



The MIDI Extractor modifier Controls tab

Controls Tab

The Controls tab is used to load the MIDI file, modify its timing, and determine which MIDI messages and events trigger changes in the Fusion parameter.

MIDI File

This browser control is used to specify the MIDI file that is used as the input for the modifier.

Time Scale

Time Scale is used to specify the relationship between time as the MIDI file defines it and time as Fusion defines it. A value of 1.0 plays the MIDI events at normal speed, 2.0 plays at double speed, and so on.

Time Offset

Time Offset adjusts the sync between the MIDI file's timing and Fusion's timing. If there is an unexpected delay, or if the MIDI file should start partway into or before some animation in Fusion, this control can be used to offset the MIDI data as required.

Result Offset, Result Scale

These sliders adjust the range of values produced by the modifier. By default, values between 0 and 1 (or -1 and 1 for PitchBend mode) are generated. This does not always suit the node/parameter, and scale can be used to make this range larger (such as * 0.0 - 2.0). Offset is used to provide some constant value as a base.

Result Curve

The Result Curve can also be used to adjust the output. However, this adjusts the curve of the result. By default, for any input MIDI data, the results fall linearly between 0.1 and 1.0 (for example, a velocity 127 note generates 1.0, whereas 63 generates approximately 0.5).

The Result Curve applies a gamma-like curve so that middle values can produce higher or lower results while still maintaining the full scale.

Mode

This menu provides Beat, Note, Control Change, Poly AfterTouch, Channel AfterTouch, or Pitch Bend, indicating from which MIDI event the values are being read. Beat mode is slightly different in that it produces regular pulses based on the tempo of the MIDI file (including any tempo maps).

The Beat mode does not use any specific messages; it bases its event timing on the tempo map contained in the MIDI file.

Combine Events

This menu selects what happens when multiple events occur at the same time. In Notes mode, this can happen easily. For other events, this can happen if Multiple Channels are selected.

Use this to take the result from the most recent event to occur, the oldest event still happening, the highest or lowest valued event, the average, sum, or the median of all events currently occurring.

Beat (Quarters) (Beat Mode Only)

This defines how often a beat occurs when in Beat mode. This is in quarter notes, so a value of 1.0 gives a beat every quarter.

Note Range (Note and Poly Aftertouch Modes Only)

This defines what range of notes causes a value to be generated. For example, use this to pick out the kick drum from a GM drum track by setting the note range between 35–36.

Pitch Scale (Note Mode Only)

Pitch Scale defines how much the result changes with pitch. A value of 1.0 causes the result to vary from 0.0 to 1.0 over the entire range.

Velocity Scale (Note Mode Only)

This defines how much the result changes with velocity. A value of 1.0 causes the result to vary from 0.0 to 1.0 over the entire range. This is added to the result from Pitch Scale for the final result.

Control Number (Control Change Mode Only)

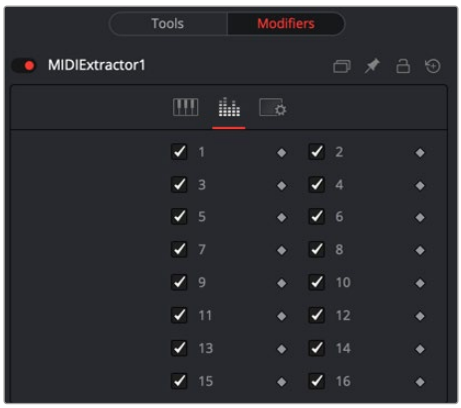
This specifies the MIDI controller number from which to extract events.

Envelope Controls (Note and Beat Modes Only)

These define an Envelope to follow for values before, during, and after the note or beat. Pre-Attack Time defines how long before the event it starts ramping up to the pre-attack level. Attack is the Time/Level to ramp to once the event has occurred, followed by the Decay ramp and Sustain, until the event stops. This stage is for Notes only. Beats have an instantaneous duration, so it goes straight to

Release. Release is the ramp-down time after the event finishes. When trying to do a Beat, set Release to some value, or there likely will not be much of a beat.

These values can be used to follow actual sounds in the MIDI sequence or just to create interesting effects. All time values used in the MIDI Extractor are in seconds.



The MIDI Extractor modifier Channels tab

Channels Tab

The Channels tab is used to select the Channels used in the modifier.

Channels

Channels checkboxes select which of the 16 channels in the MIDI file are actually considered for events. This is a good way to single out a specific instrument from an arrangement.

About MIDI

A single MIDI interface allows 16 channels. Typically, these are assigned to different instruments within a device or different devices. Usually, MIDI data is 7 bits, ranging from 0–127. In Fusion, this is represented as a value between 0–1 to be more consistent with how data is handled in Fusion.

There are numerous different MIDI messages and events, but the ones that are particularly useful with this modifier are detailed below.

MIDI Messages

Note On: This indicates that a note (on a specific channel) is being turned on, has a pitch (0–127, with middle C being 60) and a Velocity (0–127, representing how fast the key or strings or whatever was hit).

Note Off: This indicates that a note (on a specific channel) is being turned off, has a pitch (0–127, with middle C being 60) and a Velocity (0–127, representing how fast the key or strings or whatever was released).

Control Change: This message indicates that some controller has changed. There are 128 controllers (0–127), each of which has data from 0–127. Controllers are used to set parameters such as Volume, Pan, amount of Reverb or Chorus, and generic things like foot controllers or breath controllers.

MIDI Events

Channel Aftertouch: This event defines that pressure is being applied to the keys (or strings or whatever) during a note. This represents general, overall pressure for this channel, so it simply uses a pressure value (0–127).

Poly Aftertouch: This event defines that pressure is being applied to the keys (or strings or whatever) during a note. It is specific to each particular note and therefore contains a note number as well as a pressure value (0–127).

Pitch Bend

The Pitch Bend controller generally specifies the degree of pitch bending or variation applied to the note. Because pitch bend values are transmitted as a 14-bit values, this control has a range between -1 and 1 and a correspondingly finer degree of resolution.

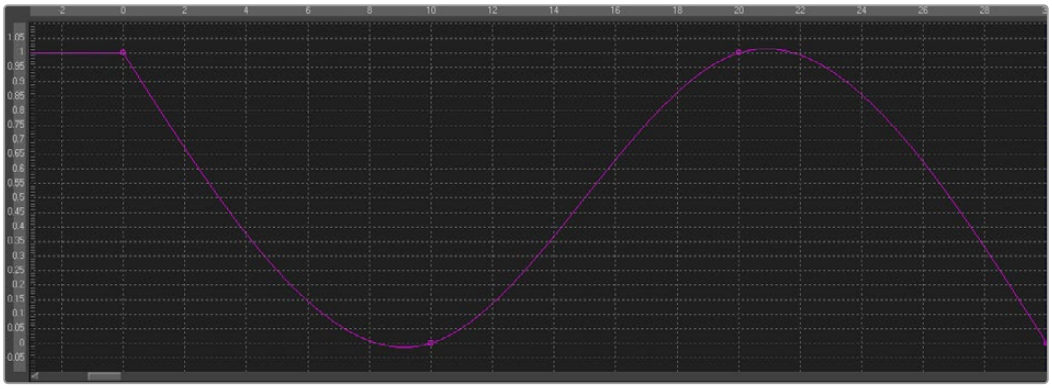
Natural Cubic Spline

The Natural Cubic Spline is one of the animation modifiers in Fusion and normally is applied to numerical values rather than point values. It can be applied by right-clicking a numerical control and selecting Modify With > Natural Cubic Spline.

NOTE: Unlike other spline types, Cubic splines have no control handles. They attempt to automatically provide a smooth curve through the control points.

Usage

Being an animation spline, this modifier has no actual Controls tab. However, its effect can be seen and influenced in the Spline Editor.



Natural Cubic Spline Editor

Offset (Angle, Distance, Position)

There are three Offset modifiers used to create variances between values. Depending on the modifier, these values relate to controls, paths, and points. The three types of Offset modifiers available in Fusion are:

- Offset Angle
- Offset Distance
- Offset Position

Offset Angle

The Offset Angle modifier outputs a value between 0 and 360 that is based on the angle between two positional controls. The Position and Offset parameters may be static, connected to other positional parameters, or connected to paths of their own. All offsets use the same set of controls, which behave differently depending on the offset type used. These controls are described below.

Offset Distance

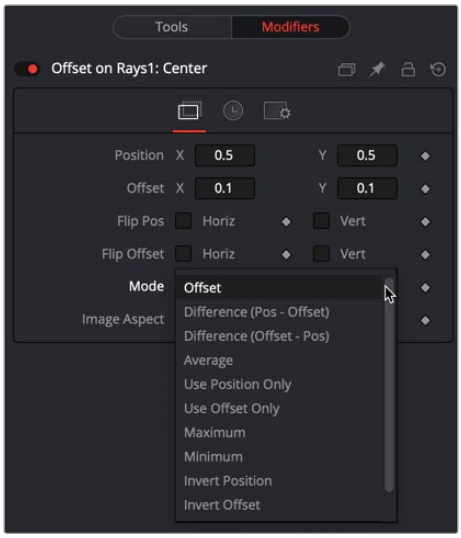
The Offset Distance modifier outputs a value that is based on the distance between two positional controls. This modifier is capable of outputting a value based on a mathematical expression applied to a position.

Offset Position

The Offset Position modifier generates a position (X and Y coordinates) that is based on the relationship between positional controls. This modifier is the equivalent of a calculation control except that it outputs X and Y coordinates instead of a value.

It can be applied by right-clicking a control and selecting Modify With > Offset.

Inspector



The Offset Position modifier controls

Offset Tab

The Inspector for all three Offset modifiers is identical. The Offset tab includes Position and Offset values as well as a Mode menu for selecting the mathematical operation performed by the offset control.

Position X and Y

The X and Y values are used by the Position to generate the calculation.

Offset X and Y

The X and Y values are used by the Offset to generate the calculation.

Flip Position Horizontal and Vertical

When these controls are selected, the Position is mirrored along the vertical or horizontal axis of the image.

Flip Offset Horizontal and Vertical

When these controls are selected, the Offset position is mirrored along the vertical or horizontal axis of the image.

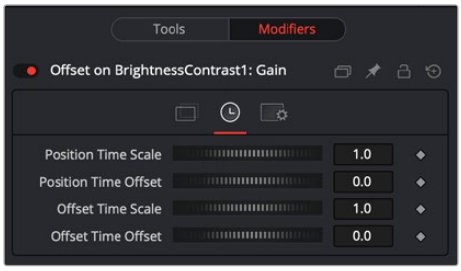
Mode

The Mode menu includes mathematical operations performed by the Offset control. Available options include:

- Offset
- Difference (Position - Offset)
- Difference (Offset - Position)
- Average
- Use Position Only
- Use Offset Only
- Maximum
- Minimum
- Invert Position
- Invert Offset
- Invert Sugar
- Random Offset

Image Aspect

Adjust the modifier's output to compensate for the image aspect (not pixel aspect) of the project. A square image of 500 x 500 would use an Image Aspect value of 1, and a rectangular image of 500 x 1000 would use an Aspect Value of 2. The default value is always based on the current frame format selected in the preferences. To calculate image aspect, divide the width by the height. This control can also be used to create the illusion of aspect.



The Offset Time tab

Time Tab

Position Time Scale

This returns the value of the Position at the Time Scale specified (for example, 0.5 is the value at half the current frame time).

Position Time Offset

This returns the value of the Position at the Time Offset specified (for example, 10 is 10 frames back).

Offset Time Scale

This returns the value of the Offset at the Time Scale specified.

Offset Time Offset

This returns the value of the Offset at the Time Offset specified.

Example

This is a simple comp to illustrate one potential use of offsets.

- Create a new Comp 100 frames long.
- Create a node tree consisting of a black background and a Text node foreground connected to a Merge.
- In the Text Layout tab, use the Center X control to animate the text from the left side of the screen to the right.
- Move to frame 0.
- In the Text tab in the Inspector, right-click the Size control and select Modify With > Offset Distance from the contextual menu.
- This adds two onscreen controls: a crosshair for the position and an X control for the offset. These onscreen controls represent the Position and Offset controls displayed in the Modifiers tab.

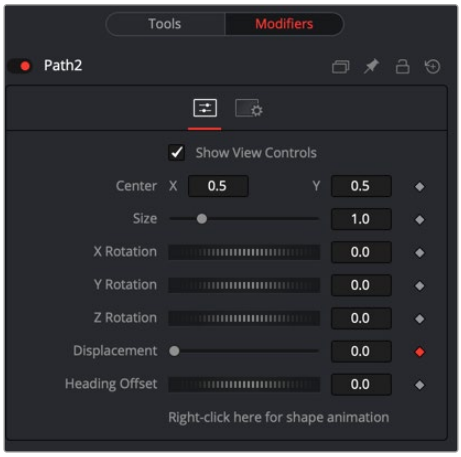
  The size of the text is now determined by the distance, or offset, between the two onscreen controls.
- Drag the X onscreen control in the viewer to see how the distance from the crosshair changes the size of the merge and by association the text.
- Both the crosshair and the X onscreen controls are animatable and can be connected to other controls.
- Position the X centered at the bottom of the viewer.
- In the Inspector, select the Modifiers tab.
- In the Offset on Text size section, right-click over Position and choose Connect To > PathConnect the position value of the Offset to the existing path by right-clicking the Position control and selecting Connect To > Path1 Position.
- Play the comp to view the animation.

Now, the text shrinks near the center of the path (when the distance between the offset and the path is at its minimum) and grows at its ends (where the distance between the offset and the path is at its maximum).

Path

The Path modifier uses two splines to control the animation of points: an onscreen motion path (spacial) and a Time spline visible in the Spline Editor (temporal). To animate an object's position control using a Path, right-click the Position control either in the Inspector or in the viewer and select Path from the contextual menu. This adds a keyframe at the current position. You can begin creating a path by moving the playhead and dragging the center position control in the viewer. The Spline Editor shows a displacement spline for editing the temporal value, or "acceleration," of the path.

Inspector



The Path modifier Controls tab

Controls Tab

The Controls tab for the path allows you to scale, reposition, and rotate the path. It also provides the Displacement parameter, allowing you to control the acceleration of an object attached to the path.

Center

The actual center of the path. This can be modified and animated as well to move the entire path around.

Size

The size of the path. Again, this allows for later modification of the animation.

X Y Z Rotation

The path can be rotated in all three dimensions to allow for sophisticated controls.

Displacement

Every motion path has an associated Displacement spline in the Spline Editor. The Displacement spline represents the position of the animated element along its path, represented as a value between 0.0 and 1.0. Displacement splines are used to control the speed or acceleration of an object's movement along the path.

To slow down, speed up, stop, or even reverse the motion of the control along the path, adjust the values of the points for the path's displacement in the Spline Editor or in the Inspector.

- A Displacement value of 0.0 in the Spline Editor indicates that the control is at the very beginning of a path.
- A value of 1.0 indicates that the control is positioned at the end of the path.
- Each locked point on the motion path in the viewer has an associated point on the Displacement spline.
- Unlocked points have a control point in the viewer but do not have a corresponding point on the Displacement spline.

Heading Offset

Connecting to the Heading adjusts the auto orientation of the object along the path. For instance, if a mask's angle is connected to the path's heading, the mask's angle will adjust to follow the angle of the path.

Right-Click Here for Shape Animation

It's possible to animate the shape of the path as well or to connect it to other path controls like Polyline Masks or Paint Strokes.

TIP: Switching Default Paths

You can change the default path type used when animating a position or center control to a path (if this is the preferred type of animation). Open the Preferences window and select the Global Settings. In the Default category, select the Point With menu and choose Path. The next time Animate is selected from a Position or Center control's contextual menu, a path is used.

Perturb

The Perturb modifier generates smoothly varying random values for a given parameter based on Perlin noise. It can be used to add jitter, shake, or wobble to any animatable parameter, even if the parameter is already animated. Its results are similar to those of the Shake modifier, although it uses a different set of controls that may be more intuitive. Unlike other random modifiers, you can apply the Perturb modifier to polylines, shapes, grid meshes, and even color gradients.

For example, to add camera shake to an existing path, right-click the crosshair and choose Insert > Perturb, and then adjust the Strength down to suit. Alternatively, right-clicking the path's "Right-click here for shape animation" label at the bottom of the Inspector lets you apply perturb to the path's polyline. This works best if the polyline has many points—for example, if it has been tracked or hand-drawn with the Draw Append pencil tool. A third usage option is to use the Insert contextual menu to insert the modifier onto the Displacement control. This causes the motion along the path to jitter back and forth without actually leaving the path.

NOTE: Perturb can only add jitter; it cannot smooth out existing animation curves.

Inspector



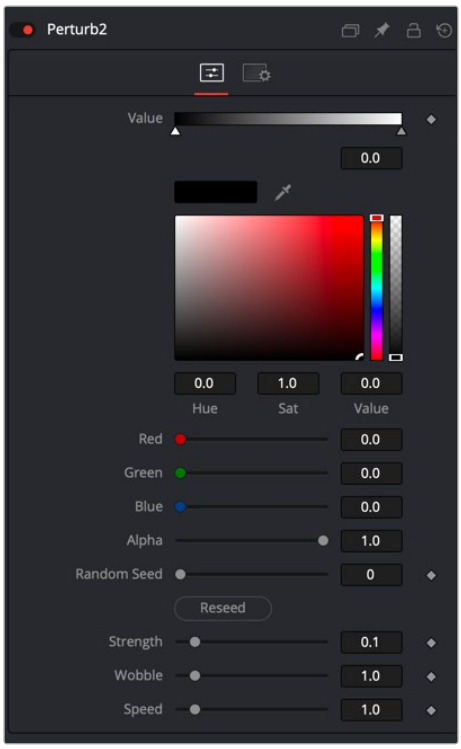
The Perturb modifier Controls tab

Controls Tab

The Controls tab for Perturb is mainly used for controlling the Strength, Wobble, and Speed parameters of the random jitter.

Value

The content of this control depends on what type of control the modifier was applied to. If the Perturb modifier was added to a basic Slider control, the Value is a slider. If it was added to a Gradient control, then a Gradient control is displayed here. Use the control to set the default, or center value, for the Perturb modifier to work on.



The Perturb modifier Gradient controls

Jaggedness

(Polylines and meshes only) This allows you to increase the amount of variation along the length of the polyline or mesh, rather than over time. Increasing Jaggedness gives a squigglier polyline or more tangled mesh, independent of its movement.

Phase

(Polylines and meshes only) Animating this can be used to move the ripple of a polyline or mesh along itself, from end to end. The effect can be most clearly seen when Speed is set to 0.0.

Random Seed Randomize

The Random Seed is used to "seed" the amount of jitter applied by the modifier. Two Perturb modifiers with identical settings, but different random seeds, produce two completely different results. Click the Randomize button to assign a random seed value.

Strength

Use this control to adjust the strength of the Perturb modifier's output, or its maximum variation from the primary value specified above.

Wobble

Use the Wobble control to determine how smooth the resulting values are. Less wobble implies a smoother transition between values, while more wobble produces less predictable results.

Speed

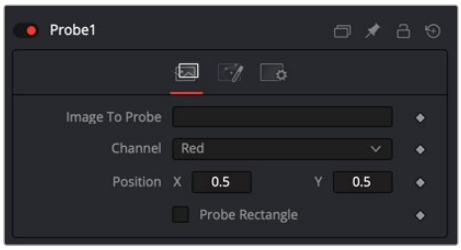
Increasing the Speed slider value speeds up the rate at which the value changes. This can increase the apparent wobbliness in a more predictable fashion than the Wobble control and make the jitter more frantic or languorous in nature.

Probe

The Probe modifier is one of the most versatile modifiers in Fusion. It allows you to control any numeric parameter by the color or luminosity of a specific pixel or rectangular region of an image. Think of driving the Brightness node by probing the pixel values of flickering lights in a shot, or measuring graded LUTs to compare values.

It can be applied by right-clicking a parameter and selecting Modify With > Probe.

Inspector



The Probe modifier Controls tab

Controls Tab

The Controls tab for the Probe modifier allows you to select the node to probe, define the channel used to drive the parameter, and control the size of the probed area.

Image to Probe

Drag a node from the Node Editor to populate this field and identify the image to probe.

Channel

Select the channel you want to probe. The usual options are:

- Red
- Green
- Blue
- Alpha

Luma

Once a Probe modifier is present somewhere in your comp, you can connect other node's values to its outputs as well. The Probe allows you to connect to its values individually:

- Result
- Red
- Green
- Blue
- Alpha

Position X Y

The position in the image from where the probe samples the values.

Probe Rectangle

By default, the Probe samples only the value of a single pixel at its position. By using the Probe Rectangle mode, you can sample from a larger area of pixels based on the Evaluation method.

Width Height Controls

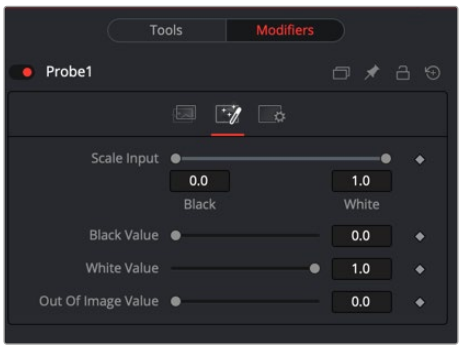
These determine the size of the area to be probed.

Evaluation

Sets how the pixels inside the rectangle are computed to generate the output value.

Options include:

- **Average:** All pixel values inside the rectangle are averaged.
- **Minimum:** The smallest value of all pixels inside the rectangle is used.
- **Maximum:** The highest value of all pixels inside the rectangle is used.



The Probe modifier Value tab

Value Tab

The Value tab controls the range or scale of the modifier adjustment, thereby adjusting the sensitivity of the Probe.

Scale Input

By default, the Probe generates the Black Value when the probed area results in a value of 0 (i.e., black), and it generates its White Value when the probed area results in a value of 1 (i.e., white). By using this range control, you can modify the sensitivity of the Probe.

Black Value

The value that is generated by the Probe if the probed area delivers the result set in Scale Input Black.

White Value

The value that is generated by the Probe if the probed area delivers the result set in Scale Input White.

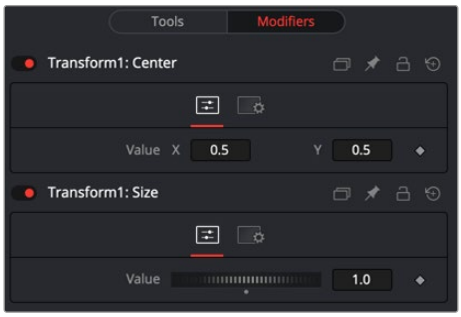
Out of Image Value

The value that is generated by the Probe if the probed area is outside the frame boundaries of the probed image. If probing a rectangle, this value does not generate before the entire rectangle is outside the frame boundaries of the image to be probed.

Publish

Only parameters that are animated will be available from the Connect To menu. To connect to non-animated parameters, you must Publish them first. Animated controls are automatically published, whereas static controls must be published manually.

To publish a static control, right-click the control and select Publish from the contextual menu.



The Publish modifier Controls tab

Controls Tab

The Controls tab shows the published control available for linking to other controls.

Published Value

The display of the published control is obviously dependent on which control is published from which node.

Resolve Parameter

The Resolve Parameter Modifier is used when creating a transition template in Fusion for use in DaVinci Resolve's Edit page or Cut page. When building a transition in Fusion, the Resolve Parameter modifier is added to any control you want to animate. The Resolve Parameter modifier automatically animates the parameter for the duration of the transition, allowing you to trim the transition in the Edit page or Cut page.

For example, to create a cross dissolve, do the following:

1. From the Effects Library, add a Fusion Composition to the Edit page Timeline.
2. In Fusion, add a Dissolve node to the Node Editor.
3. In the Inspector, right-click the Background/Foreground parameter, and then choose Resolve Parameter from the Modifier contextual menu. Adding the modifier to the Background/Foreground parameter automatically updates the slider if the transition is modified back on the Edit/Cut page.
4. In the Node Editor, right-click the Dissolve node and choose Macro > Create Macro.
5. When creating a macro that's to be used as a Fusion transition it's important that two inputs and one output are selected in the Macro Editor. In this example, under the Dissolve heading, enable the Output, Background and Foreground check boxes.
6. Give the transition a name, then save the macro from the top File menu.Save the macro in the following folder so it appears the Edit page or Cut page Effects Library:

   For MacOS: $TEMPLATE_MAC_OS_PATH/Transitions or $TEMPLATE_MAC_USER_PATH/Transitions

   For Windows: $TEMPLATE_WIN_OS_PATH\Transitions or $TEMPLATE_WIN_USER_PATH\Transitions
7. Quit and reopen DaVinci Resolve to update the list of transitions in the Effects Library.
8. On the Edit page, open the Effects Library. Navigate to Video Transitions Fusion Transitions, and the custom Fusion transition will be listed.

Shake

The Shake modifier is used to randomize a Position or Value control to create semi-random numeric inputs. The resulting shake can be entirely random. The motion can also be smoothed for a more gentle, organic feel.

To add the Shake modifier to a parameter, select Modify With > Shake from the parameter's contextual menu. The Shake modifier uses the following controls to achieve its effect. It can be applied by right-clicking a parameter and selecting Modify With > Shake.

Inspector



The Shake modifier Controls tab

Controls Tab

Random Seed

The Random Seed control contains the value used to seed the random number generator. Given the same seed, a random number generator always produces the same results. Change the seed if the results from the randomizer are not satisfying.

Smoothness

This control is used to smooth the overall randomness of the Shake. The higher the value, the smoother the motion appears. A value of zero generates completely random results, with no smoothing applied.

Lock X/Y

This checkbox is used to unlock the X- and Y-axis, revealing independent slider controls for each axis.

Minimum and Maximum

This control is used to determine the overall strength of the shake. The low values represent the lowest value that can be generated by the randomizer, and the high values represent the highest values. To create a Shake that moves a center crosshair anywhere within the image, set the Minimum to 0.0 and the Maximum to 1.0. To restrict the motion to a smaller shake in the bottom-right corner of the image, set the Minimum to 0.70 and the Maximum to 0.90.

Example

- Create a new comp, and then add and view a Text node.
- Type some text in the Text node.
- In the viewer, right-click over the Center control of the text and choose Modify With > Shake Position.
- In the Inspector, select the Modifiers tab and set the smoothing to 5.0.
- Set the Minimum to 0.1 and the Maximum to 0.9.

  This adds some chaotic movement to the text. However, we can change this over Go to frame 0 and in the Inspector click the Keyframe button to the right of both the Minimum and the Maximum controls.

- Go to frame 0 and in the Inspector click the Keyframe button to the right of both the Minimum and the Maximum controls.
- Go to frame 90 and adjust the Minimum to 0.45 and the Maximum to 0.55.
- View the results.

Now, the text starts out by flying all over the screen and tightens in toward the center of the screen as the comp plays.

Track

Although there is a standard Tracker node, you can also use a Tracker modifier to add a tracker directly to a parameter. To apply the Tracker modifier, in the viewer right-click the Center control of of any transform, text, mask, or other positionable element. From the contextual menu, choose Object x Center > Modify With > Tracker.

Then choose one of three options:

- **Tracker Position:** Tracks a point from the source.
- **Steady Position:** Stabilizes based on a single point in the source.
- **Unsteady Position:** Adds original motion back after stabilizing.

This adds a modifier in the Inspector with a set of controls almost identical to those found in the Tracker node itself.

Inspector



The Tracker modifier

For an in-depth explanation of this node, *see Chapter 59, "Tracking Nodes,"* in the Fusion Reference Manual.

Tracker Modifier vs. Tracker Node

The differences between a Tracker modifier and a Tracker node are as follows:

- The Tracker modifier can only track a single pattern.
- The Tracker modifier can only output a single value and cannot be used for complex stabilization or match-moving procedures.
- The default source image for the modifier is the node immediately upstream of the node that contains the modifier (i.e., when adding a Tracker modifier to a Glow node with a Loader as its input, the Tracker Source input defaults to the output of the Loader). Set a different source image for the Tracker modifier by typing in the name of the node. Alternatively, drag the source node from the Node Editor into the Text Box control.

Example

Imagine that you need to track an actor's eyes so that an unearthly, alien glow can be applied to the eyes.

- Select the Loader node.
- Add a Glow node.
- Add an Ellipse mask to the Glow in the shape of one of the eyes.
- Right-click the center of that mask and select Modify With > Tracker > Position.

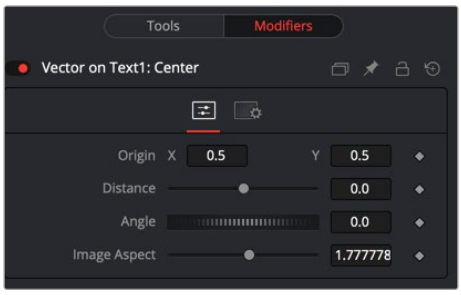
  Since the track is on the mask, the tracker takes the glow as the image for tracking. This could cause problems since the eye might be very obscured by the glow. A cleaner source will be the Loader that feeds the glow.
- Drag the Loader into the modifier Inspector's Track Source field.
- Track the actor's eye.
- Perform the same steps on the other eye.

Vector Result

The Vector Result modifier is used to offset positional controls, such as crosshairs, by distance and angle. These can be static or animated values.

It can be applied by right-clicking a control and selecting Modify With > Vector.

Inspector



The Vector Result modifier Controls tab

Controls Tab

Origin

This control is used to represent the position from which the vector's distance and angle values originate.

Distance

This slider control is used to determine the distance of the vector from the origin.

Angle

This thumbwheel control is used to determine the angle of the vector relative to the origin.

Image Aspect

This slider control is used primarily to compensate for image aspect differences. A square image of 500 x 500 would use an Image Aspect value of 1, while a rectangular image of 500 x 1000 would use an Image Aspect value of 2. The default for this value is taken from the current Frame Format preferences using width/height. It may be necessary to modify this control to match the current image.

Example

- Create a 100-frame comp.
- Create a simple node tree consisting of a black background and a Text node foreground connected to a Merge.
- Enter some text in the Text node.
- Select the Merge node.
- In the viewer, right-click the Center control of the Merge and choose Modify With > Vector Result.

  This adds a crosshair onscreen control for the Vector distance and angle. The onscreen control represents the Distance and Angle controls displayed in the Modifiers tab.
- In the Modifiers tab of the Inspector, drag the Distance control to distance the text from the Vector origin.
- Drag the Angle thumbwheel to rotate the text around the Vector origin.
- This is different from changing a pivot point, since the text itself is not rotating.
- These points are animatable and can be connected to other controls.
- In the Inspector, right-click the Origin control and choose a path to add a motion path modifier to the Origin control.
- Verify that the current frame is set to frame 0 (zero) and use the Origin controls in the Inspector or drag the Vector Origin crosshair to the bottom-left corner of the screen.
- On the Vector Angle thumbwheel, click the Keyframe button to animate this control.
- Set the Angle thumbwheel to a value of 10.
- Go to frame 100 and click at the top-left corner of the screen to move the Vector Origin crosshair.
- Set the Vector Angle thumbwheel to a value of 1000.
- Play the comp to see the results.

This causes the text to orbit around the path just created.

XY Path

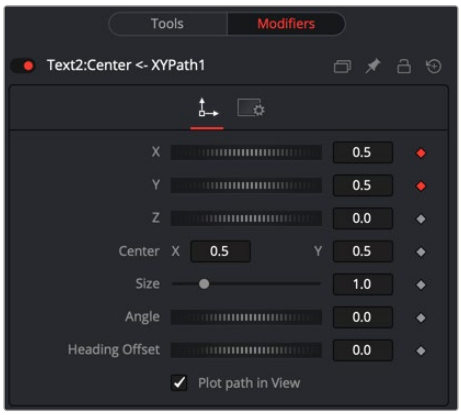
The XY Path type uses two separate splines for the position along the X-axis and for the position along the Y-axis.

To animate a coordinate control using an XY path, right-click the control and select Modify With > XY Path from the contextual menu.

At first glance, XY paths work like Displacement paths. To describe the path, change frames and position the control where it should be on that frame, and then change frames again and move the control to its new position. Fusion automatically interpolates between the points. The difference is that no keyframes are created on the onscreen path.

Look in the Spline Editor to find the X and Y channel splines. Changes to the control's values are keyframed on these splines. The advantage to the XY path is that it becomes very easy to work with motion along an individual axis.

Inspector



XY Path modifier controls

X Y Z Values

These reflect the position of the animated control using X, Y, and Z values.

Center

The actual center of the path. This can be modified and animated as well to move the entire path around.

Size

The size of the path. Again, this allows for later modification of the animation.

Angle

The angle of the path. Again, this allows for later modification of the animation.

Heading Offset

If another control (for example, a mask's Angle) is connected to the path's heading, this control allows for adding or subtracting from the calculated angle.

Plot Path in View

Toggles whether or not the actual path is displayed in the views.

Switching Default Paths

You can change the Default Path type to XY path (if this is the preferred type of animation). Open the Default category in the Global Preferences and locate the Point With drop-down menu. Change this from the current value to XY Path. The next time Animate is selected from a Coordinate control's contextual menu, an XY path is used instead of a Displacement path.

Menu Description

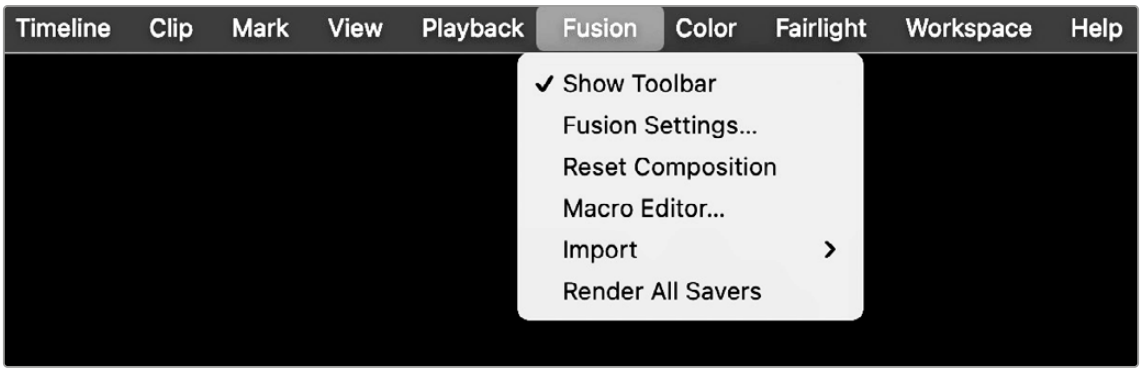
CONTENTS

Fusion 1787

Fusion 20

Menu Description

For ease of use navigating this manual, each menu item is listed here, and by clicking on the name of the menu function, you will be taken to the appropriate part of the manual that describes that function.



Fusion

Show Toolbar – Page 34
Toggles the Fusion toolbar on or off.

Fusion Settings – Page 57
Opens the Fusion Settings window.

Reset Composition
Resets a Fusion composition to its initial state.

Macro Editor – Page 166
Opens the Macro Editor.

Import
Specific file format import for Fusion.

- **Alembic Scene** – Page 702
- **FBX Scene** – Page 647
- **PSD** – Page 85
- **Shapes** – Page 81
- **SVG** – Page 81
- **Tracks** – Page 81

Render All Savers – Page 93
Renders OpenEXR files.



blackmagicdesign.com