



Reference Manual

Fusion 16

July 2019



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 DaVinci Resolve Engineering Team

A stylized, handwritten signature in black ink that reads "Grant Petty".

Grant Petty

CEO Blackmagic Design

Navigation Guide

Contents	
The Fusion User Interface	18
The Work Area	19
Interface Toolbar	19
Showing Which Panel Has Focus	20
Viewers	20
Zooming and Panning into Viewers	22
Loading Nodes Into Viewers	22
Clearing Viewers	23
Viewer Controls	23
Time Ruler and Transport Controls	25
The Playhead	25
The Current Time Field	25
Frame Ranges	25
Changing the Time Display Format	26
Zoom and Scroll Bar	27
Transport Controls	27
Keyframe Display in the Time Ruler	28
Fusion Viewer Quality and Proxy Options	29
The Fusion RAM Cache for Playback	30
Toolbar	31
Node Editor	32
Adding Nodes to Your Composition	33
Removing Nodes from Your Composition	33
Identifying Node Inputs and Node Outputs	34
Node Editing Essentials	34
Navigating the Node Editor	35
Keeping Organized	36
Status Bar	36
Effects Library	36
Inspector	37
The Tools and Modifiers Panels	38
Parameter Header Controls	38
Parameter Tabs	39

Chapter – 2 Introducing the Fusion Interface 16

Hover over and click the Contents Heading and it opens the Main Content Page at the beginning of the manual

Hover over and click each title and it opens the Page accordingly

Hover over and click the footer on each page and it opens the Chapter Content Page

PART 1

Getting Started

1	Introduction to Fusion	8
2	Introducing the Fusion Interface	15
3	Color Management	53
4	Understanding Image Channels and Node Processing	59
5	Preferences	85

PART 2

The Fusion UI

6	Working in the Node Editor	125
7	Node Groups, Macros, and Fusion Templates	168
8	Using Viewers	179
9	Editing Parameters in the Inspector	220
10	Bins	245

PART 3

Importing and Rendering

11	Preparing Projects and Importing Media	264
12	Rendering and Licensing	276
13	Fusion Connect	298

PART 4

2D Compositing

14	Learning to Composite in Fusion	314
15	Using the Tracker Node	375
16	Planar Tracking	397
17	Rotoscoping with Masks	404
18	Optical Flow and Stereoscopic Nodes	427
19	Using OFX Plug-Ins	440

PART 5

Animation

20	Keyframing in Fusion	444
21	Animating with Motion Paths	461
22	Animating with Modifiers and Expressions	474

PART 6

3D Compositing

23	3D Compositing Basics	482
24	3D Camera Tracking	538
25	Particle Systems	548

PART 7

Fusion Effects

26	3D Nodes	559
27	3D Light Nodes	646
28	3D Material Nodes	656
29	3D Texture Nodes	680
30	Blur Nodes	699
31	Color Nodes	717
32	Composite Nodes	759
33	Generator Nodes	772
34	DeepPixel Nodes	807
35	Effect Nodes	819
36	Film Nodes	836
37	Filter Nodes	848
38	Flow Nodes	858
39	Flow Organizational Nodes	862
40	Fuses	868
41	I/O Nodes	871

42	LUT Nodes	889
43	Mask Nodes	895
44	Matte Nodes	917
45	Metadata Nodes	944
46	Miscellaneous Nodes	949
47	Optical Flow	973
48	Paint Node	983
49	Particle Nodes	1012
50	Position Nodes	1064
51	Stereo Nodes	1078
52	Tracker Nodes	1100
53	Transform Nodes	1137
54	Warp Nodes	1156
55	Modifiers	1175
56	VR Nodes	1208



PART 1

Getting Started

Chapter 1

Introduction to Fusion

This introduction has been designed specifically to help users who are new to Fusion get started learning this exceptionally powerful environment for creating visual effects and motion graphics.

Fusion is a resolution-independent visual effects compositing and motion graphics application, used worldwide for high-resolution visual effects.

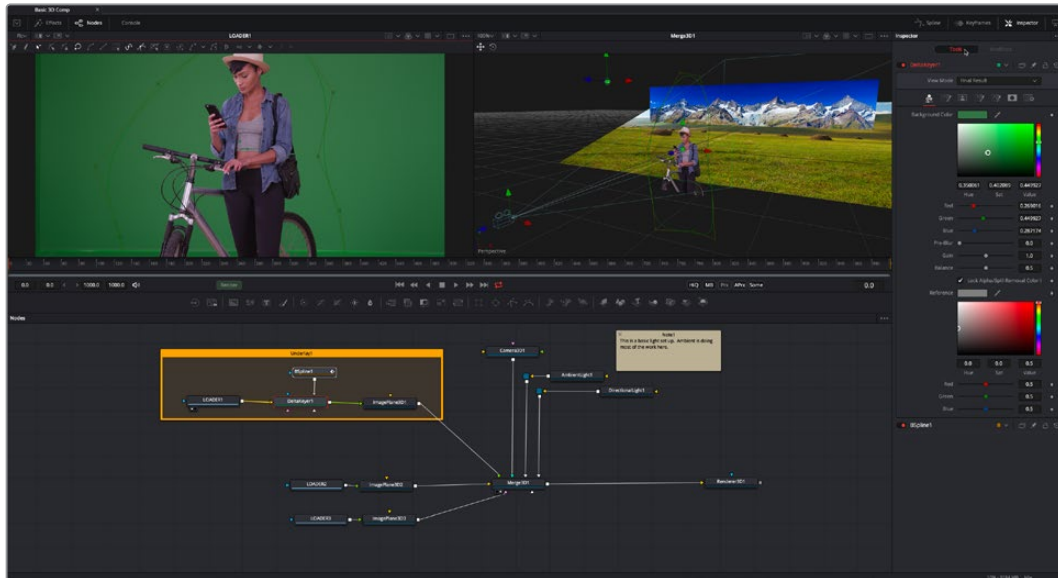
It includes hundreds of features aimed at giving you the highest quality photo-realistic composites for film and video whether you work as an individual artist or within a large studio.

Contents

Introduction to Fusion	10
How Does Fusion Work?	10
Importing Media	10
Your First Node Tree	11
Rendering Out Your Final Result	11
What Kinds of Effects Does Fusion Offer?	12
3D Compositing	12
Particles	13
Text	13
And Lots More	13
How Hard Is This Going to Be to Learn?	14

Introduction to Fusion

Fusion is a powerful 2D and 3D compositing application with over 30 years of evolution serving the film and broadcast industry, creating effects that have been seen in countless films and television series.



Fusion 16, showing viewers, the Node Editor, and the Inspector.

How Does Fusion Work?

In its simplest form, Fusion is a collection of image-processing engines called nodes. These nodes represent filters like blurs and color correctors as well as images, 3D models, and spline shapes. Similar to filter effects you may be familiar with, each node includes a set of parameters that can be adjusted and animated over time. Stringing nodes together in a graphical user interface called a node tree, allows you to create sophisticated visual effects.

To begin, you need only open Fusion, create a new project, import some clips via Loader nodes, and away you go. 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, you add a Saver node to the end of the tree of nodes you've created to render your final result.

Importing Media

The clips you want to work with can be brought into Fusion in two ways:

- Click the Loader button in the toolbar and use the dialog to locate the clips.
- Drag files from the operating system's file browser to the Node Editor in the Fusion interface.



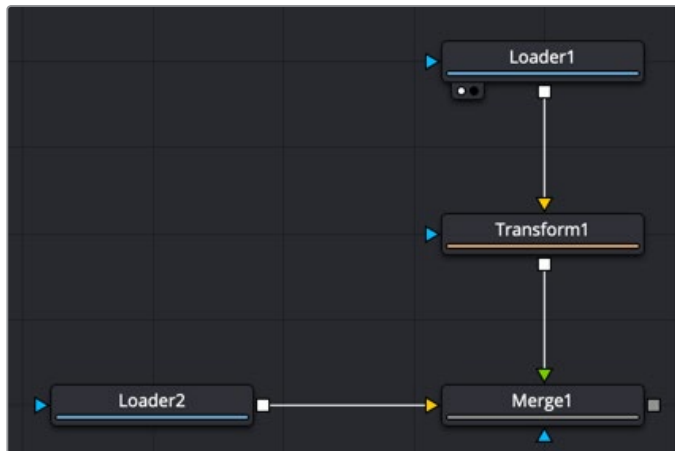
The Loader node can be added from the toolbar to ingest new media.

No matter which way you choose to import your clips, the result is that a Loader node gets added to the Node Editor. A Loader is the node that represents any clip, image file, or graphic that you bring into Fusion.

Your First Node Tree

In the example below, a very simple Fusion composition includes two Loader nodes that are used to link to two different clips on your hard drive.

Similar to connecting your cable TV box to your television with an HDMI cable, the output of Loader1 is connected to the input of the Transform node. The Transform node is used to resize, rotate, and reposition a clip. A Merge node composites the transformed foreground image (Loader 1) over the background image (Loader 2). If you look at the connection lines, they include arrows indicating the direction the image is traveling.



The Loader node can be added from the toolbar to ingest new media.

Selecting a node will display its parameters in the Inspector. For instance, selecting the Transform node will display scaling, positioning, and rotation parameters in the Inspector and in the viewer.

However, viewing the results of a node requires you to select the viewer where you want the results displayed. To see any node in a viewer, you select the node and press 1 to see it in the viewer on the left, or press 2 to see it in the viewer on the right.

Rendering Out Your Final Result

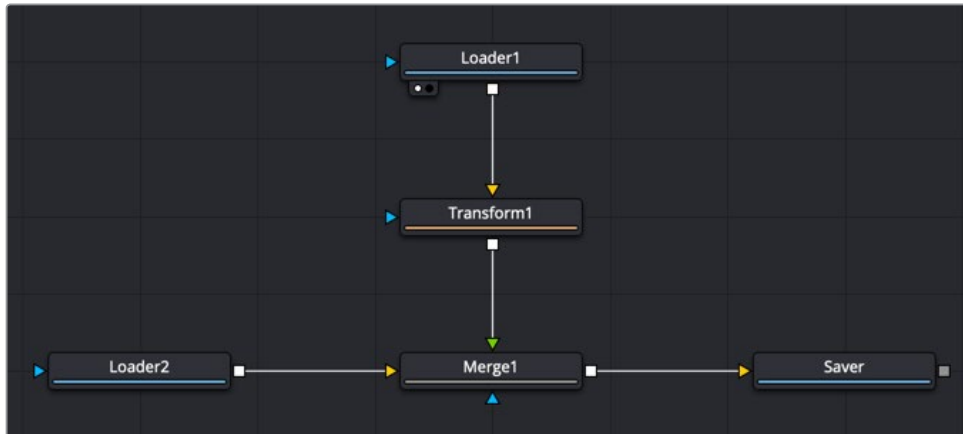
The final step when you have completed a project is to render the finished effect to disk as a movie file or image sequence. The last node in every node tree will be a Saver node. Saver nodes are used to set your output file format and render the file to disk. The Saver node gets appended to the end of your node tree. In the example of our node tree above, the Saver node would be added directly after the Merge node, since that is the final node where both images get combined.



The Saver node is added by clicking the tool in the toolbar.

To render an effect to disk, do the following:

- 1 Select the last node in your node tree.
- 2 From the toolbar, click the Saver tool.
- 3 In the dialog that appears, navigate to the location where you want the file saved.
- 4 Enter a name for your movie and click Save.
- 5 In the Inspector, choose the Output format from the menu.
- 6 In the toolbar, click the Render button or choose Render All Savers from the File menu.




The Saver node is added to the end of a node tree to render your effect to disk.

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 provides powerful 3D nodes that include modeling text and simple geometry right inside Fusion. It includes the ability to import 3D models in a variety of formats. Once you've assembled a 3D scene, you can add cameras, lighting, and 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 provides 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.



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

And Lots More

The list of effects goes on, with Stereo and VR adjustment nodes, Planar Tracking, Deep Pixel nodes for re-lighting 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 Is This Going to 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 preview was developed specifically to help users who've never worked with Fusion before learn the core concepts needed to do the basics, in preparation for learning the rest of the application on your own.

This version of Fusion is another evolution of a deep, production-driven product that's experienced 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 have experience 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.

Go on, give it a try, and remember that you have this PDF to refer to, which includes Chapter 14, "Learning to Composite in Fusion" that walks you through a broad selection of basics, showing common techniques that you can experiment with using your own footage.

Chapter 2

Introducing the Fusion Interface

This chapter provides an orientation of the user interface of Fusion, 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	18
The Work Area	19
Interface Toolbar	19
Showing Which Panel Has Focus	20
Viewers	20
Zooming and Panning into Viewers	22
Loading Nodes Into Viewers	22
Clearing Viewers	23
Viewer Controls	23
Time Ruler and Transport Controls	25
The Playhead	25
The Current Time Field	25
Frame Ranges	25
Changing the Time Display Format	26
Zoom and Scroll Bar	27
Transport Controls	27
Keyframe Display in the Time Ruler	28
Fusion Viewer Quality and Proxy Options	29
The Fusion RAM Cache for Playback	30
Toolbar	31
Node Editor	32
Adding Nodes to Your Composition	33
Removing Nodes from Your Composition	33
Identifying Node Inputs and Node Outputs	34
Node Editing Essentials	34
Navigating the Node Editor	35
Keeping Organized	36
Status Bar	36
Effects Library	36
Inspector	37
The Tools and Modifiers Panels	38
Parameter Header Controls	38
Parameter Tabs	39

Keyframes Editor	39
Keyframe Editor Control Summary	39
Adjusting Clip Timings	40
Adjusting Effect Timings	40
Adjusting Keyframe Timings	40
Spline Editor	42
Spline Editor Control Summary	42
Choosing Which Parameters to Show	43
Essential Spline Editing	43
Essential Spline Editing Tools and Modes	43
The Console	45
Customizing Fusion	46
Showing and Hiding Panels	46
Resizing Panels	46
Keyboard Remapping	47
Undo and Redo	49
Bins	50
Bins Interface	50
The Bin Studio Player	52

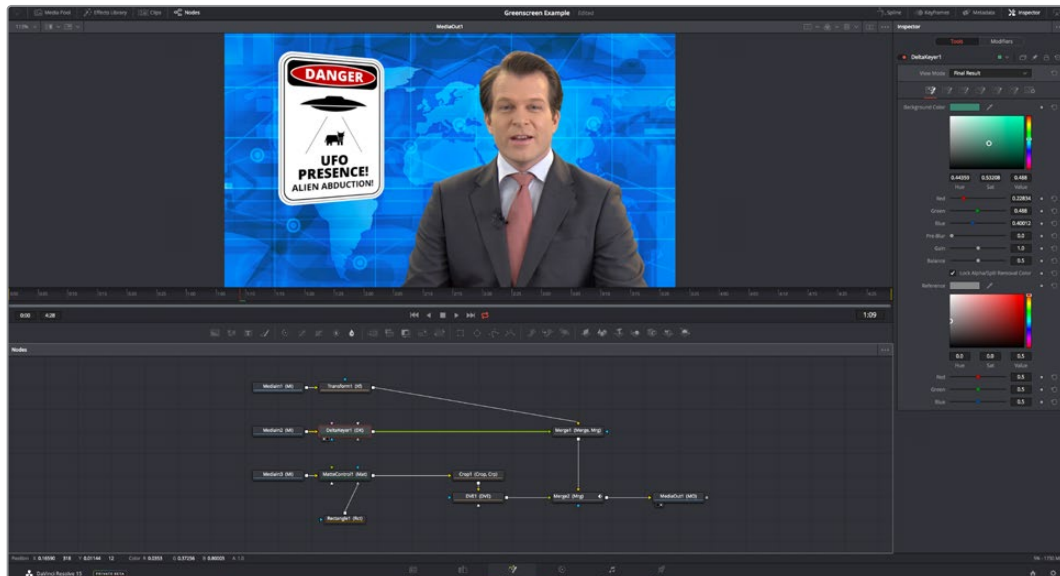
The Fusion User Interface

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 Effects Library is found at the left, the viewers are at the top, the work area is at the bottom, and the Inspector is at the right. 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 in Fusion.



The 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

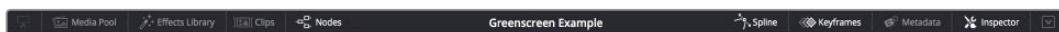
You probably won't see this term used much, in favor of the specific panels within the work area that you'll be using, but the area referred to as the work area is the region at the bottom half of the Fusion user interface, within which you can expose the three main panels used to construct compositions 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 in, but it can sit side-by-side with the Spline Editor and Keyframes Editor as necessary, and you can make more horizontal room on your display for these three panels by putting the Effects Library and Inspector into 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 Fusion user interface (UI). Buttons with labels identify which parts of the UI can be shown or hidden. If you right-click anywhere within this toolbar, you have the option of displaying this toolbar with or without text labels.



The UI toolbar of Fusion.

These buttons are as follows, from left to right:


- **Effects Library Height button:** Lets you set the area used by the 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 Effects Library is restricted to the top half of the UI along with the viewers, and the Node Editor takes up the full width of your display.
- **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.
- **Nodes:** Opens and closes the Node Editor, where you build and edit your compositions.
- **Console:** Opening the Console provides a place to get log and error messages, read FusionScript outputs, or input FusionScripts directly. The Console tab will display an icon when new content is produced. The icon will identify it as an error, log, or script output.
- **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 Keyframe Editor, which shows each clip and effects node in your Fusion composition as a layer. You can use the Keyframe Editor to edit and adjust the timing of keyframes that have been added to various effects in your composition. You can also use the Keyframe 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.
- **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 button:** 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.

Showing 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 will capture specific keyboard shortcuts to do something within that panel, as opposed to doing something elsewhere in the interface.

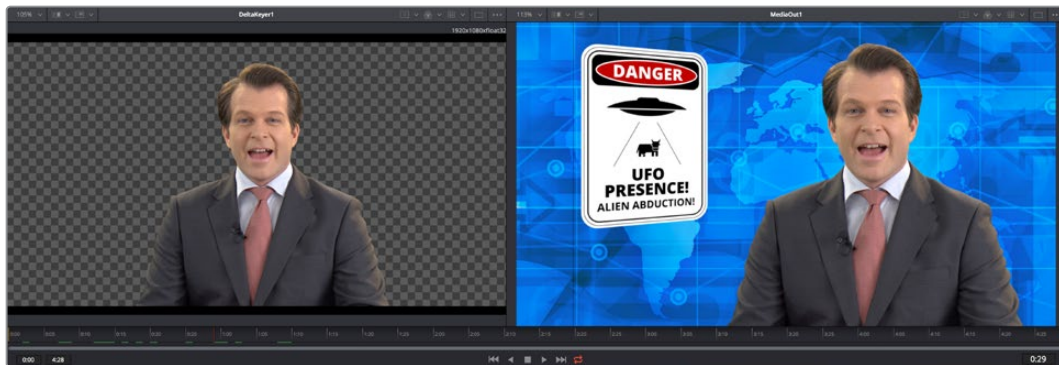
To make it easier to keep track of which panel has focus, a highlight appears at the top edge of whichever panel has focus, so you can keep track of which part of the interface is taking precedence, and you can switch focus as necessary to do what you need to do.



The optional focus indicator shown at the top edge of the Effects Library, shown next to a viewer that doesn't have focus.


Viewers

The viewer area can be set to display either one or two viewers at the top of Fusion, and this is set 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 viewer while seeing its effect on the overall composition in the other viewer.

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, including a perspective view that gives you a repositionable stage on which to arrange the elements of the world you're creating, or a quad-view that 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 Option key and 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 what you're looking at and include many more options specific to the detail-oriented work compositing entails. This section gives 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.

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 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 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-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.

Methods of spinning 3D viewers:

In 3D Perspective view, hold down the Option key and drag to spin the stage around.

Loading Nodes Into Viewers

When you first open Fusion, the two viewers are empty until you decide to view the output of a node in 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/LeftView/RightView 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 Tilde (~) 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.

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 dragging 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 on-screen control, or rotated by dragging anywhere on the dividing line on the on-screen control. Alternately, 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 way 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 turn on the quad-paned 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 limiting in the viewer, while using the menu lets you choose the region of the Roi. The Region of Interest (Roi) lets you define the region of the viewer in which which pixels actually need to be rendered. 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 on-screen 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 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, or B). 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 display on or off, while the menu lets you choose which of the many available color space conversions to apply. By default, viewers in Fusion show you the image prior to any color management. However, while you typically convert to linear color space for compositing, but you will find it desirable to make adjustments while viewing a normalized version of the image so it appears close to what the final will be. Enabling the LUT display lets you do this as a preview, without permanently applying this color adjustment to the image. The top five options let you choose Fusion controls, which can be customized via the Edit item at the bottom of this menu.
- **Option menu:** This menu contains various settings that pertain to the viewer in Fusion.
 - **Checker Underlay:** Toggles a checkerboard underlay that makes it easy to see areas of transparency.
 - **Show Controls:** Toggles whatever onscreen controls are visible for the currently selected node.
 - **Pixel Grid:** Toggles a preview grid that shows, when zoomed in, the actual size of pixels in the image.
 - **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:** Depending on the type of footage loaded, images may have pixels that are rectangular instead of square. 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.

Time Ruler and Transport Controls

The Time Ruler, located beneath the viewer area, shows the frame range of the current clip, the playback or render range, an Audio Monitoring button, a Render button for rendering the final project, transport controls, as well as Proxy and Global Motion Blur settings:

- Although the Time Ruler shows all the source frames for a clip, the yellow lines mark the In and Out points that will be rendered and played back. This is called the “render range.”



The Time Ruler displaying the render range for a clip via yellow marks (the playhead is red).

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.

The Current Time Field


The Current Time field at the far right of the transport controls shows the frame number for the position of the playhead, which corresponds to the frame seen in the viewer. However, you can also enter a frame number into this field to move the playhead to a specific frame.

When setting ranges and entering frame numbers to move to a specific frame, numbers can be entered in subframe 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 to set keyframes where they actually need to occur, rather than on a frame boundary, so you get more natural animation. Having subframe 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, the Current Time field will not correctly evaluate mathematical expressions that use the $+$ symbol, even when the display format is set to Frames mode.

Frame Ranges

As explained above, the Time Ruler displays two different frame ranges, one for the entire project, called the Global duration, and the other called the render range, which determines what to render or what to cache in memory for previews. The Global start and end range sets the total duration of a project. You cannot move the current time outside the global range. Coming from a color grading and editing application like DaVinci Resolve, the Global Start and End range is the equivalent of the Timeline duration



Composition frame range and render range fields.

Composition Start and End Range

The Composition Start and End range is simply the total duration of the current composition. You can modify these values using the outer numeric fields in the Time Ruler.

Render Range

The Render Start and End range determines the range of frames that will be used for interactive playback, disk caches, and previews. The range is normally visible in the time slider as a light gray highlighted region within the Time Ruler. Frames outside the render range will not be rendered or played, although you can still drag the playhead or set the current time to these frames to see what the image looks like.

The inner two fields at the far left of the transport controls show the first frame and last frame of this range. 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.
- 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.

Changing the Time Display Format

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

To change the time display format:

- 1 Choose Fusion Studio > Preferences.
- 2 When the Fusion Preferences dialog opens, select the Global and Default settings > Defaults panel and choose a Timecode option.
- 3 Open 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.


Zoom and Scroll Bar

A two-handled 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

There are eight transport controls underneath the Time Ruler, including Composition First Frame, One Frame Backwards, Play Reverse, Stop, Play Forward, One Frame Forwards, Composition Last Frame, and Loop.



Fusion transport controls.

Navigation Shortcuts

Many of the standard transport control keyboard shortcuts found in DaVinci Resolve also 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 Global Comp's End frame.
- **Shift-Forward Arrow:** Moves to the Global Comp's 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 Fusion 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 (see "The Fusion RAM Cache for Playback" later in this chapter).

Frame Increment Options

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

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.
- **Pingpong 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.

Keyframe Display in the Time Ruler

When you select a node that's been animated with keyframed parameters, those keyframes appear in the Time Ruler as little white tic marks, letting you navigate among and edit keyframes without being required 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.

Fusion Viewer Quality and Proxy Options

The right side of the Time Ruler also includes a number of buttons that affect quality by either enabling a high-quality playback at the expense of greater processing times, or entering various proxy modes that temporarily lower the display quality of your composition in order to speed processing as you work. This area of the toolbar also includes a global motion blur setting.



The Time Ruler buttons for controlling quality and motion blur.

High Quality (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.

Motion Blur (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 affect.

Proxy (Prx)

Turning on Proxy reduces the resolution of the images 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 Preferences, found in the Fusion Studio menu.

Auto Proxy (APrx)

Turning on Auto Proxy reduces the resolution of the image while you click and drag on a parameter's control to make an adjustment. 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 Preferences, found in the Fusion Studio menu.

The view quality and proxy settings can also be accessed by right-clicking in any area to the left or right of the transport controls.

Selective Updates Toggle

The Selective Updates button is a three way toggle providing three options for processing nodes as you work on a composite:

- **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.
- **Some:** (the default) Causes only nodes that directly contribute to the current image to be rendered, so named because only selective nodes are rendered.
- **None:** Prevents rendering altogether, which can be handy for making a lot of changes to a slow-to-render composition. While set to None, the Node Editor, Keyframes Editor and Spline Editor will be highlighted with a red border to indicate that the tools are not being updated.


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 Fusion during playback. Of course, when you play beyond the cached area of the Time Ruler, uncached frames will 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


When the size of the cache reaches the Caching Memory Limit 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 how much of the RAM cache has been used via a percentage of use indicator at the far right of the Status bar at the bottom of Fusion.



Percentage of the RAM cache that's been used at the bottom right of Fusion.

Displaying Cached Frames


All frames cached for the currently viewed range of nodes 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 line indicates 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 will be 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 will be 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 has buttons for adding commonly used nodes to the Node Editor.

The toolbar is divided into seven 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:** The Loader and Saver tools are probably the two most important nodes in Fusion. The Loader tool is used to add images to your node tree and the Saver node is used to render them out.
- **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.

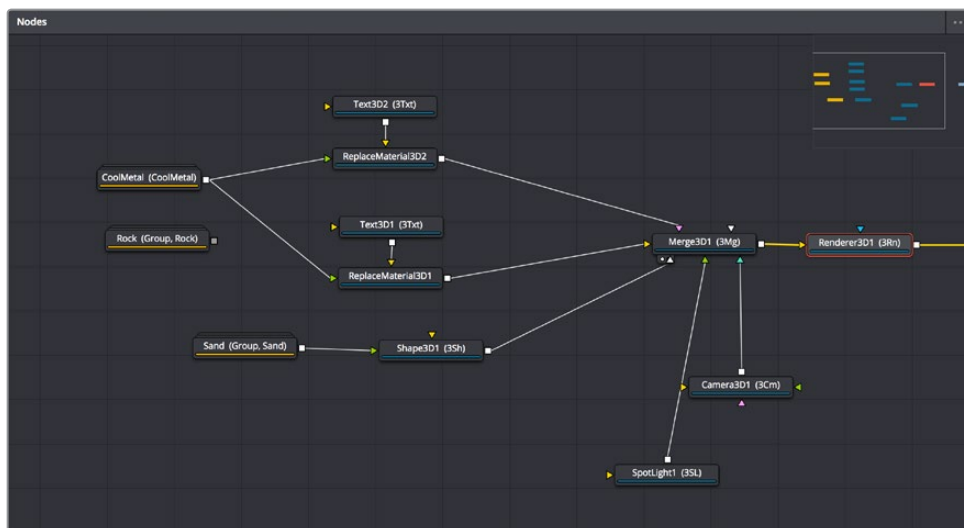
- **Particle system nodes:** Three particle nodes let you create complete particle systems when you click them from left to right. pEmitter emits particles in 3D space, while pMerge lets you merge multiple emitters and particle effects to create more complex systems. pRender renders a 2D result that 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 SpotLight lets you light the scenes in different ways, and Camera3D lets you frame the scene in whatever ways you like. 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.

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.

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.




The Node Editor displaying a node tree creating a composition.

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 to the node tree after it, and connected to the selected node. However, if there are no selected nodes, then new nodes are added as disconnected from anything else. They will be placed in the Node Editor wherever you click your mouse.

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 contextual 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

If you hover the pointer over any of a node's inputs or outputs, the name of that input or output will immediately appear in the Status bar, and if you wait for a few more moments, a floating tooltip will display the same name right over the node you're working on.

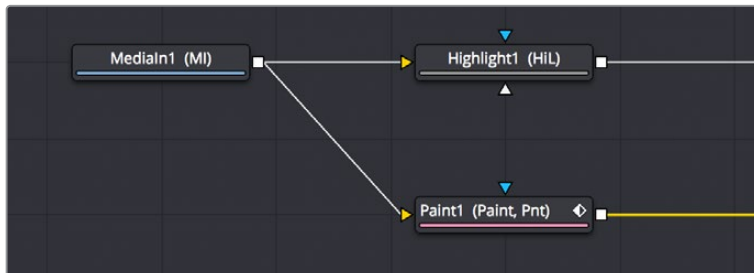
Node Editing Essentials

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”).



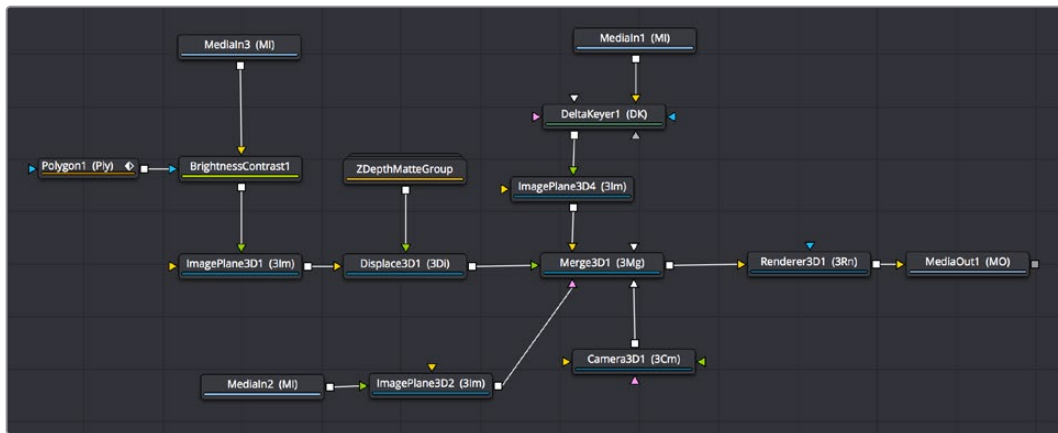
One node branching to two to split the image into two operations.

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.

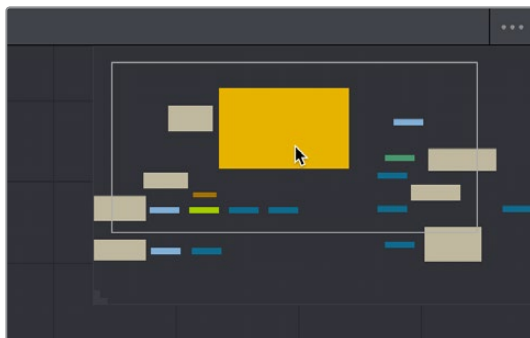
By default, new nodes are added from left to right in the Node Editor, but they can also flow from top to bottom, right to left, bottom to top, or in all directions simultaneously. Connections automatically reorient themselves along all four sides of each node to maintain the cleanest possible presentation 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 composition gets larger, parts of it will inevitably go off screen. By default, when a portion of the node tree has gone off-screen, a resizable Navigator pane appears at the upper right corner, which can be used to see 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 at the lower left corner, and you can choose to show or hide the navigator 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 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.

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, 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 at the bottom of Fusion, in the lower left corner of the interface, 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 any node displays information about that node in the Status bar (as well as in a floating tooltip), while 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.

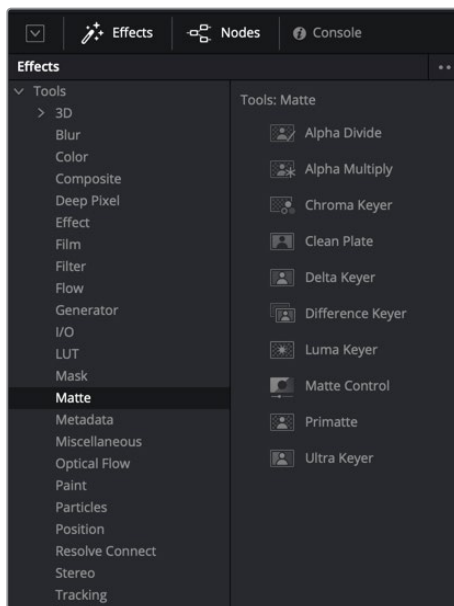
Occasionally, the user interface toolbar will display a badge to let you know there's a message in the console you might be interested in viewing. These messages can range from a simple notification that rendering is complete, to a missing third-party plug-in, to an error trying to read in an image.



A notification that there's a message in the Console.

Effects Library

The Effects Library shows all the nodes and effects available in Fusion, including effects that come with Fusion and third-party OFX, if available. 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 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 hierarchical category browser of the Effects Library is divided into Tool sections containing every node that represents an elemental image processing operation in Fusion.

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

Inspector

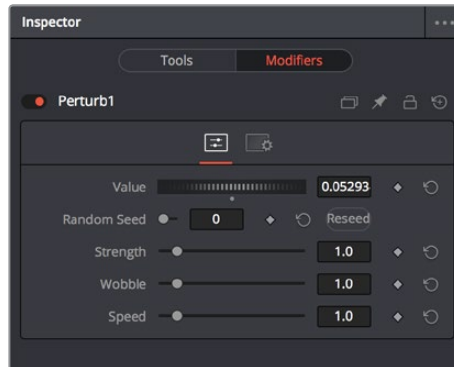
The Inspector is a panel on the right side of Fusion that you use to display and manipulate the parameters of one or more selected nodes. When a node is selected in the Node Editor, its parameters and settings appear in the Inspector.



The Inspector shows parameters from one or more selected nodes.

The 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.



The Modifier panel showing a Perturb modifier.

Other nodes display more 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 parameters for multiple nodes. 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

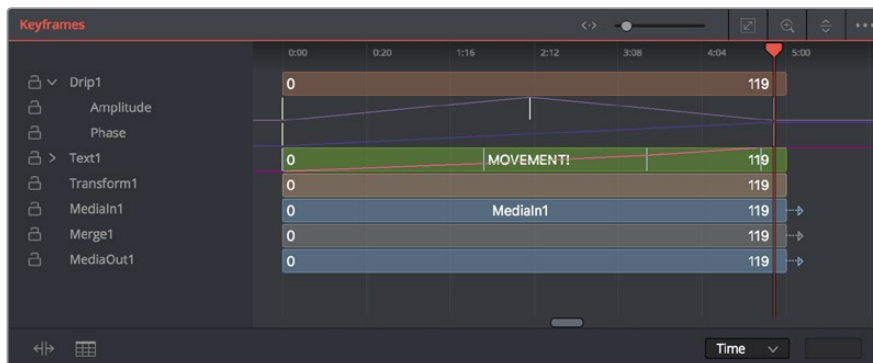
Many nodes expose multiple tabs' worth of controls in the Inspector, seen as icons at the top of the parameter section for each node. Click any tab to expose that set of controls.



Nodes with several tabs' worth of parameters.

Keyframes Editor

The Keyframes Editor displays each Loader (clip) and effects 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 dictates the order of image processing operations. You use the Keyframes Editor to trim, extend, or slide Loader (clip) and effects nodes, or to adjust the timing of keyframes, which appear superimposed over each effect node unless you open them up into their own editable track.



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

Keyframe Editor Control Summary

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

- 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 Keyframe 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 node, as well as controls governing that node.

- A lock button lets you prevent a particular node 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 nodes 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/ T Offset/ T Scale pop-up menu and value fields let you numerically alter the position of selected keyframes either absolutely, relatively, or to a scale.

Adjusting Clip Timings

Each Loader node that represents a clip used in a composition is represented as a segment on a track in this miniature timeline. You can edit a Loader'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 segment by dragging it to the left or right, to better line up with the timing of other elements in your composition.

The Keyframes Editor lets you adjust the timing of generators, filter effects, and 3D objects as well.

Adjusting Effect Timings

Each Effect node also appears as a segment, just like clips. You can resize the In and Out points of an effect segment, and slide the entire segment forward or backward in time, just like Loaders. If you trim an effect to be shorter than the duration of the composition, the effect will cut in at whichever frame the segment begins, and cut out at after the last frame of that segment, 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 in a simple way. By default, all keyframes applied to parameters within a particular node's segment appear superimposed in one flat track over the top of that segment.

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



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 contextual menu commands 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 frame offset.

Time Stretching Keyframes

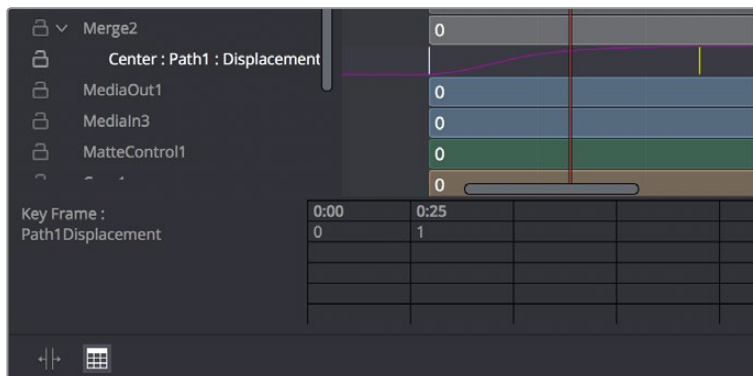
If you select a range of keyframes in a keyframe track, you can turn on the Time Stretch tool to show a box you can use to squeeze and stretch the entire range of keyframes relative to one another, to change the overall timing of a sequence of keyframes without losing the relative timing from one keyframe to the next. Alternately, 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 a 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.

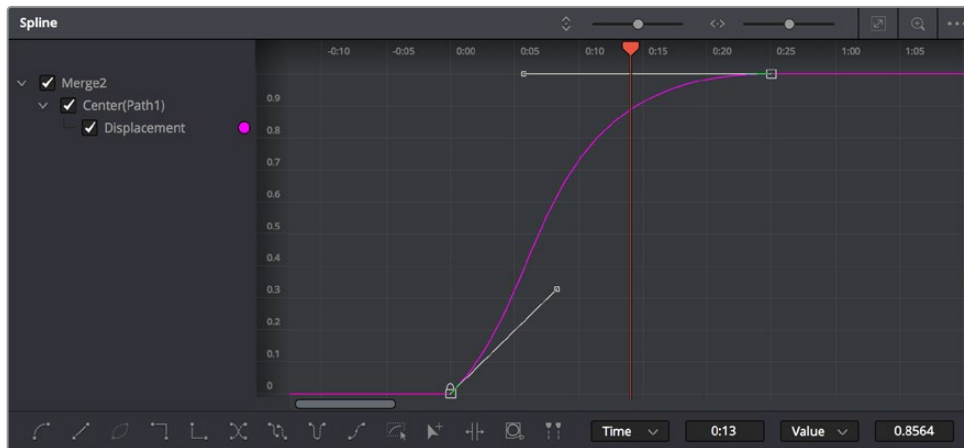


Editing keyframes in the Spreadsheet.

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.

Spline Editor

The Spline Editor provides a more detailed environment for editing the timing, value, and interpolation of keyframes that create different animated effects, using control points at each keyframe connected by splines (also called curves) that let you adjust how animated values change over time. The Spline Editor has four main areas: the Zoom and Framing controls at 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, the 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 layers 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 Sort pop-up menu lets you sort or filter the parameters in various ways.
- An Option menu provides access to many other ways of filtering parameters and controlling visible options.
- 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. By default, the horizontal axis represents time and the vertical axis represents the spline's value, although you can change this via the Horizontal and Vertical Axis pop-up menus at the bottom-right of the Spline Editor, and selected control points show their values in the accompanying edit fields.

Lastly, the toolbar at the bottom of the Spline Editor has 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 Bezier curves:

- Select any control point to make its Bezier handles visible, and drag the Bezier handles.
- Command-drag a Bezier 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 five buttons let you adjust the interpolation of one or more selected control points.



Control point interpolation controls.

- **Smooth:** Creates automatically adjusted Bezier curves to create smoothly interpolating animation.
- **Flat:** Creates linear interpolation between control points.
- **Invert:** Inverts the vertical position of selected keyframes relative to one another.
- **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 backwards.

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

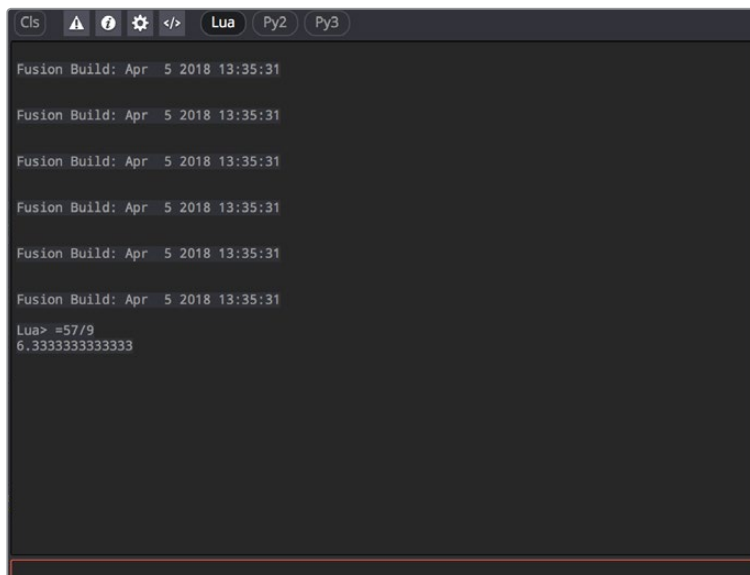
to the segment at the location where you clicked. With Step In or Step Out keyframes, clicking anywhere above or below a line segment moves that segment to where you've clicked.

- **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. Alternately, 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), cornerpin (using the corner handles), move (dragging on the box boundary), or corner stretch (Command-drag the corner handles),
- **Show Key Markers:** Turning this control on 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.

The Console

The Console, available by choosing View > Console or by clicking the Console button in the User Interface toolbar, 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.

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 badge will indicate if the message is an error, log, or script message.



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 2.7 and Python 3.3 require that you have the appropriate Python environment already installed on your computer. Since 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. There are some useful shortcuts you can perform in the Console. More information on scripting will be forthcoming as it becomes available.

Customizing Fusion

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

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 Fusion.

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.

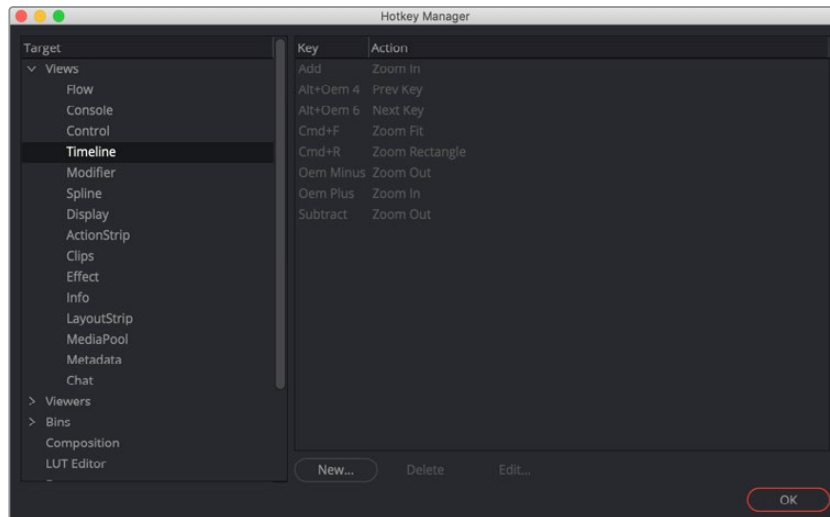


Dragging the edge between two viewers to resize it.

Pressing F4 on the keyboard will expand the interface panel your mouse pointer is over. This is incredibly helpful when you want to see greater detail in one of the two viewers or when two panels like the Node Editor and Spline Editor share the work area, and you want to give more room to one of the panels temporarily. The F4 key is a toggle; pressing it again restores the previous layout.

Keyboard Remapping

You can change and create your own Keyboard hotkeys by choosing Views > Customize Hotkeys.



The Hotkey Manager window

The 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 will display the keyboard shortcut if one exists. You can use the New button at the bottom of the dialog to add a new keyboard shortcut.

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

- 1 Open the Keyboard Hotkey Manager.
- 2 Select Views > Flow 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 Editor.
- 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.

If you are bit more technical, you can also create .fu files to customize the keyboard mapping and menu items as well. These user-created custom settings are loaded from the UserPaths:Config directory located in the Preferences.

Config files use a Menus or Hotkeys tabbed table, containing a single Target string specifying which object or view the menus are for, and a number of Before, After, or Replace tables describing the menu location, and containing a list of Action items that the menu will trigger. You can also include Sub tables to describe submenus.

For example, to create a new submenu called Special under the Help menu that's visible when no comps are open, you could create a Configs:MyMenu.fu and format it using tabs as you see below:

```
{
    Menus -- as opposed to "Hotkeys" etc
    {
        Target = 'Fusion',-- handled by the app
        After 'Help' -- place this new menu after "Help"
        {
            Sub 'Special' -- new menu is called "Special"
            {
                'Comp_New', -- this triggers the 'Comp_New'
action, creating a new comp
            }
        },
    },
}
```

To add & modify some top-level comp menus:

```
{
    Menus
    {
        Target = 'ChildFrame', -- use this for menus when a comp is
open
        Before 'File\\Start Render'-- place this above the Start Render
item on the File menu
        {
            Sub 'Stuff' -- create a submenu called 'Stuff'
            {
                'Execute { Name = "Test Item", cmd = [[
obj:Comp():Print("Test Item")]] }', -- executes 'cmd' arg string
                'RunScript { Name = "Licenses", filename = [[
Scripts:Utility/ Licensing.lua ]] }', -- runs the Licensing.lua script file
            },
            '_', --this adds a separator
        },
        Replace 'Help\\About...' -- this hides and replaces the
About... item on the Help menu
        {
            Execute { Name = "Add About Note", cmd = [[ obj:Comp().
Note{ Comments = "Fusion is the Best"} ]] }', -- this creates a Note tool with
an existing comment
        },
    },
}
```

To change context menus for a specific view, such as the Node Editor view:

```
{
    Menus
    {
        Target = 'FlowView', -- ID of the view
        After 'Add Tool' -- after "Add Tool" on the view's context menu
        {
            'Execute { Name = "Action...", cmd = [[
obj:Comp():ChooseAction(true, obj) ]] }', -- shows the ChooseAction dialog
            'AddTool { Name = "Add Overlay", id = "Overlay" }', just add a
Note tool
            '_', --add separator
        }
    },
}
```

For reference, hotkey config files are similar, but use a Hotkeys table of key identifiers and Actions:

```
{
    Hotkeys
    {
        Target = "GLView",

        V = "Viewer_SubView_Show",
        SHIFT_V = "Viewer_SubView_Swap",
        CONTROL_1 = "Viewer_Reset",
        CONTROL_2 = "Viewer_Reset{ scale = 2.0 }" -- same action but
with an optional scale argument
    }
}
```

Undo and Redo

Undo and Redo commands let you back out of steps you've taken or commands you've executed and then 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.

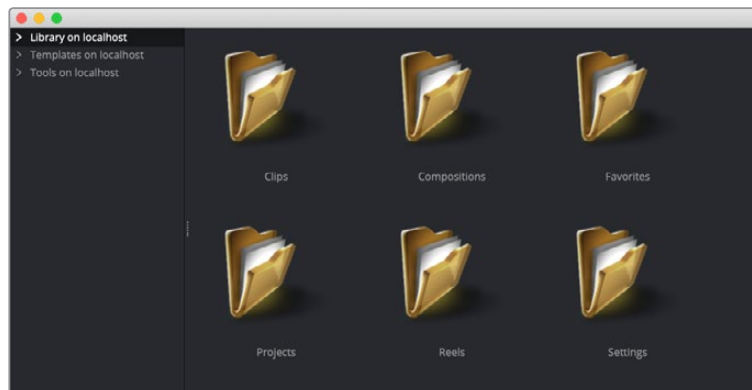
In all cases, there is no practical limit to the number of steps that are undoable (although there may be a limit to what you can remember).

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 (Command-Y) to redo to the next change.

Bins

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.




The Bins window.

To open the Bins window, choose File > Bins from the menu bar.

When adding an item to the bins, a link is created between the item and the bins. Fusion does not copy the file into its own cache or framestore, so there is no need to dedicate large amounts of drive space to Fusion bins.

Bins Interface

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

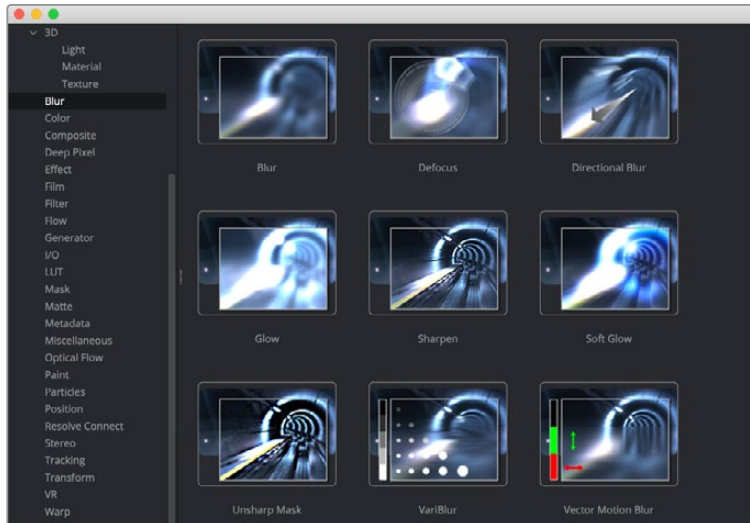


The Bins sidebar.

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, 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 folder from the sidebar, 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 a few common controls.




The Bins toolbar.

- **New Folder:** Creates a new folder in the current window.
- **New Reel:** Creates an empty reel that can contain
- **New Clip:** Opens a dialog to link new media files into a bin.
- **Studio Player:** Opens the Studio player on 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.

A contextual menu can also be 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.

The Bin Studio Player

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



The Bins Studio Player.

For more detail on Bins and the Studio Player, read Chapter 10, “Using Bins.”

Chapter 3

Color Management


This chapter discusses the value of compositing with clips using a linear gamma, and how to deal with color management in Fusion, so you can work with a linear gamma while previewing the image in the viewer using the gamma of your choice.

Contents

Color Management in Fusion	55
Color Management for Standard Dynamic Range Clips	56
Color Management for Wide Dynamic Range Clips	56
Viewer Gamma and Color Space While Working in Linear	57

Color Management in Fusion

The simplified goal of color management is to make sure that what you see as your final product on your computer screen is what your audience sees. For such a simple goal, problems arise that you, as the compositor, must deal with. Today's digital cinema cameras capture images using their own unique nonlinear gamma curves (often called Log gamma), giving more attention to the darker mid tones where the human eye is most sensitive. This allows them to save images with a wider dynamic range.



A non-linear gamma curve.

The problem is that these images do not look good on an sRGB computer monitor. Clips recorded with a Log gamma curve typically have a low contrast, low saturation appearance when viewed on an sRGB or HD video monitor.



A clip captured with non-linear gamma curve.

A larger problem is that Fusion is a linear high dynamic range compositing system with 32-bit floating point precision. That means compositing operations in Fusion handle image data, assuming the clips have linear gamma. Common operations such as those that divide an image (a.k.a. “unpremultiply”), composite modes such as “screen,” merge operations, and many other compositing tasks only work properly with a linear gamma, especially when the images or effects include bright highlights,

For example, you can apply filtering effects, such as a blur, to an image with any gamma setting, and the image will probably look fine. However, if you convert the image to a linear gamma first and then apply the blur, then images (especially those with extremely bright areas) will be processed with greater accuracy, and you should notice a different and superior result.

The answer to these problems is color management.

Images coming into Fusion via Loader nodes are not color managed. The image is displayed directly from the file to the viewer without any interpretation or conversion. However, Fusion has nodes you can use to convert the image output of each Loader node to linear gamma at the beginning of your composite, and then convert from linear back to your output gamma at the end of your composite, directly before the Saver node.

When using standard dynamic range HD clips like ProRes or DNxHD, you can use the Gamut tool, and for clips from digital cinema cameras you can use the CineonLog tool. The Gamut tool is located in the Color category of the Effects Library, while the CineonLog tool can be found in the Film category.

Color Management for Standard Dynamic Range Clips

If you are using typical HD clips in ProRes, DNxHD, or some other video range format, you can use the Gamut tool for the conversion to linear color space.

To use the Gamut tool to convert to linear gamma:

- 1 Add the Gamut tool directly after a loader.
- 2 Select the Gamut node to view its controls displayed in the Inspector,
- 3 In the Inspector, select the clip's color space in the Source Space drop-down menu.
- 4 Click the Remove Gamma checkbox. This will convert the clip to linear gamma.
- 5 There is no need to adjust the Output Space; leave that set to No Change.



A node tree using Gamut nodes with a “to linear” conversion at the beginning and a “from linear” conversion at the end.

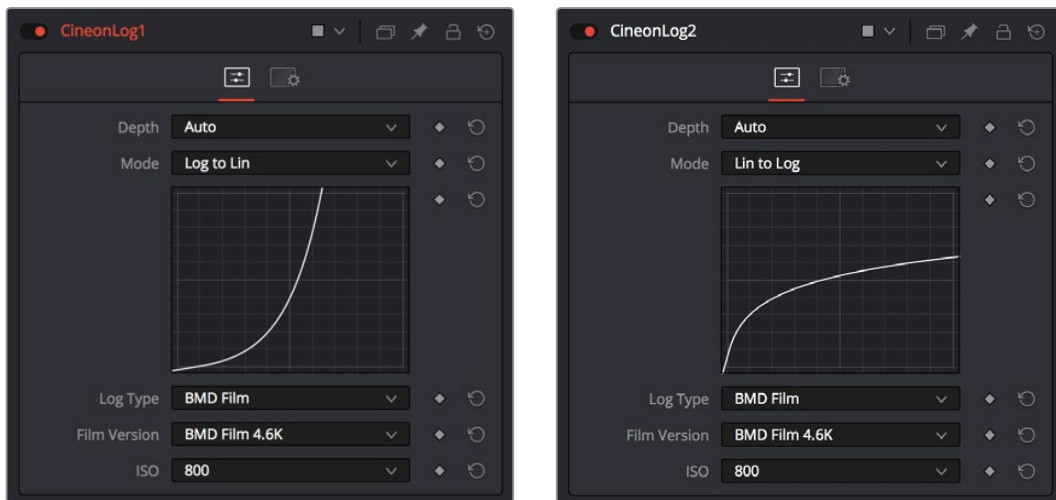
In the above example, the Loader's gamma is HD Rec 709, which is a standard dynamic range gamma, so the Loader is connected to the Gamut node, which converts the clip to linear. The second Gamut node then does a linear to Rec 709 conversion just prior to the Saver node.

Color Management for Wide Dynamic Range Clips

If you have media that comes from a digital cinema camera like the Blackmagic URSA MiniPro or others, convert to linear using the CineonLog tool. Just like the Gamut tool, the Loader node goes directly into a CineonLog node for the conversion to linear gamma. Then, depending on whether you will render the finish composite to Log encoded clip or an HD or UHD output, another CineonLog node or Gamut node is added just before the Saver node.

- 1 Add a CineonLog node directly after the Loader of any Log gamma encoded clip.
- 2 To use the CineonLog tool to convert a Log clip to linear gamma: In the CineonLog node's Inspector, set the Mode setting to Log to Lin.
- 3 Set the Log type and Film version based on the camera that recorded the content.
- 4 If you are outputting to a Log format, insert a CineonLog node just before the Saver node
- 5 Set the Saver's Cineon Log node Mode to Lin to Log.
- 6 Set the Log Type and Film Version to your desired output format.
- 7 If your output is going to UDH or HD, a Gamut node is used just before the Saver instead of the Cineon Log tool.

- 8 Leave the Gamut SourceSpace setting at No Change.
- 9 Set the Output Space to Rec 709 Display for a gamma of 2.2 or to Rec 709 Scene for a gamma of 1.96.



The first CineonLog node does a Log to Lin conversion (left); the second CineonLog node does a Lin to Log conversion (right).

TIP: The Loader and Saver nodes also include built-in functionality for doing gamma and color space conversions, so these operations can be done without needing to add additional nodes to your tree. However, adding color management nodes manually does make these operations more visible.

Viewer Gamma and Color Space While Working in Linear

Images converted to a linear gamma don't look correct. In fact, they usually look all wrong. Since all image data is converted to a linear scale for the convenience and accuracy of compositing operations, highlights usually are stretched to look extremely bright and blown out, and colors can become exaggerated and oversaturated. Happily, even though the image appears to be incorrect, the fact that Fusion works entirely with 32-float color data internally means that you're not actually clipping or losing any image data. It just looks bad when viewing the linear state of your image data directly.

It would be obviously impossible to work if you couldn't see the image as it's supposed to appear within the final gamut and gamma you'll be outputting to. For this reason, each viewer has a LUT control that lets you enable a "preview" color space and/or gamma conversion that lets you see the image in your intended color space and gamma, while the node tree is processing correctly in linear gamma.

Clicking the Viewer LUT button toggles LUT display on or off, while its accompanying pop-up menu lets you choose which color space and gamma conversions to view with.


To set the LUT appropriately for most computer monitors:

- 1 Click the LUT pop-up menu and choose the Gamut View LUT.
- 2 From the same pop-up menu select Edit.
- 3 In the ViewLUT editor that opens set the source's color space and the output display's color space. So, If you are viewing a linear image, keep the Source Space menu at No Change.

- 4 Assuming you are viewing on a standard computer monitor, set the Output Space to sRGB.
- 5 Click the Add Gamma checkbox.

You will now see a normalized image in the viewer, but all color operations will be on linear images.

TIP: If your monitor is calibrated differently, you will need to select a LUT that matches your calibration.



With the Viewer LUT button turned on, choose a LUT to use from the pop-up menu.

Whether you use the sRGB LUT or a LUT for your specific monitor calibration, you can save the viewer set up as the Default.

To Save the Gamut LUT setup as the default Viewer setup, do the following:

- Right-click in the viewer, and then choose Settings > Save Defaults.

You now have a linear color space working environment for image processing with a viewer configured for your monitor.

Chapter 4

Understanding Image Channels and Node Processing

This chapter seeks to demystify how Fusion handles image data, in the process showing you how different nodes must be connected in order 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	62
Types of Channels Supported by Fusion	62
Fusion Node Connections Carry Multiple Channels	63
Node Inputs and Outputs	64
Node Colors Tell You Which Nodes Go Together	67
Using Channels in a Composition	69
Channel Limiting	70
Adding Alpha Channels	71
How Channels Propagate During Compositing	72
Rearranging or Combining Channels	72
Understanding Premultiplication	73
The Rules of Premultiplication	74
How Do You Know You've Made a Premultiplication Mistake?	74
Setting the Premultiplied Status of Loader Nodes That Need It	75
Nodes That Affect Premultiplication	75
Control Premultiplication with Alpha Divide and Alpha Multiply	75
Understanding Auxiliary Channels	76
Image Formats That Support Auxiliary Channels	76
Creating Auxiliary Channels in Fusion	77
Auxiliary Channels Explained	77
Propagating Auxiliary Channels	81
Viewing Auxiliary Channels	82
Nodes That Use Auxiliary Channels	82
Merge	82
Depth Blur	82
Fog	82
Shader	83
SSAO	83
Texture	83
Shadow	83
Vector Motion Blur	83

Vector Distortion	83
Time Speed and Time Stretcher	83
New Eye	83
Stereo Align	83
Smooth Motion	84
Volume Fog	84
Volume Mask	84
Custom Tool, Custom Vertex 3D, pCustom	84
Lumakeyer	84
Disparity to Z, Z to Disparity, Z to WorldPos	84
Copy Aux	84
Channel Boolean	84

Channels in Fusion

If you're an old hand at compositing in Fusion, this chapter may be somewhat remedial. Otherwise, it's useful for you to know that Fusion introduces some innovative ways of working with the many different channels of image data that modern compositing workflows encompass. In particular, many shortcuts for handling different kinds of channels have been built into the way that different nodes interact with one another, making this chapter's introduction to color channels and how they're affected by different nodes and operations 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 are divided into separate channels, each of which carries a specific kind of image data. Nodes that perform different image processing operations typically expect specific channels in order to provide predictable results. This section describes the different kinds of channels that Fusion supports. Incidentally, all image data in Fusion defaults to 32-bit float.

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 a grayscale channel that represents different levels of transparency in an image. In Fusion, white denotes areas that are solid, while black denotes areas that are transparent. 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

These channels are created by Fusion whenever you create a Mask node. Mask nodes are unique in that they propagate single-channel image data that often serves a similar function as an alpha channel, defining 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 used for keying and compositing, such as the Merge node, the DeltaKeyer node, and the Matte Control node.

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 feature in an image along a Z axis (XYZ), 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 and animation, so it usually accompanies images generated by Autodesk

Maya or 3DS Max, or it may 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.

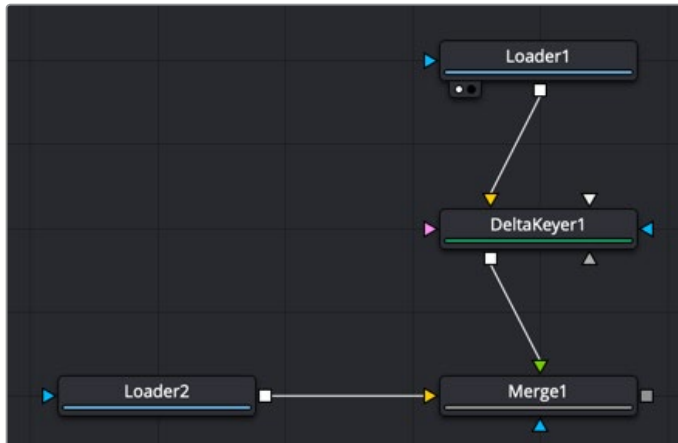
The reason to use auxiliary data is that 3D rendering is computationally expensive and time-consuming, so outputting descriptive information about a 3D image that's been rendered empowers compositing artists to make sophisticated alterations in 2D to fine-tune focus, lighting, and depth compositing that are faster (cheaper) to perform and readjust in 2D 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 control in the viewer. Clicking the Color control switches between Color (RGB) and Alpha, but clicking its pop-up menu control 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 the Node Editor of Fusion are capable of carrying multiple channels of image data along a single line. That means that a single connection may route RGB, or RGBA, or RGBAZ-Depth, or even just Z-Depth, depending on how you've wired up your node tree.

In the following example, each of the two Loader nodes output RGB data. However, the Delta Keyer adds an alpha channel to the foreground image that the Merge node can use to create a two-layer composite.



Loader1 node connected to a DeltaKeyer node, connected to a Merge node, which is connected to another Loader node to combine the two images using the alpha channel output by the DeltaKeyer.

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

Loader nodes output all available channels from the source media on disk. When you connect one node's output to another node's input, those 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 Loader node's output is connected to the input of a Highlight node to create a sparkly highlight effect.



Loader node connected to a Highlight node connected to Saver node.

When connecting nodes together, a single node output can be connected to multiple node's inputs, which is known as "branching." This is useful when you have a single node that you want to feed multiple operations at the same time.



The Loader node's output is branched to the inputs of two other nodes.

Using Multiple Inputs

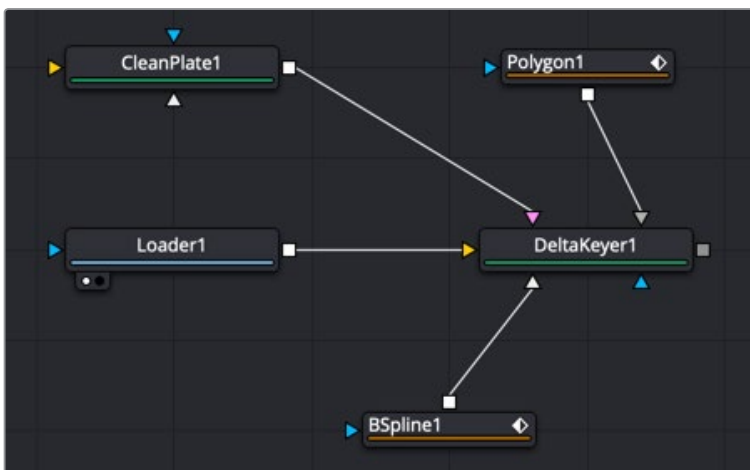
Most nodes have two inputs, one for RGBA and another for a mask that can be optionally used to limit the effect of that node to a particular part of the image. However, some nodes have three or even more inputs, and it's important to make sure you connect the correct image data to the appropriate input in order to obtain the desired result. If you connect a node 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 Matte or Mask channels to modify the alpha key being extracted from the image in different ways. If you want to perform the extremely common operation of using a MatteControl node to attach 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, since the GarbageMatte input is automatically set up to use the input mask to alter the alpha channel of the image. If you connect to any other input, your Polygon mask won't work.




Polygon node connected to a MatteControl node for roto-scoping.

In another example, the DeltaKeyer node has a primary input (labeled “Input”) that accepts RGBA channels, but it also has a CleanPlate input for attaching an RGB image with which to clean up the background (typically the CleanPlate node), and SolidMatte, GarbageMatte, and EffectsMask inputs that accept Matte or Mask channels to modify the alpha key being extracted from the image in different ways. To pull a key successfully, though, you must connect the image you want to key to the “Input” input.



Loader node connected to the main “Input” input of a DeltaKeyer node; other nodes connect to specific inputs for those particular nodes.

If you position your pointer over any node’s input or output, a tooltip will appear in the Tooltip bar at the bottom of Fusion 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 will appear in the Node Editor itself.



The node input’s tooltip in the Tooltip bar (left), and the node tooltip in the Node Editor (right).

Connecting to the Correct Input

When you’re connecting nodes together, 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 afterward to keep the node tree tidy-looking.

TIP: If you hold down the Option key 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.

Some Nodes Are Incompatible with Some Inputs

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 Merge3D node's output directly to the input of a regular Merge node, it won't work; you must first connect to a Renderer3D node that creates output appropriate for 2D compositing operations.

In other cases, connecting the wrong image data to the wrong node input won't give you any sort of error; it will simply fail 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 (or the Fusion Tool Manual for previous versions of Fusion) 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. To help you make the right corrections, 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. It usually connects to the orange-colored background input 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.

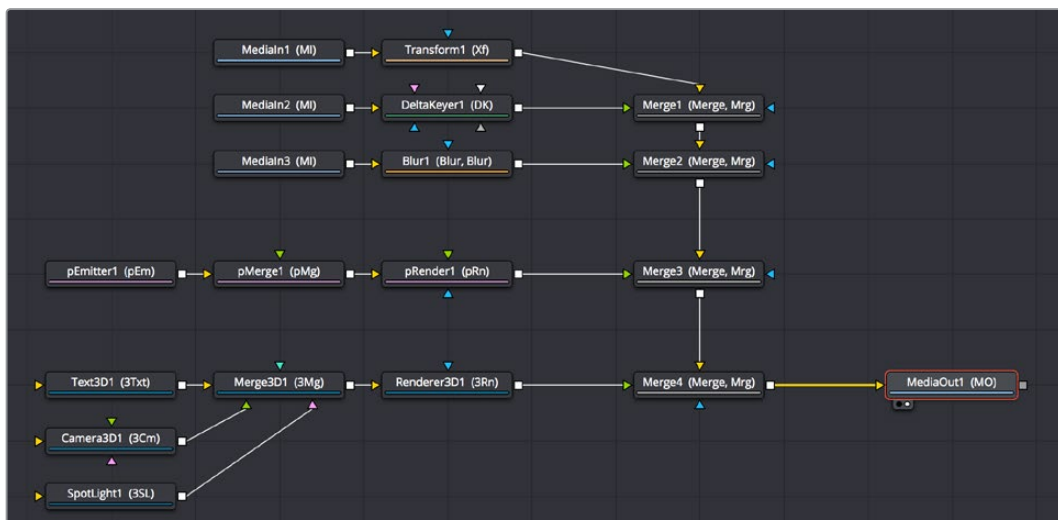
If you connect to only one input of a multi-input node and you don’t connect to the background input, you may find that you don’t get the results you wanted. This is because each multi-input node expects that the background will be connected before anything else, so that the internal connections and 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

Fusion includes over 200 nodes, and each of the many nodes 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, and it also makes it easier for artists to focus on fine-tuning specific adjustments, one at a time, when assembling the ever-growing tree of Loaders and image processing operations that make up one’s composite.

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 Loader Nodes and Green Generator Nodes

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 Loader 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.

2D Processing Nodes, Color Coded by Type

These encompass most 2D processing and compositing operations, 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, designed to let you color correct a foreground layer to match a background layer without affecting an alpha channel being used to create a composite, and the Gamut node, which 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 together to create different particle systems, and they're incompatible with other kinds of nodes until you add a pRender node which outputs 2D RGBA and auxiliary data that can be composited with other 2D nodes and operations.

Dark Blue 3D Nodes

These are 3D operations that generate and manipulate 3D data (including auxiliary channels) that is incompatible with other kinds of nodes until processed via a Renderer 3D node, which then outputs RGBA and auxiliary data.

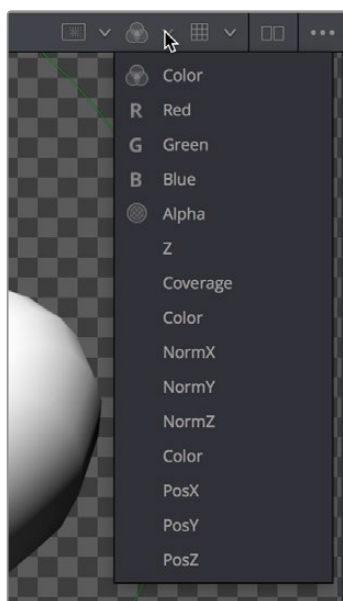
Brown Mask Nodes

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 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 that node actually uses or affects a particular channel.

Incidentally, if you want to see which channels are available for a node, you can open the Color pop-up menu in the viewer to get a list. This control also lets you view any channel on this list, so you can examine the channel data of your composite anywhere along the node tree.



All channels available to the currently viewed node can be isolated via the viewer's Color control.


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.



The Normal Z channel output by a rendered torus (left), and 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 (right).

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, have no effect on alpha or auxiliary channels. This makes them convenient for color-matching foreground and background layers you're compositing, without worrying that you're altering the transparency or depth information accompanying that layer.




Loader, DeltaKeyer, ColorCorrector, and Merge/Loader node.

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 you want to modify to another output channel that's compatible with the operation you're trying to perform, and then use another Channel Booleans node to reassign it back. 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 checkboxes in the Settings panel of that node's controls in the Inspector. These checkboxes let you exclude any combination of these channels from being affected by that node.



Setting the channel limiting checkboxes 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 Green, 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

Skiping 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 checkboxes on another tab in the node. In these cases, the Common Control channel checkboxes 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.

Adding Alpha Channels

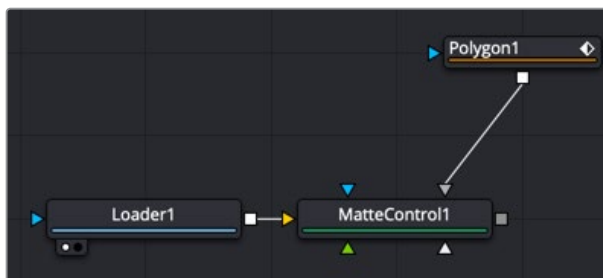
One of the reasons for compositing is to begin with a foreground image that lacks an alpha channel, add one via keying or roto-scoping, and then composite that result against other background images in order to combine a foreground subject with a background. While the methods for this are covered in detail in later chapters, here's an overview of how this is handled within Fusion.

In the case of extracting an alpha matte from a greenscreen image, you typically connect the image's RGB output to the "Input" input on one of the Keyer nodes such as the Delta Keyer, and you then use the keyer's controls to pull the matte. A Keyer node automatically inserts the alpha channel that's generated alongside the RGB channels, so the output is automatically RGBA. Then, when you connect a Keyer's output to a Merge node in order 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.

In the case of roto-scoping using a Polygon or B-Spline node, you'll typically connect the image being roto-scoped to the background input of a MatteControl node, and a Polygon or B-Spline node to its garbage matte input (which you invert in the Inspector, unless you invert the Polygon's output). This lets you view the image while drawing, using the controls of the Polygon or B-Spline node, and the resulting alpha channel is merged together with the RGB channels so the Merge Alpha node's output is RGBA, which can be connected to a Merge node to composite the roto-scoped subject over another image.



A simple roto-scoping node tree

In both cases, you can see how Fusion node tree's ability to carry 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 will 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 greenscreen plate with subjects or elements that are meant to be combined against the background.

Many compositions use multiple Merge nodes to bring together many differently processed branches of a large node tree, so 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 is compared. Which one is in front depends on its Z-Depth and not which input it's connected to.

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 three 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. Alternately, 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:** Is a 3D operation 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 a Blinn. SpecularExponent, you can remap the channels here.
- **Channel Booleans:** A more standard 2D version, the Channel Booleans is used to shuffle YRGB/auxiliary channels within a single input image, or among two input images, to create a single output image. If you connect only a single image to this node, it must be connected to the background input to make sure everything works.
- **Matte Control:** Designed to do any combination of the following: (a) recombining 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 image, alpha, and auxiliary channels in Fusion, it's time to learn a little something about premultiplication, to make sure you always combine RGB and alpha channels correctly to get the best results from Merge node composites.

Premultiplication is an issue whenever you find yourself compositing multiple images together, and when at least one of them contains RGB with an alpha channel. For example, if a motion graphics artist gives you a media file with an animated title graphic that has transparency rendered into it to accommodate later compositing, or if an animator gives you an isolated VFX plate of a spaceship coming in for a landing with the transparency baked in, you may need to consider the premultiplied state of the RGBA image data as you use these images.


Most computer-generated images you'll be given should be premultiplied. A premultiplied alpha channel means that, for every pixel of an image, the RGB channels are multiplied by the alpha channel. This is standard practice in VFX workflows, and it guarantees that translucent parts of the rendered image, such as flares, smoke, or atmospheric effects, are correctly integrated into the background black areas of the isolated image, so that the image appears correctly when you view that layer by itself.

NOTE: Computer generated 3D images that were rendered anti-aliased are almost always premultiplied.

So-called "straight" alpha channels, where the RGB channels have not been multiplied by the alpha channel, will appear weirdly bright in these same translucent areas, which tells you that you probably need to multiply the RGB and A channels prior to doing specific tasks.



Premultiplied alpha image.



Straight alpha image.

The Rules of Premultiplication

In general, 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:

- Always color-correct images that are not *premultiplied*.
- Always filter and transform images that are *premultiplied*.

How Do You Know You've Made a Premultiplication Mistake?

Improper handling of premultiplication manifests itself in a few obvious ways:

- You see thin fringing around a subject composited with a Merge node.
- You notice a node adjustment affecting parts of the image that shouldn't be affected by that operation.
- You've combined RGB and alpha channels from different sources and the checkerboard background pattern in the viewer (if enabled) is only semi-transparent when it should be fully transparent.


If you spot these sorts of issues, the good news is they're easy to fix using either the internal settings of the nodes causing the problem, or with dedicated nodes to force the premultiplied state of the image at specific points in your node tree.

Setting the Premultiplied Status of Loader Nodes That Need It

When you select a Loader node, the Import panel in the Inspector has a group of checkboxes that let you determine how an alpha channel embedded with that image should be handled. There are checkboxes to make the alpha channel solid (to eliminate transparency), to Invert the alpha channel, and to Post-Multiply the RGB channels with the alpha channel, should that be necessary.

Nodes That Affect Premultiplication

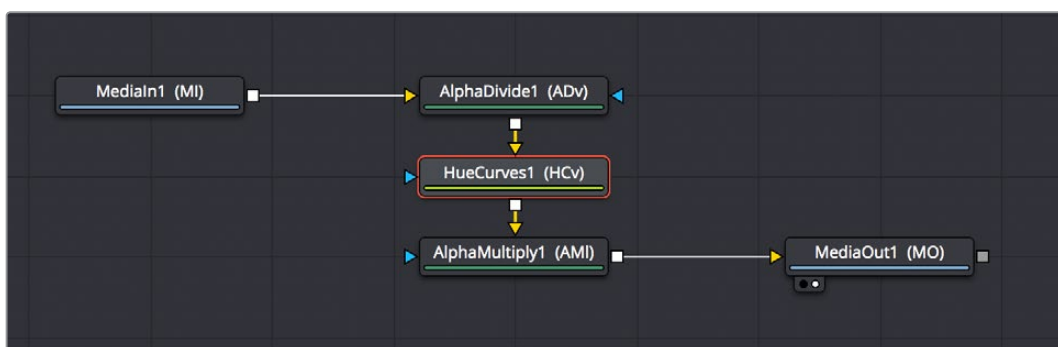
Most nodes that require you to explicitly deal with the state of premultiplication of RGBA image input have a “Pre-Divide, Post-Multiply” checkbox. This includes simple color correction nodes such as Brightness Contrast and Color Curves, as well as the Color Correct node, which has the “Pre-Divide/Post-Multiply” checkbox in the Options panel of its Inspector settings.



The Pre-Divide/Post-Multiply checkbox of the Color Curves node, seen in the Inspector.

Control Premultiplication with Alpha Divide and Alpha Multiply

The Alpha Divide and Alpha Multiply nodes, found in the Matte category of the Effects Library, are provided whenever you need to do operations on RGBA image data where you need explicit control over the pre-multiplied state of an image’s RGB channels against its alpha channel. Simply add the Alpha Divide node when you want the RGBA image data to not 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 manually control premultiplication before and after such an adjustment.



A node tree with explicit Alpha Divide and Alpha Multiply nodes.

Understanding Auxiliary Channels

Auxiliary channels describe a family of special-purpose image data that typically describes 3D position, orientation, and object information for use in 2D composites. For example, Z-Depth channels describe the depth of each region of an image along a Z axis (XYZ), while an XYZ Normals channel describes the orientation (facing up, facing down, facing to the left or right) of each pixel in an image. Auxiliary channel data is generated by rendering 3D data, so it may accompany images generated by Autodesk Maya or 3DS Max, or it may 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.

One of the most common reasons 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 affecting focus, lighting, and depth compositing that are faster to perform and readjust in 2D than it would be to re-render the 3D source material over and over.

There are two ways of obtaining auxiliary channel data:

- First, auxiliary data may be embedded within a clip exported from a 3D application that's in a format capable of containing auxiliary channels. 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, using the Optical Flow node, or using the Disparity node.



An RGBA 3D rendered scene that can also generate auxiliary channels.

Image Formats That Support Auxiliary 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 can 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) 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)

These image formats are 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, so different layers of the image can be loaded for very complex depth composites.

Fusion RAW (*.RAW)

Fusion's native RAW format is able to contain all 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 Fusion 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.

Auxiliary Channels Explained

Fusion is capable of using auxiliary channels, where available, to perform depth-based compositing, to create masks and mattes based on Object or Material IDs, and for texture replacements. Tools that work with auxiliary channel information have been specifically developed to work with this data.

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.

Z-Coverage

The Z-Coverage channel is used to indicate pixels in the Z-Depth that contains two objects. The value is used to indicate, as a percentage, how transparent that pixel is in the final depth composite.

WARNING: Depth composites in Fusion that are based on images that lack a Z-Coverage channel, as well as a background RGBA channel, will not be properly anti-aliased.



Z-Coverage channel.

Background RGBA

This channel contains the color values from the objects behind the pixels described in the Z-Coverage.



Background RGBA.

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.



Object ID.

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 the creation of masks based on materials.



Material ID.

UV Texture

The UV Texture channels contain information about mapping coordinates for each pixel in the scene. This is used to apply textures wrapped to the object.



UV Texture.

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.



XYZ Normals.

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 as a post process or to track pixels over time for retiming. The XY Vector points to the next frame, while the XY BackVector points to the previous frame.



XY Vector.

XYZ Position


The XYZ Position channels indicate where each pixel is assigned; the XYZ position of its location is in 3D space, typically in world coordinates. This can be used, like Z-depth, for compositing in depth but can also be used for masking based on 3D position, regardless of camera transforms.



XYZ Position.

XY Disparity

The XY Disparity 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.



XY Disparity

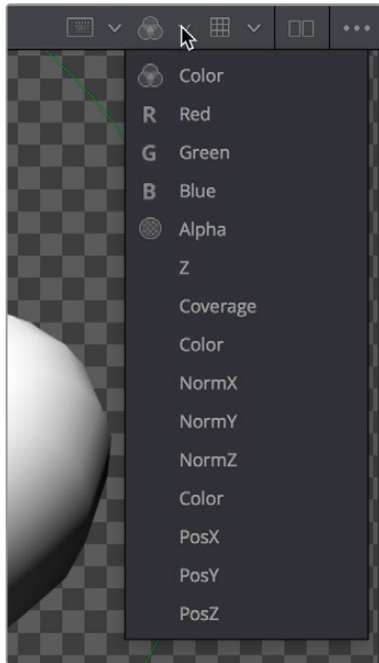
Propagating Auxiliary Channels

Ordinarily, auxiliary channels will be 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 will propagate (and potentially manipulate) auxiliary channels without problems.

However, when you composite two image layers using the Merge node, auxiliary channels will 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 greenscreen plate with subjects or elements that are combined against the background, which lack auxiliary channels.

Viewing Auxiliary Channels

You can view the auxiliary channels by selecting the desired channel from the viewer's toolbar or from the viewer's contextual menu. The Color Inspector SubV can also be used to read numerical values from all of the channels.



Selecting a channel from the viewer's toolbar.

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.

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.

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.

Fog

The Fog tool makes use of the Z-Depth to create a fog effect that is thin when closer to the camera and that 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.

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.

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.

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.

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.

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.

Vector Distortion

The forward XY Vector channels can be used to warp an image with this tool.

Time Speed and Time Stretcher

These tools can use the Vector and BackVector channels to retime footage.

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.

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.

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.

Volume Fog

Volume Fog is a raymarcher that uses the 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 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.

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.

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.

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.

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 Boolean

The Channel Boolean tool can be used to combine or copy the values from one channel to another in a variety of ways.

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.

Chapter 5

Preferences

This chapter covers the various options that are available from the preferences panel.

Contents

Preferences Overview	87
Categories of Preferences	87
Preferences In Depth	89
3D View Preferences	89
Defaults	91
Flow	92
Frame Format	94
General	96
GPU	98
Layout	99
Loader	100
Memory	102
Network	103
Path Maps	104
Preview	105
Quicktime	106
Script	107
Spline Editor	108
Splines	109
Timeline	110
Tweaks	111
User Interface	114
Video Monitoring	115
View	116
VR Headsets	117
Bins/Security	119
Bins/Server	120
Bins/Settings	121
EDL Import	122
Customization	123
Toolbar Customization	123
Shortcuts Customization	123

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 dialog. The Preferences dialog can be opened from a menu at the top of the interface.

To open the Preferences window, do one of the following:

- On Mac OS X, Choose Fusion Studio > Preferences.
- On Windows, Choose File > Preferences.

The Preferences window is divided into a category sidebar on the left and the settings panel on the right. There are two levels of preferences: Global and Composition. 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 can further modify the currently open composition without affecting the Global preferences or any other flow 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 preferences disclosure arrow reveals the following sections.

3D View

The 3D View preference offers control over various parameters of the 3D Viewers including grids, default ambient light setup, and Stereoscopic views.

Defaults

The Defaults preference is 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 preference 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 preference is 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 preference contains options for the general operation, such as auto save, and gamma settings for color controls.

GPU

The GPU preference includes 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

You can use the Layout preference to save the exact layout of Fusion's windows.

Loader

Using the Loader preference, you can set options for the default Loader's depth and aspect ratio as well as define the local and network LoaderCache settings.

Memory

Memory management for multi frame and simultaneous branch rendering is configured in the Memory preference.

Network

The Network rendering preference is 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 preference is 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

The Preview preference is where you configure the Preview creation and playback options.

QuickTime

This section lets you preconfigure the QuickTime codec used for rendering.

Script

The Script preference includes 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 preference allows 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 preference is where you create and edit Timeline/Spline filters and set default options for the Key Frames Editor.

Tweaks

The Tweaks preference handles 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

The Video Monitoring preference is where you can configure your Blackmagic video display hardware for monitoring on an HD, Ultra HD, or DCI 4K display.

View

The View preference is 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 preference allows configuration of any connected Virtual Reality headsets, including how stereo and 3D scenes are viewed.

Bins

There are three panels as part of the Bins preference: 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.

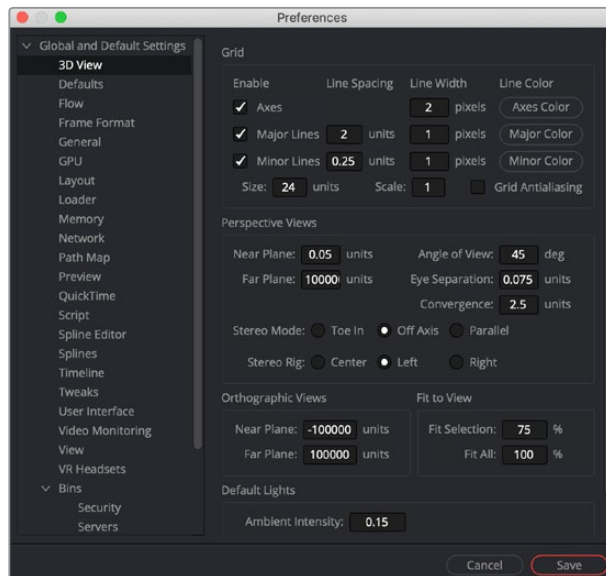
Import

The Import settings contains options for EDL Import that affect how flows are built using the data from an EDL.

Preferences In Depth

3D View Preferences

The 3D View preferences contain settings over 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 preference 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 the grid off 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.

PerspectiveViews

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 view, the Orthographic Views (front, top, right, and left views) section sets the nearest and furthest point any object can get to or from the view before it is clipped.

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.

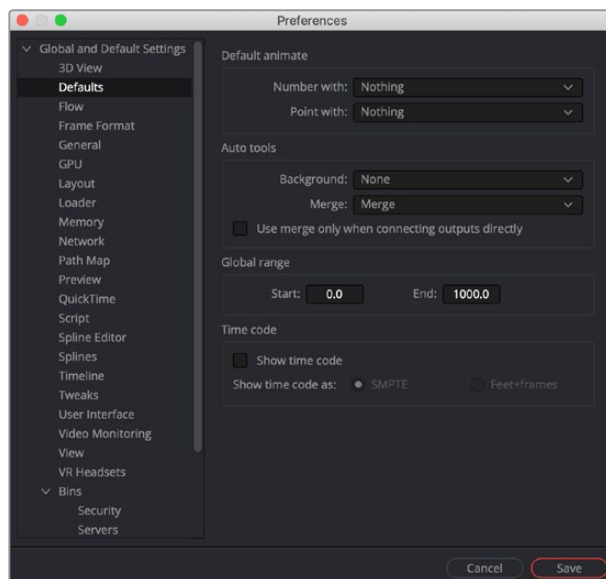
Default Lights

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.

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 Bezier 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

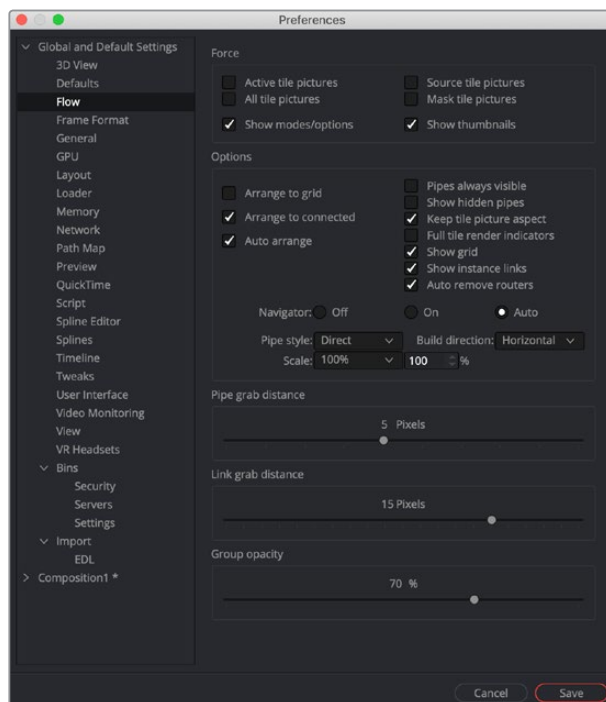
Using the Start and End fields, you can define the Global Start and End frame used when creating new compositions.

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 has a number of 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 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

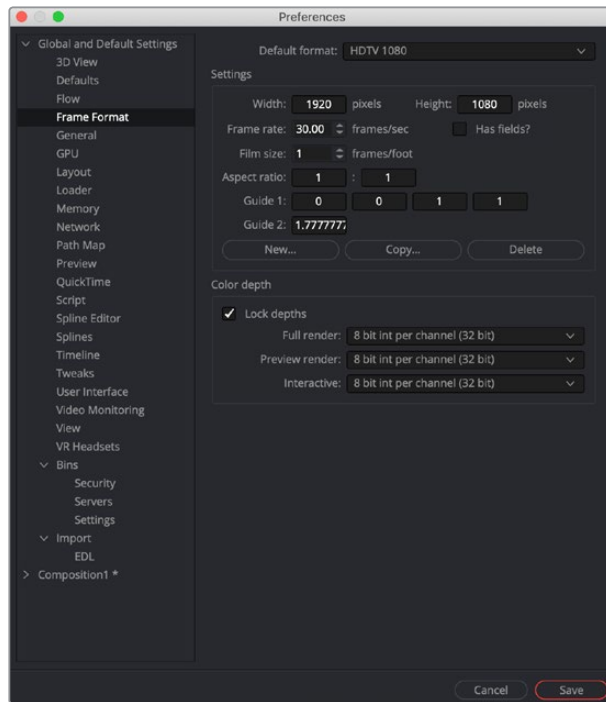
This Pipe Grab Distance slider allows you to choose how close the cursor must be (in pixels) to a connection line in the node tree when selecting them.

Group Opacity

This slider controls the opacity of an expanded group's background in the Node Editor.

Frame Format

Frame Format preferences select the format settings for the interface. For more detail, read Chapter X “Frame Formats” in this manual.



The Frame Format preferences.

Default Format

This drop-down menu is used to select the default resolution for tools and media from a list of presets. This is only a default setting; Loaders and Creators can be adjusted to different resolutions.

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 timecode 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 on screen. 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 Button:** 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 Button:** The Copy button copies the current setting to create a new one for customization.
- **Delete Button:** The Delete button will remove the current setting from the default drop-down list.

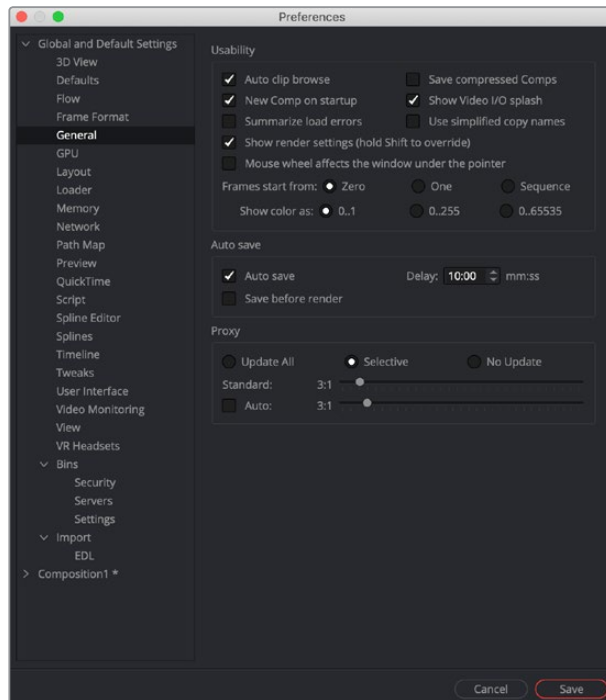
Color Depth

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 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 flow.
- **New Comp on Startup:** When checked, a new, empty project is created each time the application is launched.
- **Summarize Load Errors:** When loading node trees or “comps” that contain unknown tools (e.g., comps that have been created on other computers with plug-ins 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.
- **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 Render Settings dialog will be displayed every time a render is started. 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 track pad 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

When enabled, 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

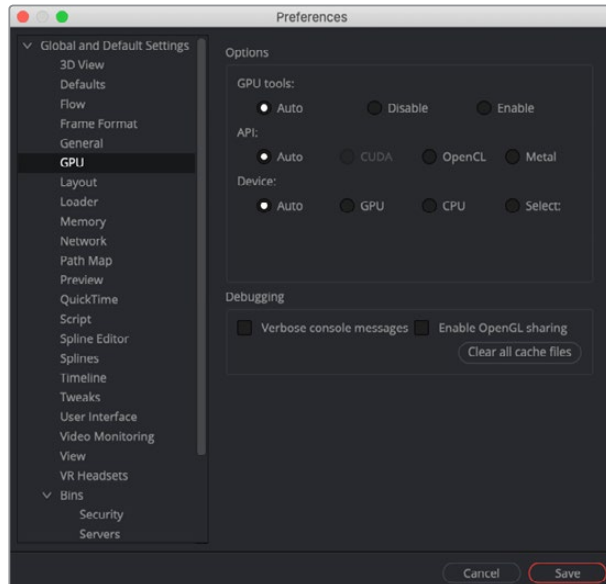
- **Update All, Selective, No Update:** The Update mode button is located above the toolbar. You can use this preference to determine the default mode for all new comps. Selective is the usual default. It renders only the tools needed to display the images in the Display view. All will render all tools 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 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

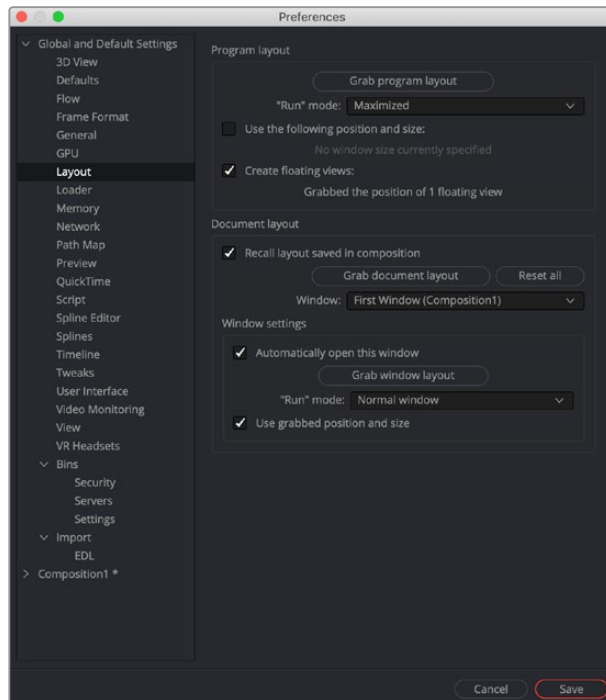
- **GPU tools:** This preference has three settings: Auto, Disable, and Enable. When set to Disable, no GPU acceleration is used for tools or third-party plug-ins. Fuses may still require GPU acceleration. If Enable is selected, GPU acceleration is available for tools and plug-ins, if appropriate drivers 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, the platform default will be used. 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

- **Verbose Console Messages:** Enabling this option causes information to be shown in the Console. For example, Startup Logs, Compiler Warnings, and Messages.
- **OpenGL Sharing:** No Information was available at the time of this writing.
- **Clear Cache Files:** This option will clear already compiled GPU code and then recompile the kernels.

Layout

These 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 layout.

Program Layout

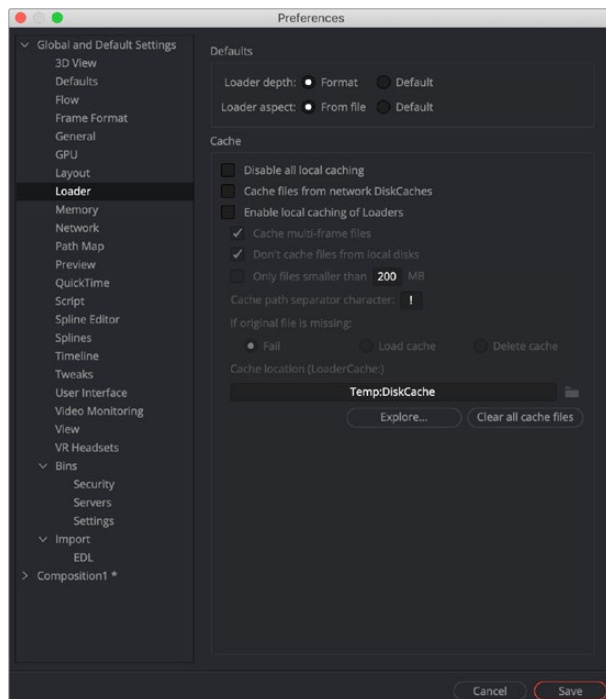
- **Grab Program Layout:** Pressing this button stores the program's overall current position and size.
- **Run Mode:** This menu is used to select the program's default mode at startup. You choose between a Maximized program, a Minimized program, or a Normal program display.
- **Use the Following Position and Size:** When checked, the values stored when Grab Program Layout was selected will be used when starting Fusion.
- **Create Floating Views:** When checked, the position and size of the floating display viewers will be saved when the Grab Program Layout button is used.
- Document Layout
- **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

- **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 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

Using the Loader preference, 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.

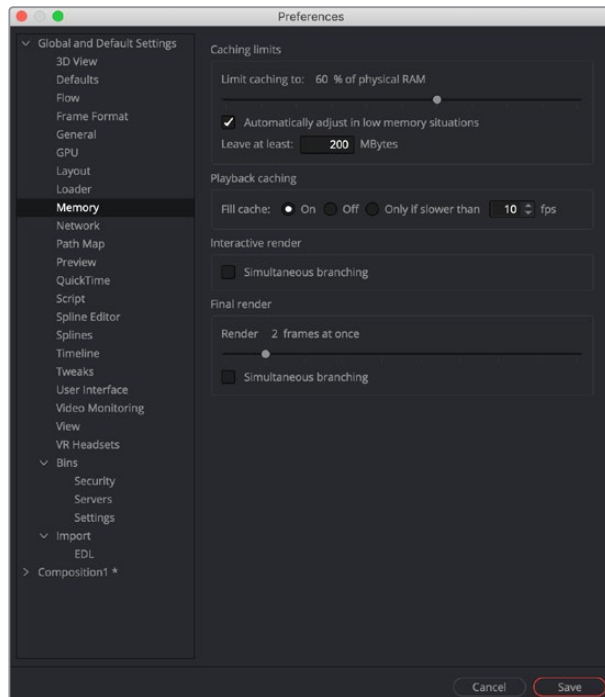
Cache

- **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 Mac OS X Finder window or a Windows Explorer window.
- **Clear All Cache Files:** This button deletes all cached files present in the LoaderCache path.

Memory

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.



The Memory preferences.

Caching Limits

- **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.

Interactive Render

- **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.

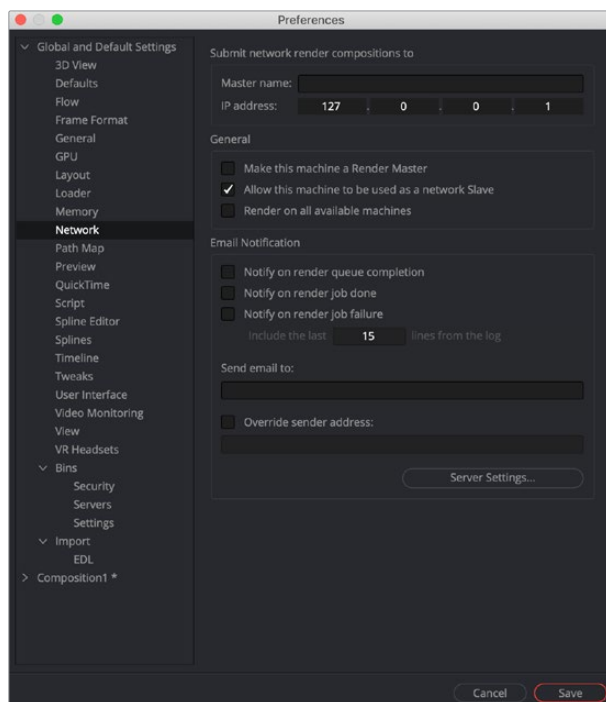
Final Render

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

These preferences are used to set up and control Fusion's network rendering. 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

- **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

- **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.

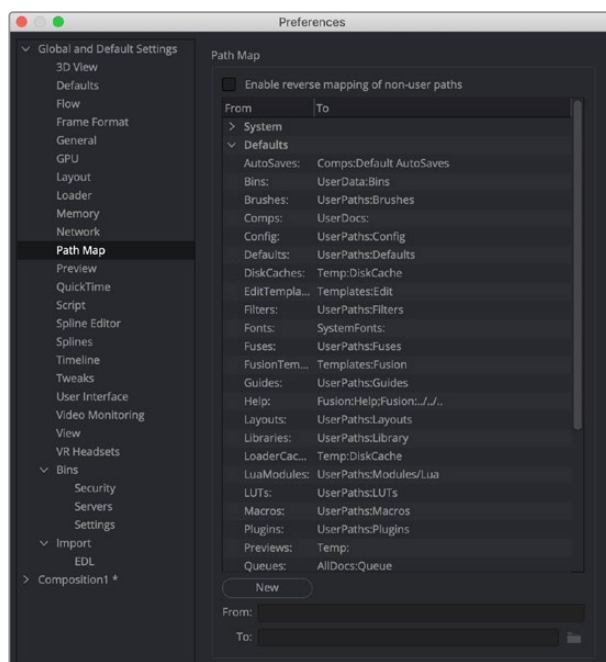
Server Settings

This section covers Clustering and Network Rendering.

For more detail on these settings and clustering, refer to ChapterX, “Network Rendering.”

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.



The Path Map preferences.

There are two main advantages to virtual Path Maps instead of actual filenames. One is that you can easily change the path to footage (for example, copying from one share to another), without needing to make any changes at all to their compositions.

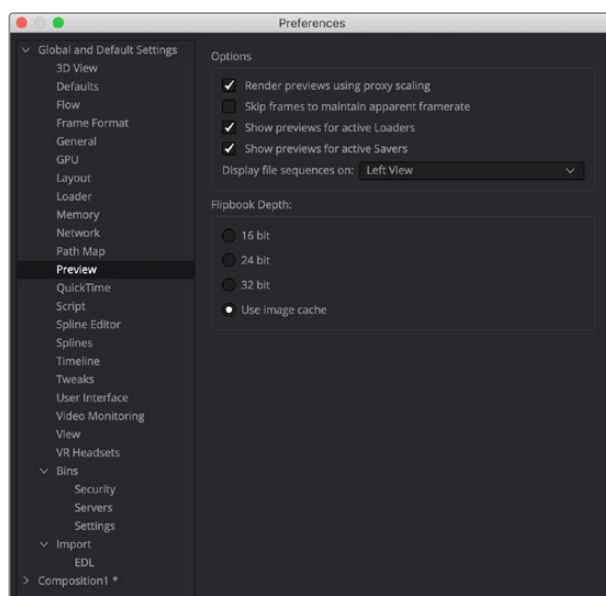
Also, path maps are used when network rendering to bypass the different filename conventions.

- **Built in Path Maps:** There are several built in path maps. Comp refers to the folder where the current composition is saved. Temp refers to the system's temporary folder. Fusion refers to the folder where Fusion is installed.
- **Global and Composition Path Maps:** Both Global and Composition preferences show a Path Maps preferences panel. The Global preferences are applied to all compositions, while Composition path maps are saved with the composition and only apply to it.

Composition path maps will override Global path maps with the same name. The built-in Comp.Path Map refers to the default Composition folder when used in a Global path map.
- **Nesting Path Maps:** You can use a path map in a path map's definition, provided it has been defined first. For example, define a path map called 'Episode' that maps to x\shows\Episode1. Then create path maps called Renders and Stills that mapped to Episode\Renders_v1 and Episode\Stills_v1.
- **Enable Reverse Mapping of Paths Preferences:** When this checkbox is selected, the built-in path maps for entries in the Paths preferences will be taken into account when applying mapping to existing filenames.
- **Creating a 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 on the Delete button. Built-in path maps cannot be deleted. Only path maps created by the user can be removed from the path maps list.

Preview

The Preview preference is where you configure the creation and playback options for preview renders.



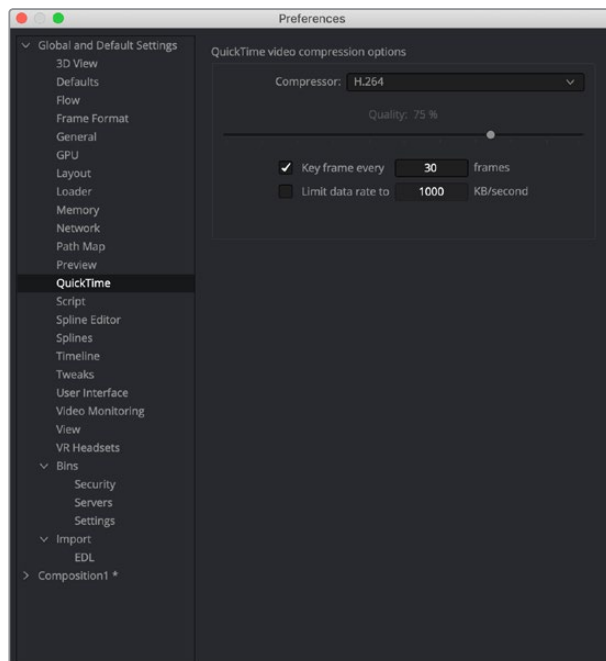
The Preview preferences.

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 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 preference configures the default QuickTime codec settings when QuickTime is selected as the rendering file format in the Saver node.



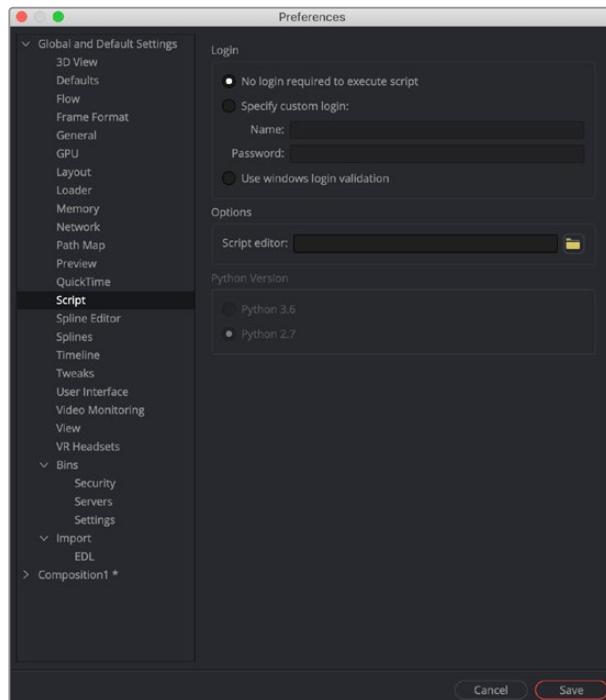
The QuickTime preferences.

- **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 on to the workstation in order to run.
- **Specify Custom Login:** If a username and password is assigned, Fusion will refuse to process incoming external script commands (from FusionScript, for example), unless the Script first logs on 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 detail, 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

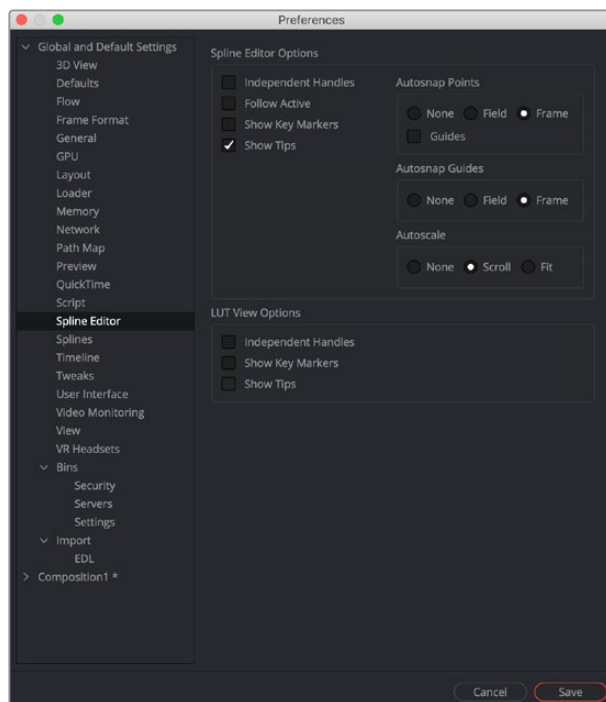
- **Script Editor:** Use this preference to select an external editor for scripts. This preference is used when selecting Scripts > Edit.

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 preference allows 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 Spline Editor preferences.

Spline Editor Options

- **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.

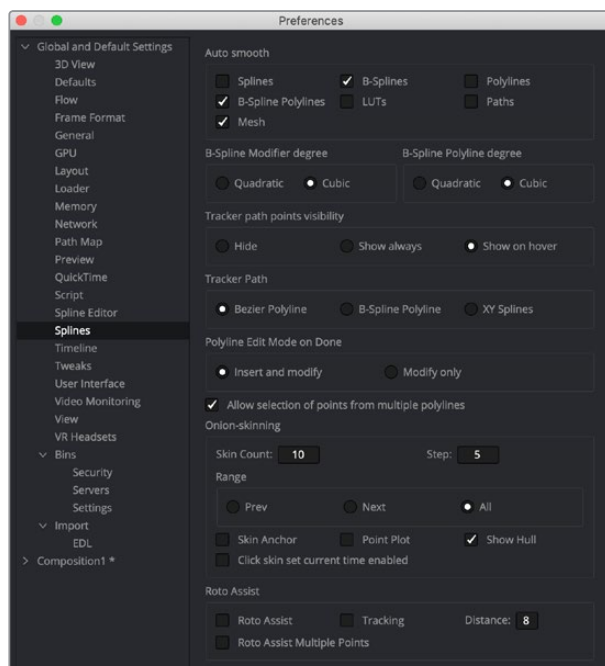
- **Autosnap Points:** When moving points in the Spline Editor, these will snap to the Fields or Frames or can be moved freely. This option is also available via the Options submenu when right-clicking in the Spline editor graph.
- **Guides:** When moving points in the Spline Editor, these will snap 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

- **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 if tooltips are displayed or not.

Splines

Options for the handling and smoothing of animation splines, tracker path defaults and roto-scoping are found in the Splines preferences.

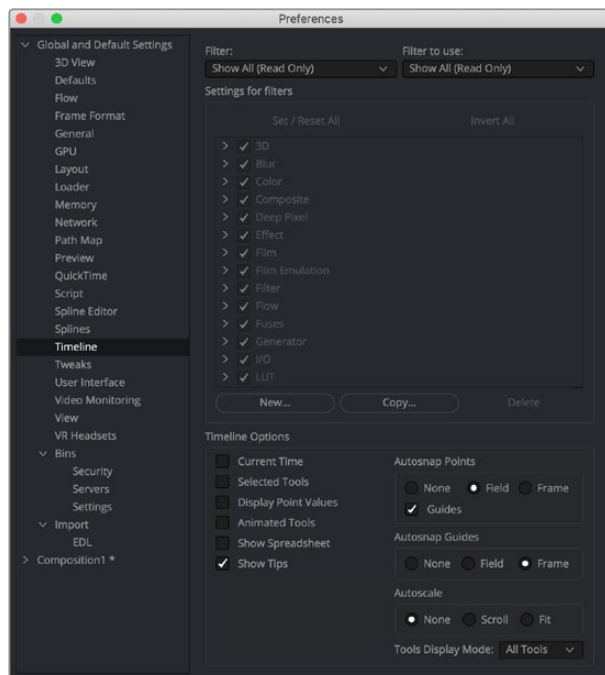


The Splines preferences.

- **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-Spline used for masks.
- **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 Bezier 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 frame prior to the current frame, after the current frames, or split between the two.

Timeline

The Timeline preference is where you create and edit Timeline/Spline filters and set default options for the Keyframe 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 Keyframe 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 Keyframe 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.

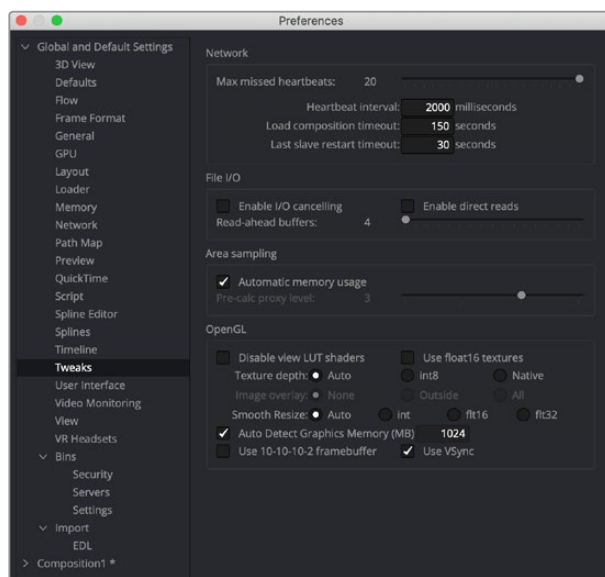
Timeline Options

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 detail on the Timeline functions, read Chapter 7, “The Keyframe Editor.”

- **Autosnap points:** When moving points in the Keyframe Editor, the points will snap to the Fields, to the Frames, or can be moved freely.
- **Guides:** When moving points in the Keyframe 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, to the Frames, or 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 Keyframe editor. The default can be changed using the Sort order menu in the upper right of the Keyframe editor.

Tweaks

The Tweaks preferences handle a collection of settings for fine tuning Network rendering and graphics hardware behavior.



The Tweaks preferences.

Network

The Network section is used to control and monitor the health of communication packets over a TCP/IP when rendering over a network.

- **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 manager 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 slave to respond before using another render slave.

File I/O

- **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 used 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

- **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.

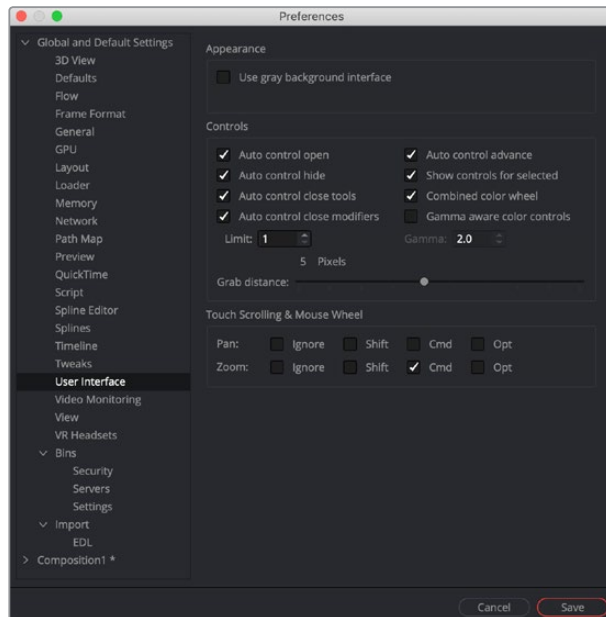
- **Disable view LUT Shaders:** OpenGL shaders can often dramatically accelerate View LUTs, but this can occasionally involve small tradeoffs in accuracy. This setting will force Fusion to process LUTs at full accuracy using the CPU instead. Try activating this if View LUTs do not seem to be giving the 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.

NOTE: 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

These 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.

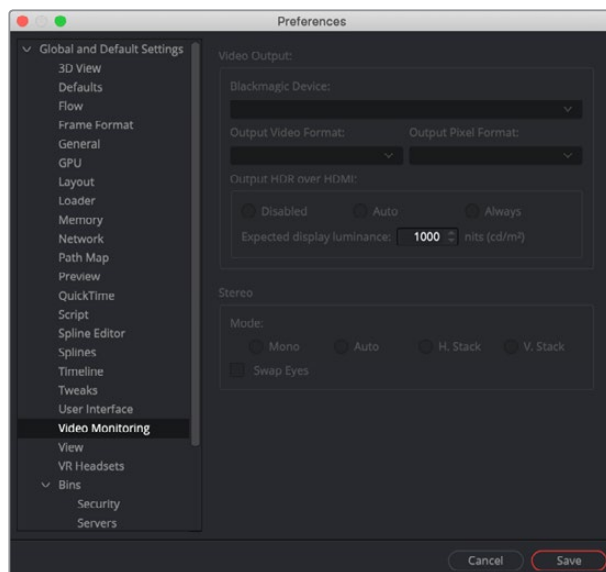
- **Show Controls for Selected:** When this option is disabled, only the active tool's parameters are shown in the Inspector. By default, it is enabled showing controls for the active tool as well as all selected tools.
- **Combined Color Wheel:** When the Color Corrector tool is displayed in the Inspector, 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.** There was no information available at the time this document was published.
- **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

The Video Monitoring settings 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 HDMI 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 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/m2) to whatever peak luminance level your HDMI connected HDR display is capable of.

Stereo Mode

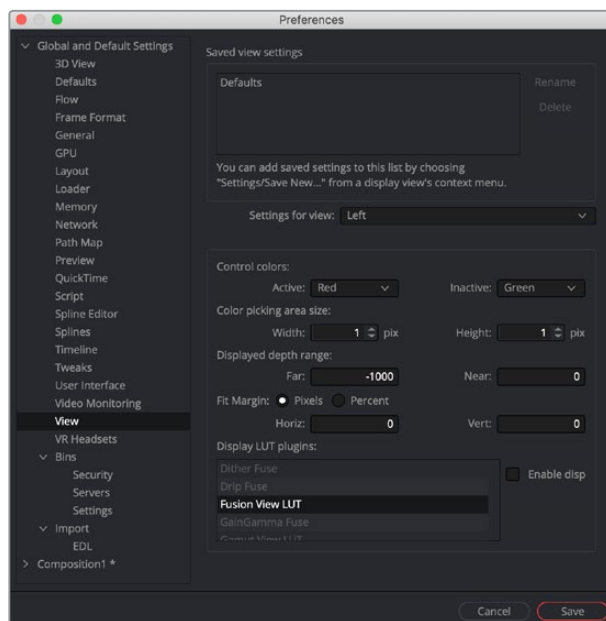
This group of settings configures the output hardware for displaying stereo 3D content.

- 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 preference is used to manage settings and default controls for viewers.



The View preferences.

View Settings

The area at the top of the viewer preference lists the currently saved settings that you create from the viewer's contextual menu. For more details on the viewer and its contextual menu, read Chapter 5 on the viewers. You can use the Rename and Delete buttons to manage the selected entries in the list.

Settings for View

Each viewer has its own preferences. The Settings for the 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

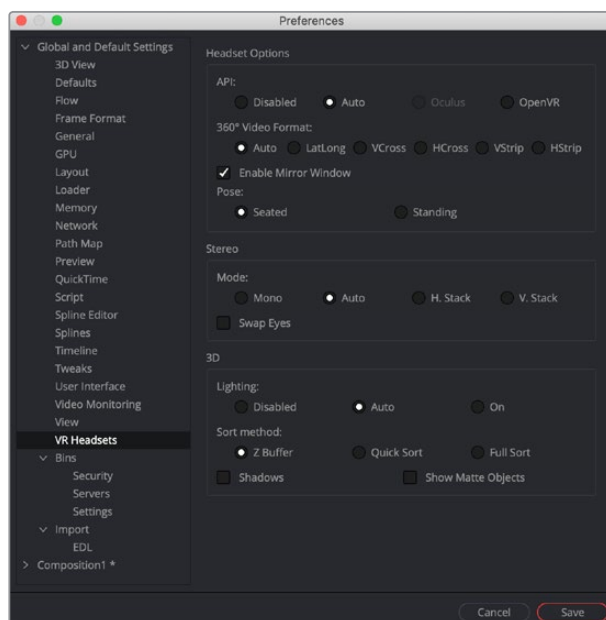
The Fit Margin determines how much padding is left around the frame when the Fit button is pressed or Fit is selected from the viewer's contextual menu.

Display LUT Plug-Ins

This list shows the available Display LUTs and activates the selected one as default.

VR Headsets

The VR Headsets preference allows configuration of any connected Virtual Reality headsets, including how stereo and 3D scenes are viewed.



VR Headsets preferences

Headset Options

API

- **Disabled** turns off and hides all usage of headsets.
- **Auto** will detect which headset is plugged in.
- **Oculus** will set the VR output to the Oculus headset.
- **OpenVR** will support a number of VR headsets like the HTC Vive.

360° Video Format

- **Auto** will detect the incoming image layout from the metadata and image frame aspect.
- **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** 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** is a single 2:1 image in equirectangular mapping.
- **Enable Mirror Window** will show a window displaying the headset user's live view.

Stereo

Mode

- **Mono** will output a single non stereo eye.
- **Auto** will detect the method with which the stereo images are stacked.
- **Vstack** stereo images are stacked vertically as left on top and right at the bottom.
- **Hstack** stereo images are stacked horizontally as left and right.
- **Swap eyes** will swap the eyes if stereo is reversed.

3D

Lighting

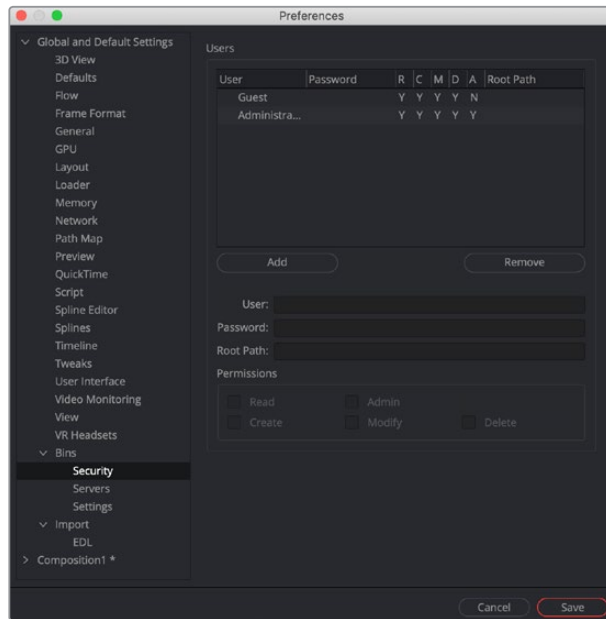
- **Disabled** lighting is off.
- **Auto** will detect if lighting is on in the view.
- **On** will force lighting on in the VR view.

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

This preference is 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.

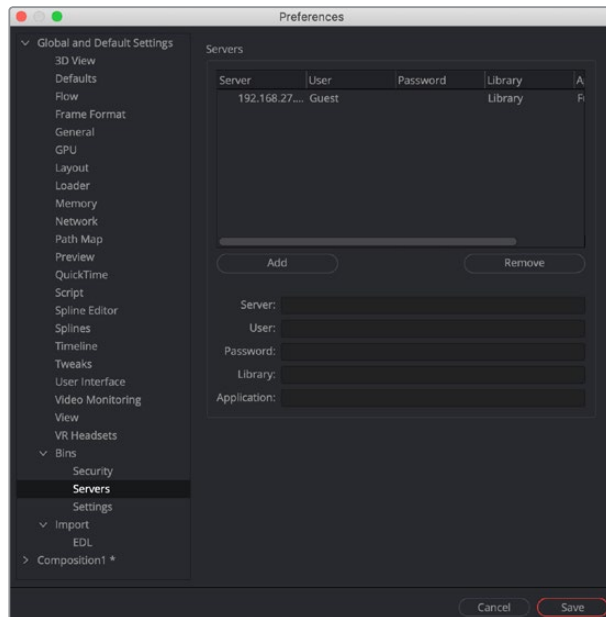
Permissions

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

This preference is 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 on 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

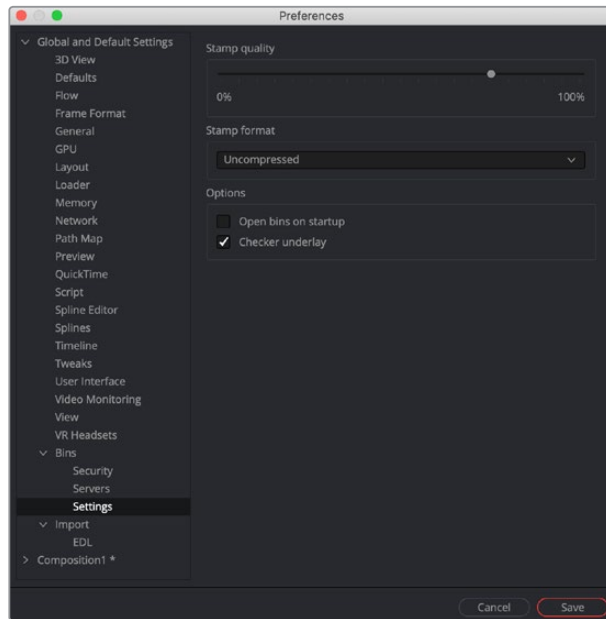
The Library field lets you name the bins. If you wanted to create a bin for individual projects, you would name it in the Library field and each project would get its own bin.

Application

The Application field allows larger studios to specify some other program to serve out the Bin requests.

Bins/Settings

This preference is 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

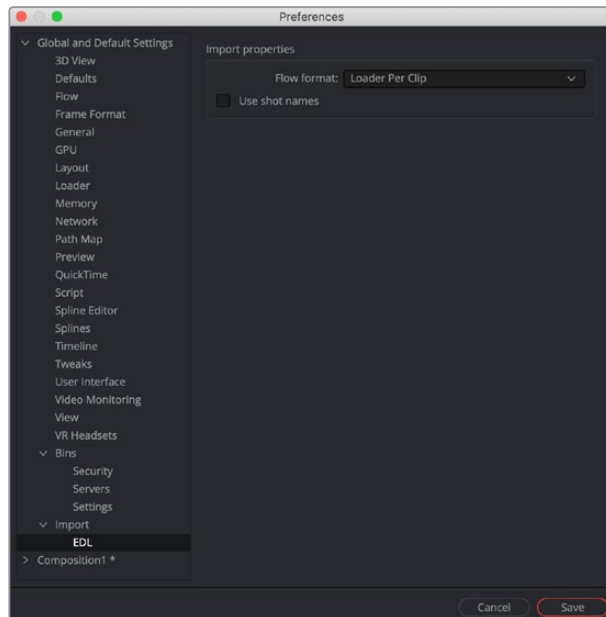
This drop-down list determines whether the Stamp thumbnails will be saved as compressed or uncompressed.

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 Check 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.

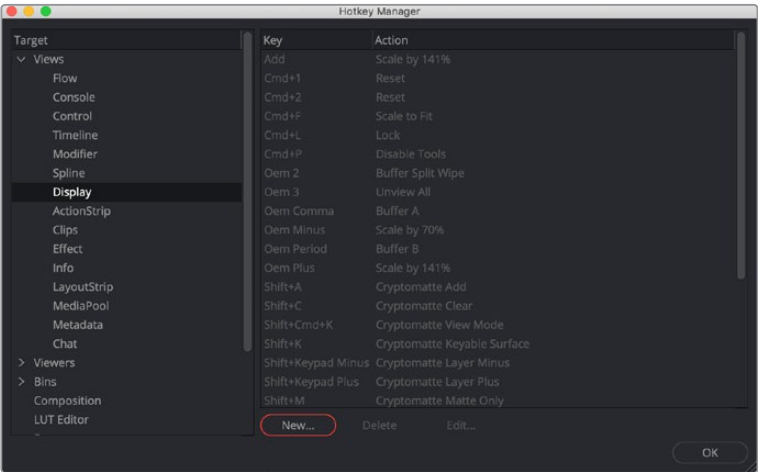
Use Shot Names

When checked, shot names stored in the EDL file are used to locate the footage.

Customization

Shortcuts Customization

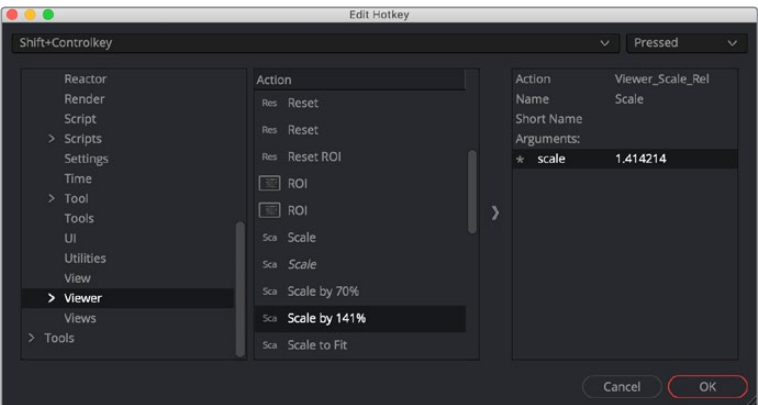
Keyboard shortcuts can be customized in Fusion. You can access the Hotkey Manager from the View, menu.



Hotkey Manager

Fusion has active windows to focus attention on those areas of the interface, like the Node Editor, the viewers, and the Inspector area. When selected, a grey border line will outline that section. The shortcuts for those sections will only work if the region is active. For example, Ctrl+F in the View will scale the image to fit the view area; in the flow view, Ctrl+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.



Edit Hotkey

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.



PART 2

The Fusion UI

Chapter 6

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.

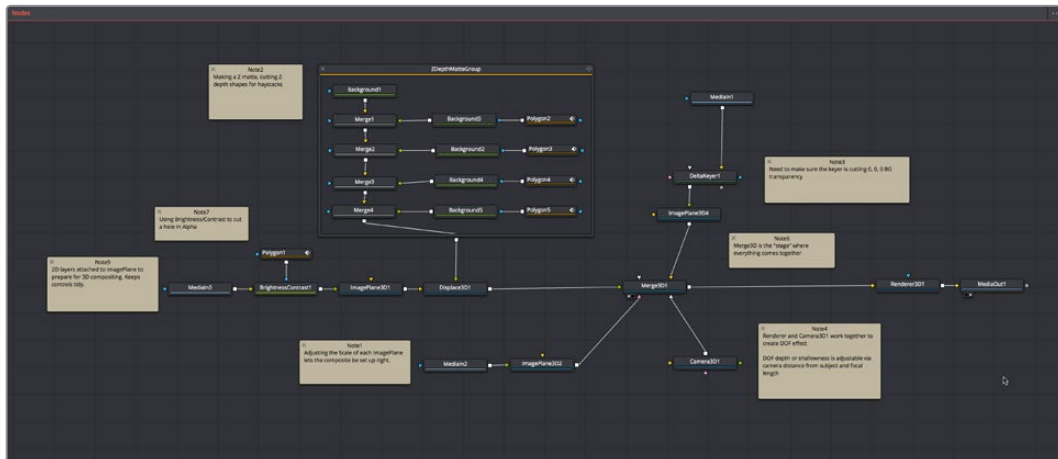
Contents

Learning to Use the Node Editor	128
Navigating within the Node Editor	128
Automatic Node Editor Navigation	129
Using the Node Navigator	129
Adding Nodes to a Composition	130
Adding, Inserting, and Replacing Nodes Using the Toolbar	130
Adding Nodes Quickly Using the Select Tool Window	131
Adding Nodes from the Effects Library	132
Adding, Inserting, and Replacing Nodes Using the Contextual Menu	135
Deleting Nodes	136
Disconnected Nodes	136
Selecting and Deselecting Nodes	136
Selecting Nodes	136
The Active Node	137
Deselecting Nodes	137
Loading Nodes into Viewers	137
Viewed Nodes When You First Open Fusion	138
Node View Indicators	138
Drag and Drop Nodes into a Viewer	139
Using the Contextual Menu	139
Clearing Viewers	139
Create/Play a RAM Flipbook	139
Connecting and Disconnecting Nodes	140
Node Basics	140
How to Connect Nodes	141
Identifying Node Inputs	143
Node Order Matters	143
Upstream and Downstream Nodes	144
Disconnecting and Reconnecting Nodes	145
Tracing Connections Through the Node Tree	146
Branching	146
Connecting Merge Nodes	147
Automatically Creating a Merge Node When Adding Nodes	148
Automatically Creating a Merge Node by Connecting Two Outputs	148
Connection Options and Routers	149
Using Routers to Reshape and Branch Connections	149

Swapping Node Inputs	151
Extracting and Inserting Nodes	151
Cut, Copy, and Paste Nodes	152
Cut, Copy, and Paste in the Node Editor	152
Pasting Node Settings	153
Copying and Pasting Nodes To and From Any Text Editor	153
Instancing Nodes	154
De-Instancing and Re-Instancing Specific Parameters	156
Keeping Node Trees Organized	156
Moving Nodes	156
Renaming Nodes	157
Changing Node Colors	158
Using Sticky Notes	158
Using Underlay Boxes	159
Node Thumbnails	160
Choosing Which Nodes Show Thumbnails	161
Switching Thumbnails Between Images and Icons	161
Finding Nodes	162
Performing Simple Searches	162
Using Regular Expressions	163
Custom Node Settings	164
Managing Saved Settings	164
Resetting Defaults	164
Saving and Loading Alternate Node Settings	165
Adding Saved Settings From the File System	165
Node Modes Including Disable and Lock	166
Node Editor Options	166
Node Tooltips and the Status Bar	167

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.

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.

Methods of zooming the Node Editor:

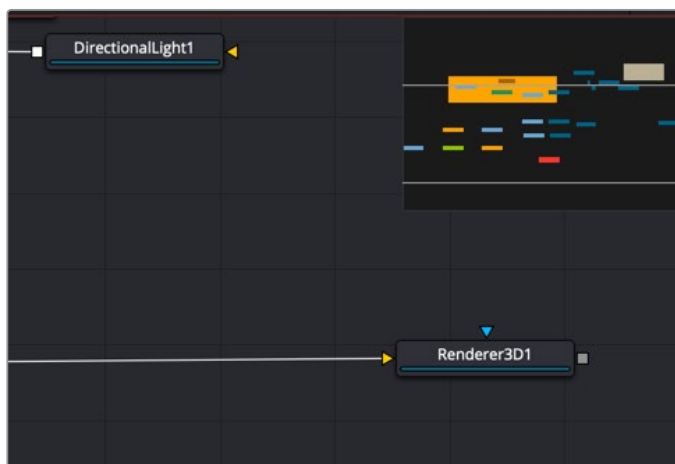
- Press the Middle and Left buttons simultaneously and drag to resize the Node Editor.
- Hold down the Command key and use your pointer's scroll control to resize the Node Editor.
- Right-click the Node Editor and choose an option from the Scale submenu of the contextual menu.
- Press Command-1 to reset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

Automatic Node Editor Navigation

If a node 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:

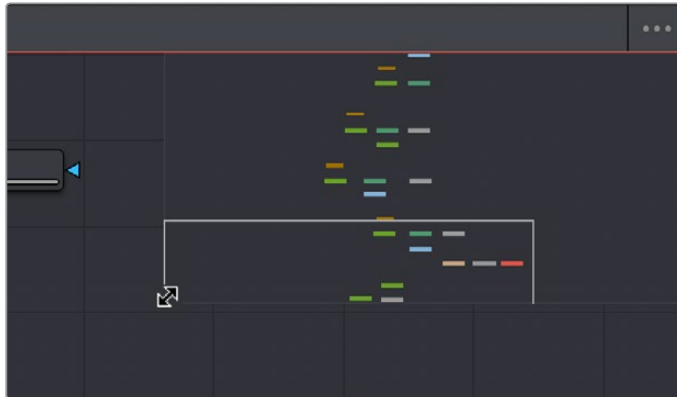
- Right-click in an empty area of the Node Editor, and then choose Options > Show Navigator.
- Press the V key.

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 the 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.

Adding Nodes to a Composition

You can add nodes to the Node Editor in a variety of different 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 most any 2D operation. However, if you tried 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 a Loader, 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 toolbar.

TIP: If you don't know which node a particular icon corresponds to, you can hover the 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 will 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.

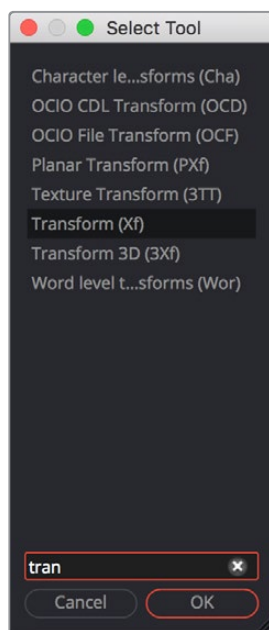
TIP: When you replace one node with another, any settings that are identical between the two nodes are copied into the new node. For example, replacing a Transform node with a Merge will copy the existing center and angle values from the Transform to the Merge.

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:
 - a. 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.
 - b. If you want to create a disconnected node, then 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).
- 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.

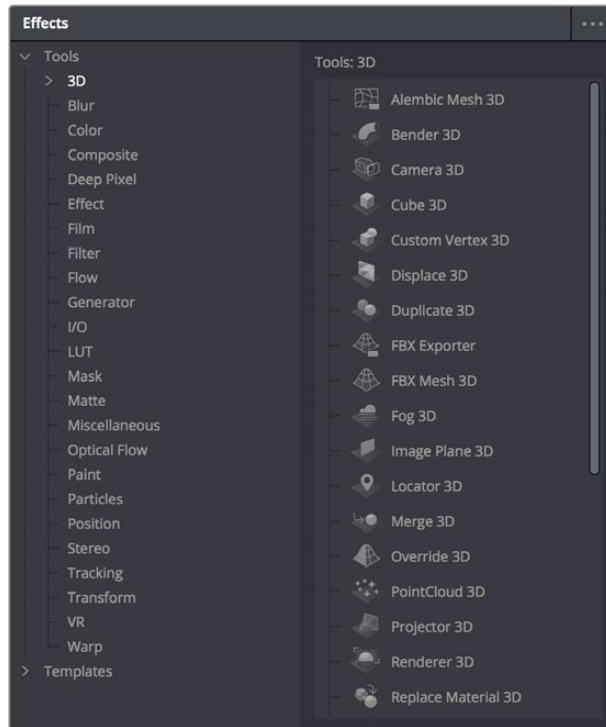
Adding Nodes from the Effects Library

While the toolbar shows many of the most common nodes you'll use 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 Fusion.

The Effects Library appears at the upper-left corner of Fusion and consists of two panels. A category list at the left shows all categories of nodes and presets that are available, and a browser 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 two primary sets of effects, Tools and Templates, with disclosure controls to the left that hierarchically show all subcategories within each category. The top two 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. If you have third-party OFX plug-ins on your workstation, those appear in here as well.
- **Templates:** The 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.

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.

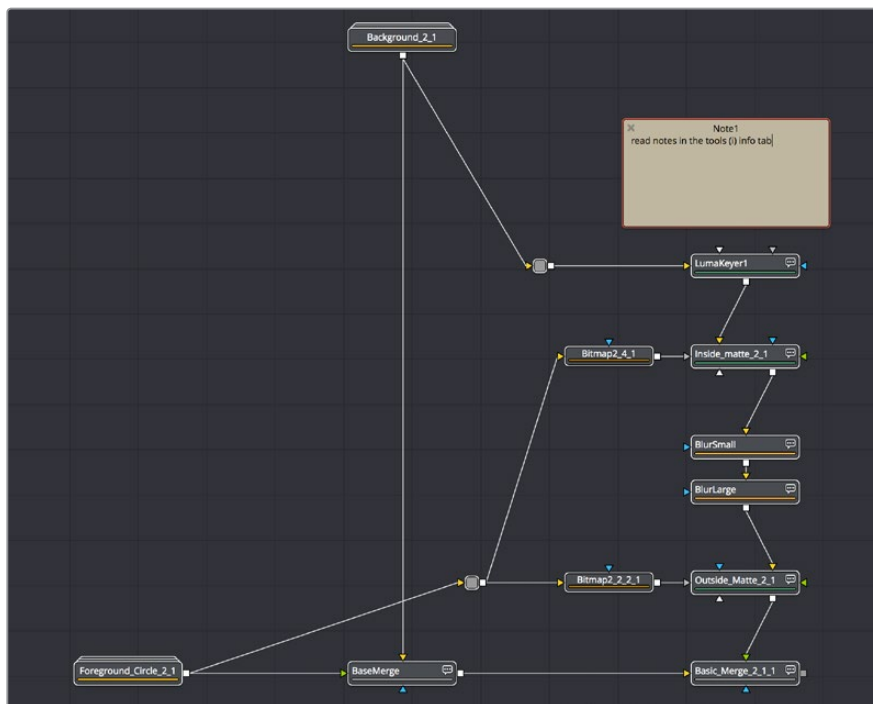
Adding, Inserting, and Replacing Templates Using the Bin

Adding items from the Templates category in the Bin 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 of the Bin, 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 Bin.

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 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 deleting the tree, 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 decide that you don't need them. 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.

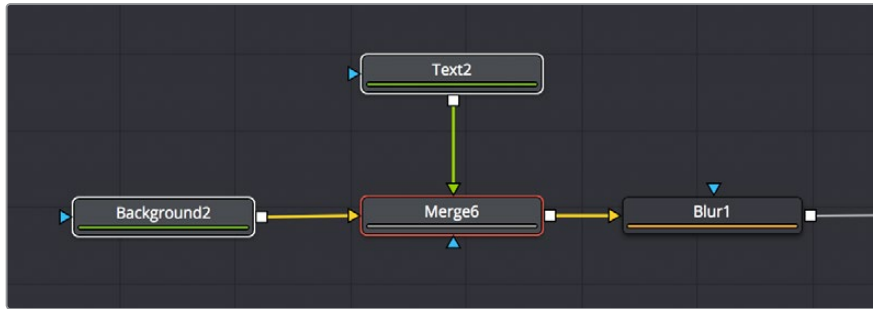
Methods of selecting nodes:

- **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. So you can tell 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 mostly as you would expect.

Methods of deselecting nodes:

- Click once in the background of the Node Editor 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 on-screen 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 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 add Loaders or any node to a new project, nothing is displayed in the viewers by default. There are a number of 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. Secondly, 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 in Fusion. The left and right indicators correspond to the left and right viewers, regardless of whether both viewers are visible or just one is visible.



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 cursor 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.

On 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 cursor 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 a keyboard oriented person, 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

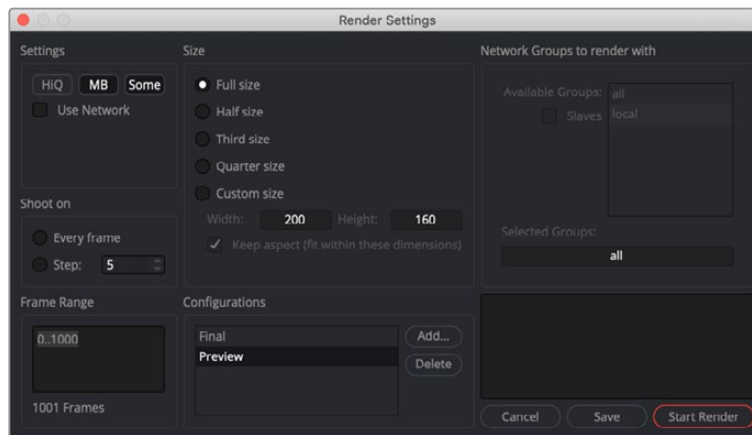
Whenever you load a node into a viewer, you require that node, all upstream nodes, and other related nodes to be rendered. If you load nodes into both viewers, this is doubly true. If you want to prevent your computer from processing views that aren't currently necessary, you can clear each viewer.

Methods of clearing viewers:

- Press 1 or 2 to empty the left or right viewers if they're filled.
- Press ` (the tilde key) to empty both viewers.

Create/Play a RAM Flipbook

RAM Flipbook Previews are preview renders that exist entirely within RAM. To create a RAM Preview, you right-click a node you want to preview, and choose Create/Preview Play On Left/Right View, depending on the viewer where you want the preview to play. After making a selection, a render dialog will appear where you can change the settings that determine the quality and resolution of the preview.



The RAM Preview dialog.

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.


For more detail on RAM Flipbook Previews, see Chapter 8, “Using Viewers.”

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 shapes around the edges. One or more arrows represent inputs, and the square represents 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.

An output of a node connects to an Input of another node for image processing. For instance, by connecting a Loader node’s output to a Transform node, you move image data from the Loader node to the Transform node, which does something to process the image before the Transform 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 one, click on one node's output and drag a connection line out to drop on another node's input. The direction 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.



Dragging a connection line 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 body 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 Loader node and dropped onto the body of a Defocus node, and the background input is connected first.



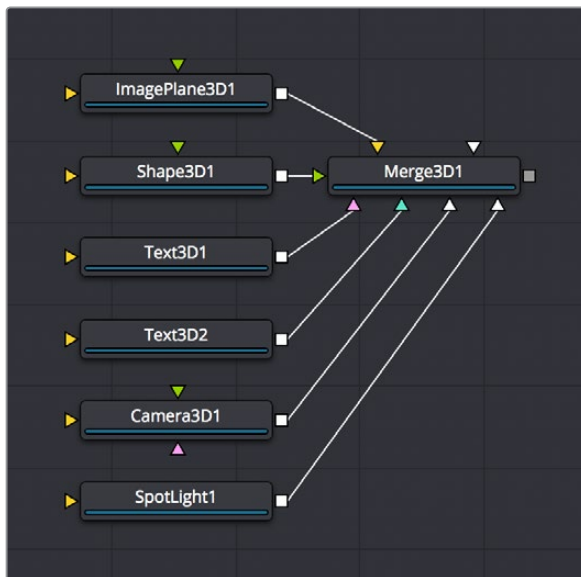
Before (top) and after dragging a connection line and dropping it on top of a node (bottom).

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 (top) and after dragging a connection line and dropping it on top of a node that has the background input already connected (bottom).

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.

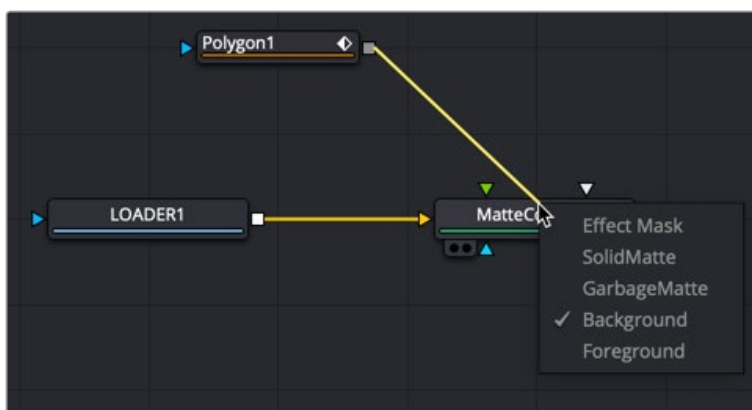


After dragging a connection line and dropping it on top of a Merge3D node.

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 somehow limit the node's effect. 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 dragging a connection from a mask node and dropping it on top of a MatteControl node (right).

Identifying Node Inputs

While you are still figuring out all the nodes and their inputs, hovering the cursor over any knot will display a node tip with the knot's name.

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 Loader 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 Saver node.



Adding a Defocus effect first, and then the TV node.

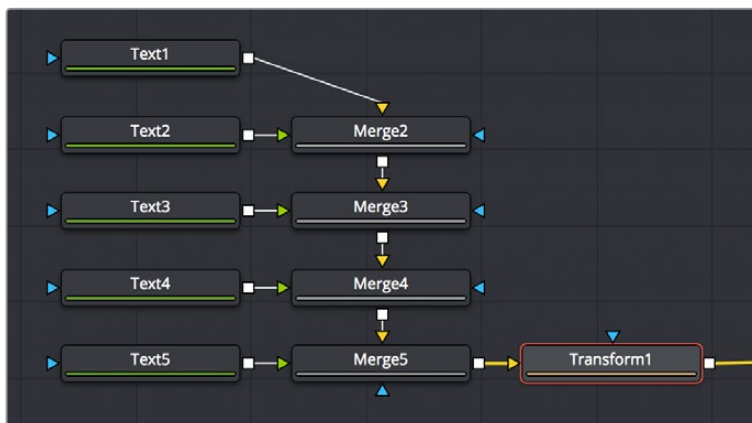
As you can see above, connecting the Defocus node first, followed by the TV node, means that while the initial image is softened, the TV effect is sharp. However, if you reverse the order of these two nodes, then the TV effect distorts the image, but the Defocus node now blurs the overall result, so 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 then Defocus second.

As you can see, the node tree that makes up 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 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 (at left) and downstream (at 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 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, and 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:

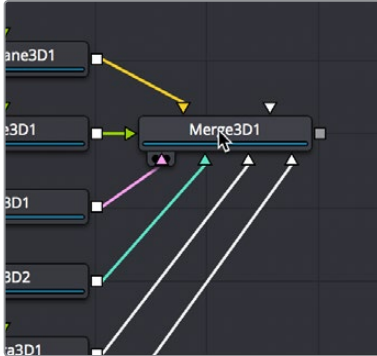
Drag the output or input half of a connection, and drop it directly onto another node's input or output. This simultaneously disconnects the previous connection and connects the one you're dragging.

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 are attached. Additionally, highlighted connection lines display the color of the input to which they are connected. This 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.

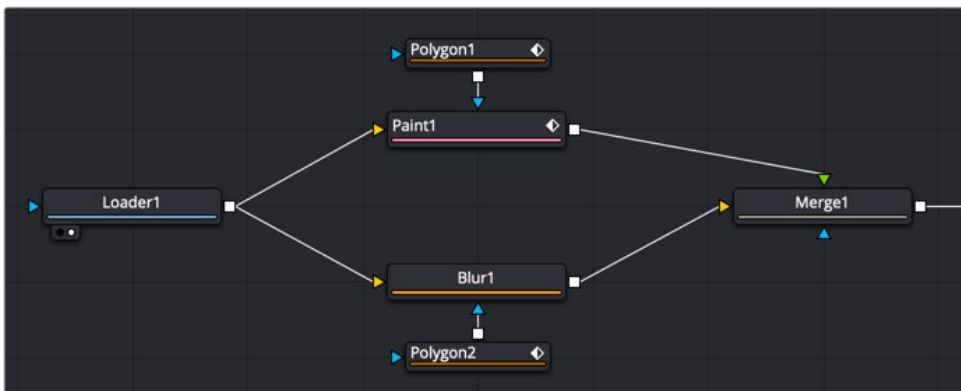
Additionally, positioning the pointer over a connection line causes a tooltip to appear that displays the output and input that connection is attached to.



Hovering the pointer over a node highlights the connection between it and other nodes.

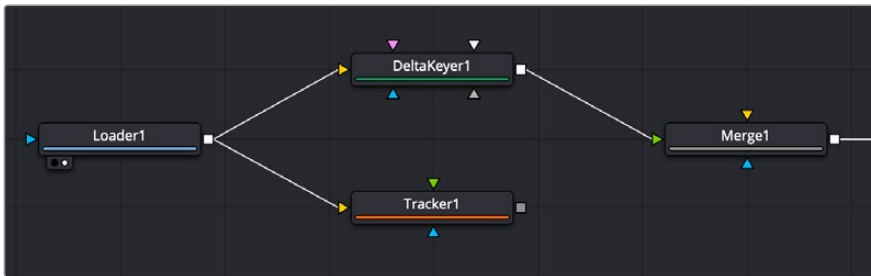
Branching

Any single node input can have only one connection attached to it. However, a tool's output can be connected to many different inputs. 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 Loader node branched to two node operations and then recombined using a Merge node.

Alternately, 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 Loader 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. This node is covered more extensively elsewhere, but for now all you need to know is that if you attach a background image to the yellow Background input (such as a landscape), and a foreground image with an alpha channel to the green Foreground input (such as a graphic with an alpha channel), the Merge node will combine them into a single image for further compositing.



Two Loader nodes and a DeltaKeyer node attached to the foreground input of a Merge node creating a composite.

Each Merge node has three inputs:

- **Background (green):** The default input. Whichever image is connected to this input defines the output resolution of the Merge node.
- **Foreground (yellow):** 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 background input is automatically 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 Loader immediately to your node tree as the top foreground layer of a composite, and that's to select the node that will be the background input to your merge, and then click the Loader tool from the Toolbar to import the foreground.


This automatically creates a Merge node, the background is the selected node, and the foreground input is the new loader you've just added.



The selected node is the background to a merge (left).
The added Loader becomes the foreground to a merge (right).

Automatically Creating a Merge Node by Connecting Two Outputs

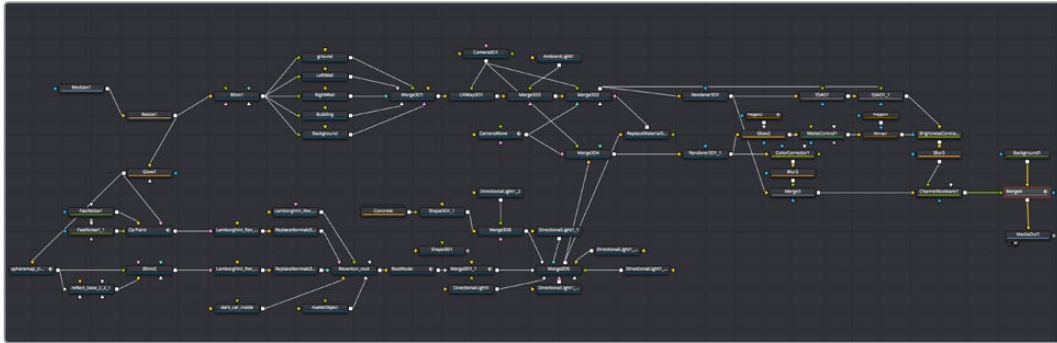
Here's an endlessly useful shortcut for when you have a disconnected node that you want to composite against another node. Drag a connection from the output of the node you want to be the foreground layer, and drop it on top of the output of the node you want to be the background layer, and a Merge node will be automatically created to build that composite.



Dragging a connection from a disconnected node to another node's output (left),
and dropping it to create a Merge node composite (right).

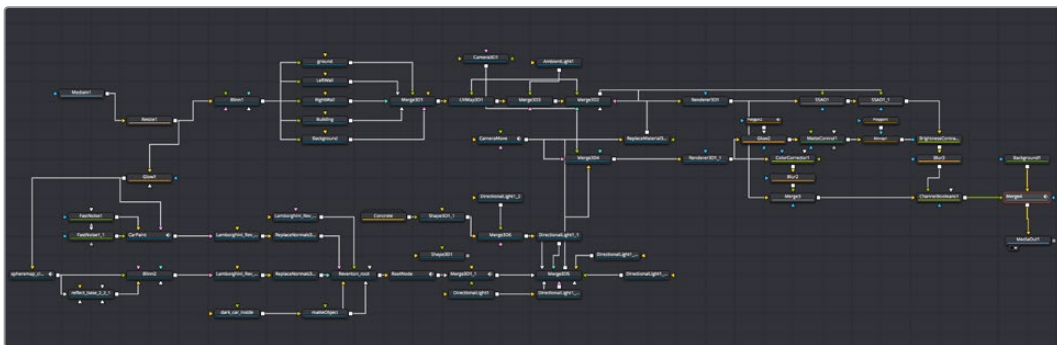
Connection Options and Routers

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.

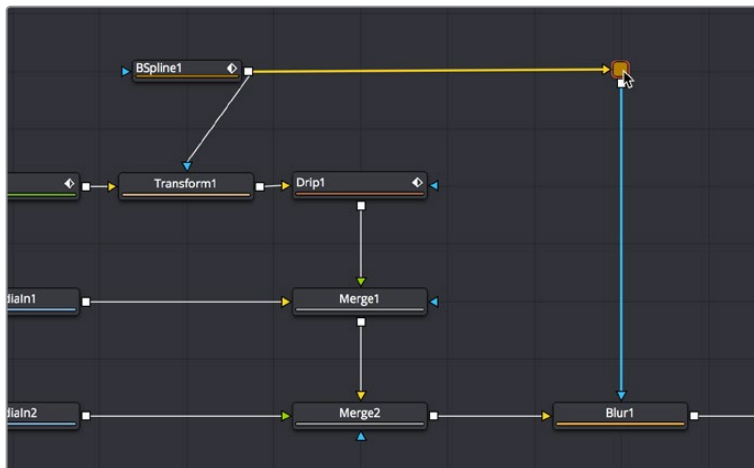
Functionally, there's no difference to your composition; this only affects how your node tree appears.

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

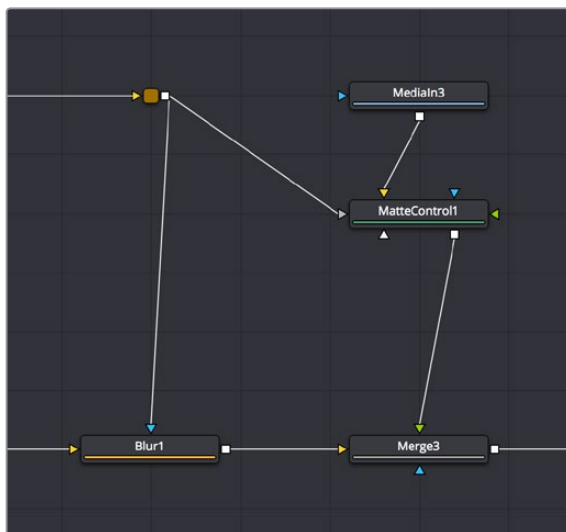
If you want to force a particular connection to be drawn at an angle, to keep your node tree tidy, you can add a router to either linear or orthogonal connections to force an angle so it will be drawn however you like.



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.

You can also 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 needed.
- **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 (left), and after swapping node inputs (right); connections don't move but the colors change.


Because inputs in Fusion can freely move around the node, swapping two inputs doesn't move the connection lines, instead the inputs change color to indicate you've reversed the background (yellow) and foreground (green) connections.

Extracting and Inserting Nodes

When building a composition, you'll often find 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 unattached, 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, then 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 unattached (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 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 node(s) 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 node(s) 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 a disconnected node(s) in a specific area of the Node Editor:** Deselect all nodes, then click the place in the Node Editor where you want pasted nodes 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 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 Default Preferences panel in the Global preferences.

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.

You should 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 in fact a simple text format. The implications of that may not have been obvious, but one benefit example 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 as script 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:

```
{
  Tools = ordered() {
    MediaIn3 = MediaIn {
      ExtentSet = true,
      CustomData = {
        MediaProps = {
          MEDIA_HEIGHT = 1080,
          MEDIA_MARK_IN = 0,
          MEDIA_MARK_OUT = 244,
          MEDIA_NAME = "86803838-aerial-abstract-waves-crystal-PRORESHD1080.mov",
          MEDIA_NUM_FRAMES = 245,
          MEDIA_NUM_LAYERS = 1,
          MEDIA_PAR = 1,
          MEDIA_PATH = "/Volumes/Media Raid Too/DaVinci Media/Fusion Manual Example Media/Pond5 Media/Textures/86803838-aerial-abstract-waves-crystal-PRORESHD1080.mov",
          MEDIA_SRC_FRAME_RATE = 25,
          MEDIA_START_FRAME = 0,
          MEDIA_WIDTH = 1920
        },
      },
    },
    Inputs = {
      GlobalOut = Input { Value = 244, },
      MediaID = Input { Value = "40086dc-2c6f-4fa8-be00-52729b48f6d5", },
      Layer = Input { Value = "", },
      ClipTimeEnd = Input { Value = 244, },
    },
    ViewInfo = OperatorInfo { Pos = { 825, 115.5 } },
    Merge3 = Merge {
      Inputs = {
        Background = Input {
          SourceOp = "MediaIn3",
          Source = "Output",
        },
        PerformDepthMerge = Input { Value = 0, },
      },
      ViewInfo = OperatorInfo { Pos = { 825, 181.5 } },
    },
    Glow1 = Glow {
      CtrlWZoom = false,
      Inputs = {
        Blend = Input { Value = 0.2, },
      },
      ViewInfo = OperatorInfo { Pos = { 935, 181.5 } },
    },
  }
}
```

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.

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 need to have the exact same 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.

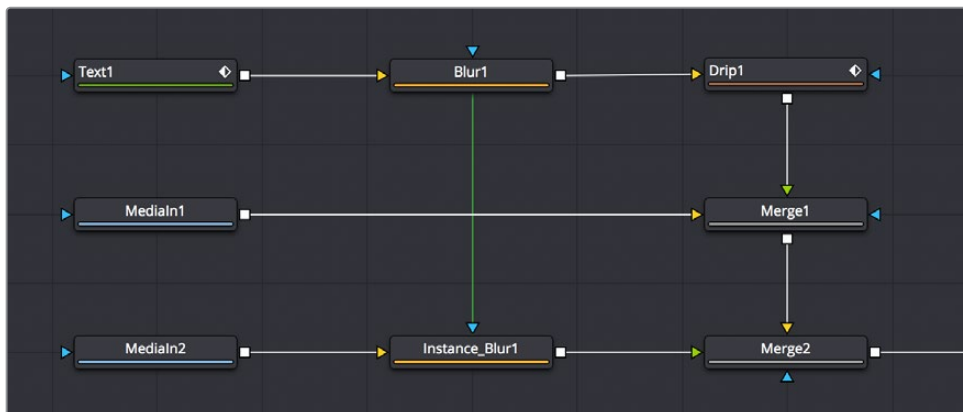
Using Instanced Nodes

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. In addition, a green outline surrounds the edit boxes in the Inspector of an instanced tool. You have the option to hide the green link lines in the Node Editor to reduce visual clutter.

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.

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

Moving nodes around is one of the most basic ways you can stay organized in the Node Editor, simply by dragging nodes and collections of nodes around the Node Editor to better arrange them for future connections, as well as to improve visual organization. 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, if you're really into straight node trees, you can also have nodes you're dragging automatically snap to the grid, making it easier to keep things aligned.

To have nodes snap 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, and turning on the Fusion > Node Editor > Arrange To Grid or Arrange to Connected checkboxes.

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 the number of nodes that have already been created of a given node type).

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 can 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 will automatically be replaced with underscores (_) and invalid characters will be automatically deleted.

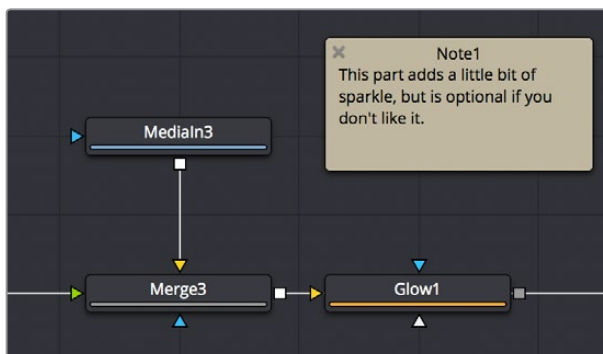
Changing Node Colors

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.

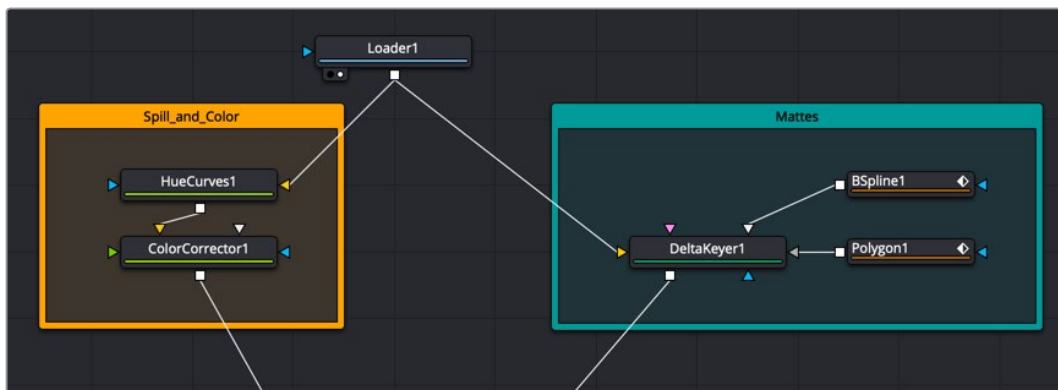
Methods of working with Sticky Notes:

- **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 > Flow 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, 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, and then 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 set the color of a Sticky Note:** Right-click a Sticky Note, choose Set Color, and then choose the color from the submenu.
- **To delete a Sticky Note:** Click any Sticky Note to select it, 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 will move along with it.



Underlays 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.

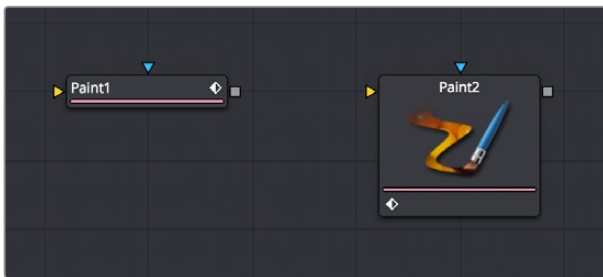
Methods of working with Underlay Boxes:

- **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 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 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 of those nodes.
- **To move an Underlay Box and all of its nodes:** Once nodes have been placed inside an Underlay Box and are 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, then press the Delete key. The nodes within remain where they were.

Node Thumbnails

Once a clip 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 both with and without 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 of 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 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 tell apart 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 or not 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 Timeline 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 result in compositions that have 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 in 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 or not 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 of 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 in the Find dialog, which lets you enter some simple expressions with which to create more complex find operations. Some useful examples of valuable regular expressions 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; will find nodes with a, b, c, or d

[Tt]

Finds: Every node with an uppercase T or a lowercase t

[aeiou]

Finds: Every vowel

[0-9]

Finds: Every numeral

[5-7]

Finds: Every numeral from five to seven (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. There are two ways of doing this.

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.

Managing Saved Settings

Custom node default settings are saved to a directory on your hard drive that's based on the Path Map > Defaults preference in the Fusion Settings. This path is customizable, so you can save them somewhere that's commonly accessible to multiple Fusion artists. The default paths are:

- For macOS systems, this path defaults to /UserName/Library/Application Support/Blackmagic Design/Fusion/Defaults.
- For Windows systems, this path defaults to C:\Users\<username>\AppData\Roaming\Blackmagic Design\Fusion/Defaults.
- For Linux systems, this path defaults to ~/.fusion/BlackmagicDesign/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 plug-in 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 Control Panel, 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, 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 displays a series of commands in the Modes submenu, some with accompanying keyboard shortcuts that let you set control visibility, and 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 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.
- **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 either be built 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 if connections between nodes are drawn as Direct (straight) lines or Orthogonal (bent) lines.

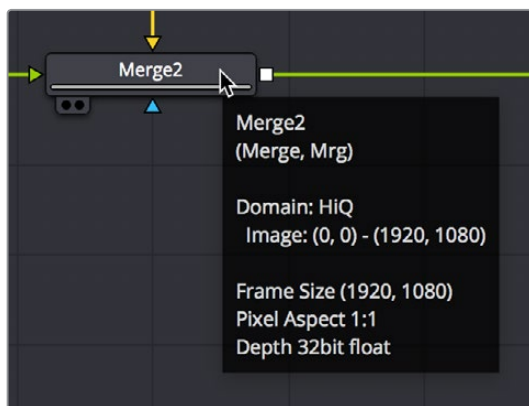
Node Tooltips and the Status Bar

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.

Merge2 - Frame Size (1920, 1080), Pixel Aspect 1:1, Depth 32bit float

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 provides 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 7

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	170
Creating Groups	170
Deleting Groups	170
Expanding and Collapsing Groups	170
Panning and Scaling within Open Group Windows	171
Ungrouping Nodes	171
Saving and Reusing Groups	172
Macros	172
Creating Macros	173
Using Macros	173
Re-Editing Macros	173
Other Macro Examples	174
Creating Fusion Templates	174
Getting Started Creating a Fusion Template	174
Create a Macro	175
Open DaVinci Resolve and Use Your New Template	177

Groups

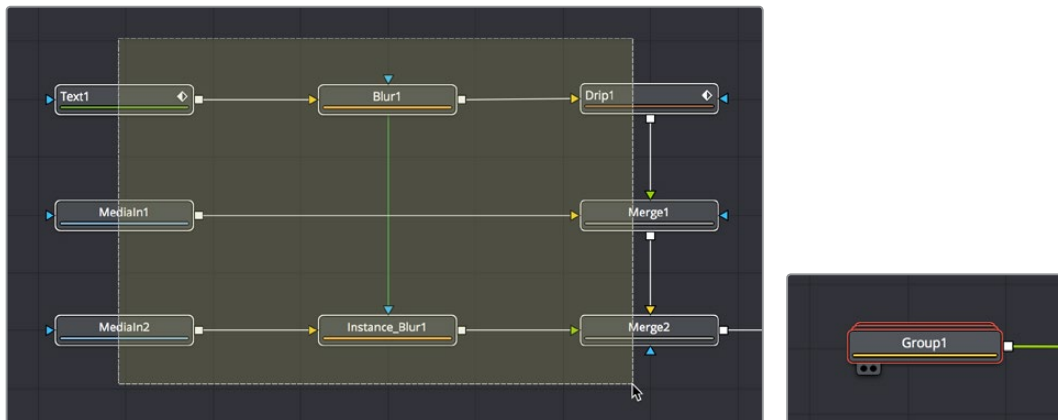
When you work on complex visual effects, node trees can become sprawling and unwieldy, so grouping items together can help you better organize all the nodes and pipes. Groups are containers in your node tree that can hold multiple nodes, similar to the way a folder on your Desktop holds multiple files. 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 (Left), and the resulting group (right).

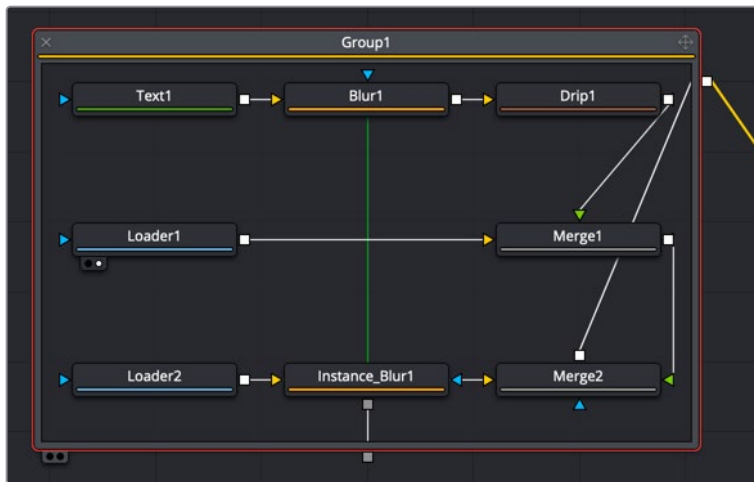
The selected nodes are collapsed into a group, which is displayed as a single node in the Node Editor. The Group node can have inputs and outputs, depending on the connections of the nodes within the group. The Group node only displays inputs for nodes that are already connected to nodes outside the 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 an expanded group.

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.

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.



An open group window showing the Position button.

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.
- **To add a group to a Bin:** Save the group as a setting, then drag the setting file from your computer's file browser.

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.

If you want to control the order in which each node's controls will appear in the macro you're creating, then 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.
- 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. On macOS, the Macintosh HD/Library/Application Support/Blackmagic Design/Fusion/Macros/ directory lets that macro appear in the Effects Library and the Macro submenu of the contextual menu.
- 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/Fusion/Macros.

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.

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

Fusion has the ability to create Title templates for DaVinci Resolve. Creating 2D or 3D animated titles in Fusion can be saved as a setting and show up as an editable title template in

Creating Fusion Templates

Fusion has the ability to create Title templates for DaVinci Resolve. Creating 2D or 3D animated titles in Fusion can be saved as a setting and show up as an editable title template in DaVinci Resolve. This section shows you how it's done.

Getting Started Creating a Fusion Template

The first part of creating a Fusion 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.

Getting Started Creating a Fusion Template

The first part of creating a Fusion 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.

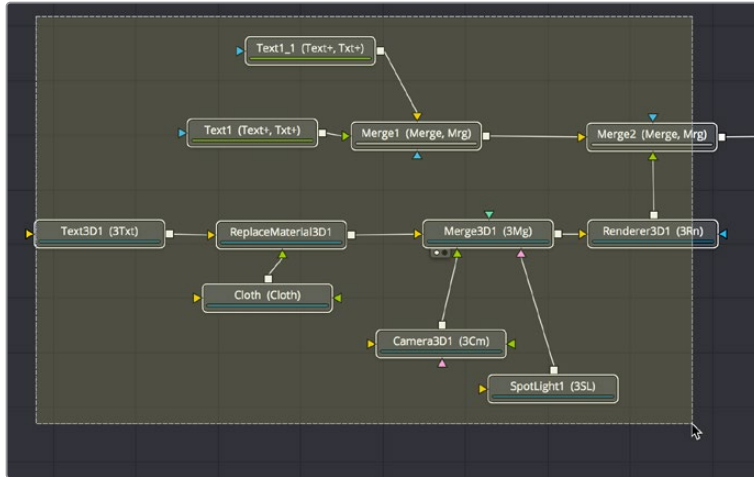
The screenshot displays the Blender 2.80 interface. The top-left 3D Viewport shows a perspective view of a scene with the word "CHEESE" rendered in large, golden, serif letters on a dark grey floor. The top-right 2D Viewport shows a checkerboard background with the text "The CHEESE Show" in a similar golden font. The bottom half of the screen features the Outliner and Dope Sheet editors. The Outliner lists objects like "Text_1 (Plain Text)", "Merge (Merge_Mat)", and "MediaOut (Mat)". The Dope Sheet editor shows a timeline from 0:00 to 1:21 with various keyframes and animation curves. On the right side, the Properties panel is open for the "Text" object, displaying settings for "Stylized Text", "Font" (Oswald Plain), "Color" (white), "Size" (1.0), "Tracking" (0.0), "Line Spacing" (1.0), "V-Anchor" (Top), "H-Anchor" (Left), "Justify" (Left), "Direction" (Automatic), "Line Direction" (Automatic), "Wrap On" (Start), "Extrusion" style (Classic), "Extrusion Depth" (0.0095), "Bevel Width" (0.0), "Smoothing Angle" (16.0), and "Back" button.



Create a Macro

Using Fusion's macro functionality, we can also build templates for DaVinci Resolve.

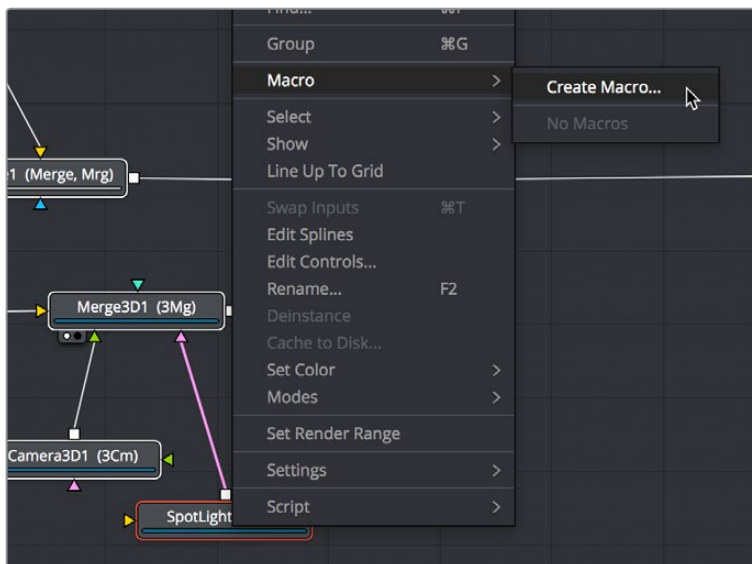
Having built your composition in Fusion, select every single node you want to include in that template except for the Saver node.



Selecting the nodes you want to turn into a template.

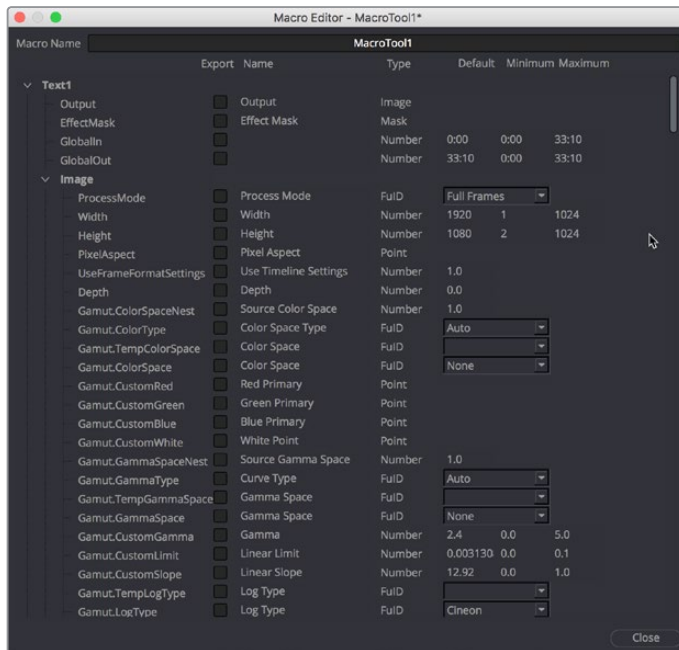
TIP: If you want to control the order in which node controls will be displayed later on, you can Command-click each node you want to include in the macro, one by one, in the order in which you want controls from those nodes to appear. This is an extra step, but it keeps things better organized later on.

Having made this selection, right-click one of the selected nodes and choose Macro > Create Macro from the contextual menu.



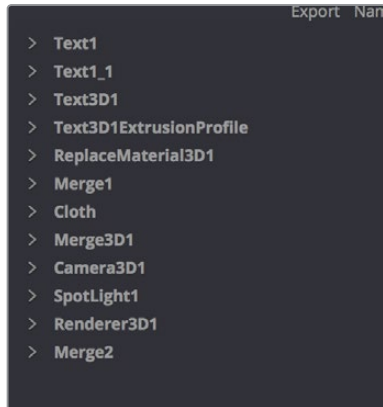
Selecting the nodes you want to turn into a template.

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.

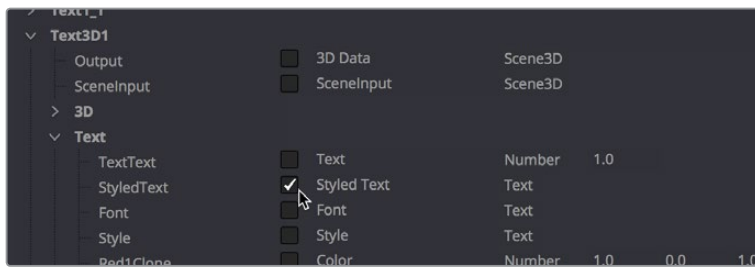
This list may look intimidating, but closing the disclosure control of the top Text1 node shows us what’s really going on.



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).



Turning on the checkboxes of parameters we'd like to edit when using this as a template.

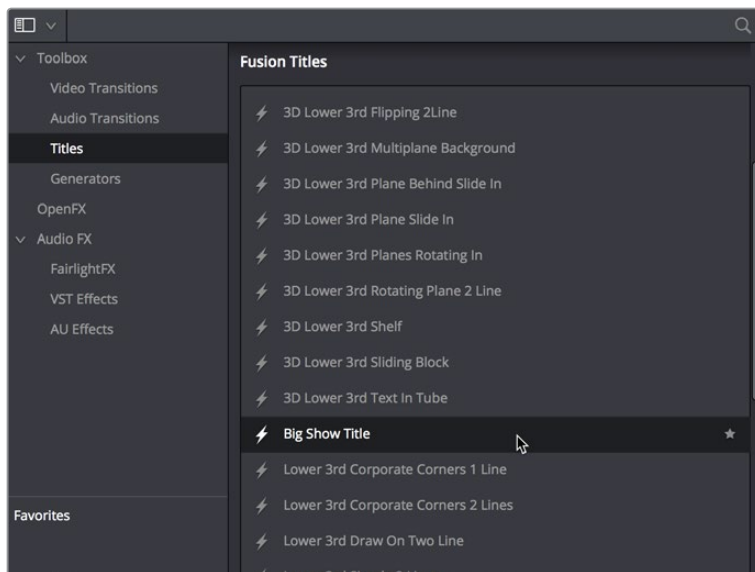
Once we've turned on all the parameters we'd like to use in the eventual template, we click the Close button, and a Save Macro As dialog appears. You can save the setting someplace easy to locate on your computer and then copy it over to a computer that has Davinci Resolve installed. Where you place the .settings file on the DaVinci Resolve computer is key to having it show up correctly. If we're using macOS, we navigate to the /Library/Application Support/Blackmagic Design/DaVinci Resolve/Fusion/Templates/Edit/Titles directory, and copy the .settings file to that location. On Windows, navigate to C:\Program Files\Blackmagic Design\DaVinci Resolve\Fusion\Templates\Edit\Titles.



Choosing where to save and name the Macro.

Open DaVinci Resolve and Use Your New Template

After you've copied your macro to the DaVinci Resolve computer, you'll need to quit and re-open DaVinci Resolve. When you open the Effects Library of the DaVinci Resolve Edit page, you will 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 DaVinci Resolve 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!

Chapter 8

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	182
Single vs. Dual Viewers	182
Video Output	182
Loading Nodes Into Viewers	183
Clearing Viewers	183
Position and Layout	184
The Viewer Divider	184
Zooming and Panning into Viewers	185
Flipbook Previews	185
Creating Flipbook Previews	185
Playing Flipbook Previews	186
Removing Flipbook Previews	187
Flipbook Preview Render Settings	187
Onscreen Controls	188
Showing and Hiding Onscreen Controls	189
Making Fine Adjustments to Onscreen Controls	189
Toolbars	189
Viewer Toolbar	189
Node Toolbars	190
A/B Buffers	190
Flipping Between Buffers	190
Split Wipes Between Buffers	191
Moving the Wipe Divider	192
Subviews	192
Showing and Hiding Subviews	192
Changing the Subview Type	193
Swapping the Subview with the Main View	193
Viewer and Subview Types	193
Viewing Selective Channels	197
Viewing Color Channels	197
Viewing Auxiliary Channels	198

The 3D Viewer	198
Panning and Scaling and Rotating a 3D Viewer	198
Viewing Objects Via Wireframe	199
Changing the POV of a 3D Viewer	199
Changing Cameras in a 3D Viewer	199
Copying a Viewer's POV to a Camera	199
Lighting and Shadows in 3D Viewers	200
Transparency in 3D Viewers	202
Grid	202
Vertex Normals	203
Quad View	203
Quad View Layouts	204
Using Quad Views for 2D Scenes	204
Guides	204
Frame Format Settings	205
Domain of Definition and Region of Interest	206
Domain of Definition (DoD)	206
Region of Interest (RoI)	207
Managing Viewer Lookup Tables (LUTs)	209
How Lookup Tables Work in Fusion	209
Types of Viewer LUTs	210
Using Viewer LUTs	211
Editing Viewer LUTs	211
LUT Processing Order	214
Applying Multiple LUTs	214
Saving Custom LUTs	215
Viewer Preferences and Settings	217
Viewer Settings	217
The Viewer Options Menu	217
Locking the Viewer (Command-L)	218
Additional Viewer Options	218
Status Bar Information	219


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 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 a single viewer being shown, and two viewers being shown side by side.



The single/dual viewer toggle button.

Video Output

Blackmagic Design capture and playback products like DeckLink PCIe cards, can be used for monitoring your work on an HD, Ultra HD or DCI 4K display. You can select the hardware and output format using the Video Monitoring Preferences. The Video monitor acts as a third viewer, so pressing 3 on any node will display that node's output on the video monitor.

Loading Nodes Into Viewers

When you first open Fusion, the viewers will remain empty even after you bring in clips with a Loader or add nodes to the node tree. It is up to you to decide when and in which viewer you want to display an image.

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/LeftView/RightView 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 Tilde (~) 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.

Position and Layout


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 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 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.

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 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 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-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 Option key and drag to spin the stage around.

Flipbook Previews

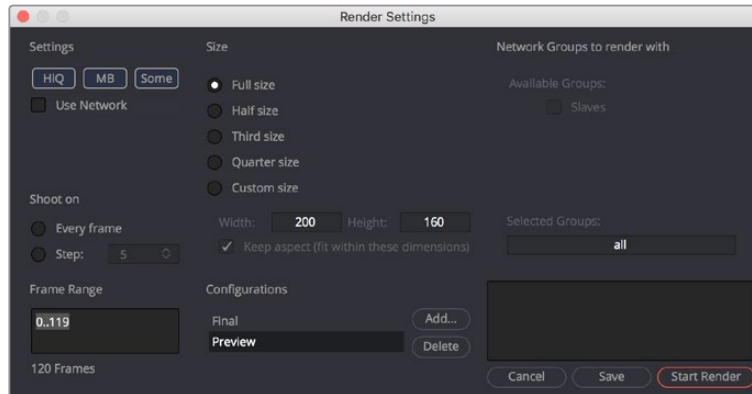
As you build more and more complex compositions, and you find yourself needing to preview specific branches of your node tree to get a sense of how various details you're working on are looking, you may find it useful to create targeted RAM previews at various levels of quality right in the viewer by 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 pretty 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 Shift-Option down 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 fullscreen 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 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 Fullscreen 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 not currently available.

Shoot On

Sometimes you may not want to render every single frame, but 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

This field 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

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

An option designed for the interactive frame-by-frame work of rotoscoping and paint is to set a Preview to Update from its contextual menu. When active, any frames modified on the previewed node are automatically updated in the preview's playback. This lets you reserve the RAM to do 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 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 Angle preview control.

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 Control Panel, 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 of the wipe. 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. It's usually used to show different information about the image.



The Subview menu with the Histogram subview displayed.

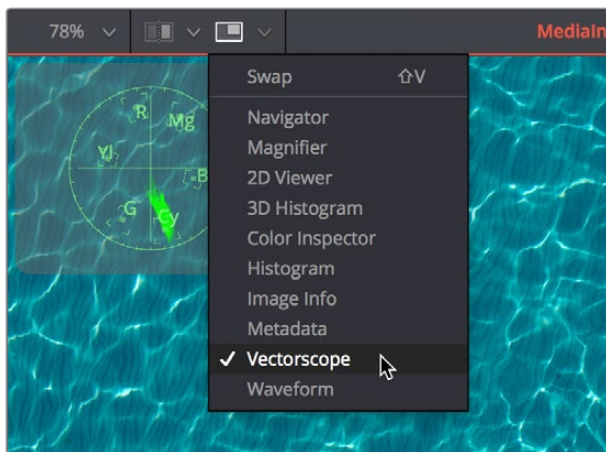
For example, the RGB channels can be viewed in the main viewer, while the alpha channel is displayed in a subview. For the most part, the subview is a fully functional miniature viewer, with its own contextual menu and options. It responds to the same keyboard shortcuts and navigation controls as any 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 popup menu and contextual menu show all the available subview types. Once you choose an option from the list, that is the view that 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 only available as subviews. Below is detailed information about the different view types available.

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.

Histogram

The Histogram is an analysis node that can be used to identify problems with contrast and the 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 titlebar into the viewer.



The Histogram Viewer type for evaluating the contrast and color cast in an image.

3D Histogram

This more advanced histogram type 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.

Vectorscope


The Vectorscope 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 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.

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 only be used 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.

Image Info

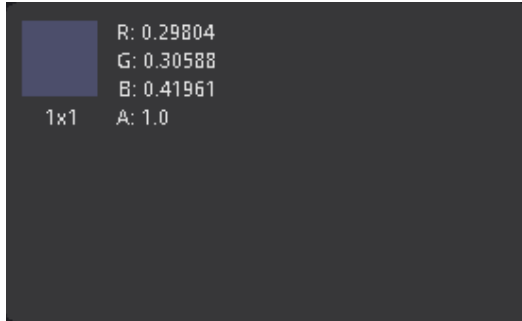
The Image Info view can only be used in a subview. The Image Info tab shows a horizontal bar across the top of the image with information about the frame size, pixel aspect, and color depth of the viewed image.



The Image Info subview for viewing size, pixel aspect, and color depth information.

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.

Metadata

The contents of this subview is entirely based on the amount of metadata in your image. Most Loaders will give the colorspace and file path for the image. Much more information can be displayed if it exists in the image.



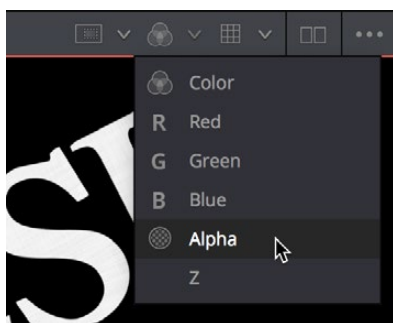
The Metadata subview for viewing embedded metadata.

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.



The Channel toolbar button for switching between color, alpha, and depth channels.

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.

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 the available channels for 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 Open EXR 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 pop-up menu or right-click in the viewer and choose an option from the Channels submenu of the contextual menu.

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 detail on 3D controls, see Chapter 23, “3D Compositing Basics.”

Panning and 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 but then want that point of view 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:

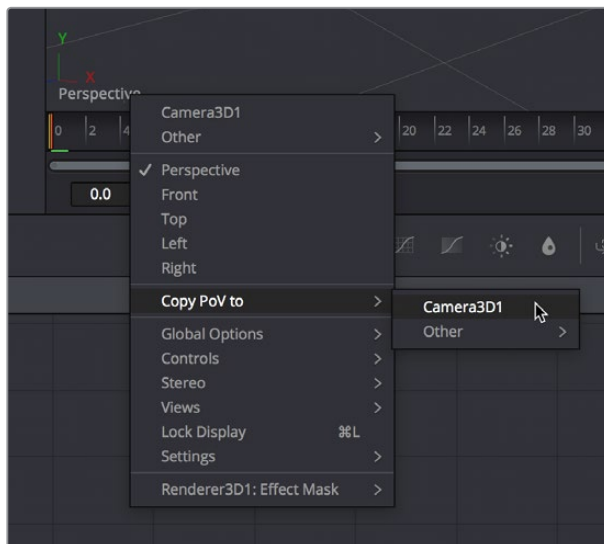
- 1 Set up a 3D Viewer with the point of view you want by zooming, panning, and rotating the Viewer.
- 2 Add a camera to your 3D scene.
- 3 Right-click anywhere within the 3D Viewer, and choose Camera > Copy PoV to > Camera3DNameOfCamera from the contextual menu.

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.

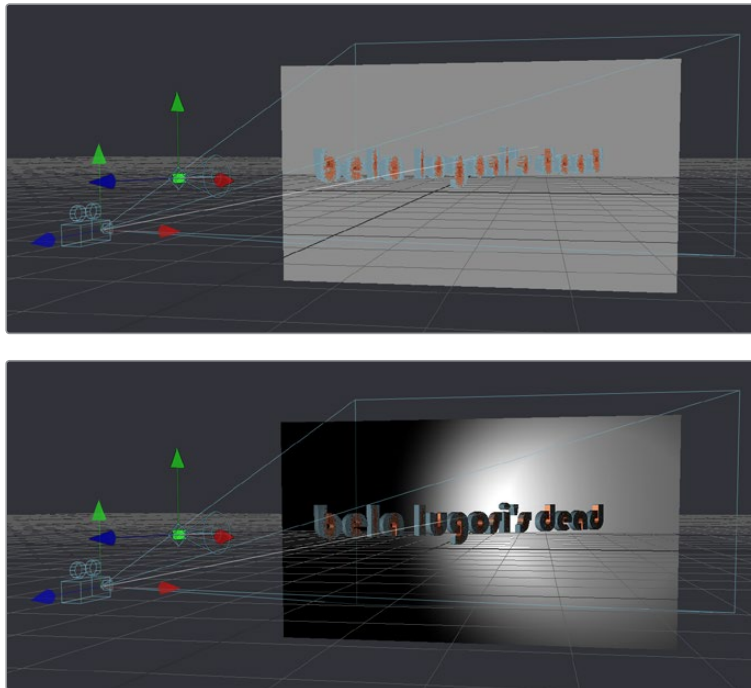
To see the effects of the default light on the scene:

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.



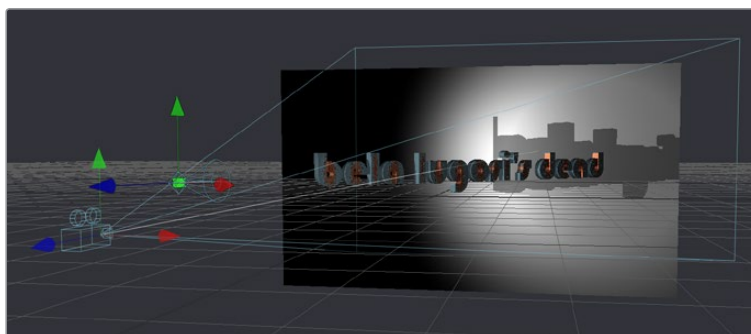
A 3D scene using default lights (top), and the same scene with lighting turned on (bottom).

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.

NOTE: The shadows shown in the 3D viewer are always hard edged. Soft shadows are available for output to the rest of your composition in the software renderer of the `Renderer3D` node.

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:


- **Quick Sort:** Reorders the polygons in the scene serially, from back to front, to produce a reasonably accurate rendering of transparency.
- **Full Sort:** Renders every polygon in Z order to produce the most 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.



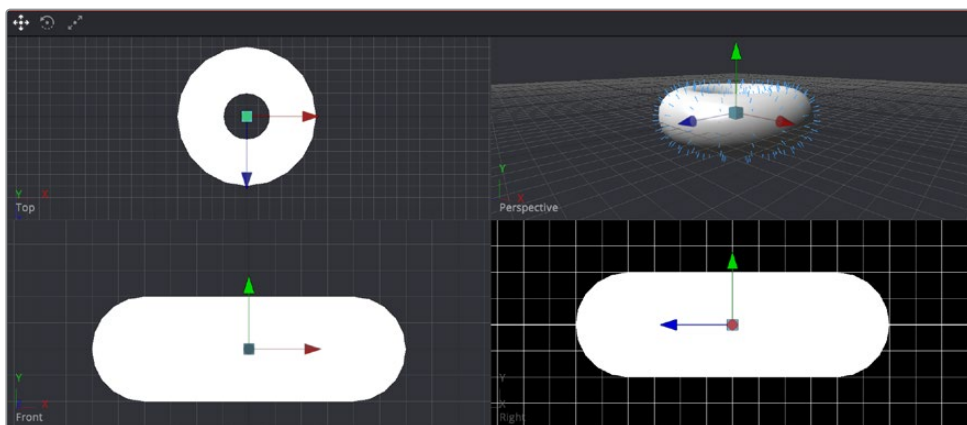
The Normals viewed in a 3D scene.

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 with which to divide 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 Enable the 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 Turn the Subview button on, then choose a subview, and press Shift- to force it into the pane you clicked.


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.

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 cross hair 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.



The Guides submenu in the viewer's contextual menu.

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, Fusion uses the standard compositing conventions of Region of Interest (RoI) and Domain of Definition (DoD) to dramatically improve rendering 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 do not get affected by the node. This helps the renderer to optimize its performance. Secondly, Fusion can now keep track of, and apply a nodes 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 only occupies 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 of 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 .raw and OpenEXR formats are capable of storing the data window of the image, and Fusion will apply this as the DoD when loading such an image and will write out the DoD through the Saver node.

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 only occupy 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 (Rol)

The Region of Interest, frequently abbreviated to Rol, 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 Rol tells the node which pixels actually need to be rendered. When a node renders, it intersects the current Rol with the current DoD to determine what pixels should be affected.

Enabling Rol Controls

You can turn on the Rol 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 Rol 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 Rol controls, do one of the following:

- Click the Rol button in the 2D Viewer toolbar.
- Right-click in a viewer and choose Region > Show DoD from the contextual menu.

When Rol is enabled and Show is selected from the menu, a rectangular Rol control appears in the viewer. If this is the first time Rol has been enabled, it will be set to the full width and height of the image. Otherwise, the last known position of the Rol for that view is used. However, if you want to set the Rol to a custom area within the frame, you can do one of the following.

To adjust the Rol controls, do one of the following:

- Drag any edge of the Rol rectangle to adjust one side of the Rol.
- Drag a corner to adjust the size of the Rol rectangle from that corner.
- Drag the small circle found at the top left corner of the Rol rectangle to move the Rol without adjusting its dimensions.

Sometimes, it's faster to simply draw a rectangle where you want the Rol 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're finished needing 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 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. 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 Loader 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: Image Overlays for showing on-screen controls will defeat 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, while 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 primarily used 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 3D scenes. Image LUTs are routinely used to get from one scene referred colorspace to another. For example, if you're working with log-encoded media but you want to see how the image will look in the final color space, you can choose a LUT to make this image transform as a preview.

Buffer LUTs

The Buffer LUT is applied to all 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 particular to the display you're using. For instance, making a DCI-P3 projector show the image as it would look on an sRGB monitor.

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 final output.

However, in more elaborate production pipelines, you may have the 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, 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 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

This LUT is most often used when working in Linear gamma but needing to view content in sRGB or REC 709. Similar to the Gamut node, this LUT converts a source color space to an output color space, with options to deal with gamma settings, alpha channels, and premultiplication.

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. On macOS, the Macintosh HD/Library/Application Support/Blackmagic Design/Fusion/LUTs/ directory lets that macro appear in the viewer LUT menu. Windows uses C:\Program Files\Blackmagic Design\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 detail about creating Macros, see Chapter 7, "Node Groups, Macros, and Fusion Templates."

Preset LUTs

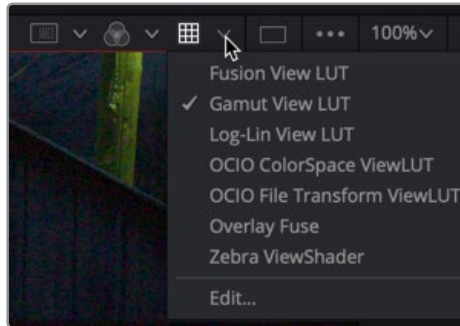
All LUTs that are available appear, by category, in the Viewer LUT menu.

Fuse LUTs

Fuses are scriptable plug-ins 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 detail about fuses, see Chapter 7, "Node Groups, Macros, and Fusion Templates."

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 that are available to Fusion, including custom LUTs you've installed yourself.



The Viewer LUT button and menu.

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 viewers'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

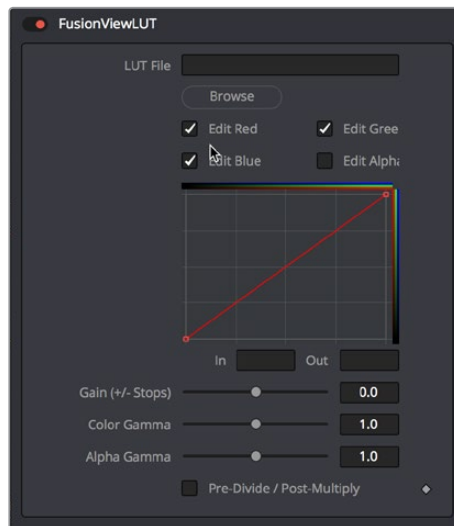
The viewers are the primary area where composites are assessed, so it's crucial that they provide an accurate representation of what the content will look like when it's played for an audience. The LUT Editor allows you to customize your viewer's output to match the gamma and color characteristics of 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 bottom 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.



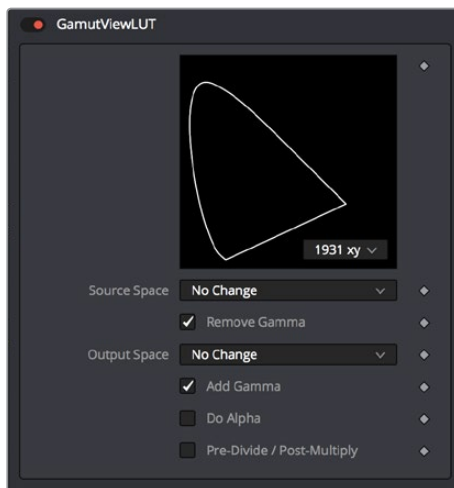
The LUT Editor for the default Fusion View LUT.

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 non-linear gamma, or they let you simply remove or add the appropriate gamma values without changing the color space.

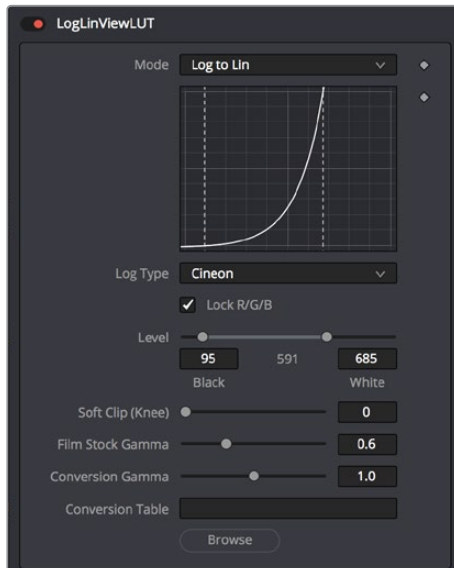
Selecting the Pre-Divide/Post-Multiply checkbox will cause the image's pixel values to be divided by the alpha values prior to this conversion, and then re-multiplied by the alpha value after this conversion. This helps to avoid the creation of illegally additive images, particularly around the edges of a blue/green key or when working with 3D rendered objects.



The Gamut View LUT Editor.

Editing the Log-Lin 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 pop-up, and whether or not 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 being made. 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 high dynamic range file in Log color space you may often apply three LUTs. First, a Log-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 on-screen 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 > File name 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 window 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. For details on the ASCII LUT file format, see the ASCII Import Appendix.

To export a LUT, do the following:

- 1 Click the viewer LUT button to enable it.
- 2 Click the viewer LUT menu, 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 pop-up 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 pop-up 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 DaVinci Resolve .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.

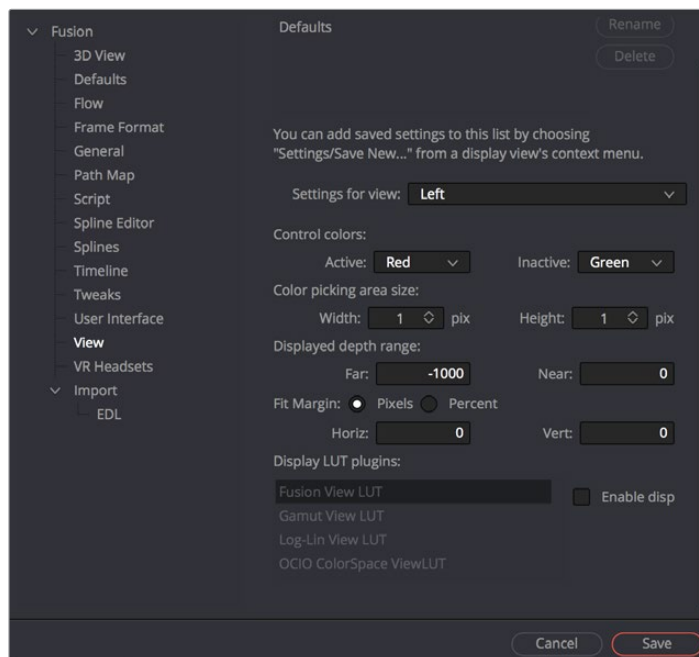
To use a node, group or macro as a Viewer LUT, do the following:

- 1 Select the node, group, or macros.
- 2 Right-click over the selected node, 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 View panel of the Fusion Settings window. Clicking the Enable Display LUT checkbox allows you to select a LUT from the Display LUT plug-ins 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 > Settings name.

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 on-screen 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.

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 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 on-screen 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

Lets you toggle the display of the text that sometimes accompanies on-screen 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 9

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	222
The Tools and Modifiers Panels	222
Customizing the Inspector	223
Inspector Height	223
Inspector Display Preferences	223
Opening Nodes in the Inspector	224
Pinning Multiple Nodes in the Inspector	225
Hiding Inspector Controls	226
Using the Control Header	226
Selecting and Viewing Nodes in the Inspector	226
Using Header Controls	227
Versioning Nodes	227
Parameter Tabs	228
The Settings Tab	228
Inspector Controls Explained	232
Fusion Slider Controls	232
Thumbwheel	232
Range Controls	233
Checkboxes	233
Pop-Up Menus	233
Button Arrays	234
Color Wheel and Picker	234
Modifiers	238
Animating Parameters in the Inspector	239
Removing Animation From a Parameter	239
Attaching a Parameter to an Existing Animation Curve	240
Connecting Parameters	240
Connecting Parameters by Publishing	240
Connecting Parameters by Pick Whipping	240
Contextual Menus	241
Customizing Node Parameters with UserControls	242
An Example of Customizing DirectionalBlur	242

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 (previously called the Control Panel) 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 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.

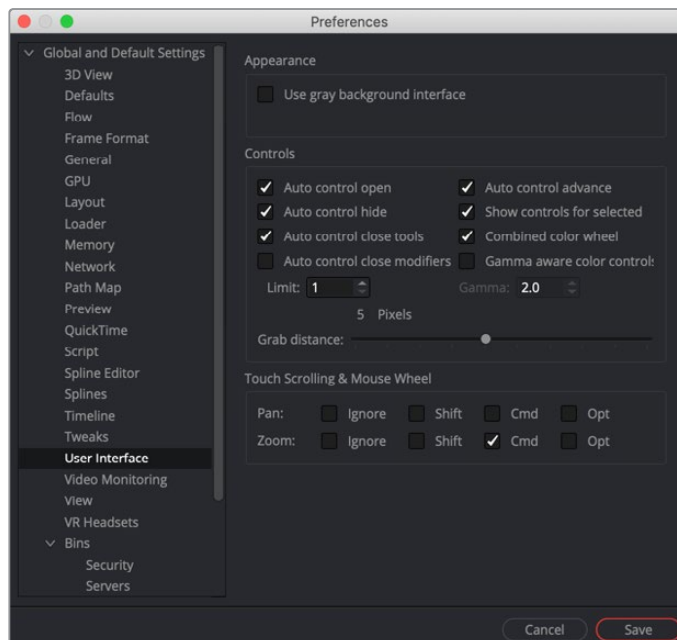


The Maximize button on the left side of the Inspector.

In maximized height mode, the Inspector takes up 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 Studio > Preferences 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 control header in the Inspector, but the parameters remain hidden unless you click the control 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.

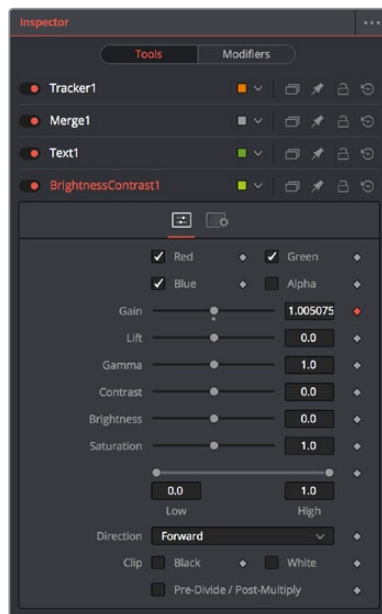
Opening Nodes in the Inspector

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, Keyframe 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, control 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 control 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 control 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 control 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 Control 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 composers 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 on or off:

Right-click on the node in the Node Editor, or on the Control Header, and choose Modes > Show Controls from the contextual menu.

Using the Control Header

When you select a node, it populates the Inspector with a title bar, or control header, that displays that node's name as well as other controls that govern that node. A node's control header itself has a variety of controls, but clicking (or double-clicking) on a control header also exposes that node's parameters.



A node's control header.

When you select multiple nodes at once, you'll see multiple control 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

Control headers are click targets for selecting nodes, opening and closing node parameters, and other things.

Methods of using control headers:

- **To select a node using the control header:** When multiple nodes are selected, you can make a node the active node by clicking its control header in the Inspector. As the actively selected node, the control 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 control header:** You can view a node by dragging its control 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 control 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 control header makes it fast to do simple things.

- **To turn nodes off and on:** Each control 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 control header name:** The name of the node corresponding to that control header is displayed next. You can change the name by right-clicking the control 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 a control 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 control 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 a control 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 control 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.

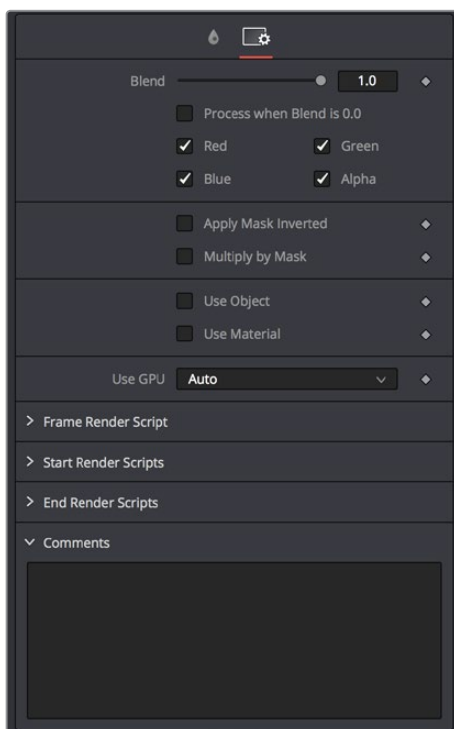
More complicated nodes have more tabs containing more groups of parameters. For example, the Delta Keyer has seven tabs: separating Key, Pre-Matte, Matte, Fringe, Tuning, and Mask parameters, along with the obligatory Settings tab. These tabs keep the Delta Keyer from being a giant scrolling list of 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 DeltaKeyer 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 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 plug-ins 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.

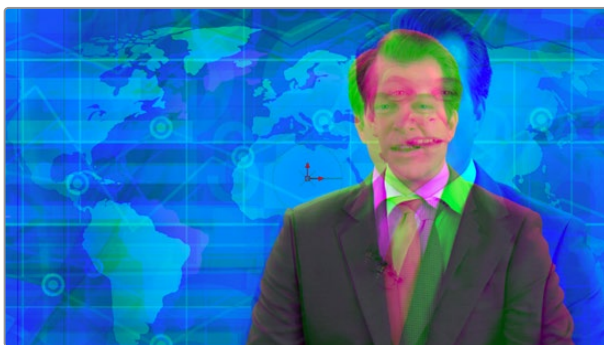
Red/Green/Blue/Alpha Channel Checkboxes

Most nodes have a set of RGBA checkboxes in the Settings tab. These checkboxes let you exclude any combination of these channels from being affected by that node.



The channel limiting checkboxes 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.

Skiping 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 checkboxes on another tab in the node.



Channel checkboxes on the Controls tab of the Blur node indicates that disabled channels won't be processed at all, to save rendering time.

In these cases, the Common Control channel checkboxes 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 Open EXR 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.

Pick Controls

The Pick 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 Pick button to grab IDs from the image in the viewer, the same way you use the Color Picker to select a color. 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

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

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

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 details 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 detail, see Chapter 33, "Generator Nodes."


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 (macOS) or Ctrl (Windows) 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.

Once you click directly on a slider handle you can make changes to its value using the Left and Right Arrow keys. The Command and Shift keys can again be used to modify the value in larger or 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 by entering 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.

You can use the arrowheads at either end of the control to fine tune your adjustments. Once the thumbwheel has been selected either by dragging or using the arrow keys, you can use the Left and Right 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 (macOS) or Ctrl (Windows) 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.

Pop-Up Menus

Pop-up 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.



Pop-up menu in the TimeSpeed node.

Pop-up 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

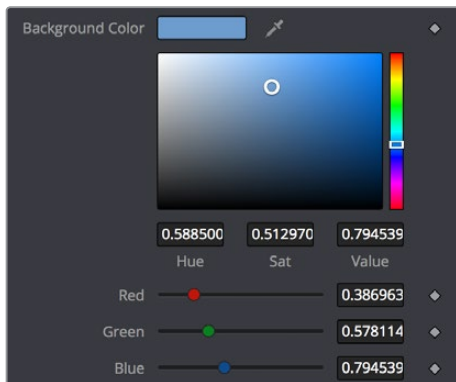
Button arrays are groups of buttons that allow you to select from a range of options. They are almost identical in function to drop-down menu controls, except that in the case of a button array it is possible to see all of the available options at a glance. Often button arrays use icons to make the options more immediately comprehensible.



The Lens Type button array in the Defocus node

Color Wheel and Picker

The Color Picker 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 below the Pick button. The swatch has two halves; the left half shows the color, and the right half shows the color overlaid on a checkerboard background to preview transparency.



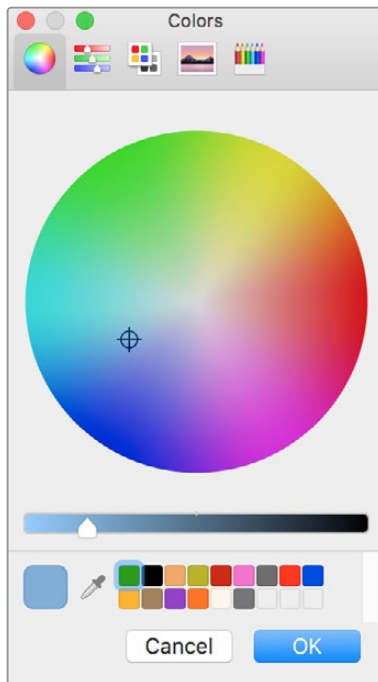
The color swatch with transparency preview.

The Color Picker is extremely flexible and has four different techniques for selecting and displaying colors.

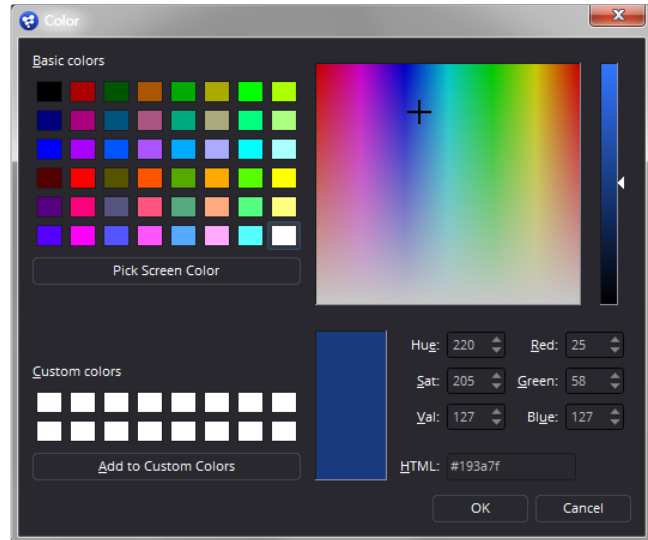
TIP: Color can be represented by 0–1, 0.255, or 0–65000 by setting the range you want in the Preferences > General panel.

MacOS and Windows Color Nodes

Clicking on the Pick button will display the operating system's standard Color Selection node.



macOS Colors panel.




Windows Color dialog.

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.

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 Pick button, and then drag the cursor into the viewer. The cursor will change to an Eye Dropper, 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 eye dropper 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 Eye Dropper, 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 draws the gradient along a straight line from the starting color stop to the ending color stop.




Linear gradient.

Reflect draws the gradient by mirroring the linear gradient on either side of the starting point.




Reflect gradient.

Square draws the gradient by using a square pattern when the starting point is at the center of the image.




Square gradient.

Cross draws the gradient using a cross pattern when the starting point is at the center of the image.



Cross gradient.

Radial draws the gradient in a circular pattern when the starting point is at the center of the image.



Radial 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 on-screen 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 Eye Dropper 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 Eye Dropper 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 of 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.



KeyframeControlsFusionPageInspector.

Orange Keyframe buttons in the Inspector show there's a keyframe at that frame

Once you've keyframed one or more parameters, 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 Keyframe Editor and/or Spline Editor. For more detail about keyframing in Fusion, see Chapter 20, "Keyframing in 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 Bezier, 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.

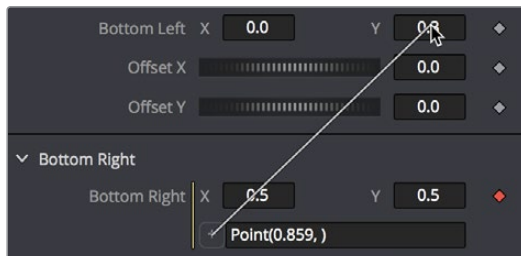
Connecting Parameters by Pick Whipping

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.

To link two parameters using a pickwhip:

- 1 Double-click the field of a parameter you want to pickwhip 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.



Pickwhipping 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 detail on Pick Whipping and Expressions, see Chapter 22, “Animating With Modifiers and Expressions.”

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 control header. The node’s contextual menu includes the same menu options that are accessed by right-clicking on a node in the Node Editor. See Chapter 6, “Working in the Node Editor,” for more details on these options.

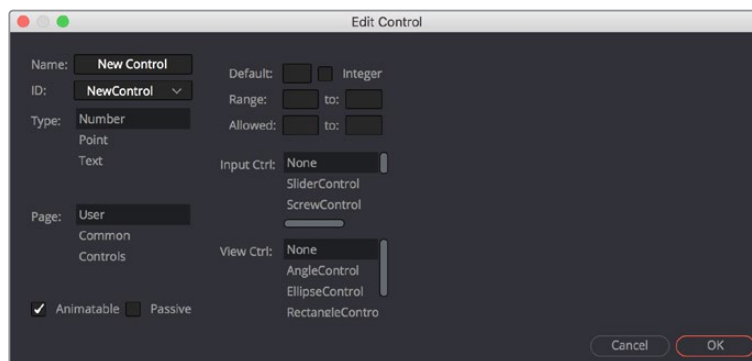
Parameter Contextual Menus

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 UserControls

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.



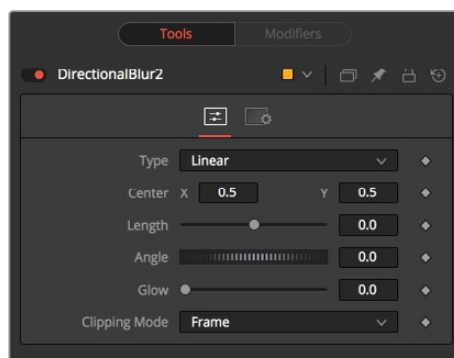
The Edit Control window.

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 DirectionalBlur

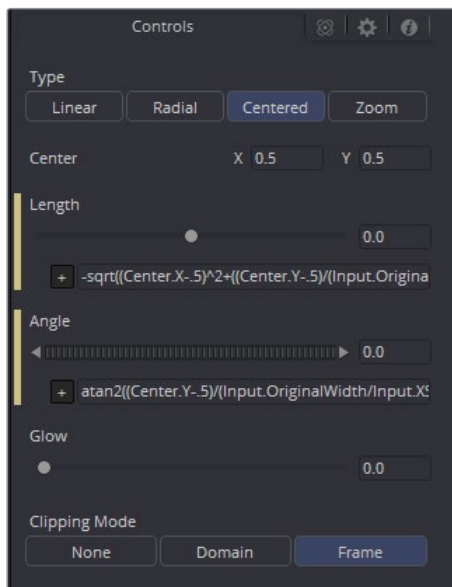
In the following example, let's suppose we wanted to create a more intuitive way of controlling a linear blur than using the Length and Angle sliders independently.



Default DirectionalBlur 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`



DirectionalBlur 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.



DirectionalBlurs with UserControls applied.

Chapter 10

Bins

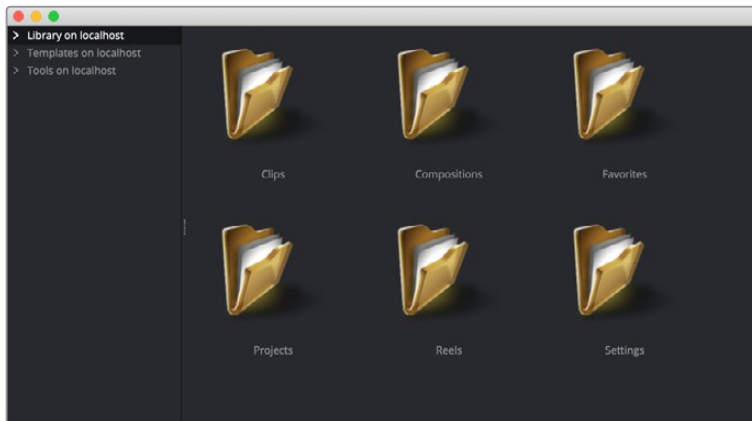
This chapter covers the Bin system in Fusion. Bins allow for storage and organization of clips, compositions, tool settings, and macros. It includes a built-in Studio Player for creating a playlist of multiple shots and their versions. Bins can be used in a server configuration for organizing shots and collaborating with other team members across the studio.

Contents

Bins Overview	247
Organizing Bins	249
Adding and Using Content	250
File Type Details	251
Using Content from Bins	251
Using the Studio Player	252
Jog and Shuttle	260
Stamp Files	260
Connecting Bins Over a Network	261
Adding a Remote Bin Entry	261
Accessing Remote Bins	262
Permissions	262
Studio Player and Bin Server	262


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 Interface

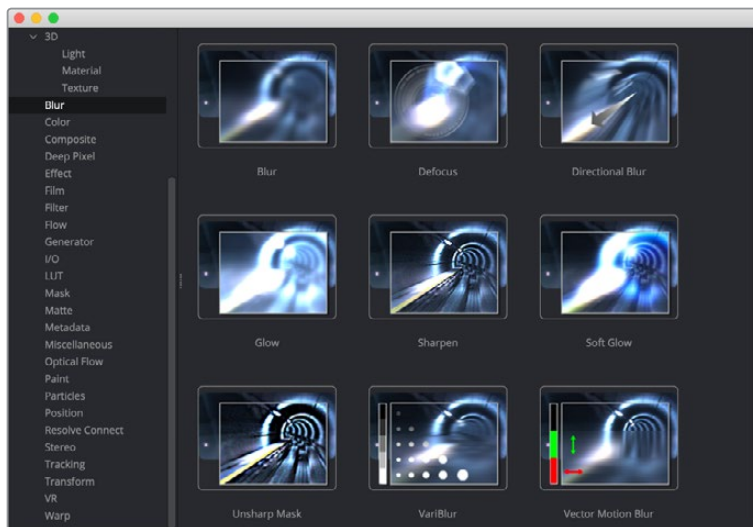
The Bins window is separated into two panels. The sidebar on the left is a list of the bins, while the panel on the right displays the selected bin's content.



The Bins sidebar.

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, Settings, and Tools. Parent folders contain subfolders that hold the content. For instance, the Tools bin is parent folder to all the categories of Tools. Parent folders can be identified by the disclosure arrow to the left of the 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.

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.

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, do the following:

- 1 Right-click in an empty area of the Contents panel.
- 2 From the contextual menu, choose View > Details.


Name	File name
ImagePlane_Mountain_BG.setting	/Users/dion/Documents/Princess_Package/Setti
ChromaKeyer_Mountain.setting	/Users/dion/Documents/Princess_Package/Setti
GarbageMask.setting	/Users/dion/Documents/Princess_Package/Setti
ImagePlane_Sky.setting	/Users/dion/Documents/Princess_Package/Setti
ImagePlane_Princess.setting	/Users/dion/Documents/Princess_Package/Setti
UltraKeyer_Princess.setting	/Users/dion/Documents/Princess_Package/Setti
ImagePlane_Mountain_FG.setting	/Users/dion/Documents/Princess_Package/Setti
Camera.setting	/Users/dion/Documents/Princess_Package/Setti

The Bins 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.

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 maybe easier to view as icons.

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

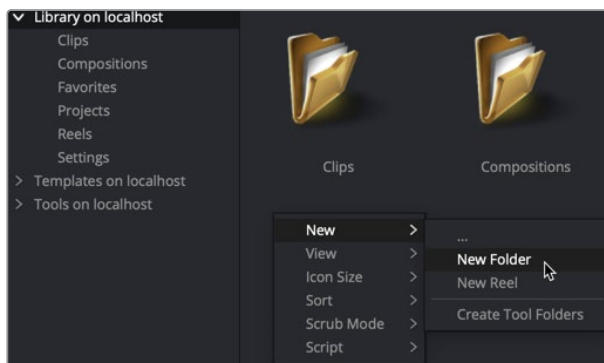
Once you begin adding your own categories and content, you can have hundreds or thousands of items that need to be organized. To keep the bins accessible, you'll want some basic organizational skills, just like the skills you use to keep files and documents organized on your computer.

To open the Bins window, do one of the following:

- Choose File > Bins from the menu bar.
- Press Command-B .

To create a new folder in the sidebar, do the following:

- 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, do the following:

- 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, do the following:

- 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, the folder you are dragging is nested under the targeted folder. 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 an item from the bins, do the following:

- 1 Select the folder in the Contents panel.
 - 2 Right-click on the item and choose Delete from the contextual menu.
- or
- 1 Select the folder in the Contents panel.
 - 2 Press Command-Delete (Backspace on Windows).



The Delete function in the contextual menu.


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 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, do the following:

- 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 footage, files, or Comps in the file browser, and then click Open.



Add Item 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 the 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 due to the file type and some because of where they are located.

Projects and Media

In addition to using the Add Item contextual menu as explained earlier, Fusion 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.

Tools

An additional way to add a tool with the current settings to a bin is to drag the tool's tile from the Node Editor into the bin's Contents panel.

Tool Settings

If you want to add a tool or tools with custom settings, you can select the tool(s) in the Node Editor and drag it/them into the desired bin's Contents panel. A dialog allows you to choose where on disk the new tool settings file will be saved and then adds it to the bin.

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 in to a bin.

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 creates a new Loader that points to the media on disk. Still files or photos are automatically set to loop. The media that you add to the bins is referred to as a clip.

Compositions

To add a composition you must double-click it in the bin to open it. 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 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 pipe, 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.

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 can ingest, like DPX, ProRes, BMD RAW, QuickTime, and others. Clips are cached into RAM, so even large formats like EXR will loop playback from 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. Annotation notes can be added to the project, and when a Bin Server is set up the Studio Player can be accessed by multiple artists.

Here are some Studio Player highlights:

- The Timeline interface can create a playlist for reviewing multiple shots.
- 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.
- Per-shot color adjustment controls allow for consistent display of shots from different formats.
- Annotation notes can be typed on each shot and version, as well as the entire project.
- Audio Scratch track can be enabled for each clip during playback.
- Shot 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 features and clips of the Studio Player.

Playing a Single Clip

Clips created from image sequences, .MOV files, and AVI files can be previewed using the Studio Player without having to first add the clip to a Node tree.

To play back 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 back.

To scrub a clip in Studio Player, do the following:

- 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 by clicking the Option menu in the lower-left corner of the Studio player and choosing 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, a new Reel item can be created to hold multiple clips. This Reel can be identified by the multi image border around its thumbnail.

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.

Double-clicking the Reel will open the Studio Player interface along the bottom of the Bin window. Without a clip selected, the top half of the Studio Player shows the bin in place of a viewer.

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.

- Set Loop In/Out sets the start and end frame for playing a section of the timeline in a loop.
- Shot sets the loop the current clip.
- Reset disables the loop mode.
- The M button is used to show Metadata of the image.
- RGB and alpha view toggles between color and alpha of the clip.
- Brightness Gamma adjusts the brightness and gamma of the overall display and is applied to all clips. Individual clip color controls can also be applied using another menu.
- The Video button outputs playback to Blackmagic Design DeckLink and UltraStudio devices.
- Transport controls are used to play forward, backward, fast forward, and fast backward, as well as go to the start and end of a clip.
- The Sync button is 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.
- Three buttons control the visibility of three customizable Guide settings.

Adding Clips and Comps

Clips or comps from the Bin can be dragged to the storyboard area of the reel.




Inserting Shots and Versions

There are two options when dragging clips and comps to the storyboard playlist. You can insert a shot between existing clips by positioning the new clip in between existing items in the reel.



Alternatively, you can add a version to an existing clip by dragging the new item on top.

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. In the example below, the first shot has two versions, the second shot has four versions, and the last clip has only one version.



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.



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 point, add notes, adjust the color, and add an audio soundtrack.




- **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.



When notes are added, they are time and date stamped as well as named 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 Audio menu option can import an audio .wav file that will play back along with the selected clip.


Option Menu

The three-dot Option 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 Option 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

Customizable guide overlays are used to display monitor/title safety and crops. Guides can be loaded and activated in the Studio Player viewer by clicking the guide switches in the toolbar. Up to three guides can be active simultaneously for any given project.



Three toolbar guide buttons enable guide overlays in the viewer.

Guides can be customized using the Guides folder in the Fusion package. A simple XML file, as explained below, defines each guide. This makes it easy to create and share guides.

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 out of four hex values and determines the visual appearance of the line.

Examples for such patterns are:

```
>>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="FOFO" 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"

Guide

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%" },
    },
}
```

```

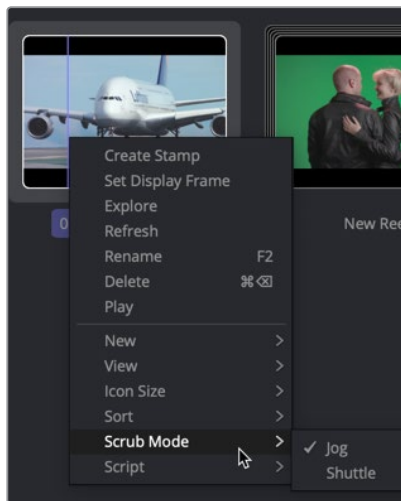
        HLine { Y1="50%", Pattern = 0xF0F0, Color = { R = 1.0, G = 0.75,
B = 0.05, A=1.0 } },
        VLine { X1="50%", Pattern = 0xF0F0, Color = { R = 1.0, G = 0.75,
B = 0.05, A=1.0 } },
    },
}

```

Jog and Shuttle

Without opening the Studio Player, 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 is paused.

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, do the following:

- 1 Right-click on the clip in the bin.
- 2 Choose Create Stamp from the contextual menu.

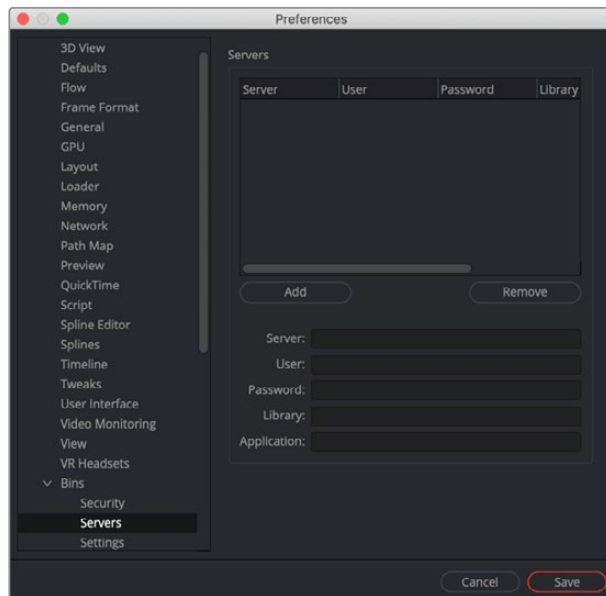
Connecting Bins Over a Network

The Status bar at the top of the Bins window shows the progress as the stamp is created. Since the stamp is created in the background, other stamps can be queued as well and you can continue working with the composition.

You can share bins among computers running Fusion on the network. These shared bins are called Remote bins, and one or more Remote bins can be shared by everyone in a studio.

To connect to a remote system and display its bins, do the following:

- 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 a 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 get 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, 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.



Master enables controlling other slave players on the network.



Slave disables the local transport controls and syncs playback to the master.



PART 3

Importing and Rendering

Chapter 11

Preparing Projects and Importing Media

This chapter provides information about handling timing, bit depth, and color space conversions for Loader nodes. It covers many of the important steps you'll need to know when starting a project.

Contents

Preparing Projects and Images	266
Overview	266
Setting Up a Project	266
Loading Images	267
Setting Global and Trim Ranges	268
Color Bit Depths	268
Setting Color Depth	269
Converting Images to Linear Color Space	273

Preparing Projects and Images

Overview

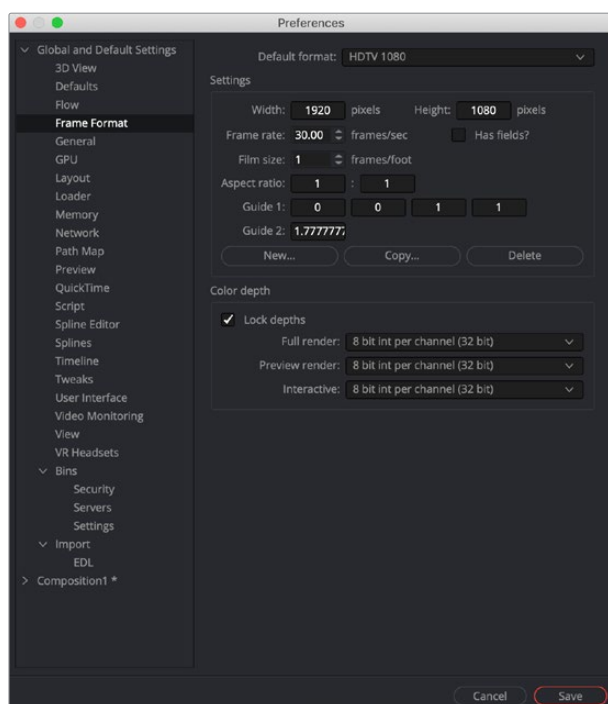
This chapter is about preparing a project and adding clips to the Node Editor. Fusion can read single frames, image sequences, or movie files at various resolutions and bit depths. Knowing which files you can load in, how to set up a project to handle them, and finally, reading those files in are the first steps in beginning your composition.

Setting Up a Project

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 configures and combines materials of different formats when loading and merging them together.

The first thing you do when starting on a new project 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 compositions.

The Frame Format preferences are used to determine the default resolution used for new Creator tools (i.e., text, background, fractals, etc.), pixel aspect for display and rotation, as well as the frame rate used for playback.



Frame Format preferences for images generated in Fusion.

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 new Comp 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 compositions you create. They do not affect any existing compositions. 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.

The preferences are described in greater detail in Chapter 5, "Preferences."

Loading Images

Once the Frame Format preferences are set, you typically begin to composite by importing or, more accurately, reading in source media. When Fusion reads in media, it doesn't convert or move the original files. It just reads the files in place, so the concept of importing source media into a Library or some kind of media vault is not applicable in Fusion. 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 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.

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 a Windows Explorer file browser or macOS Finder window into the Node Editor.

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.


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.

When you drag from the operating system, a Loader is automatically added to the Node Editor, and the media file is directly added to the comp. If multiple files are dragged into the Node Editor, a separate Loader is added for each file. This comes in handy when you want to read in photographs from a digital camera that are numbered sequentially.

For more detail about the Loader node, see Chapter 41, "I/O Nodes."

Setting Global and Trim Ranges

The Inspector for a Loader node includes settings for how to interpret the content in terms of color bit depth and time. 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 compositions global time.



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 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 have an In on frame 10, then frame 10 will appear at the comps starting point on frame 0.

Instead of using the Inspector to adjust timing, it is visually more obvious if you use the Keyframe Editor. For more detail on the Keyframe Editor and adjusting a clip's time, see Chapter 20, "Keyframing in Fusion."

Color Bit Depths

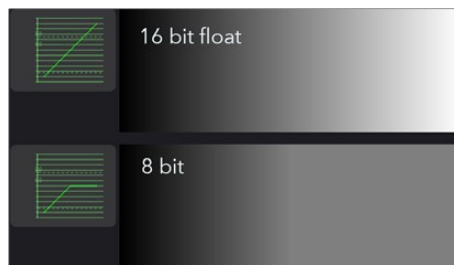
The term bit depth describes how many colors are available in the color palette that is used to make up an image. The higher the bit depth, the greater the precision of color is in the image, and therefore the greater the reproduction of color. The higher precision is most obvious in gradients with subtle changes. Lower bit depth gradients have noticeable banding artifacts, whereas higher bit depth images can reproduce more colors, so fewer, if any, banding artifacts are seen. You can set Fusion to process images with 8-bit integer, 16-bit integer, 16-bit float and 32-bit float bits per channel.

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. 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 over 8-bit integer images, eliminating problems with stepping in gradients and improving the resolution of the color. 16-bit integer color comfortably includes all the colors that can be represented in standard digital film processing. Although you can select 16-bit integer processing for an 8-bit clip, it will not reduce banding that already exists in the original file, but it can provide more precision when adding additional effects to the clip. Additionally, many 10-bit YUV formats will see a benefit when loaded with 16 bits of RGBA rather than 8 bits. However, 16-bit integer limits your shadows and highlights to a scale of 0 to 1, with blacks at 0 and whites at 1.

This sounds perfectly fine 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 and specular highlights above 1.

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. 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 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.

32-bit float can work with shadow areas below 0 and highlights above 1, similar to 16-bit float, except with a much greater range of precision.

Setting Color Depth

Fusion automatically uses the color depth that makes the most sense for the file format of an image. For example, if you read in a TGA file from disk, then the color depth for the Loader will be set to 8-bits per channel. The TGA format is an 8-bit format, so loading the image at a greater color depth would generally be wasteful. If a 16-bit TIFF is loaded, the color depth will be 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 tool's Inspector. The Loader's Inspector, as well as the Inspector for images generated in Fusion (i.e., text, gradients, fractals, 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.

The Default setting will force the tool to process based on the settings that are configured in the Node Editor's Frame Format preferences.

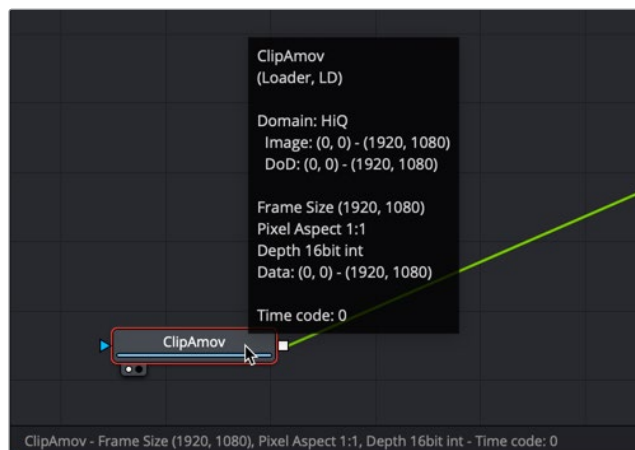
The Frame Format preferences are used to set a default value for color depth, applied when a source tool is added to the Node Editor. There are three drop-down menus to configure color depth in the preferences. They specify the different color depths for the interactive session, final renders, and preview 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 will be 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 that are set to integer options.



The Frame Format Color Bit Depth settings.

If you aren't sure what the color depth at which a tool processes is, place the mouse over the tool's tile in the flow and a tooltip will appear, listing the color depth for that tool. It will also appear on the status bar.



Hover over a tool 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 will preserve 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 foreground input of the tool will be adjusted to match the color depth of the background.

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. The second benefit is the preservation of shadow and highlight values that go beyond the normal tonal range.

Greater Accuracy

Let's look at a simple example of how floating-point values prevent loss of accuracy caused by integer rounding. Imagine an 8-bit pixel that has a red value of 75. An 8-bit color value ranges from 0 (black) to 256 (white), so our pixel is a bit dark.

Imagine that the gain of that pixel is halved using a Color Correction tool. Now the pixel's red value is half of 75, or 37.5. You cannot, however, store decimal or fractional values in 8-bits, so you must round that value to 37. Now, double the brightness of the pixel with another Color Correction tool. The pixel's red value is 37 multiplied by 2, or 74. We lost a full value 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 come up when merging images together, or transforming them. The more operations that are applied to an image, the more color precision is lost to rounding.

Floating-point values range from 0.0 to 1.0. The value for our example pixel was 75 in 8-bit processing, but the same pixel processed in floating-point color depth would have a value of 0.2941176 (75 divided by 255).

Because floating-point processing allows decimal or fractional values for each pixel, 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.

Accessing Extended Highlights and Shadows

More and more productions are capturing out-of-range images thanks to cost effective cinema cameras like the Blackmagic URSA Mini Pro and even the Pocket Cinema 4K camera. These cameras capture very high dynamic range RAW images and maintain color detail even in heavily over or under exposed 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, to do that you must clip all 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.

An example is to imagine an 8-bit pixel that has a red value of 200 (bright red). Now add a Color Gain tool and double the brightness of the red channel. The result is $200 \times 2 = 400$. As mentioned above, however, 8-bit color values are limited to a range of 0 through 255. So the pixel's value will be clipped to 255, or pure red. If now the brightness of the result is halved, the result will be half of 255, or 127 (rounded), instead of the original value of 200.

When processing floating-point colors, pixel values greater than white or darker than black are maintained. There is no value clipping. The pixel is still shown in the display view 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.

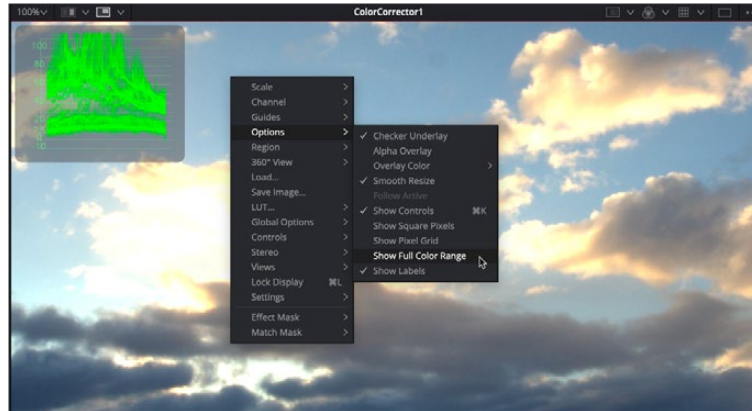
There is also some value to using float color depths with 8-bit HD video. A good time to use float processing with video is when there will be a lot of color correction. Using float will help maintain precision by avoiding the rounding errors common to 8-bit processing as described above.

Detecting Extended Highlight and Shadow Values

Although floating-point processing preserves extended values below 0 and greater than 1, 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 if there are out-of-range values in a viewed image, do the following:

Right- click in the viewer and choose Options > on the Show Full Color Range button in the Viewer toolbar.



ViewerShowFillRange

Use the Show Full Color Range contextual menu to detect out-of-range images.


Enabling this display mode will rescale the color values in the image so that the brightest color in the image is remapped to white and the darkest is remapped to black. Out-of-range colors will be brought back into visible range, and the overall contrast of the image will be lowered as a result. If there are out-of-range values in the image, the image will appear washed out while displayed in this mode.

The 3D Histogram view type can also be helpful in visualizing out-of-range colors in an image. For more details, see Chapter 8, “Using Viewers.”

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 tool set to Clip Black and Clip White, with only the Alpha checkbox selected.




Clip White, Clip Black settings in Brightness/Contrast can be used to clip mattes.

Alternatively, you can clip the range by adding a Change Depth tool and switching to 8-bit or 16-bit integer color depths.


Converting Images to Linear Color Space

The simplified goal of color management is to make sure that what you see as your final product on your computer screen is what your audience sees. For such a simple goal, problems arise that you, as the compositor, must deal with. To understand the problem, you must understand that human beings have a nonlinear bias in how we perceive brightness. We are much more sensitive to changes in the shadow and mid tone areas of an image than we are to changes in the highlights area. Digital cinema cameras also capture images using their own nonlinear color space. In fact, a Quicktime movie, a JPEG sequence, and a DPX sequence will all have different kinds of nonlinear color space they use to store color information.



A nonlinear, logarithmic color conversion.

But color correction and compositing operations work best when everything is linear. This means if we color correct an image to be twice as bright as it started, its pixel values should be exactly two times bigger in number, consistently from shadows to highlights.

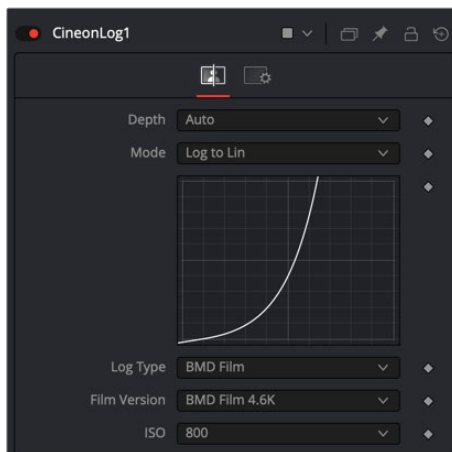


A linear color conversion

When loading images that use a nonlinear gamma into Fusion, you must normalize them by converting them into a linear color space.

To convert an image into linear color space:

- 1 In the Node Editor, select the Loader for the image.
- 2 In the Effects Library, select CineonLog from the Film category and add it after the Load node.
- 3 In the Inspector, set the Mode menu to Log to Lin.
- 4 Set the Log type and Log version to the camera and gamma setting used when capturing the clip.



The CineonLog node setting for a BlackMagic RAW clip.


The image will appear darker or lighter depending on the viewer LUT setting.

TIP: If the image is coming from a different software application, be sure to confirm that the correct color space has been identified. Sometimes applications will embed incorrect metadata in the files.

Computer displays also have their own color space bias that affect the appearance of images. Typically, computer displays use the sRGB color space. So in addition to normalizing the image file's color space, we must apply a Lookup table, or LUT, to the viewer to compensate for the computer's color conversion.

To apply a Gamut LUT to the viewer, do the following:

- 1 Click the LUT button above the viewer to enable the LUTs.
- 2 Click the viewer LUT pop-up menu next to the viewer LUT button.
- 3 Select Gamut View LUT from the menu.
- 4 Click the View LUT pop-up menu again and choose Edit.
- 5 Choose sRGB as the Output Space.



The Viewer LUT pop-up menu.

TIP: If your monitor is calibrated differently, you will need to select a LUT that matches your calibration.

Whether you use the sRGB LUT or a LUT for your specific monitor calibration, you can save the viewer setup as the Default.

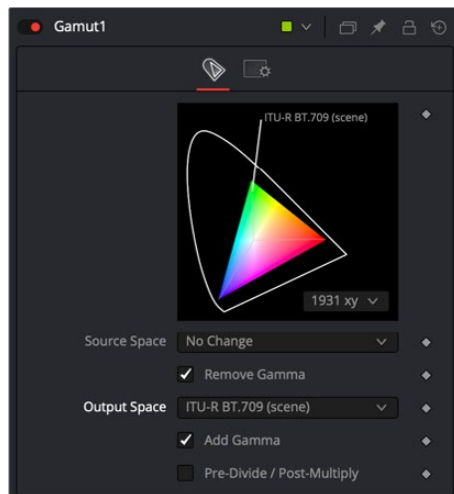
To save the Gamut LUT setup as the default viewer setup do the following:

Right-click in the viewer, and then choose Settings > Save Defaults.

You now have a linear color space working environment for the viewer and one Loader. When you are ready to render, you must output the images back to the appropriate color space for the file type you are saving. If you are outputting to HD, you can use a Gamut node directly before the Saver node.

To set the output render color space:

- 1 Select the last tool in your comp.
- 2 From the Color category in the Effects Library, click Gamut tool.
- 3 In the Inspector, ensure the Source Space is set to No Change.
- 4 From the Output Space menu, select Rec 709 (scene) for rendering out HD content.
- 5 Add a Saver node directly after the Gamut node.



The Gamut node setting for rendering HD content.

Chapter 12

Rendering and Licensing

This chapter covers how to render compositions on your local computer and using multiple computers over a network.

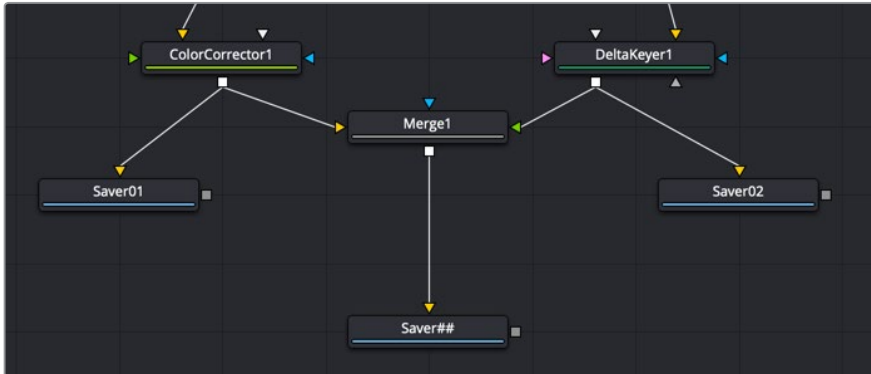
Contents

Rendering with a Saver Node	278
Network Rendering	279
Setting Up Network Rendering	280
The Render Manager	282
The Render Queue	284
The Render Log	287
Preparing a Composition for Network Rendering	288
Submitting a Composition	289
Flipbook Previews	289
Disk Cache	289
When Renders Fail	290
Managing Memory Use	291
Things You Need to Know	292
Troubleshooting	293
How to Set Up Headless RenderNode	294
Requirements	294
Setting Up CentOS	294
Starting Xorg Server with XDummy Driver	294
Starting Fusion RenderNode	294
Appendix I	295
Sample xorg.conf	295
Fusion Studio Network Licensing	296
Setup	296
Technical Info	296

Rendering with a Saver Node

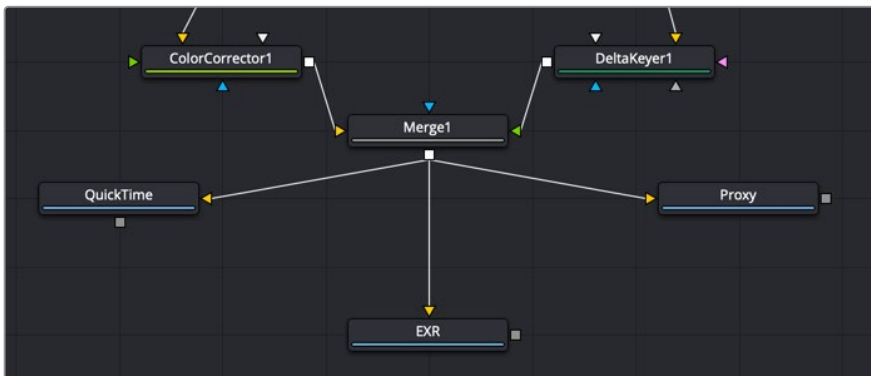
When you have finished your composite, you will need to add one or more Saver nodes to export out a media file. You can attach multiple Saver nodes anywhere along the node tree to render out different parts of a composite. Most of the time, you will place at least one Saver node at the very end of your tree.

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 file out nodes stemming from the same node in order to create servers 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 dialog where you can decide where the exported file is saved and also name the file. You can then use the Inspector to configure the output format.

For more detail on the Saver node, see Chapter 41, “I/O Nodes.”

File Names 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 configure the inspector to output an EXR file. If you decide to change the file type later, you can change it in the Inspector. If you decide to output an image sequence, a four-digit frame number is automatically added before the filename extension. For our file example from above, the name would be `image_name0001.exr`, `image_name0002.exr`, and so on. You can specify the frame padding by adding several `#` to indicate number of digits. For example `###` signifies 001.

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

Network Rendering

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

Fusion can submit compositions to be rendered by other copies of Fusion, as well as to one or more Fusion render nodes.

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, but it can be installed at any time using the Fusion Render Node Installer located in the Blackmagic Fusion Studio installer.dmg on macOS and the Blackmagic Fusion Studio.zip on Windows.

Render Node Installation (Windows)

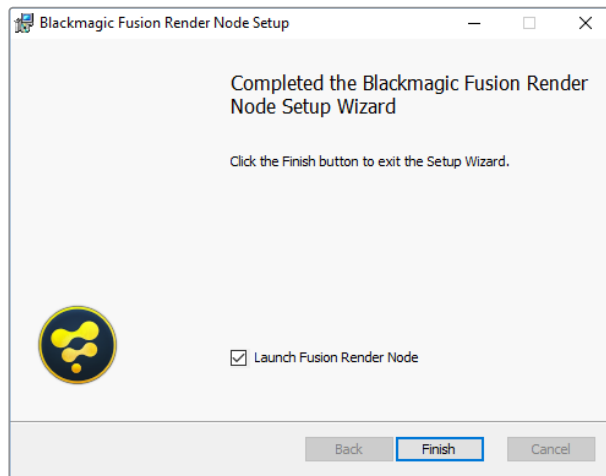
- 1 Locate the Render Node installer in the Blackmagic Fusion Studio directory once it is unzipped.
- 2 Copy the `Install_Fusion_Render_Node_[version].msi` to each computer on the network that you want to perform rendering operations.
- 3 Install the Render node.

Render Node Installation (Mac OS)

- 1 Locate the Render node installer located in the disk image that was downloaded for the application.
- 2 Copy the `Install_Fusion_Render_Node` 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 Startup folder on Windows and to the menu bar in macOS. Each time you log in, 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 or Windows system tray.

To disable the launching, uncheck the checkbox in the last page of the installer. When Render node is running, this icon is visible in the system tray.



Using a Third-Party Render Manager

Many studios make use of a Third-Party Render Manager to control render farms. This allows for efficient sharing of the farm's resources between the many applications that make up a studio's pipeline. Examples of such managers are Smedge, Spider, Rush, and Deadline. Several facilities also use render managers developed in-house. Generally, these render managers expect to find a command line renderer. Fusion Studio includes a script called RenderTool. Lua to integrate Fusion with third-party render managers.

Instructions for using the script in the Fusion > Scripts > Utility folder are found in a text file in the same folder. Alternatively, the instructions are available by opening the script using a text editor and reading the comments at the top.

The script is required because the Fusion Render Slave is a GUI application. When launched from the command line, Fusion will detach from the command line session. This interferes with the method most managers use to monitor the progress of a render.

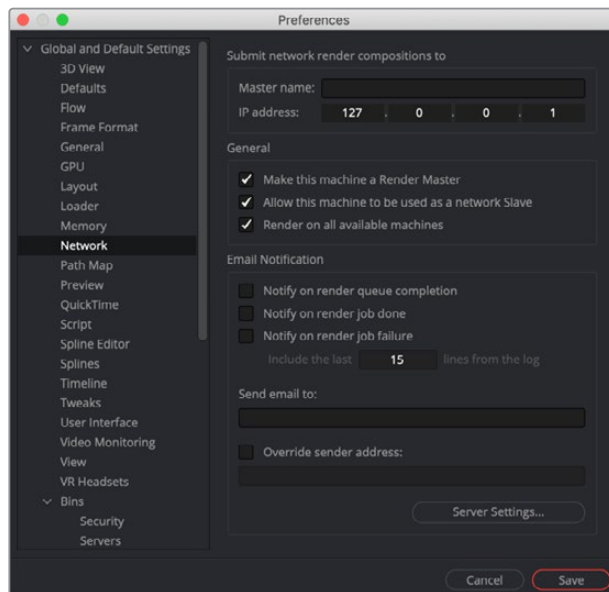
Keep in mind that using a third-party Render Manager will prevent the use of some of Fusion's network rendering features, such as the ability to create net-rendered Flipbook Previews and disk caches.

Setting Up Network Rendering

What Is the Render Master?

The Render Master manages the list of compositions to be rendered (the queue) and allocates frames to slaves for rendering. The Render Manager is also used to maintain the slave list and to push updates to the various slaves 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 using the Network Preferences. The Render Manager application can also be used as a standalone Render Master.



The Network Preferences allow setting up computers as Render Masters.

Render Manager

Acting as a Render Master has no significant impact on render performance. The system resources consumed are insignificant. Many smaller facilities use one of the Render nodes as a Render Master, while freelancers with just one license often use their own interactive license of Fusion as a Render Master.

Setting Up the Render Master

Select the computer to be used as a Render Master and either install a copy of Fusion or the Render node.

In Fusion:

In Fusion, select Fusion Studio > Preferences and open the Global > Network Preferences window.

Enable the Make This Machine a Render Master checkbox. If this machine is to participate in rendering compositions, enable the Allow This Machine to Be Used as a Network Slave checkbox as well.

To have the Render node act as the Render Manager:

Right-click on the Node icon in the Windows system tray or click the icon in the macOS menu bar. Then, select Preferences from the context menu. Open the Global Network preference dialog and enable the checkbox marked Make This Machine a Render Master. If this machine is to participate in rendering compositions, enable the Allow This Machine to Be Used as a Network Slave checkbox as well.

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

Preparing Render Slaves

It is important to ensure that the Render Slaves will accept instructions from the Render Master.

In Fusion:

Select Allow Network Renders from the File menu, or enable the Allow This Machine to Be Used as a Network Slave in the Global > Network Preferences.

For Render Node and Render Manager:


Right-click on the icon for the Render node in the system tray on Windows or click the icon in the macOS menu bar. Then, select Allow Network Renders from the contextual menu, or enable the Allow This Machine to Be Used as a Network Slave in the Global Network Preferences.

Choosing the Render Master Used by Slaves

To submit a composition to a Render Master from a workstation or node, you must tell Fusion which master will manage the render queue. Open the Global Preferences page and locate the Network tab. Type the name of the Render Master in the Master Name field or enter the IP address of the master. Fusion will automatically try to fill in the alternate field by performing a lookup on the local network. Workstations that are not configured as slaves can only set the Render Master with this preference option.

The Render Manager

The Render Manager dialog is used to monitor the progress of rendering to reorder, add, or remove compositions from a queue and to manage the list of slaves used for rendering. Before a Render Master can be useful, add the slaves it will manage to the Slave List, as described below.



Use the Render Manager to add slaves to the list.

Opening the Render Manager Dialog

In Fusion:

Select Render Manager from the File menu or press Command-M to open the Render Manager Dialog.

In the Render node and Render Manager:

Right-click on the icon for the Render node in the system tray and select Render Manager from the context menu.

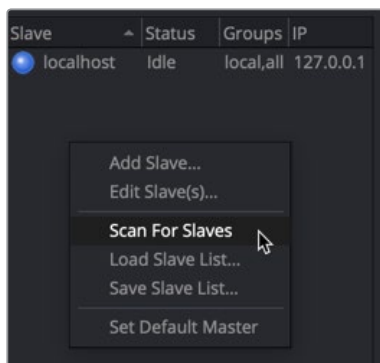
Adding Computers to the Slave List

The Render Manager always starts with one slave in the Slave List, which is itself. This allows the Render Manager to render local queues without using the network. For the Render Master to control additional slaves, add them to the Slave List.

Add slaves into the Slave List by entering the slave's name or IP address manually, or scan for slaves on the local network.

Scanning for Slaves

Select Scan For Slaves from the Slave menu at the top of the dialog on Windows. Alternatively, right-click in the Slave List and select this option from the contextual menu on Mac OS.




Search through all IP addresses on the Subnet for render slaves.

Scanning looks through all IP addresses in the subnet to determine if any other computers in the local network are actively responding on the port Fusion uses for network rendering. A copy of Fusion Studio or the Fusion Render Node must be running on the remote machine in order for it to be detected by the scanning.

Manually Adding Slaves

To manually add a slave to the Slave List, select Add Slave from the Slave menu.



Choose Add Slave from the contextual menu to add a new slave to the list.


Enter the name or the IP address of the remote slave in the dialog box that appears. The manager will attempt to resolve names into IP addresses and IP addresses into names automatically. Use this method to add slaves to the list, slaves that are not currently running or available on the network.

Removing Computers from the Slave List

To remove a computer from the Slave List, select the slave or slaves to be removed from the list and choose Remove Slave(s) from the Slave menu. Use the Command key to select multiple slaves for deletion.

Loading and Saving Slave Lists

The list of slaves is automatically saved in the file Slaves.slv in the Fusion > Queue folder when the Render Manager is exited. In addition, save and load other alternate slave lists by selecting Save Slave List and Load Slave List from the Slave menu.



Open and Save Slave lists for adding them to multiple computers.

The Render Queue


The Render Queue is a list of compositions to be rendered, in the order in which they are to be rendered. The top entry in a queue list will be rendered first, followed by the next one down, and so on. Multiple entries in a queue may render at the same time, depending on the slave group and priority of the composition.

Adding Compositions to the Queue

To add a composition to the queue, use one of the following methods:

From the Render Manager


To add a composition to the queue from within the Render Manager dialog, click on the Add Comp button near the top left of the dialog or right-click in the queue list and select Add Comp from the menu.



Click the Add Comp button to add a save Fusion Comp to the render manager.

From the Composition

When starting a preview or final render, selecting the Use Network checkbox from the Render Settings dialog will add the composition to the end of the current queue in the Render Master's queue manager. The Render Master used is the one configured in the network preferences of the Fusion application that is submitting the composition.



Enable the Use Network setting for previews and final renders.

From a File Browser

Add one or more compositions to the queue by dragging the Fusion .comp project file into the Render Manager's queue list from the Finder (macOS) or File Explorer (Windows).

From the Command Line

You can also add compositions to the queue using the Renderslave utility, which can be found in Fusion > Utilities. A sample Command Line might be `FRender localhost c\compname.comp` queue. This would submit the composition `compname.comp` to the Render Manager on the local machine.

Removing Compositions from the Queue

To remove a composition from the Queue list, select the composition and press the delete key on macOS or the Del key on Windows. Alternatively, right-click on the entry in the Queue list and select Remove Composition from the contextual menu. If the composition is currently rendering, the render will be automatically halted and the next composition in the list will immediately begin rendering.

Saving and Loading Queues


It may be useful to save a Queue list for later reuse. The current Queue list is normally saved as `queue.slv` in the Fusion Queue directory. To save the current queue with a new name, select Save Queue from the File menu. To load a queue from disk, select Load Queue from the File menu.

Reordering the Queue

While building or rendering a queue, priorities for a composition may change. Shifting deadlines may require that a composition further down the Queue list be rendered sooner. Drag an entry to a new position in the list to rearrange the order of the compositions within the queue. If a composition that is already complete (status is done) is reordered so that it is below the currently rendering compositions, that composition will not re-render. Reset the status by right-clicking on the composition in the list and selecting Clear Completed Frames from the contextual menu.

Groups

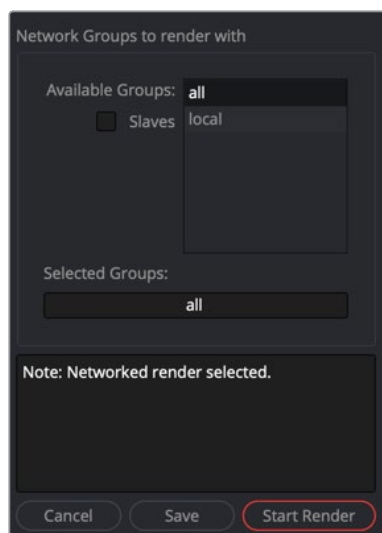
Slaves can be configured into Groups, which are then used when submitting compositions. For example, imagine that there are five Render Slaves. All the slaves are members of the group All, but two of them have more memory than the other slaves, so these two slaves are also added to a group called Hi-Mem. By default, new slaves and compositions in the queue are automatically assigned to All.



The Choose Groups window can be opened from the Render Manager Comp menu.

When a render is submitted to the network, you can choose to submit it to the All group or the Hi-Mem group. When a render is submitted to the Hi-Mem group, only the two machines that are part of that group will be used to render the composition. If a composition is then submitted to the All group, the remaining three machines will start rendering the new composition. When the two slaves in the Render Master group stop rendering the first composition, they will then join the render in progress on the All group.

Groups are optional and don't have to be used. However, managing large render farms across an installation will become easier if they are used.



Select which group to render to from the Render settings dialog.

Assigning Slaves to Groups

To assign a slave to a specific group, right-click on the slave and select Assign Group from the contextual menu. Type the name of the group in the dialog that appears. If the slave is to be a part of multiple groups, separate the name of each group with a comma (e.g., All, Local, Hi-Mem). The order of the groups determines the priority. See “Using Multiple Groups” below for details.

Assigning Compositions to Groups

When a composition is added to the Render Queue, assign it to a group using one of the following methods:

From the Render Settings Dialog

Select the Use Network checkbox in the Render Settings dialog and a section of the dialog will become available to select a group or groups from a list. The list is filled with whatever groups are currently configured for the Render Master. If the Render Master is not the local machine, the list will be filled by querying the remote machine. Use the Command key to select multiple groups for the submitted composition.

Alternatively, select individual slaves using the method of bypassing the group system entirely.

From the Render Manager

When a composition is added to the queue from the Render Manager, the composition is automatically submitted to the group All. To change the group assigned to the composition, right-click on the composition’s entry in the Queue list and select Assign Group from the contextual menu.

Using Multiple Groups

A single slave can be a member of multiple groups. A single composition can also be submitted to multiple groups. Submitting a composition to multiple groups has a relatively straightforward effect: The composition will render on all slaves in the listed groups.

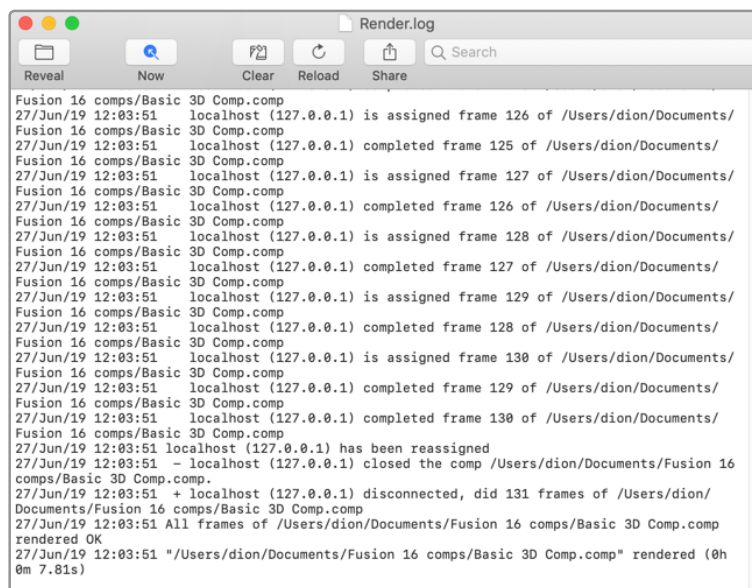
When a slave is a member of multiple groups, the ordering of the groups becomes important because the order defines the priority of the group for that slave.

For example, imagine that a slave is a member of two groups, All and Hi-Mem. If the groups are assigned to the slave as All, Hi-Mem renders submitted to the All group will take priority, overriding any renders in progress that were submitted to the Hi-Mem group.

If the order is changed to Hi-Mem and All, the priority reverses. The Render Slave will participate in rendering compositions submitted to the all group as long as no Hi-Mem renders are in the queue. When no Hi-Mem renders are waiting to be completed in the queue, that slave will render any composition submitted to the All group.

The Render Log

The Render Log is displayed beneath the Queue list in the Render Manager. The text shown in this window displays the Render Manager’s activities, including which frame is assigned to what slave, which slaves have loaded the compositions in the queue, together with statistics for each render after completion.



The Render Log can be opened from within the Render Manager.

There are two modes for the Render Log: a Verbose mode and a Brief mode. Verbose logging logs all events from the Render Manager, while Brief mode logs only which frames are assigned to each slave and when they are completed. Toggle between Verbose and Brief logging in the Misc menu of the Render Manager. Verbose mode is on by default.

Preparing a Composition for Network Rendering

Paths

The Paths used to load a composition and its footage, and to save the composition's results, are critical to the operation of network rendering. Ensure that Loaders use filenames for footage that will be valid for all slaves expected to render the composition. Savers should save to folders that all slaves can see and to which all slaves have write access. The composition should be saved in a folder accessible to all slaves and added to the Queue list using a path visible to all slaves.

For example, imagine that the composition `c:\compositions\nettest1.comp` is open in Fusion. Click on Network Render in the File menu to add the composition to the queue. The Render Manager detects the new entry in the queue and sends a message to each slave to load the composition and render it.

The problem in the above scenario is that each computer is likely to have its own `c:\drive` already. It is extremely unlikely that the composition will be present on each node and, if by some strange chance it is, it is unlikely to be up to date. In most cases, the nodes will report that they are unable to load the composition and quickly fail to join the render.

This will also affect the slave's ability to make use of any pre-rendered disk caches in the composition. If the path to the disk cache is not valid for the slave, the slave will not be able to load the cached frame and will be forced to re-render the cached tools. The solution is to use UNC or mapped paths when loading and saving footage and compositions to disk. These two methods are described below.

Using UNC Names

```
\\file_server\shared_drive\Project6\clipC\ clipC_0001.tga
```

Use UNC (Universal Naming Convention) names when loading or saving compositions and when adding compositions to the Render Manager's queue. UNC names always start with a double backslash (\\), rather than the usual drive letter (for example, c:\ or Macintosh HD/). For more information, see UNC names in the Windows or macOS documentation.

Using Mapped Drives

Mapped drives assign a shared network resource to a letter of the alphabet. For example, the file server's drives could appear as the letter Z on each of the slaves. This requires more setup effort but may allow flexibility when changing to which machine or shared directory the drive letters are mapped, without requiring the compositions themselves to be changed.

Fonts

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 slaves that do not have the font installed.

Plug-Ins

All third-party plug-ins and tools used by a composition must be installed in the plug-ins directory of the slaves. A slave that attempts to render a composition that contains a plug-in it does not have installed will fail to render. Beware of plug-ins that have a demo mode. For example, imagine a composition that uses a plug-in, but the nodes used as slaves in the render farm also require a license and do not have one. In certain cases, the comp will fail to load on those slaves. In certain cases, when the plug-ins are installed in the slaves but not licensed, the comp will load without complaint but produces frames that contain a watermark.

Submitting a Composition

Once the Render Master is set up and some machines have been added to the master's Slave List, start using the network to render compositions, previews, and disk caches. The Start Render dialog used for Flipbook Previews and final renders contains a Use Network checkbox. When this checkbox is enabled, the render task will be submitted to the Render Manager configured in Global > Network Preferences instead of rendered locally.

Select the Use Network checkbox to enable the controls for choosing what group will be used.

Flipbook Previews

Fusion is able to use slaves 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 slaves will render the preview frames to the directory set in the Preferences Global > Path > Preview Renders. This directory should be valid for all slaves 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 slave are spooled into memory on the local workstation. As each frame is copied into memory, it is deleted from disk.

Disk Cache

The dialog used to create disk caches for a tool provides a Use Network checkbox. If this option is selected, click on the Pre-Render button to submit the disk cache to the network.

When Renders Fail

It is a fact of life that render queues occasionally fail. Software crashes, the power goes out, or a computer is accidentally disconnected from the network. These are only a few of the things that can happen in a network. The more complex the topology of the network, and the larger the compositions, the more likely some failure or another will interrupt the render in progress. If someone is monitoring the render as it progresses, this may not be too catastrophic. 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 employs 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

A slave can become unavailable to the Render Master for several reasons. The computer on which the slave operates may lose power or become disconnected from the network. Perhaps the Allow Network Rendering option is disabled. If something like this occurs, the slave will no longer respond to the regular queries sent by the master. Frames assigned to that slave will be reassigned among the remaining slaves in the list.

When the slave is restarted, reconnected to the network, or 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 slave.

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

Select Last Slave Restart Timeout in the Network Preferences to set how long Fusion will wait after the last slave goes offline before aborting that queue and waiting for direct intervention.

FusionServer

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

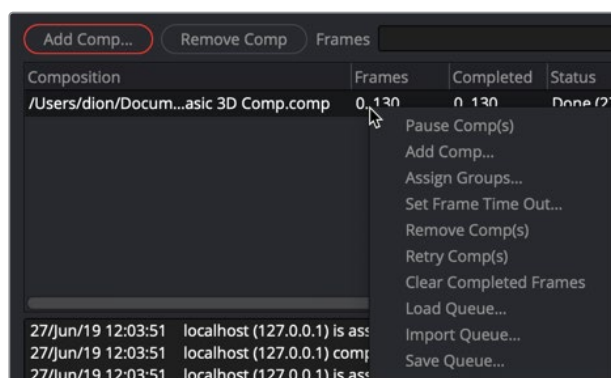
FusionServer monitors the slave to ensure that the slave is still running during a render. It consumes almost no CPU cycles and very little RAM. If the monitored executable disappears from the system's process list without issuing a proper shutdown signal, as can happen after a crash, the FusionServer will relaunch the slave after a short delay, allowing the slave to rejoin the render.

FusionServer will only detect situations where the slave has exited abnormally. If the slave is still in the process list but has become unresponsive for some reason, the FusionServer 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 slave's render if a frame takes longer than the specified time (default 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.

If it is expected to never have a single frame take longer than 10 minutes (for example, 4K DCI sized footage is never rendered), reduce the timeout to 10 minutes instead. To change the frame timeout value, select Set Frame Timeout from the Render Manager's Misc menu.

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 node asking if the node is still responsive and healthy. If the slave fails to respond to several consecutive heartbeats, Fusion will assume the slave is no longer available. The frames assigned to that slave will be reassigned to other slaves in the list.

The number of heartbeats in a row that must be missed before a slave 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 preferred to increase these two settings somewhat to compensate for the sluggish response time until more RAM can be added to the slave.

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 Slaves in the network. Fusion Render Node 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 slave.

To access preferences for a node, right-click on the icon in the taskbar and select Show Network and Memory Preferences.

Override Composition Settings

Enable this option to use the slave'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.

Things You Need to Know

There are a few important issues 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 tool 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 Multiframe Formats

Multiple machines cannot render a single QuickTime file due to the limitations of the file format itself. QuickTime (or AVI) movies are formats that are a single stream containing multiple frames. These frames must be written one after the other and not in the out-of-order sequence typically produced by network rendering. 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 multiframe formats to render successfully. Only final renders are affected by this limitation.

Video Hardware

When network rendering in Fusion, always render to separate sequential file formats like DPX or OpenEXR, for more efficiency, and then convert the image sequence into a format supported by the capture and playback card. Blackmagic DeckLink, UltraStudio, and Intensity devices can be used to output the composition to an external display.

Troubleshooting

There are some common pitfalls when rendering across a network. If some difficulties are encountered, try these steps. If the problem is still not resolved, contact Blackmagic Design's technical support. Ideally, email a copy of the render.log file to support-usa@blackmagicdesign.com to help resolve the problem. Virtually all problems with network rendering have to do with path names or plug-ins. Verify that all slaves can load the compositions and the footage, and that all slaves have the plug-ins used in the composition installed.

Check the Render Log

The log file shown in the Render Manager dialog displays messages that can assist with diagnosing why a render or slave has failed. The Render Log shows a step-by-step account of what happened (or didn't happen) during a render. If a slave 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 Slaves Could Be Found

On the Preferences Network tab, make sure that there is at least one slave available, running and enabled. If all slaves are listed as Offline when they are not, check the network.

The Composition Could Not Be Loaded

Some slaves may not be able to load a composition while others can. This could be because the slave could not find the composition (check that the path name of the composition is valid for that slave) or because the composition uses plug-ins that the slave does not recognize.

The Slave(s) Stops Responding

If a network link fails, or a slave goes down for some reason, the slave will be removed from the active list and its frames will be reassigned. If no more slaves 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 slaves failed and why.

The Slave(s) Failed to Render a Frame

Sometimes a slave simply cannot render a particular frame. This could be because the slave could not find all of the source frames it needed, or the disk it was saving to became 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 slave. If no slave can render the frame, the render will fail. Try manually rendering that frame on a single machine and observe what happens.

Check the Slaves

Fusion's Render Manager incorporates a number of methods to ensure the reliability of network renders. Heartbeats are used to detect network or machine failures. In this event, failed slave's outstanding frames are reassigned to other slaves where possible. In rare cases, a slave may fail in a way that the heartbeat continues even though the slave is no longer processing. If a slave may have 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 slave 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 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 slave 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 slaves manually. If the remote systems do not respond when they are up and running, the network is not functioning and should be examined further.

How to Set Up Headless RenderNode

Requirements

- 1 Headless Linux CentOS 6.6 or later
- 2 Fusion RenderNode 8.2.1 or later

Setting Up CentOS

Make sure the following packages are installed on CentOS

- 1 xorg-x11-server-Xorg
- 2 xorg-x11-drivers
- 3 mesa-dri-drivers
- 4 mesa-libEGL
- 5 mesa-libGL
- 6 mesa-libGLU

Starting Xorg Server with XDummy Driver

For running Fusion RenderNode on a headless machine, we have to use Xorg server with xorg- x11-drv-dummy driver. Follow these steps for that:

- 1 Make sure that none of the proprietary GPU drivers (Nvidia, AMD Radeon, etc.) are installed.
- 2 Make sure that another instance of Xorg server is not running. (Disable autostart of Xorg server if needed.)
- 3 Make sure that the Xorg configuration file is configured to use the dummy driver instead of the default one. Sample Xorg configuration with dummy driver can be found at <http://xpra.org/xorg.conf> (also in APPENDIX I). Save the xorg.conf file to:
/etc/X11/xorg.conf.d/xorg.conf
- 4 Start Xorg server by running the following command as SuperUser # Xorg -noreset +extension GLX

Starting Fusion RenderNode

Before starting Fusion RenderNode, we need to set the DISPLAY environment variable.

```
$ export DISPLAY=:0
```

Once the DISPLAY environment variable is set, start FusionRenderNode using the following command:

```
$ /opt/BlackmagicFusion/FusionRenderNode8/FusionRenderNode
```

Appendix I

Sample xorg.conf

```
Section      "ServerFlags"

    Option "DontVTSwitch" "true" Option "AllowMouseOpenFail" "true"
    Option "PciForceNone" "true" Option "AutoEnableDevices" "false" Option
    "AutoAddDevices" "false"

EndSection


Section "InputDevice"

    Identifier "dummy_mouse" Option "CorePointer" "true" Driver "void"

EndSection


Section      "InputDevice"

    Identifier "dummy_keyboard" Option "CoreKeyboard" "true" Driver "void"

EndSection


Section "Device"

    Identifier "dummy_videocard" Driver "dummy" Option "ConstantDPI" "true"

    #VideoRam      4096000

    VideoRam 256000

    #VideoRam 192000 EndSection


Section "Monitor"

    Identifier "dummy_monitor" HorizSync      5.0      - 1000.0
    VertRefresh 5.0 - 200.0

    Modeline "1920x1200" 26.28 1920 1952 2048 2080 1200 1229 1231 1261
    Modeline "1920x1080" 23.53 1920 1952 2040 2072 1080 1106 1108 1135
    Modeline "1680x1050" 20.08 1680 1712 1784 1816 1050 1075 1077 1103
    Modeline "1600x1200" 22.04 1600 1632 1712 1744 1200 1229 1231 1261
    Modeline "1600x900" 33.92 1600 1632 1760 1792 900 921 924 946

EndSection Section "Screen"

    Identifier "dummy_screen" Device
    "dummy_videocard" Monitor
    "dummy_monitor" DefaultDepth 24
```

Fusion Studio Network Licensing

As of Fusion 9 Studio, a simplified network licensing system via a dongle server setup can be used.

Render nodes are unlimited and the number of Fusion Studios running is determined by the dongle.

Fusion will auto search to find dongle licenses, making it easy to set up. You can specify server addresses as well. There are no license files. This has been automated as much as possible. Fusion can act as a license server as well as a more permanent server-based setup.

These dongles are programmed in 10, 20, and 50 instances of Fusion, and can be combined to tailor the number of Fusion seats in a facility. For example, 3 dongles of 10 Fusions would be 30 licenses served. This also allows for redundancy, as a number of servers can act as license servers or multiple dongles can be plugged into a single machine.

Setup

Install 1 or more dongles on a machine to act as the license server. The Render Node installer will install (and set up to run on startup) Fusion Server, which acts as the license server. The Render Node does not have to run on that machine. Only the Fusion Server will set up as a service/daemon, ready to serve licenses and bins.

Technical Info

- Fusion Studio needs either a single-license dongle on that machine or a multi-license dongle with a spare license on a machine anywhere on that subnet.
- Fusion Studio can act as the license server if the dongle is plugged into that machine.
- This works cross -platform without restriction.
- Render nodes are unlimited and need a single or multi-license dongle somewhere on the subnet.
- Fusion Server is what serves multi-license dongles, so it must be running (it is started automatically by Fusion).
- By default, Fusion Server will stick around as long as Fusion is running locally or if a license is being used by another machine (including Render Node). If nothing is using Fusion Server, it will quit in a few seconds.
- “Fusion Server -S” (capital S) will stick around forever, until force-quit.
- “Fusion Server -i” or “-u” will install/uninstall it as a service or daemon, running it on startup before user login.
- Fusion Server also serves shared bins, which can be connected to remotely.
- Multiple dongles can be installed on a single machine, and licenses will add up.
- Multiple license servers can be present, and Studio will try all of them as needed.
- The dongle, Fusion, and Render Node need to be on the same subnet for Auto License searching.
- The FUSION_LICENSE_SERVER environment variable can now take a semicolon-separated list of IP addresses and host names to scan for license servers—for example:
 - “bobs-mac.local;10.0.0.23;*.license.mystudio.com”
 - including * to indicate broadcast search of the local subnet.

Like most environment variables, you can put them in Fusion.prefs under Global.EnvironmentVars:

```
fu:SetPrefs("Global.EnvironmentVars.FUSION_LICENSE_SERVER",  
"192.168.1.12;*) fu:SavePrefs()
```

- Multiple license servers (semicolon-separated) in FUSION_LICENSE_SERVER env var, like:

```
fu:SetPrefs("Global.EnvironmentVars.FUSION_LICENSE_SERVER", "192.168.1.12;  
192.168.10.55;*)
```
- Removing a dongle or breaking the network will eventually (30-60 secs) drop the licenses of the Studio copies using them.
- Upon losing its license, Studio will start searching for another license, locally or on a different machine.
- If no license can be found, Studio will pause rendering and display a modal dialog, with Retry/Exit options. Retry searches again. Exit will autosave the comp first.
- Render Node just checks on startup for a dongle on the network but does not check afterward. This means it will not be affected by dongle removal or network issues.

NOTE: Users may need to configure their firewall on the server (if they have one) to allow connections for ports. For studios that need network licensing, shared bins, network renders, or multi-machine scripts, the following incoming ports must be opened:

- TCP 1144
- UDP 1144
- TCP 49152-49159
- UDP 49152-49159

Chapter 13

Fusion Connect

Contents

Fusion Connect Overview	300
The Effect Palette	300
The Layer Input Dialog	301
Applying Fusion Connect to a Transition Point	301
Export Clips	301
Edit Effect	302
Browse for Location	302
Auto Render in Fusion	302
Red on Missing Frames	303
Rename Shot	303
Compress Exported Frames	303
Edit Effect also launches Fusion	303
Versioning	303
Create New Version	303
Version	303
About RAW images	303
About Color Depth	304
Warning Overview	304
Manual vs Auto-Render	305
AVID/Fusion Layer to Comp Relationship	306
Fusion/AVID Project Relationship	306
Memory Considerations between Fusion and AVID	307
Rendering with Fusion	308
Directory Structure of Fusion Connect Media	308
Advanced Project Paths	309
Configuring Paths on OS X	309
Configuring Paths on Windows	310
Fields and Variables	310
Environment Variables	311
Definitions	311

Fusion Connect Overview

Fusion Connect is an AVX2 plug-in for AVID Media Composer.

It allows the editor to create a conduit between Avid editing products and the Fusion procedural node visual effects products, connected on the Avid timeline.

Fusion Connect will export clips from the Avid timeline as image sequences and assemble Fusion compositions which then allow you to work your magic on the footage.

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

- **Supported Avid products:** Media Composer 8.x
- **Supported product:** Fusion 8.1 or later
- **Memory recommendation:** 8 GB minimum 16 GB or more recommended
- **Recommended Monitor Resolution:** 1920 x 1080
- **Installation:**
Two files will be installed in your Media Composer:


Fusion Connect.avx

BlackmagicFusionConnect.lua
- **Avid's default directory:** \Avid\AVX2_Plug-ins

The Effect Palette

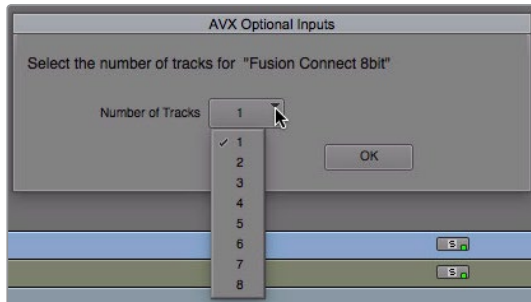
After launching Media Composer, you will find Fusion Connect in the Blackmagic effect palette category.

You can apply (drag and drop) the Fusion Connect AVX2 plug-in to any clip. This includes filler*, edit transition point, or layer stack on the Avid timeline.



The Layer Input Dialog

When you apply the Fusion Connect AVX plug-in to a clip or a layer stack, you are presented with a layer AVX Optional Input dialogue box.




Once the layer count is selected, Fusion Connect will be applied to the timeline.

- Select the layer count to the equal number of layers you want to ingest into Fusion.
- Filler can be used as a layer.
- Fusion Connect will allow a maximum of 8 layers.

Applying Fusion Connect to a Transition Point

If you apply Fusion Connect to a transition point, you will not receive any dialog box, and the AVX plug-in will simply be applied to the timeline transition point.



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

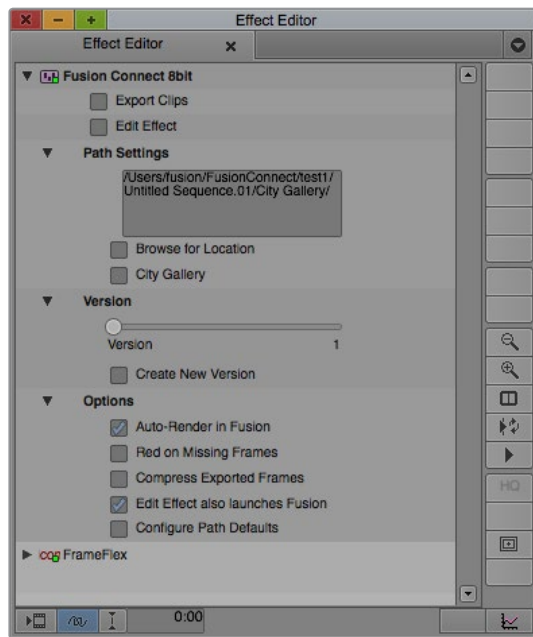
By pressing the Export clips button, Fusion Connect will export all the associated clip(s) as sequences of frames, to provide access to them in Fusion.

Fusion Connect is always live.

Fusion Connect will save 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.

Pressing Export Clips will overwrite any existing images, assuring that all media needed by Fusion is accessible.



Edit Effect

This button performs three subsequent functions.

- 1 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.
- 2 Launch Fusion (if installed on the machine), and if it is not already launched.
- 3 Open the Fusion comp associated with created effects.

Browse for Location

The Fusion Connect media directories will default to the drive where the associated Avid Media resides, defaulting to the root level of that drive. If you wish to choose another location to store and access the media, select the Browse for Location button.

A standard browse requester will appear, allowing you to make your location choice. This includes the ability to create additional directories. The window below will show you the path.

The last selected media path is remembered and is used as a default location when adding a new instance of Fusion Connect. You will be required to browse to any unique directory names associated with the Avid project.

The Path Settings section of the controls may need to be closed and reopened (refreshed) in order to see the updated directory path.

Auto Render in Fusion

This toggle button allows you to automatically render your Fusion comp from within Avid.

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 will also export the necessary media without having to manually execute the Export Clips function first.

Red on Missing Frames

This toggle button will display red images within the Avid timeline viewer (timeline monitor) if no rendered frames from Fusion can be found, or if the rendered frames are not of a high enough resolution.

If disabled, the original untouched frames will be shown instead of the red frame.

Rename Shot

This button opens a standard dialog box, allowing you to change the original Avid shot name to a unique user-defined name.

Compress Exported Frames

When enabled, this button will allow for Fusion RAW files with compression, 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 will launch or switch over to Fusion when the Edit Effect button is clicked, allowing for a more direct Avid/Fusion workflow.

When disabled, this button will not launch Fusion when the Edit Effect button is clicked, but a Fusion .comp file will still be created. This is beneficial for preparatory work prior to a Fusion session or for Avid workstations that are not running Fusion on the same machine but require a remote artist to access the comp and media on a shared storage network.

Versioning

Create New Version

This checkbox will take your current comp and create an exact copy of it, without affecting the original.

Any changes that are rendered in the copy will be written to a new directory and become another version of the rendered result to be played back on the Avid timeline. Previous versions of the comp and their rendered results are preserved.

Version

This slider allows the editor to choose which version of the comp is used to populate in the 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 RAW file image sequence for both directions between Avid Media Composer and Fusion, in order to preserve all the image information. This benefits the editor in numerous ways because the images are not sitting in RAM.

The benefits 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 remotely

About Color Depth

Fusion Connect derives its images directly from the RGB data within Avid MXF 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 16bit float to maintain the full color fidelity supported by Media Composer.


Warning Overview

If your Avid product is in proxy half or quarter resolution mode (notably when you are using a software-only Media Composer), you will receive the above dialog warning when clicking Export Clips.

This indicates that exported frames will not be full resolution, which will mean that rendered frames from Fusion will not be full resolution either.

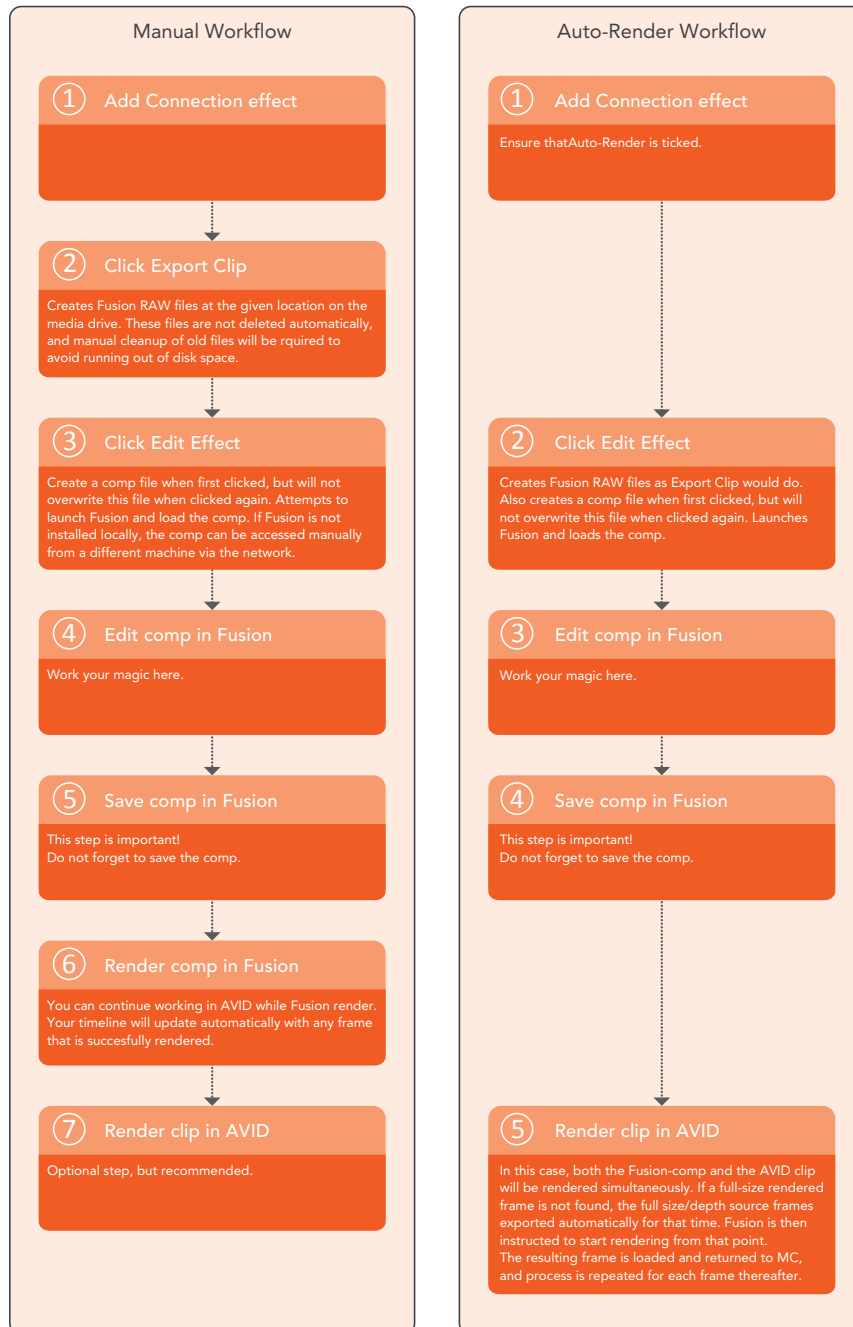
To clear the warning and export full resolution media, click Cancel to abort the export and clear the Fusion Connect warning message.

In the Media Composer timeline window, set the timeline to full resolution, and then click on the Export button in the Effect Editor once again.



Manual vs Auto-Render

The following diagram shows typical workflows for manual and automatic renders.



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 is not required to be installed on the AVID system itself but can sit on any other machine on the network, and artists other than the AVID editor can work on the compositions.

For Auto-Render, Fusion must be installed on the local machine.

AVID/Fusion Layer to Comp Relationship

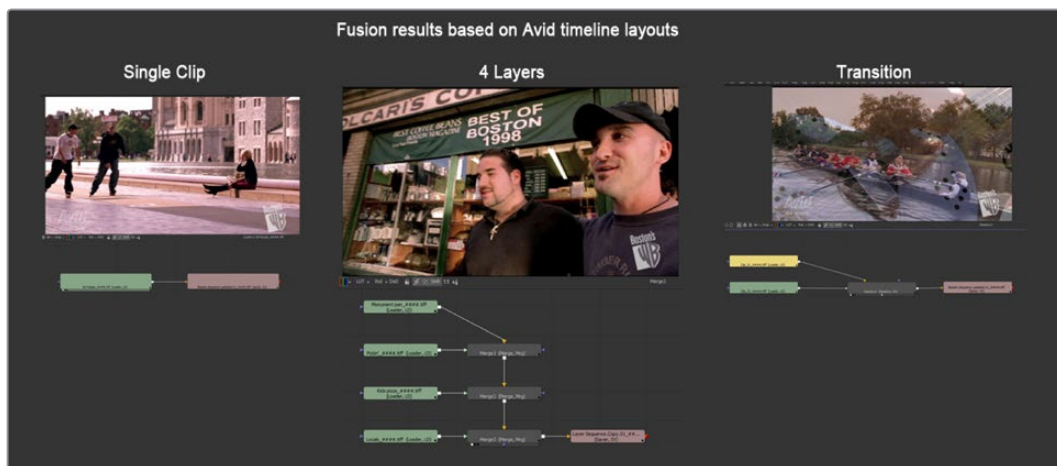
Once the initial trip from Avid to Fusion is complete, depending on the type of clip to which you assigned the effect, in Fusion you will be presented with either:

- A Loader representing a single clip.
- Two or more Loaders connected to Fusion Merges representing layers.
- Two Loaders connected to a Dissolve representing a transition.

You will also see a Saver. The Saver is directly linked to the directory that is connected back to the Media Composer timeline.

Do not change the file format or directory path of the Saver.

Due to the design of the AVX2 standard, hidden embedded handles are not supported. To add handles, increase the length of your clip in the Media Composer timeline to include the handle length.



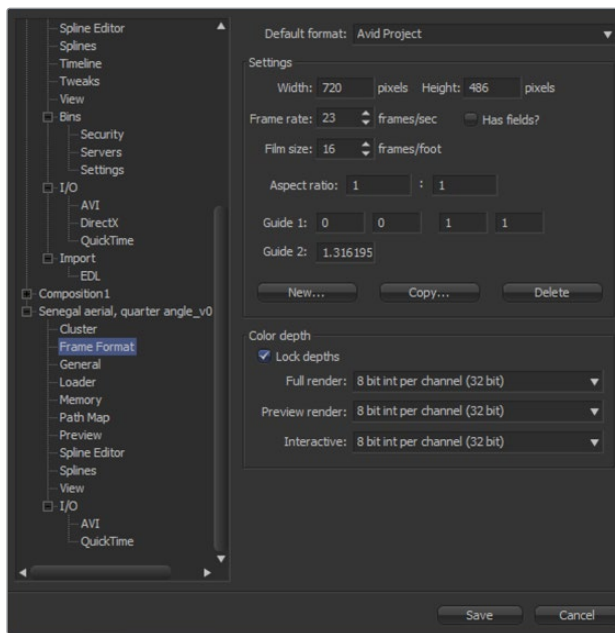
TIP: move your longest clip to the top layer and apply Fusion Connect to that clip.

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.

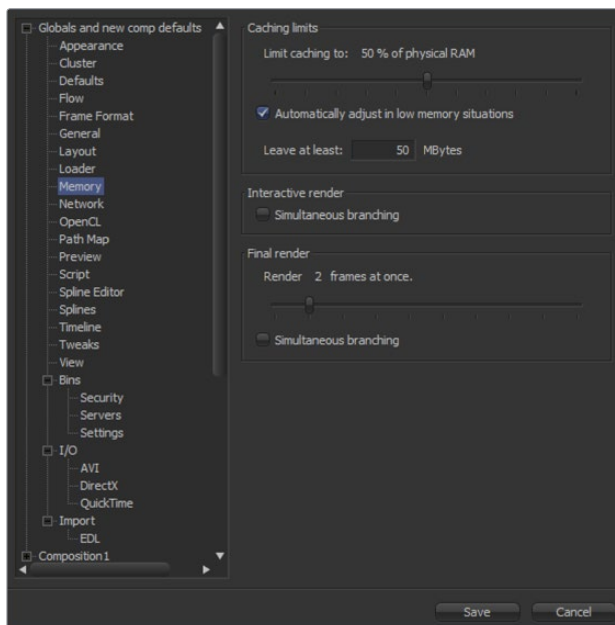


Memory Considerations between Fusion and AVID

The frame rate and image size preferences created in Avid products 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.

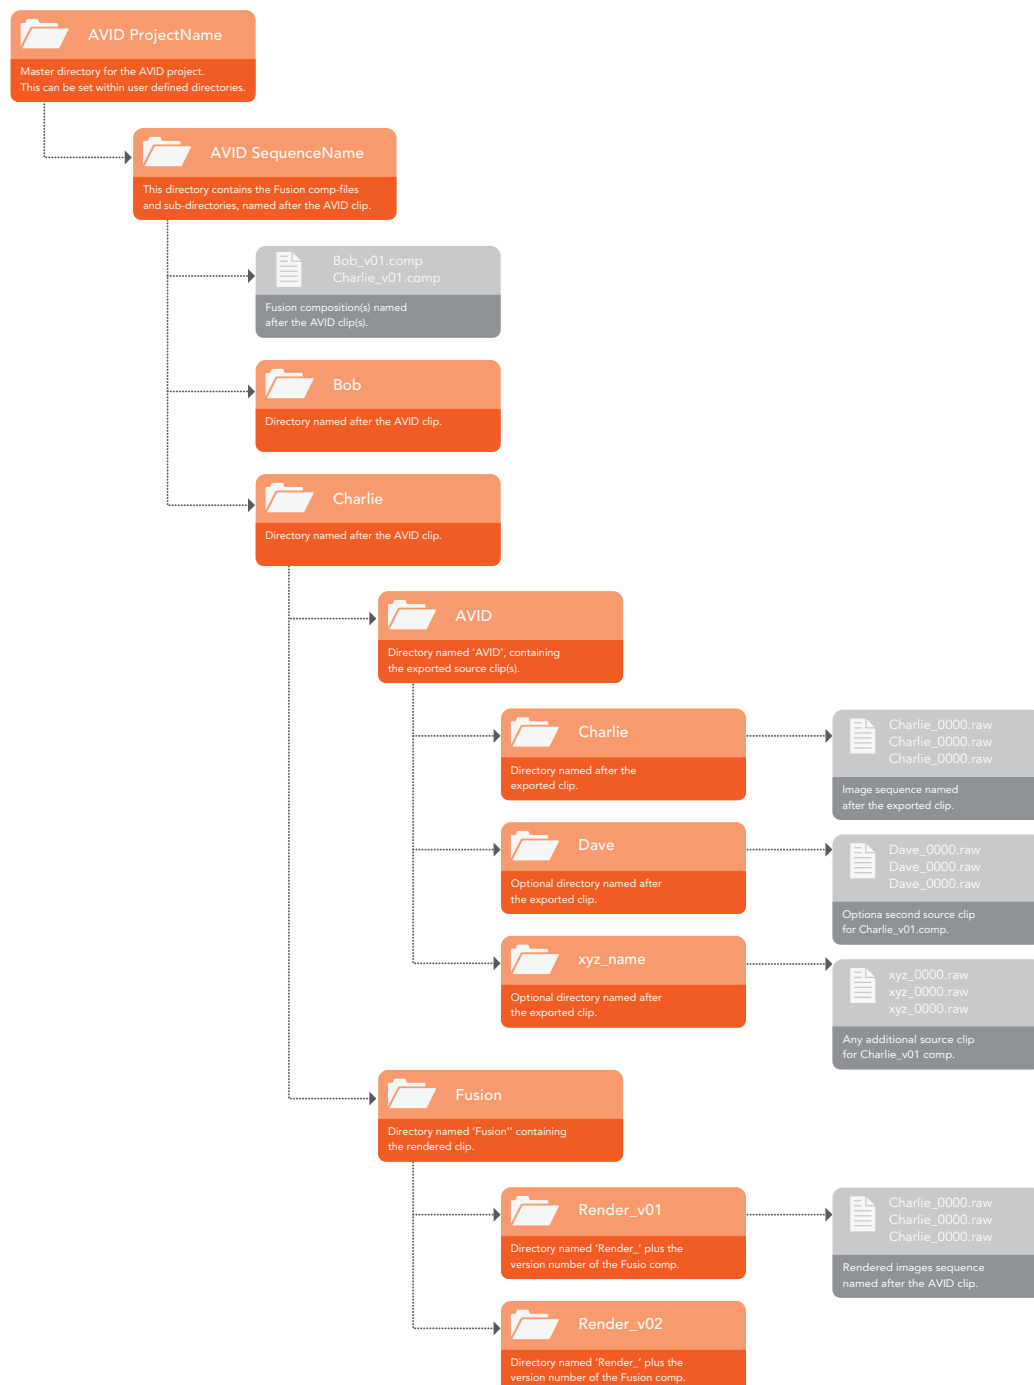
The format settings do not prevent you from using or mixing any other sized imaging within the composition as Fusion is truly resolution independent.



Rendering with Fusion

When you perform a render of your comp inside Fusion, the results are rendered to the output directory created by Fusion Connect during the initial application of the plug-in to the timeline. Upon rendering, you will immediately be able to see the results of the rendered Fusion comp if you switch back to Avid. 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 directory 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 directory hierarchy is automatically created based on Avid projects, sequences, and clips. This structure allows for multiple instances of Fusion to access the media and multiple instance 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 directories with matching names. If such directories 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 directory name, if a directory of that name already exists. The corresponding comp file will be named accordingly.



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.

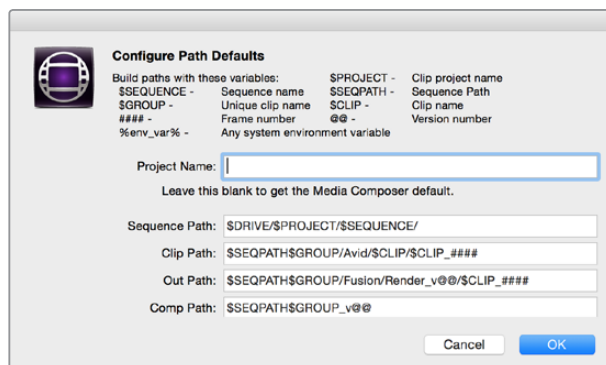
It's recommended that you render the green dot effect, which will force an MXF precompute to be created to guarantee real-time playback.

Advanced Project Paths

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 Fusions running on other systems.

Configuring Paths on OS X

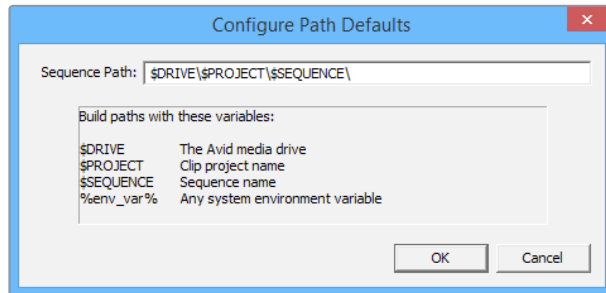
When using Fusion Connect on OS X, the Configure Path Defaults dialog looks like this:



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:



In 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. It is located in the Options section of the Fusion Connect AVX2 plug-in. You must click on the triangle to see it.

Fields and Variables

The fields and variables that can be used on OS X and Windows are described below:

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 Variable on Windows

The quickest way to access it is through your Windows control panel and search 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.”

TIP: System variables control the environment throughout the operating system, no matter which user is logged in.

Definitions

- **Variable** is the control title defined specifically by the application that is being controlled.
- **Value** is the instructions that tell the variable what to do.
- **Variables** and what they mean.

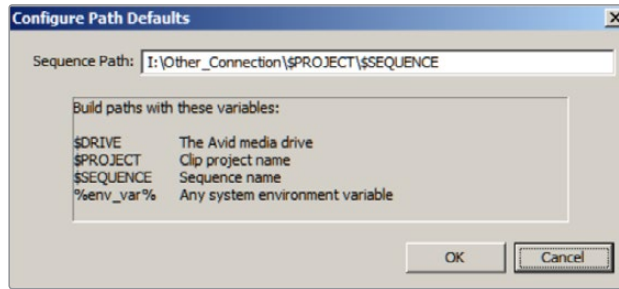
When typing the variable it must be typed exactly as shown.

TIP: If you type this 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 clear (get rid of) a variable, you must exit the Media Composer and clear the environment variable in the Windows interface and restart the Media Composer.


Other Values you can control derived from your AVID Bin:

- **\$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 setup to support project and sequence names within your directory.



Here is the same example in the Windows environment variable editor





PART 4

2D Compositing

Chapter 14

Learning to Composite in Fusion

This chapter is a grand tour of the basics of Fusion, walking you through the process of ingesting a clip into Fusion, and then working in the Node Editor to create some simple effects. Subsequent topics build upon these basics to show you how to use the different features in Fusion to accomplish common compositing and effects tasks. In the process you'll learn how node trees are best constructed, and how to use the different panels of Fusion together to work efficiently.

Contents

What's a Composition?	317
How Nodes Are Named	317
Applying and Masking Effects	317
Adding a Node to the Tree	317
Editing Parameters in the Inspector	318
Replacing Nodes	319
Adjusting Fusion Sliders	320
Masking Node Effects	321
Compositing Two Clips Together	324
Adding Additional Media to Compositions	324
Automatically Creating Merge Nodes	324
Adjusting the Timing of Clips	325
Fixing Problem Edges in a Composite	326
Apply Modes and the Corner Positioner	328
Setting Up the Initial Composite	328
Controlling the Viewers	329
Displaying Images in the Viewers	329
Adding the Corner Positioner Node with a Search	330
Warping the Image with the Corner Positioner Node	332
Toggling On-Screen Control Visibility	332
Navigating the Viewer	333
Using the Screen Apply Mode in the Merge Node	333
Tweaking Color in the Foreground Layer	334
Creating and Using Text	336
Creating Text Using the Text+ Node	336
Styling and Adjusting Text	337
Using One Image's Alpha Channel in Another Image	338
Using Transform Controls in the Merge Node	343
Match Moving Text With Motion Tracking	343
Adding a Layer We Want to Match Move	343
Setting Up to Track	344
A Simple Tracking Workflow	345
Connecting Motion Track Data to Match Move	348
Offsetting the Position of a Match-Moved Image	350

Using Paint and Planar Tracking	351
Using a Planar Tracker to Steady a Subject to Paint	351
Painting Over Blemishes	354
Inverting the Steady Effect to Put the Motion Back In	357
Fixing the Edges by Using Only the Fixed Part of the Frame	359
Building a Simple Green Screen Composite	363
Pulling a Greenscreen Key Using the DeltaKeyer	363
Using the Transform Node to Resize a Background	367
Masking a Graphic	369
Animating an Image Using Keyframes	371
Animating a Parameter in the Inspector	372
Using the Spline Editor	373

What's a Composition?

A “composition” describes the collection of nodes that create an effect in Fusion. This typically includes Loaders that are used to access source elements, tools that are used to manipulate and combine the elements in various ways, and finally one or more Savers to output the final shot.

Once you bring your footage into Fusion via a Loader node, selecting the node displays its parameters in the Inspector to the right. Pressing 1 or 2 will display it in viewer 1 (on the left) or viewer 2 (on the right).

At this point, you're ready to start compositing.

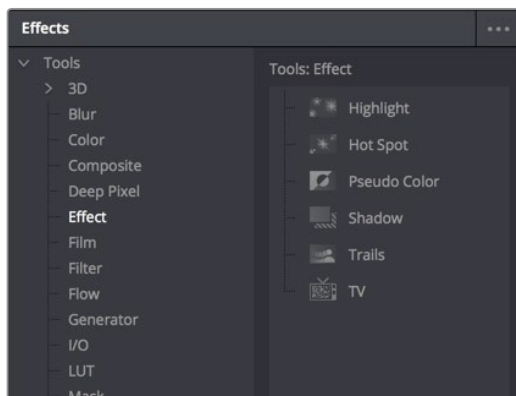
How Nodes Are Named

While the documentation refers to nodes by their regular name, such as “Medialn,” 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.

Applying and Masking Effects

Let's begin by looking at some very simple effects, and build up from there. Opening the Effects Library, and then clicking the Disclosure control to the left of Tools reveals a list of categories containing all the effects 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. For now, we're interested in the TV effect.



Browsing the Effect category to find the TV node.

Adding a Node to the Tree

Assuming the Loader node is selected in the Node Editor, clicking once on the TV node in the Effects Library automatically adds that node to the node tree to the right of the selected node. With the TV node selected, pressing 2 will display the results of that node in viewer 2.



A new node added from the Effects Library

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

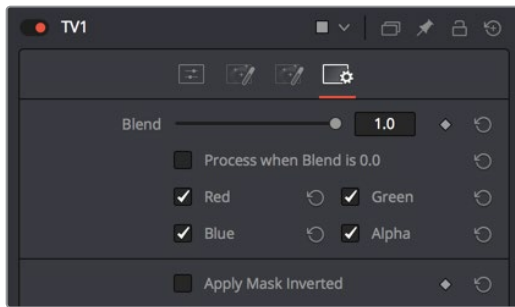
Looking at the TV effect in the viewer, you may notice a lot of transparency in the image because of the checkerboard pattern. If you don't see the checkerboard pattern in the viewer, it might be turned off. You can turn it on by clicking the viewer Option menu and choosing Checker Underlay.

To improve the effect, we'll make an adjustment to the TV1 node's parameters in the Inspector at the right. Whichever node is selected shows its controls in the Inspector, and most nodes have several panels of controls in the Inspector, seen as little icons just underneath that node's title bar.



The Inspector showing the parameters of the TV effect.

Clicking the last panel 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 checkboxes, are found. These let you choose which image channels are affected, and let you blend between the effect and the incoming source.



The Settings panel, which has channel limiting and mask handling controls that every node shares.

In our case, the TV effect 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 the Alpha checkbox off results in a more solid image, while opening the Controls panel (the first panel) 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

That was fun, but having previewed this effect, we decide we want to try something different with this shot. Going back to the Effect category of the Effects Library, there is a Highlight node we can use to add some pizzazz to this shot, instead of the TV node.

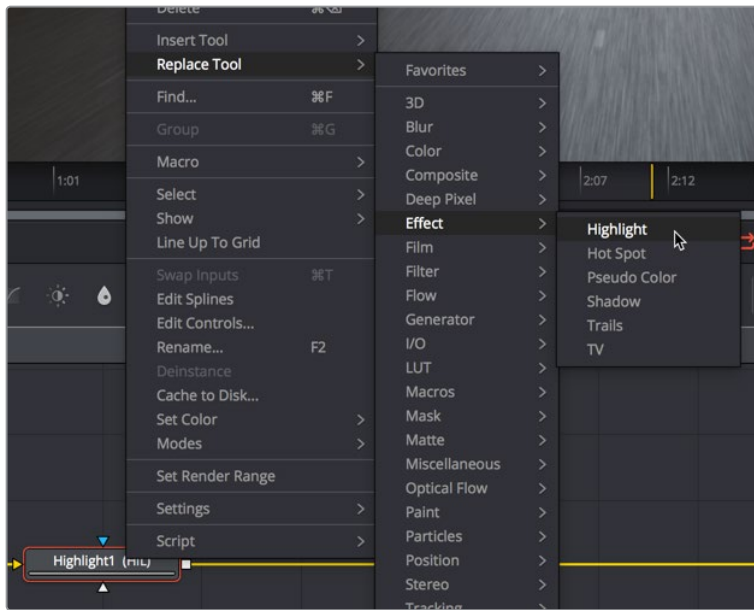
Instead of clicking the Highlight node, which would add it after the currently selected node, we'll drag and drop it on top of the TV1 node in the Node Editor. A dialog appears asking "Are you sure you want to replace TV1 with Highlight?" and clicking OK makes the replacement.



Dragging a node from the Effects Library onto a node in the Node Editor to replace it.

A 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 image's case consists of star highlights over the lights in the image.

Incidentally, another way you can replace an existing node with another type of node in the Node Editor is to right-click a node you want to replace and choose the new node you want from the Replace Node submenu of the contextual menu that appears.



Right-clicking a node to use the contextual menu Replace Node submenu.

It's time to use the Inspector controls to customize this effect, but first let's take a look at how sliders in Fusion work.

Adjusting Fusion Sliders

When you drag a slider in the 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 particular to the parameter you're adjusting. In this case, the Number of Points slider maxes out at 24. However, you can re-map the range of many (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.

Masking Node Effects


Going back to the Length slider, increasing its value gives us a nice big flare.



The Highlight effect with a raised Length value (zoomed in).

This is a nice effect, but maybe we only want to apply it to the car in the foreground, rather than to every single light in the scene. This can be accomplished using a Mask node connected to the Effect input of the Highlight node. The Effect Mask input is a blue input that lets you use a mask or matte to limit that node's effect on the image. Most nodes have an Effects Mask input, and it's an enormously useful technique.

However, there's another node input that's more interesting, and that's the gray Highlight Mask input on the bottom of the node. This is an input that's specific to the Highlight node, and it lets you use a mask to limit the part of the image that's used to generate the Highlight effect.



The blue Effect input of a node is on top, and the gray Highlight Mask input that's specific to the Highlight node is on the bottom.

Adding a Mask Node

To see the results of using either of these two inputs, let's add a mask, this time using the toolbar, which presents a collection of frequently-used mask nodes that we can quickly create.



Clicking the Ellipse button on the toolbar.

With the Highlight node selected already, clicking the Ellipse button (the circle) automatically creates an Ellipse1 node that's connected to the blue Effect Mask input. Creating new masks while a node is selected always auto-connects to that node's Effect Mask input as the default behavior.




Automatically connecting an Ellipse node to the blue Effect Mask.

Adjusting Mask Nodes

Masks, in Fusion, are shapes you can either draw or adjust that have a special single-channel output that's meant to be connected to specialized mask inputs, to either create transparency or limit effects in different ways as described above. With the Ellipse1 node connected and selected, a round on-screen control appears in the viewer that can be adjusted in different ways.

- Drag on the edges of the mask to reshape it.
- Drag the center handle to reposition it freely.
- Drag the up or right arrows to reposition it constrained vertically or horizontally.
- Drag the top, bottom, left, or right sides of the ellipse to stretch it vertically or horizontally.
- Drag any of the corners of the ellipse to resize it proportionally.

Resizing the ellipse to hug only the headlights of the main car, you can see that using the Effect Mask cuts off the long flares we've created, because this masks the final effect to reveal the original image that's input into that node.



The result of using the Effect Mask input.


Reconnecting Node Connections to Different Inputs for a Different Result

This isn't satisfactory, so we drag the connection line attaching the Ellipse node off the Effect Mask input and onto the Highlight Mask input underneath. It's easy to reconnect previously connected nodes in different ways simply by dragging the second half of any connection (it highlights when you hover the pointer over it) to any other node input you want to connect to.




Dragging a connection from one node input to another.

After you make the connection, the connection line goes back to the top of the node, and the top connection is now gray. This is because node inputs in Fusion automatically rearrange themselves to keep the node tree tidy, preventing connection lines from overlapping nodes unnecessarily and creating a mess. This may take a bit of getting used to, but once you do, you'll find it an indispensable behavior.



The Ellipse node now connected to the Highlight Mask input, which has moved to the top of the node to keep things neat.

Now that the Ellipse1 node is connected to the Highlight Mask, the tight mask we've created just around the car headlights restricts the node in a different way. The Highlight Mask lets you restrict which part of the image is used to trigger the effect, so that only the masked car headlights will generate the Highlight effect in this filter. The result is that the flares of the Highlight effect themselves are unhindered, and stretch well beyond the boundaries of the mask we've created.



The Highlight effect is uncropped because the effect is being limited via the Highlight Mask input, rather than the Effect Mask input.

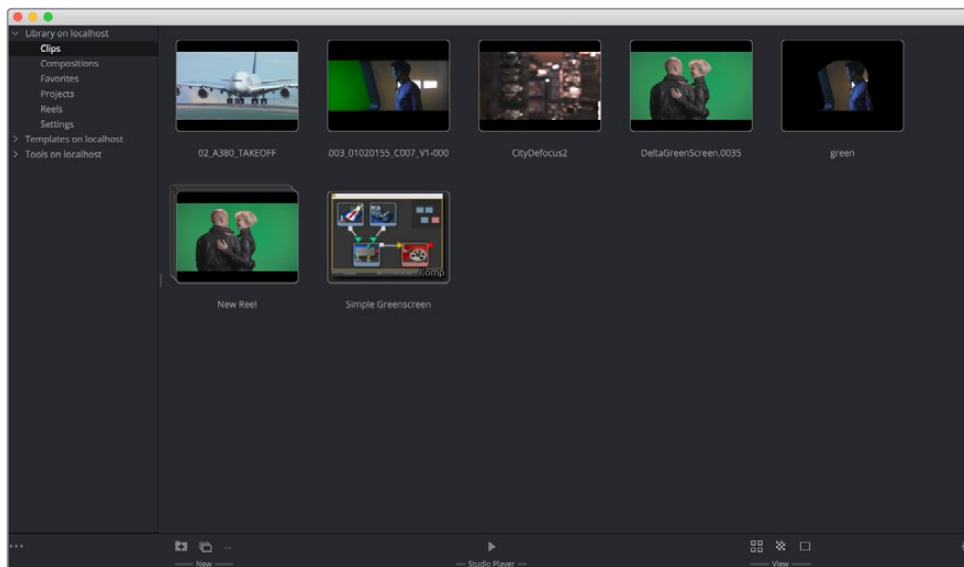
Nodes may have any number of inputs particular to what that node does. This example should underscore the value of getting to know each node's unique set of inputs, in order to best control that node's effect in the most appropriate way.

Compositing Two Clips Together

As entertaining as it is adding individual nodes to create simple effects, eventually you need to start adding additional layers of media in order to merge them together as composites. Let's turn our attention to another composition in which we need to combine a background clip with a foreground clip that's been handed to us that already has a built-in alpha channel, to see a simple composite 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. For this reason, you can add as many Loaders as you require by clicking the Loader icon in the toolbar or by dragging clips directly from the macOS Finder or a Windows Explorer to the Node Editor. You can also add clips to the Bin window and drag them to the node Editor as you need them.




The Bin window can store clips for later use.

If no other node is selected in the Node Editor when you add a Loader node or drag a clip from the operating system to an empty area of the Node Editor, you'll add an unconnected Loader node (incremented to keep it unique) that you can then connect in any way you want.

Automatically Creating Merge Nodes

However, there's a shortcut 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 the clip right on top of any connection line. When you drop the resulting node, this automatically creates a Merge1 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 MediaIn2 node that represents the clip you've just added.



The selected node is the background to a Merge (left), and an added Loader becomes the foreground to a Merge (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, and a Merge node will be automatically created to build that composite. Remember, background inputs are yellow, 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).

Adjusting the Timing of Clips

In our example composite, the Loader1 node that's connected to the Merge1 node's foreground input has an alpha channel. This simple Merge node composite automatically creates a result that we can see in the viewer, but the way the two clips line up at the beginning of the composition's range in the Time Ruler is not great, because the Loader1 node's clip is being lined up with the very first frame of the Loader2 clip. All clips added to a composite have their start at the beginning of the comp on frame 0. This is not always desirable based on the different durations and timing of elements.

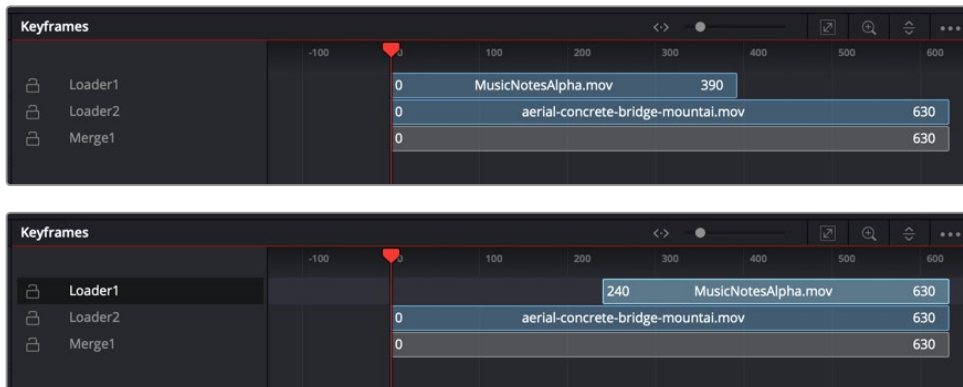


The composite is good, but the timing of the foreground clip relative to the background clips is not ideal.

Since it is rare to have all clips equal in length, you can slip clips and trim their In and Out points using the Keyframes Editor, which can be opened via a button in the UI toolbar.

The Keyframes Editor shows each Loader and Effect node as a bar in a vertical stack that shows you the relative timing of each clip and effect. Keep in mind that the vertical order of these layers is not indicative of which layers are in front of others, as that is defined by layer input connections in the Node Editor. The layers displayed in the Keyframes Editor are only intended to show you the timing of each composited clip of media.

In this case, we can see that the Loader 1 node is offset to the left, so it's easy for us to drag it to the right, watching the image in the viewer until the frame at the composition In point is what we want.



The original stack of layers (top), and sliding the MediaIn2 layer to line it up better with the MediaIn1 layer (bottom).

As a result, the Loader2 clip lines up much better with the Loader1 clip.



After sliding the MediaIn2 clip to improve its alignment with the other clip in the composite.

Fixing Problem Edges in a Composite

Most of the time, 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. 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 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 going to be 0. This obscures the background (by multiplying with the inverse of the foreground alpha), 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 an operation that 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.

Apply Modes and the Corner Positioner

In this next compositing example, we'll explore how you can use the Corner Positioner node to cornerpin warp a composited layer into place as a screen replacement. Then we'll use an Apply mode in the Merge node to refine the screen replacement effect to incorporate real reflections from the scene.

Setting Up the Initial Composite

The base image in the Loader1 node will be the background for our composite. A second Loader will be used as a new screen for our computer. You can merge the two elements together by dragging the output of Loader2 (foreground) to the output of Loader1 (background). A Merge node will automatically be created with the two elements connected for you.



Adjusting the Edit sizing of a clip before moving it into the Fusion for compositing.

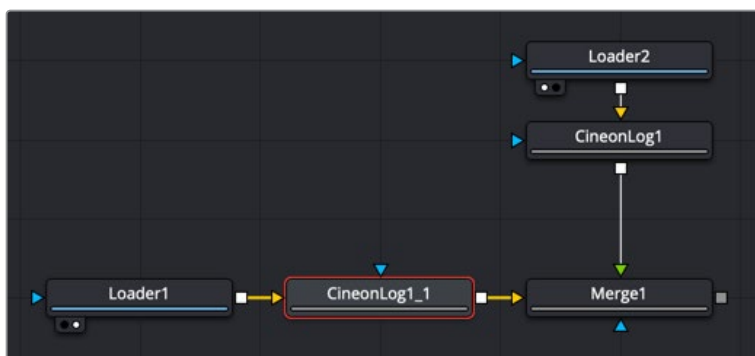
Next, we'll start by setting up some routine color management in the node tree, to illustrate how this should be handled.

Removing Log Gamma Curves

The foreground loader2 has been converted to linear gamma by adding a Cineon Log node. You commonly use this node to remove the Log gamma curve found in RAW clips with wide dynamic range. However, since the background Loader1 is also a RAW file with a log gamma curve, we find we need to add another copy of the same Cineon Log node after the new Loader2 node.

Happily, this is easy to do by selecting and copying the Cineon Log node that's connected to the Loader2 node (Command-C), then selecting the Loader1 node, and pasting (Command-V). When you paste one or more nodes while a node is selected in the Node Editor, the nodes you paste are inserted onto the connection line from the selected node's output. You can tell when a node has been copied and pasted because it shares the same name as the copied name, but with a “_#” appended to it.





Copying a node from one part of the node tree (top), and pasting to insert it after a selected node (bottom).

If you need to output this composite to HD, you need to convert the final composite to the HD rec 709 color space and gamma. You can do this by adding a Gamut node just before the Saver node, and setting the Gamut node Output Space pop-up menu to Rec 709 (Display) or (Scene) depending on whether you want gamma set to 1.98 or 2.2, respectively.

TIP: If you paste one or more nodes while no nodes are selected, then you end up pasting nodes that are disconnected. However, to control where disconnected nodes will appear when pasted, you can click the place in the Node Editor where you'd like pasted nodes to appear, and when you paste, the nodes will appear there.

Controlling the Viewers

Viewing the Gamut node will allow you to see a normalized HD image in the viewer; however, displaying any other node in the viewer will show a linear gamma image. Since viewing linear gamma is not suitable when working on these images, the viewer must be adjusted with a LUT. Above viewer 1, click the LUT button, and then using the pop-up menu next to the LUT button, select Gamut View LUT. Again from the pop-up menu, choose Edit to open the View LUT edit dialog. In the dialog, set the Output Space to ITU-R-BT 709 (Scene) and make sure the Add Gamma checkbox is enabled. Now viewer 1 is set to normalize the linear images, and viewer 2 is configured to display HD images. In some cases you may decide not to add a Gamut tool before the Saver until you are ready to render the final composite. This way, you can set both viewers to display normalized linear images, making it easier to view any image from the composite in either viewer.

Displaying Images in the Viewers

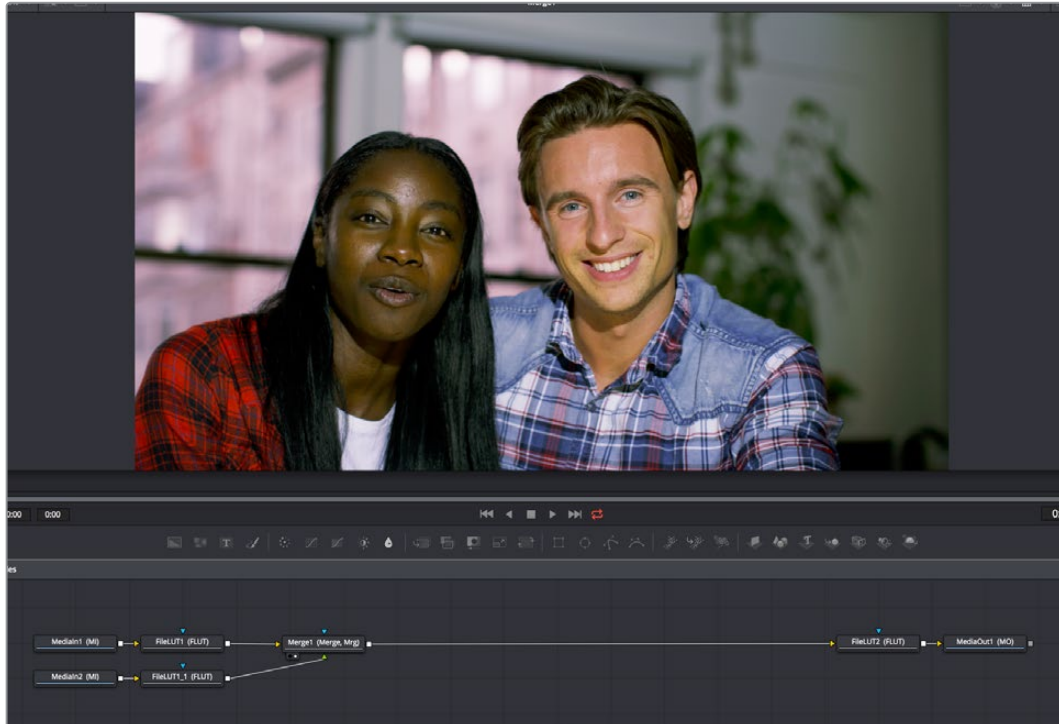
Now that our setup is complete, we can begin looking at our images by loading them into the viewers.

Happily, there are a wide variety of ways you can load a particular node into the viewer to see what you're doing as you work:

- 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/LeftView/RightView in 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).

Using any of these methods, we load Merge1 into the viewer.



Loading a node from the middle of the node tree into the viewer to see a specific node you're working on.


We can tell which node is loaded into the viewer because of the viewer indicators/buttons at the bottom left of the node. Not only is this a visual indication of which node is being viewed, but these buttons can be clicked to load that node into the left or right viewer, if you go into Dual-viewer mode.



A pair of buttons at the bottom-left of nodes that are loaded into the viewer let you see which node is being viewed, as well as giving you a click-target for reassigning that node to another viewer.


Adding the Corner Positioner Node with a Search

Now that we have a foreground image composited over the background image of the computer screen, it's time to reposition the foreground layer to fit into the screen. To do so, we'll use the Corner Positioner node, from the Warp category, which is the main node for doing cornerpinning. To add this to the node tree, we'll use a different method to search for the node we want right from the Node Editor. First, select the node you want to insert a new node after. In this case, we want to cornerpin the image from the Loader2 node, so we'll select the CineonLog node that's attached to it.




Selecting a node you want to add another node behind.

Next, pressing Shift-Spacebar opens the Select Tool dialog. Once it appears, just start typing the first few letters of the name of the node you're looking for to find it in a master list of every node in Fusion. In this case, you're looking for the CornerPositioner node, so type "cor" and the list of nodes will shrink to two, with the one we're looking for being selected.



Pressing Shift-Spacebar opens the Select Tool dialog for quickly finding and adding new nodes.

With the node we're looking for found and selected in the Select Tool dialog, pressing the Return key inserts the new Corner Positioner node after the previously selected node and closes the Select Tool dialog.



The CornerPositioner node added to cornerpin the foreground image prior to the Merge operation.

Warping the Image with the Corner Positioner Node

The Corner Positioner node is a node in the Warp category of the Effects Library that lets you do absolute positioning at four corner points to fit an image within a rectangular region into a scene. Immediately upon adding this node, a default cornerpin operation is applied to the image to show that it's being manipulated.



The Corner Positioner node adds a default transform to the image.


Using the on-screen control points, we can now warp the image by dragging each corner to fit within the computer screen.



Using the CornerPositioner node to fit the video image to the screen it's replacing.

Toggling On-Screen Control Visibility

It's worth knowing that you can toggle the visibility of on-screen controls using Show Controls in the viewer Option menu or by pressing Command-K. You might find it useful to hide on-screen controls if they're getting in the way of seeing the image you're adjusting, but if you've added an effect and you don't see any controls available for adjusting it, you'll know you need to turn this option on.



Show Controls in the Option menu toggles on-screen control visibility on and off.

Navigating the Viewer

As you work, you may find that parts of the image you want to work on extend off screen. To deal with this, there are a few ways of panning and zooming around the viewer.

- Middle-click and drag to pan around the Node Editor.
- Press Command and Shift and drag to pan around the Node Editor.
- Hold the Middle and Left buttons down simultaneously and drag to zoom in to or out of the Node Editor.
- Hold down the Command key and use your scroll wheel to zoom in and out of the Node Editor.
- Swipe with two fingers on a multi-touch trackpad.

Using the Screen Apply Mode in the Merge Node

Once the foreground input image is fit to the computer screen, we have an opportunity to create a more convincing composite by taking advantage of the reflections of the scene on the front of the screen, and using the screen Apply mode to make the 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, including an Apply mode pop-up menu that has a selection of composite modes you can use to combine the foreground and background layers together, and a Blend slider you can use to adjust how much of the foreground input image to merge 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 Apply mode is perfect for simulating reflections, and lowering Blend a bit lets you balance the coffee cup reflections from the display in the background with the image in the foreground. It's subtle, but helps sell the shot.



The original composite (left), and the composite using the Screen Apply mode (right).

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.

Tweaking Color in the Foreground Layer

It's as important to make sure that the color matches between two images being composited together as it is to create a convincing blend, and for this reason Fusion has a whole category of color adjustment tools available in the Color category of the Effects Library. In fact, the ColorCorrector, ColorCurves, HueCurves, and Brightness/Contrast nodes are considered so important that they appear in the toolbar.



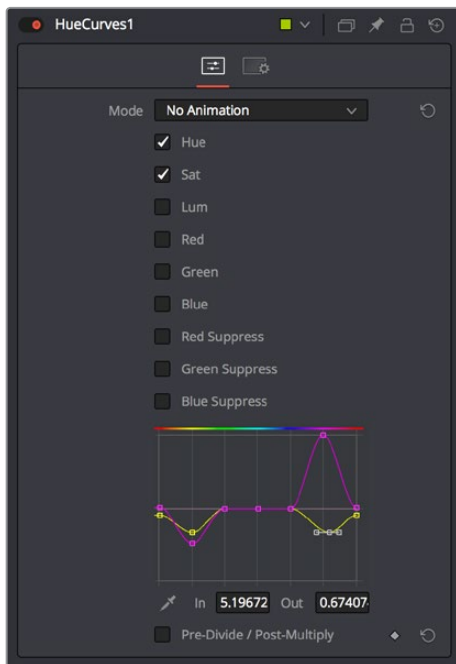
Frequently used Color nodes in the toolbar.

In this particular image, the color of the foreground image on the computer screen is just a bit green and oversaturated, and the view out the window is a bit magenta. However, these problems are easily overcome using a Hue Curves node from the toolbar. Selecting the Corner Positioner node we added, and clicking the Hue Curves button on the toolbar adds that node between the Corner Positioner and the Merge node.



Adding the HueCurves node to make a correction to the foreground image.

The Hue Curves node exposes a curve control in the Inspector with options for adjusting nine kinds of curves, each overlapping the others for simultaneous adjustment. By first turning on the Hue checkbox to make adjustments, and then the Sat checkbox in the Inspector, these two curves can be simultaneously adjusted to push the green toward a healthier red in the skin tones of both the man and the woman, to desaturate the red, yellow, and green a bit, and to push the magenta outside the window to more of a warm orange light, to make the foreground seem like a natural part of the image.



The controls of the HueCurves node, adjusted to correct the screen replacement image.

The result is subtle, but it's a much more convincing composite.



The uncorrected foreground (left), and using a Hue Curves node to adjust the image for a better composite (right).

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 Fusion. You can access it from the Generator category in the Effects Library or more directly from the toolbar. The Text+ node is an incredibly deep tool for creating text effects, with six panels 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 next demonstration.

With the Loader1 node that will serve as our background selected in the Node Editor, clicking the Text+ button automatically creates a new Text+ node connected to the foreground input of a Merge node.



Selecting a node you want to append another node to (left). Clicking the Text+ button on the toolbar automatically creates a Merge composite with the text as the foreground input connection (right).

Selecting the Text1 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 us type directly into the viewer, so we type “SCHOOLED” into the Styled Text field, since that’s the title of the program we’re working on (in case you didn’t know).



The viewer toolbar for the Text node with tools for text entry, kerning, and outline controls.

The text appears in the viewer, superimposed against the background clip. On-screen controls appear that let us 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. At this point, we have our basic title text.



Text that's been typed into the viewer, with on-screen text transform controls.

Styling and Adjusting Text


Now we need to style the text to suit our purposes, so we'll use the controls in the Inspector, increasing Size and decreasing Tracking to move the letters closer together so they can be made larger.



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.


The result has somewhat uneven kerning, but we can adjust that. Selecting the Manual Kerning tool in the viewer toolbar (second tool from the left) reveals small red dots underneath each letter of text.



The Manual Kerning tool in the viewer toolbar.

Clicking a red dot under a particular letter puts a kerning highlight over that letter. Here are the different methods you can use to make manual kerning adjustments:

- Press the Left or Right Arrow keys or hold Command and press the left or right arrows for moving in smaller increments.
- 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.
- 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.



Once the little red dot is selected, Left and Right Arrow keys manually adjust kerning left or right.

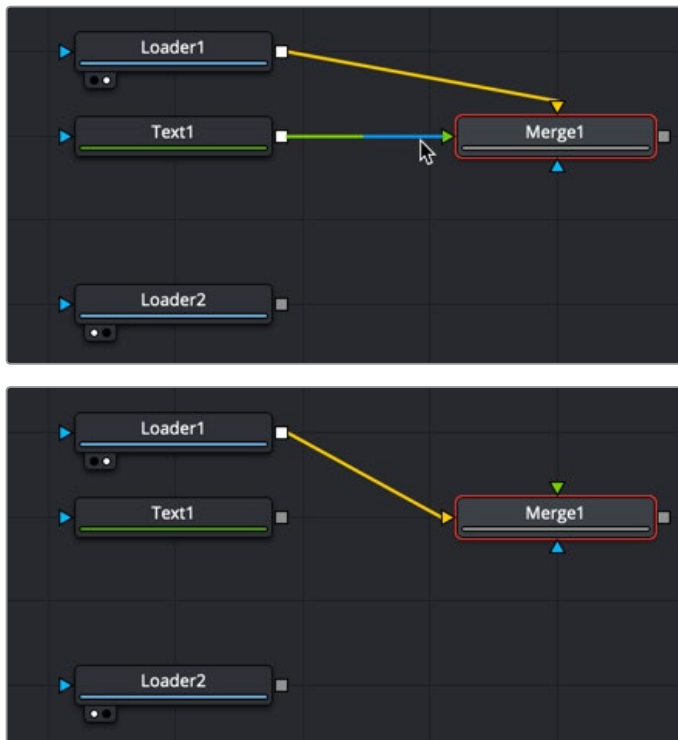
So there we go; we now have a nice title, styled using the Viewer tools and Inspector controls on the Text panel. This looks good, but we have much grander designs.

Using One Image's Alpha Channel in Another Image

We're not going to use the text we've created as a title directly. Instead, we'll use the text as a matte to cut these letters out of another layer we'll be using for texture. So, first we'll add another Loader to ingest our fill element.

Disconnecting and Reconnecting Nodes

Now we need to do a little rearranging. We'll move the Merge1 node up, and then click the last half of the connection from the Text1 node to the Merge foreground input to disconnect it.



Clicking the second half of a connection to disconnect it (left), and the result with the text node disconnected (right).

Next, we'll drag a connection from the Loader2 node onto the Merge1 node's foreground input, so the entire viewer becomes filled with the chalkboard (assuming we're still viewing the Saver node or the Merge). At this point, we need to insert the Text1 node's image as an alpha channel into the Loader2 node's connection, and we can do that using a MatteControl node.

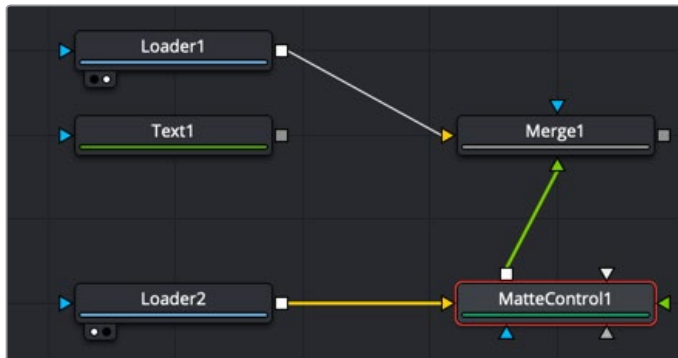


The updated composite, with two video images connected and the text node disconnected.

Using Matte Control Nodes

Selecting the Loader2 node, we click the Matte Control button of the toolbar to add it between the Loader2 and Merge1 nodes (to tidy things up, I've moved the nodes around a bit in the screenshot).

The MatteControl node has many, many uses. Among them is taking one or more masks, 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 (yellow) input of the MatteControl node, as seen in the following screenshot, or the MatteControl node won't work.

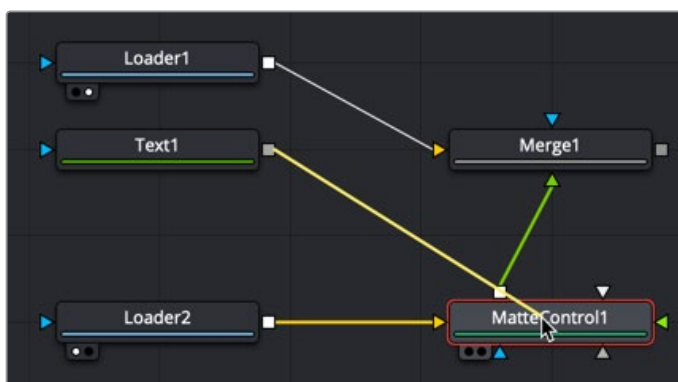


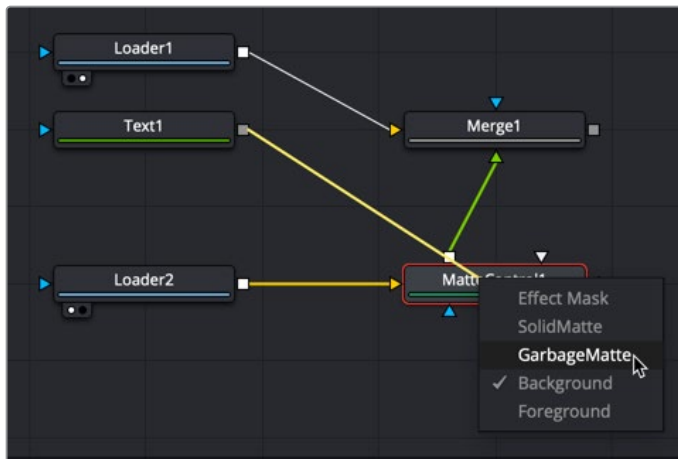
The second image properly connected to the Matte Control node's background input.

With this done, we'll connect the Text1 node's output, which has the alpha channel we want to use, to the MatteControl node's garbage matte input, which 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. Since 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.


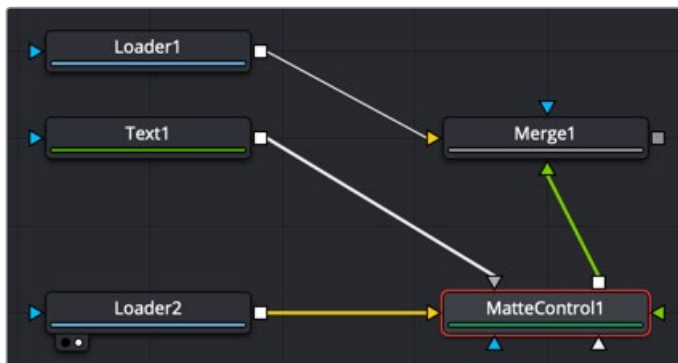
One alternate method of connecting nodes together is to hold down the Option key while dragging a connection from one node's output and dropping it onto the body of another node. This opens a pop-up menu from which you can choose the specific input you want to connect to, by name. Note that the menu only appears after you've dropped the connection on the node and released your pointing device's button.





Before (top) and after (bottom) Option-dragging a node connection to drop onto another node exposes a node input menu.

Once the Text1 node is properly connected to the MatteControl node's Garbage Matte input, you should see a text-shaped area of transparency in the graphic if you load the MatteControl node into the viewer.




Connecting the Text node to the Matte Control node's garbage matte input (top), and the resulting hole punched in the image (bottom).

Customizing Matte Control Nodes

With this accomplished, we need to use the Inspector to change some parameters to get the result we want. In the Inspector controls for the Matte Control node, click the disclosure control for the Garbage Matte controls to expose their parameters. Because we actually have a

Using Transform Controls in the Merge Node

Fortunately, there's an easy fix that doesn't even require us to add another node. Selecting the Merge1 node, we can see 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.



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.

Dragging the Size slider to the left shrinks the text to create the effect we really want, and at this point we've got the composite we need.



The final composite.

Match Moving Text With Motion Tracking

This next example introduces motion tracking, and 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 We Want to Match Move

In this example, 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 on-screen 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 off-screen can still be seen as an outline showing us where it is.

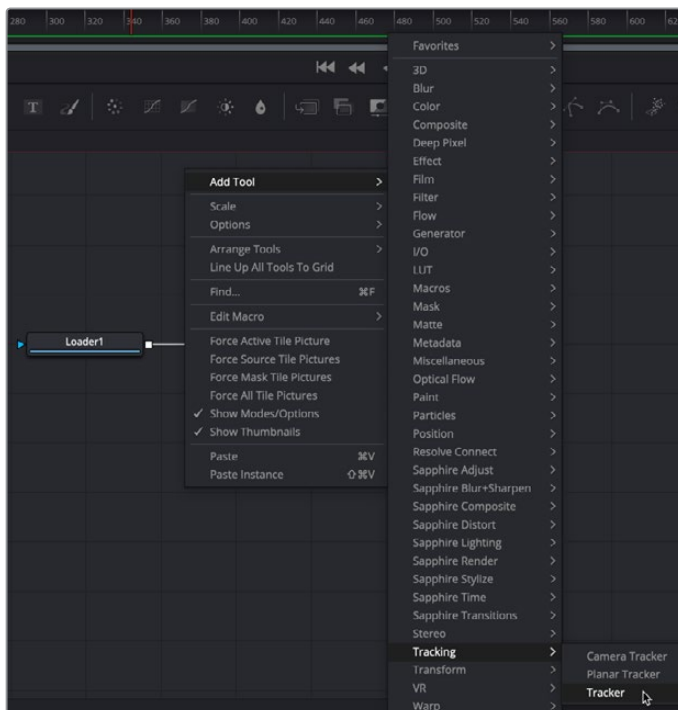


Some text superimposed against a background, ready to track.

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 to Track

To set up for doing the motion track, we'll begin by creating a disconnected Tracker node, using another method 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 Loader node



Creating a new node using the Node Editor contextual menu

Next, we'll drag a connection from the Loader1 node to the Tracker1 node to automatically connect the source clip to the Tracker1 background input. This branches the output from the Loader1 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 another 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 in Fusion, 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 On-Screen Control

When the Tracker node is selected, a single green box appears in the viewer, which is the default on-screen 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 on-screen controls for nodes that are selected, so if you don't see the on-screen tracker controls, you know you need to select the tracker you want to work with. Loading the tracker you want to work on 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 on-screen control for that tracker appears, and if you click the on-screen 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 that you want to follow the motion of. The pattern box has a tiny handle at its upper-left-hand 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 on-screen controls of a selected tracker seen in isolation.

It's worth saying a second time, the handle for moving a tracker's on-screen control is a tiny dot at the upper-left-hand corner of the inner pattern box. You must click on this dot to drag the tracker around.



The handle for dragging the tracker boxes to move them around.

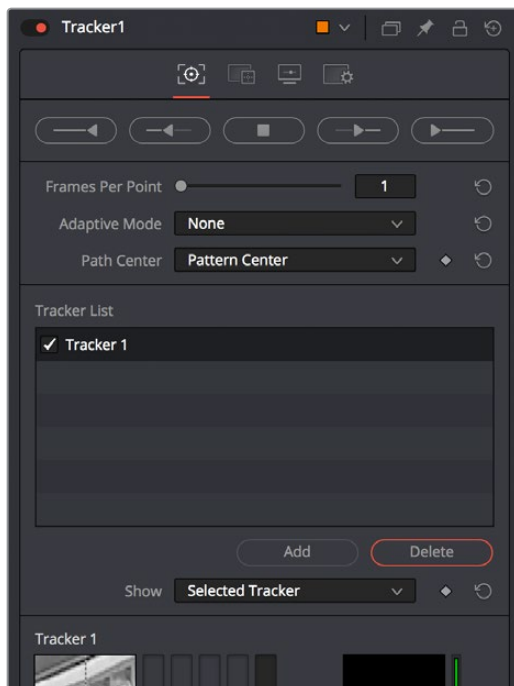
In this example, we'll drag the on-screen control so the pattern box overlaps a section of the bridge right over the leftmost support. As we drag the on-screen control, we see a zoomed-in representation of the part of the image we're dragging over, to help us position the tracker with greater precision. 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 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 pop-up 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, all we need to do is to use the tracker analysis buttons at 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 on-screen controls appropriately at the last or first frame.



The analysis buttons, left to right: Track from Last Frame, Track Backward, Stop Tracking, Track Forward, and Track from First Frame.

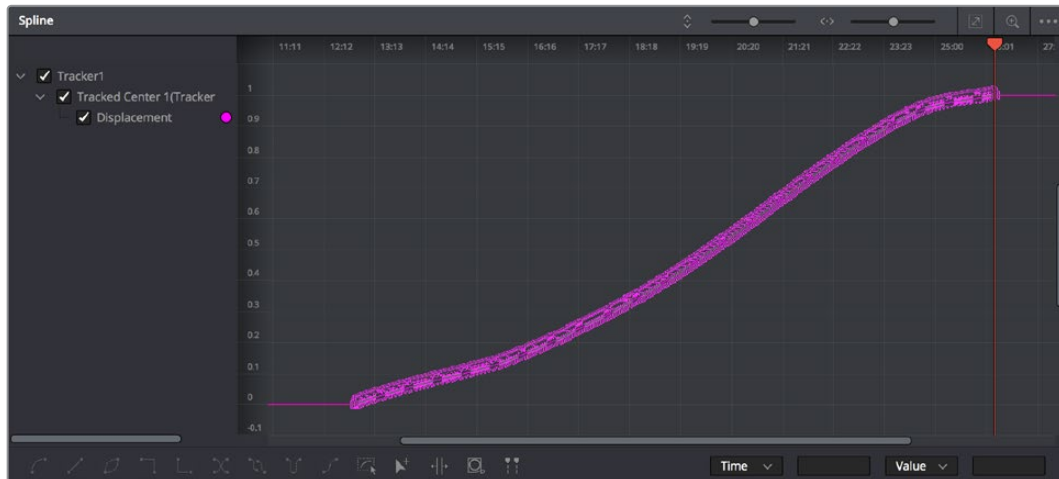
For now, clicking the Track from Beginning button will analyze the entire range of this clip, from the first frame to the last. A dialog lets you know when the analysis is completed, and clicking the OK button dismisses it so you can see the nice clean motion path that results.



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.

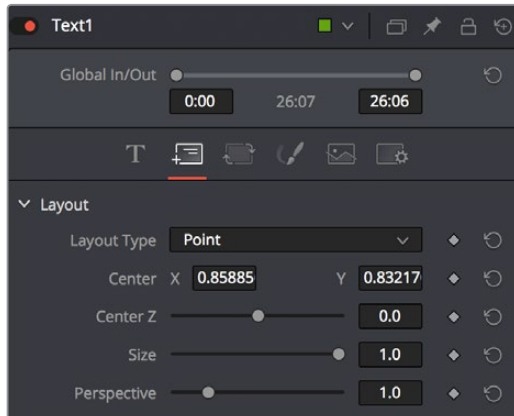
Connecting Motion Track Data to Match Move

Now that we've got 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 makes 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.

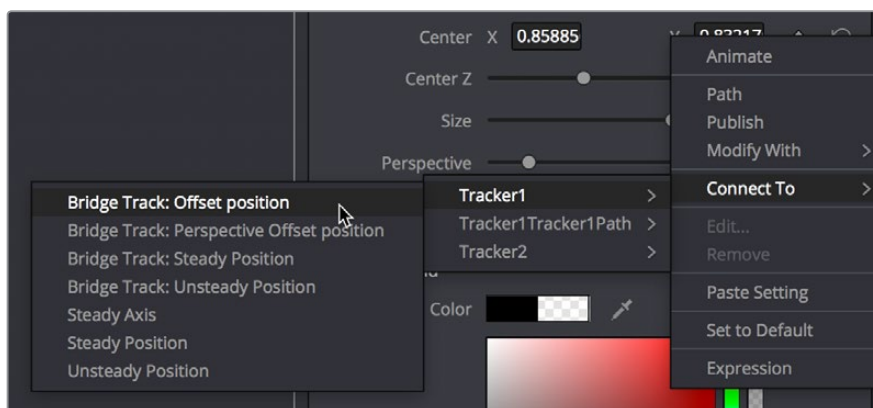
Now it's time to connect the track we've just made to the text in order to start it in motion. Loading the Merge1 node into the viewer to see the text in context with the overall composite we're creating, we'll select the Text1 node to open its parameters in the Inspector, and click the Layout panel 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 on-screen 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.

Using Paint and Planar Tracking

In this next example, we'll take a look at a paint example in which we eliminate some facial scars on an actor's forehead in a commercial. This workflow combines the Paint node with the Planar Tracking node, illustrating a common way of using these two powerful tools.



The actor has some scars on his forehead that the director would like painted out.

Using a Planar Tracker to Steady a Subject to Paint

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.

The best way to do this in Fusion is to use the Planar Tracker, so we'll add the PlanarTracker node after the Loader1 node, such that the image we want to track is connected to the background input of the PlanarTracker node. As always, it's important to be careful about which input you connect the image to for the effect to work properly.



Adding a PlanarTracker node to analyze and steady the part of the image we want to paint on

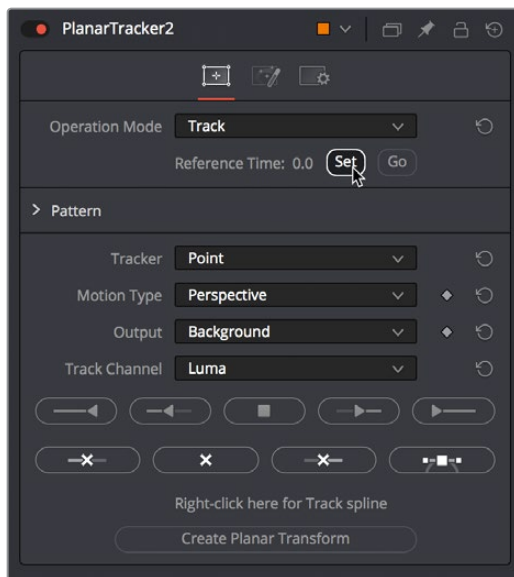
With the PlanarTracker node selected, and either it or the Loader1 node loaded in the viewer, a viewer toolbar appears with a variety of tools for drawing shapes and manipulating tracking data. The Planar Tracker works by tracking “planar” (read: flat) surfaces that you define by drawing a shape over the feature you want to track. When you first create a PlanarTracker node, you’re immediately put into a mode for 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 on each corner of the man’s forehead to create control points, clicking the first one we created once again to close the shape.



Drawing a shape over the man's forehead to prepare for Planar Tracking

Turning our attention to the Inspector, we can see that the PlanarTracker node has tracking transport controls similar to those of the Tracker, but with one difference. There are two buttons, Set and Go, underneath the Operation Mode pop-up, 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 should 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 toward 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 pop-up menu is perhaps the most important. In this particular case, Perspective tracking is exactly the analysis we want, but 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, and a green progress bar at the bottom of the Timeline Ruler lets you see how much of the shot is remaining to track.



After 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 you click one of the Track buttons to begin tracking and nothing happens, or if you track for a few frames and then stop, 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 be 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.


Painting Over Blemishes

Adding a Paint node after the PlanarTracker node gets us ready to paint.



Adding a Paint node after the PlanarTracker to paint onto the steady surface

With the Paint node selected and the Loader1 node loaded in the viewer, we can see the paint tools in the viewer toolbar. The first thing we want to do is to click on the fourth tool from the left, the Stroke tool, which is the preset tool for drawing strokes that last for the duration of the clip. The default Multi-Stroke tool is intended for frame by frame work such as painting out falling raindrops, moving dust and dirt, or other things of limited duration. The Stroke tool is much more appropriate when you want to paint out features or paint in fixes to subjects within the frame that need to remain in place for the whole shot.



Choosing the Stroke tool from the Paint node's tools in the Viewer toolbar.

Next, we need to go to the Inspector controls for the Paint node and choose the Clone mode from the Apply Controls. We're going to clone part of the man's face over the scars to get rid of them, and choosing the Clone mode switches the controls of the Paint node to those used for cloning.



Choosing the Clone mode in the Inspector.

There are additional controls located in this palette, however, that you should be familiar with.

- Brush Controls (at the top) contain the Brush Shape, Size, and Softness controls, as well as settings for how to map these parameters for tablet users.
- Apply Controls (in the middle) let you choose a paint mode, which includes Color, Clone, Emboss, Erase, Merge, Smear, Stamp, and Wire Removal. In this example, we'll use Clone. The mode you choose updates what controls are available below.

- Stroke Controls (at the bottom) are intended to let you adjust strokes after they've been painted, and include controls for animating them with "write-on" effects, transforming strokes with standard sizing parameters, and adjusting brush spacing.

With the Stroke tool selected in the viewer toolbar, and Clone mode selected in the Inspector controls, 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 you want to clone from. In this example, we'll sample from just below the first scar we want to paint. After Option-clicking to sample part of the image, clicking to begin painting sets an offset between where we're sampling from and where we're painting to, and dragging to draw paints a clone stroke.



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 (right).

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.

Before moving on, we'll open the Modifiers panel of the Inspector, where we can see that every single paint stroke we've made appears as an item on the Modifiers list. This gives us access to the strokes we've painted for further modification. We don't need to do anything at the moment, but when the time comes when you want to start making changes to strokes you've made, this is where they appear.




Each stroke made appears as an entry with controls in the Modifiers panel of the Inspector.

Keep in mind that the last stroke on the Modifiers list isn't really a stroke, it's a placeholder for the next stroke you're about to make, which might explain the numbering of the strokes if you're new to Fusion.

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 the Planar Tracker applied to put the clip back to the way it was, only with the painted fix attached in the process. This ends up being a two part process, but the first part is the simplest.



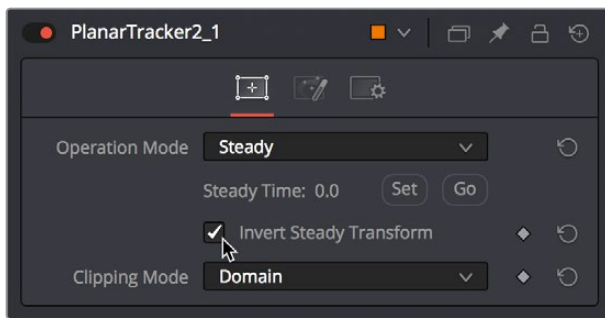
Scrubbing through the steadied clip shows the paint fix is "sticking" to the man's forehead.

Selecting and copying the PlanarTracker node coming before the Paint node, we select the Paint node and paste a copy of it after. This copy has all the analysis and tracking data of the



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 in theory inverts the steady warp transform to put the image back to the way it was. However, in practice, the more the image needs to be warped to steady it, the more likely that inverting the warp will introduce other problems.




Turning on Invert Steady Transform to try to put the image back to the way it was.

While the initial result appears to have just another warp applied to it, this time in reverse, the truth is that the region of the image centered on the shape used to do the planar analysis, the forehead, has gone back to the way it was before being steadied. It's just the edges of the frame that are distorted.

Using the Viewer's Split Wipe Control


This is a good example of a situation that can be tested using the Split Wipe control in the viewer title bar.



Opening the Split Wipe pop-up menu in the viewer.

Using the Split Wipe pop-up, switch to B View (the current image is A View), then drag the second PlanarTracker node into the viewer to load it into the B buffer, and then switch back to A View and drag the Loader1 node into the viewer to load it into the A buffer.

Turning on the Split Wipe button displays a split screen of the original image (A) against the transformed image (B). You can drag the handle of the green split control to adjust the split, and you can drag the line to change the angle of the split (holding Shift lets you snap the angle to 45° angles).



Comparing the "Invert Steady" version of the image with the original image to see the forehead is the same in both frames.

So, the forehead is fine, but the rest of the image is now warping in an unusable way because of the extremity of the warp needed to steady the region we wanted to paint. That's fine, because there's an easy fix that's a necessary part of this technique.


Fixing the Edges by Using Only the Fixed Part of the Frame

At this point, we're ready for the second part of this fix, which is to mask and composite just the fixed forehead against the original clip.

Isolating the Painted Forehead

First, we need to mask out just the man's painted forehead. We can do this by connecting a Polygon node to the Garbage Matte input of a MatteControl node, and then connecting the second PlanarTracker node's output (with the fixed forehead) to the MatteControl node's background 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.

The placement of these two new nodes can be seen in the following screenshot. We can wire this up before drawing the shape. In fact, it's essential because otherwise you want to trace the image being fed to the MatteControl node using the Polygon node.



Adding a Polygon node, a MatteControl node, and a Merge node to composite the painted forehead on the original clip

TIP: When it comes to using Masks to create transparency, there are a variety of ways to do this. For example, (a) attaching the image you want to mask to the background input of a Brightness/Contrast node with Alpha enabled to darken a hole in the alpha channel by lowering the Gain slider while the Polygon node is attached to the effect mask input, or (b) using ChannelBooleans to copy channel data to alpha from a Polygon node attached to the foreground input while the image you want to mask is attached to the background layer. However, the MatteControl node is flexible enough and useful enough to merit learning about it now.


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 fixed forehead. Loading the MatteControl1 or Saver node 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 almost every Bezier-style shape drawing in other graphics applications:

- Clicking once draws a corner control point.
- Clicking and dragging creates a Bezier curve.
- Click the first control point you created to close 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 GarbageMatte 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 (top), and the resulting inverted mask inverting the forehead (bottom).

Compositing the Painted Forehead Against the Original Image

Almost finished. We'll add one more node, a Merge node, that we'll use to actually layer the fixed forehead against the original image being output by the Loader node.

Creating a disconnected Merge node, we reconnect the MatteControl's output to the green foreground input of the Merge node, and then pull out a second branch from the Loader1 node's output to connect to the Merge node's orange background input. 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

So now we've got the best of both worlds, with 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 last issue to deal with. Happily, 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 Loader node, and temporarily choosing Track from the Operation mode pop-up menu, we can see there's a Create Planar Transform button at the bottom of the listed controls. Clicking this button creates a new, disconnected node in the Node Editor that uses the planar track as a transform operation for doing easy match moving. We click the Create Planar Transform button, and then set Operation Mode back to Steady.



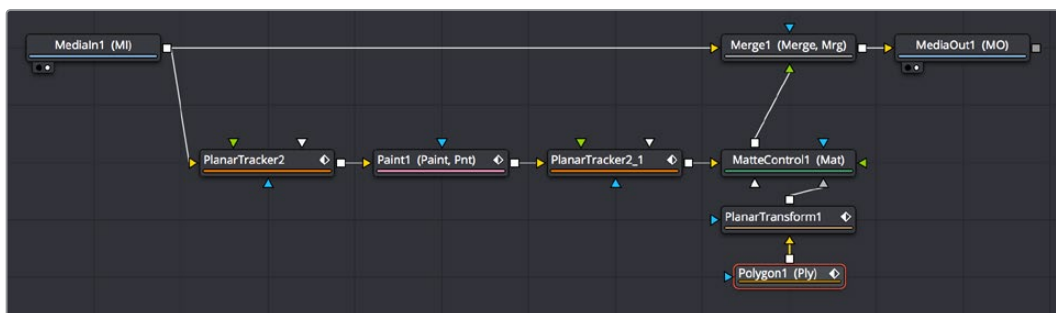
Creating a PlanarTransform node you can use to Match Move other images.

We can insert this new node into the node tree to use it by holding down the Shift key and dragging it 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 mode (right)

With the new PlanarTransform node inserted, the Polygon is automatically transformed 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.

NOTE: While on-screen controls are only visible when you select the node they belong to, on-screen controls only appear transformed properly when you load a node into the viewer that's downstream of operations that will transform the image.

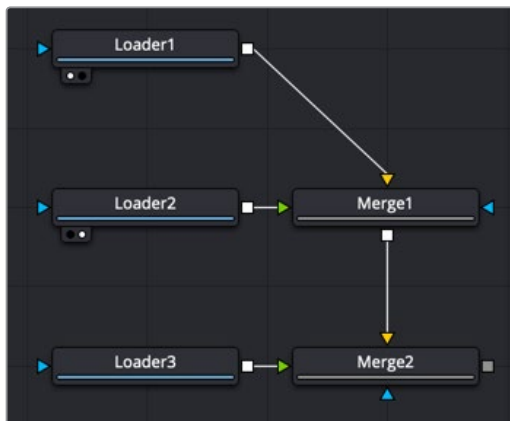
Building a Simple Green Screen Composite

In this next example, we'll take a look at how to build a simple composite using a greenscreen key and two other layers to create a news story. This effect involves a greenscreen clip, a background graphic, and a foreground graphic. All three clips are ingested as individual loaders. Loader1 is the background for our composite with Loader2 being the greenscreen foreground into Merge1. The output of Merge1 is connected as the background of Merge2, and Loader3 is a graphic that is the foreground for Merge2. This will create a three-layer effect. If you are more familiar with timeline-based compositing, Loader1 is the equivalent of video track 1, Loader2 is video track 2, and Loader3 is video track 3.



A background is Loader1 (left), a greenscreen is Loader2 (center), and a graphic is Loader3 (right).

This cascade of Loaders and Merge nodes (one Merge node for each pair of clips) takes care of combining each layer of video the way they would be on a Timeline.

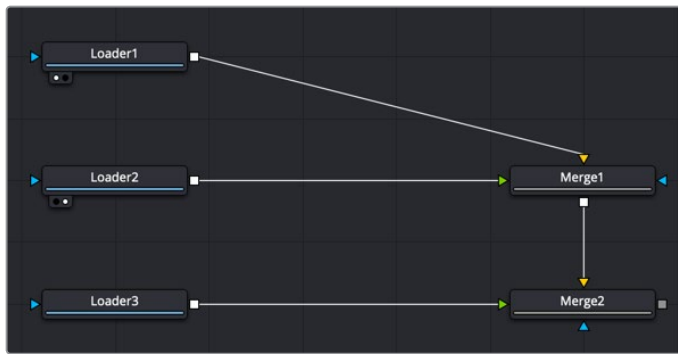


The initial node tree of three clips combined with multiple Merge nodes.

With this node tree assembled, we can now focus on adding the nodes we'll need to each branch of this tree.

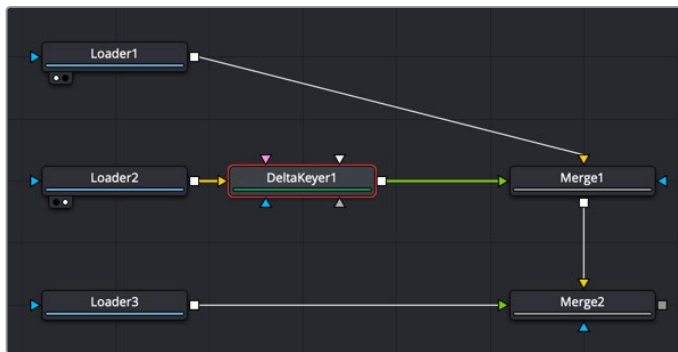
Pulling a Greenscreen Key Using the DeltaKeyer

First, we'll pull the greenscreen key we'll need to create transparency behind the newscaster. To prepare, we'll pull the Merge nodes off to the right to make room for the additional nodes we'll add after the Loader nodes as we work.



Creating space after the Loader nodes, and selecting the second one in preparation for adding a node.

Selecting the Loader2 node and loading the Merge1 node into the viewer lets us see the greenscreen clip and makes it easy for us to add a DeltaKeyer node inline by pressing Shift-Space to open the Select Tool dialog with which to search for and insert any node.

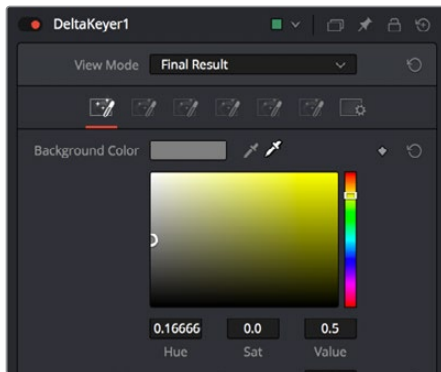


Adding a DeltaKeyer node inline after the Loader2 node.

The DeltaKeyer node is the main greenscreen/bluescreen keyer in Fusion 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 greenscreen, it's not limited to only keying green or blue; 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, using an Eye Dropper to sample the shade of green from the background of the image.

In the inspectorInspector, drag the pointereye dropperEye Dropper over the green of the background in the vViewer.




Dragging the Eye Dropper to the viewer to sample the Background Color.

As we drag in the viewer, an analysis of the color picked up by the location of the Eye Dropper appears within a floating tooltip, giving us some guidance as to which color we're really picking. Meanwhile, we get an immediate preview of the transparency we'll get at that pixel, and since we're viewing the Merge1 node, this reveals the image we've connected to the background.




The original image before (left) and after sampling the greenscreen using the Eye Dropper 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.



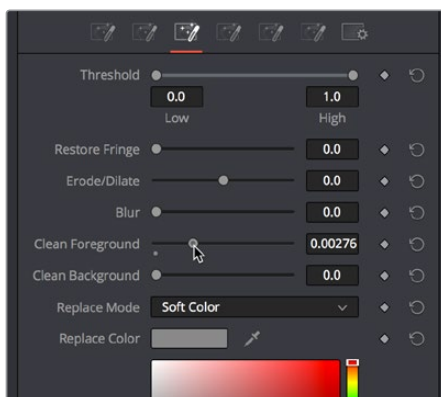
The DeltaKeyer Inspector updates with the sampled color.

Now that we've selected a background color to pull a key with, we can load the DeltaKeyer node into the viewer itself, and click the Color button in the viewer toolbar, or select the viewer and press A to switch the viewer between the RGB color channels of the image and the alpha channel to evaluate the quality of the key.



Loading the DeltaKeyer into the viewer and clicking the Color button to view the alpha channel being produced.

A close examination of the alpha channel reveals some fringing in the white foreground of the mask. Happily, the DeltaKeyer has integrated controls for doing post-processing of the key being pulled, found in the third of the seven panels of controls available in the DeltaKeyer. Clicking the Matte panel opens up a variety of controls for manipulating the matte, and since the fringing we don't like is on the foreground (white) part of the key, we'll use the Clean Foreground slider to make the fix.



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 compromising the edges of the key.



The original key before (left), and the key after using the Clean Foreground slider (right).

With this accomplished, we're good with the key, so we load the Merge1 node back into the viewer, and press A to toggle 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.

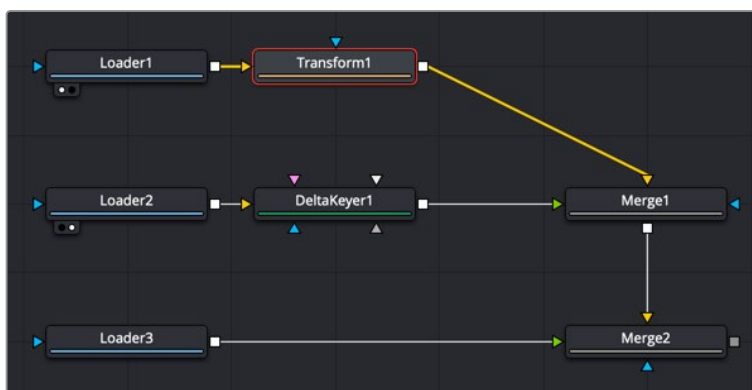


The final key is good, but now we need to work on the background.

Using the Transform Node to Resize a Background

Since the background isn't covering up the whole frame, we need to transform it. It's a high-resolution image, so that's not a problem, however it's connected to the background input of the Merge1 node, and although Merge nodes have built-in transform controls, they only work on the foreground input (on the premise that the foreground will need to be fit to the background).

This means that we need to add a Transform node to the Loader1 node to take care of this. Selecting the Loader1 node and clicking the Transform button in the toolbar takes care of this, and we're ready to work.



Adding a Transform node to change the sizing of the Loader1 image connected to the background.

While there are slider controls in the Inspector for Center, Size, and Angle (among other parameters), there are on-screen controls that give more satisfyingly direct control. Zooming out of the viewer a bit by holding the Command key and using the scroll control of your pointer, we drag the side border of the graphic to proportionally enlarge the blue background until it fills the screen (there's still a black border at the top and bottom of the clip, but that's burned into the news clip we have).



Enlarging the background to fill the frame using the viewer's on-screen controls.

At this point, we decide to make room for the graphic we know we'll be putting into the frame at left, so we take advantage of the built-in transform controls in the Merge1 node that affect the foreground input. Selecting the Merge1 node, we drag the left arrow of the on-screen controls that appear to move the man to the right, and we take advantage of knowing the image of the man is high-resolution relative to our project resolution by dragging the side edge to proportionally enlarge the foreground image to crop out the black bars.



Using the Merge1 node's on-screen transform controls to reposition and enlarge the image to prepare for adding another element.

NOTE: You may have noticed that there are both Transform and Resize buttons in the toolbar. It's important to be aware that while the Transform node always refers to the original source resolution of the image for resolution-independent sizing in which multiple Transform nodes can scale the image down and up repeatedly with no loss of image resolution, the Resize node actually decreases image resolution when you shrink an image, or increases image resolution (with filtering) when enlarging. In most situations, you want to use the Transform node, unless you specifically want to alter and perhaps reduce image resolution to create a specific effect.

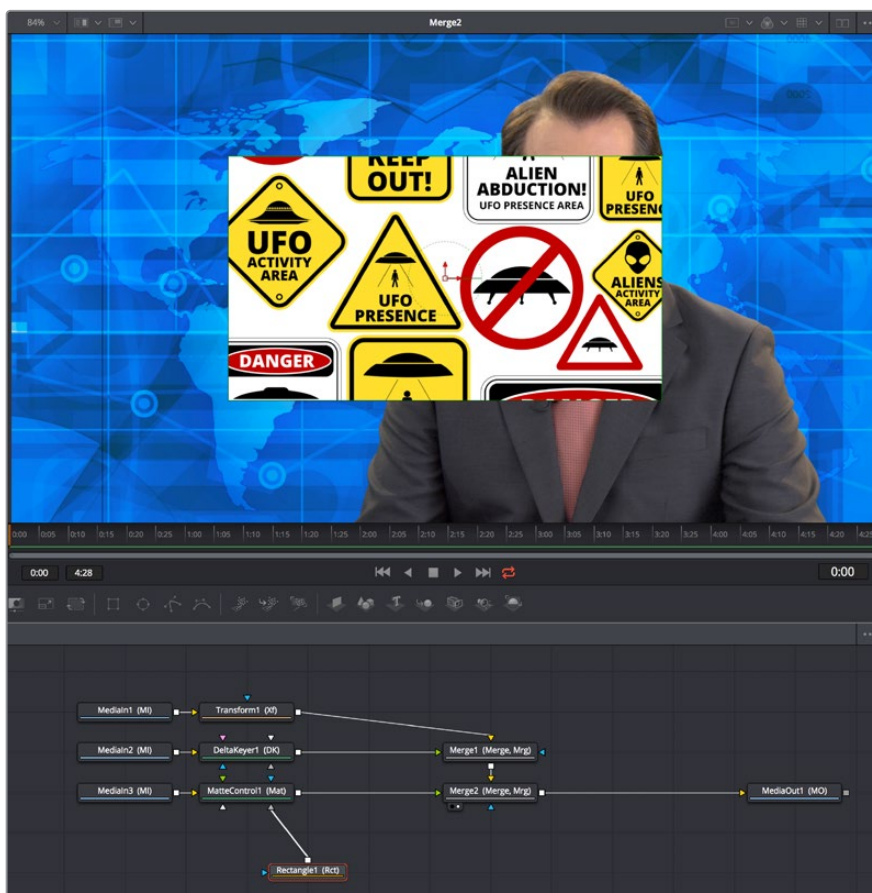
Masking a Graphic

Next, it's time to work on the news graphic that will appear to the left of the man. If we load the Merge2 node, that combines the blue background and newscaster we just finished working on with the logo (Loader3). We can see that the logo layer is actually a sheet of different logos that appear on top, so we need to cut one out using a mask and fit it into place.



We need to mask out a single logo from this sheet to use in our composition.

Selecting the Loader3 node that's feeding the logo layer, we click the MatteControl button of the toolbar to add a MatteControl node, and then we add a Rectangle mask, manually connecting the Rectangle mask's output to the gray garbage mask input of the MatteControl node. Finally, we select the Rectangle node, and click its Invert checkbox to invert the Rectangle Mask's output, so it's cropping the logo layer correctly.



Masking the logo using a Rectangle mask connected to a MatteControl node.

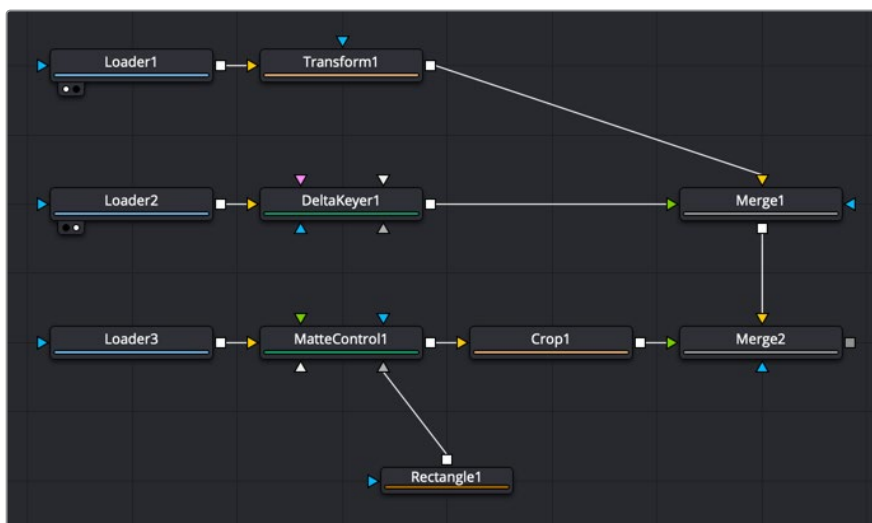
Now, all we need to do is to use the on-screen controls of the Rectangle mask to crop the logo we want to use, dragging the position of the mask using the center handle, and resizing it by dragging the top/bottom and left/right handles of the outer border.

As an extra bonus, we can take care of the fact that the logo has rounded borders by using the Corner Radius slider in the Inspector controls for the Rectangle matte 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.

Now that we've masked the logo, we'll crop the unused parts of this image so that the logo we're using is centered on the frame, which will make subsequent transform operations much easier. Selecting the MatteControl1 node, we add the Crop node from the Tools > Transform category of the Effects Library, and load the new node into the viewer.



Adding a Crop node after masking the image to center the cropped logo on the frame

With the Crop node selected, we can click the Crop tool in the viewer toolbar.



Selecting the Crop tool in the viewer toolbar.

This lets us crop the image by dragging a bounding box around it.



Dragging a bounding box using the Crop tool (left), and the cropped logo now centered on the frame (right).

NOTE: The Cropping node discards resolution, just like the Resize node does, so use it with care.

At this point, we're all set to move the logo into place, so we select the Merge2 node and load it into the viewer, and once again avail ourselves of the built-in transform controls for foreground inputs, using the on-screen controls to put the logo where we want it and make it a suitable size.



Placing the logo using the foreground input transform controls of the Merge2 node.


Animating an Image Using Keyframes

We're almost done with this grand tour of Fusion functionality, but we have one last task to accomplish. Now that we've positioned the logo appropriately, we need to animate it coming into frame to open the segment. To do this, we'll use the keyframe controls in the Inspector to begin keyframing, then we'll use the controls in the viewer to create a motion path, and finally we'll use the Spline Editor to refine the result.

Animating a Parameter in the Inspector

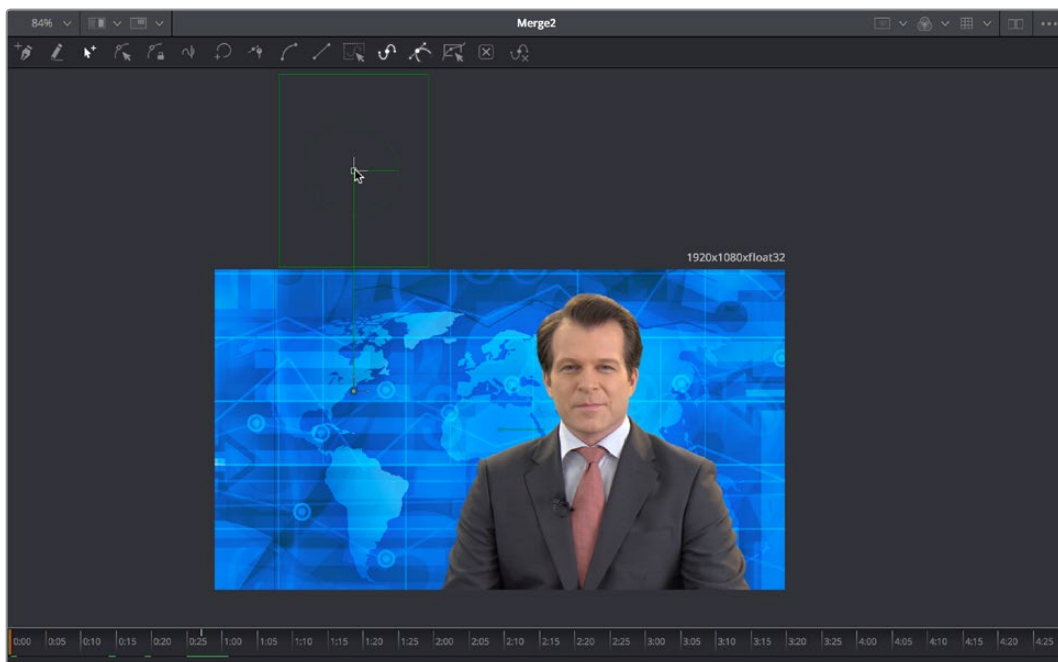
Before beginning to keyframe, it's always good to think your way through what you want to do before starting anything, just to make sure you're taking the right approach. In this case, we just want to slide the logo down from the top of the screen to where we've positioned it, so it's probably best to start adding keyframes at the end point of the animation we want to create by moving the playhead in the Time Ruler 24 frames forward from the beginning of the composition.

Selecting the Merge2 node, in which we used transform controls to position the logo, we click the small diamond control to the right of the Center parameter to create a keyframe for that parameter, in the process setting up that parameter so that every alteration we make on a different frame adds a keyframe.



Adding a keyframe to begin animating a parameter.

Next, we move the playhead back to the beginning of the composition, then zoom out of the viewer so there's more room around the frame before dragging the center handle of the logo up until we've dragged it off-screen. In the process, a second keyframe appears next to the Center parameter in the Inspector to show there's a keyframe at that frame, and a motion path appears in the viewer showing you the route the now animated logo will take.

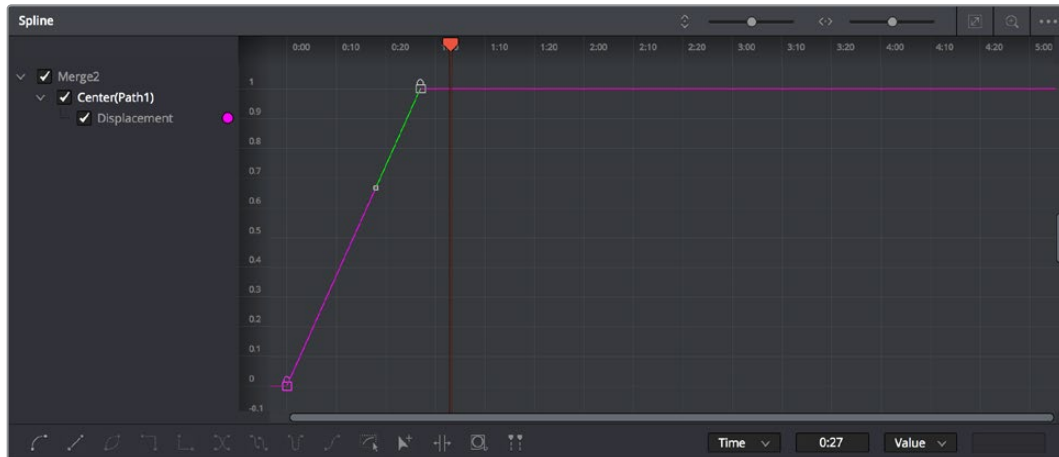


Moving an object in the viewer to create animation via a motion path.

At this point, if we play through the animation, it's functional but not exciting. The motion is linear so it comes into the frame and stops with a nearly audible “thunk.” Happily, we can fix this using the Spline Editor.


Using the Spline Editor

Clicking the Spline button in the UI toolbar opens the Spline Editor at the right of the Node Editor. The Spline Editor is a keyframe graph where you edit and finesse the curves created by animated parameters. By default, each animated parameter from every node in the current composition appears in the parameter list to the left of the curve graph. Turning on the Displacement checkbox shows our animated curve in the graph so we can work on it.



The Displacement curve from the animated Center parameter of the Merge2 node in the Spline Editor.

Drag a bounding box over the second of the two control points that are shown in the graph, so it's highlighted.




Selecting a control point to modify.

With that control point selected, click the Smooth button in the toolbar at the bottom of the Spline Editor to turn that keyframe into a Bezier curve (this also works for multiple selected keyframes). This has the effect of easing the motion to a stop at that second keyframe.



Clicking the Smooth button to turn the selected control point in the graph into a Bezier curve.

Playing through the animation, the logo does ease to a stop, but it's subtle. We up the ante by dragging the Bezier handle of the final keyframe to the left, making the curve steeper and resulting in the logo coasting to a stop more gradually.



Editing the spline to create a steeper curve, making the logo coast more gradually to a stop.

Congratulations

At this point, we're finished with our tour. As many things as we've covered, this is still only scratching the surface of what Fusion is capable of. However, this introduction should have given you a solid look at how to work in Fusion so that you can explore further on your own.

Have fun!

Chapter 15

Using the Tracker Node

This chapter shows the many capabilities of the Tracker node in Fusion, starting with how they can be connected in your node trees, and finishing with the different tasks that can be performed.

Contents

Introduction to Tracking	377
Overview of the Tracker Node	377
An Overview of the Tracker Node Workflow	377
Tracker Inspector Controls	381
Motion Tracking Workflow In Depth	382
First: Connect the Image You Need to Track to a Tracker Node	382
Second: Add Trackers to the Tracker List in the Inspector	382
Third: Position the Pattern Boxes of Each Tracker	383
Fourth: Define the Search Area of Each Tracker	384
Fifth: Perform the Track Analysis	385
Sixth: Use the Track	386
Using The Pattern Flipbooks	387
Using Adaptive Pattern Tracking	387
Dealing with Obscured Patterns	388
Dealing with Patterns That Leave the Frame	389
Setting Up Tracker Offsets	389
The Many Tasks of Match Move	390
Stabilizing Motion (Using the Match Move Mode)	391
Smoothing Motion	392
Using the Tracker Node for Match Moving	392
Simple Match Moving	392
Corner Positioning Operations	393
Perspective Positioning Operations	393
Connecting Trackers to Other Operations	394
Steady Position	394
Steady Angle	394
Offset Position	394
Unsteady Position	395
Steady Size	395
Using Tracker Modifiers	396

Introduction to Tracking

Tracking is one of the most powerful automation nodes available to a compositor, and Fusion has a variety of different tracking nodes available to let you analyze different kinds of motion, and use the resulting data to match the motion of one image to that of another, stabilization, motion smoothing, and a host of other essential tasks.

Overview of the Tracker Node

The tracker node has four operation modes for different situations.

Stabilizing

You can use one or more trackers 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 only stabilize 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 an image sequence using one or more patterns. Instead of removing that motion, it is applied to another image that matches the movement 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.

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.


An Overview of the Tracker Node Workflow

All tracking workflows consist of three fundamental steps.

First, Attach a Tracker Node's Background Input to the Image You Want to Analyze

You attach a Tracker node's background input to the output of the node with the image you want to track, as seen in the following screenshot. The Tracker node only analyzes the state of the image that's attached to its background input.

If you intend to use the Tracker node itself to do a transform operation using the tracking data it analyzes, you can connect it serially. However, if you're just using a Tracker node to analyze data for use elsewhere, 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 set to have no effect on the image as well.



Tracker node connected serially so it can both track and transform the input image (top), and Tracker connected as a branch to indicate it only contains tracking data, without using it directly (bottom).

Second, Analyze the Image to be Tracked

One or more features in the image that you wish to track (referred to as patterns) are identified by creating trackers (there's one by default). Tracker on-screen controls appear in the viewer that you can position over the patterns you need to track. After the Tracker node analyzes the shot, 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 that is now available.



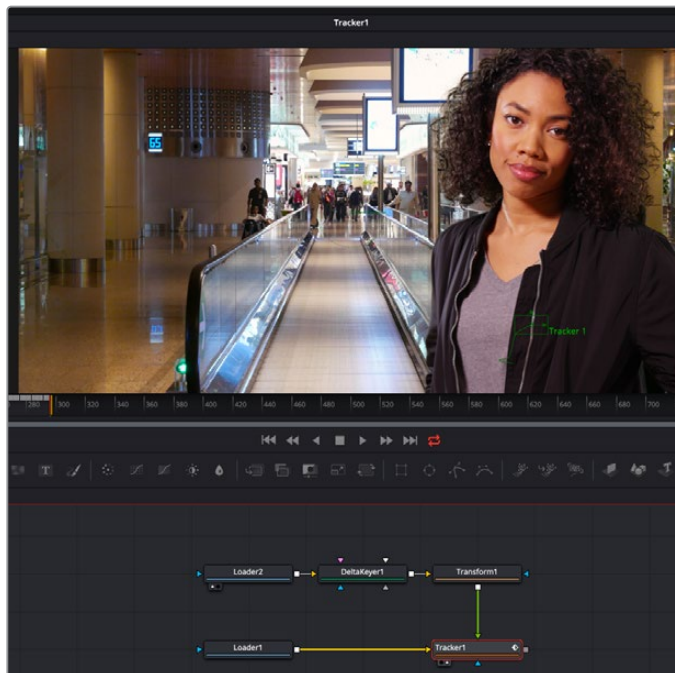
A motion path that indicates the tracked motion path, and tic marks that indicate tracking data keyframes.

Third, Apply the Tracking Data

The resulting tracking data stored within the Tracker node is used to stabilize, match move, cornerpin, or perspective position that or other images, in one of two ways.

Use the Tracker node itself to match move and merge

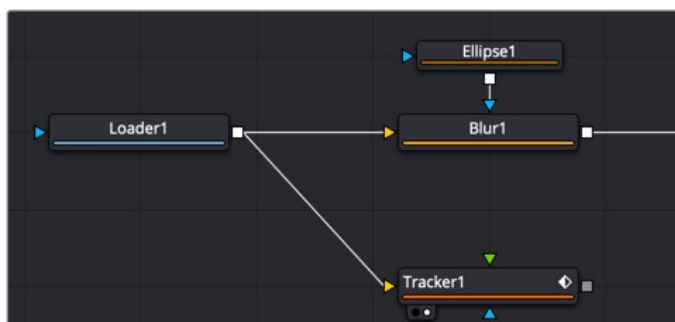
You can connect the image you want to match the motion of the analyzed background image to the foreground input of the Tracker node, and then set the Operation parameter in the Operation tab in the Inspector to Match Move, Corner Position, or Perspective Position, as necessary. This is an easy workflow for simple situations. In this case, the Tracker node then also does what the Merge node does as well.



Using a Tracker node to do a match move and merge, all in one


Connect the Tracker you analyzed to specific parameters:

Alternatively, you can apply 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. For example, in the following node tree, an Ellipse node is being used to limit a Blur effect to a woman's face for a documentary. However, the ellipse needs to follow the motion of the woman's face, so a Tracker node is used to analyze the movement of the woman's nose so that track data can be used to animate the ellipse by connecting the track data to a specific parameter.



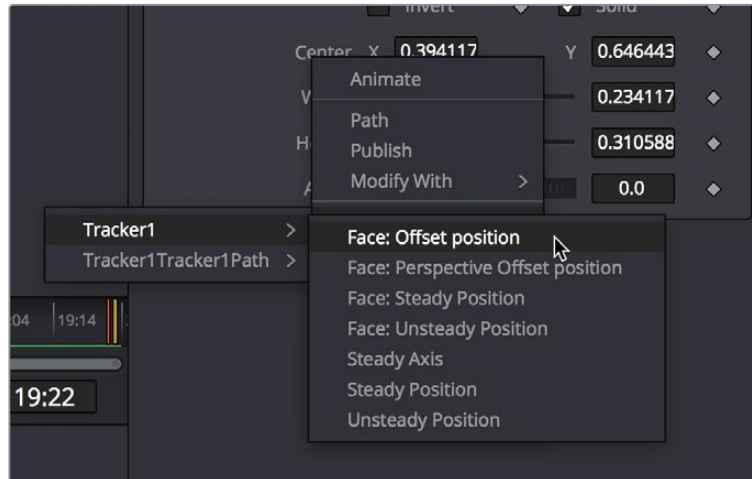
Blurring an actor's face by connecting tracking data to the center position of an Ellipse node.

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 something descriptive text.

Once the nose on the woman's face has been tracked, this tracking data is then connected to the Center parameter of an Ellipse node that's limiting a Blur node to the woman's face by right-clicking the label of the Center parameter in the Inspector, and choosing Tracker1 > Face: Offset position from the Connect to submenu of the contextual menu. All the data from every Tracker node in your node tree appears within this submenu, and since we named the Tracker we want, it's easy to find. We choose Offset position because that's the value that gives us the most control.



Right-clicking a parameter's label lets you connect tracking data to it to animate it.

You can connect any Tracker node's data to any other node's parameter; however, you'll most typically connect track data to center, axis, or corner X/Y style parameters. When you use tracking data this way, it's not necessary to connect the Tracker node itself to anything else in your node tree; the data is passed from the Tracker to the Center parameter as an expression.

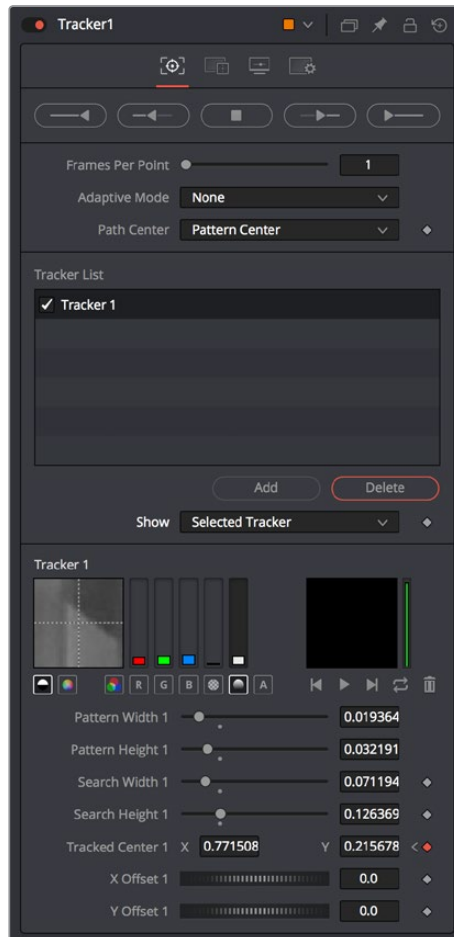


Connecting the center of the Ellipse node to the Face tracker animates the ellipse.

Tracker Inspector Controls


The layout of the Tracker node's tabs in the Inspector reflects this workflow. It's divided into three main tabs:

- **The Tracker tab:** Where you create on-screen 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:** Where you decide how the tracking data is used.



The Tracker Operations tab.

- **The Display Options tab:** Where you can customize how the onscreen controls look in the viewer.



The Tracker Display Options tab.

Motion Tracking Workflow In Depth

Analyzing motion using one or more trackers within the Tracker node is easy.

First: Connect the Image You Need to Track to a Tracker Node

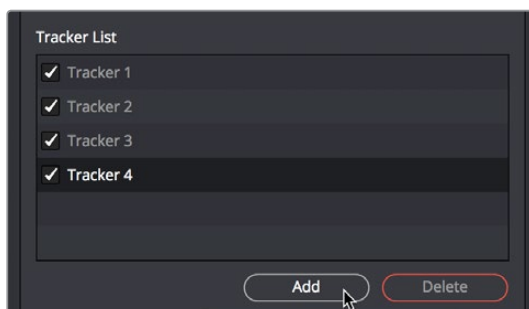
Tracker nodes serve two purposes. They provide the nodes you need to analyze an object you want to follow, and they serve as a container for the resulting track data. Regardless of whether you actually use the Tracker node itself to do anything with the tracking data, the image you want to track must 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's ignored for purposes of tracking analysis, so you must connect the background for tracking to work successfully.



Connecting a MediaIn node's output to a Tracker node's background input.

Second: Add Trackers to the Tracker List in the Inspector

Although each Tracker node starts with a single tracker to get you started, a single node is capable of analyzing multiple trackers 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 and corner or perspective positioning. They also help to keep the Node Editor from becoming cluttered by collecting into a single node what would otherwise require several nodes.


Methods of managing 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.

Third: Position the Pattern Boxes of Each Tracker

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 trackers, each targeting different patterns. Each tracked pattern will produce its own motion path.

Clicking any part of a tracker's on-screen controls will select it. Selected pattern boxes are red, while de-selected 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 it 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 eye.

Fourth: Define the Search Area of Each Tracker

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.

Fifth: 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 Track Backward from Current Frame to track backward from the current playhead position.
- Click the Track Forward button to track from the very start of the render range.
- Click Track Forward from Current Frame 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.

Sixth: Use the Track

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.



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 pattern that is selected, 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 buttons 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, 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 nose is being offset to the actor's eye.

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.

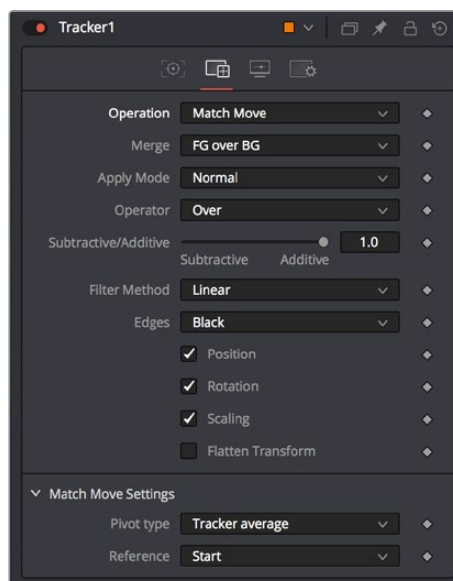
The Many Tasks of Match Move

When a Tracker node is set to Match Move in the Operations tab, it is capable of a variety of functions:

- Stabilizing footage to completely remove motion from the scene or smooth existing motion.
- Applying the motion from one clip to another, basically matching the movement and stabilizing one shot with reference to another.

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 button.


Stabilizing Motion (Using the 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 Match Move button is selected in the Tracker's Operation tab, the Tracker can use the data from its patterns for stabilization. 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 buttons determine to what input connection on the Tracker node the stabilization data is applied. When stabilizing an image to remove all motion, or smooth the motion, the Merge button should be set to BG Only.

Edges

The Edges buttons determine 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. These controls are described in greater detail in the Node References for the Tracker later in this manual.

Position/Rotation/Scaling

Use the Position, Rotation, and Scaling checkboxes to select what aspects of the motion are corrected.

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 will show 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.

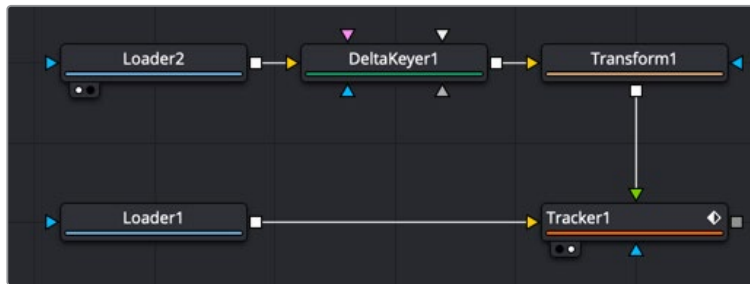
Simple Match Moving

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 buttons to BG Only, set it to FG Over BG, FG Only, or in rare occasions, BG Over FG.



Set the Merge buttons to BG Only, FG Over BG, or BG Over FG.

When using this Merge buttons, 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 buttons are 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.

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 buttons are set to BG Only.

Connecting Trackers to Other 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, each pattern in a tracker outputs five values for use as connections that are available for use by other nodes.

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 > Unsteady Position).

There are five connection types available 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 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

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.

The connections above are output by pattern in the Tracker node. Each node itself also outputs a Steady Position, Angle, and Size output. The values of these outputs are calculated using all the patterns in that tracker, as configured by the Stabilize Settings controls in the Tracker's Operation tab.

As an example, to use the Connect To menu to perform a match move, do the following:

- 1 Add a Transform node to the clip you want to match move.
- 2 Right-click over the Transform's Center and choose Connect to > Steady Position.
- 3 Set the Transform node's Edges mode to Mirror so that pixels will not get cropped from the image when it is stabilized.
- 4 Add another Transform node to the Node Editor after the merge.
- 5 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 Tracker Modifiers

Another technique for adding a tracker directly to a control is to add it 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.

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 a Glow node.
- 2 Create an ellipse mask in the shape of the eye.
- 3 In the Inspector, right-click on the mask's Center parameter.
- 4 From the contextual menu choose Modify With > Tracker Position.

Choosing the Tracker from the Modify with Contextual menu adds a modifier in the Inspector with a set of parameters almost identical to those found in the Tracker node itself.

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 will default to the output of the Loader). You can set a different source image for the Tracker modifier by typing in the name of the node at the top of the Inspector Modifier tab. Alternately, drag and drop the source node from the Node Editor into the Text Box control or use Connect To from the contextual menu.

For more detail on the Tracking controls please see Chapter 52, "Tracker Nodes."

Chapter 16

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 detail about the Planar Tracker node, see Chapter 52, “Tracker Nodes.”

Contents

Using the Planar Tracker	399
Different Ways of Using the Planar Tracker Node	399
Setting Up to Use the Planar Tracker	400
Check for Lens Distortion	400
A Basic Planar Tracker Match Move Workflow	400
Tips for Choosing Good Planes to Track	403

Using the Planar Tracker

The Planar Tracker node is designed to deal with match moving issues that commonly arise during post-production. Examples include moving footage containing planar surfaces such as a license plate, a road sign, or a brick wall that often needs 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 reapplying 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, it's not 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 or the current image in different ways.

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 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

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.

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, click the Go button to 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.



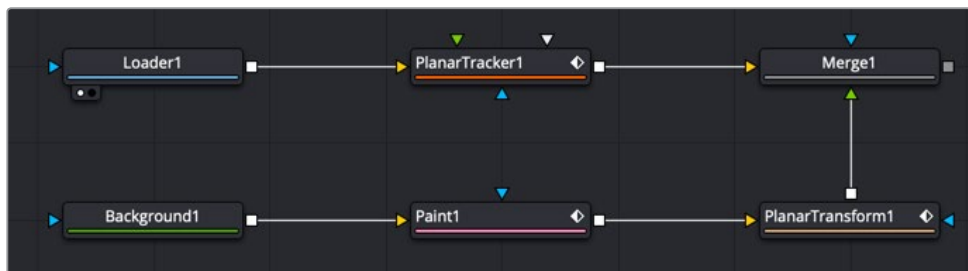
During tracking, you can see Track Markers and Trails to follow how well the track is going.

- 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. Where the planar tracker handles rectangular images for corner pinning, the planar transform can be used for irregular shapes like polygon masks.



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 end 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 transparent paint strokes on the wall, you can attach a Paint node to a Background node set to 100 transparency, as seen in the node tree above. The resulting image will be whatever paint strokes you make against transparency—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 big 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 17

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	406
Mask Nodes	406
Converting Polylines from One Type to Another	408
How to Use Masks with Other Nodes	408
Attaching Masks to an Image for Rotoscoping	409
Combining Multiple Masks	410
Mask Inputs on Other Nodes	411
Creating and Editing Polylines In-Depth	414
The Polyline Toolbar	414
Selecting a Specific Polyline	414
Polyline Creation Modes	414
Protection Modes	416
Closing Polylines	416
Selecting and Adjusting Polylines	417
Polyline Points Selection	417
Moving Polyline Points	417
Smoothing a Polyline Segment	417
Linearizing a Polyline Segment	418
Transforming Individual or Multiple Points	418
Deleting Selected Points	418
Editing Bezier Handles	419
Point Editor	419
Reduce Points	420
Shape Box	420
Showing and Hiding On-Screen Polyline Controls	421
Stop Rendering	422
Roto Assist	422
Creating Softness Using Double Polylines	422
Converting a Single Polyline to a Double Polyline	423
Adding Softness to a Segment	424
Adding Additional Points to the Shape	424
Locking/Unlocking Point Pairs	424
Animating Polyline Masks	424
Removing Animation from a Polyline Mask	425
Adding and Removing Points from an Animated Mask	425
Publishing Specific Control Points	425

Introduction to Masks and Polylines

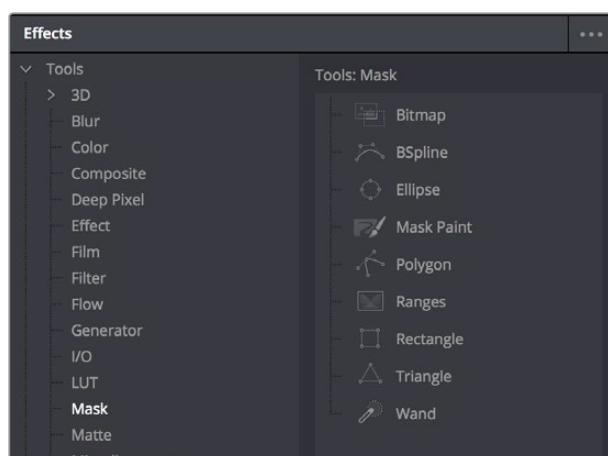
Polylines are splines that are used whenever a control is animated with a motion path or when a node's effect is masked with a drawn shape. They are also used in the Paint and Grid Warp nodes. In a more basic form, polylines are used to control the animation in the Spline Editor. Since these splines 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 sections on how to create paths or effect masks.

Mask Nodes

Masks are like other creator nodes in the Node Editor, with the exception that they create a single channel image rather than a full RGBA image. Most of these Mask nodes are also located under the viewers, with the exception of Ranges Mask and Triangle Mask.

For more details on the controls for each Mask node, see Chapter 43 “Node Reference.”



The available nodes in the Mask bin of the Effects Library.

Polygon Mask

Polygon Masks are user-created Bezier shapes. This is the most common type of polyline and the basic workhorse of rotoscoping.

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.

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. See Rectangle, Ellipse, and Triangle in the Node Reference.

Ranges Mask


Similar to 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 the Color Corrector node.

Polyline Types

You can draw polylines using B-Spline or Bezier spline types. Which you choose depends on the shape you want to make and your comfort with the spline style.


Bezier Polylines

Bezier Polylines are shapes composed of control points and handles. Several points together are used to form the overall shape of a polyline.



Bezier 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.




Bezier 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 Bezier polylines.

B-Spline Polylines

A B-Spline polyline is similar to a Bezier 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

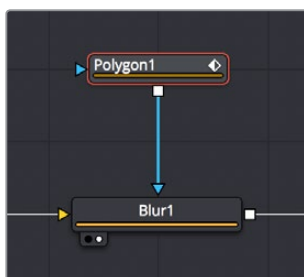
Just because a shape using a B-Spline or polyline has been created does not mean that the spline type has to be used. You can convert from B-Spline to Bezier, or Bezier to B-Spline, as needed. Once you create a polyline, you can right-click in the viewer and choose Convert Bezier Spline to B-Spline or Convert B-Spline to Bezier from the spline's contextual menu. Only the appropriate option will be displayed.

When converting from one type to another, the 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.

Although animation is preserved, the conversion process will not yield perfect results every time. It is 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 another 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, and 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 do 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 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 Loader node directly into the viewer while selecting the Mask node, you can draw a spline to rotoscope the image.



Rotoscoping a Loader 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 Loader node to use a spline as an alpha channel.

TIP: If you connect a Mask node to a 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 Loader node becomes completely blank. This is why when you want to rotoscope by connecting a mask to the input of a 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 Loader's input, and you're good to go.

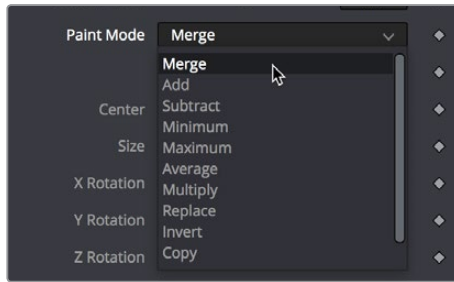
Combining Multiple Masks

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 pop-up appears, which allows you to choose how you want to combine the two masks.



The Paint Mode parameter 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

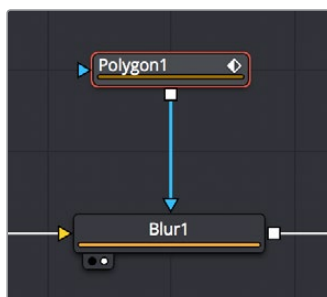
Mask Inputs on Other Nodes

Masks can be used for a variety of reasons, so there are several categories of mask inputs that different nodes have 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 a mask input and add a mask, it'll automatically be connected to that input.


Effects Mask Inputs

Almost every node in Fusion has an Effect 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.



A Polygon node used as a mask to limit the Blur node's effect.

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.

Pre-Masking Inputs

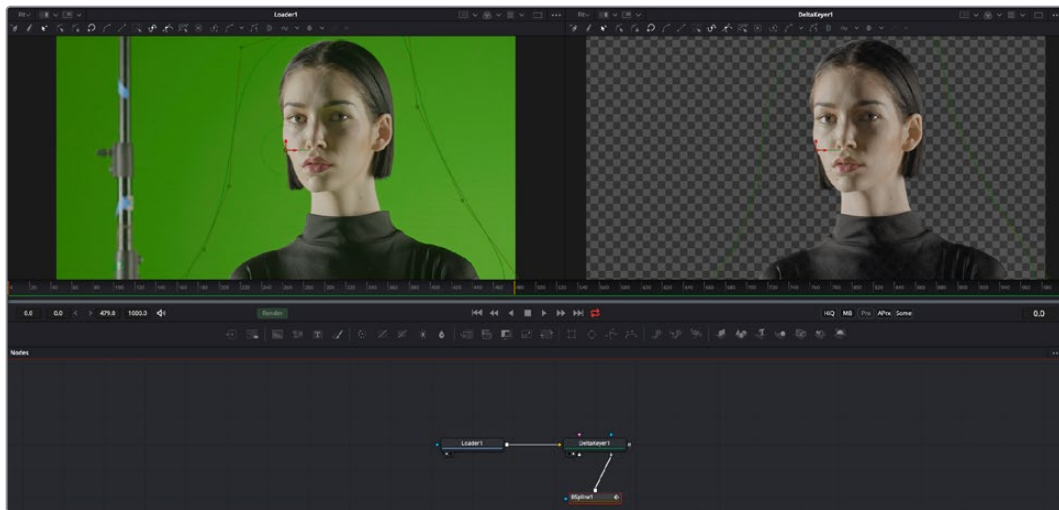
Unlike effect masks, a pre-mask input (the name of which is usually specific to each node using them) is used by the node before the effect is applied. This usually causes the node to render more quickly and to produce a more realistic result. In the case of the Highlight and the Glow nodes, a pre-mask 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 lighting instruments, rigging, and boom microphones that intrude upon masks being pulled via bluescreen and greenscreen 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.



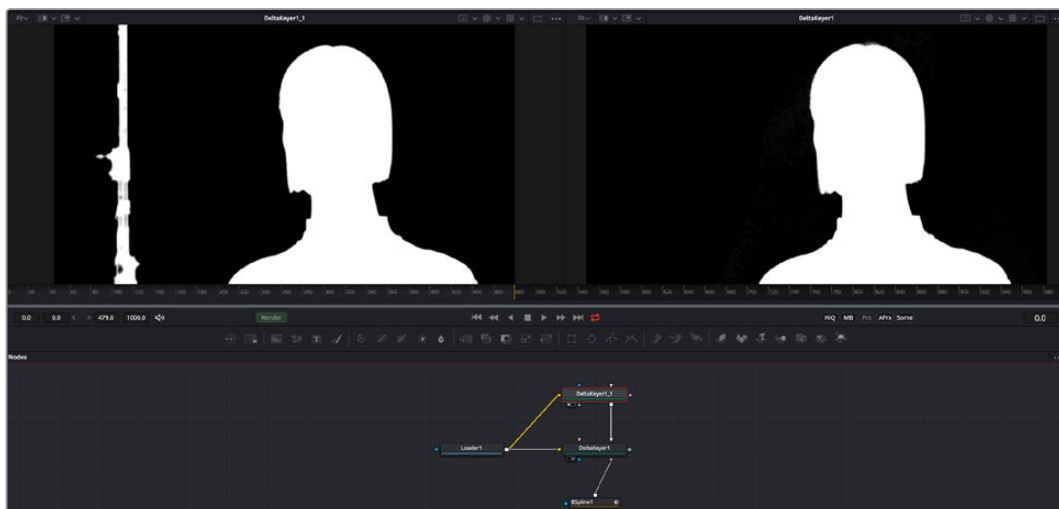
A B-Spline node is connected to the Garbage Matte input of a DeltaKeyer node to eliminate a light stand at the left of frame.

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 (at left) with another, harder but eroded key in DeltaKeyer2 that's connected to the SolidMatte input of DeltaKeyer1.

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.

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 cursor 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 be 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 cursor where you want to start the shape.
- 3 Move and click the cursor 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 cursor 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 know the timing yet.

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.

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.



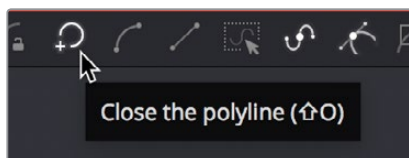
The Done toolbar button (Shift-N).

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 cursor 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.
- Choose Closed from the polyline's contextual menu.



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 cursor, 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 Left, Right, Up, or Down Arrow keys on the keyboard.
- Hold Command-Arrow keys to move in smaller increments.
- Hold Shift-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 Bezier 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 needs to be linearized. A linear segment aligns the Bezier 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 Smooth button in the toolbar (Shift-S).

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 cursor 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 Bezier Handles

For Bezier 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.

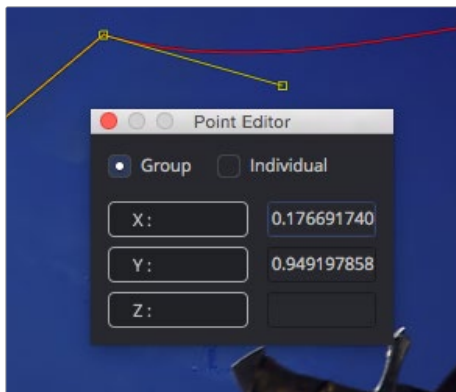


Hold Command to adjust one handle independently.

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 On-Screen 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 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. If you want to select the outer shape, use the Tab key to cycle between the on-screen controls until the dashed outline is visible, or you can select the outer polyline using the contextual menu's 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 to change the shape of the mask, and 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 Bezier 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 on-screen 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 18

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	429
Stereoscopic Overview	429
Optical Flow Overview	429
Toolset Overview	429
Working with Aux Deep Channels	430
Optical Flow Workflows	431
OpticalFlow	431
TimeSpeed, TimeStretcher	431
SmoothMotion	432
RepairFrame, Tween	432
Stereoscopic Workflows	432
Stereo Camera	433
Stereo Materials	434
Disparity	434
NewEye, StereoAlign	434
DisparityToZ, ZToDisparity	434
Separate vs. Stack	434
Setting Up Stereo in the Node Editor	435
About the Disparity Channel	435
Viewing of Disparity and Vector Channels	436
Stereo and Optical Flow Best Practices	436
Semi-Transparent Objects	436
Motion Blur	437
Depth of Field	437
Where to Calculate Disparity and Optical Flow?	437
Cropping the Source	437
Nodes with Multiple Outputs	437
Picking from Aux Channels	438
Vector and Disparity Channels	438

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.

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:** for combining stereo images to create a single anaglyph image for viewing.
- **Stereo > Combiner:** stacks 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:** for adjusting 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 Back Vector 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 46, "Miscellaneous Nodes."

SmoothMotion

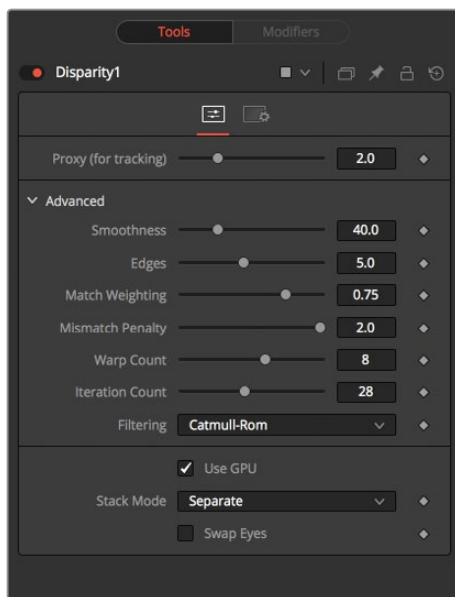
SmoothMotion can be used to smooth the Vector and BackVector channels or smooth the disparity in a stereo 3D clip. This node passes through, modifies, or generates new aux channels, but does not destroy them.

RepairFrame, 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 47, "Optical Flow."

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.

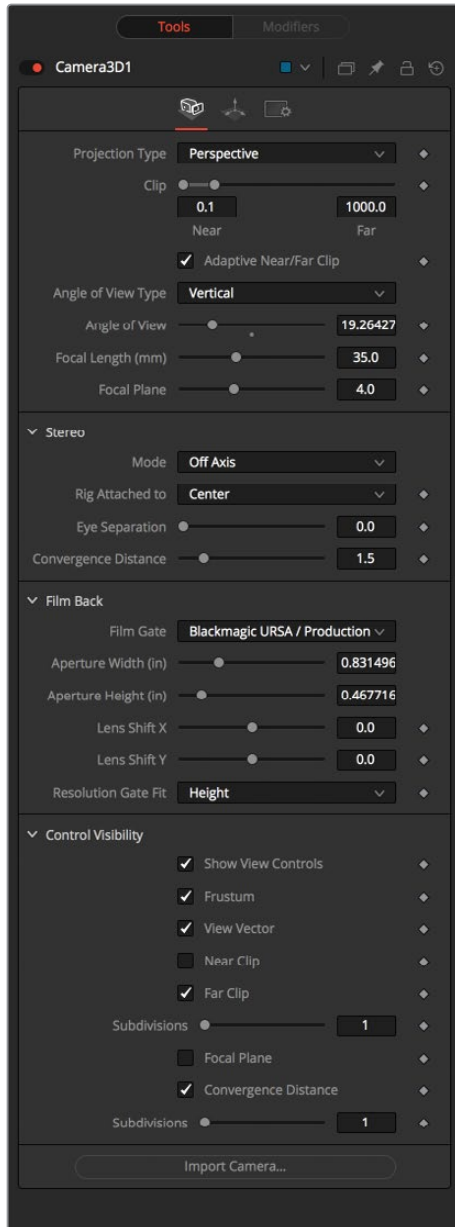
TIP: When connecting Stereo pairs in the node tree, make sure that the left and right images are connected to the left and right inputs of the Disparity node.

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.




Stereoscopic cameras can be done with a single camera or two connected cameras.

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.

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

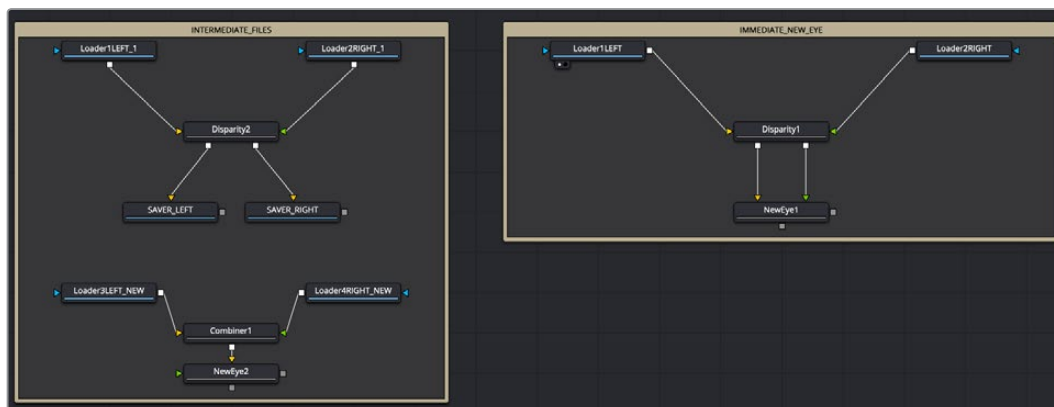
Stereo nodes can work in “Separate” or “Stack” modes. When in Stack mode, the left/right eyes are stacked horizontally or vertically, forming one image with double width or height, respectively.

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 to happen. 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:

$(x_{\text{left}}, y_{\text{left}}) + (D_{\text{left}}.x, D_{\text{left}}.y) \rightarrow (x_{\text{right}}, y_{\text{right}})$
 $(x_{\text{right}}, y_{\text{right}}) + (D_{\text{right}}.x, D_{\text{right}}.y) \rightarrow (x_{\text{left}}, y_{\text{left}})$

You would expect for non-occluded pixels that $D_{\text{left}} = -D_{\text{right}}$, 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. 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 $(\text{image.Width}/\text{image.OriginalWidth}, \text{image.Height}/\text{image.OriginalHeight})$.

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 float32 1080p image containing just RGBA uses about 32MB of memory, but with all the aux channels enabled it consumes around 200MB 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 your input images down 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:

$(x_{\text{left}}, y_{\text{left}}) + (D_{\text{left. x}}, D_{\text{left. y}}) \rightarrow (x_{\text{right}}, y_{\text{right}})$ $(x_{\text{right}}, y_{\text{right}}) + (D_{\text{right. x}}, D_{\text{right. y}}) \rightarrow (x_{\text{left}}, y_{\text{left}})$

You would expect for non-occluded pixels that $D_{\text{left}} = -D_{\text{right}}$, 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 $(\text{image.Width}/\text{image.OriginalWidth}, \text{image.Height}/\text{image.OriginalHeight})$.

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 them to contain flow forward/back by 1 frame.

NOTE: Currently DoD/Rol is not supported for all Fusion nodes.

Chapter 19

Using OFX Plug-Ins

Fusion supports third-party OFX plug-ins for expanding the list of visual effects tools.

Contents

Using OFX	442
-----------	-----

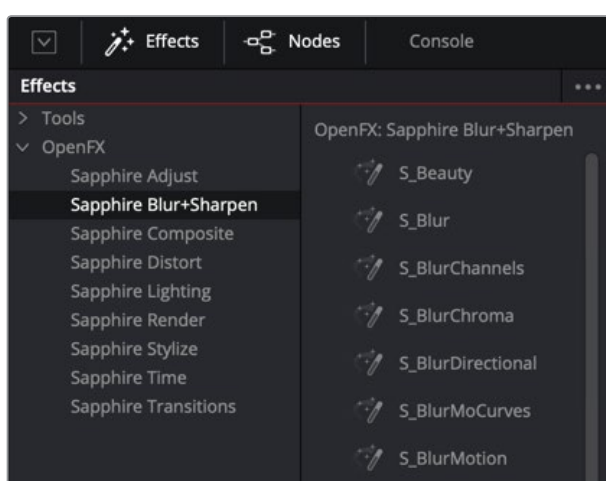
Using OFX

Fusion is able to use compatible OpenFX (OFX) plug-ins from other companies that are installed on your workstation.

When you open Fusion, during the startup time, it looks in specific folders to locate compatible OFX plug-ins. Where it looks for the plug-ins varies, depending on your operating system:

- **For macOS systems:** /Library/OFX/Plugins
- **For Windows system:** Fusion Studio 16: C:/Program Files/Common Files/OFX/Plugins
- **For Linux system:** /usr/OFX/Plugins


After the plug-ins are found, they are displayed in the Effects Library under the OpenFX category.



OpenFX plug-ins in the Effects Library.

Plug-ins work the same as native nodes in Fusion. You add them to the Node Editor and connect them the same way. Their parameters appear in the Inspector where you can modify and animate them.

Some plug-ins provide a custom user interface beyond the sliders and controls in the Inspector.



Custom controls button for Mocha Pro plug-in.



PART 5

Animation

Chapter 20

Keyframing in Fusion

This chapter covers how you can keyframe effects in the Inspector, and how you can edit clips, effects, and keyframes in the Keyframe Editor.

Contents


Keyframing in the Inspector	446
Removing Animation in the Inspector	446
Attaching a Parameter to an Existing Animation Curve	447
Keyframe Editor Overview	447
Keyframe Editor Tracks	448
The Timeline Header	448
The Playhead	448
Spreadsheet	449
Scaling and Panning the Timeline	449
Working with Segments in the Timeline	450
Moving Segments in the Timeline	450
Trimming Segments	450
Holding the First or Last Frame	451
Working with Keyframes in the Timeline	451
Drag and Drop Keyframe Editing	451
Keyframe Editing Using the Time Editor	452
The Keyframe Spreadsheet	452
Duplicating Spline Keyframes	452
Time Stretching Keyframes	453
Showing Keyframe Values	453
Timeline Filters	453
Selected Filtering	455
Sorting in the Timeline	455
The Sort Menu	455
Guides	456
Working with Guides	456
Autosnap	457
Autosnap Points	457
Autosnap Guides	458
The Spreadsheet Editor	458
Selecting a Node to Edit	458
Inserting Keyframes	459
Selecting Multiple Nodes to Edit	459
Customizing the Keyframe Editor	459
Line Size	459
Display Point Values	460

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 are available within the Inspector next to each keyframable parameter. These controls are:

- A gray Keyframe button to the right of 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, 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 Keyframe 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 Path1 from the contextual menu (Path1 may be numbered differently depending on how many parameters are animated).

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:

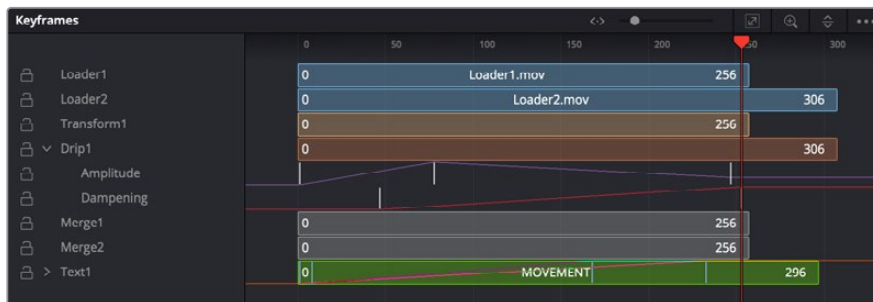
- 1 Right-click on the second parameter you want to attach.
- 2 In the contextual menu, hover over the Connect To submenu.
- 3 In the Connect To submenu, choose the name of the animated parameter.

Keyframe Editor Overview

The Keyframe 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 Keyframe 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 Keyframe Editor.

To show the Keyframe Editor, do one of the following:

- Click the Keyframe Editor button in the UI toolbar to toggle visibility of the Keyframe Editor on and off.
- Press F7 on the keyboard.

Keyframe 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.



The Timeline header area.

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 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 Keyframe 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.

- 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 Keyframe Editor.
- A Zoom to Rect tool lets you draw a rectangle to define an area of the Keyframe 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 discontinuous 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 Keyframe 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 and Effect nodes:

- Trimming a Loader 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 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 discontinuous 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 pop-up and editing field at the bottom right of the Keyframe 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 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 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.

	0						
Text1							
CharacterSpacing							
Center : Path1 : Displacement							
Key Frame :	8.0	168.0					
Text1CharacterSpacing	1	1.386					

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 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 Keyframe 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 Keyframe 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.

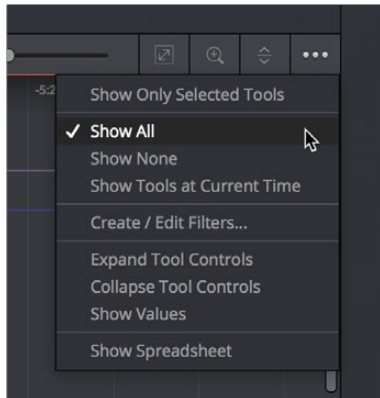
To use a Timeline filter:

Open the Keyframe 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:

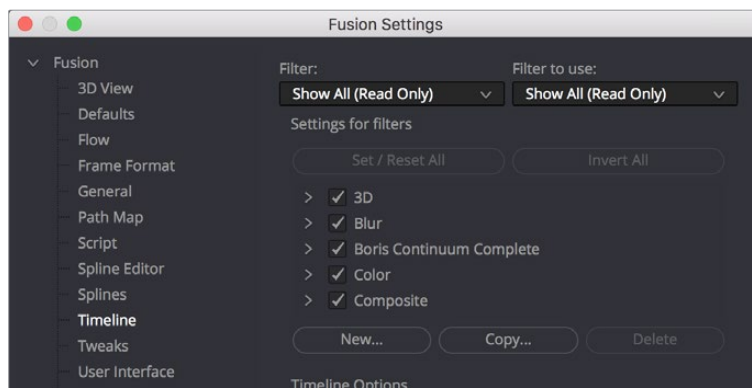
- Choose Show All from the Keyframe Editor Option menu. All layers will reappear.



Choosing a Timeline filter.

To create a Timeline filter:

- 1 Choose Create/Edit Filters from the Keyframe 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 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 Keyframe Editor Option menu.

To delete a filter:

- 1 Choose Create/Edit Filters from the Keyframe 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 Keyframe 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 Keyframe Editor.

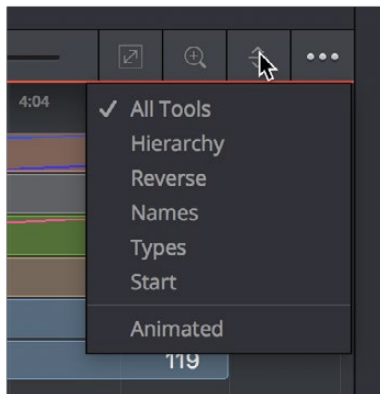
Sorting in the Timeline

There are a couple ways you can change the order in which the nodes are displayed from top to bottom in the Timeline.

- You can freely drag them into any order you like.
- You can use the Sort pop-up menu.

The Sort Menu

The Sort menu reorders how the layers of each node appear in the Keyframe 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.

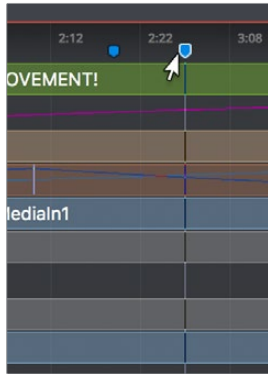


The Timeline Sort Order menu.

- **All Tools:** Forces all tools to be displayed in the Keyframe 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:** A filter that 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.

Guides

Guides are designed to 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. Guides are created in the Timeline Ruler, where they appear as a small marker with a line extending vertically through the graph view.



A guide being moved in the Timeline.

To create a guide, do the following:

Right-click at a frame in the Timeline Ruler of the Keyframe Editor and choose Add Guide from the contextual menu.

Working with Guides


The most important attribute of a guide is its position. For it to add value, it needs to be placed on the frame you intended it to be on. Hovering the cursor over a guide displays a tooltip with its current frame position. If it is on the wrong frame, you can drag it along the Timeline to reposition it.

Jumping to Guides

Double-clicking a guide jumps the playhead to that guide's position.

Renaming Guides

By default, a guide uses the frame number it's on as its only 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 guide, right-click at the top of the guide and choose Rename Guide from the contextual menu. Enter a name in the dialog and click OK.

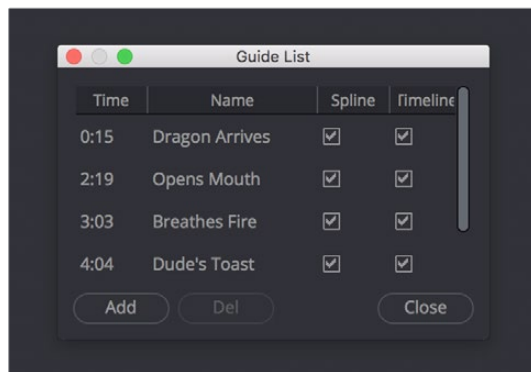


The Guide contextual menu is accessed by right-clicking on the top of the guide.

Show Guide List

Guides can be used to jump to specific locations in a composition using the Guide List. If you right-click the top of a guide to bring up the contextual menu, you can choose Show Guide List, or press Shift-G, to display the Guide List dialog. The Guide List is a floating dialog that will remain on top of the main window until closed.

The Guide List shows all the current guides in the composition, listed according to their position in time along with any custom name you've given them. If you double-click a guide's name from the list, the playhead jumps to the guide's location.



The Guide List dialog allows you to navigate through a composition using guides.

There is a pair of checkboxes beside the names of each guide. One is for the Spline Editor, and one is for the Keyframe Editor (labeled Timeline). By default, guides are shown in both the Spline Editor and Keyframe Editor, but you can deselect the appropriate checkbox to hide the guides in that view.

Deleting Guides

You can delete a guide by dragging it up beyond the axis labels and releasing the mouse. You can also use the Guide's contextual menu to choose Delete Guide.

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 guides. The Autosnap option is accessed through the Timeline'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 guides.

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 Keyframe 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 guides in the Timeline.


Autosnap Guides

When you click to create a new guide, the default behavior is that it will snap to the closest frame. If you reposition the guide, it also snaps to the nearest frame as you drag. This behavior can be changed in the Keyframe Editor's contextual menu by choosing from the Options > Autosnap Guides submenu. The options are:

- **None:** Guides can be placed anywhere with subframe accuracy.
- **Frame:** Frame forces all guides to snap to the nearest frame.
- **Field:** Field forces all guides to snap to the nearest field.

The Spreadsheet Editor

The Spreadsheet Editor is a separate panel that can be displayed beneath the Keyframe Editor. It is used to compactly show the numeric values of the keyframes for selected parameters in the Timeline header, via a table with rows and columns, showing time and value.



The screenshot shows the Spreadsheet Editor interface. On the left is a list of parameters: Drip1, Amplitude, Dampening, Merge1, Merge2, and Text1. The 'Amplitude' parameter is selected. The main area displays a table of keyframe data for 'Drip1Amplitude'. The table has columns for time (1.0, 40.0, 78.0, 213.0, 247.0, 269.0) and values (0.031, 0.055, 0.11, 0.055, 0.055, 0.039). Above the table, a timeline visualization shows keyframes as vertical lines and segments as horizontal bars. The 'MOVEMENT' segment is highlighted in green.

Key Frame :	1.0	40.0	78.0	213.0	247.0	269.0	
Drip1Amplitude	0.031	0.055	0.11	0.055	0.055	0.039	

The Spreadsheet Editor showing editable data for six keyframes.

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 Timeline interface.

Selecting a Node to Edit

To display a node's timing in the Spreadsheet, select the node's name in the Timeline 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 Timeline 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.

78.0	213.0	247.0
0.11	0.055	0.055

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 Command, you can click on additional parameters on different nodes to add to the Spreadsheet.

TIP: You can use the Tab and Shift-Tab key shortcuts to move the selection right or left in the Spreadsheet Editor.

Customizing the Keyframe Editor

There are a few ways you can change the appearance of the Keyframe Editor to better fit your needs. All these options are found by right-clicking anywhere within the Keyframe 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 height of just one segment:** Right-click anywhere within the Keyframe 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 Keyframe Editor and choose a size from the All Line Size submenu. The options are Minimum, Small, Medium, Large, and Huge.

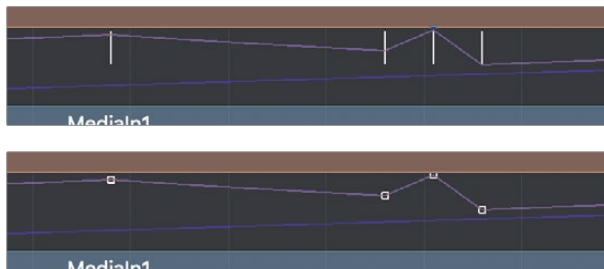
Display Point Values

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 Keyframe 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).

Chapter 21

Animating with Motion Paths


Moving layers and 3D objects can have their animation edited and controlled using motion paths. This chapter discusses how you can use motion paths in Fusion.

Contents

Animating Using Motion Paths	463
Types of Motion Paths	463
Poly Paths	463
Path Modifier	465
Displacement Splines and Path Timing	466
Types of Control Points	467
Locked Points	467
Unlocked Points	467
Locking and Unlocking Points	471
XY Paths	471
Tips for Manipulating Motion Paths	472
Path Centers	472
Copying and Pasting Motion Paths	472
Removing Motion Paths	473
Recording Motion Paths	473
Importing and Exporting Polylines	473
Native Format	473

Animating Using Motion Paths

Motion Paths are polylines that define the movement for two- and three-dimensional controls, such as the Center X/Y and Pivot X/Y parameters of Transform, mask, and effect nodes. Keyframing these kinds of parameters creates visible motion paths in viewers with which you can visually adjust how different effects move. For all motion paths, the coordinate control represents the position of an object or effect, such as a Merge node's center or published polyline points on a mask. Coordinate controls are represented on-screen with a crosshair or an X.



A Center Offset on-screen control for an Ellipse node.

It's not possible to add a motion path to a one-dimensional value, like a 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: Poly paths, XY paths, and 3D motion paths for 3D scenes.

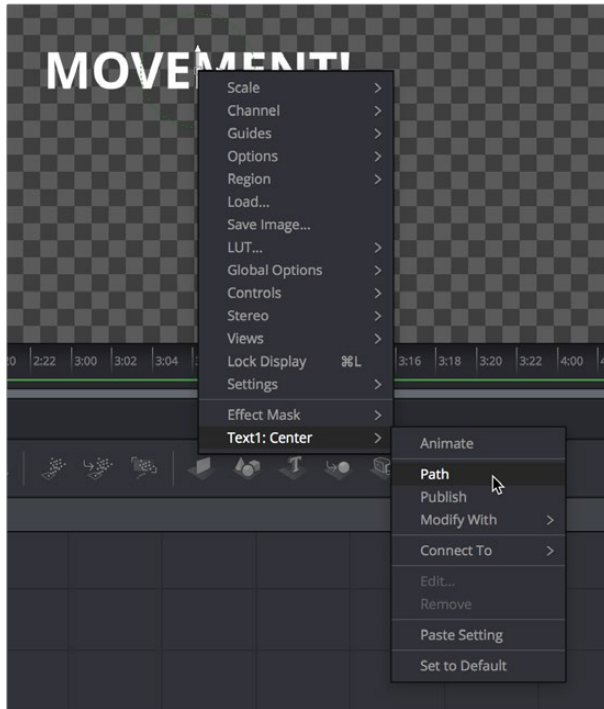
- A Poly path uses two splines to record the path, one for shape of the path and a displacement spline for the timing along the path. The Poly path is the default type for motion paths and most documentation in this chapter assumes that this type is used.
- The XY path type employs a spline for the X position of the point and another for the Y position. The XY paths are explained in detail toward the end of this chapter.
- 3D motion paths pertain only to positional controls within 3D scenes.

Poly Paths

Poly paths are reasonably easy to work with. They're similar to X/Y paths when you're working in the viewer, but they provide a simpler curve editing experience in the Spline Editor.

To create a Poly path, do the following:


- 1 Position the playhead on the frame where the motion will begin.
- 2 Position the control for the layer, effect, or mask at its starting position.
- 3 Right-click the on-screen center control in the viewer, and choose Path from the contextual menu for that control.



Initiating path keyframing for the center position of a text title.

The Center X/Y parameter for that node will display keyframes at those frames to indicate that the parameter is now animated. Any changes made to the control will cause a new keyframe to be created.

- 4 Move the playhead to a new frame.
- 5 Drag the on-screen control or adjust the Offset or Center values in the Control panel. A keyframe is automatically created on the motion path, and a polyline is drawn from the original keyframe to the new one.



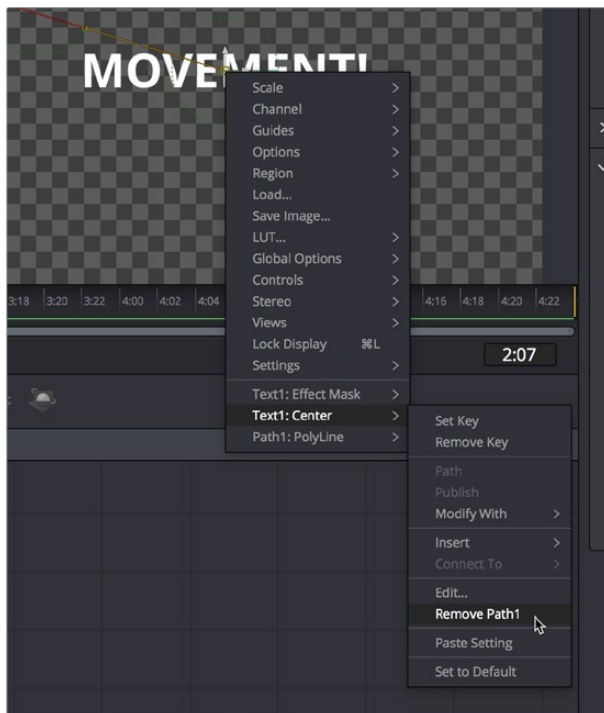
Two keyframes spaced several frames apart display a motion path showing the direction of animation.

- 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 control 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 pressing Command-I or clicking the 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.

To remove a Poly path:


- 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.

Path Modifier

There's also an XY Path modifier that lets you add a motion path to any compatible control. Simply right-click on the name of a compatible parameter in the Inspector (such as Center X/Y) and choose Modify With > X/Y Path to add that modifier.




Adding the XY Path modifier to a Center parameter in the Inspector.

When you create a path this way, controls for that path appear in the Modifiers tab of the Inspector.



XY Path modifier controls in the Modifier tab of the Control panel.


Now you can use the modifier controls to create a motion path, while using the object's original controls as an offset to this motion path, if you need to. Using the XYZ, Center X/Y, Size, Angle, and Heading Offset controls automatically sets keyframes, and an editable motion path appears in the viewer once you've set two or more keyframes to create animation.



Using an XY Path modifier to animate a piece of text.

Displacement Splines and Path Timing

Every Poly path has an associated displacement spline in the Spline Editor. The displacement spline represents acceleration, or the position of the animated control 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.

Displacement splines are used to control the speed of the movement along a path. To slow down, speed up, stop, or even reverse the motion of the control along the path, adjust the points for the path's displacement in the Spline Editor.

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.


Types of Control Points

Displacement 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 will have an associated point on the displacement spline; unlocked points will not have a corresponding point on the displacement spline. 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 in the Spline Editor.



A locked point in the viewer.

Deleting a locked point from the motion path will change the overall timing of the motion.

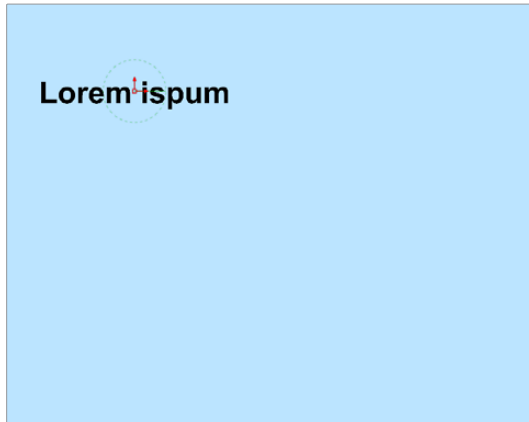
Unlocked Points

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.

Unlocked points do not have corresponding points on the path's displacement spline. They are represented in the viewer as smaller, solid square points.

To experience the difference between locked and unlocked points, do the following:

- 1 Add a Text node to the Node Editor and type a word in the Text node.
- 2 Position the text's center in the upper left corner of the frame.




Text placed in the upper left corner of the frame.

- 3 Set the Playhead at frame 0.
- 4 In the viewer, right-click on the text center and choose Animate from the contextual menu.

This creates the first locked point of the path.


- 5 Position the playhead on frame 45.
- 6 Move the text center to the center of the screen.



Moving the playhead and repositioning the text adds a locked point.

This sets the second locked point.

- 7 View the Spline Editor and display Path 1's displacement spline.



The path's displacement spline.

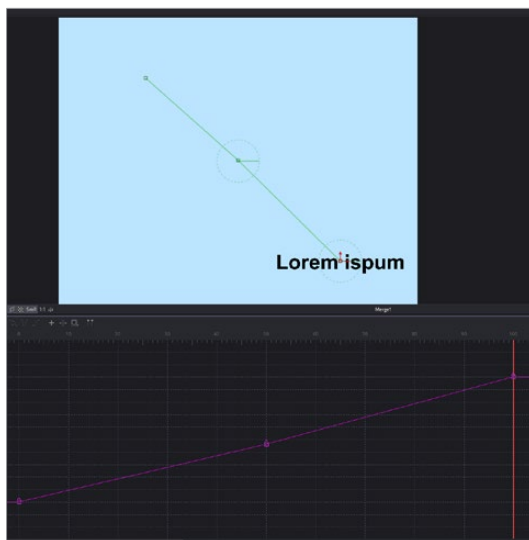
At a value of 0.0, the control will be located at the beginning of the path. When the value of the displacement spline is 1.0, the control is located at the end of the path.

- 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 text center to the bottom right corner of the screen.



Position the playhead at 100 and drag to reposition the text, creating another locked point.

This will create an additional locked point and set a new ending for the path.

- 10 Select the motion path spline by using the Tab key to cycle controls until the path is selected.

Currently, the path is in Click Append mode.


- 11 To insert some points, click the Insert and Modify button in the toolbar.

- 12 Click on the path and create two new points: one part-way between the first and the second points, and the other part-way between the second and the third.



Add two new points to the path.


- 13 Smooth these new points to create a curve in the path.



Adjust the path using the newly added 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. This behavior can be used to make the layer's motion pause briefly.

- 14 Select the point at frame 50 in the displacement spline and hold down the Command key while dragging it to frame 65. The point is copied at frame 65.



Create a pause in the motion by copying locked points.

- 15 Select the point at frame 0 and at frame 50 and press Shift-S while in the Spline Editor to smooth the displacement curve.

This will cause the motion to accelerate slightly at the beginning and slow to a stop at the center.

- 16 Render a preview of frames 0 to 100 to see the results so far.
- 17 Remove the preview when done.
- 18 In the viewer, delete the unlocked point added between the first point and the middle point. The spline segment of the motion path will become linear.
Deleting this point did not change the amount of time it takes the spline to reach the center of the screen, only the path it takes to get there.
- 19 Step through a few of the frames in the scene to confirm this.
- 20 Now delete the locked point in the center of the screen.
Removing this point changes the timing of the animation rather profoundly because the key point in the displacement spline is also removed.

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.

XY Paths

The XY path type uses a separate spline in the Spline Editor to calculate position along the X-axis, as it does for the position along the Y-axis.

To animate a coordinate control using an XY path, do the following:

- Right-click on the on-screen control and choose Control name > Modify With > XY Path from the contextual menu.

At first glance, XY paths work like displacement paths. To create the path, position the playhead and drag the on-screen control where you want it. Position the playhead again and move the on-screen control to its new position. The difference is that the control points are only there for spatial positioning. There are no locked points for controlling temporal positioning.

The Spline Editor for the XY path displays the X and Y channel splines. Changes to the controls position will be keyframed on these splines. The advantage to the XY path is that you can explicitly set an XY coordinate at a specific time for more control.

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.

Switching Default Paths

If you want to change the Default Path type to XY path, you can choose Fusion > Preferences > Globals and select the Defaults category. In the Point With pop-up menu, choose XY path. The next time Animate is selected from a coordinate control's contextual menu, an XY path will be used instead of a displacement path.

Tips for Manipulating Motion Paths

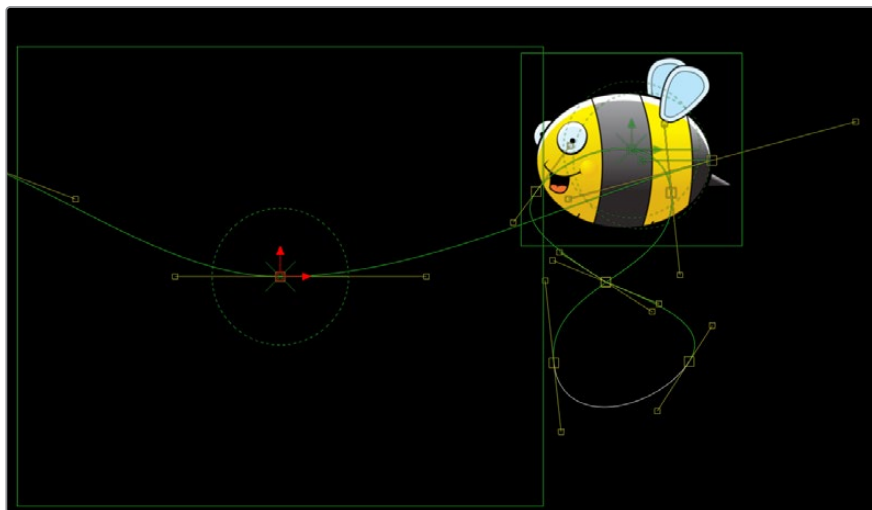
There are a variety of ways you can create and edit motion paths in the viewer.

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.

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.



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 Modifier 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.

In all cases, the old motion path will be overwritten with the one in the clipboard.

Removing Motion Paths

To remove a control from a motion path, right-click on the motion path's control header and select Delete from the motion path's context menu. Alternatively, right-click on the animated control and select Remove "Name of Modifier."

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 Polyline

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, you 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 22

Animating with Modifiers and Expressions

This chapter discusses how modifiers and SimpleExpressions can be used to control parameters and automatically create animations that would be difficult to achieve manually.


Contents

Modifiers	476
Adding the Right Modifier for the Job	476
Knowing Which Parameters Are Modified	476
Using the Modifier Tab	477
Connecting Multiple Parameters to One Modifier	477
Adding Modifiers to Modifiers	477
SimpleExpressions	478
Pick Whipping	480

Modifiers

Parameters can be controlled with modifiers to help you 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 can be complex expressions, procedural functions, external data, third-party plug-ins, or fuses.

You can add modifiers to an input through the Inspector's contextual menu or on the Preview control in the viewer.



Inspector contextual menu Modify With submenu.

Adding the Right Modifier for the Job

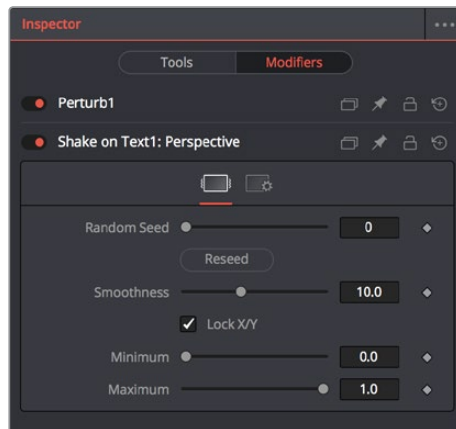
Which modifiers are available depend on the type of parameter you're trying to add one to. Numeric values, text, polylines, gradients, and points each have different sets of modifiers that will work with them.

Knowing Which Parameters Are Modified

The Fusion user interface will alert you that an input is modified by adding a keyframe highlight for that parameter in the Inspector. Some parameters, like those for numbers and points, will also color the numerical entry field.

Using the Modifier Tab

Modifiers with additional UI will be displayed in the Modifiers tab of the Inspector. When a selected node has modifiers applied, a marker will appear on the Modifiers tab as an indication.



The Modifiers tab.

Modifiers appear with header bars and header controls just like nodes. A modifier's title bar can also be dragged into a viewer to see its output.

Connecting Multiple Parameters to One Modifier

Once a modifier has been applied to a parameter, you can connect another parameter to that modifier so that the modifier can affect both parameters. This can be handled through the Connect To contextual menu. As with modifier assignment, the list is filtered by the type of the parameter. This connection is bi-directional; editing either parameter will cause the other parameter to change.

The Publish modifier does nothing by itself, but it does let you connect parameters together without animation or having to use another modifier.

Adding Modifiers to 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.

For more detail on all Modifiers available in Fusion, see Chapter 55, "Modifiers."

SimpleExpressions


SimpleExpressions 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.

SimpleExpressions can be added from the parameter's contextual menu.



Right-clicking on a parameter to add an Expression from the contextual menu.

A text entry box will appear below the control, and a yellow indicator will appear to the left of the control. The current value of the parameter will be entered into the text box.



A parameter with a SimpleExpression applied.

Inside this text box, you can enter one-line scripts in Lua with some Fusion-specific shorthand. Some examples of SimpleExpressions and their syntax include:

time

This returns the current frame number.

Merge1.Blend

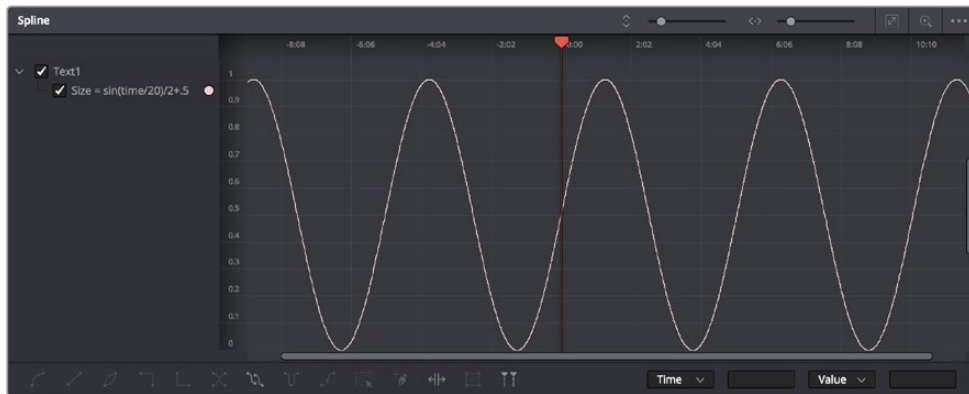
This returns the value of another input, Blend, from another node, Merge1.

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. If you look in the Spline Editor, you can see 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.

```
iif(Merge1.Blend == 0, 0, 1)
```

This returns 0 if the Blend value is 0, and returns 1 if it is not. The `iif()` 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. Another way to get the same result would be:

```
Text1.Center - Point(0,.1)
```

```
Text("Colorspace: "..(Merge1.Background.Metadata.ColorSpaceID )
```

This SimpleExpression returns Text instead of a Number or Point. 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:

```
Rendered Nov 12, 2015 at 15:43
```

```
on the computer Rn309 running Windows_NT from the comp \\SRVR\Proj\
Am109\SlateGenerator_A01.comp
```

The OS library can pull various information about the computer and environment. Here `os.date` with different formatting arguments gets the date and time. Any environment variable can be read by `os.getenv`, in this case the computer name and the operating system. Various attributes from the comp can be accessed with the comp variable, like the filename, expressed as a UNC path. To get a new line in the Text, the escape sequence `\n` is used. When working with long SimpleExpressions, it may be helpful to drag the Tool Control panel out to make it wider or to copy/paste from a text editor or the Console.

For more details about writing SimpleExpressions, see the Fuse Guide, Scripting Guide, and the official Lua documentation.

Pick Whipping

To the left of the SimpleExpression sits a button with a + on it. Clicking and dragging the button onto another control, or “pick whipping,” will let you easily get the name of that control. Hovering over a tab while pick whipping will open that tab.



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 Expression from the contextual menu. The SimpleExpression will be plotted in the Spline Editor, allowing you to see the result over time.

A quick way of setting a SimpleExpression in a Number or Point input is to type = into the text box to the left of the number already there.



PART 6

3D Compositing

Chapter 23

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	485
3D Compositing Fundamentals	486
Creating a Minimal 3D Scene	486
The Elements of a 3D Scene	488
Geometry Nodes	488
The Merge3D Node	490
The Renderer3D Node	493
Software vs. GPU Rendering	494
Software Renderer	494
OpenGL Renderer	495
OpenGL UV Renderer	495
Loading 3D Nodes into the Viewer	496
Navigating the 3D View	498
Transforming Cameras and Lights Using the Viewers	498
Transparency Sorting	499
Material Viewer	500
Transformations	501
On-Screen Transform Controls	501
Pivot	502
Target	502
Parenting	504
Cameras	504
Quickly Viewing a Scene Through a Camera	505
Plane of Focus and Depth of Field	505
Importing Cameras	506
Lighting and Shadows	507
Enabling Lighting in the Viewer	507
Enabling Lighting to Be Rendered	507
Controlling Lighting Within Each 3D Object	508
Lighting Types Explained	508

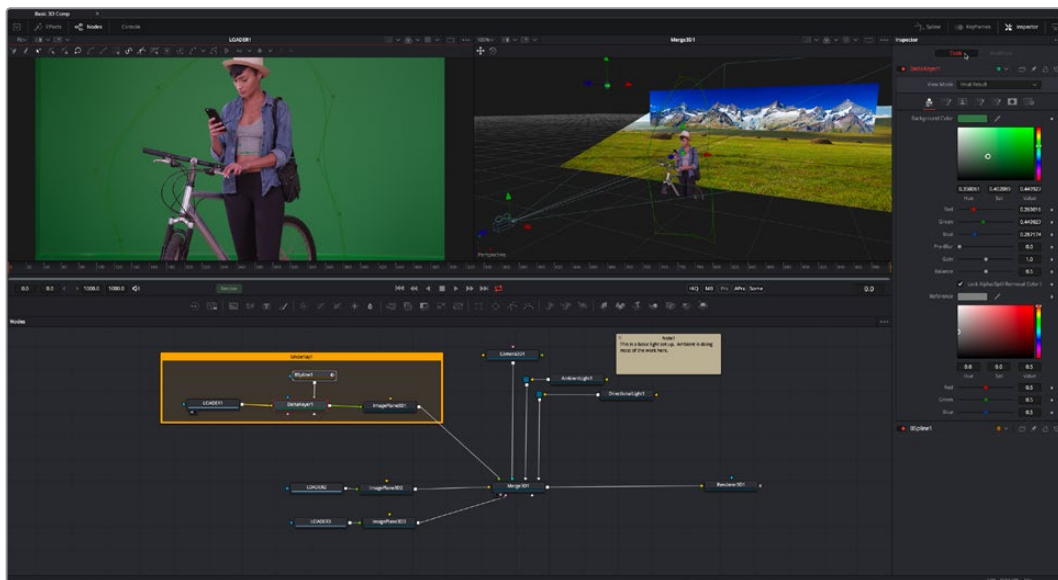
Materials and Textures	512
Material Components	513
Alpha Detail	515
Illumination Models	517
Textures	518
Reflections and Refractions	519
Bump Maps	521
Projection Mapping	522
Project a Texture onto a Catcher Material	523
Project Using the UVMap Node	524
Geometry	525
Common Visibility Parameters	525
Adding FBX Models	526
Using Text3D	527
Fog 3D and Softclipping	531
Material and Object IDs	533
World Position Pass	533
Point Clouds	535

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 Node Editor, you have a complete 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
- Importing matched cameras and point clouds from applications such as SynthEyes or PF Track
- Importing cameras, lights, and materials from 3D applications such as Maya, 3ds Max, or LightWave
- 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



An example 3D scene in Fusion.

3D Compositing Fundamentals

The 3D category or 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.

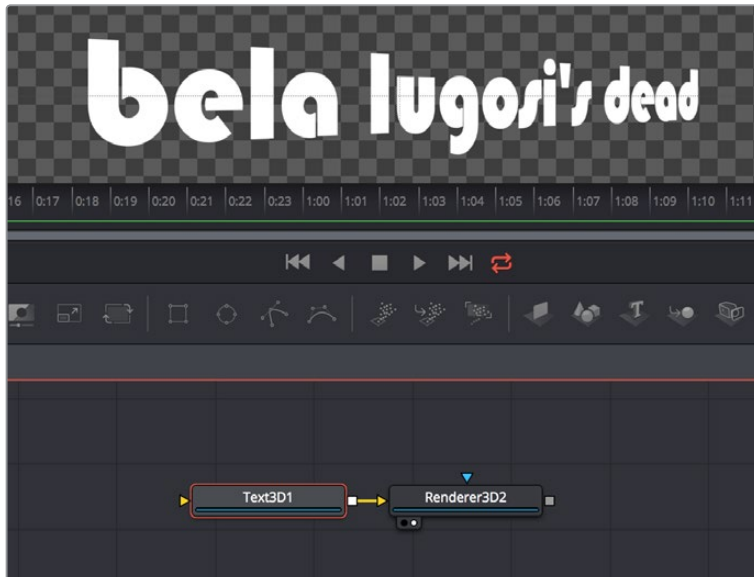


The 3D category of nodes in the Effects Library.

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.



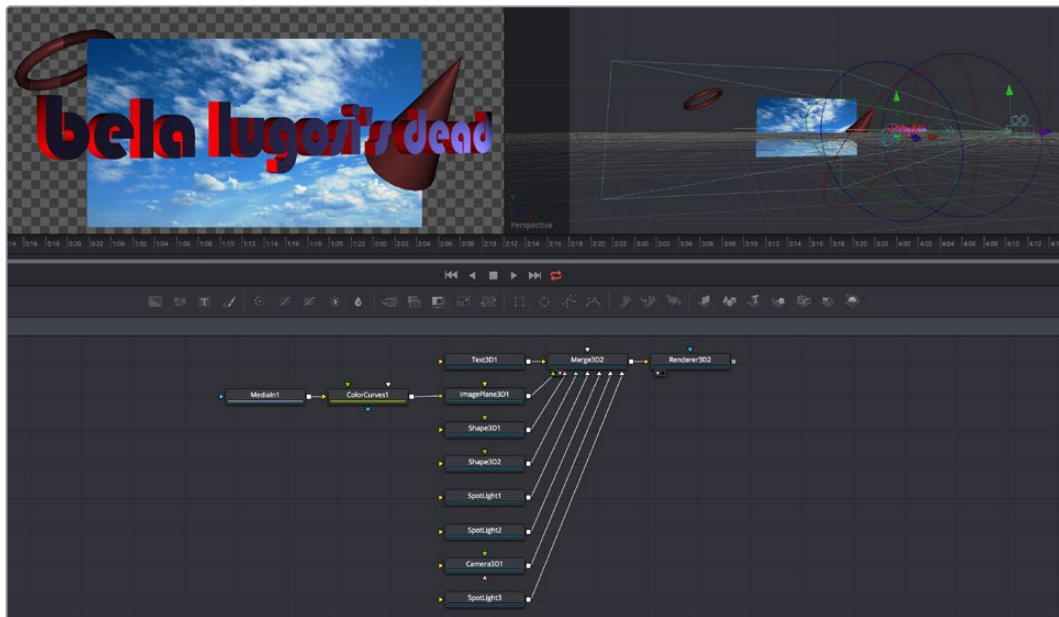
The 3D nodes available from the toolbar include the ImagePlane3D, Shape3D, Text3D, Merge3D, Camera3D, SpotLight3D, and Renderer3D nodes.

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 Taurus, with the image (left), and the shaded taurus (right).

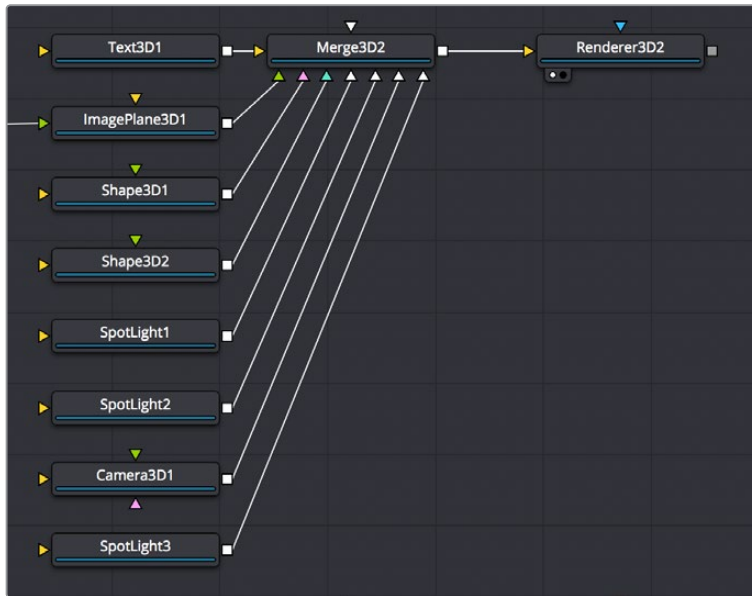
If you're shading or texturing Text3D nodes, you need to add a texture in a specific way since each node is actually a scene with individual 3D objects (the characters) working together. In the following example, the RustyMetal shader preset is applied to a Text3D node using the ReplaceMaterial3D node. 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

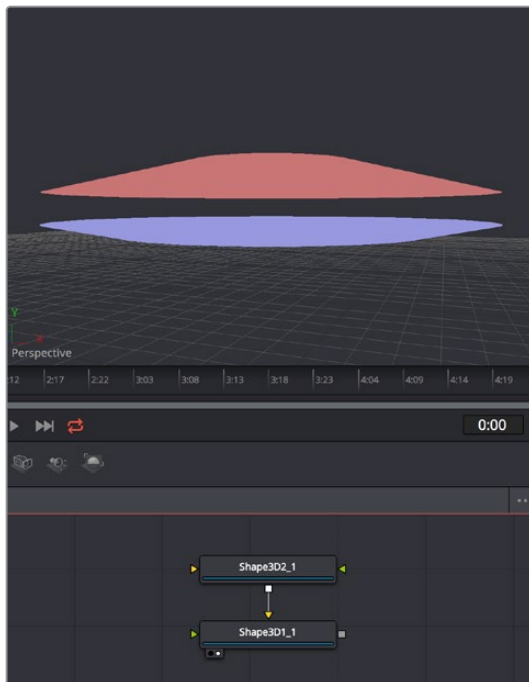
The Merge3D node combines the output of one or more 3D nodes into a single scene. Unlike the 2D Merge node, the ordering of elements in the scene is not restricted to only background and foreground inputs. Instead, the Merge3D node lets you connect an unlimited number of inputs, with the resulting output 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

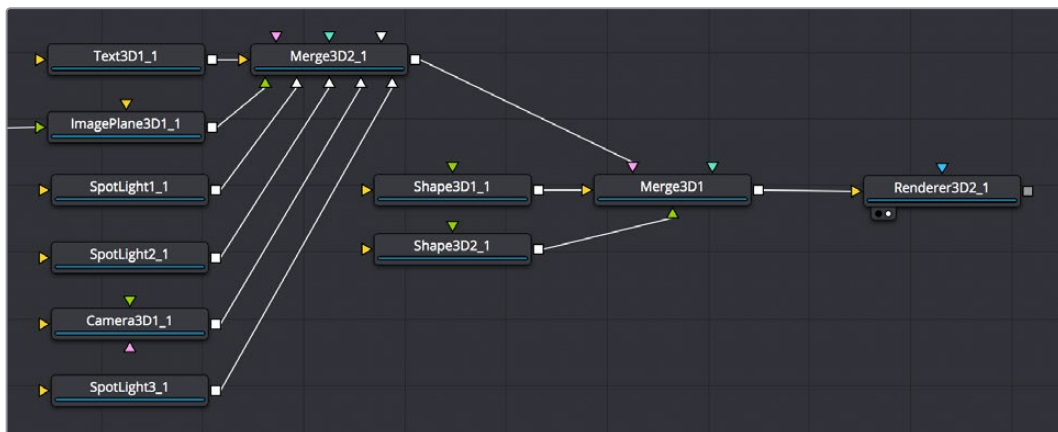
While the Merge3D node provides a structured way of combining objects, you can also combine 3D objects such as Text3D and Shape3D nodes by connecting the output of one 3D object node to the input of another, as seen in the following screenshot. When you do this, you must use each node's internal transform 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 taurus 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 taurus 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.



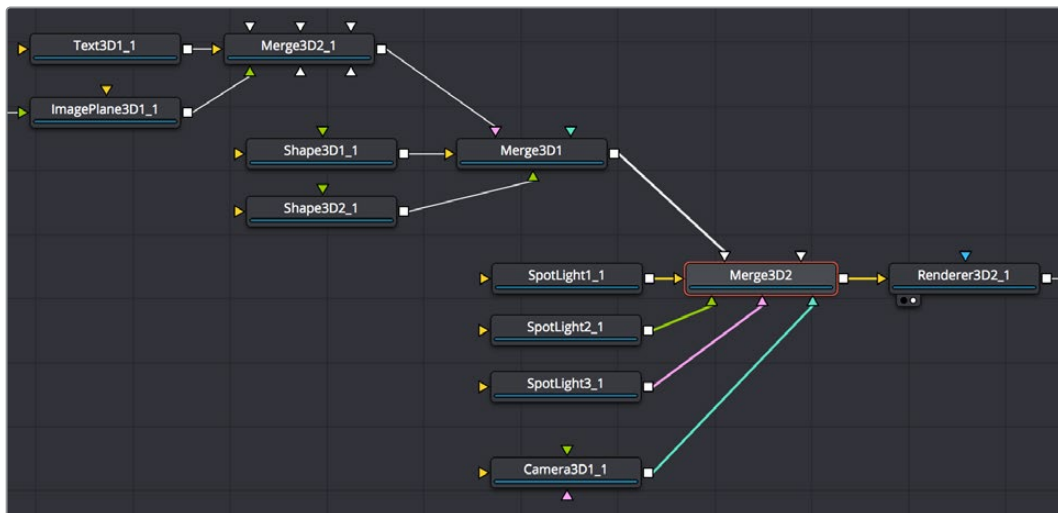
The Transform tab of a Merge3D node.

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.

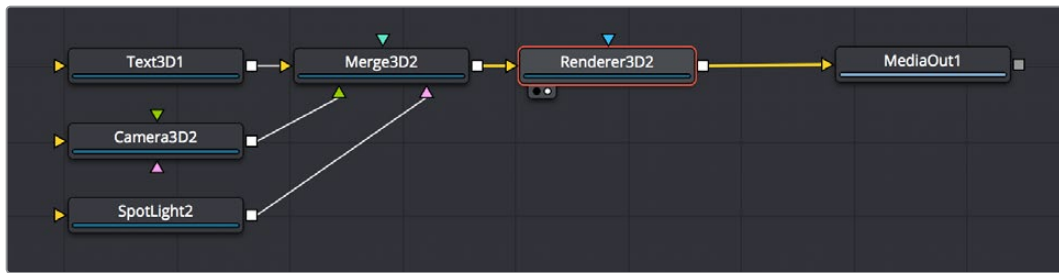


An example of a 3D scene using multiple Merge3D nodes working together; the upstream Merge3D nodes arrange the 3D objects placed within the scene, while the last Merge3D node (orange) lights and frames the scene.

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 connected to a Saver node to output 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” pop-up is set to OpenGL Renderer, you cannot render soft shadows or excessively large textures (left). When the Renderer3D node “Renderer Type” pop-up 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.

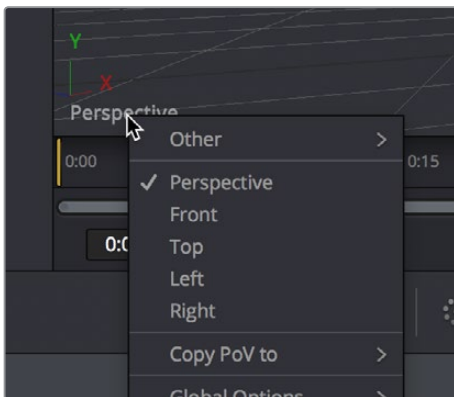
The interactive 3D Viewer is highly dependent on the computer's graphics hardware, relying on support from OpenGL. The amount of onboard memory, as well as the speed and features of your workstation's GPU, make a huge difference in the speed and capabilities of the 3D Viewer.

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.



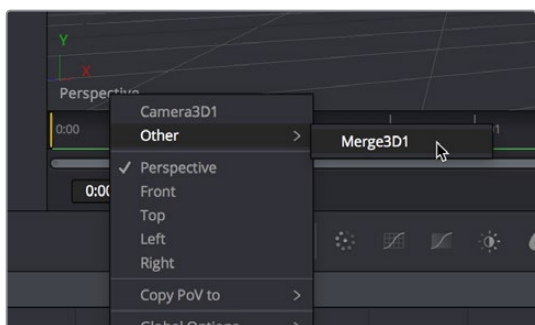
A 3D Viewer's default perspective view

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 pop-up 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 detail about the options available in the 3D Viewer, see Chapter 8, “Using Viewers.”

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.

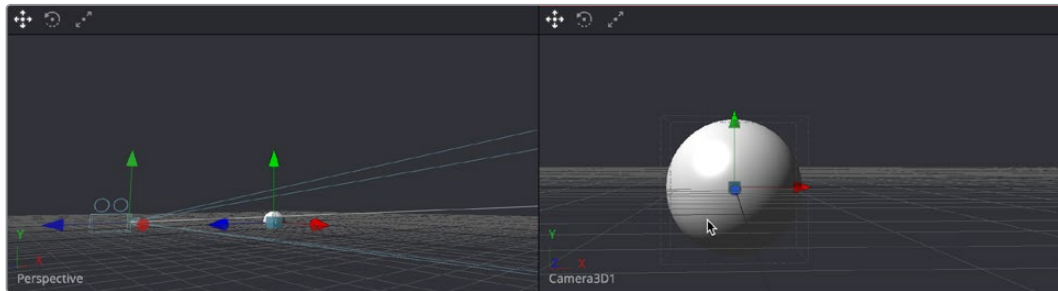
Furthermore, selecting a 3D node in the Node Editor also selects the associated object in the 3D Viewer.

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 two-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.



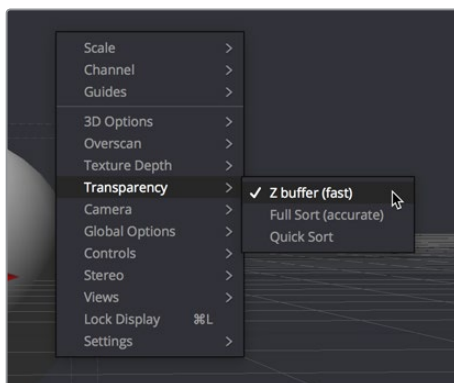
When a viewer is set to display the view of a camera or light, panning, zooming, or rotating the viewer (seen at right) actually transforms the camera or light you're viewing through (seen at left)

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

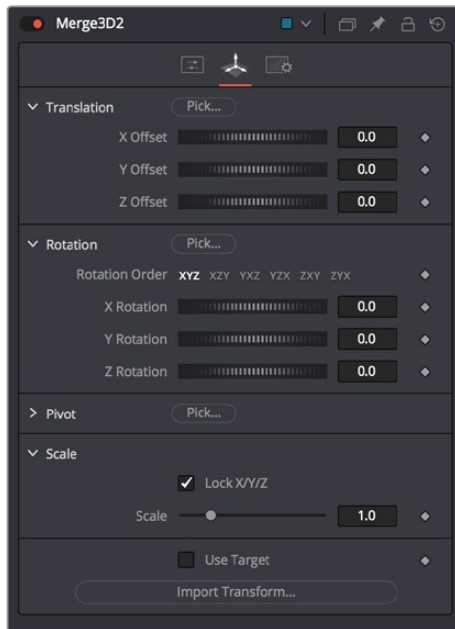
The type of geometry, the renderer, and the state of the lighting can all be set by right-clicking the viewer and choosing options from the contextual menu. Each viewer supports A and B buffers to assist with comparing multiple materials.

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 on-screen controls.

On-Screen Transform Controls

When an object is selected, it displays on-screen 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 On-Screen 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 on-screen control to constrain the transform to that axis, or if you drag the center

of the on-screen 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 on-screen control, while the green and blue portions of the on-screen 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 on-screen 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 an object's Shape3D node.


- 4 Use the X/Y/Z Target Position controls in the Inspector or the Target on-screen 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 taurus facing its on-screen 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.

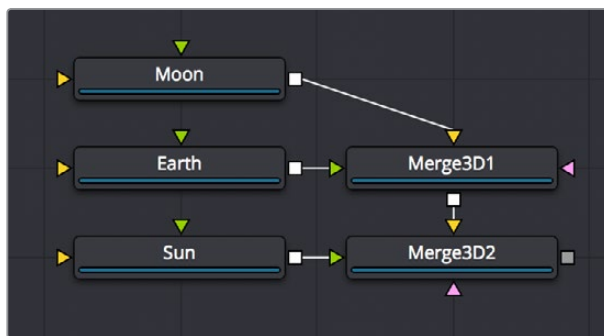


A light made to face the wall using its enabled target control.

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.

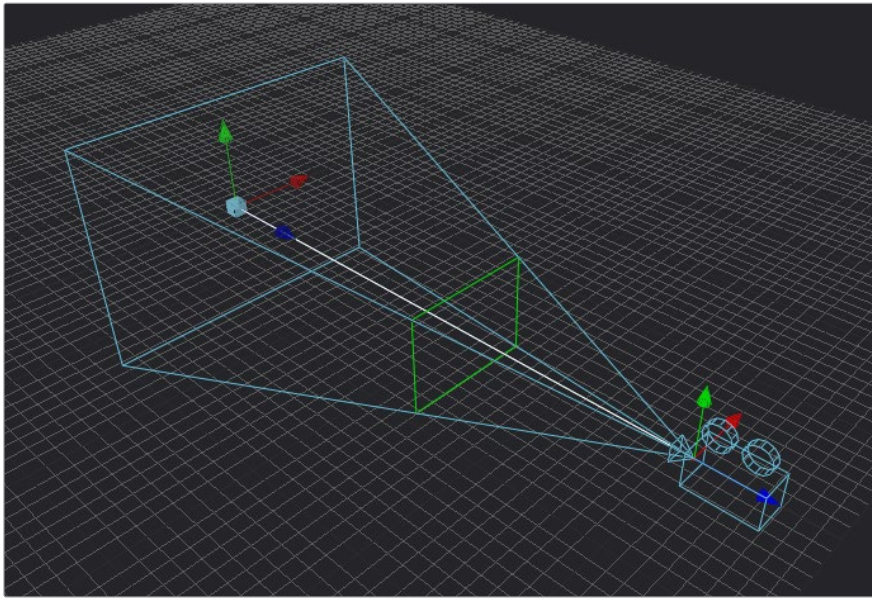
Here are the two simple rules of transforming parented Merge3D nodes:

- 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 on-screen 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.

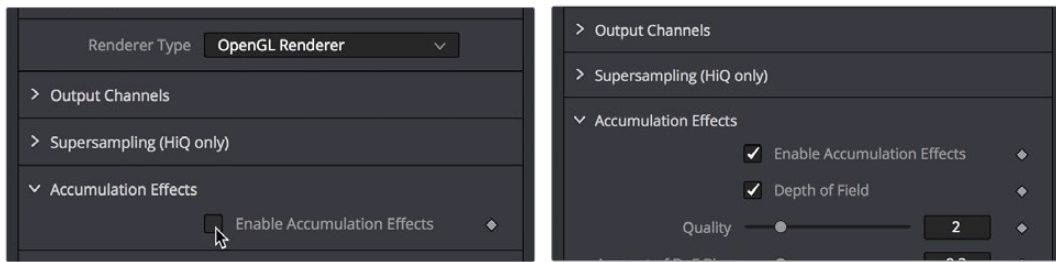
The viewer's frame may be different from the camera frame, so it may not match the true boundaries of the image that will be rendered by the Renderer3D node. If there is no Renderer3D node added to your scene yet, you can use Guides that represent the camera's framing. For more details about Guides, see Chapter 8, "Using Viewers."

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:

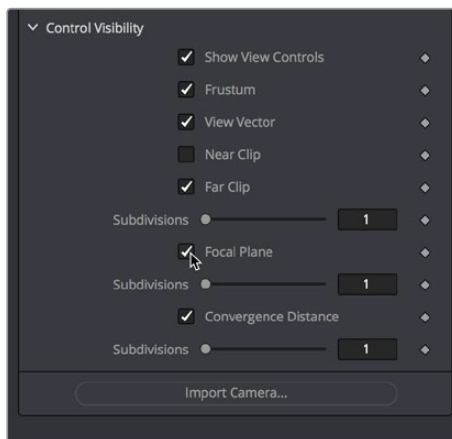
- 1 You must add a Renderer3D node at the end of your 3D scene.
- 2 Select the Renderer3D node, and set the Renderer Type to OpenGL Renderer.
- 3 Open the Accumulation Effects disclosure control that appears, and turn on the Enable Accumulation Effects checkbox in the OpenGL render.



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 the Camera3D section in Chapter 66, “3D Nodes.”


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 Import Camera Control 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.

Controlling 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.

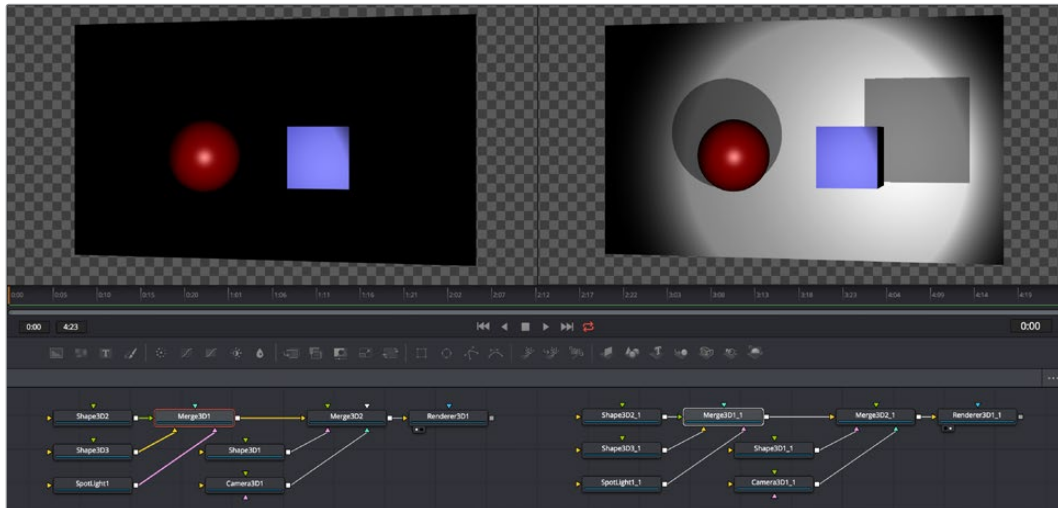


From left to right: Directional light, point light, and spotlight.

All of the Light nodes display on-screen 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.




3D objects have individual lighting controls that let you control how each object interacts with light and shadows.

- **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.

Shadows

The only light that can cast shadows is the spotlight. Spotlight nodes 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

See the “Spotlight” section of Chapter 27, “3D Light Nodes,” for a more detailed description of the shadow controls.

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 a



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 detail, see “Spotlight” in Chapter 27, “3D Light Nodes.”


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 (top), and the corrected shadow shown using Biasing (bottom).

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 RGB = (1, 1, 1), which means 100% of green, blue, and red light pass 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 this 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. There are four different illumination models that 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 20.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

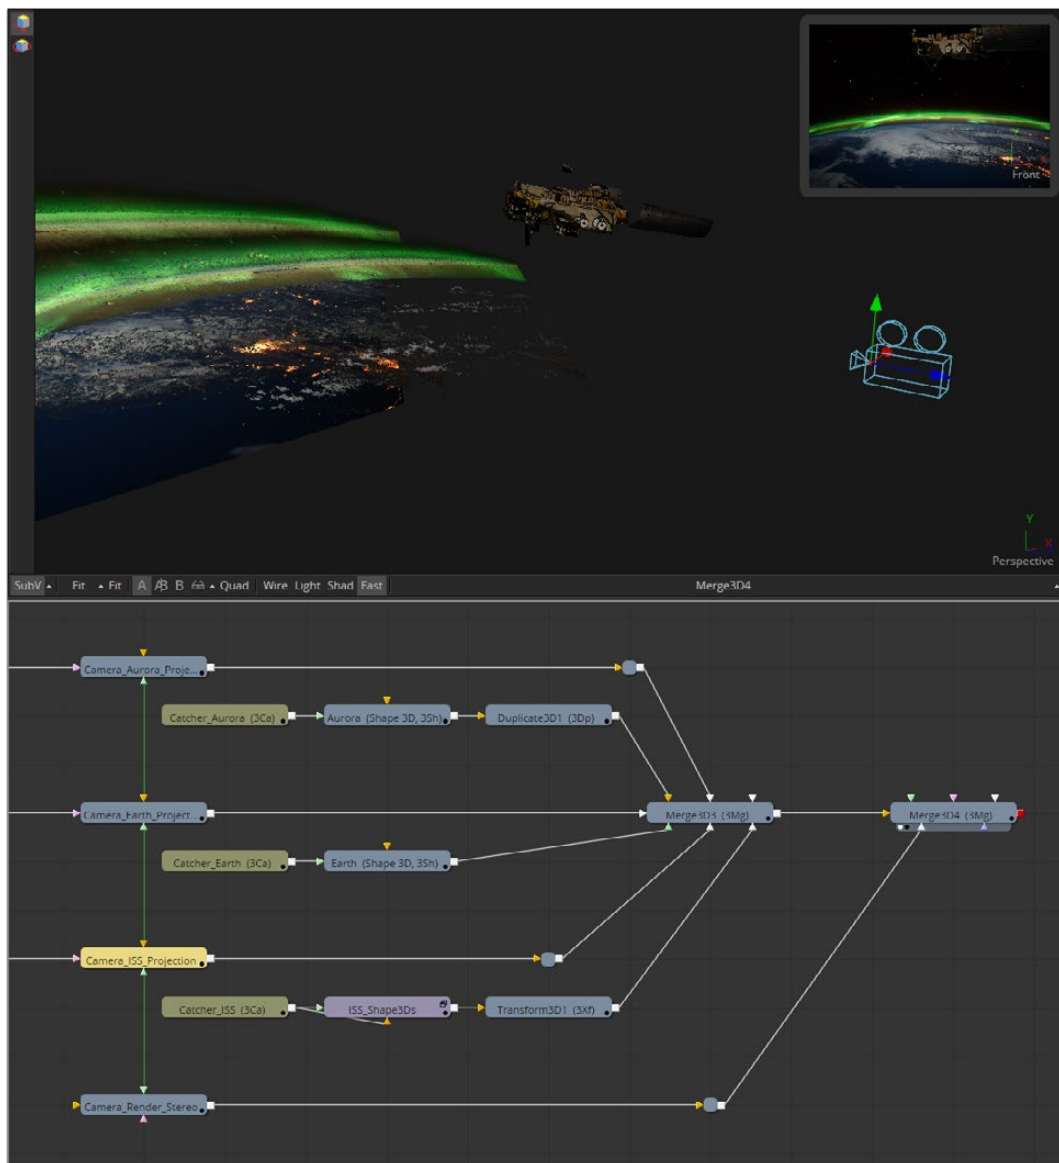
When lighting is enabled, a Camera 3D or Projector 3D can act as a light with all the lighting features. When Camera Projection is enabled or you use a projector, you can choose whether the projection behaves like a spotlight or an ambient light; however, alpha channels cannot be projected. Overlapping 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 or destroyed or reordered (e.g., projection locking will not work on particles because they get created/destroyed, a Cube3D with its subdivision level slider animated, or a Duplicate3D node with its time offset set to non zero so additional meshes get created over time).

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 Part 7, “Node Reference,” 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. See Chapter 89, “Particle Nodes,” for more detail.

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

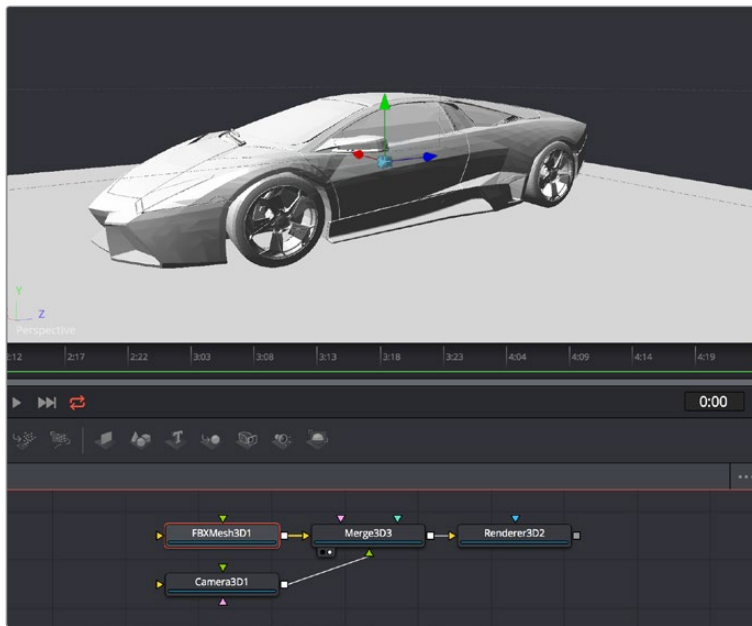
For any piece of geometry, the `Renderer3D` rejects transparent pixels in the auxiliary image channels. The reason this is the default is to prevent aux channels (e.g., normals, Z-channel, UVs) from filling in where there should be blank space or full transparency. For example, suppose in post you want 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

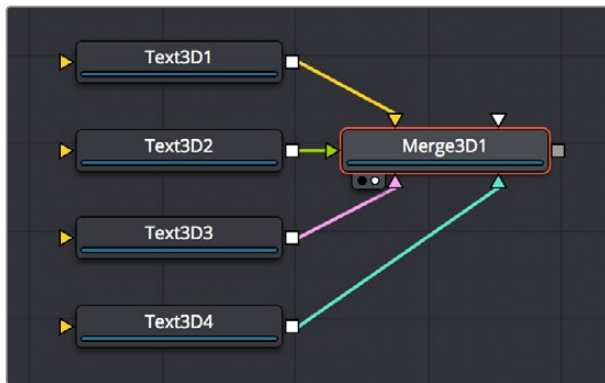
You can export a 3D scene from Fusion to other 3D packages using the FBX Exporter node. On render, it saves geometry, cameras lights, and animation into different file formats such as .dae or .fbx. The animation data can be included in one file, or it can be baked into sequential frames. Textures and 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 detail, see Chapter 26, "3D Nodes" in Part 7, "Node Reference."

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.

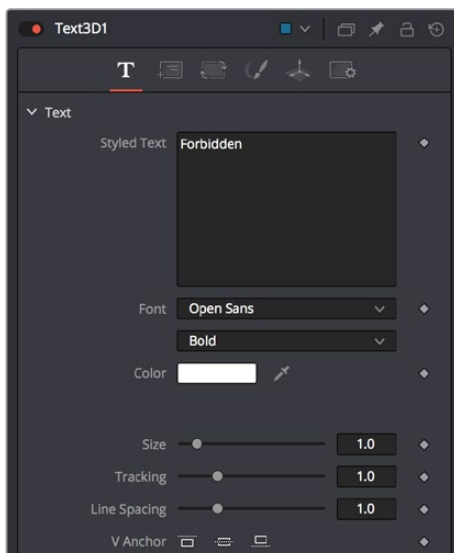


Merging multiple text objects to create an intricately styled scene.

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.

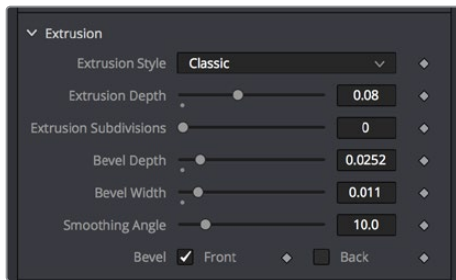
Entering Text

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 on-screen. 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.

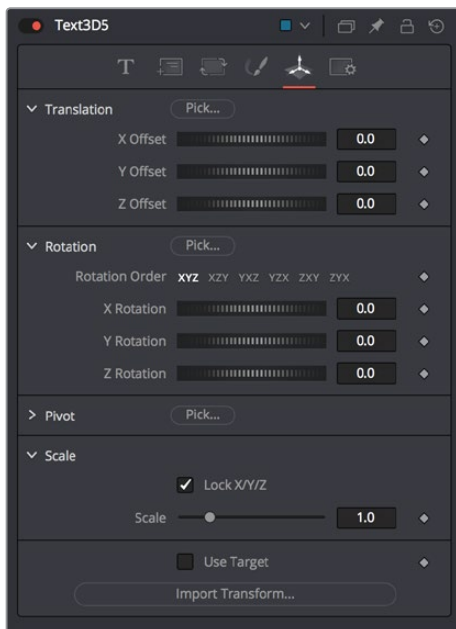
By default all text created with the Text3D node is flat, but you can use the Extrusion Style, Extrusion Depth, and various Bevel parameters to give your text objects thickness.



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 on-screen 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

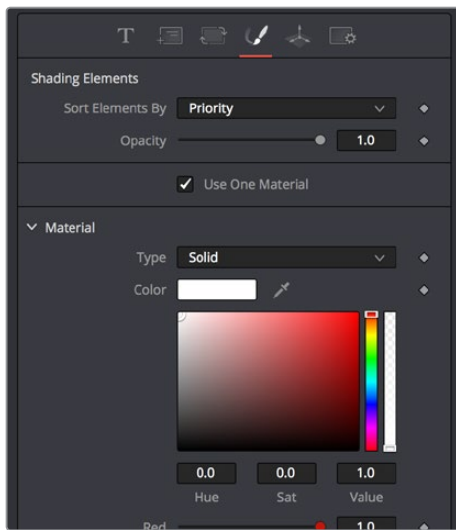
Another Transform tab (which the documentation has dubbed the “Sub” Transform tab) lets you apply a separate level of transform to either characters, words, or lines of text, which lets you create even more layout variations. For example, choosing to Transform by Words lets you change the spacing between 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 Softclipping

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 IsMatte option, which when turned on enables two more options.



Matte parameters in the Shape3D node; enabling IsMatte reveals additional options.

IsMatte

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 IsMatte 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 IsMatte 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 Render3D, 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”

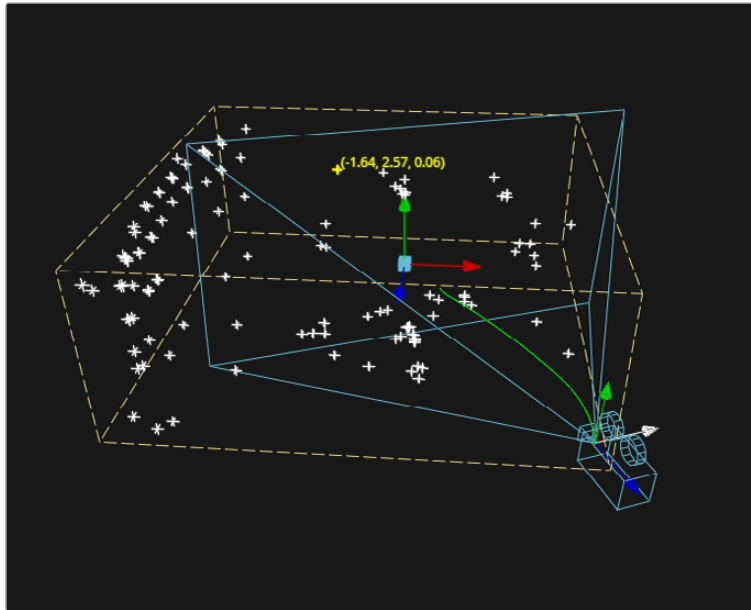
Empty regions of the render will have the Position channel incorrectly initialized to (0,0,0). To get the correct Position data, add a bounding sphere or box to your scene to create distant values and allow the Position nodes to render correctly.



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.

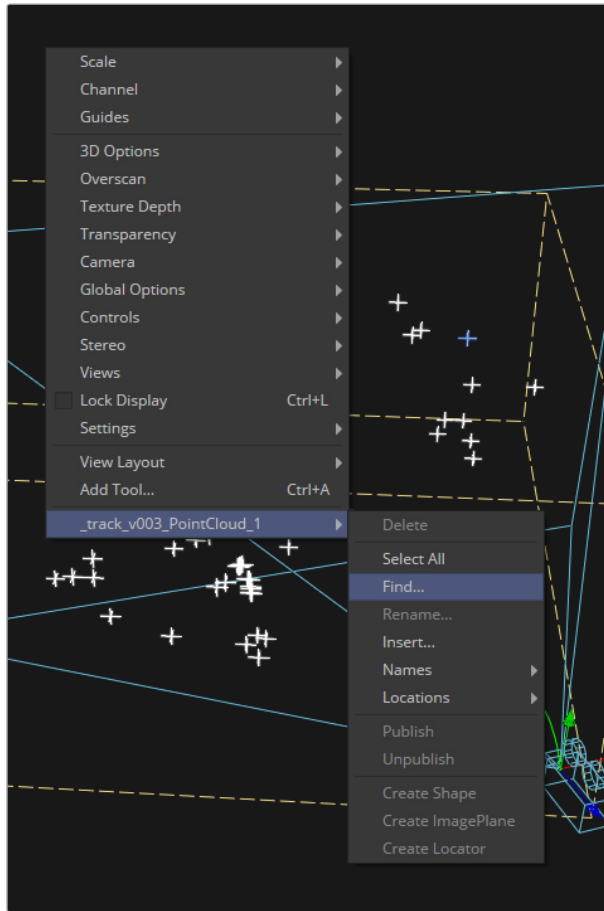
The entire point cloud is imported as one object, which is a significantly faster approach.

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.

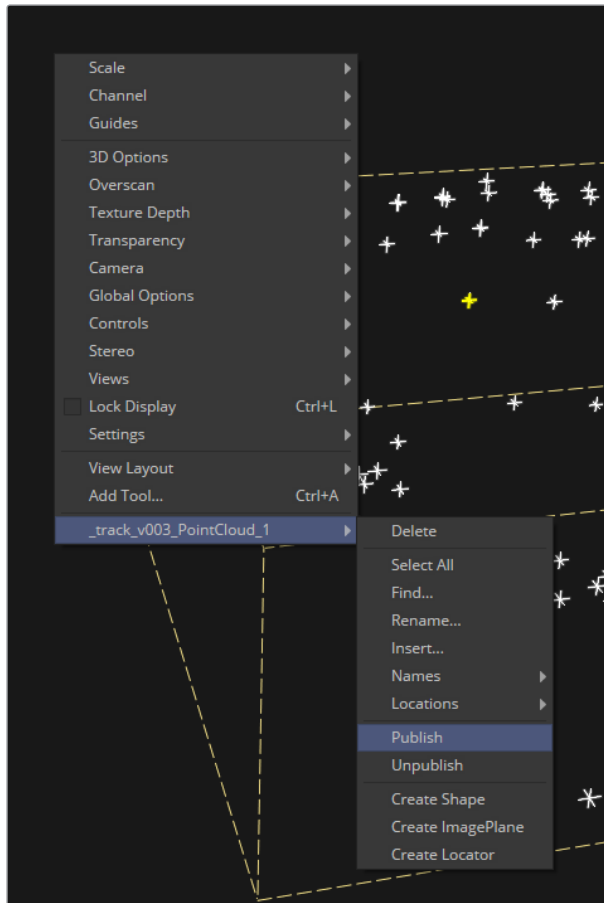
TIP: The Point Cloud Find function is a case-sensitive search. A point named tracker15 will not be found if the search is Tracker15.

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 24

3D Camera Tracking

This chapter presents an overview of using the Camera Tracker node and the workflow it involves.

Contents

Introduction to 3D Camera Tracking	540
The 3D Camera Tracking Workflow	540
Tracking	541
Camera	542
Solving	542
Cleaning Up Camera Solves	544
Export	545
Outputting from the Camera Tracker	546
2D View	546
3D View	547

Introduction to 3D Camera Tracking


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. The Camera Tracker will also create a point cloud in 3D space, which can be used to align objects and other 3D models to the live-action scene.



An example of 3D elements integrated in a live-action scene.

The 3D Camera Tracking Workflow

The Camera Tracker is a complete workflow in one tool. By tracking the frames from a camera moving in a scene, the 3D environment of the location and the camera's motion can be reconstructed. To do this, the track needs to be of a scene that has tracking features in a fixed frame of reference. Moving objects or people must be masked out from the tracker to get a good solve, as these false tracks will cause inaccuracies when solving for the virtual camera. When solving for the virtual camera, it is helpful to provide certain additional pieces of information, such as the camera sensor size, the focal length of the lens, and measured distances of known tracking marks. This will help guide the solver toward generating a more accurate 3D camera and point cloud.



Track Data and a solved camera path.

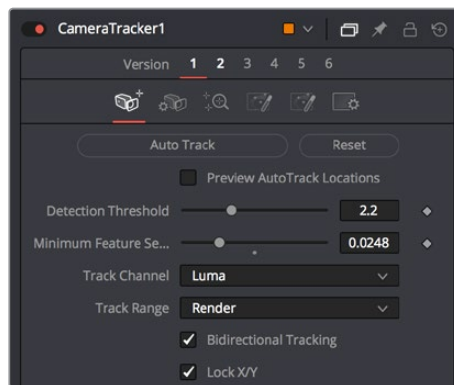
The Camera Tracker's task is to create a 3D animated camera and point cloud of the scene. To generate a camera, the basic approach to camera tracking has two steps. Tracking, which is the analysis of a scene, and solving, which calculates the virtual camera and point cloud. Once these steps are taken, an animated camera can be exported from the tool.

There are five tabs in the Camera Tracker node: Track, Camera, Solve, Export, and Options. To define the workflow where an image sequence is tracked, set basic Camera parameters, Solve, and then Export the 3D virtual camera and point cloud. The top buttons of each tab section are the operational process that will trigger the actions for each process.

Tracking

Tracking is the term used to describe the task of observing or analyzing the shot. 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. To do this, the tool searches for features, which are high-contrast patterns within the shot, and assigns trackers automatically to those features. A wide spread of tracking points across the scene will result in the best possible track, and adjusting controls such as the Minimum Feature Separation will help populate the scene with more points. It is recommended when tracking a scene to avoid features on moving objects or features that are caused by parallax differences at different depths.

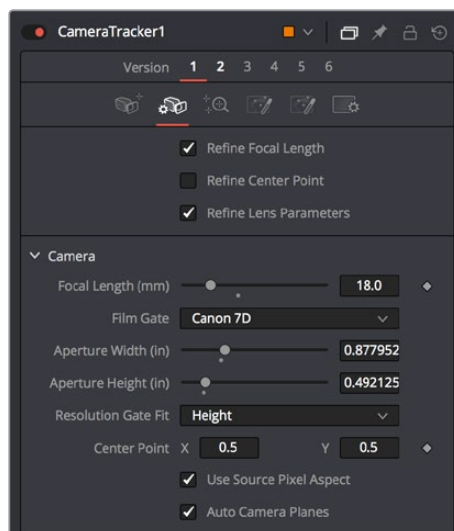
One way of avoiding these problem areas is masking. Applying a mask to the Camera Tracker will isolate areas of a scene that the Camera Tracker can analyze. For example, suppose some footage shot with a moving camera of a subject standing in front of a greenscreen needs camera tracking. Apply a mask to the subject, invert that mask (so the subject is black and the environment is white), and attach it to the Camera Tracker Track Mask input. By doing this, the tool will look for features within the mask area and not on the subject. Alternatively, tracks can be selected and deleted from the viewer.



The Camera Tracker tab.

Camera

This section is where the basic settings of the camera are set, such as film gate size, focal length, and lens distortion parameters. If the actual values are not known, try a best guess. The solver will attempt 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 3D track will be. At a minimum, try to enter the correct focal length of the lens.

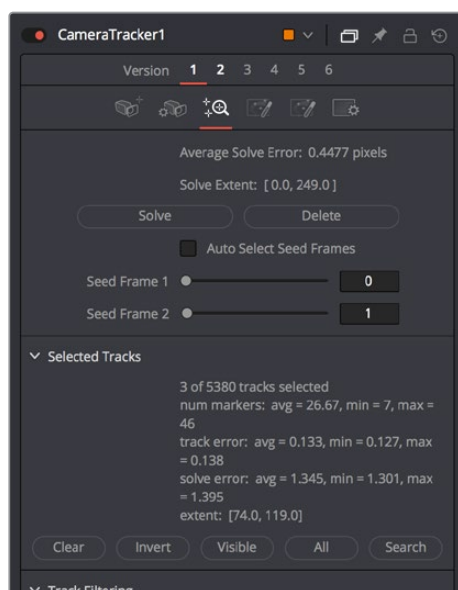


The Camera tab.

Solving

Solving is the process in which the tool takes the currently existing tracks and calculates from them the movement of the live-action camera. From a user's point of view, solving should be thought of as an iterative process:

- 1 Run the solver.
- 2 Delete poor tracks and/or adjust the initial guess provided in the Camera tab.



The Camera Solver tab.

At the end of the solving process, a solve error (sometimes called a reproduction error) is displayed. This is the key value that tells you how well the calculation has gone. A good solve error for HD content is below 1.0. This can be interpreted, since at any given time the track could be off by 1 pixel. The higher the resolution, the lower the solve error should be. If you are working with 4K material, a goal should be 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 greatly 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.

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 on one side of the image.
- Track starts and ends staggered in time with not too many tracks ending on the same frame.

Sometimes There's Nothing You Can Do

Some shots, if they 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 the set with proper usage of track markers and ensuring that camera moves have enough perspective for the solver to “latch” onto.

Cleaning Up Camera Solves

Sometimes the first solve will be good enough. Other times, it may take many hours of cleaning up tracks to get a good solve, and other times 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 solve quality to decrease, as the solver has too little information to work with. In particular, if there are less than eight tracks on any frame, mathematically there is not enough information to solve the shot. However, it is strongly recommended to use a lot more than eight tracks to get a robust and accurate solve.

IMPORTANT It may be tempting for users unfamiliar with working with camera trackers to try to directly edit the 3D splines of the resulting camera in order to improve a solved camera's motion path. This option should be used as an absolute last resort, preferring instead to modify the 2D tracks being fed into the solver.

How to Judge Track Accuracy

Cleaning up false tracks is facilitated by showing the tracks overlaid on top of the live-action footage you're analyzing. Under the Options tab, the length of these tracks can be lengthened to show longer motion paths. Tracks can be selected either individually or in multiples via a bounding box, and the Delete button will remove tracks you decide are erroneous.

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 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.

Tips for Deleting Tracks

When deleting tracks, it's good practice to note the current "average solve error," then rerun the solver, and note whether the changes increased or decreased the average solve error. In addition to looking at the average solve error to judge the quality of a camera solve, it is also useful to:

- Look at the camera path in 3D views to see if it contains any unexpected jumps, breaks, or jitter.
- Look through the camera in a 3D view at the locators in the 3D point cloud to see how well they stick to features in the live-action footage.

Using Filters to Delete Problem Tracks

The Solve tab includes filters that can be used to delete tracks based on track length, track error, and solve error. These can be used to quickly remove poorly performing tracks that may be misleading to the resulting camera, leaving a concise list of accurate tracks. For more detail on the Filtering options, see Chapter 52, "Tracker Nodes."

Export

Before exporting, it is a good idea to line up the virtual ground plane in Fusion's 3D environment with the physical ground plane in the live-action footage. Camera Tracker provides various translation, rotation, and scale controls to accomplish this. By selecting tracking points on-screen, the ground plane can be aligned to these points, as well as rotation. By using two points of known distance, the scale of the scene can be set.

The Export will maintain a link to the exported tools, so adjustments and new solves will auto update these tools.



The Camera Tracker tab.

NOTE: Camera Tracker saves all its 2D tracks into the composition, sometimes resulting in a rather large file on disk. If there are too many 2D tracks over a long shot, the saved composition can reach over a gigabyte in size. In some cases, this can make compositions containing Camera Tracker nodes cumbersome to load and work with. While it is possible to work directly with the Camera Tracker tool via the camera coming out of the 3D output, once the quality of the solve is satisfactory, consider instead using the Export functionality to take a "low memory" snapshot that can be cut and pasted into another composition.


Outputting from the Camera Tracker

The Camera Tracker node has two outputs: The primary output is 2D, and there is also a 3D output for viewing the camera path and point cloud in 3D space. When refining tracks to increase the accuracy of the solve, it can be helpful to simultaneously view the 2D and 3D outputs in side-by-side views. Note that 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 functionality of this tool.

2D View

The 2D view is the primary display for the node. By dragging and dropping the node to the view, it will display the image being tracked as well as overlay tracker markers and plotted paths of the tracker motion.

A dedicated toolbar gives you access to the common features used to track and solve a shot.

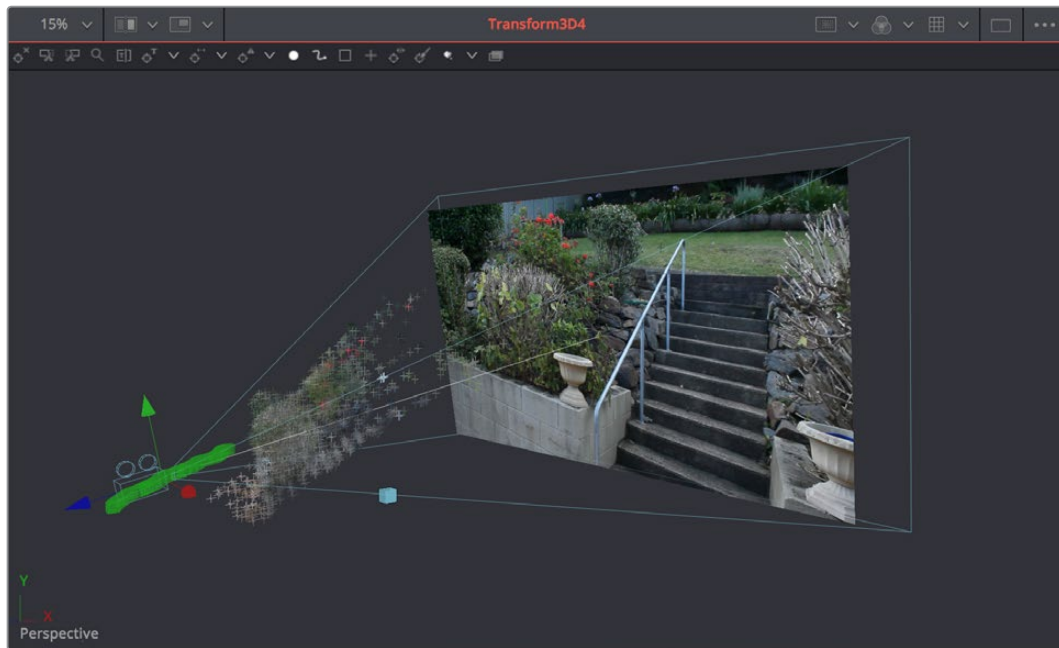


A 2D view of the camera track data.

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 node and view that tool. The 3D output will display the point cloud and the camera along with the image connected to it.

Selecting points will invoke the on-screen menu that will give control of various functions, such as displaying Frame Ranges, Solve Error, and Name, as well as Renaming, Deleting, and changing the colors.



A 3D view of the camera track data.

Chapter 25

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 detail on each Particle System node that’s available, see Chapter 49, “Particle Nodes.”

Contents

Introduction to Particle Systems	550
Anatomy of a Simple Particle System	551
Particle System Distribution	553
Particle Nodes Explained by Type	555
Emitters	555
Forces	556
Compositing	556
Rendering	556
Example Particle Systems	557


Introduction to Particle Systems

Particle systems are computer simulations that use customizable rules to automatically generate and animate 3D geometry that's designed to act like smoke, dust, fire, leaves, sparks, or any other animated system of shapes. As Fusion is a full-featured 3D compositing environment, particle systems are fully 3D systems, which makes them incredibly flexible and capable of producing all kinds of visual effects or abstract animated content for use in motion graphics.



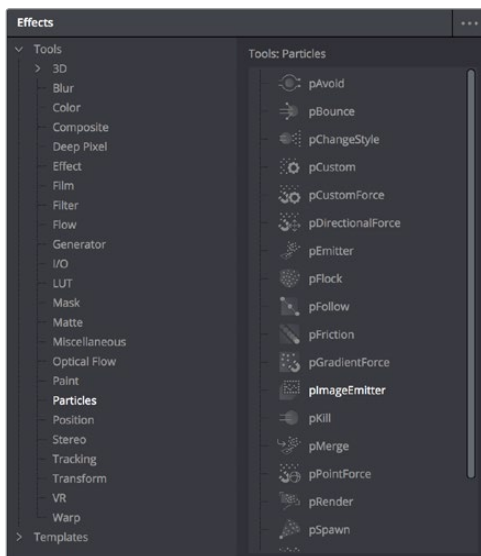
A 3D particle system, also created entirely within Fusion.

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 three Particle System nodes available from the toolbar.

However, these three nodes are only the tip of the iceberg. Opening the Particle bin of 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 all the 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.



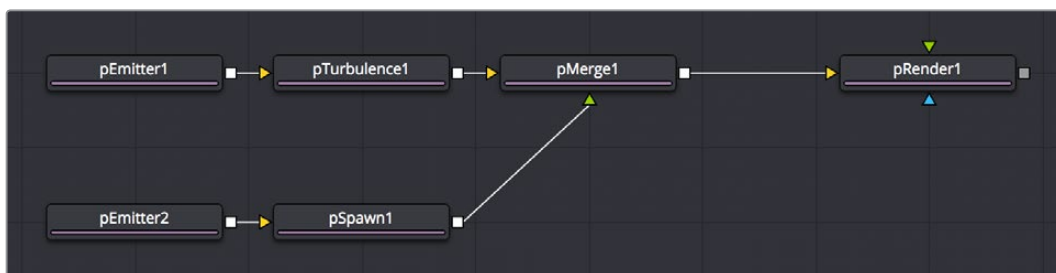
The minimum node tree required to create a simple particle system.

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.

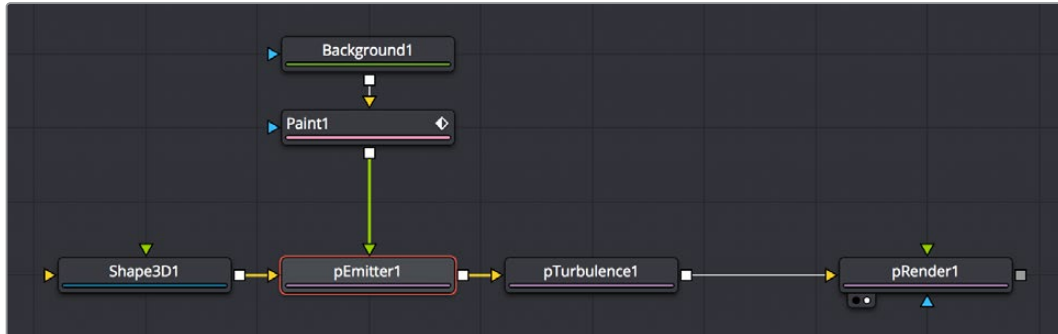


Customizing the effect of pEmitter nodes using different forces to add complexity to the particle animation.

You can also attach the following types of nodes to a pEmitter node to deeply customize a particle system:

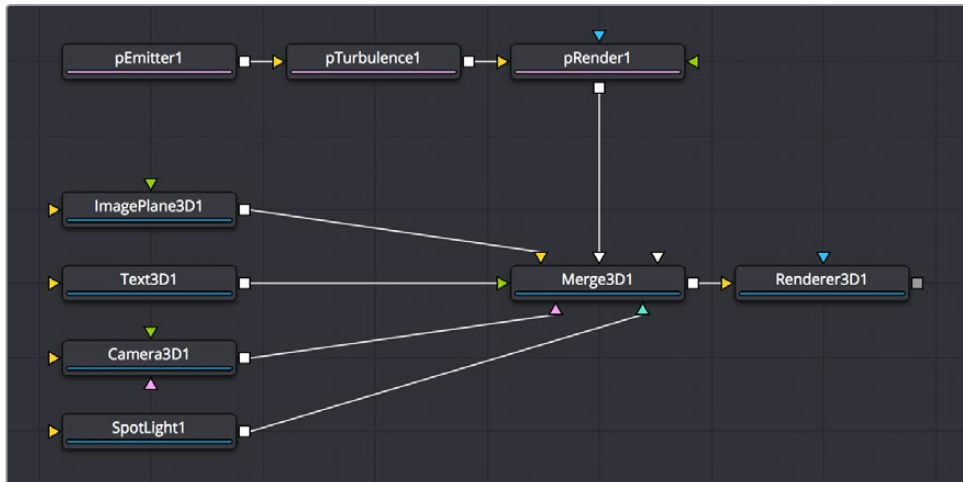
Attach a 2D image to a pEmitter node to create highly customized particle shapes. Make sure your image has an appropriate alpha channel.

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).



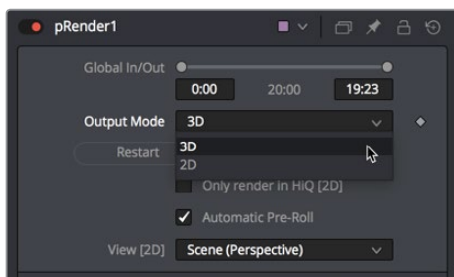
Customizing pEmitter nodes using mesh geometry to define regions and 2D images to define particle shape.

The above examples assume you'll be outputting 2D renders to combine into the rest of a 2D composition. However, because particle systems in Fusion 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.

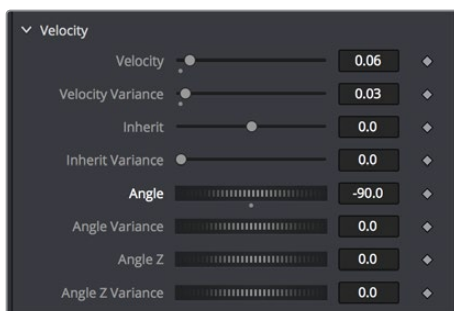
If you connect a pRender node to a Merge3D node, the Output Mode is locked to 3D, meaning that 3D geometry is output by the pRender node for use within the Merge3D node's scene. This means that the particles can be lit, they can cast shadows, and they can interact with 3D objects within that scene.



The result of using a particle system within a 3D scene.


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.



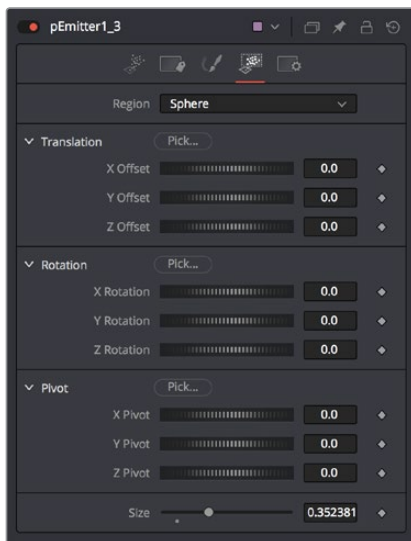
A pEmitter node's Velocity Angle and Angle Variance controls let you adjust the direction and width of particle distribution.

Particle systems can be positioned and rotated by loading the pEmitter nodes that generate particles into a viewer and using the on-screen 3D position and Rotation controls provided to move the particle system around.



A pEmitter node loaded into the viewer with the rotation on-screen 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.

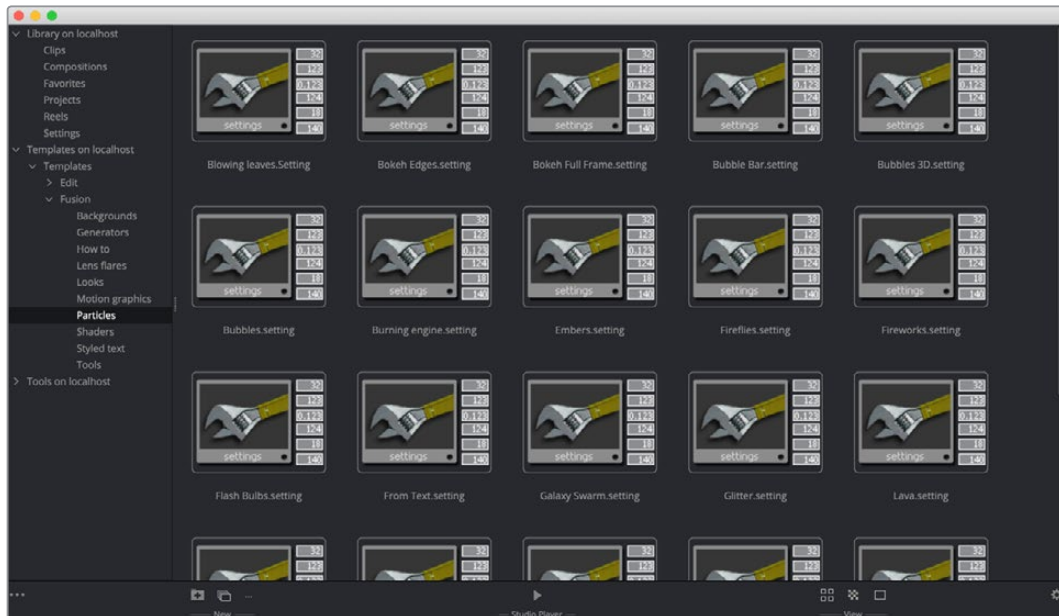
Rendering

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.

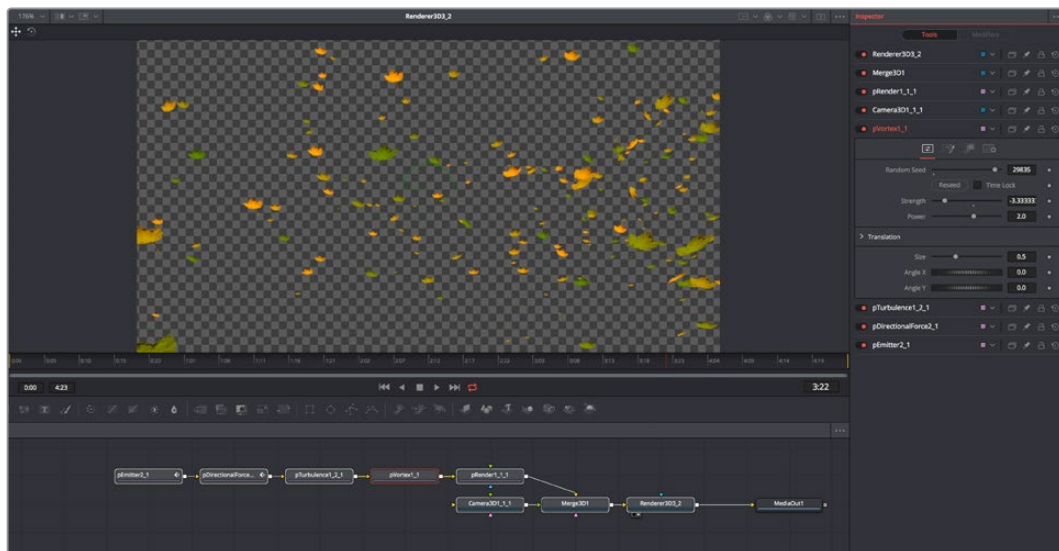
Example Particle Systems

The Templates category of the Bins window 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.

Simply drag and drop any of the particle presets from the Bins window 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 of the Bins window.



PART 7

Fusion Effects

Chapter 26


3D Nodes

This chapter covers, in great detail, the nodes used for creating 3D composites.

Contents

Alembic Mesh 3D [ABC]	561
Bender 3D [3BN]	564
Camera 3D [3CM]	565
Cube 3D [3CB]	573
Custom Vertex 3D [3CV]	577
Displace 3D [3DI]	578
Duplicate 3D [3DP]	580
FBX Exporter 3D [FBX]	583
FBX Mesh 3D [FBX]	585
Fog 3D [3FO]	590
Image Plane 3D [3IM]	592
Locator 3D [3LO]	596
Merge 3D [3MG]	598
Override 3D [3OV]	599
Point Cloud 3D [3PC]	600
Projector 3D [3PJ]	604
Renderer 3D [3RN]	609
Replace Material 3D [3RPL]	618
Replace Normals 3D [3RPN]	619
Replicate 3D [3REP]	621
Ribbon 3D [3RI]	626
Shape 3D [3SH]	627
Softclip [3SC]	632
Text 3D [3TXT]	633
Transform 3D [3XF]	637
Triangulate 3D [3TRI]	639
UV Map 3D [3UV]	640
Weld 3D [3WE]	643
Modifier	644

Alembic Mesh 3D [ABC]



There are two way to import Alembic files:

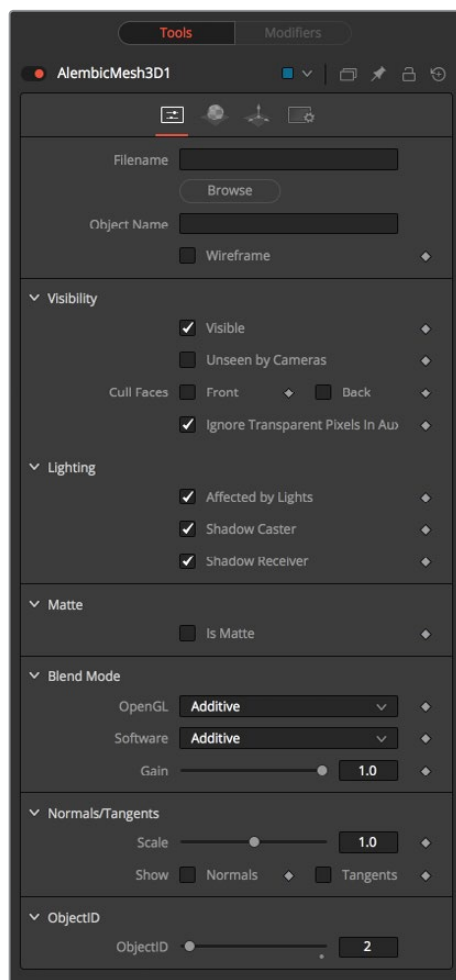
- Using the File > Import > Alembic menu option
- Manually adding an AlembicMesh3D node to the node tree

The first method is strongly recommended.

The Alembic format allows for arbitrary user data to be stored within the file. Fusion ignores most of this metadata for various reasons. No conventions have been defined yet for how this metadata is named and metadata might change between different ABC exporters. When the Alembic file is imported through the menu option, the transforms are read into splines and into the inputs on the nodes, which get saved with the comp.

This means that when re-loading the comp, the transforms are loaded from the comp and not the Alembic file. The meshes are handled differently; they are always reloaded from the Alembic file.

Controls



Filename

The name of the imported Alembic file.

Object Name

This input shows the name of the mesh from the Alembic file that is being imported.

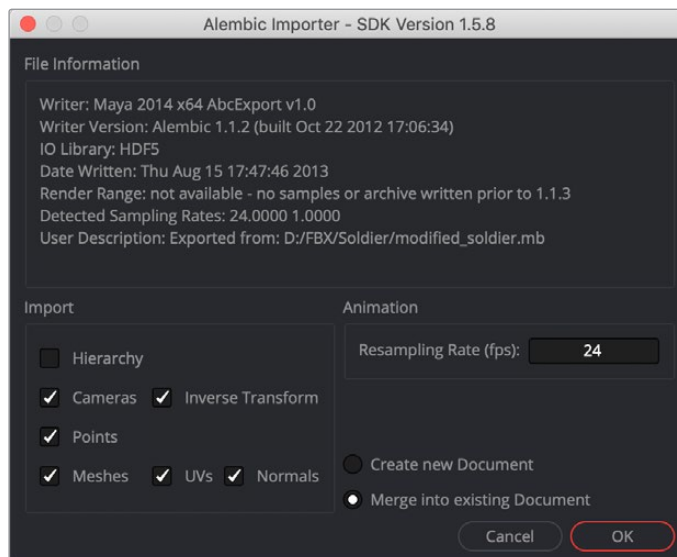
If this field is blank, the entire contents of the FBX geometry will be imported as a single mesh. This input is not editable by the user; it is set by Fusion when importing Alembic files via the File > Import > Alembic utility.

Sampling Rate

The framerate is set when importing the file. It can be altered using this slider to create effects like slow motion.

Dump File

Opens the resulting ASCII in the preferred text editor.



Alembic Import Dialog

- **Writer:** The name of the plug-in/application that created/wrote out the Alembic file
- **Writer Version:** The version of the Alembic sdk that was used to write out the Alembic file
- **RenderRange:** This gives you an idea of the duration of the animation in the Alembic file in seconds
- **DetectedSamplingRates:** Fusion examines the framerates in the file and reports them here. This is useful to determine the value at which to set the resampling rate.

Hierarchy

If disabled, the transforms in the Alembic file are flattened down into the cameras and meshes. This results in a number of meshes/cameras connected to a single merge node in Fusion. When enabled, you get the full parenting hierarchy.

Orphaned Transforms

Transforms that do not parent a mesh or camera will not be imported if this option is unchecked. For example, if you have a skeleton and associated mesh model, the model will be imported as an Alembic mesh and the skeleton will be imported as a tree of Merge3Ds. Disabling this option causes the Merge3Ds not to be imported.

Cameras

Near/Far/Apertures/Angles of View/Plane of Focus are imported. The resolution Gate Fit may be imported; it depends if the Writer correctly tagged the resolution Gate Fit metadata. If your camera does not import correctly, you should check to see if Camera3D.ResolutionGateFit is set correctly. Stereo information is not imported.

InverseTransform

Imports the Inverse Transform (World to Model) for cameras.

Points

Alembic supports a Points type. This is a collection of 3D points without orientation. Some 3D apps export particles as points, but keep in mind the direction and orientation of the particles are lost; you just get positions. It very well may be possible that the exocortex Alembic plug-ins write out extra user data that contains orientation.

Meshes

Optionally import UVs and normals.

ResamplingRate

When an Alembic animation is exported, it is stored on disk in seconds, not in frames. When the Alembic data is brought into Fusion, you need to supply a framerate with which to resample the animation. Ideally, you should choose the same framerate that it was exported with so your samples match up with the original samples. Detected Sampling Rates can give an idea of what to pick if unsure.

Lights

Import currently not supported. There is no universal convention on Alembic light schemas.

Materials

Import currently not supported. There is no universal convention on Alembic material schemas.

Curves

Import currently not supported.

Multiple UVs

Import currently not supported. There is no universal convention yet.


Velocities

Import currently not supported.

Cyclic/Acyclic Sampling

Currently not implemented. Uniform sampling, which is the most common mode, works fine. We recommend the use of FBX for lights/cameras/materials and Alembic for meshes only. If cameras and Alembic work for you, then go for it. The reason is that our Alembic plug-in doesn't support lights/materials, but FBX has good support. Alembic import of cameras has problems with Resolution Gate Fit and doesn't import stereo options.

Bender 3D [3BN]



The Bender 3D node is used to bend, taper, twist, or shear the geometry in a 3D scene based upon its bounding box. It takes a 3D scene as input and outputs a modified 3D scene. 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; only existing vertices in the geometry are altered. As a result, when applying the Bender 3D node to primitives created in Fusion, it is a good idea to increase the value of the Subdivision controls in the original primitives to give a higher quality result.

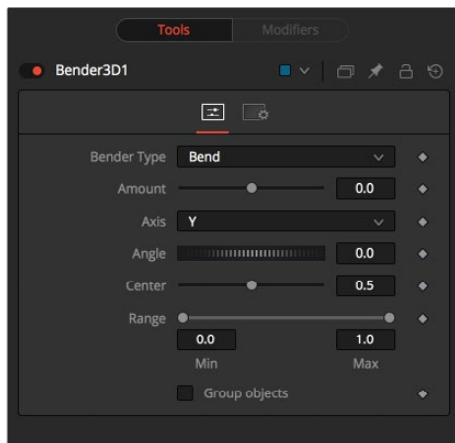
External Inputs

The following inputs appear on the node in the Node Editor.

Bender3D.SceneInput

[gold, required] This input expects a 3D scene.

Controls



Bender Type

Use the Bender Type to select the type of deformation to apply to the geometry. There are four modes available: Bend, Taper, Twist, and Shear.



Amount

Adjust the Amount slider to change the strength of the deformation.

Axis

The Axis control determines the axis along which the deformation is applied and has a different meaning depending on the type of deformation. When bending, this determines the axis that is bent, 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 that 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 visible when the Bender Type is set to Shear.

Group Objects

When this is checked, all the objects in the input scene are grouped together into a single object and that object is deformed around the common center, instead of deforming each component object individually.

Camera 3D [3CM]



The Camera 3D node generates a virtual camera through which the 3D environment can be viewed. It closely emulates the settings used in both real and virtual cameras in an effort to make matching the cameras used in other scene elements as seamless as possible.

The camera should be added to the scene using a Merge 3D node. 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 scene from the Merge 3D node that introduces the camera, or any node downstream of that Merge 3D. Then right click in the view and select Camera > Cameraname from the contextual menu. Right clicking on the axis label found in the bottom corner will display the Camera sub-menu directly.

The aspect of the Viewer may be different from the aspect of the camera, so that the view through the camera interactively may not match the true boundaries of the image which will actually be rendered by the Renderer 3D node. To assist you in framing the shot, guides can be enabled that represent the portion of the view the camera actually sees. Right click in the Viewer and select an option from the Guides > Frame Format sub-menu. The default option will use 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 view is active.

The Camera 3D node can also be used to perform Camera Projection, where a 2D image is projected through the camera into 3D space. This 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 exactly with the camera. The Image Plane, Projection and Materials tabs will not appear until a 2D image is connected to the Camera 3D node in the Node Editor.

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 input. Additionally the plane of focus control for depth of field rendering is also available here.

If you add a camera by means of dragging the 3Cm icon from the toolbar onto the 3D view, it will automatically merge it with the scene you are viewing. In addition, it will be automatically set to the current viewpoint, and the view will be set to look through the new camera.

Alternatively, it is possible to copy the current viewpoint to a camera (or spotlight or any other object) by means of the Copy PoV To option in the Viewer's contextual menu, under the Camera submenu.

External Inputs

The following inputs appear on the node in the Node Editor.

Camera3D.SceneInput

[gold, required] This input expects a 3D scene.

Camera3D.RightStereoCamera


[green, optional] This input should be connected to another Camera 3D node. It is used to override the internal camera used for the right eye in stereoscopic renders and viewers.

Camera3D.ImageInput

[magenta, optional] This input expects a 2D image. The image is used as a texture when camera projection is enabled, as well as when the camera's image plane controls are used to produce parented planar geometry linked to the camera's field of view.

Controls

The options in this tab are used to set the camera's clipping, field of view, focal length and stereoscopic properties.



Projection Type

Use the Projection Type button to choose 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 only present controls for the near and far clipping planes, and a control to set the viewing volume.

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 which 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 object more than 20 units distant from the camera will be invisible to the camera. A near clipping plane of 0.1 means that any object closer than 0.1 units will also be 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.

Adaptively Adjust Near/Far Clip

When selected, the renderer will automatically adjust 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 only appears 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.

Angle of View Type

Use the Angle of View Type button array 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 will cause 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 much 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.

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 $\text{angle} = 2 * \arctan[\text{aperture} / 2 / \text{focal_length}]$.

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)


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

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 will be 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.

Film Back

Film Gate

The Film Gate menu shows a list of preset camera types. Selecting one of the options will automatically set the aperture width and aperture height to match the selected camera type.

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. Aperture is generally measured in inches, which are the units used for this control.

Resolution Gate Fit

Determines how the film gate is fit 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. This setting corresponds to the Maya Fit Resolution Gate. The modes Overscan, Horizontal, Vertical and Fill correspond to Inside, Width, Height, and Outside.

- **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 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 will be scaled uniformly until its width (X) fits the width of the 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 will be scaled uniformly until its height (Y) fits the height of the 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 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 image source and mask, either the image source's width or height may be cropped or not fit the respective dimension 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.

Control Visibility

Allows you to selectively activate the onscreen controls that are displayed along with the camera.

- **Frustrum:** 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 plane of focus according to the respective 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.

Import Camera

The Import Camera button displays a dialog to import a camera from another application.

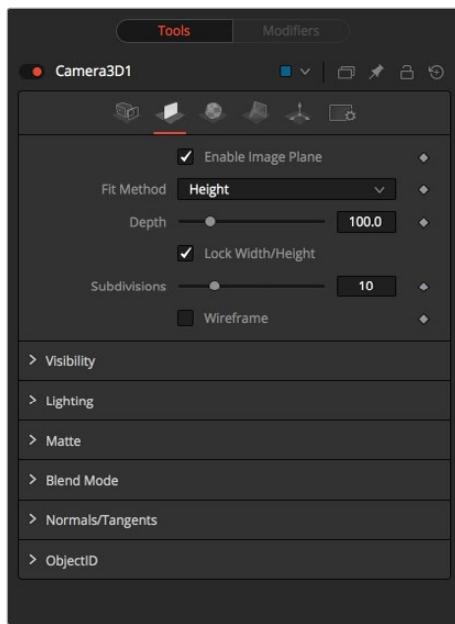
It supports the following file types:

*LightWave Scene	.lws
*Max Scene	.ase
*Maya Ascii Scene	.ma
*dotXSI	.xsi

Image

When a 2D image is connected to the camera, an Image Plane is created that is always oriented so that the image fills the camera's field of view. The Image Plane tab is hidden until a 2D image is connected to the Camera 3D's input on the node tree.

With the exception of the controls listed below, the options presented in this tab are identical to those presented in the Image Plane node's control tab. Consult that node documentation for a detailed description.



Enable Image Plane

Use this checkbox to enable or disable the creation of the Image Plane.

Fill Method

Describes how to deal with the input image if the camera has a different aspect ratio.

- **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 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 will be scaled uniformly until its width (X) fits the width of the 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 will be scaled uniformly until its height (Y) fits the height of the 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 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 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.

Projection

If a 2D image is connected to the camera it becomes possible to project the image into the scene. A projection is different from an image plane in that the projection will fall onto the geometry in the scene exactly as if there was a physical projector present in the scene. The image is projected as light, which means the renderer must be set to enable lighting for the projection to be visible. See the Projector 3D node for additional information.



Enable Camera Projection

Select this checkbox to enable projection of the 2D image connected to the Camera node.

Projection Fit Method

This button array can be used to select the method used to match the aspect of 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 which can be relighted using other lights. Needs a Catcher node connected to the appropriate input ports of the specific material.

Tips for Camera 3D

When importing a camera from a 3D application that will also be used as a projector, make sure that the Fit Resolution Gate options on the main Controls tab as well as the Projection tab are in sync. Only the first one will automatically be set to what the 3D app was using. The latter might have to be adjusted manually.

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 onto. You need to use a Replace Material node after your Camera node.

To achieve real Parallel Stereo mode you can:

- Connect an additional external (right) camera to "Right Stereo Camera" input of your camera.
- Create separate left and right cameras
- Set the ConvergenceDistance slider to a very large value of 999999999.

Rendering with Overscan from Fusion's 3D Space

If you want to render an image with overscan you also have to 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 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.

External Inputs


Cube3D.SceneInput

[orange, optional] This input expects a scene from a 3D node output.

Cube3D.NameMaterialInput

These 6 inputs are used to define the materials applied to the six faces of the cube. They will accept either a 2D image or a 3D material as valid.

Controls



Lock Width/Height/Depth

This checkbox locks the Width, Height, and Depth dimensions of the cube together, so that they are always the same size. 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 will also be 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.

If the Open GL viewer and renderer are set to Vertex lighting, the more subdivisions in the mesh, the more vertices will be available to represent the lighting. For this reason, high 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 will cause the Mesh to render only the Wireframe for the object when rendering with the OpenGL renderer.

Visibility

- **Visible:** If the Visibility checkbox is not selected, the object will not be visible in the Viewers, nor will it be rendered into the output image by the Renderer 3D node. A non-visible object does not cast shadows.
- **Unseen by Cameras:** If the unseen by cameras checkbox is selected, the object will be visible in the Viewers (unless the Visible checkbox is turned off), except when viewed through a camera. The object will not be rendered into the output image by the Renderer 3D node. Shadows cast by an unseen object will still be visible when rendered by the Software renderer, though not by the OpenGL renderer.
- **Cull Front Face/Back Face:** Use these options to cull (eliminate) rendering and display of certain polygons in the geometry. If Cull Back Face is selected, all polygons facing away from the camera will not be rendered, and will not cast shadows. If Cull Front Face is selected, all polygons facing toward the camera will likewise be dropped. Selecting both checkboxes has the same effect as deselecting the Visible checkbox.
- **Ignore Transparent Pixels in Aux Channels:** In previous versions of Fusion, transparent pixels were rejected by the software/GL renderers. To be more specific, the software renderer rejected pixels with R=G=B=A=0 and the GL renderer rejected pixels with A=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 in post 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, then it would be useful to have transparent areas set aux channels (in particular UVs). As another example, suppose you are doing post DoF. You will probably not want the Z-channel to be set on transparent areas, as this will give you a false depth. Also, keep in mind that this rejection 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 will not affect it.

Lighting

- **Affected by Lights:** If this checkbox is not selected, lights in the scene will not affect the object, it will not receive nor cast shadows, and it will be shown at the full brightness of its color, texture or material.
- **Shadow Caster:** If this checkbox is not enabled, the object will not cast shadows on other objects in the scene.
- **Shadow Receiver:** If this checkbox is not enabled, the object will not receive shadows cast by other objects in the scene.

Matte

Enabling the Is Matte option will apply 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 will override all textures. See the Matte Objects section of Chapter 60, “3D Compositing Basics,” for more information.

- **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 only visible when the Is Matte option is enabled.
- **Infinite Z:** Sets the value in the Z-channel to infinite. This checkbox is only visible when the Is Matte option is enabled.

Blend Mode

A Blend mode specifies which method will be 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.

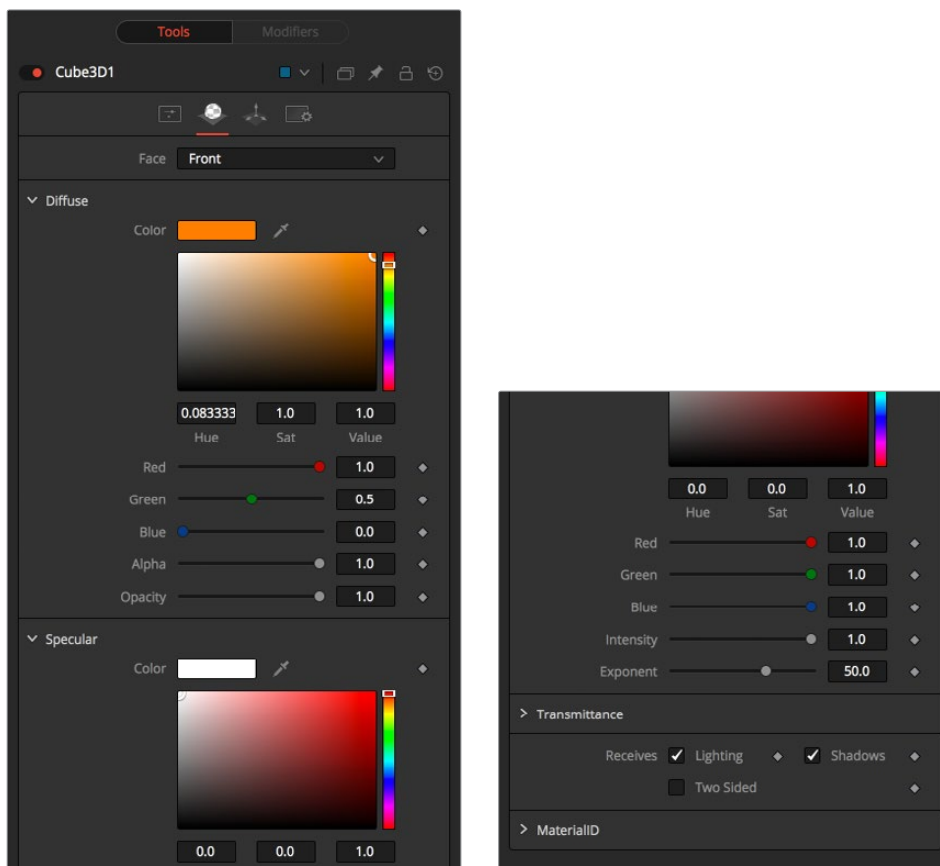
The blending modes were originally designed for use with 2D images. Using them in a lit 3D environment can produce undesirable results. For best results, use the Apply modes in unlit 3D scenes rendered in software.

- **OpenGL Blend Mode:** Use this menu to select the blending mode that will be used when the geometry is processed by the OpenGL renderer. This is also the mode used when viewing the object in the Viewers. Currently the OpenGL renderer supports three blending modes.
- **Software Blend Mode:** Use this menu to select the blending mode that will be used when the geometry is processed by the software renderer. Currently, the software renderer supports all of the modes described in the Merge node documentation, except for the Dissolve mode.

Material Tab

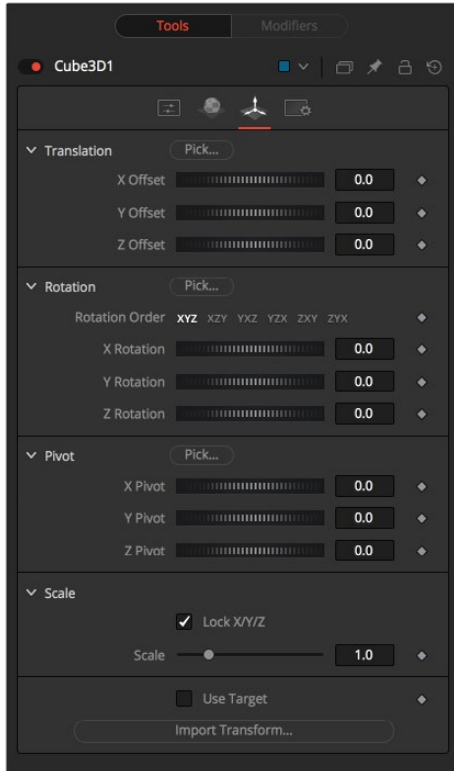
The options which appear in this tab determine the appearance of the geometry created by this node. Since these controls are identical on all nodes that generate geometry, these controls are fully described in the Common 3D Controls section of this documentation.

If an external 3D material is connected to the node's material input, then the controls in this tab will be replaced with the "Using External Material" label.




Transform Tab

The options which appear in this tab determine the position of the geometry created by this node. Since these controls are identical on all nodes that generate geometry, these controls are fully described in the Common 3D Controls section of this documentation.



Custom Vertex 3D [3CV]



This is a custom node for 3D geometry that can be used to do per vertex manipulations, for example on an image plane like: $(px, py, \sin(10 \cdot (px^2 + py^2) + n1))$. Other vertex attributes like normals, vertex color, texture coordinates, tangents, and velocity can be modified as well.

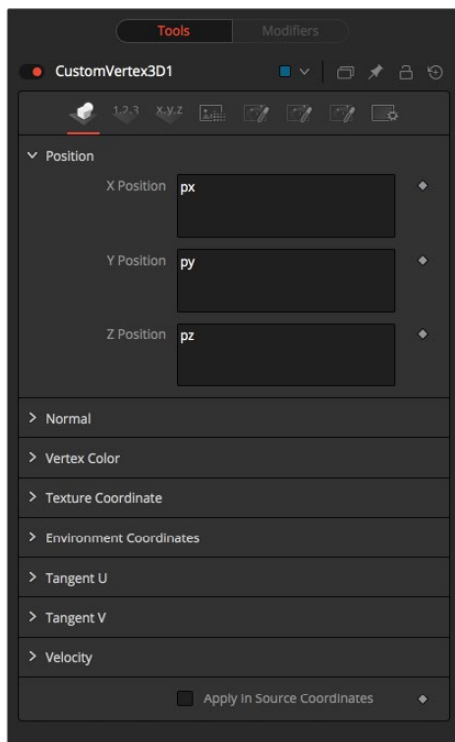
NOTE: Modifying the X, Y and Z positions of a 3D object does not modify the normals/tangents. You can use a ReplaceNormals node afterwards to recompute the normals/tangents.

TIP: Not all geometry has all vertex attributes. For example, most Fusion geometry does not have vertex colors, with the exception of particles and some imported FBX/Alembic meshes.


No geometry has envcoord currently. Only particles have velocities. If a stream is not present on the input geometry it is assumed to have a default value.

NOTE: Missing streams on the input geometry are created if the expression for a stream is non-trivial. The values for the streams will be as given in the above point. For example, if the input geometry does not have normals, then the values of (nx, ny, nz) will always be (0,0,1). To change this, you could use a ReplaceNormals node beforehand to generate them.

Controls



Displace 3D [3DI]



The Displace 3D node is used to displace the vertices of an object along their normals based upon 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 will only displace existing vertices and will not tessellate objects. 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.

Passing a particle system through a Displace 3D node will disable 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.

External Inputs


Displace3D.SceneInput

[orange, required] This input expects to receive a 3D scene.

Displace3D.Input

[green, optional] This input expects a 2D image to be used as the displacement map. If no image is provided, this node will effectively pass the scene straight through to its output.

Controls



Channel

Determines which channel of the image is connected to Displace3D. Input 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 afterwards.

Camera Displacement

- **Point to Camera:** When the Point to Camera checkbox is enabled, each vertex is displaced towards the camera rather than 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 drop-down box is used to select which camera viewer in the scene is used to determine the camera displacement when the Point to Camera option is selected.

Duplicate 3D [3DP]



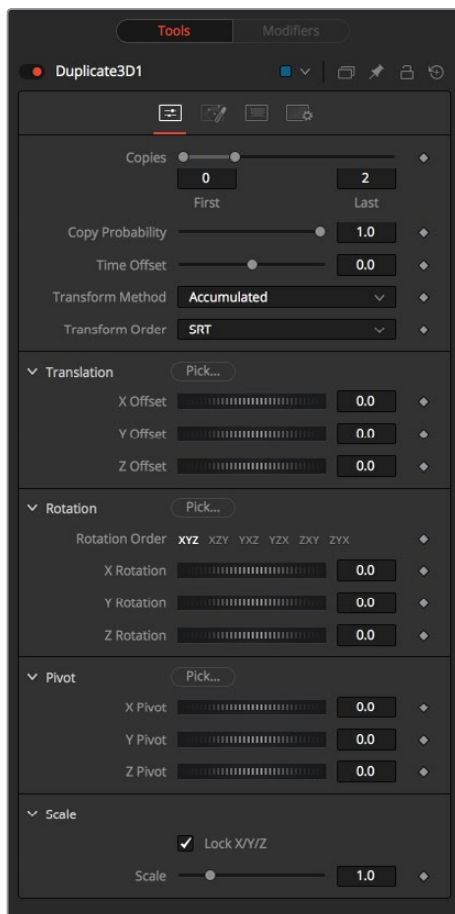
The Duplicate 3D node can be used to quickly duplicate any geometry in a scene, applying a successive transformation to each, and creating repeating patterns and complex arrays of objects. The options in the Jitter tab allow for non-uniform transformations, such as random positioning or sizes.

External Inputs

Duplicate3D.SceneInput

[orange, required] This input expects a 3D scene.

Controls



First/Last Copy

Use this range control to set how many copies of the geometry to make. 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 for some interesting effects when transformations are applied to each copy using the controls below.

Using a value for both the First Copy and the Last Copy will show only the original input. Setting the First Copy to a value greater than 0 will exclude the original input and show 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 will show the animation from a frame earlier. The second copy will show 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.

Transform Method

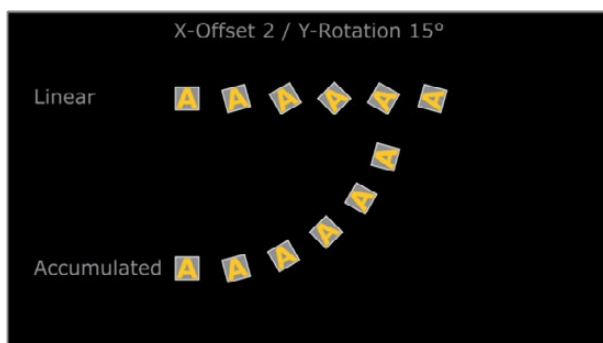
Accumulated

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.



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.



Transform Order

With these buttons, the order in which the transforms are calculated can be set. It defaults to Scale-Rotation-Transform (SRT).

Using different orders will result in different positions of your final objects.

Jitter

XYZ Offset

These three sliders tell the node how much offset to apply to each copy. An X offset of 1 would offset each copy 1 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, then the Z-axis rotation.

XYZ Rotation

These three Rotation sliders tell the node how much rotation to apply to each copy.

XYZ Pivot

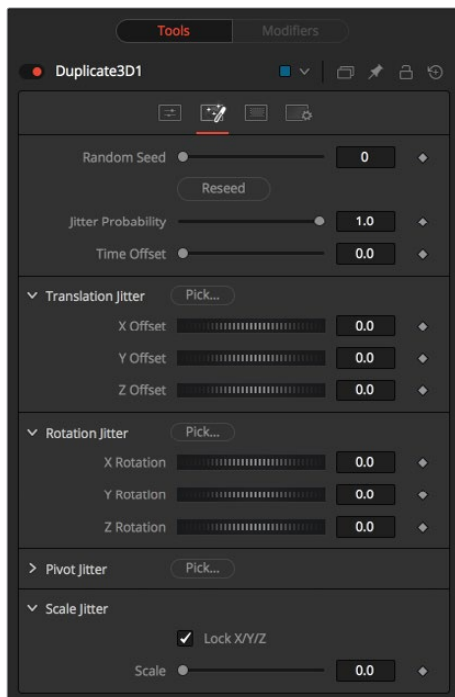
The pivot controls determine the position of the pivot point used when rotating each copy.

Lock XYZ

When the Lock XYZ checkbox is selected, any adjustment to the duplicate scale will be applied to all three axes simultaneously. If this checkbox is disabled, the scale slider will be replaced with individual sliders for the X, Y, and Z scale.

Scale

The scale controls tell Duplicate how much scaling to apply to each copy.



Random Seed/Randomize

The Random Seed is used to 'seed' the amount of jitter applied to the duplicated objects. Two Duplicate nodes with identical settings but different random seeds will produce two completely different results. Click on 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 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 will show the animation from a frame earlier. The second copy will show 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.

Translation XYZ Jitter

Use these three controls to adjust the amount of variation in the translation of the duplicated objects.

Rotation XYZ Jitter

Use these three controls to adjust the amount of variation in the rotation of the duplicated objects.

Pivot XYZ 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 XYZ Jitter

Use this control to adjust the amount of variation in the scale of the duplicated objects. Uncheck the Lock XYZ checkbox to adjust the scale variation independently on all three axes.

FBX Exporter 3D [FBX]



The FBX Exporter 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 will 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.


In addition to 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.

External Inputs

FBXExporter.Input

[orange, required] This input expects a 3D scene.

Controls



Filename

This file browser control can be used to set the file which will be output by the node. Click on the yellow folder icon to open a file browser dialog.

Format

This control is used to set the format of the output file.

Not all of the formats support all of the features of this node. For example, the obj format does not handle animation.

Version

The Version drop-down menu shows the available versions for the format selected by the control above. The menu's contents will change dynamically to reflect the available versions for that format. If the selected format only provides a single option, this menu will be hidden.

Choosing Default for the FBX formats uses FBX200611.

Geometry/Lights/Cameras

These three checkbox controls determine whether the node will attempt 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.

Reduce Constant Keys

Enabling this option will automatically remove keyframes if the adjacent keyframes have the same value.

File Per Frame (No Animation)

Enabling this option will force the node to export one file per frame, resulting in a sequence of numbered files. This will disable the export of animation.

Set Sequence Start

Normally, Fusion will use the render range of a composition to determine the numeric sequence used when rendering a file sequence to disk. Enable this checkbox to reveal the Sequence Start Frame control to set the number of the first frame in the sequence to a custom value.

Sequence Start Frame

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 will normally be numbered 0001–0030. If the Sequence Start frame is set to 100, the rendered output would be numbered from 100–131.

FBX Mesh 3D [FBX]



The FBXMesh3D node is used to import polygonal geometry from scene files that are saved using 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 of the geometry contained in the FBX file will be combined into one mesh with a single pivot and transformation. The FBXMesh node will ignore any animation applied to the geometry.

The File > Import > FBX utility can be used to import an FBX and create individual nodes for each camera, light and mesh contained in the file. This utility can also be used to preserve the animation of the objects.

If the Global > General > Auto Clip Browse option is enabled (default) then adding this node to a composition from the toolbars or menus will automatically display a file browser.

External Inputs

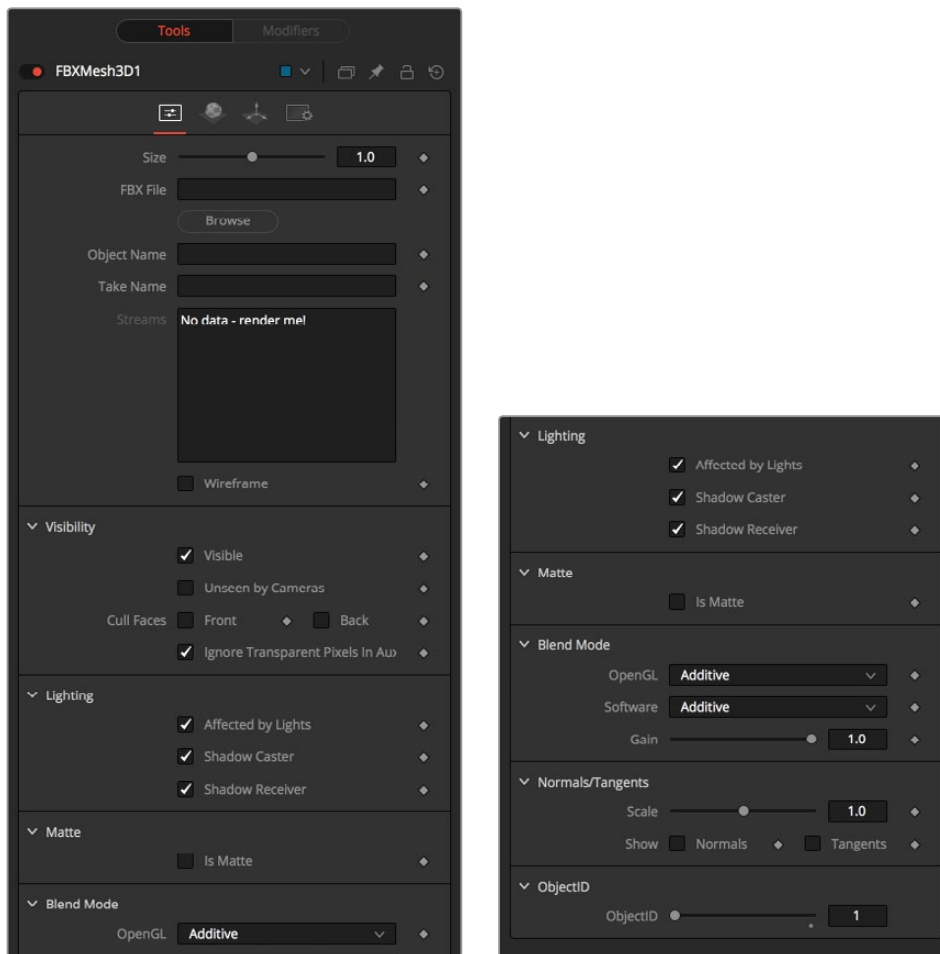
FBXMesh3D.SceneInput

[orange, required] This input expects a 3D scene as its input.

FBXMesh.MaterialInput

[green, optional] This input will accept either a 2D image or a 3D material. If a 2D image is provided, it will be used as a diffuse texture map for the basic material built into the node. If a 3D material is connected, then the basic material will be disabled.

Controls



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 control shows the filename of the currently loaded FBX. Click on the icon of the yellow folder to open a file browser that can be used to locate an 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 entire contents of the FBX geometry will be imported as a single mesh. This input is not editable by the user; it is set by Fusion when importing FBX files via the File > Import > FBX utility.

Take Name

This input shows the name of the animation take to use from the FBX file. If this field is blank, then no animations will be imported. This input is not editable by the user; it is set by Fusion when importing FBX files via the File > Import > FBX utility.

Wireframe

Enabling this checkbox will cause the mesh to render only the Wireframe for the object. Currently, only the OpenGL renderer supports wireframe rendering.

Visibility

- **Visible:** If the visibility checkbox is not selected, the object will not be visible in the Viewers, nor will it be rendered into the output image by the Renderer 3D node. A non-visible object does not cast shadows.
- **Unseen by Cameras:** If the Unseen by Cameras checkbox is selected, the object will be visible in the Viewers (unless the Visible checkbox is turned off), except when viewed through a camera. The object will not be rendered into the output image by the Renderer 3D node. Shadows cast by an Unseen object will still be visible when rendered by the software renderer, though not by the OpenGL renderer.
- **Cull Front Face/Back Face:** Use these options to cull (eliminate) rendering and display of certain polygons in the geometry. If Cull Back Face is selected, all polygons facing away from the camera will not be rendered, and will not cast shadows. If Cull Front Face is selected, all polygons facing toward the camera will likewise be dropped. Selecting both checkboxes has the same effect as deselecting the Visible checkbox.
- **Ignore Transparent Pixels in Aux Channels:** In previous versions of Fusion, transparent pixels were rejected by the software/GL renderers. To be more specific the software renderer rejected pixels with R=G=B=A=0 and the GL renderer rejected pixels with A=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 in post 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, then it would be useful to have transparent areas set aux channels (in particular UVs). As another example, suppose you are doing post DoF. You will probably not want the Z-channel to be set on transparent areas, as this will give you a false depth. Also keep in mind that this rejection 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 will not affect it.

Lighting

- **Affected by Lights:** If this checkbox is not selected, lights in the scene will not affect the object, it will not receive nor cast shadows, and it will be shown at the full brightness of its color, texture or material.
- **Shadow Caster:** If this checkbox is not enabled, the object will not cast shadows on other objects in the scene.
- **Shadow Receiver:** If this checkbox is not enabled, the object will not receive shadows cast by other objects in the scene.

Matte

Enabling the Is Matte option will apply 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 will override all textures. See the Matte Objects section of Chapter 60, “3D Compositing Basics,” for more information.

- **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 only visible when the Is Matte option is enabled.
- **Infinite Z:** Sets the value in the Z-channel to infinite. This checkbox is only visible when the Is Matte option is enabled.

Blend Mode

A bBlend mode specifies which method will be 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.

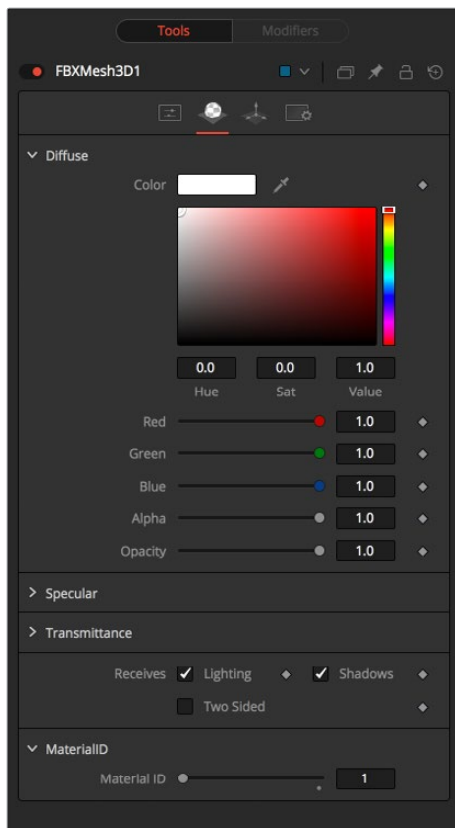
The blending modes were originally designed for use with 2D images. Using them in a lit 3D environment can produce undesirable results. For best results use the Apply modes in unlit 3D scenes rendered in software.

- **OpenGL Blend Mode:** Use this menu to select the blending mode that will be used when the geometry is processed by the OpenGL renderer. This is also the mode used when viewing the object in the Viewer. Currently the OpenGL renderer supports three blending modes.
- **Software Blend Mode:** Use this menu to select the blending mode that will be used when the geometry is processed by the software renderer. Currently the software renderer supports all of the modes described in the Merge node documentation, except for the Dissolve mode.

Material Tab

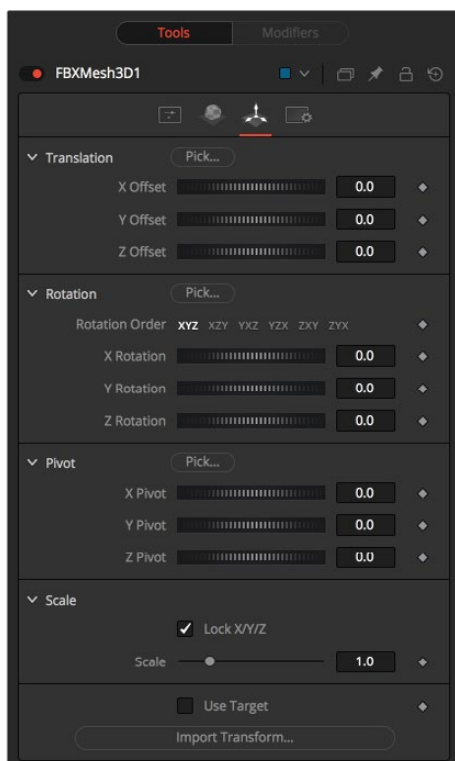
The options that appear in this tab determine the appearance of the geometry created by this node. Since these controls are identical on all nodes that generate geometry, these controls are fully described in the Common 3D Controls section of this documentation.

If an external 3D material is connected to the node’s material input, then the controls in this tab would be replaced with the Using External Material label.



Transform Tab

The options that appear in this tab determine the position of the geometry created by this node. Since these controls are identical on all nodes that generate geometry, these controls are fully described in the Common 3D Controls section of this documentation.



Fog 3D [3FO]



The Fog 3D node applies depth cue based fog to the scene. 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 antialiasing and depth of field effects during rendering.

The Fog 3D node essentially re-textures 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.

External Inputs

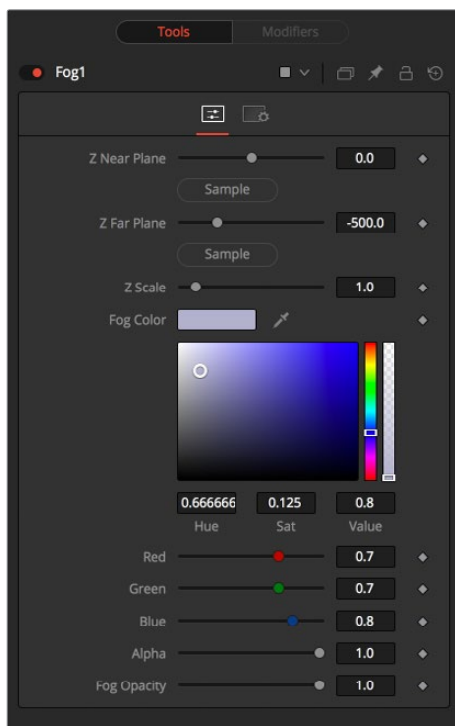
Fog3D.SceneInput

[orange, required] This input expects a 3D scene.

Fog3D.DensityTexture

[green, optional] This input expects a 2D image. The color of the fog created by this node is multiplied by the pixels in the image. When creating the image for the density texture, keep in mind that the texture is effectively projected onto the scene from the camera.

Controls



Enable

Use this checkbox to enable or disable the node.

Show Fog in View


By default, the fog created by this node is only visible 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 has been provided.

Radial

By default, the fog is done based upon 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 even though they remain the same distance from the eye. Radial fog fixes this. Sometimes Radial fog is not desirable. For example, if you are fogging an object that is close to the camera, like an image plane, the center of the image plane could be unfogged while the edges could be fully fogged.



Fog 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.

Image Plane 3D [3IM]



The Image Plane node produces 2D planar geometry in 3D space. The most common use of the node is to represent 2D images in the 3D space. An image input on the node provides the texture for the rectangle from another source in the composition. The aspect of the image plane is determined by the aspect of the image used for its diffuse texture. If planar geometry, whose dimensions are not relative to the texture image, is required, then use a Shape 3D node instead.

External Inputs

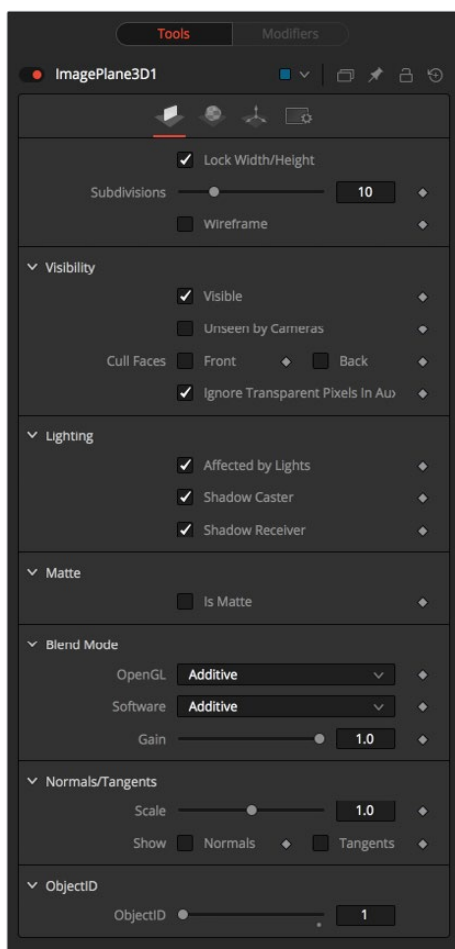
Imageplane3D.SceneInput

[orange, optional] This input expects a 3D scene. As this node creates geometry, it is not required.

Imageplane3D.MaterialInput

[green, optional] This input will accept either a 2D image or a 3D material. If a 2D image is provided, it will be used as a diffuse texture map for the basic material built into the node. If a 3D material is connected, then the basic material will be disabled.

Controls



Lock Width/Height

When checked, the subdivision of the plane will be applied evenly in X and Y. When unchecked, there are two sliders for individual control of the subdivisions in X and Y. 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 will be available to represent the lighting. For this reason, high subdivisions can be useful when working interactively with lights.

Wireframe

Enabling this checkbox will cause the Mesh to render only the wireframe for the object when using the OpenGL renderer.

Visibility

- **Visible:** If the Visibility checkbox is not selected, the object will not be visible in the Viewer, nor will it be rendered into the output image by the Renderer 3D node. A non-visible object does not cast shadows.
- **Unseen by Cameras:** If the Unseen by Cameras checkbox is selected, the object will be visible in the Viewers (unless the Visible checkbox is turned off), except when viewed through a camera. The object will not be rendered into the output image by the Renderer 3D node. Shadows cast by an Unseen object will still be visible when rendered by the software renderer, though not by the OpenGL renderer.
- **Cull Front Face/Back Face:** Use these options to cull (eliminate) rendering and display of certain polygons in the geometry. If Cull Back Face is selected, all polygons facing away from the camera will not be rendered, and will not cast shadows. If Cull Front Face is selected, all polygons facing toward the camera will likewise be dropped. Selecting both checkboxes has the same effect as deselecting the Visible checkbox.
- **Ignore Transparent Pixels in Aux Channels:** In previous versions of Fusion, transparent pixels were rejected by the software/GL renderers. To be more specific, the software renderer rejected pixels with R=G=B=A=0 and the GL renderer rejected pixels with A=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 in post 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. then it would be useful to have transparent areas set aux channels (in particular UVs). As another example suppose you are doing post DoF. You will probably not want the Z-channel to be set on transparent areas, as this will give you a false depth. Also keep in mind that this rejection 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 will not affect it.

Lighting

- **Affected by Lights:** If this checkbox is not selected, lights in the scene will not affect the object, it will not receive nor cast shadows, and it will be shown at the full brightness of its color, texture, or material.
- **Shadow Caster:** If this checkbox is not enabled, the object will not cast shadows on other objects in the scene.
- **Shadow Receiver:** If this checkbox is not enabled, the object will not receive shadows cast by other objects in the scene.

Matte

Enabling the Is Matte option will apply 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 will override all textures. See the Matte Objects section of Chapter 60, “3D Compositing Basics,” for more information.

- **Is Matte:** When activated, objects whose pixels fall behind the matte objects pixels in Z do not get rendered.
- **Opaque Alpha:** Sets the alpha value of the matte object to 1. This checkbox is only visible when the Is Matte option is enabled.
- **Infinite Z:** Sets the value in the Z-channel to infinite. This checkbox is only visible when the Is Matte option is enabled.

Blend Mode

A blend mode specifies which method will be 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.

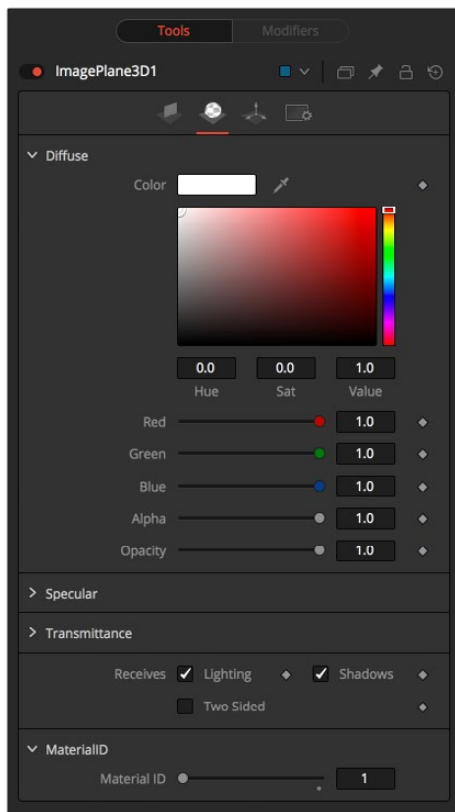
The blending modes were originally designed for use with 2D images. Using them in a lit 3D environment can produce undesirable results. For best results use the apply modes in unlit 3D scenes rendered in software.

- **OpenGL Blend Mode:** Use this menu to select the blending mode that will be used when the geometry is processed by the OpenGL renderer. This is also the mode used when displaying the object in the Viewers. Currently the OpenGL renderer supports three blending modes.
- **Software Blend Mode:** Use this menu to select the blending mode that will be used when the geometry is processed by the software renderer. Currently the software renderer supports all of the modes described in the Merge node documentation except for the Dissolve mode.

Material Tab

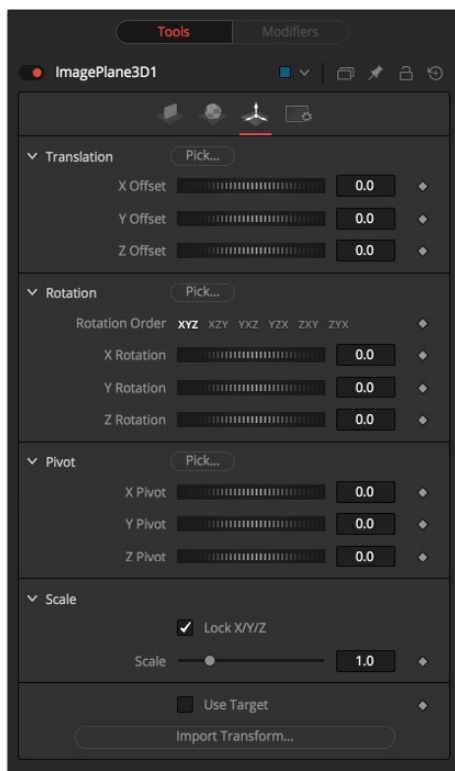
The options that appear in this tab determine the appearance of the geometry created by this node. Since these controls are identical on all nodes that generate geometry, these controls are fully described in the Common 3D Controls section of this documentation.

If an external 3D material is connected to the node’s material input, then the controls in this tab would be replaced with the Using External Material label.




Transform Tab

The options that appear in this tab determine the position of the geometry created by this node. Since these controls are identical on all nodes that generate geometry, these controls are fully described in the Common 3D Controls section of this documentation.



Locator 3D [3LO]



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 will transform the coordinates of a 3D control into 2D screen space. The 2D position is exposed as a numeric output which 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.

The scene provided to the Locator's input must contain the camera through which the coordinates are projected. As a result, 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 second input, the Locator will be positioned at the object's center, and the Transformation tab's Offset XYZ sliders will function in the object's local coordinate space rather than global scene space. This is useful for tracking an object's position regardless of any additional transformations applied further downstream.

External Inputs

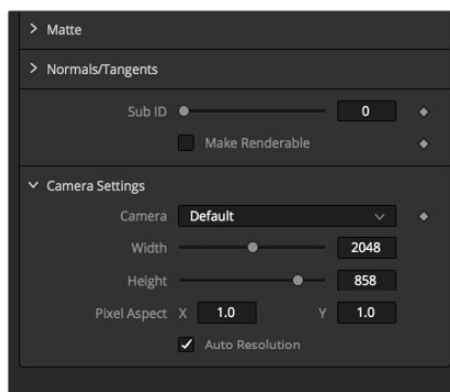
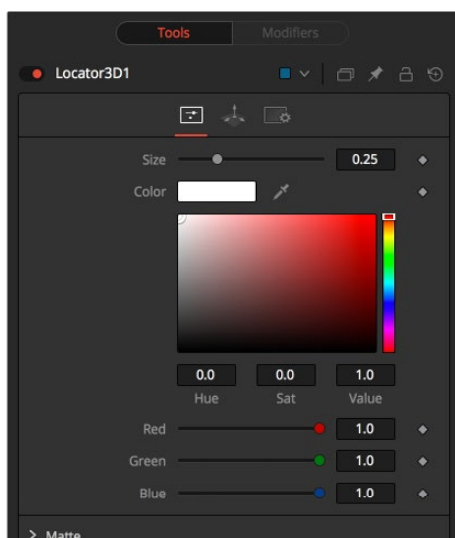
Locator3D.SceneInput

[orange, required] This input expects a 3D scene.

Locator3D.Target

[green, optional] This input expects 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.

Controls



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 will apply 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 will override all textures. See the Matte Objects section of Chapter 60, "3D Compositing Basics," for more information.

- **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 only visible when the Is Matte option is enabled.
- **Infinite Z:** Sets the value in the Z-channel to infinite. This checkbox is only visible when the Is Matte option is enabled.

Sub id

The Sub ID slider can be used to select an individual sub-element 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 if the Locator is rendered as a visible object by the OpenGL renderer. The software renderer is not currently capable of rendering lines and hence will ignore 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 will be 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 will display a contextual menu containing the frame formats configured in the composition's preferences.

Merge 3D [3MG]



The Merge 3D node is used to combine separate 3D elements into the same 3D environment.

For example, with a scene that is 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 node displays only two inputs at first, but as each input is connected a new input will appear on the node, assuring there is always one free to add a new element into the scene.

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 of the objects connected to the merge. This behavior forms the basis for all parenting in Fusion's 3D environment.


External Inputs

Merge3D.SceneInput[#]

[any, see description] These inputs expect a 3D scene. When the node is constructed it will display two inputs. 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.

Controls

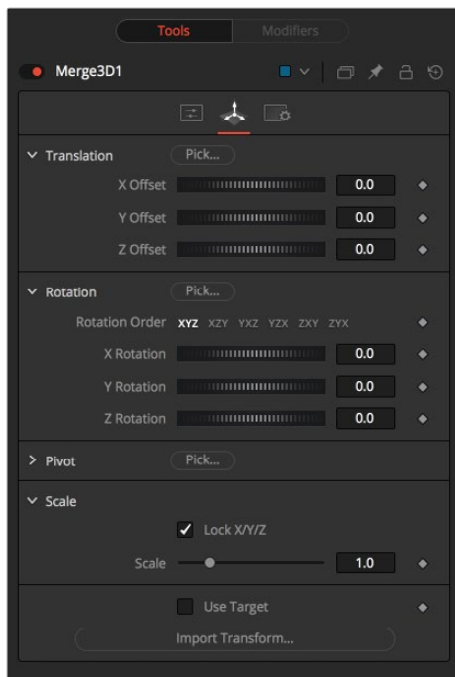


Pass Through Lights

When the Pass Through Lights checkbox is selected, lights will be passed through the merge into its output so they can 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.

Transform Tab

The options that appear in this tab determine the position of the geometry created by this node. Since these controls are identical on all nodes that generate geometry, these controls are fully described in the Common 3D Controls section of this documentation.



Override 3D [3OV]



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.

It is frequently used in conjunction with the Replace Material node to produce isolated passes. For example, a scene can be branched out to an Override node which turns off the Affected by Lights property of each node, then connected to a Replace Material node that applies a Falloff shader to produce a falloff pass of the scene.

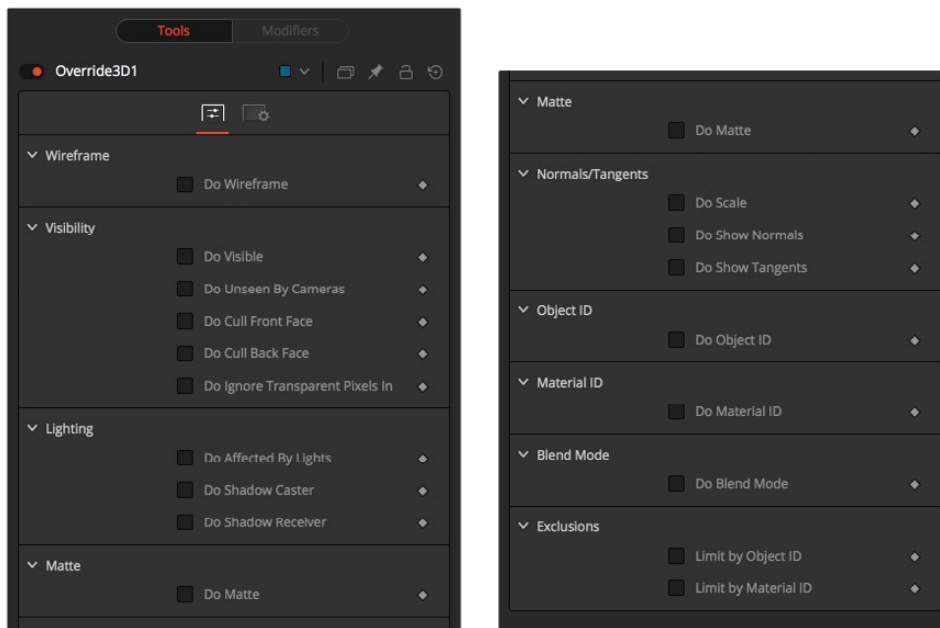
External Inputs

Override3D.SceneInput

[orange, required] This input expects a 3D scene.

Controls

The philosophy of the controls found in the Override node is fairly straightforward. First, you select the option to override using the Do [Option] checkbox. That will reveal 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, 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.

Point Cloud 3D [3PC]



A Point Cloud is generally a large number of points 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 by importing a 3D scene.

External Inputs

Pointcloud3DSceneInput

[orange, required] This input expects a 3D scene.

Controls



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 filetypes are:

* Alias's Maya	.ma
* 3DS Max ASCII Scene Export	.ase
* NewTek's LightWave	.lws
* Softimage XSI's	.xsi.


Make Renderable

Determines if the point cloud is visible in the OpenGL viewport 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 will be visible in the Viewers but not rendered into the output image by the Renderer 3D node.

Onscreen Contextual Menu



Frequently, one or more of the points in an imported point cloud will have been manually assigned in order 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 pop-up will appear with the name of the point.

When the Point Cloud 3D node is selected, a submenu will be added to the display 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 display Viewer contextual menu will open a dialog that can be used to search for and select a point by name. Each point that matches the pattern will be selected.

Rename

Rename one or more points by selecting Rename from the contextual menu. Type the new name into the dialog that appears and hit enter. The point will now have that name, with a four-digit number added to the end. For example, the name window will be 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).

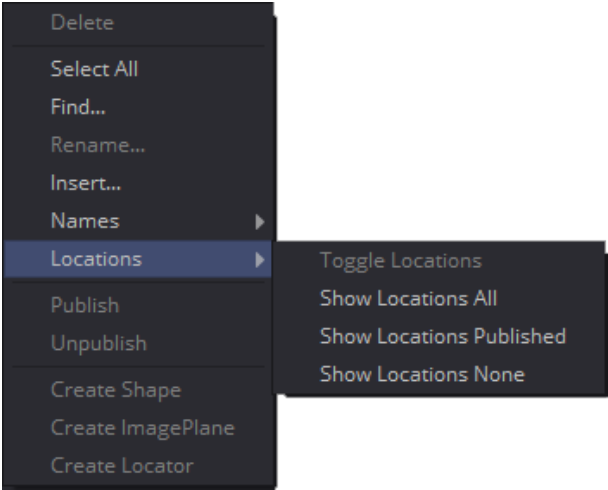
Delete

Selecting this option will delete the currently selected points.

Publish


Normally, the exact position of a point in the cloud is not exposed. To expose the position, select one or more points then select the publish option from this contextual menu. This will add a coordinate control to the control panel for each published point that displays the point's current location.

Additional Toolbar and 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 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 much more. The Projector node is just one of several nodes capable of projecting images and textures. Each method has advantages and disadvantages. See the Projection section of Chapter 60, "3D Compositing Basics," for more information about each approach.

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 together with a Merge 3D they can be moved as one and the texture will remain locked to the object.

The Projector 3D nodes'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 will be affected by the surface normals and can cause specular highlights. If this is undesirable, set Projector 3D to project into the Ambient Light channel.
- Enabling Shadows will cause Projector 3D to cast shadows.
- Just as with other lights, the light emitted by a Projector 3D will only affect objects that feed into the first Merge 3D that is downstream of the Projector 3D node on the node tree.
- Enabling Merge 3D's Pass Through Lights checkbox will allow 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 will not clip geometry in Light or Ambient Light mode. Use Texture mode instead.
- If two projections overlap, their light contributions will be added together.

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 will 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 will clip the geometry according to the alpha channel of the projected image.

See the section for the Catcher node for additional details.

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. Projector 3D 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.

When using projectors in an OpenGL view or in the OpenGL renderer, the lighting must be set to Per-pixel (the default) in order to see the results of the projection in the main view. The Per-pixel lighting mode requires a certain level of shader support from your graphics card. Even if the graphics card does not support pixel shading, the projection will still be visible when using the software renderer.

External Inputs

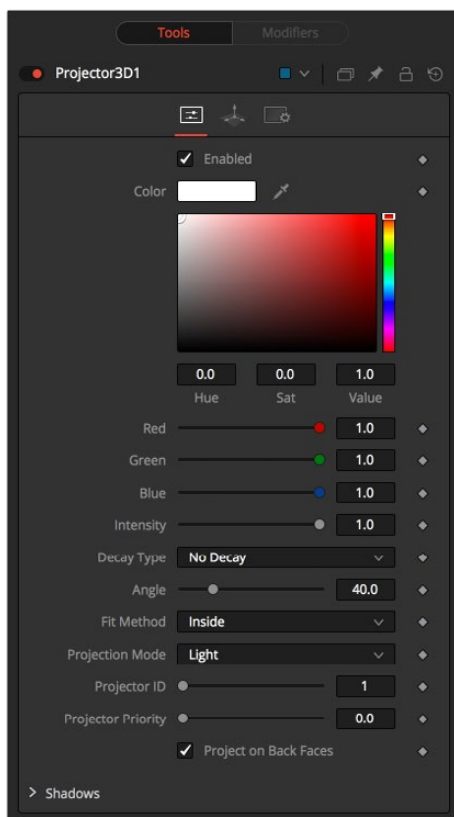
Projector3D.SceneInput

[orange, optional] This input expects a 3D scene. If a scene is connected to this input, then transformations applied to the spotlight will also affect the rest of the scene.

Projector3D.ProjectiveImage

[white, optional] This input expects a 2D image to be used for the projection.

Controls



Enabled

When this checkbox is selected the projector affects the scene. Clear the checkbox to turn it off.

Color

The input image is multiplied by this color before 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, regardless of 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 fit within the projection cone.

The first thing to know is that although this documentation may call it a 'cone,' the Projector3D node and the Camera3D nodes do not project actual cones; it's more of a pyramid of light with its apex at the camera/projector. The Projector3D 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 Camera3D 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 Projector3D/Camera3D 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 Camera3D) are different, the fit options become important.

For example, 'Fit by Width' will fit the width of the image across the width of the Camera3D pyramid. In this case, if the image has a greater aspect ratio than the aspect of the pyramid, some of the projection will extend 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 will ever receive 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 will always get 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 will only strike 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 will cause 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 how opaque the shadow will be. A density of 1.0 will produce 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 will 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 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 will occasionally be Z-fighting, where the portions of the object that should be receiving the shadows render over top of the shadow instead.

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.

Make adjustments to the multiplicative bias first, then fine tune the result using the additive bias control.

Force All Materials Non-Transmissive

Normally, a RGBAZ shadowmap is used when rendering shadows. By enabling this option you are forcing the renderer to use a Z-only shadowmap. This can result in 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 shadowmap when it is sampled. Fusion has three separate filtering methods available when rendering shadows which produce different effects.

- **None:** Shadows will have a hard edge. No filtering of the shadowmap is done at all. The advantage of this method is that you only have to sample one pixel in the shadowmap, so it is fast.
- **Constant:** Shadows edges will have a constant softness. A filter with a constant width is used when sampling the shadowmap. Adjusting the Constant Softness slider controls the size of the filter. Note that the larger you make the filter, the longer it will take to render the shadows.
- **Variable:** The softness of shadow edges will grow 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.

Constant Softness

If the Softness is set to constant, then this slider will appear. 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. To be more precise, it controls how fast the shadowmap filter size grows based upon the distance between shadow caster and receiver. Its effect is mediated by the values of the Min and Max Softness sliders.

Min Softness

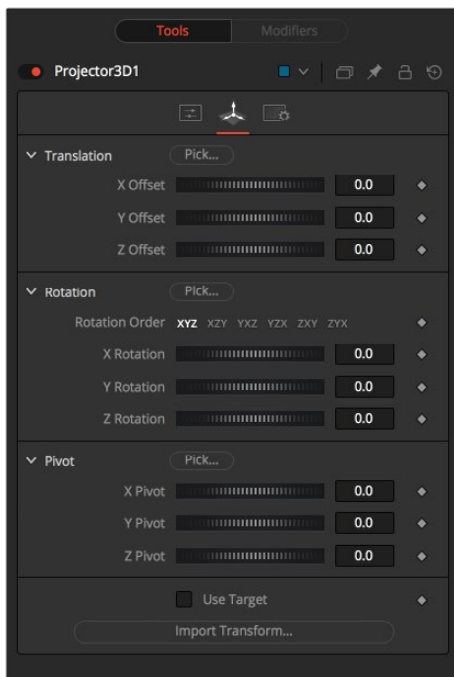
The Min Softness slider appears when the Softness is set to variable. This slider controls the Minimum Softness of the shadow. The closer the shadow is to the object casting the shadow, the sharper it will be up to the limit set by this slider.

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 will be up to the limit set by this slider.

Transform Tab

The options that appear in this tab determine the position of the geometry created by this node. Since these controls are identical on all nodes that generate geometry, these controls are fully described in the Common 3D Controls section of this documentation.



Renderer 3D [3RN]



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 should terminate with at least one Render 3D node. The Renderer node can use either of the software or OpenGL render engines to produce the resulting image. Additional render engines may also be available via third-party plug-ins.

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 Control tab. Be aware that scenes containing one or more particle systems will require that the Motion Blur settings on the pRender nodes exactly match the settings on the Renderer 3D node.

Otherwise, the subframe renders will conflict producing unexpected (and incorrect) results.

NOTE: The GL renderer respects the Color Depth option in the Image tab of the Renderer3D node. This can cause slowdowns on certain graphics cards when rendering to int16 or float32.

External Inputs

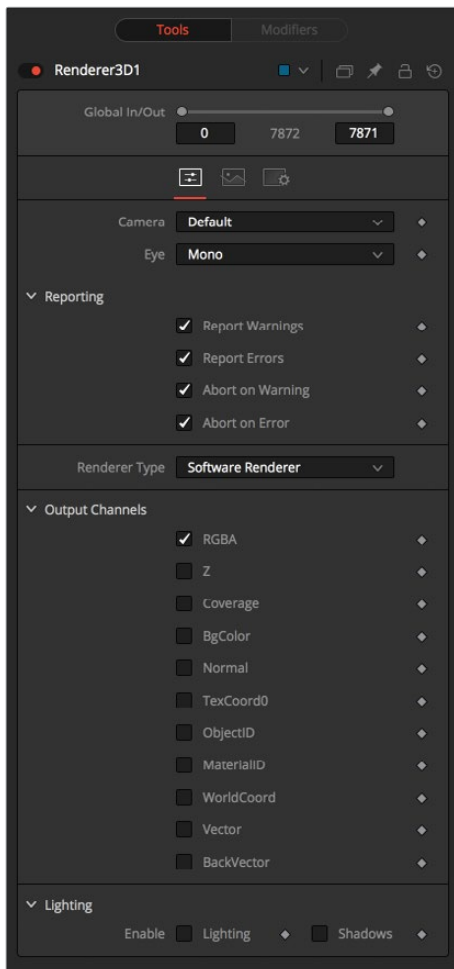
Renderer3D.SceneInput

[orange, required] This input expects a 3D scene.

Renderer3D.EffectMask

[violet, optional] This input uses a single or four channel 2D image to mask the output of the node.

Controls



Camera

The camera drop-down list is used to select which camera from the scene is used when rendering. The default option is Default, which will use the first camera found in the scene. If no camera is located, the default perspective view will be used instead.

Eye

The Eye control tells the node how to render the image in stereoscopic projects. The Mono option will ignore the stereoscopic settings in the camera. The Left and Right options will translate the camera using the stereo Separation and Convergence options defined in the camera to produce either left or right eye outputs.

Reporting

The first two checkboxes in this reveal can be used to determine whether the node will print warnings and errors produced while rendering to the console. The second row 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 drop-down menu lists the available render engines. Fusion provides two: the software and OpenGL render engines (described above), and additional renderers can be added via third-party plug-ins.

All of 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. As a result, each renderer is described in its own section below.

Software Controls

Output Channels

In addition to the usual Red, Green, Blue and Alpha channels, the software renderer can also embed the following channels into the image. Enabling additional channels will consume additional memory and processing time, so these should be used only when required.

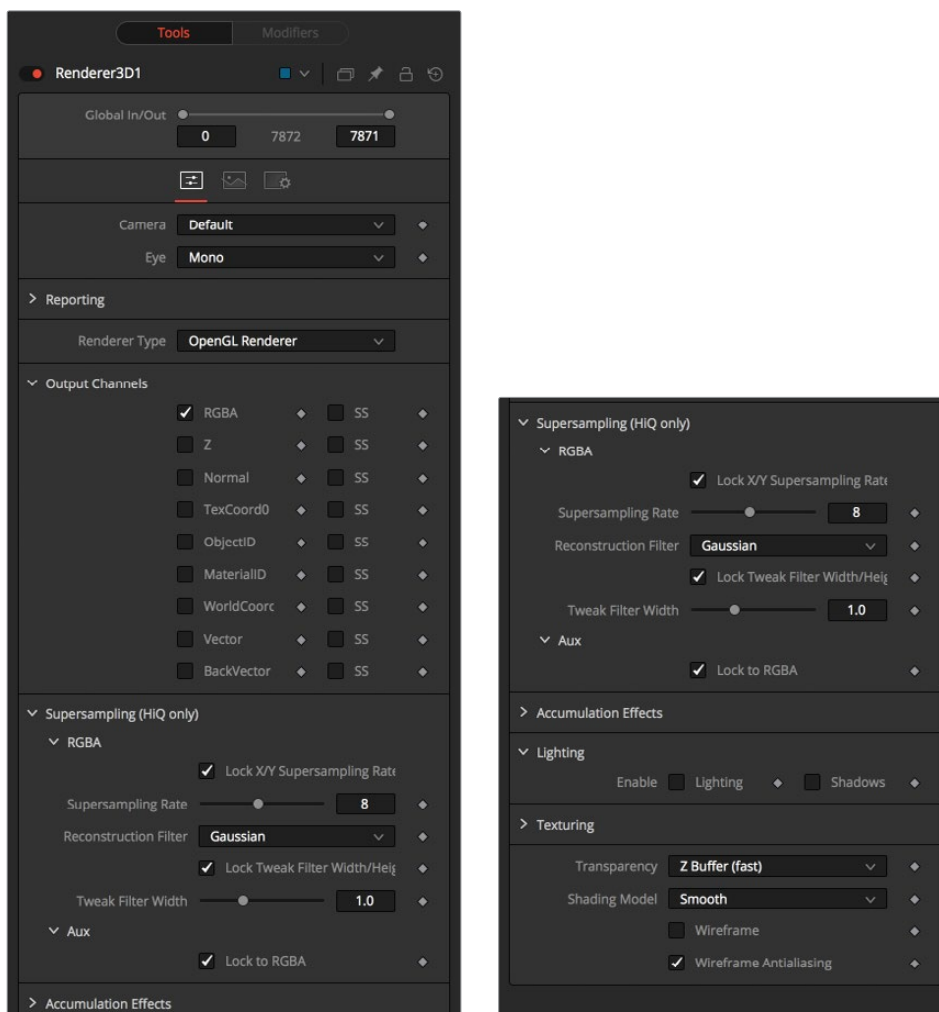
- **RGBA:** This option tells the renderer to produce the Red, Green, Blue and Alpha color channels of the image. These channels are required and they cannot be disabled.
- **Z:** This option enables rendering of the Z-channel. The pixels in the Z-channel contain a value that represents the distance 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 antialiasing. 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 3-component UVW, Fusion images only store 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 MaterialID 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 of the pixels that comprise a brick texture.

Lighting

- **Enable lighting:** When the Enable Lighting checkbox is selected, objects will be lit by any lights in the scene. If no lights are present, all objects will be black.
- **Enable Shadows:** When the Enable Shadows checkbox is selected, the renderer will produce shadows, at the cost of some speed.

OpenGL 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 will consume 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 3-component UVW, Fusion images only store 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 MaterialID 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 of the pixels that comprise a brick texture.

Anti Aliasing

Anti Aliasing in Fusion is divided into two methods Supersampling and Multi sampling. These two methods work together to provide refined control over speed and quality. You can enable the processing to be done at high quality (HiQ) or a lower quality (lowQ) for speed.

Super sampling renders the scene at a multiple higher resolution than the output image and resizing it down to create anti aliasing. Multi sampling antialiases the geometry of the scene so that super sampling can be done at a lower multiple. A good setting for rendering is Multi Sampling HiQ set to 8 and Super Sampling HiQ set to 3.

Anti Aliasing is enabled for a selected channel by means of the Channel menu.

The Anti Aliasing Preset provides separate options for color and aux channels since supersampling of color channels is a quite a bit slower than aux channels. You may find that a super sampling rate of 3 Rate is sufficient for color, but for world position or Z, you may require rate setting of 16 to get adequate results. The reasons color anti aliasing is slower are that the OpenGL 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.

Enable (LowQ/HiQ)

These two check boxes are used to enable anti aliasing. If the HiQ check box is enabled, high quality anti aliasing is used. If HiQ is disabled and LowQ is enabled lower quality anti aliasing is used. If both check boxes are disabled no anti aliasing is performed.

The GL renderer can perform the rescaling of the image directly on the GPU much more quickly than the CPU can manage it. In general, the more GPU memory the graphics card has the faster the operation will be.

Due to hardware limitations, point geometry (particles) and lines (locators) are always rendered at their original size, independent from supersampling. This means that these elements will be scaled down from their original sizes, and will likely appear much thinner than expected.

NOTE: For some things, sometimes using an Super sampling Z-buffer will improve 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 MaterialID and ObjectID channels. We recommend disabling super sampling on TexCoord, Normal, BackVector, and Vector channels. The issue arises when you have multiple 3D surfaces with radically different TexCoord values in one pixel. The super sampling does not restrict itself to sampling the main surface but will sample both surfaces. For example, if one surface has TexCoords that are approx (u,v)=(0, 0) within that pixel and the other surface has (0.5, 0.5), you will 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.


Super sampling 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 will be rendered and then scaled back to produce the target image. Set the multiplier higher to get better edge antialiasing at the expense of render time. Typically 8x8 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 will often depend 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 check boxes, then adjust the quality and Amount sliders.

The flurries you want the out of focus areas to be, the higher the quality setting you will need. A low amount setting will cause 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 will create rack of focus effects.

Lighting

- **Enable Lighting:** When the Enable Lighting checkbox is selected, any lights in the scene will light objects. If no lights are present, all objects will be black.
- **Enable Shadows:** When the Enable Shadows checkbox is selected, the renderer will produce 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 amount of detail in the scene's geometry for lighting, so generally produces superior results.

While using per-pixel lighting in the OpenGL renderer produces results closer to that produced by the more accurate software renderer, it still has some disadvantages. Even with per-pixel lighting, the OpenGL renderer is less capable of dealing correctly with semi transparency, soft shadows and colored shadows. The color depth of the rendering will be 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. As a result, semi-transparent objects may not be shown correctly, depending on their ordering within the scene.
- **Sorted (accurate):** This mode will sort 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 views and Flat produces a simpler and faster shading model.

Wireframe

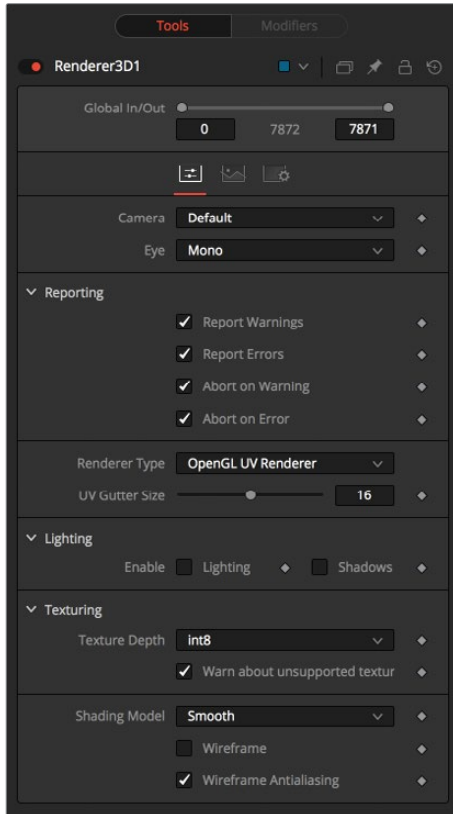
Renders the whole scene as Wireframe. This will show the edges and polygons of the objects. The edges are still shaded by the material of the objects.

Wireframe Antialiasing

Enables anti-aliasing for the Wireframe render.

OpenGL UV Unwrap Renderer

Takes a model with existing textures and renders it out via the `Renderergluv3D` to produce an unwound flattened 2D version of the model. Optionally lighting can be baked in. You can then paint on the texture and reapply it.




NOTE: After you have baked lighting into a models texture, you need to be careful to turn lighting off on the object in the future when you render it with the burnt-in lighting texture.

- **Issue:** 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 of this is the texture for a person where the artist mirrored the left side mesh/uvs/texture to produce the right side. Trying to burn in lighting in this case won't work.

Unwrapping more than one mesh at once can cause problems. The reason is that most models are authored so they make maximal usage of (u,v) in $[0,1] \times [0,1]$, so that in general models will 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.

Replace Material 3D [3RPL]



The Replace Material 3D node replaces the material applied to all of 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. The scope can also be limited to individual channels, making it possible to use a completely different material on the Red channel, for example.

External Inputs

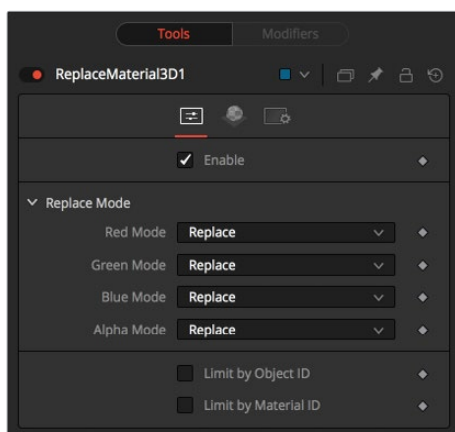
Replacematerial3D.SceneInput

[orange, required] This input expects a 3D scene.

Replacematerial3D.MaterialInput

[green, optional] This input will accept either a 2D image or a 3D material. If a 2D image is provided, it will be used as a diffuse texture map for the basic material built into the node. If a 3D material is connected, then the basic material will be disabled.

Controls



Enable

Enables the material replacement.

Replace Mode

Red, Green, Blue, and Alpha Mode

Offers several replacing modes for each RGBA channel:

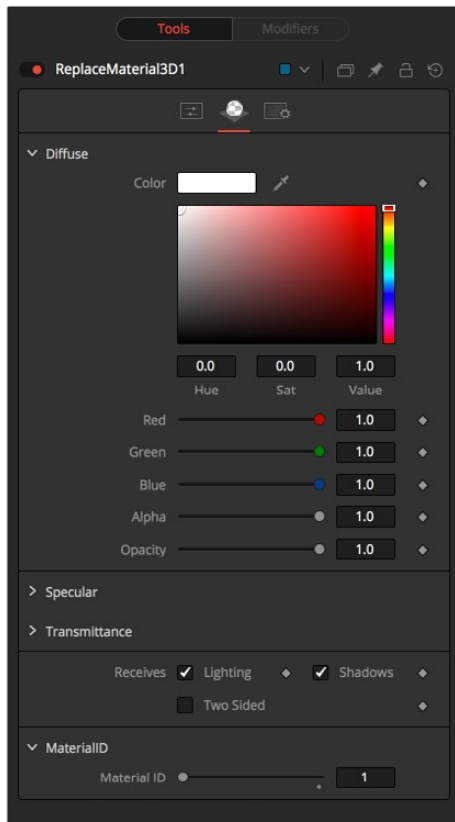
- **Keep:** Prevents the channel from being replaced by the input material.
- **Replace:** Replaces the material for the according color channel.
- **Blend:** Blends the materials together.
- **Multiply:** Multiplies the according channels of both inputs.

Limit by Object id/Material id

When enabled, a slider appears where the desired IDs can be set. All other objects will keep their materials. If both options are enabled an object must satisfy both conditions.


Material Tab

The options that appear in this tab determine the appearance the geometry created by this node. Since these controls are identical on all nodes that generate geometry, these controls are fully described in the Common 3D Controls section of this documentation.




If an external 3D material is connected to the nodes's material input, then the controls in this tab will be replaced with the "Using External Material" label.

Replace Normals 3D [3RPN]



ReplaceNormals is used to replace the Normals/Tangents on incoming geometry. All geometry in the input scene 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 FBX import because they were not present on the FBX.

Controls



Pre-Weld Position Vertices

Pre-welds the 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 will always be recomputed.
- **If Not Present:** The normals on the mesh are recomputed only if they are not present.
- **Never:** The normals will never be computed. This option is useful when animating.

Smoothing Angle

Adjacent faces with angles in degrees smaller than this will 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. 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 will not be 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 will be smoothed across. There is currently no way to visualize Smooth Groups within Fusion.

Flip Normals

Flipping of tangents can sometimes be confusing. Flip will have an effect if the mesh has tangent vectors. Most meshes in Fusion don't have tangent vectors until they reach a `Renderer3D`, though. Also, when viewing tangent vectors in the Viewers, the tangent vectors will be created if they don't exist. The confusing thing is if you view a `Cube3D` that has no tangent vectors and press the `FlipU/FlipV` button, nothing happens. This is because there were no tangent vectors to create, but later the GL renderer created some (unflipped) tangent vectors.

NOTE: The FBX importer will recompute the normals if they don't exist, but you can get a higher quality result from this node.

Bumpmaps can sometimes depend on the model's normals. In particular, when you simplify a complex high poly model to a low poly model + bumpmap, the normals and bumpmap can become 'linked.' Recomputing the normals in this case can make the model look funny. The bumpmap was intended to be used with those normals.

Most primitives in Fusion are not generated with tangents; when needed they are generated on the fly by a `Renderer3D` and cached.

Tangents currently are only needed for bumpmapping. If a material needs bumpmapping, then tangents are created. These tangents are created with some default settings (e.g., `SmoothingAngle`, and so on), and if you don't want Fusion automatically creating tangents you should manually create them using `ReplaceNormals`.

All the computations are done in the local coordinates of the geometries rather than in the coordinate system of the `ReplaceNormals3D` node. This can cause problems when there is a non-uniform scale applied to the geometry before `ReplaceNormals3D` is applied.

Replicate 3D [3REP]



The Replicate 3D node replicates input geometry at positions of destination vertices. This includes 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 for non-uniform transformations, such as random positioning or sizes.

External Inputs

Replicate3d.Destination

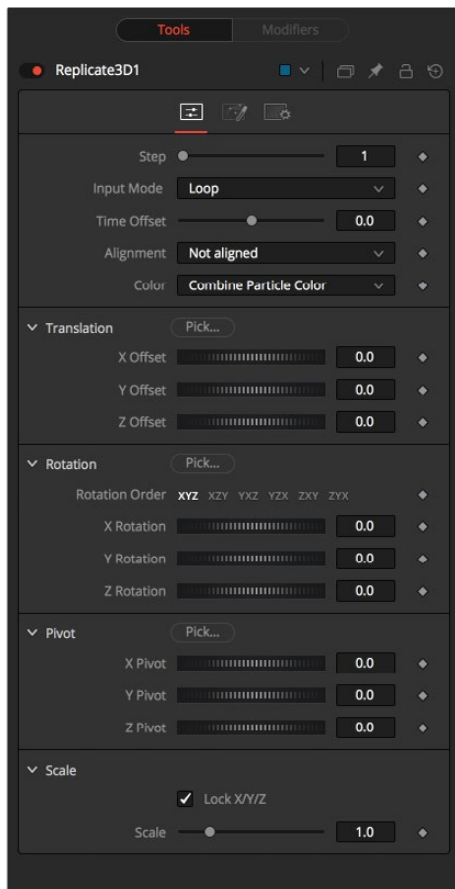
[orange, required] This input expects a 3D scene with vertex positions, either from meshes or 3D particle animations.

Replicate3d.Input[#]

[any, required] This input expects a 3D scene that will be replicated. Once connected, a new input for alternating 3D scenes will be created.

At least one connected input is required.

Controls



Step

Defines how many positions are skipped. For example, a step of 3 means that only every third position of the destination mesh will be used, while a step of 1 means that all positions will be used.

The step 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.

Pointclouds are internally represented by six points once the Make Renderable option has been set. In order to get a single point, use a step of 6 and set a X offset of -0.5 in the Replicate3D to get to the center of the pointcloud. Use -0.125 for Locator3Ds. Once these have been scaled, the offset may differ.


Input Mode

This parameter defines in which order multiple input scenes are replicated at the destination. Both parameters won't have a visible effect if only one input scene is supplied.

With Loop, the inputs are used successively. The first input will be at the first position, the second at the second, and so on. If there are more positions in the destination present than inputs, the sequence will be looped.

Random will use a definite but random input for each position based on the seed in the Jitter tab. This input mode can be used to simulate variety with few input scenes.

Death of Particles causes their 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 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 will show the animation from a frame earlier. The second copy will show 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.


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.



- **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 will be noticeable. For best results, it is recommended to use this method at the origin before any transformations.



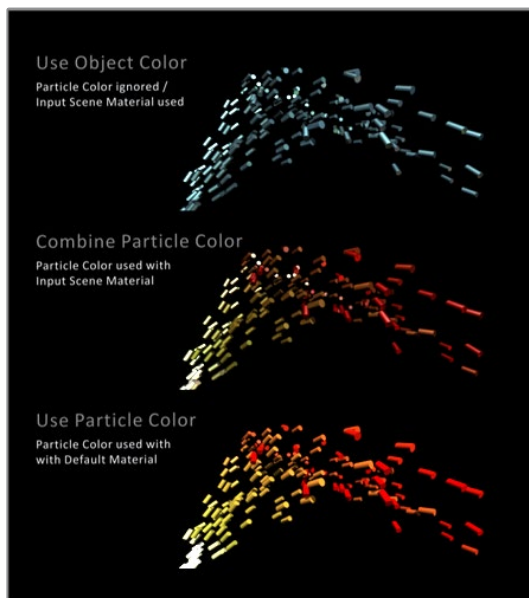
- **Aligned TBN:** Results in a more accurate and stable alignment based on the tangent, binormal, and normal of the destination point. Works best for particles and geometric shapes. On unwelded meshes, two copies of multiple unwelded points at the same position may result in different alignments due to their individual normals.



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.



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, 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

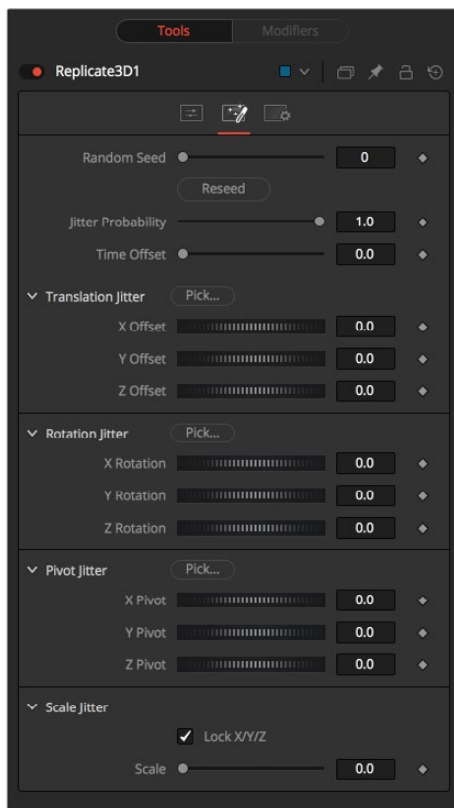
When the Lock XYZ checkbox is selected, any adjustment to the scale will be applied to all three axes simultaneously.

If this checkbox is disabled, the scale slider will be replaced with individual sliders for the X, Y, and Z scale.

Scale

The Scale control sets how much scaling to apply to each copy.

Jitter



Random Seed/Randomize

The Random Seed is used to 'seed' the amount of jitter applied to the replicated objects. Two replicate nodes with identical settings but different random seeds will produce two completely different results. Click on 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 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 will show the animation from a frame earlier. The second copy will show 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.

Translation XYZ Jitter

Use these three controls to adjust the amount of variation in the translation of the replicated objects.

Rotation XYZ Jitter

Use these three controls to adjust the amount of variation in the rotation of the replicated objects.

Pivot XYZ Jitter

Use these three controls to adjust the amount of 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 amount of variation in the scale of the replicated objects. Uncheck the Lock XYZ checkbox to adjust the scale variation independently on all three axes.

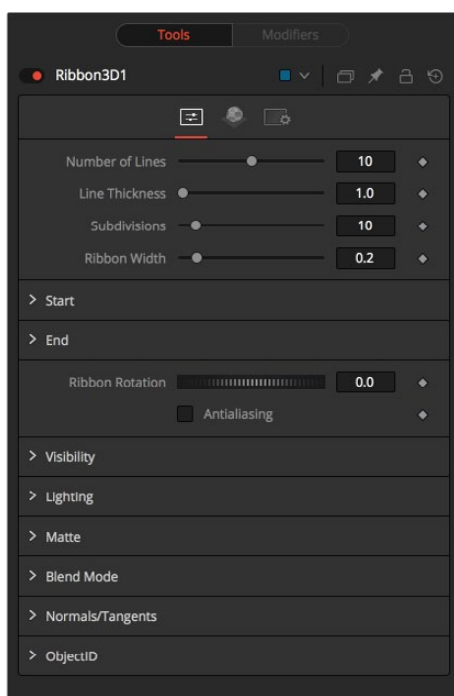
Ribbon 3D [3RI]



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, UVMap3D can be used to alter the texture coordinates. This node heavily relies on certain OpenGL features and will not produce any visible result in the software renderer.

Furthermore, the way lines are drawn is completely up to the graphic card vendor, so any artifacts may vary from card to card.

Controls



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 will only allow 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 endpoint. The higher the number, the more precise any 3D displacement can be applied.

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 point.

Antialiasing

Allows you to apply antialiasing to the rendered lines. Using antialiasing isn't necessarily advised. When activated, there will 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 graphic card vendor, which means that these artifacts can vary from card to card.

Shape 3D [3SH]



The Shape 3D node is used to produce several basic primitive 3D shapes, including planes, cubes, spheres and cylinders.

External Inputs

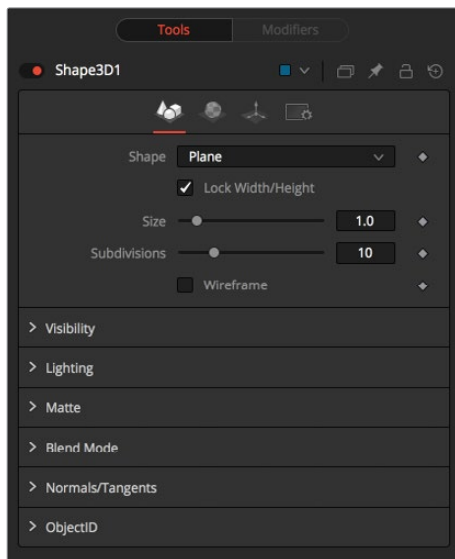
Shape3d.SceneInput

[orange, required] This input expects to receive a 3D scene.

Shape3d.MaterialInput

[green, optional] This input will accept either a 2D image or a 3D material. If a 2D image is provided it will be used as a diffuse texture map for the basic material built into the node. If a 3D material is connected, then the basic material will be disabled.

Controls



Shape

Select one of these options to determine which geometry primitive will be produced by the Shape 3D node. The controls below will 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 control over the size of the shape along each axis is provided.
- **Size Width/Height/Depth:** [plane, cube] Used to control the size of the shape.

Cube Mapping

[Cube] Uses cube mapping to apply the Shape node's texture (a 2D image connected to the Texture input).

Radius

[Sphere, Cylinder, Cone, Torus] Sets the radius of the selected shape.

Top Radius

[Cone] This control is used to define a radius for the top of a cone, making it possible to create truncated cones.

Start/End Angle

[Sphere, Cylinder, Cone, Torus] This range control determines how much of the sweep 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

[Sphere, Torus] This range control is used to slice the object by defining a latitudinal sub-section of the object.

Bottom/Top Cap

[Cylinder, Cone] Used for cylinder and cone shapes only, 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

[Torus] Used for torus only, section controls the thickness of the tube making up the torus.

Subdivision Level/Base/Height

[All shapes] Used for all shapes, the Subdivision controls are used to determine the tessellation of the mesh composing the object. The higher the subdivision, the more vertices each shape will have.

Wireframe

Enabling this checkbox will cause the mesh to render only the wireframe for the object.

Visibility

- **Visible:** If the Visibility checkbox is not selected, the object will not be visible in the Viewers, nor will it be rendered into the output image by the Renderer 3D node. A non-visible object does not cast shadows.
- **Unseen by Cameras:** If the Unseen by Cameras checkbox is selected, the object will be visible in the Viewers (unless the Visible checkbox is turned off), except when viewed through a camera. The object will not be rendered into the output image by the Renderer 3D node. Shadows cast by an unseen object will still be visible when rendered by the software renderer, though not by the OpenGL renderer.
- **Cull Front Face/Back Face:** Use these options to cull (eliminate) rendering and display of certain polygons in the geometry. If Cull Back Face is selected, all polygons facing away from the camera will not be rendered, and will not cast shadows. If Cull Front Face is selected, all polygons facing toward the camera will likewise be dropped. Selecting both checkboxes has the same effect as deselecting the Visible checkbox.
- **Ignore Transparent Pixels in Aux Channels:** In previous versions of Fusion, the software/OpenGL renderers rejected transparent pixels. To be more specific, the software renderer rejected pixels with R=G=B=A=0 and the GL renderer rejected pixels with A=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 in post 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, then it would be useful to have transparent areas set aux channels (in particular UVs).

As another example, suppose you are doing post DoF. You will probably not want the Z-channel to be set on transparent areas, as this will give you a false depth. Also keep in mind that this rejection 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 will not affect it.

Lighting

- **Affected by Lights:** If this checkbox is not selected, lights in the scene will not affect the object, it will not receive nor cast shadows, and it will be shown at the full brightness of its color, texture or material.
- **Shadow Caster:** If this checkbox is not enabled, the object will not cast shadows on other objects in the scene.
- **Shadow Receiver:** If this checkbox is not enabled, the object will not receive shadows cast by other objects in the scene.

Matte

Enabling the Is Matte option will apply 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 will override all textures. See the matte objects section of the 3D chapter for more information.

- **Is Matte:** When activated, objects whose pixels fall behind the matte objects pixels in Z do not get rendered.
- **Opaque Alpha:** Sets the alpha value of the matte object to 1. This checkbox is only visible when the Is Matte option is enabled.
- **Infinite Z:** Sets the value in the Z-channel to infinite. This checkbox is only visible when the Is Matte option is enabled.

Blend Mode

A blend mode specifies which method the renderer will use when combining this object with the rest of the scene. The blend modes are essentially identical to those listed in the documentation for the 2D Merge node. For a detailed explanation of each mode, see the documentation for that node.

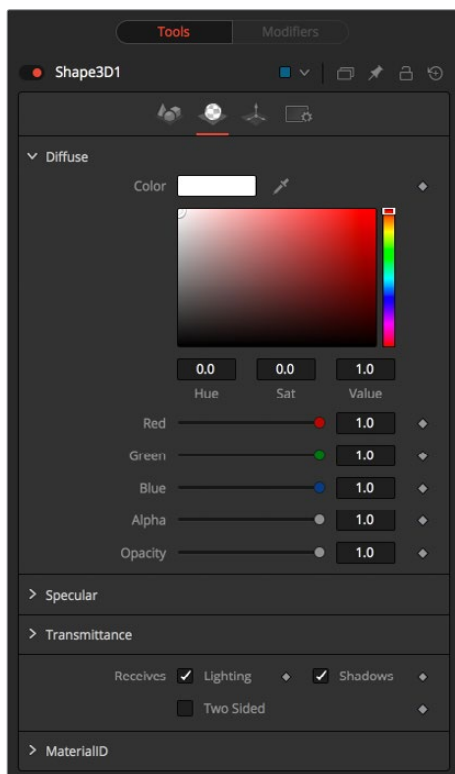
The blending modes were originally designed for use with 2D images. Using them in a lit 3D environment can produce undesirable results. For best results, use the Apply modes in unlit 3D scenes rendered in software.

- **OpenGL Blend Mode:** Use this menu to select the blending mode that will be used when the geometry is processed by the OpenGL renderer. This is also the mode used when viewing the object in the Viewers. Currently the OpenGL renderer supports three blending modes.
- **Software Blend Mode:** Use this menu to select the blending mode that will be used when the geometry is processed by the software renderer. Currently, the software renderer supports all of the modes described in the Merge node documentation, except for the Dissolve mode.

Material Tab

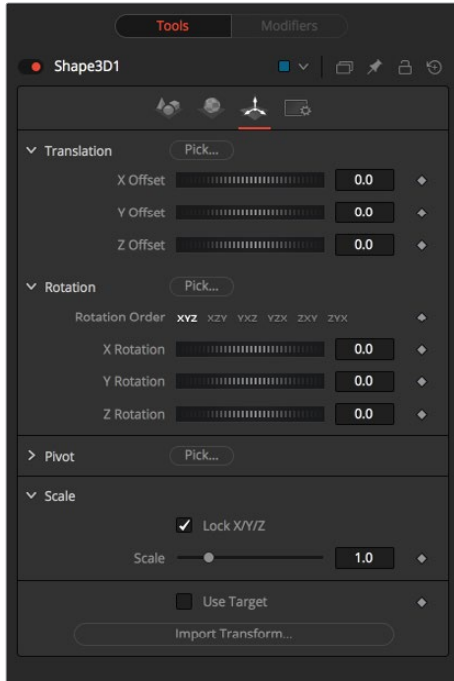
The options that appear in this tab determine the appearance the geometry created by this node. Since these controls are identical on all nodes that generate geometry, these controls are fully described in the Common 3D Controls section of this documentation.

If an external 3D material is connected to the nodes's material input then the controls in this tab will be replaced with the "Using External Material" label.



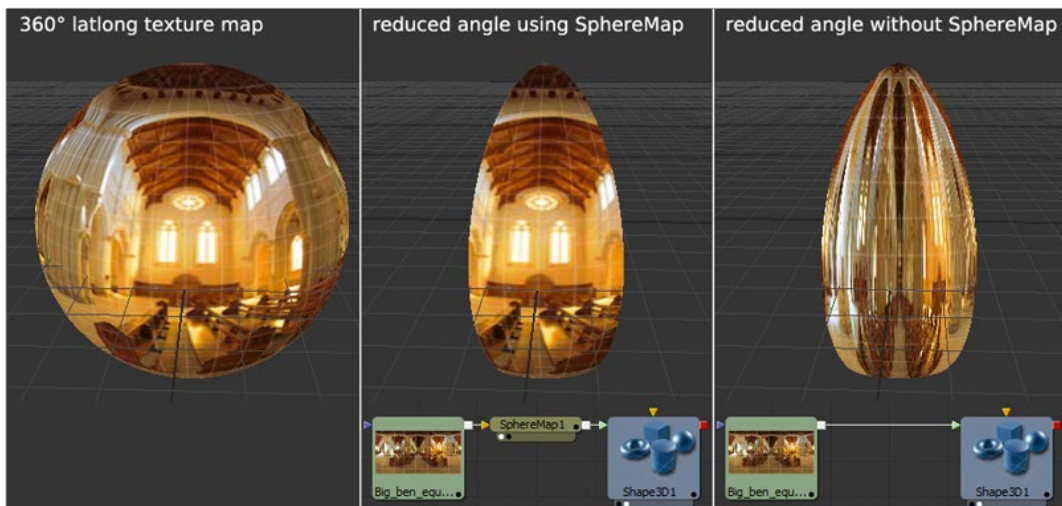
Transform Tab

The options that appear in this tab determine the position of the geometry created by this node. Since these controls are identical on all nodes that generate geometry, these controls are fully described in the Common 3D Controls section of this documentation.




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 will be squashed. When using the Sphere Map node, the texture will be cropped. Compare:



NOTE: If you pipe the texture directly into the sphere, it will also be mirrored horizontally. You can “fix” this by using a Transform node first.

Softclip [3SC]



The Softclip 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.


The Softclip node is usually placed just before the renderer 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.

External Inputs

Softclip.SceneInput

[orange, required] This input expects a 3D scene.

Controls



Enable

This checkbox can be used to enable or disable the node.

Smooth Transition

Changes the linear gradient to a curve shaped one.

Radial

By default, the softclipping is done based upon 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 softclipping 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, even though they remain the same distance from the eye. Radial softclip fixes this. Sometimes Radial softclipping is not desirable.

For example, if you apply softclip 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 softclip becomes visible in the scene from all points of view.

Transparent/Opaque Distance

Defines the range of the softclip. The objects will 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.

Text 3D [3TXT]



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 that it only supports one shading element. See Text Plus for more details on the controls found in this node.

The Text 3D node was based on a tool that pre-dates the Fusion 3D environment. As a result, 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.

One thing to be aware of when network-rendering a comp that contains a Text3D node is that each machine that does not have the font installed will fail. Fusion cannot share or copy fonts to Render Slaves due to legal issues.

External Inputs

Text3d.SceneInput

[orange, required] This input expects a 3D scene.


Text3d.ColorImage

[orange, required] This input expects a 2D image. It is only visible when the Image mode is enabled in the Material section of the Shading tab.

Text3d.BevelTexture

[orange, required] This input expects a 2D image. It is only visible when the Image mode is enabled in the Bevel Material section of the Shading tab.

Text Extrusion



Extrusion Depth

An extrusion of 0 produces completely 2D text. Any value greater than 0 will extrude the text to generate text with depth.

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 will be smoothed as well. If the keypoint is linear, there will be a sharp shading edge. The first and last keypoint on the spline defines the extent of the text.

- **Custom Extrusion Subdivisions:** Controls the number of subdivisions within the smoothed portions of the extrusion profile.

TIP: Remember that the spline can also be edited from within the Spline Editor tab and you do not have to work with the limited abilities and size of the text3D.

Extrusion profile spline control: Do not try to go to zero size at the Front/Back face. This will result in Z-fighting due to self-intersecting faces. To avoid this problem, make sure the first and last point have their profile set to 0.

Shading


Opacity

Reducing the material's opacity will decrease 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 will reveal a second set of Material controls for the beveled edge of the text.

Material Tab



Type

To use a solid color texture, select the Solid mode. Selecting the Image mode will reveal 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 will appear on the node that can be used to apply the output of a 2D node as the texture. Selecting Clip will expose a file browser that can be used to select an image or image sequence from disk. The Brush option will provide a list of clips found in the Fusion\brushes folder.

Bevel Material

This reveal only appears when the Use One Material checkbox control is selected. The controls under this reveal are an exact copy of the Material controls above, but are applied only to the beveled edge of the text.

Transform

These controls can be used to transform the material applied to the text. See the Text+ node for a complete description of these nodes.

Tips for Text3D


Character Level Styling

The Text 3D node doesn't support Character Level Styling directly. You have to 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).

Uncapped 3D Text

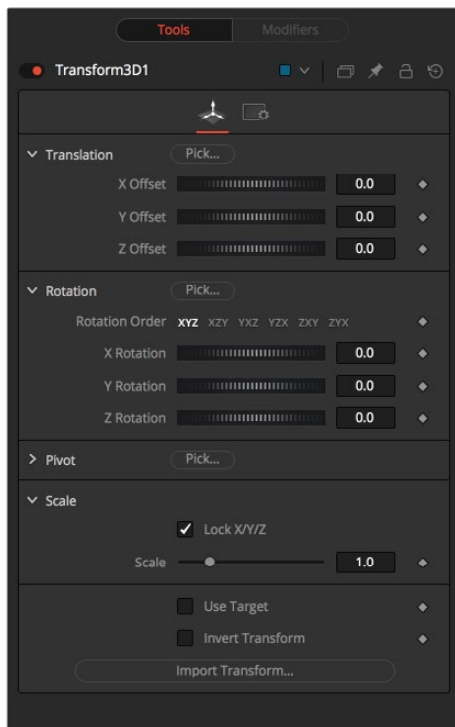
To hide the front face of extruded text, uncheck Use One Material on the Shading tab and reduce the first material's color to black including its alpha value.

Transform 3D [3XF]



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.

Transform 3D Tab



Translation

- **X, Y, Z Offset:** These controls can be used to position the 3D element.

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 followed by 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 will rotate 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 will be shown. This adjusts the overall size of the object. If the Lock checkbox is unchecked, individual X, Y, and Z sliders will be 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

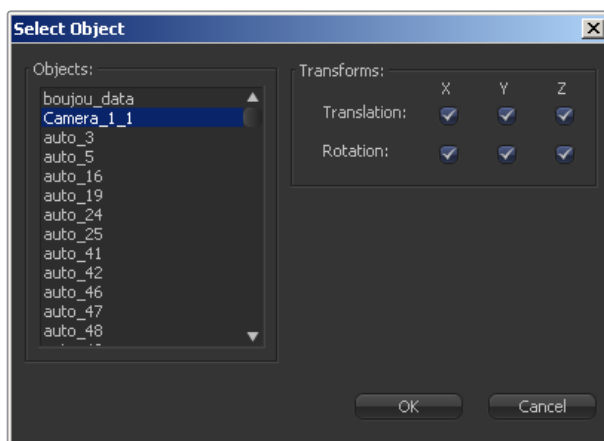
Selecting the Use Target checkbox enables a set of controls for positioning an XYZ target. When Target is enabled, the object will always rotate to face the target. The rotation of the object becomes relative to the target.

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 will only import transformation data. For 3D geometry, lights and cameras, consider using the File > FBX Import option from the menus.



On Screen Transformation Widget

A Transformation Widget represents most of the controls in this tab in the Viewer, with modes for transformation, rotation, and scaling. To change the mode of the widget, 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 defaults 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.




Triangulate 3D [3TRI]



This node has no controls. It triangulates any quads or convex polygons.



UV Map 3D [3UV]



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. The node provides 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 for 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.

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.

External Inputs

UVMap3d.SceneInput

[orange, required] This input expects to receive a 3D scene.

UVMap3d.CameraInput

[green, optional] This input expects the output of the Camera 3D node. It is only visible when the Camera Map mode is used.

Camera Projections with UV Map 3D

The Camera Mapping mode makes it possible to project texture coordinates onto geometry through a camera. Enable the Camera Mapping mode on the UV Map 3D node. Directly connect the camera that will be used to create the UV coordinates to the new 'Camera' input that appears on the UV Map 3D node in the node tree.

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 will use the UV coordinates created by the camera. Because this is a texture projection and not light, the alpha channel of the texture will correctly set the opacity of the geometry.


See the Camera 3D and Projector 3D node for alternate approaches to projection.

The projection can optionally be locked to the vertices as it appears on a selected frame.

This will fail 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. For this reason, projection locking does not work for many particle systems, for primitives with animated subdivisions, and 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 will use. This may be confusing because the material usually sits upstream. For example, a composition containing the nodes: Loader > Blinn > Shape 3D (cube) > UV Map 3D > Renderer 3D.

Controls



Map Mode

Defines how the texture coordinates are created. Think of it as a virtual geometry which 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 port of the node. After connecting a camera to it, 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 will fit the Map mode to the bounding box of the input scene.

Center

Clicking this button will move 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. Enabling this option will do so and also reveal 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 uvmap 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 followed by the Z-axis.

Rotation X/Y/Z

Sets the orientation of the projection object for each axis, dependent from the rotation order.

Tile U/V/W

Defines how often a texture fits into the projected UV space on the according axis. Note that the UVW coordinates are transformed, not a texture. Works best when used in conjunction with the Create Texture node.

Flip U/V/W

Mirrors the texture coordinates around the according axis.

Flip Faces (CubeMap 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 will take 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 is 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 will cause 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.

Weld 3D [3WE]




Sometimes 3D geometry has vertices that should have been joined but haven't. This can cause artifacts, especially when the two vertices have different normals.

For example, you may find:

- The different normals will produce a hard shading/lighting edge where none was intended.
- If you try to Displace3D the vertices along their normals, a crack will appear.
- Missing pixels or doubled up pixels in the rendered image.
- Particles will pass through the tiny invisible cracks.

Rather than roundtripping back to your 3D modeling application to fix the 'duplicated' vertices, the Weld3D node allows you to do this in Fusion. Weld3D 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 even though the positions of two vertices have been made the same, their normals will still have their old values. This can lead to hard edges in certain situations.

Controls



Fracture

Fracturing is the opposite of welding, so all vertices are unwelded. This means that all polygon adjacency information is lost. For example, an Imageplane3D 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 Weld3D when issues occur with the geometry. Don't use it everywhere just because it's there, as it will influence render time.

Weld3D is intended to be used as a mesh robustness tool and not as a mesh editing tool to merge together vertices. If you can see the gap between the vertices you want to weld in the 3D view, you are probably misusing Weld3D. Unexpected things may happen when you do this; do so at your own peril.

Current Issues

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 ISS and there are lots of big polygons and lots of really tiny polygons. If you set the tolerance too large, small polygons will be merged that shouldn't; if you set the tolerance too small, some large polygons won't be merged.

Vertices that are far from the origin can fail to be merged correctly. This is because $\text{bignumber} + \text{epsilon}$ can exactly equal bignumber in float math. This is one reason it may be best to merge in local coords and not in world coords.

Sometimes Weld3-ing a mesh can make things worse. Take Fusion's cone as an example. 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 will merge together and will only have one normal, making the lighting look weird.

Weld3D is not multithreaded.

Warning

Do not misuse Weld3D 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.

Modifier

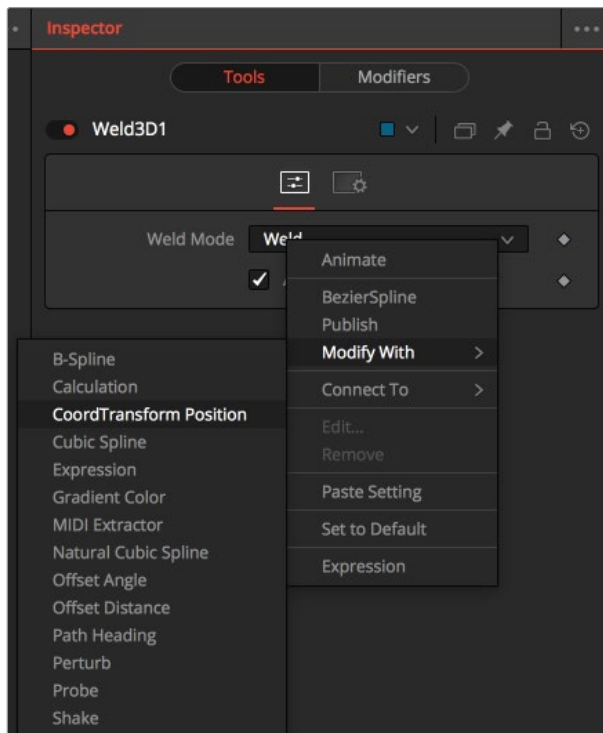
Coordinate Transform 3D

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 will calculate the current position of a given object at any point in the scene hierarchy.

To add a Coordinate Transform modifier, simply right-click on the a numeric input on any node, and select Modify With/CoordTransform Position from the Controls contextual menu.

Controls



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 on 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 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 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 on the control and select an object from the Connect To pop-up menu.

Chapter 27


3D Light Nodes

This chapter details the 3D Light nodes available in Fusion.

Contents

Ambient Light [3AL]	648
Directional Light [3DL]	649
Point Light [3PL]	650
Spot Light [3SL]	652

Ambient Light [3AL]



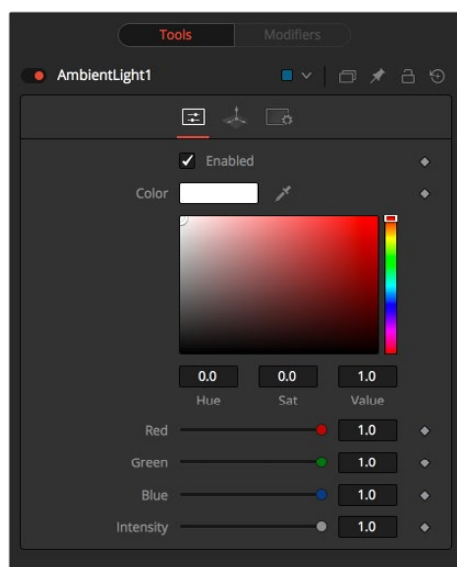
An Ambient Light is a directionless light that globally illuminates a scene. It has no real position or rotation, although an onscreen widget will appear in the views to indicate that a light is present in the scene. Position controls for the widget are provided to make it possible to move the widget out of the way of other geometry, if necessary.

External Inputs

AmbientLight.SceneInput

[orange, optional] This input expects a 3D scene. If a scene is provided, the Transform controls in this node will apply to the entire scene provided.

Controls



Enabled

When the Enabled checkbox is selected, the ambient light affects the scene. Clear the checkbox to turn off the light.

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 .2 indicates 20% percent light. A perfectly white texture lit only with a .2 ambient light would render at 20% gray (.2, .2, .2).

Transform Tab

The options that appear in this tab determine the position of the geometry created by this node. Since these controls are identical on all nodes that generate geometry, these controls are fully described in the Common 3D Controls section of this documentation.



Directional Light [3DL]



A Directional Light is a light with a clear direction but without a clear source. This light shows an onscreen widget, but the position of the widget has no meaning. The rotation of the widget is used to determine from where in the scene the light appears to be coming.


External Inputs

DirectionalLight.SceneInput

[orange, optional] This input expects a 3D scene. If a scene is provided, the Transform controls in this node will apply to the entire scene provided.

Controls

This tab contains all parameters for the node.



Enabled

When the Enabled checkbox is selected, the directional light affects the scene. Clear the checkbox to turn off the light.

Color

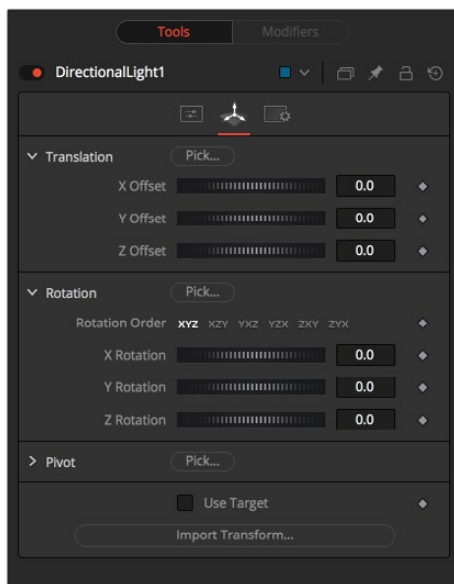
Use this standard Color control to set the color of the light.

Intensity


Use this slider to set the Intensity of the ambient light. A value of .2 indicates 20% percent light.

Transform Tab

The options that appear in this tab determine the position of the geometry created by this node. Since these controls are identical on all nodes that generate geometry, these controls are fully described in the Common 3D Controls section of this documentation.



Point Light [3PL]



A Point Light is a light source with a clear position in space that emits light in all directions. A light bulb is a point light, as is the sun, although light from the sun can appear to be ambient due to scattering from the atmosphere.


This light shows an onscreen widget, although only the position of the widget affects the light. Since the light is a 360-degree source, the rotation of the widget has no meaning. Unlike both ambient and directional lights, a point light may fall off with distance.

External Inputs

PointLight.SceneInput

[orange, optional] This input expects a 3D scene. If a scene is provided, the Transform controls in this node will apply to the entire scene provided.

Controls



Enabled

When the Enabled checkbox is selected, the point light affects the scene. Clear the checkbox to turn off the light.

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 .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 either to Linear or Quadratic modes.

Transform Tab

The options that appear in this tab determine the position of the geometry created by this node. Since these controls are identical on all nodes that generate geometry, these controls are fully described in the Common 3D Controls section of this documentation.



Spot Light [3SL]



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 will recognize the spotlight as being very similar to standard luminaries that are used in live productions. This is the only type of light capable of casting shadows.

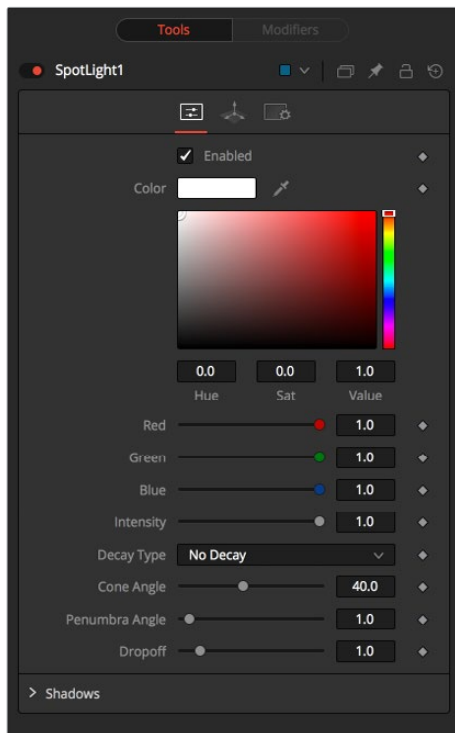
External Inputs

SpotLight.SceneInput

[orange, optional] This input expects a 3D scene. If a scene is provided, the Transform controls in this node will apply to the entire scene provided.

Controls

This tab contains all parameters for the node.



Enabled

When this checkbox is selected, the spotlight light affects the scene. Clear the checkbox to turn off the light.

Color

Use this standard Color control to set the color of the light.

Intensity

Use this slider to set the Intensity of the spotlight.

Decay Type

A spotlight defaults to No Falloff, meaning that its light has equal intensity on geometry regardless of the distance from the light to the geometry. To cause the intensity to fall off with distance, set the Decay type either to 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. See the Lighting and Shadows section of Chapter 60, "3D Compositing Basics," for details.

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 how opaque the shadow will be. A density of 1.0 will produce 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 will 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 will occasionally be Z-fighting, where the portions of the object that should be receiving the shadows render over top of the shadow instead. 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. Make adjustments to the Multiplicative Bias first, then fine tune the result using the Additive Bias control.

See the Multiplicative and Additive Bias section of Chapter 60, “3D Compositing Basics,” for examples and more information.

Force All Materials Non-Transmissive

Normally, a 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 result in 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 has three separate filtering methods available when rendering shadows which produce different effects.

NOTE: Shadows will 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 will 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 will take to render the shadows.
- **Variable:** The softness of shadow edges will grow 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.

Constant Softness

If the Softness is set to constant, then this slider will appear. 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. To be more precise, it controls how fast the shadow map filter size grows based upon the distance between shadow caster and receiver. Its effect is mediated by the values of the Min and Max Softness sliders.

Min Softness

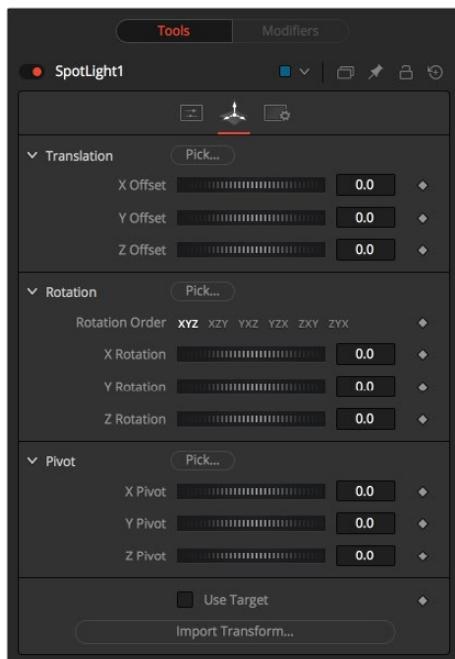
The Min Softness slider appears when the Softness is set to variable. This slider controls the Minimum Softness of the shadow. The closer the shadow is to the object casting the shadow, the sharper it will be up to the limit set by this slider.

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 will be up to the limit set by this slider.

Transform Tab

The options that appear in this tab determine the position of the geometry created by this node. Since these controls are identical on all nodes that generate geometry, these controls are fully described in the Common 3D Controls section of this documentation.



Chapter 28

3D Material Nodes

This chapter details the 3D Material nodes available in Fusion.

Contents

Blinn [3BI]	658
Channel Boolean Material [3BOL]	661
Cook Torrance [3CT]	664
Material Merge 3D [3MM]	667
Phong [3PH]	668
Reflect [3RR]	672
Stereo Mix [3SMM]	674
Ward [3WD]	675

Blinn [3BI]



The Blinn node is a basic illumination material that can be applied to geometry in the 3D scene. It describes how the object will respond to light, and provides a large number of texture map inputs to allow fine control over the diffuse, specular, and bumpmap components of the material.

The standard basic material provided in the Material 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 implemented by Fusion calculates the highlight as the dot product of the surface normal and the half angle vector between lightsource and viewer ($\text{dot}(\mathbf{N}, \mathbf{H})$). This may not always match the Blinn model illumination model used by other 3D applications.

External Inputs

Blinn.DiffuseTex

[orange, optional] This input will accept a 2D image or a 3D material to be used as a diffuse texture map.

Blinn.SpecularColorTex

[green, optional] This input will accept a 2D image or a 3D material to be used as a specular color texture map.

Blinn.SpecularIntensityTex

[magenta, optional] This input will accept a 2D image or a 3D material to be used as an intensity map for the materials specular highlights. When the input is a 2D image, the alpha channel is used to create the map while the color channels are discarded.

Blinn.SpecularExponentTex

[light blue, optional] This input will accept a 2D image or a 3D material to be used as a falloff map for the materials specular highlights. When the input is a 2D image, the alpha channel is used to create the map while the color channels are discarded.

Blinn.BumpmapTex

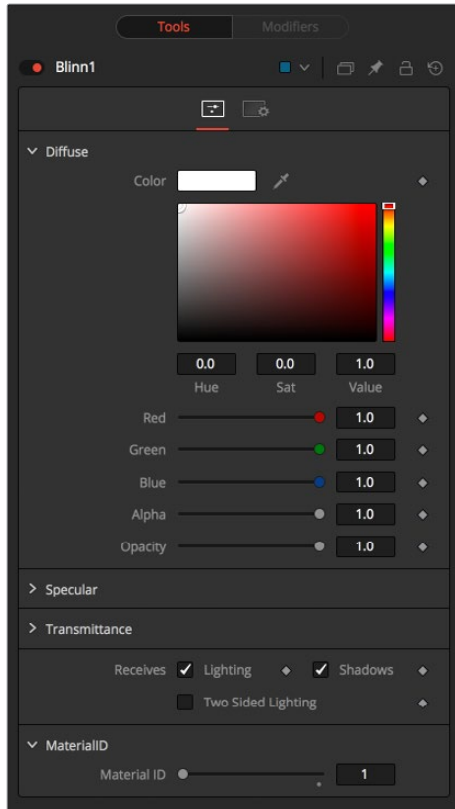
[white, optional] This input will accept a 2D image or a 3D material, then uses the RGB information as texture-space normals.

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 the Option (macOS) or Alt (Windows) key down 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 menu listing all of the inputs provided by the node will appear. 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.

Controls



Diffuse

Diffuse describes the base surface characteristics without any additional effects like reflections or specular highlights. In addition to 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 of any object the material is applied to 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 will decrease 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 will cast 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 for adjusting both opacity and transmittance separately. This might be a bit counter-intuitive to those 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 passed 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 pass 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. This allows for “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 is 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 will result 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 will always be 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 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 does arises because Fusion does not cull back-facing polygons by default. If you revolve around a one-sided plane in Fusion, you will 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, which is counter-intuitive. If you view from the frontside a transparent two-sided surface illuminated from the backside, it will look unlit.

Material ID

This slider sets the numeric identifier assigned to this material. This value will be rendered into the MatID auxiliary channel if the according option is enabled in the renderer.

Channel Boolean Material [3BOL]



The Channel Boolean Material 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.

External Inputs

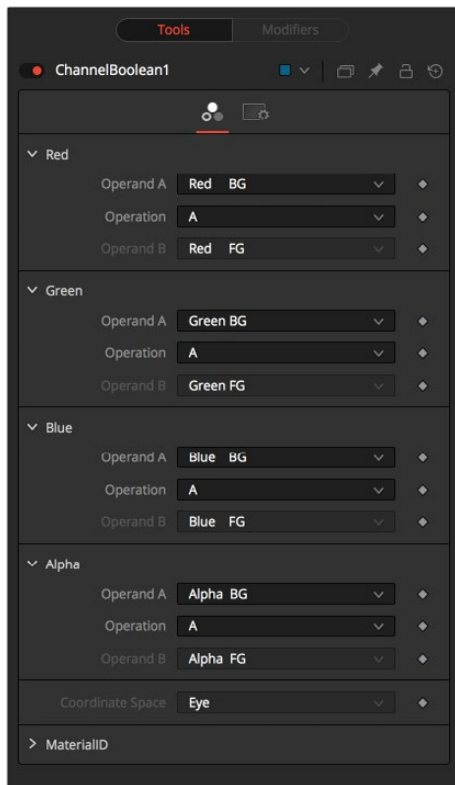
ChannelBooleanMaterial.BackgroundMaterial

[orange, optional] This input will accept a 2D image or a 3D material.

ChannelBooleanMaterial.ForegroundMaterial

[green, optional] This input will accept a 2D image or a 3D material.

Controls



Operand A/B

The Operand menus, one for each output RGBA channel, allow the user to set the desired input information for the according 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, 1 or 0.5.
- **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 according 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 according 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**
Set 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 will be rendered into the MatID auxiliary channel if the according option is enabled in the renderer.

Cook Torrance [3CT]



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.

External Inputs

CookTorrance.DiffuseTex

[orange, optional] This input will accept a 2D image or a 3D material to be used as a diffuse texture map.

CookTorrance.SpecularColorTex

[green, optional] This input will accept a 2D image or a 3D material to be used as a specular color texture map.

CookTorrance.SpecularIntensityTex

[magenta, optional] This input will accept a 2D image or a 3D material to be used as an intensity map for the materials specular highlights. When the input is a 2D image, the alpha channel is used to create the map while the color channels are discarded.

CookTorrance.SpecularRoughnessTex

[light blue, optional] This input will accept 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.

CookTorrance.SpecularRefractiveIndexTex

[white, optional] This input will accept a 2D image or a 3D material to be used as a map for modifying the roughness of the specular refractive index. The alpha of the texture map is multiplied by the value of the refractive index.

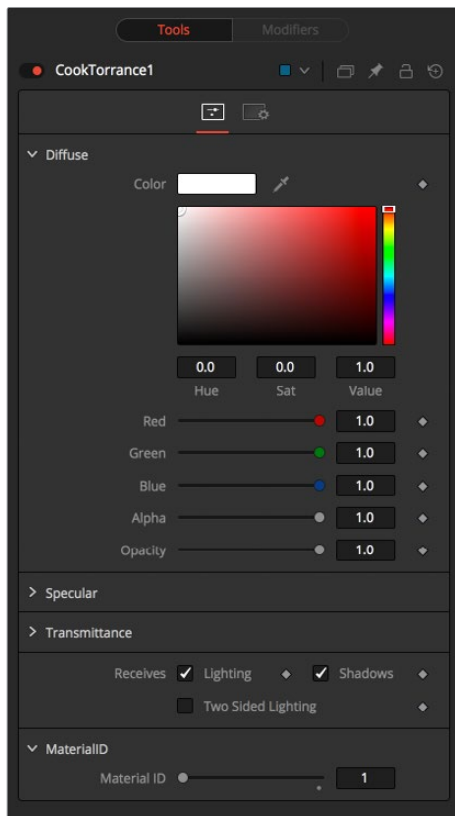
CookTorrance.BumpmapTex

[white, optional] This input will accept a 2D image or a 3D material, then 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 the Option (macOS) or Alt (Windows) key down 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 menu listing all of the inputs provided by the node will appear. Click on the desired input to complete the connection.

Controls



Diffuse

Diffuse describes the base surface characteristics without any additional effects like reflections or specular highlights. In addition to 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 of any object the material is applied to 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 will decrease 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 will add Fresnel calculations to the materials illumination model. This will provide 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 will cast 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 for adjusting both opacity and transmittance separately. This might be a bit counter-intuitive to those 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 passed 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 pass 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. This allows for '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 is 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 will result 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 will always be 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 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 counter-intuitive. If you view from the frontside a transparent two-sided surface illuminated from the backside, it will look unlit.

Material ID

This slider sets the numeric identifier assigned to this material. This value will be rendered into the MatID auxiliary channel if the according option is enabled in the renderer.

Material Merge 3D [3MM]



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 (Bumpmap, Reflection) to create complex shader networks.

The node also provides a mechanism for assigning a new material identifier to the combined material.

External Inputs

MtlMerge3D.BackgroundMaterial


[orange, required] This input will accept a 2D image or a 3D material to be used as the background material. A 2D image will be treated as a diffuse texture map in the basic shading model.

MtlMerge3D.ForegroundMaterial

[green, optional] This input will accept a 2D image or a 3D material to be used as the foreground material.

A 2D image will be treated as a diffuse texture map in the basic shading model.

Controls



Blend

The Blend behavior of the Material Merge is similar to the Dissolve (DX) node for images. The two materials/textures are mixed together using the value of the slider to determine the percentage each input contributes. While the background and foreground inputs can be a 2D image rather than a material, the output of this node will always be a material.

Unlike the 2D Dissolve node, both foreground and background inputs are required.

MaterialID

This slider sets the numeric identifier assigned to the resulting material. This value will be rendered into the MatID auxiliary channel if the according option is enabled in the renderer.

Phong [3PH]



The Phong node is a basic illumination material that can be applied to geometry in the 3D scene. It describes how the object will respond to light and provides a large number of texture map inputs to allow fine control over the diffuse, specular, and bumpmap 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.

External Inputs

Phong.DiffuseTex

[orange, optional] This input will accept a 2D image or a 3D material to be used as a diffuse texture map.

Phong.SpecularColorTex

[green, optional] This input will accept a 2D image or a 3D material to be used as a specular color texture map.

Phong.SpecularIntensityTex

[magenta, optional] This input will accept a 2D image or a 3D material to be used as an intensity 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.

Phong.SpecularExponentTex

[light blue, optional] This input will accept 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.

Phong.BumpmapTex

[white, optional] This input will accept a 2D image or a 3D material, then 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 the Option or Alt key down 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 menu listing all of the inputs provided by the node will appear. Click on the desired input to complete the connection.

Controls



Diffuse

Diffuse describes the base surface characteristics without any additional effects like reflections or specular highlights. In addition to 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 of any object the material is applied to 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 will decrease 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 will cast 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 for adjusting both opacity and transmittance separately. This might be a bit counter-intuitive to those 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 passed 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 pass 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. This allows for “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 is 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 will result 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 will always be 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 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 does arises because Fusion does not cull back-facing polygons by default. If you revolve around a one-sided plane in Fusion, you will 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, which is counter-intuitive. If you view from the frontside a transparent two-sided surface illuminated from the backside, it will look unlit.

Material ID

This slider sets the numeric identifier assigned to this material. This value will be rendered into the MatID auxiliary channel if the according option is enabled in the renderer.

Reflect [3RR]



The Reflection 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.

The Reflection node is usually combined with a standard lighting material such as Blinn, Cook Torrance, Phong or Ward by connecting the output of that node to the Reflection node's background material input. Then a reflection texture is connected to the reflection texture input. This can be a 2D image, but is more frequently an environmental map created by the Sphere Map or Cube Map nodes.

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. 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 but rather the infinite environment map). It also means that if you use the same cubemap on multiple objects in the scene, those objects will not inter-reflect each other (e.g., two neighboring object would not reflect each other). If you want objects to reflect each other you need to render a cubemap for each.

For more information see Reflections and Refractions in this manual.

External Inputs

Reflect.BackgroundMaterial

[orange, optional] This input expects a 2D image or a 3D material. If a 2D image is provided, the node will treat it as a diffuse texture map applied to a basic material.

Reflect.Reflection.ReflectionTex

[green, optional] This input expects a 2D image or a 3D material. The RGB channels are used as the reflection texture, and the alpha is ignored.

Reflect.Reflection.ReflectionIntensityTex

[magenta, optional] This input expects a 2D image or a 3D material. The alpha channel of the texture is multiplied by the intensity of the reflection.

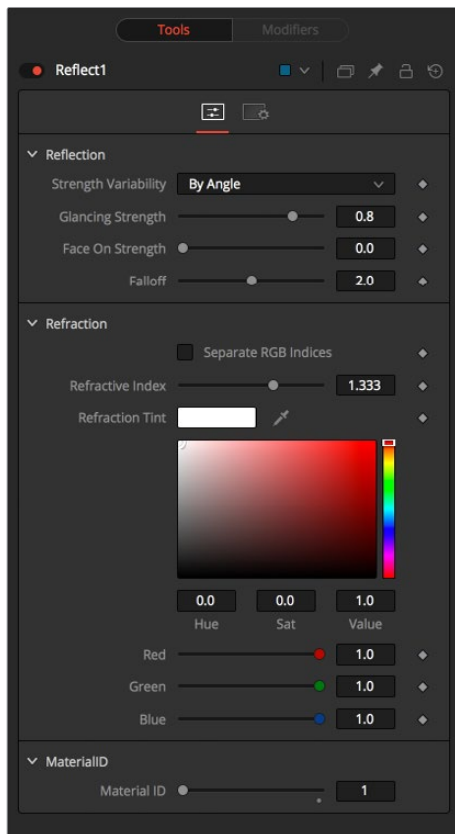
Reflect.Refraction.RefractionTex

[light blue, optional] This input expects a 2D image or a 3D material. The RGB channels are used as the refraction texture.

Reflect.BumpmapTex

[white, optional] This input will accept a 2D image or a 3D material, then uses the RGB information as texture-space normals.

Controls



Reflection

Reflection Strength Variability

This multi-button control can be set to Constant or By Angle for varying the reflection intensity, according to the relative surface orientation to the viewer. The following three controls are only visible 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 to be analogous to applying gamma correction to a gradient between the Face On and Glancing values.

Constant Strength

[Constant Angle] This control is only visible when the reflection strength variability is set to Constant. In this case, the intensity of the reflection is constant regardless of 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 will be hidden and three sliders for adjusting the refraction index of the Red, Green and Blue channels will appear in its place. This allows for simulation of the spectral refraction effects commonly seen in thick imperfect glass, for example.

Refraction Index

This slider controls how strongly the environment map will be 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.

Stereo Mix [3SMM]



External Inputs

StereoMix.LeftMaterial


[orange, required] This input will accept 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 will be converted to a diffuse texture map using the basic material type.

StereoMix.RightMaterial

[green, required] This input will accept 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 will be converted to a diffuse texture map using the basic material type.

While the inputs can be either 2D images or 3D materials, the output will always be a material.

Controls



Swap

This option will swap both inputs of the node.

Material ID

This slider sets the numeric identifier assigned to this material. This value will be rendered into the MatID auxiliary channel if the according option is enabled in the renderer.

Ward [3WD]



The Ward node is a basic illumination material that can be applied to geometry in the 3D scene. It describes how the object will respond to light, and provides a large number of texture map inputs to allow fine control over the diffuse, specular and bumpmap components of the material.

In particular, the Ward node is ideal for simulating brushed metal surfaces, as the highlight can be elongated in along the U or V directions of the mapping co-ordinates. 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.

External Inputs

Ward.DiffuseTexture

[orange, optional] This input will accept a 2D image or a 3D material to be used as a specular color texture map.

Ward.SpecularColorTexture

[green, optional] This input will accept a 2D image or a 3D material to be used as a specular color texture map.

Ward.SpecularIntensityTexture

[magenta, optional] This input will accept a 2D image or a 3D material to be used as an intensity 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.

Ward.SpreadUTexture

[light blue, optional] This input will accept a 2D image or a 3D material. The value of the Spread U option in the node's controls will be multiplied against the pixel values in the material's alpha channel.

Ward.SpreadVTexture

[white, optional] This input will accept a 2D image or a 3D material. The value of the Spread V option in the node's controls will be multiplied against the pixel values in the material's alpha channel.


Ward.BumpmapTexture

[white, optional] This input will accept a 2D image or a 3D material, then 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 the Option or Alt key down 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 menu listing all of the inputs provided by the node will appear. Click on the desired input to complete the connection.

Controls



Diffuse

Diffuse describes the base surface characteristics without any additional effects like reflections or specular highlights. In addition to 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 of any object the material is applied to 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 will decrease 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 appears 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 will cast 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 for adjusting both opacity and transmittance separately. This might be a bit counter-intuitive to those 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 passed 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 pass 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. This allows for “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 is 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 will result 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 will always be 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 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 does arises because Fusion does not cull backfacing polygons by default. If you revolve around a one-sided plane in Fusion you will 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 counter-intuitive. If you view from the frontside a transparent two-sided surface illuminated from the backside, it will look unlit.

Material ID

This slider sets the numeric identifier assigned to this material. This value will be rendered into the MatID auxiliary channel if the according option is enabled in the renderer.

Chapter 29


3D Texture Nodes

This chapter details the 3D Texture nodes available in Fusion.

Contents

BumpMap [3BU]	682
Catcher [3CA]	684
CubeMap [3CU]	686
Falloff [3FA]	688
Fast Noise Texture [3FN]	690
Gradient 3D [3GD]	692
Sphere Map [3SPM]	694
Texture [TXR]	696
Texture Transform [3TX]	697

BumpMap [3BU]



The BumpMap node either converts a grayscale (height map) image into a bump map or takes a bump map created by the Create BumpMap node directly. The node outputs a material.

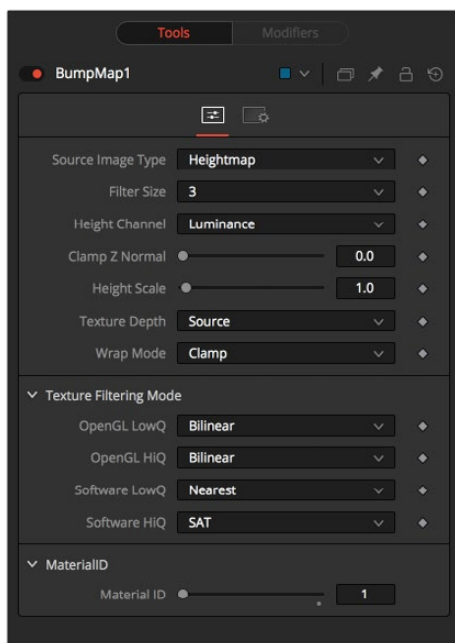
External Inputs

Bumpmap.ImageInput (white)

Receives the RGBA channels from an image for the bump calculation or an existing bump map.

Controls

This tab contains all parameters for the node.



The Source Image is a...

Toggle between HeightMap, which will create a bump map similar to the CreateBumpMap node, and BumpMap, which expects a bump map created by the CreateBumpMap node.

Filter Size

The process of generating the bump information is basically a Custom Filter. This multi-button control sets the filter size.

Extract height information from...

Set the channel from where to extract the grayscale information.

Clamp Normal.Z

Clips the lower values of the Blue channel in the resulting bump texture.

Filter Wrap Mode

Basically “wraps” the image at the borders, so the filter produces correct result when using seamless tileable textures.

Height Scale


Changes the contrast of the resulting values in the bump map. Increasing this value yields in a more visible bump map.

Bump map Texture Depth


Optionally converts the resulting bump texture into the desired bit depth.

Notes on Bump maps

There is some confusion of terminology with bumpmapping, depending on which papers/books/people you are reading/talking to. 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)

Normal 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 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 purpose of the Catcher node it helps to understand the difference between light-based projections and texture-based projections. A light-based projection 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 matte out the windows. This will make 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 matted 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 from Light or Ambient Light mode to Texture mode, Fusion knows to apply the projected image as a texture map. When using this technique, the windows would become transparent, and it would be possible to see the geometry behind the window.

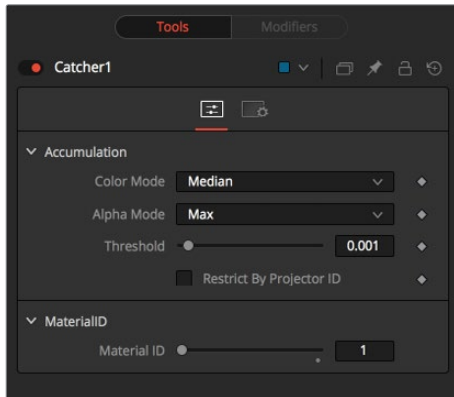
The main advantages of this approach over light projection are that the Catcher can be used to project alpha onto an object, and it doesn't require lighting to be enabled. Another advantage is that the Catcher is not restricted to the diffuse input of a material, making it possible to project specular 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 will simply make the object transparent and invisible.

External Inputs

This material node does not have any external inputs.

Controls



Enable

Use this checkbox to enable or disable the node.

Color Accumulation Mode

The Color Accumulation mode is used to control how the Catcher will combine the light from multiple projectors. It will have no effect on the results when only one projector is in the scene. This control is designed to work with the software renderer, and will have no effect when using the OpenGL renderer.

Alpha Accumulation Mode

The Alpha Accumulation mode is used to control how the Catcher will combine the alpha channels from multiple projectors. It will have no effect on the results when only one projector is in the scene. This control is designed to work with the software renderer, and will have no effect when using the OpenGL renderer.

Accumulation Threshold

The Accumulation 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 will only receive light from projectors with a matching ID. Projectors with a different ID will be ignored.

Material ID

This slider sets the numeric identifier assigned to this material. This value will be rendered into the MatID auxiliary channel if the according option is enabled in the renderer.


CubeMap [3CU]



The CubeMap 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 6 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 CubeMap node is typically used to produce environment maps for distant areas (such as skies or horizons) or reflection and refraction maps.



External Inputs


CubeMap.CrossImage

[white, required] This input is only visible when the node is set to the Vertical Cross or Horizontal Cross orientation. It expects a 2D image.

CubeMap.[DIRECTION]

[white, required] These six inputs are only visible when the node is set to the Separate Images orientation mode. Each input expects an image aligned to match the left, right, up, down, front and back faces.

Controls



Orientation

This multi-button control tells the node which type of input to expect 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 will be 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 will crop it down so it matches the appropriate 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 will crop it down so that matches the appropriate aspect ratio.

Coordinate System

This multi-button control sets the coordinate system 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.


Warn about bad dimensions

Selecting this checkbox will print a warning message into the console if the dimensions of the image provided do not meet the requirements of the selected orientation mode.


Material ID

This slider sets the numeric identifier assigned to this material. This value will be rendered into the MatID auxiliary channel if the according option is enabled in the renderer.

Falloff [3FA]



The Falloff node blends two materials together based on the incidence angle between the object the material is applied to 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.



External Inputs

Falloff.FaceOnMaterial

[orange, optional] This input expects a 2D image or a 3D material. If a 2D image is provided, it will be turned into a diffuse texture map using the basic material shader.

Falloff.GlancingMaterial

[green, optional] This input expects a 2D image or a 3D material. If a 2D image is provided, it will be turned into a diffuse texture map using the basic material shader.

While the inputs for this node can be images, the output will always be a material.

Controls



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

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.

Face On Opacity

Reducing the material's opacity will decrease the color and alpha values of the Face On material, making the material transparent.

Glancing Color

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.

Glancing Opacity

Reducing the material's opacity will decrease 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 will be rendered into the MatID auxiliary channel if the according option is enabled in the renderer.

Fast Noise Texture [3FN]



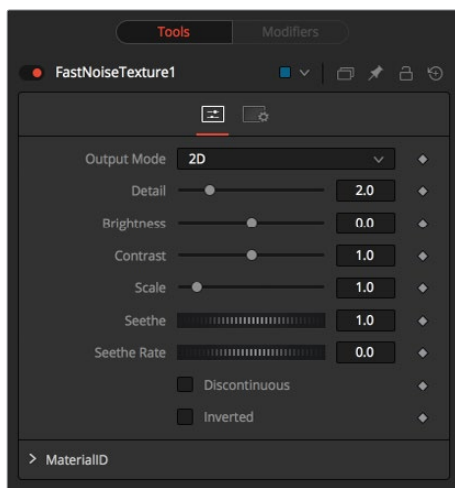
The Fast Noise Texture node is the procedural resolution-independent version of the 2D Fast Noise node. It creates a noise texture directly as a material for usage with 3D nodes. It offers a 3D volumetric mode for creating seamless textures in conjunction with nodes providing UVW texture coordinates (like the UV Map node set to XYZtoUVW or Camera).

External Inputs

FastNoiseTexture.SourceMaterial

[orange, required] This input will accept a 2D image or a 3D material. This is modulated by the noise pattern.

Controls



Output Mode

- **2D:** Calculates the noise texture based on 2D texture coordinates (UV). This mode allows smoothly varying the noise pattern.
- **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. Does not support animation of the noise pattern.

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.

Invert

Select this checkbox to Invert the noise, creating a negative image of the original pattern. This is most effective when Discontinuous is also enabled.

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 feature scale 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. 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.

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.

Material ID

This slider sets the numeric identifier assigned to this material. This value will be rendered into the MatID auxiliary channel if the according option is enabled in the renderer.

Gradient 3D [3GD]



With the Gradient 3D node it is possible to texture objects with a variety of gradient types. It offers many of the controls of 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

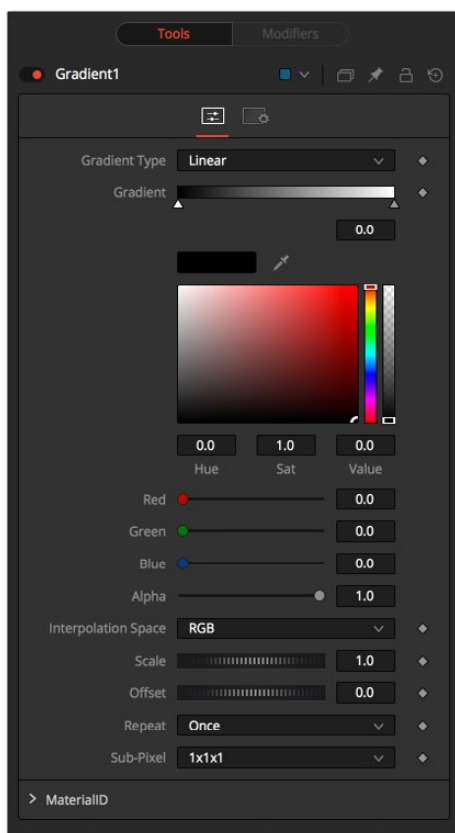
The Texture Transform node can be used to adjust the mapping per pixel.

UVMaP 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.

Working with the Gradient node may be a bit confusing at first. The gradient defaults to a linear gradient that goes from -1 to +1 along the Z-axis. All primitives in Fusion (Shape 3D) can output a third texture coordinate for UVW mapping.


Controls



Gradient Type

Determines the type of or pattern used for the gradient.

- **Linear:** A simple linear gradient.
- **Reflect:** Based on the Linear mode, this gradient will be mirrored at the middle of the textured range.
- **Square:** The gradient is applied using a square pattern.
- **Cross:** Similar to the Reflect mode, but it will use two axes to apply the gradient.
- **Radial:** The Radial mode uses a circular pattern to apply the gradient.



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 own 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 linear interpolated from point to point in RGB color space by default. This can result in unwanted colors sometimes. Choosing another color space may provide a better result.

Offset

Allows panning through the gradient.

Repeat

Defines how the left and right border of the gradient is treated.



Gradients set to Once, Repeat, and Ping Pong from top to bottom respectively and shifting the gradient to the left.

- **Once:** When using the Gradient Offset control to shift the gradient, the border colors will keep their values. Shifting the default gradient to the left will result in a white border on the left, while shifting it to the right will result in a black border on the right.
- **Repeat:** When using the Gradient Offset control to shift the gradient, the border colors will be wrapped around. Shifting the default gradient to the left will result in a sharp jump from white to black, while shifting it to the right will result 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 will result in the edge fading from white back to black, while shifting it to the right will result in the edge fading from black back to white.

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 will be rendered into the MatID auxiliary channel if the according option is enabled in the renderer.

Sphere Map [3SPM]

The Sphere Map node allows the creation of a spherical texture map from an image. The input image should represent the texture information in a longitude/latitude format, where the X-axis represents 0–360 degrees longitude and the Y-axis represents –90 to +90 degrees latitude.

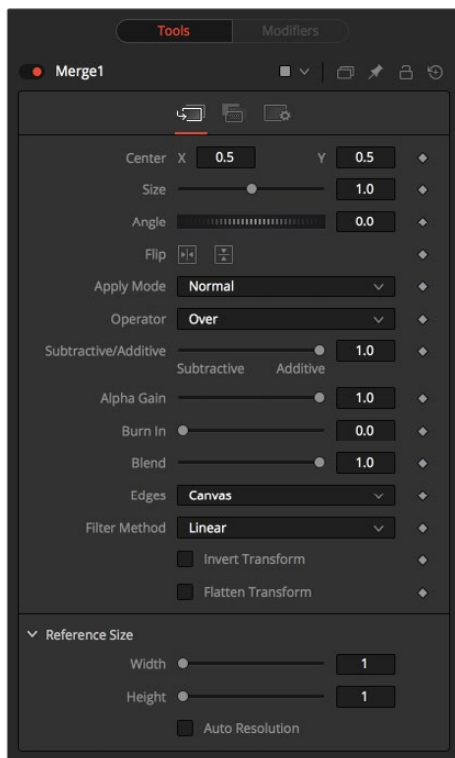


External Inputs

SphereMap.ImageImage

[white, required] Receives the RGBA channels from an image output.

Controls



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. In other words, 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

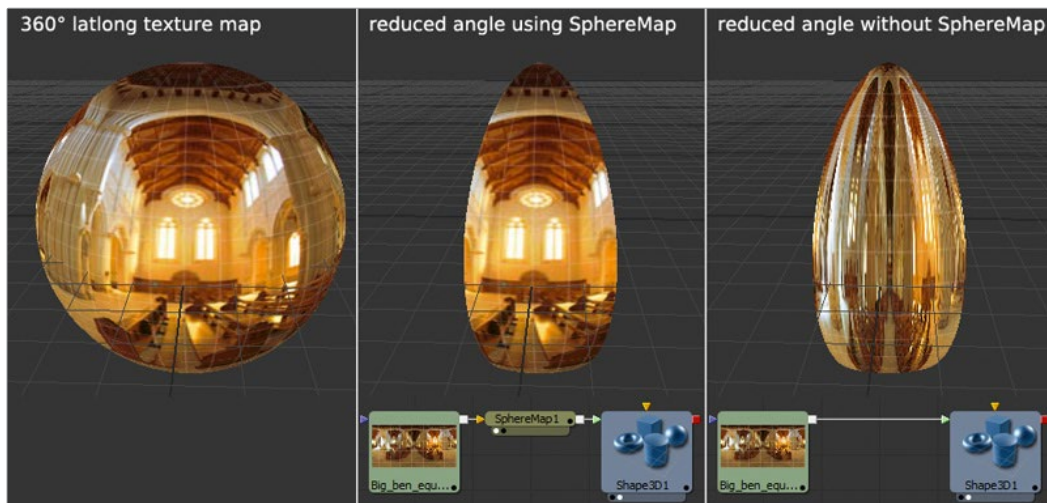
This slider sets the numeric identifier assigned to this material. This value will be rendered into the MatID auxiliary channel if the according option is enabled in the renderer.

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 fit onto the sphere and the poles will display clamped edges.
- **2 * width < height:** The height is fit onto the sphere and there will be clamping about the 0 degree longitude line.


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 will be squashed. When using the Sphere Map node, the texture will be cropped. Compare:



NOTE: If you pipe the texture directly into the sphere, it will also be mirrored horizontally. You can “fix” this by using a Transform node first.

Texture [TXR]



The Texture node can control the texture mapping of elements in a rendered image. The Texture-map image (connected to the green input) can be wrapped around objects to replace

the current texture. 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 will 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.

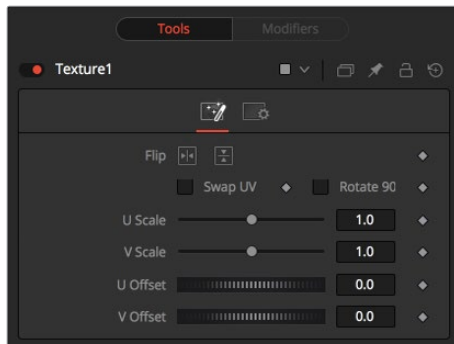
For more information, see Chapter 56, “Understanding Image Channels and Noise Processing.”

External Inputs

CreateTexture.ImageInput

[orange, required] This input expects a 2D image.

Controls



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 are 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.


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.

Texture Transform [3TX]



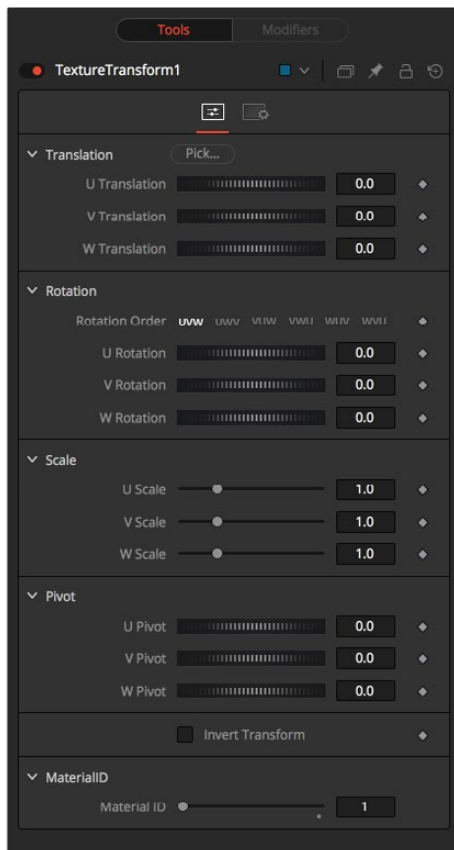
The TextureTransform node can be used to translate, rotate, and scale image textures on the input material. While the input can also be an image, the output will always be a material.

External Inputs

TextureTransform.MaterialInput

[orange, optional] This input expects a 2D image or 3D material.

Controls



NOTE: Not all Wrap modes will be supported by all graphics cards.

Translation

U, V, W Translation

These sliders will shift the texture along U, V, and W axes.

Rotation

Rotation Order

Use these buttons to set the order in which the rotation is applied.

U, V, W Rotation

In conjunction with the Rotation Order, these settings define the rotation around the UVW axis.

Scale

U, V, W Scale

Scales the texture along the according UVW axis.

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 will be rendered into the MatID auxiliary channel if the according option is enabled in the renderer.

Chapter 30


Blur Nodes

This chapter details the Blur nodes available in Fusion.

Contents

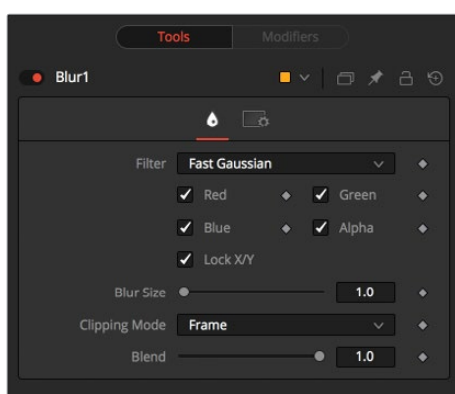
Blur [BLUR]	701
Defocus [DFO]	703
Directional Blur [DRBL]	705
Glow [GLO]	706
Sharpen [SHRP]	709
Soft Glow [SGlo]	711
Unsharp Mask [USM]	713
VariBlur [VBL]	714
Vector Motion Blur [VBL]	715

Blur [BLUR]



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.

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 give rise to visible ringing around the edges in certain extreme cases. This ringing only appears 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 viable alternative.

Filter Type

The Filter Type button array allows for the selection of the filter to be applied to the image.

- **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.
- **Soften:** Soften applies a general softening filter effect. This filter method is slower than the Gaussian filter and produces a lower quality result. It is included for compatibility with older node trees only.
- **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. 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 the checkboxes beside each channel to make them active or inactive.

NOTE: This is not the same as the RGBA checkboxes found under the common controls. The Blur node takes these selections into account before it processes the image, so deselecting a channel will cause 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 checked 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 sets the mode used to handle the edges of the image 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 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.


Blend

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.

Examples

This is a comparison of Blur filters visualized as “cross-sections” of a filtered edge. As you can see, Box will create a linear ramp, while Bartlett creates a somewhat smoother ramp. Multi-box and Gaussian result in even smoother ramps that are virtually indistinguishable unless you zoom in really close on the slopes. As mentioned above, Gaussian will overshoot slightly and may result in negative values if used on floating-point images.

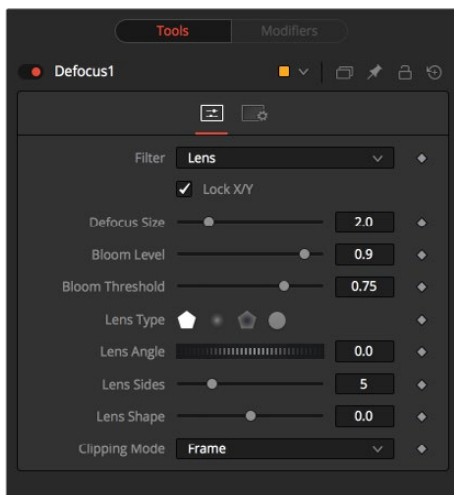


Defocus [DFO]



The Defocus node simulates the effects of an out-of focus camera lens, including blooming and image flaring. It provides a fast but relatively inaccurate Gaussian mode, as well as a more realistic but much slower Lens mode.

Controls



Filter

Use this menu to select the exact method applied to create the defocus. Gaussian applies a fairly simplistic effect, while Lens mode will create a much more realistic defocus. Lens mode will take 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.

Lens Type

The basic shape that is 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. Due to the round nature of a circle, this slider will have no visible effect when the Lens Type is set to Circle.

Lens Sides

Defines how many sides the NGon shapes will have. Best visible with NGon lens types. Due to the round nature of a circle, this slider will have no visible effect when the Lens Type is set to Circle.

Lens Shape

Defines how pointed the NGons are. Higher values will create a more pointed, starry look. Lower values create smoother NGons. Best visible with NGon lens types and Lens Sides between 5 and 10. Due to the round nature of a circle, this slider will have no visible effect when the Lens Type is set to Circle.

Clipping Mode

This option sets the mode used to handle the edges of the image 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 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

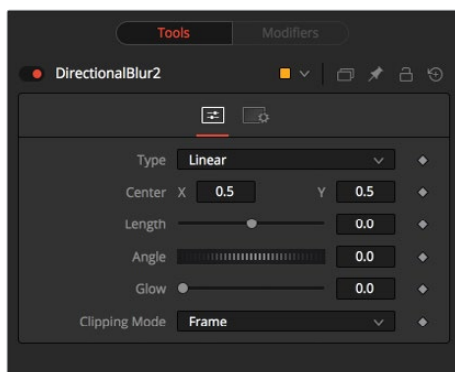
Setting this option to None will not perform any source image clipping at all. This means that any data required to process the node’s effect that would normally be outside the upstream DoD will be treated as black/transparent.

Directional Blur [DRBL]



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).

Controls



Type

This button array is used to select the Type of directional blur to be supplied 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 will create 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 crosshair affects the Radial and Zoom Motion blur types only. It is used to calculate the position from where the blurring effect starts.

Length

Length adjusts the strength and heading of the effect. Values lower than zero cause blurs to oppose the angle control. Values greater than the slider maximum may be typed into the slider's edit box.

Angle

In both Linear modes, this control will modify the direction of the directional blur. In the Radial and Zoom modes, the effect will be similar to that of the camera spinning while looking at the same spot. If the setting of the length slider is other than zero, the effect will create a whirlpool effect.

Glow

This will add 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 sets the mode used to handle the edges of the image 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 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.

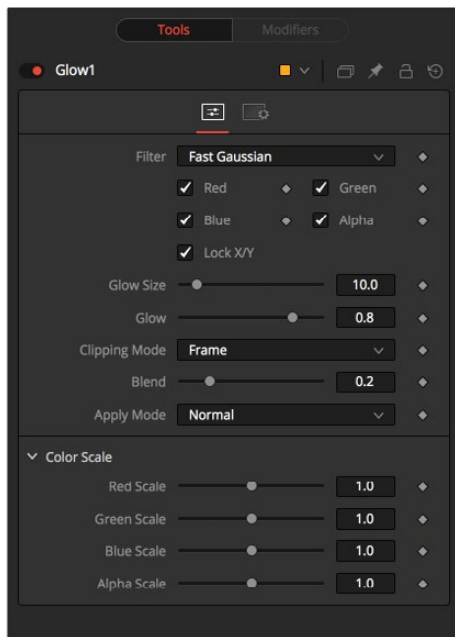
Glow [GLO]



A Glow is basically created by blurring an image, then brightening the blurred result and mixing it back with the original.

The Glow node provides a more convenient way to accomplish this effect, as well as a variety of variations on the theme. For example, a Bartlett glow is a high quality glow with smoother drop off, however, it is more processor-intensive at larger sizes.

Controls



Filter

Use these buttons 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.

Blend

Blend adds a non-linear 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)

The filter defaults to operating on R, G, B, and A channels. Selective channel filtering is possible by clicking the checkboxes beside each channel to make them active or inactive.

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 will cause 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 will be locked. Otherwise, separate amounts of blur 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 result in a smoother distribution of the effect but also increase render times. It's good to find that thin 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 sets the mode used to handle the edges of the image 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 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.

Blend

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.

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 will appear. Pixels in the glowed areas with values below the low value will be pushed to black. Pixels with values greater than high will be pushed to white.
- **High-Low Range Control:** Only available in Threshold mode. Pixels in the glowed areas with values below the low value will be pushed to black. Pixels with values greater than high will be pushed to white.

Color Scale (RGBA)

These Scale sliders are normally a reveal control labeled Color Scale. They can be used to adjust the amount of glow applied to each color channel individually, thereby tinting the glow.


Glow Pre Mask

The Glow node supports pre-masking using the glow mask. 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.

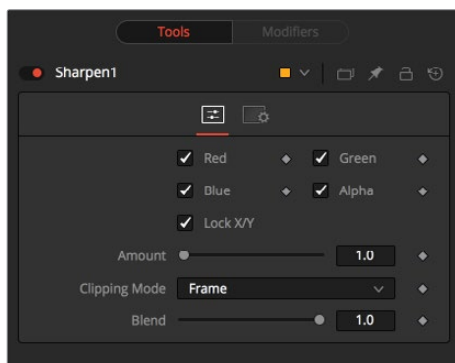
To apply a glow mask, select Glow Mask instead of Effect Mask from the contextual menu when a Glow node is active. Glow masks are identical to Effects masks in every other respect.

Sharpen [SHRP]



The Sharpen node uses a convolution filter to enhance detail in an image.

Controls



Color Channels (RGBA)

The filter defaults to operating on R, G, B, and A channels. Selective channel filtering is possible by clicking the checkboxes beside each channel to make them active or inactive.

This is not the same as the RGBA checkboxes found under the common controls. The node takes these controls into account before it processes, so deselecting a channel will cause 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 sets the mode used to handle the edges of the image 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 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.

Blend

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.

Soft Glow [SGlo]



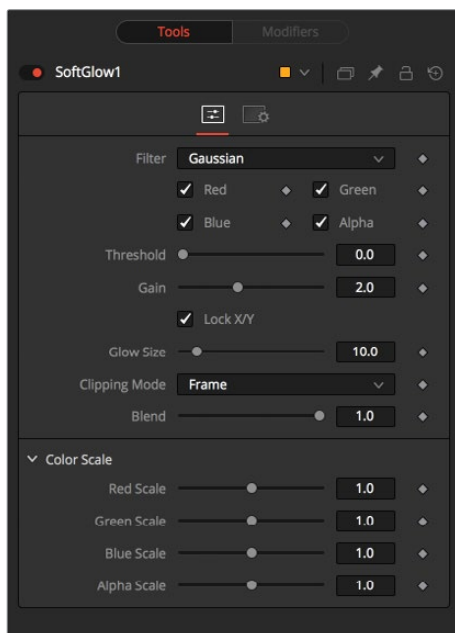
The Soft Glow node is similar to the Glow node but performs additional processing of the image to create a much softer, more natural glow.

This node is perfect for atmospheric haze around planets, skin tones, and simulating dream-like environments.

NOTE: The Glow node supports pre-masking using the Glow Pre Mask input on the node. A pre-mask limits the image before applying the glow. The glow is then combined with the original image. This is different from a regular effect mask that limits 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.

Controls



Filter

Use these buttons 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 checkboxes beside each channel to make them active or inactive.

This is not the same as the RGBA checkboxes found under the common controls. The node takes these controls into account before it processes, so deselecting a channel will cause 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 affect 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 will be locked. Otherwise, separate amounts of glow may be applied to each axis of the image.

Glow Size

Amount 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 result in a smoother distribution of the effect but also increase render times. It's good to find that thin line between desired quality and acceptable render times.

Clipping Mode

This option sets the mode used to handle the edges of the image when performing domain of definition rendering. This is profoundly important for node 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 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.


Blend

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.

Color Scale (RGBA)

These Scale sliders are normally a reveal control labeled Color Scale. They can be used to adjust the amount of glow applied to each color channel individually, thereby tinting the glow.


Unsharp Mask [USM]



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 far-away 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.

Controls



Color Channels (RGBA)

The filter defaults to operating on R, G, B, and A channels. Selective channel filtering is possible by clicking the checkboxes beside each channel to make them active or inactive.

This is not the same as the RGBA checkboxes found under the common controls. The node takes these controls into account before it processes, so deselecting a channel will cause 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 will be 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 will be identified as detail.

Gain

The Gain control adjusts how much gain is applied to pixels identified as detail by the mask. Higher values will create a sharper image.

Threshold Low and High

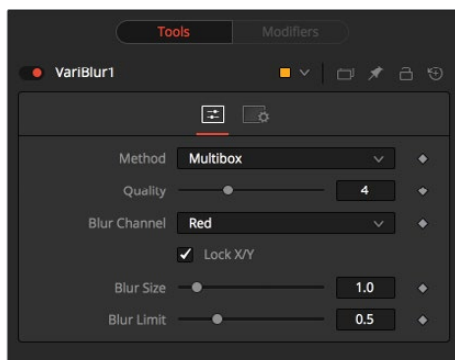
This range control determines the frequencies from the source image to be extracted. Raising low values will eliminate dark pixels from the comparison. Pixels above the high value will have the full effect of the gain applied.

VariBlur [VBL]



The VariBlur 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 cleaner results in many cases.

Controls



NOTE: The Blur Image input must be connected, or no blurring will be performed.

Method

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.

Multibox

Similar to Soften, this gives a better Gaussian approximation at higher Quality settings.

Defocus

Not really a true defocus, this gives 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. 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 will increase the overall amount of blur applied to each pixel. Those pixels where the Blur Image is black or non-existent will never be blurred, regardless of Blur Size.

Vector Motion Blur [VBL]



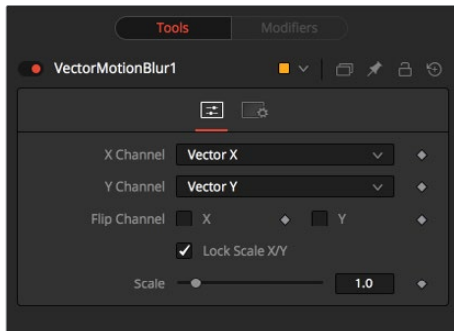
This node is used to create Directional blurs based on a Vector Channel.

This node will perform a 2D blur on the image, using a vector map produced by a 3D application. 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 the image in the case of 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.

Controls



X Channel

Use this multi-button array to choose which channel of the image will provide the vectors for the movement of the pixels along the X-axis.

Y Channel

Use this multi-button array to choose which channel of the image will provide the vectors for the movement of the pixels along the Y-axis.

Flip X Channel

This checkbox can be used to flip, or invert, the X-vectors. A value of 5 for a pixel in the X-vector channel would become -5 when this checkbox is selected.

Flip Y Channel

This checkbox can be used to flip, or invert, the Y vectors. A value of 5 for a pixel in the Y-vector channel would become -5 when this checkbox is selected.

Lock Scale X/Y

Selecting this checkbox will provide access to separate sliders for X and Y Scale. By default only a single Scale slider is provided.

Scale/Scale X

This slider will be labeled Scale if the Lock Scale X/Y checkbox is not selected, otherwise it will be labeled Scale X. The vector channel value for a pixel is 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.

Scale Y

This slider will only appear if the Lock Scale X/Y checkbox is selected. Otherwise, it will be hidden, and use the same value set in the Scale slider above.

Chapter 31


Color Nodes

This chapter details the Color nodes available in Fusion.

Contents

Auto Gain [AG]	719
Brightness Contrast [BC]	720
Channel Booleans [BOL]	722
Color Corrector [CC]	724
Color Curves [CCV]	733
Color Gain [CLR]	736
Color Matrix [CMX]	739
Color Space [CS]	742
Copy Aux [CPA]	744
Gamut [GMT]	747
Hue Curves [HCV]	749
OCIO CDL Transform [OCD]	750
OCIO ColorSpace [OCC]	753
OCIO FileTransform [OCF]	754
Set Canvas Color [SCV]	755
White Balance [WB]	755


Auto Gain [AG]



The Auto Gain node is used to automatically adjust the color range of an image so that the darkest and brightest pixels are set to user-selected values. By default, the darkest pixels in the image are pushed down to black, the brightest pixels are pushed to white and all of the pixels in between are stretched to cover the color 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 (though the Viewer's View Normalized Image button is generally more suitable for this).

Controls



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 will get suddenly brighter, as the remaining darker values are 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-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 is used to set the lowest and highest possible pixel value in the image. All color values in the image are rescaled to fit within this range.

Examples

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 to 0.5. This will cause the brightest pixels to be pushed down to 0.5 and the darkest pixels will get pushed to black. The remainder of the pixel values will be scaled between those limits.

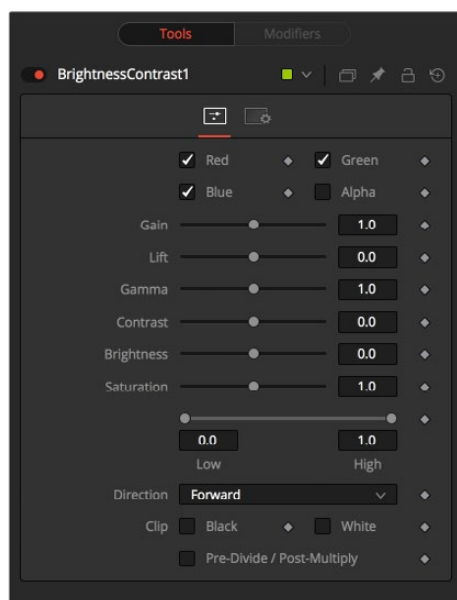
Brightness Contrast [BC]



The Brightness Contrast node is used to adjust the gain, brightness, contrast, gamma and saturation of an image. The order of the controls represents the order in which each operation is applied (for example, gamma is 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, make sure that your image is processed in 32bit floating point.

Controls



Color Channels (RGBA)

The filter defaults to operating on R, G, B, and A channels. Selective channel filtering is possible by clicking the checkboxes beside each channel to make them active or inactive.

This is not the same as the RGBA checkboxes found under the common controls. The node takes these controls into account before it processes, so deselecting a channel will cause 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.

Gain

The pixel values are multiplied by the value of this control. A Gain of 1.2 will make a pixel that is R0.5 G0.5 B0.4 into R0.6 G0.6, B0.48 (i.e., $0.4 * 1.2 = 0.48$). Gain affects higher values more than it affects lower values, so the effect will be strongest in the midrange and top range of the image.

Lift

While Gain basically 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 will make 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 will be strongest in the midrange and low range of the image.

Gamma

Values higher than 1.0 will raise the Gamma (mid gray), whereas lower values will decrease it. The effect of this node is not linear and existing black or white levels will not be affected at all. Pure gray colors will be affected the most.

Contrast

Contrast is the range of difference between the light to dark areas. Increasing the value of this slider will increase the contrast, pushing color from the midrange toward black and white. Reducing the contrast will cause 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 is added to the value of each pixel in the image. This control's affect on an image is linear so the effect will be applied identically to all pixels regardless of value.

Saturation

This control is used to increase or decrease the amount of Saturation in the image. A saturation of 0 has no color. All colors are grayscale.

Low and High

This range control is similar to the Gain control in some respects. If Low is anchored at 0.0 and the High value is reduced from 1.0, the effect is identical to increasing the gain. High values are multiplied by the inverse of the high value. (e.g., if high is 0.75, each pixel will be multiplied by 1/0.75 or 1.3333).

Leaving the high anchored at 1.0 and increasing the low is exactly 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 are used to 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 will 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 will cause 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 to prevent the creation of illegally additive images, particularly around the edges of a blue/green key or when working with 3D rendered objects.

Channel Booleans [BOL]

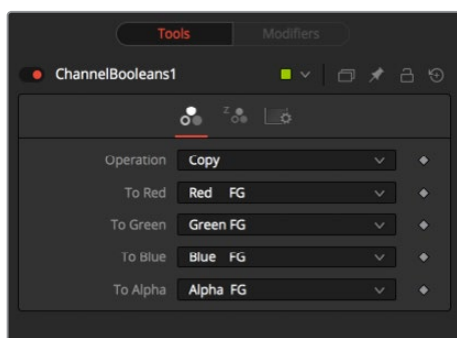


The Channel Booleans node can be used to apply a variety of mathematical and logical operations on the channels in an image. This node works by using one image's channels to modify another image's channels. If a foreground input is not available, selecting options that use color channels from the foreground will end up using the background input's color channels instead.

On the Color Channels Tab the Node Controls are Divided as Follows:

On the left side are target channels for the image piped into the Channel Booleans (background input). The drop down to the right lets you choose whether you want to modify the BG image with its own channels (suffix BG after list name) or with the channels from another image, which must be piped into the foreground input on the Channel Booleans node (suffix FG in the drop-down list).

Controls



Operation Type

This drop-down box is used to select the mathematical method applied to the selected channels. Its settings are as follows:

Copy

Copy the value from one color channel to another. For example, copy the foreground red channel into the alpha channel to create a matte.

Add

Add the color values from color channel to 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 will generally remove 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 will generally add 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 will generally flip bits from the color channel of the background image.

Multiply

Multiply 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.

Divide

Divide the values of a color channel. This will give the appearance of lightening the image as the values are scaled 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, Green, Blue, 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 will be used to produce the output channel.

The default setting simply 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.

Auxiliary Channels

There are several auxiliary channels. Use these menus to select a source for the auxiliary channels of the output image. See Chapter 74, “3D Nodes,” for further details on auxiliary channels.

Enable Extra Channels

When the Enable Extra Channels checkbox is selected, the Channel Boolean node will be able to output images with channels beyond the usual RGBA. The remaining controls in this tab will become active and can be used to copy data into the auxiliary channels.

Examples

To copy the alpha channel of one image to its own 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 can also be performed using the Matte Control node.

To combine any type of mask into an alpha for 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.”

Color Corrector [CC]



The Color Corrector node is a comprehensive color node with histograms, matching and equalization, hue shifting, tinting and color suppression. The Color Corrector has two image inputs. Connect the image to be corrected to the primary input. The secondary input can be used as a reference for histogram matching.

Overview

Controls in the Color Correction node are separated into four separate categories: colors, levels, histogram and suppress. Selecting one of the category buttons from the array at the top of the Correction tab will cause that category’s controls to appear. Each category is described in detail below.

Colors Mode



Master/Shadows/Midtones/Highlights

This array of buttons determines the range of colors affected by the controls in this tab. For example, when the Shadows range is selected, any color adjustments made will affect only the darker pixels of the image.

The selected state of this button 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 will 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 display 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 on the display, or by entering values in the text boxes to the right of the color wheel.

Hue

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 0.0 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.

The degree of hue shifting can be entered directly into the text control, or by placing the mouse above the outer ring of the color control and dragging the mouse up or down. The outer ring will always show the shifted colors compared to the original colors shown in the center of the control.

Saturation

The Saturation control is used to adjust the intensity of the color values. A saturation of 0 produces gray pixels without any chroma or color component, whereas a value of 1.0 produces no change in the chroma component of the input image. Higher values will generate over-saturated values with a high color component.

Saturation values can be entered directly into the text control, or by dragging the mouse to the left and right on the outer ring of the color wheel control.

Tint/Strength

The Tint control is used to tint an image or selected color range. The values in this control go from 0 to 1.0, which indicate the angle of the tint color on the color wheel.

A value of 0.25 would indicate 90 degrees, which would be midway between green and yellow on the color wheel.

The Strength control determines how much tint is applied to the selected range of colors.

The tinting is represented in the Color Wheel control by small circles that show the color and strength of the tint. The Highlight Ranges marker is a black outline of a circle. The Midtones and Shadows are represented by gray circles. The Master Tint Marker is also black, but it has a white M in the center to distinguish it from the others.

The mouse can position the marker for each range only when the appropriate range is selected. For example, the Highlight Marker cannot be moved when the Master range is selected.

Holding down the Command or Ctrl key while dragging this control will allow you to make finer adjustments by reducing the control's sensitivity to mouse movements. Holding the Shift key down will limit the movement of the marker to a single axis, allowing you to restrict the effect to either tint or strength.

Tint Mode

Fast/Full/Better

These three buttons are used to select the speed and quality of the algorithm used to apply the tint, 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 shown in the color wheel above. The slider makes it easier to make small adjustments to the value with the mouse.

Saturation

This slider is a clone of the Saturation control shown in the color wheel above. The slider makes it easier to make small adjustments to the value with the mouse.

RGB/Red/Green/Blue

These buttons are the same buttons seen in the Histogram, Color and Levels sections of the Color Corrector node. When the red channel is selected, the controls in this tab will affect the red channel only, and so on.

The controls are independent, so switching to blue will not remove or eliminate any changes made to red, green or master. The animation and adjustments made to each channel are separate. These buttons simply determine what controls to display.

Master RGB Contrast

Contrast is the range of difference between the light to dark areas. Increasing the value of this slider will increase the contrast, pushing color from the midrange toward black and white. Reducing the contrast will cause the colors in the image to move toward midrange, reducing the difference between the darkest and brightest pixels in the image.

Master RGB Gain

The pixel values are multiplied by the value of this control. A Gain of 1.2 will make a pixel that is R0.5 G0.5 B0.4 into R0.6 G0.6, B0.48 (i.e., $0.4 * 1.2 = 0.48$). Gain affects higher values more than it affects lower values, so the effect will be strongest in the midrange and top range of your image.

Master RGB Lift

While Gain basically 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 will make 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 will be strongest in the midrange and lowrange of the image.

Master RGB Gamma

Values higher than 1.0 will raise the Gamma (mid gray) while lower values will decrease it. The effect of this node is not linear, and existing black or white levels will not be affected at all. Pure grays will be affected the most.

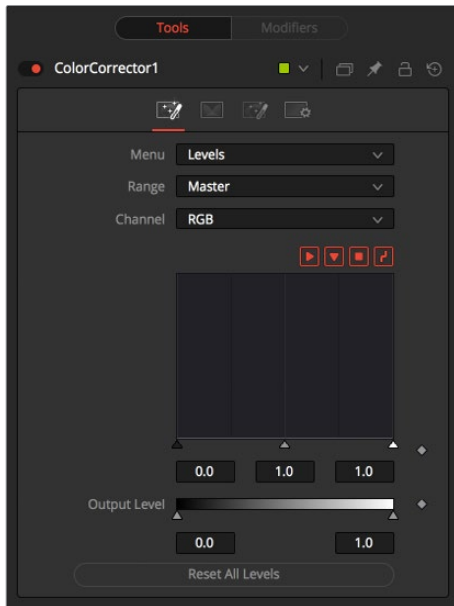
Master RGB Brightness

The value of the Brightness slider is added to the value of each pixel in your image. This control's affect on an image is linear, so the effect will be applied identically to all pixels regardless of value.

Reset All Color Changes

Selecting this button will return all color controls in this section to their default values.

Levels Mode



Master/Shadows/Midtones/Highlights

This array is described in the Colors mode above.

Histogram Control

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 will be.

Luminance is calculated per channel, therefore, the red, green, and blue channels all have their own histogram and the combined result of these comprise 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.

RGB/Red/Green/Blue

These buttons are used to select and display the histogram for each color channel or for the master channel.

Display Selector Toolbar

The Display Selector toolbar 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.

Low/Mid/High

These controls 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) will cause the histogram to slant toward white, shifting the image distribution toward white. The low value will have a similar effect in the opposite direction, pushing the image distribution toward black.

Threshold Output Low/High

The Threshold control can apply clipping to the image, compressing the histogram. Decreasing the level of the High control will reduce the value of pixels in the image, sliding white pixels down toward gray and gray pixels toward black.

Adjusting the Low control toward High will do the opposite, sliding the darkest pixels toward white.

If the low value was set to 0.1, pixels with a value of 0.0 would be set to 0.1 instead, and all 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 on this button will reset all of the controls in the Levels section to their defaults.

Histogram Mode

This Color Corrector mode produces a Histogram display 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.

Float Images and Histogram Equalization or Matching

Use the histogram Matching or Equalization methods on a float image and the color depth of the output image will be 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 will always revert to 16-bit integer processing.

Histogram Control

The Histogram Control is described in detail earlier in this node documentation, under the heading of the Levels Mode.

Keep/Equalize/Match Buttons

Each of these buttons 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 of 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 will appear similar.

When selected, the Equalize and Match modes reveal the following controls.

Match/Equalize Luminance

This slider affects the degree to which the Color Corrector node will attempt 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 will be 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) at the same time.

Lock Red/Green/Blue

When this checkbox is selected, color matching will be applied to all color channels equally. When the checkbox is not selected, individual controls for each channel will appear.

Equalize/Match Red, Green, Blue

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 amount of 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.

8-Bit, 10-Bit, 16-Bit Buttons

This array of buttons determines the level of 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 Out Correction Curves

Often, color equalization and matching operations will 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

Click this button to take a freeze 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 on this button to release the current snapshot of the histogram and return to using the live reference input.

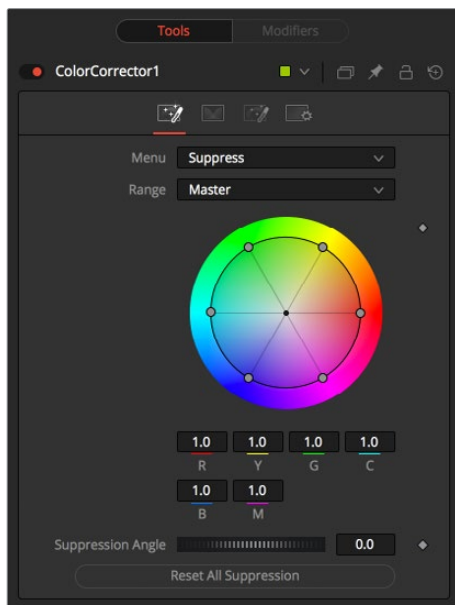
Reset All Histogram Changes

Selecting this button will remove all changes made to the histogram, returning the controls to default and setting the mode back to keep.

Suppress Mode

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 will be suppressed from the image.



Suppression Angle

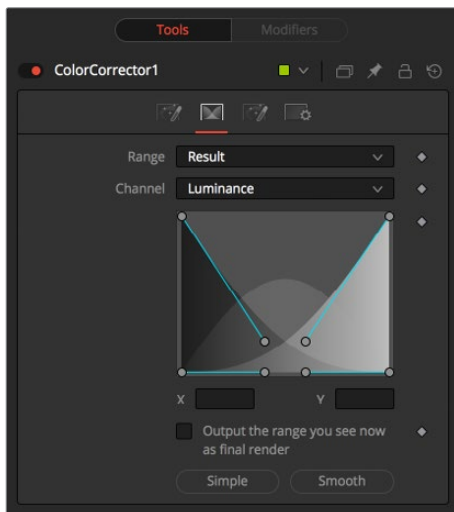
Use the Suppression Angle control to rotate the controls on the suppression wheel and zero in on a specific color.

Reset All Suppression

Clicking on 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 any pixels not already included in either the shadows or the highlights.



Result/Shadows/Midtones/Highlights

These buttons are used to select the color range displayed in the Viewers. They help to visualize the actual pixels that will be included in the range. When the Result button is selected, the image displayed by the color corrector in the views will be that of the color corrected image. This is the default.

Selecting one of the other buttons will switch 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 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 will not receive the full effect of any color adjustments to that range.

Spline Display

The extent of the ranges is selected by manipulating the spline handles. There are four spline points, each with one Bezier 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 Bezier 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 Bezier point or handle.

Channel

The Channel selection buttons shown 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.

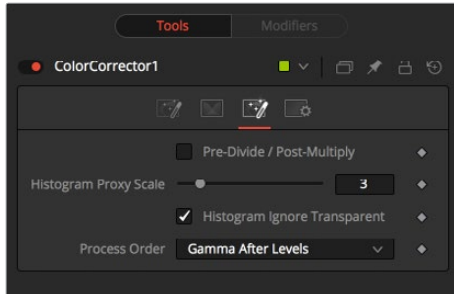
Output the Range You See Now as Final Render

Selecting this checkbox will cause the monochrome display of the range shown in the Viewers to be output as the final render. Normally, the Color node will output the full RGBA image, even if the node was 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 Mode



Pre-Divide/Post-Multiply

Selecting this option will divide 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 against black.

Histogram Proxy Scale

The Histogram Proxy Scale determines the level of 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.

Color Curves [CCV]

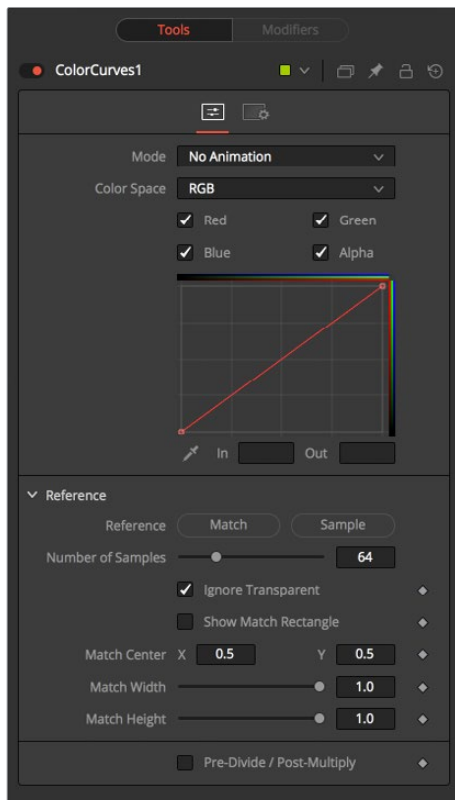


The Color Curves node is a spline-based node for performing Look Up Table (LUT) color manipulations. A separate spline is provided for each color channel. The effect can be animated or dissolved and can be applied to the image using RGB, YUV, YIQ, CMY or HLS color spaces.

The LUT view in the color corrector can be scaled using the + and - keys on the numeric keypad. The color curves LUT has full support for out of range values, 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 a greater precision be required when adjusting the controls.

Controls



Mode

The Mode options change between Animated and Dissolve modes. The default mode is static, 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 can 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 can be found in the online reference documentation for the Color Space node.

Color Channels (RGBA)

Use the Color Channel controls to select which channel's spline is currently active for editing. The labels of these controls will change to reflect the names of the channels for the current color space. Normally, they will read as Red, Green and Blue. If the Color Curves node is operating in YUV color space, they will 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 does not unintentionally affect a different channel's spline.

Spline Window

The Spline Window displays splines for each RGBA channel. These can be edited individually or as a group, depending on the color channels selected above.

The Spline default to a linear range, from 0 in, 0 out at the bottom left to the 1 in, 1 out Top right. In the default a color will process to the same value to the output. If a point is added in the middle at 0.5 in 0.5 out, and the point is moved up, this will raise the mid color of the image brighter.

The Spline curves allow for 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.

Pick

Click on the Pick button and select a color from an image in the display to automatically set keypoints on the spline for the selected color. The new points will be 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.

Match Reference

Clicking on the Match Reference button will automatically set points on the curve to match an image provided in the second (reference) input of the Color Curves node.

Sample Reference

Clicking the Sample Reference button will sample the center scanline of the background image and create a LUT of its color values

Number of Samples on Match Curve


This slider determines how many points are used to match the curve to the range in the reference image.

Pre-Divide/Post-Multiply

Selecting this checkbox will cause 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.

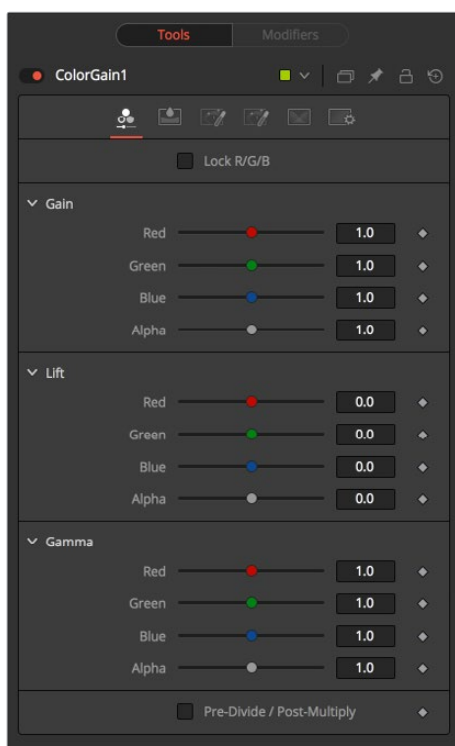
This helps to avoid the creation of illegally additive images, particularly around the edges of a blue/green key or when working with 3D rendered objects.

Color Gain [CLR]



The Color Gain node contains options for adjusting the gain, gamma, saturation and hue of the image. Many of the 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 controls. These can be used to adjust the tinting of the colors in the high, mids and lows.

Controls



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 will be larger on bright pixels and smaller on dark. Black pixels will not be changed ($x * 0 = 0$).

Lift RGBA

While Gain basically 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 will make 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 will be strongest in the midrange and lowrange of the image.


Gamma RGBA

The Gamma RGBA controls affect the brightness of the mid-range in the image. The effect of this node is non-linear. 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 value changes to this control will tend to push mid-range pixels into black or white, depending on the value used.

Pre-Divide/Post-Multiply

Selecting this checkbox will cause 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 to avoid the creation of illegally additive images, particularly around the edges of a blue/green key or when working with 3D rendered objects.

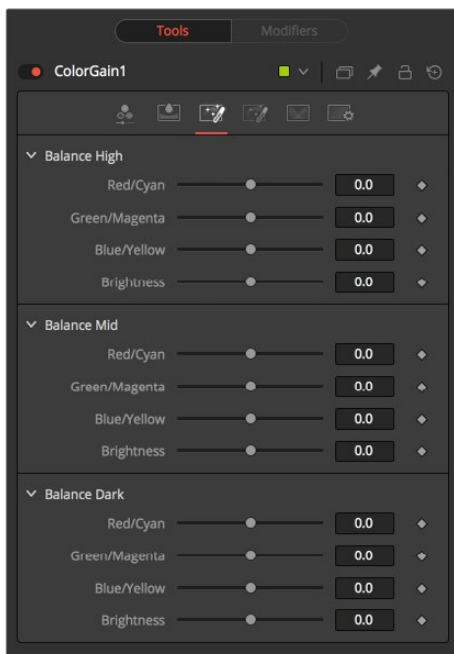
Saturation Tab



RGB Saturation

This setting controls the intensity of the colors in the image channels. A value of 0.0 will strip all of the color out of an image channel. Values greater than one will intensify the colors in the scene, pushing them toward primary colors.

Balance Tab




CMY Brightness Highs/Mids/Darks

This section of 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 to Magenta and Blue to Yellow. Brightness can be raised or lowered for each of the channels.

By default, the Balance sliders can be adjusted by -1 to +1, but values outside of 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



High/Mid/Dark Hue

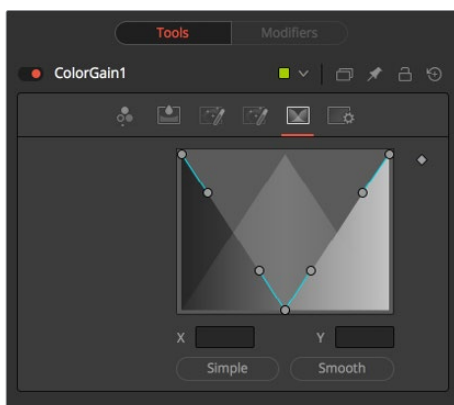
Use the Hue section of the Color Gain node to shift the overall hue of the image, without affecting the brightness or saturation. Independent control of the High, Mid and Dark ranges is offered by three sliders.

The following is the order of the hues in the RGB color space: Red, Yellow, Green, Cyan, Blue, Magenta and Red. Values above 0 push the hue of the image toward the right (red turns to yellow). Values below 0 push the hue toward the left (red turns to 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 of 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 any pixels not already included in either the shadows or the highlights.



Spline Display


The extent of the ranges is selected by manipulating the spline handles. There are four spline points, each with one Bezier 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 Bezier 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 Bezier 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.

Color Matrix [CMX]



The ColorMatrix allows for a vast number of operations to modify values individually in the different color channels.

Controls



Update Lock

When this control is selected, Fusion will not render the node. This is useful for setting up each value of the node, then turning Update Lock off in order to render it.

Matrix

This defines what type of operation actually takes place. The horizontal rows define the output values of the node, the vertical columns the input values. The “add” column allows for simple adding of values to the individual color channels. By default the output values are identical to the input values.

- 100% of the Red channel input is copied to the Red channel output.
- 100% of the Green channel input is copied to the Green channel output.


- 100% of the Blue channel input is copied to the Blue channel output.
- 100% of the Alpha channel input is copied to the Alpha channel output.
- We can also write the default settings as mathematical equations
- $[R \text{ out}] = 1 * [R \text{ in}] + 0 * [G \text{ in}] + 0 * [B \text{ in}] + 0 * [A \text{ in}] + 0$
- $[G \text{ out}] = 0 * [R \text{ in}] + 1 * [G \text{ in}] + 0 * [B \text{ in}] + 0 * [A \text{ in}] + 0$
- $[B \text{ out}] = 0 * [R \text{ in}] + 0 * [G \text{ in}] + 1 * [B \text{ in}] + 0 * [A \text{ in}] + 0$
- $[A \text{ out}] = 0 * [R \text{ in}] + 0 * [G \text{ in}] + 0 * [B \text{ in}] + 1 * [A \text{ in}] + 0$

Invert

Enabling this option will invert 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

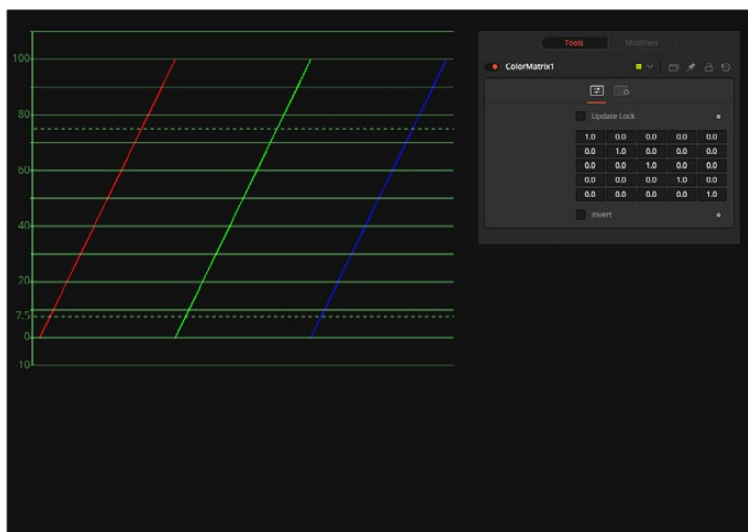
If we want to do a simple invert or negative of the color values, but keep our alpha channel as it is, the matrix would look like this.




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.

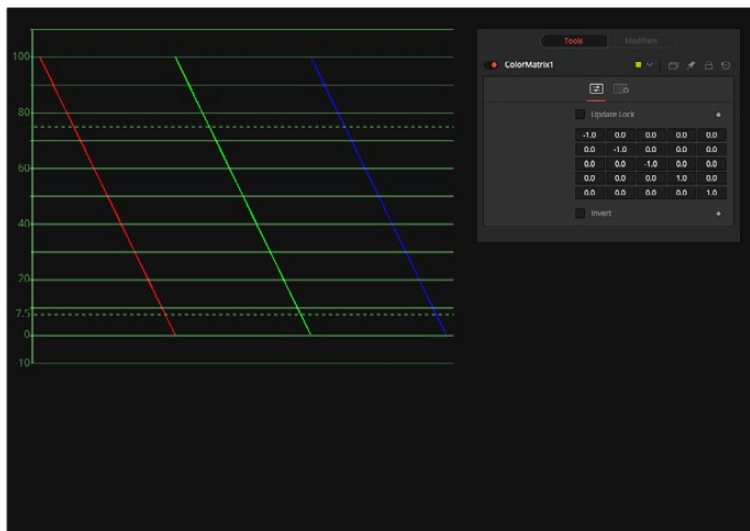
- **01:** The original grayscale



- **02:** RGB set to -1. The values get inverted but fall below 0.



- **03:** Adding 1 to each channel keeps the inversion but moves the values back into a positive range.



Example 2 – Brightness per Channel

Let's influence 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.

Example 3 – Copying Values

Of course we can also copy color values back and forth between individual channels. Let's make the red channel contain the luminance values of the image based on thirds and the green channel contain the luminance values based on the proper black-and-white conversion method, whereas in the blue channel we use a third method based on getting more information from red and less from blue. We also lower the blue channel's brightness by 0.1 and replace the alpha channel with the original blue channel.

Color Space [CS]




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 will cause most images to look odd, as Fusion's viewers will still interpret the primary channels as Red, Green and Blue. For example, viewing an image converted to YUV in one of the Viewers will show the Y channel as Red, the U channel as Green and the V channel as Blue.

Several elements of the Fusion interface refer to the RGB channels directly. The four checkboxes used 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 elements remain R, G and B, but the values they represent are from the current color space. (Red is Hue, Green is Luminance, Blue is Saturation for a RGB to HLS conversion. The Alpha value is never changed by the color space conversion.)

Controls



Color Space Conversion

This button array has three options.

None

The Color Space node has no effect on the image.

To Color

The input image will be converted to the color space selected in the Color Type control found below.

To RGB

The input image will be converted back to the RGB color space from the type selected in the Color Type control below (for example, YUV to RGB).

Color Type

These buttons are used to select the color space conversion applied when the To Color conversion is selected.

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. This format is often used to color correct images, due to 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 more rare 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 non-linear.

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, rather than a non-linear one, unlike the other available color spaces. Non-linear 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.

Copy Aux [CPA]

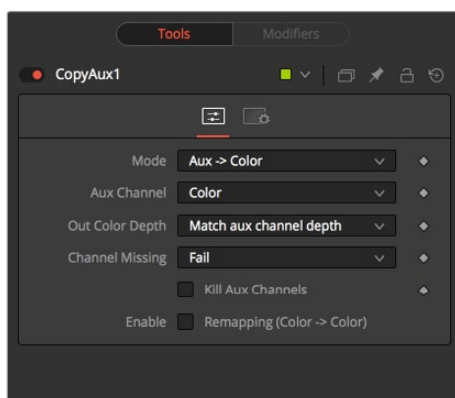


CopyAux copies aux channel groups into RGBA channels. It is mostly a convenience node as the copying can also be accomplished with more effort using a ChannelBoolean node. Although CopyAux has quite a few options, most of the time you will only adjust the channel to copy and ignore the rest.

Before Fusion 6.31, to access aux channels you used a ChannelBoolean to copy them into the RGBA channels. Often this would also involve a ChangeDepth node to make sure that the RGBA channels that were receiving the aux channel were float32. Now in Dimension, CopyAux accomplishes the same result in fewer mouse clicks, allowing you to work faster. Where ChannelBooleans deals with individual channels, CopyAux deals with channel groups. By default, the CopyAux node will automatically promote the depth of its output to match the depth of the aux channel.

CopyAux 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 will 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 will turn 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.

Controls



Mode

Mode determines whether the aux channel is copied into the RGBA color channel or vice versa. Using this option, you can use one CopyAux to bring an aux channel into color, do some compositing operations on it, and then use another CopyAux to write the color back in the aux channel. When the Mode is set to "Color>Aux," all the inputs except AuxChannel are hidden.

Aux Channel

The Aux Channel is to be copied from or written to depending on the current mode. When the aux channel abcd has one valid component, it will be copied as aaa1, two valid components as ab01, three valid components as abc1, and four components as abcd. For example, the Z-channel will be 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 about copying float channels into integer image formats, as they can get clipped if you do not set up CopyAux correctly. For the purpose of this node, all aux channels are considered to be float32 except ObjectID/MaterialID, which are considered to be int16.

Match Aux Channel Depth

The bit depth of the RGBA channels of the output image will be increased to match the depth of the aux channel. In particular, this means that the RGBA channels of the output image will be either int16 or float32. It is wise to be careful when using this option as, for example, if you normally have int8 color channels, you will now be using 2x or 4x more memory for the color channels. In particular, the Z, Coverage, TextureCoordinate, Normal, Vector, BackVector, WorldPosition, and Disparity channels will always be output as float and the Material/ObjectID channels will be output as int16.

Match Source Color Depth

The bit depth of the RGBA channels of the output image will be 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 which are floating-point numbers in the $[-1, 1]$ range will be clipped to a non-negative numbers $[0, 1]$ range. As a more extreme example, consider what will happen to Z values. Z values are floating point numbers stored in the $[-1e30, 0]$ range and they will all get truncated to the $[0, 1]$ range, which means your Z-channel will be full of zeroes.

Force Float32

The bit depth of the RGBA channels of the output image will always be 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

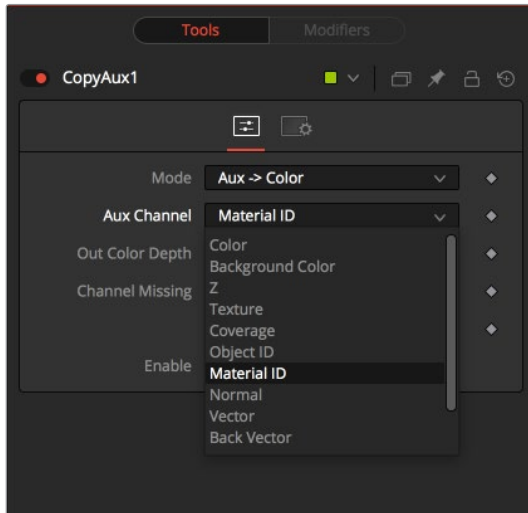
The node fails and prints an error message to the console.

Use Default Value

This fills the RGBA channels with the default value of zero for everything except Z for which it is $-1e30$.

Kill Aux Channels

When this is checked, CopyAux will copy the requested channel to RGBA and then output a resulting image that is purely RGBA with all other channels being killed. This is useful if you want to increase the number of frames of CopyAux 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 Copy Color > Color for a longer color playback.



Enable Remapping

When remapping is enabled, the currently selected aux channel will be 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/From. 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 will be remembered. One way this could be useful is that you can set up all of your 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 in order to store them in an int8 image.

From. Min

This is the value of the aux channel that will correspond to To. Min.

From. Max

This is the value of the aux channel that will correspond 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

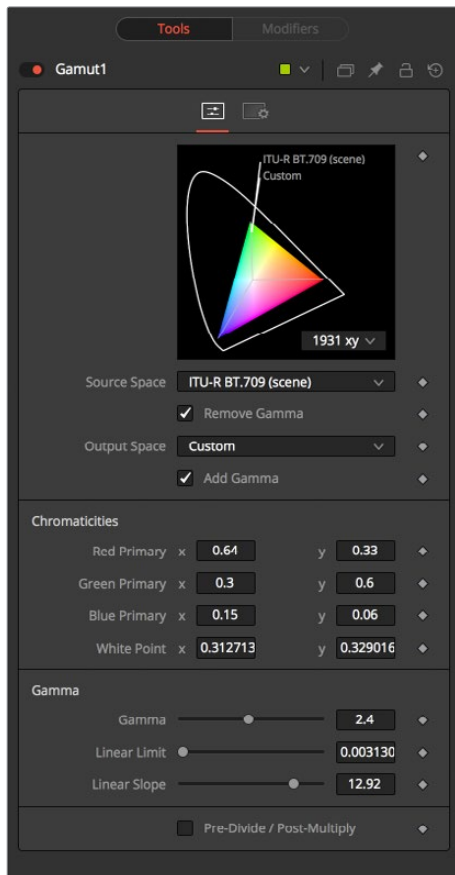
After the values have been rescaled into the [To. Min, To. Max] range, this inverts/flips the range.

Gamut [GMT]



The Gamut node converts color from different gamuts to other color gamuts as well as removing or adding the appropriate gamma to entirely linearize incoming images.

Controls



Source Space

Source Space determines the input color space of the image. Leave this at “No Change” if you want to just add Gamma using the Output Space control.

DCI-P3

The DCI-P3 color space is most commonly used in association with DLP projectors, and is frequently provided as a color space available with 2K DLP projectors, and as an emulation mode for 10-bit LCD monitors such as the HP Dreamcolor. 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 is the converted gamut to the desired color space. Leave this at No Change if you want to just remove Gamma using the Source Space control.


Remove/Add Gamma

Select these checkboxes to do the gamut conversion in a linear or non-linear gamma, or simply remove or add the appropriate gamma values without changing the color space.

Pre-Divide/Post-Multiply

Selecting this checkbox will cause 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.

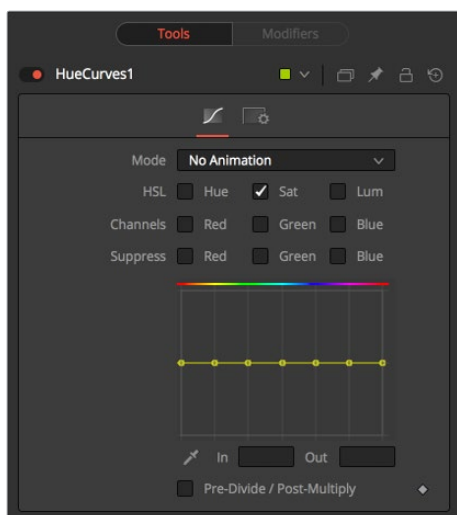
Hue Curves [HCV]



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.

Controls



Mode

The Mode options changes between Animate and Dissolve modes. The default mode is Static, where adjustments to the curves are applied consistently over time. Setting the Mode to Animate or Dissolve allows for 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 in the Pick Color process.

Any number of activated splines can be edited simultaneously, however in most cases it's more convenient to have only the currently modified spline active to avoid unwanted changes to other splines.

When using the Pick Color button a point will be created on all active splines, representing the selected color.

Spline Window

This Look Up Table (LUT) control 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 in this case 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 basically 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 key points placed horizontally at each of the primary colors. From left to right these are: Red, Yellow, Green, Cyan, Blue, and Magenta. Due to the cyclical nature of the hue gradient, the leftmost key point in each curve is connected to the rightmost key point of the curve.

Right clicking in the LUT Control will display a contextual menu containing options for resetting the curves, importing external curves, adjusting the smoothness of the selected key 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.

Pick

Left clicking and dragging from the Pick Color button will change the current mouse cursor to an eyedropper. While still holding down the left mouse button, drag the cursor to a viewer to pick a pixel from a displayed image. This will cause key points, which are locked on the horizontal axis, to appear on the currently active curves. The key points will 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 will cause 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.

OCIO CDL Transform [OCD]

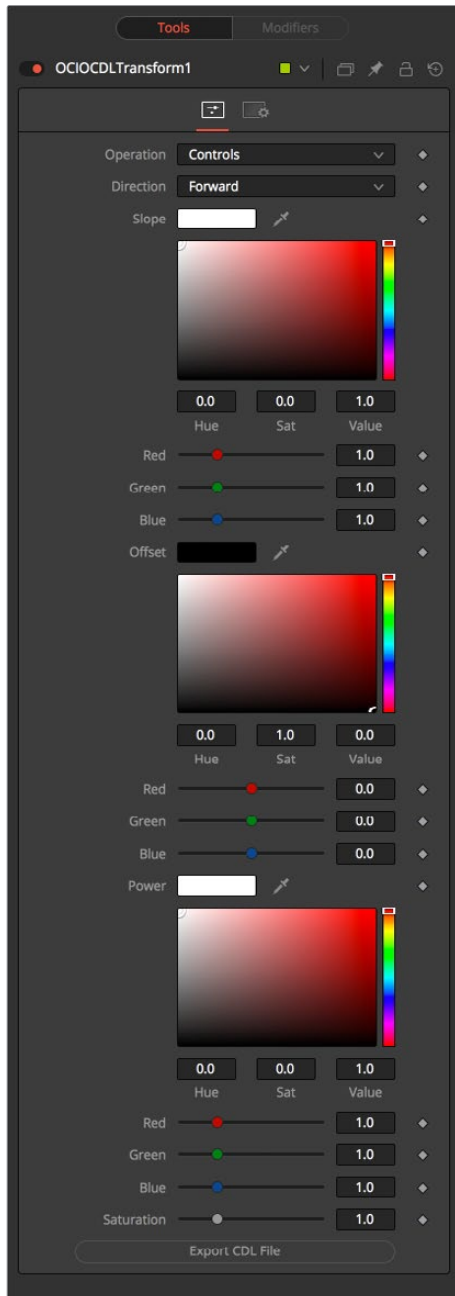


Fusion supports the Open Color IO workflow specified by Sony Imageworks. In general, the color pipeline is made up from a set of color transformations defined by OCIO-specific config files, commonly named with a ".ocio" extension, which allows users to easily share color settings within or between facilities. 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 will be used.

For in-depth documentation of the format's internals, please refer to the official pages on opencolorio.org.

The OCIO CDLTransform allows users to create, save, load, and apply CDL files.

Controls



Operation

Toggles 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.

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 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



Multiplies the color values; this is the same as Gain in the BrightnessContrast node

Offset



Adds to the color values; this is the same as Brightness in the BrightnessContrast node

Power

Applies a Gamma Curve. This is an inverse of the Gamma function of the BrightnessContrast node.


Saturation

Enhances or decreases the color saturation. This works the same as Saturation in the BrightnessContrast node.

Export File

Allows the user to export the settings as a CDL file.

OCIO ColorSpace [OCC]




Fusion supports the Open Color IO workflow specified by Sony Imageworks.

In general, the color pipeline is made up from a set of color transformations defined by OCIO-specific config files, commonly named with a “.ocio” extension, which allows users to easily 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 will be used. For in-depth documentation of the format’s internals, please refer to the official pages on opencolorio.org. The OCIOColorSpace allows for sophisticated color space conversions, based on an OCIO Config File. Sample configs can be obtained from opencolorio.org/downloads.html

The functionality of the OCIOFileTransform node is also available as a ViewLUT node from the ViewLUT menu.

Controls



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 will be listed here.


The content of this list is solely based on the loaded profile and hence can vary immensely.

Output Space

Based on the Config file, the available output color spaces will be listed here.

The content of this list is solely based on the loaded profile and hence can vary immensely.

OCIO FileTransform [OCF]




Fusion supports the Open Color IO workflow specified by Sony Imageworks.

In general, the color pipeline is made up from a set of color transformations defined by OCIO-specific config files, commonly named with a “.ocio” extension, which allows users to easily 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 will be used. For in-depth documentation of the format’s internals, please refer to the official pages on opencolorio.org. The OCIOFileTransform allows the user to load and apply a variety of Look Up Tables.

The functionality of the OCIOFileTransform node is also available as a ViewLUT node from the ViewLUT menu.

Controls



LUT File

Displays a File > Open dialog to load the desired LUT.

CCC ID

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 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 user to select the color interpolation to achieve the best quality/render time ratio.

The ViewLUT Version of the Node


Set Canvas Color [SCV]



Set Canvas Color is used to set the color of the workspace – all the area beyond the defined pixels within an image (the DoD). This area usually extends to infinity. By default, the canvas color used is black/no alpha (transparent).

Some nodes may change an image's canvas color, for example, inverting a mask will change the mask's canvas from black to white. Set Canvas Color allows you to control and override this.


Controls



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.

White Balance [WB]

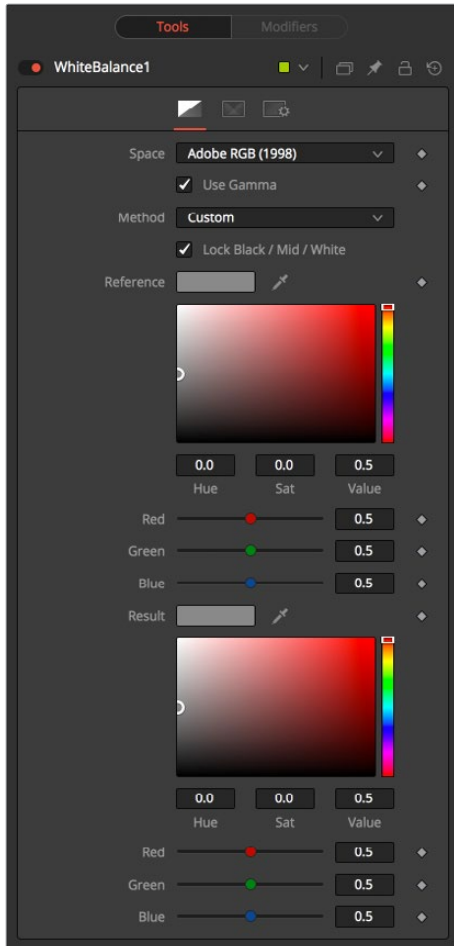


The White Balance node can be used to automatically remove color casts in the image caused by the incorrect setup of a camera, problems in a camera's CCD, or bad lighting conditions.

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.

Balance Tab



Color 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 will be able to 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 and 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 that correction is applied to the entire image, it generally 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, Mid and White points together so that the entire image is affected equally. Unchecking the control will provide individual controls for white balancing each range separately. This control affects both methods equally.

Black/Mid/White Reference

These controls only appear if the Custom method is selected. They are used to select a color from a pixel in the source image. The White Balance node will color correct 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 is 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, even though the color is not white. Similarly, a really dark blue-gray pixel might be 0, 2, 10. It is clipped in red as well, even though it is not black.

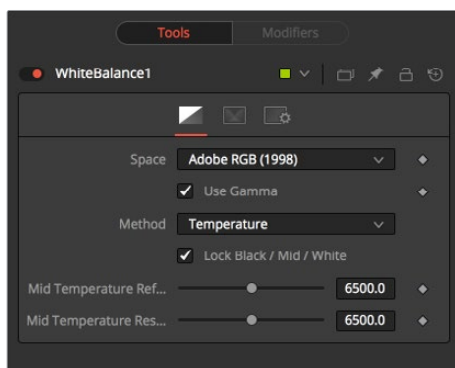
Neither example would be a good choice as a reference pixel because there would not be enough headroom left for the White Balance node.

Black/Mid/White Result

These controls only appear if the Custom method is selected. They are used to select the color to which the node will 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.

Temperature Control



Temperature Reference

Use this control 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.

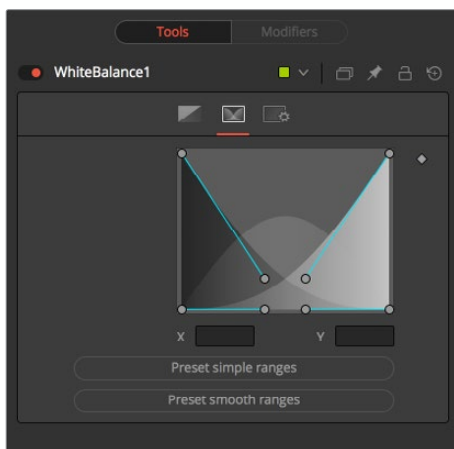
Temperature Result

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 will take the gamma of the image into account when applying the correction, using the default gamma of the color space selected in the menu at the top of the tab.

Ranges Tab



Ranges

Use the controls in the Ranges tab to customize the range of pixels in the image considered to be shadows, midtones, and highlights by the node. The use of the controls in this tab is documented in detail in the Color Corrector node documentation.

Chapter 32

Composite Nodes

This chapter details the Dissolve and Merge nodes available in Fusion.

Contents

Dissolve [DX]	761
Dissolve Node Introduction	761
Inputs	761
Typical Node Structure	762
Resolution Handling	762
Controls	762
Merge [MRG]	764
Merge Node Introduction	764
Inputs	764
Typical Node Structure	765
Resolution Handling	765
Controls	765

Dissolve [DX]



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 amount of 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 all 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.

In addition to 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 Fusion 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 together. Unlike most other nodes, it is not necessary to connect the background input before connecting the foreground input.
- **Foreground:** The second of two images you want to switch between or mix together. 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.

Typical Node Structure

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


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 together by setting the Background/Foreground slider somewhere in between, the output resolution will be set to the larger of the two input resolutions, to make sure there's enough room to encompass 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 in between the foreground and background.

Controls



Primary Controls

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 together 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 regardless of 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) The 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:** (Only appears 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:** (Only appears when Border is turned on) Use Border Color to select the color used in the border.

Merge [MRG]



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 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 is 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 in back.

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; that's 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.

Typical Node Structure

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.



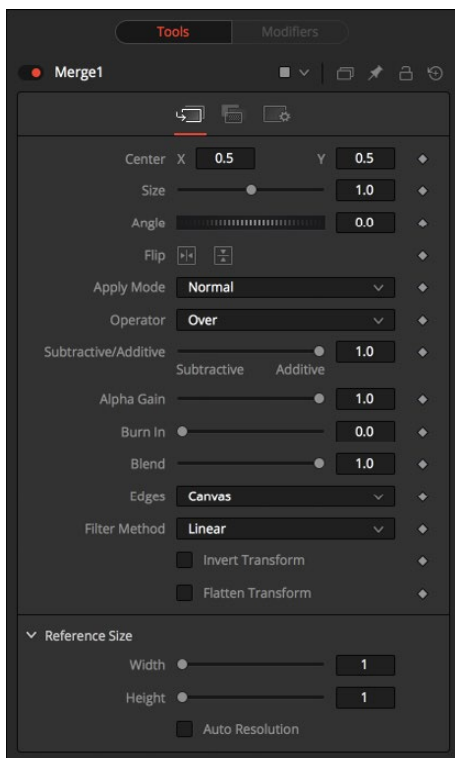
A typical Merge node structure

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 the Resize node to change both the resolution and the size of the image.

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 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:** Multiply 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 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.
- **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 only visible 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, T., and T. Duff, SIGGRAPH 84 proceedings, pages 253-259. Essentially, the math is as described below. Note that some modes not listed in the Operator drop-down (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.
 $x = 1, y = 1 - [\text{foreground alpha}]$
- **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.
 $x = [\text{background alpha}], y = 0$

- **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.
 $x = 1 - [\text{background alpha}], y = 0$

- **ATop:** ATop places the foreground over the background only where the background has a matte.

$$x = [\text{background alpha}], y = 1 - [\text{foreground alpha}]$$

- **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.

$$x = 1 - [\text{background alpha}], y = 1 - [\text{foreground alpha}]$$

- **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 pre-multiplied (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 pre-multiplied, 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 going to be 0. This obscures the background (by multiplying with the inverse of the foreground alpha), 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 pre-multiplied 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.


- **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 by, 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. In fact, 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. Defaults to 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 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.
 - **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 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. 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 of the frame is set to the current color/opacity of the canvas. If you want to change what value this is, 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 out up, down, left, and right from each side to reach the end of the DoD.
 - **Mirror:** Similar to duplicate, except that 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 the purpose of 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 the Reference Size reveal 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 for 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 $5/100 = 0.05$.

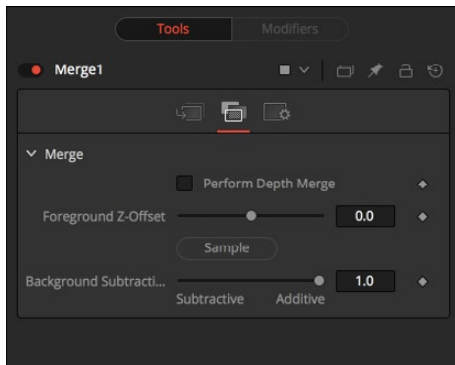
If you specify the dimensions of the background image in the Reference Size controls, this changes the way the Center control values are displayed so that it shows the actual pixel positions in its X and Y fields.

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 were to be queried via scripting or the center control were 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



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 checkbox:** 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 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 slider:** 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 semi-transparent areas. An Additive merge will maintain the transparencies of the image but will add their values to the background.

Chapter 33


Generator Nodes

This chapter details the Generator nodes available in Fusion.

Contents

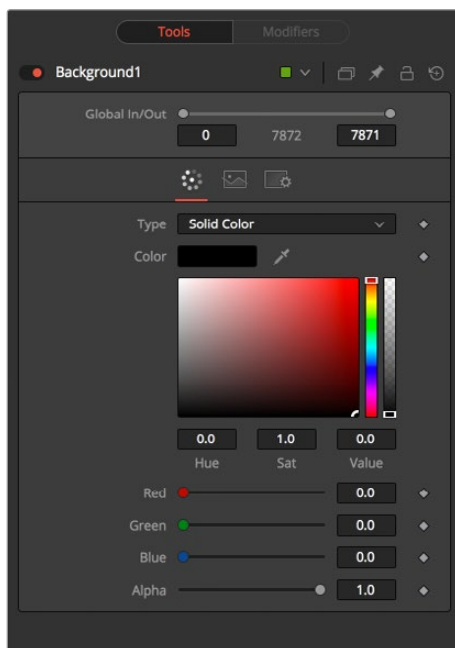
Background [BG]	774
Day Sky [DS]	777
Fast Noise [FN]	780
Mandelbrot [MAN]	784
Plasma [PLAS]	787
Text Plus [TXT+]	790
Modifiers	802
Character Level Styling	802
Comp Name	803
Follower	803
Text Scramble	805
Text Timer	805
Time Code	806

Background [BG]



The Background node can be used to produce anything from simple color backgrounds to complex loopable gradients.

Color Tab




Mode

This control is used to select the mode used by the Background node when the image is generated. 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.

Gradient

This creates a background from a custom gradient.




Color

These controls are used to select the color of the Background node. Depending on the mode selected, one to four color controls will be displayed to create linear color ramp backdrops. Select one and pick a color from the preset menu or create them. Alternatively, enter values in RGB via the keyboard by clicking in the values box and typing in the value. Each color has its own alpha value slider to adjust the transparency of each color.

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 on which frame that the clip starts and Global Out to specify on which frame this clip ends (inclusive) within the project's Global Range.

The node will not produce an image on frames outside of 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).

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 4 times 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.

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.

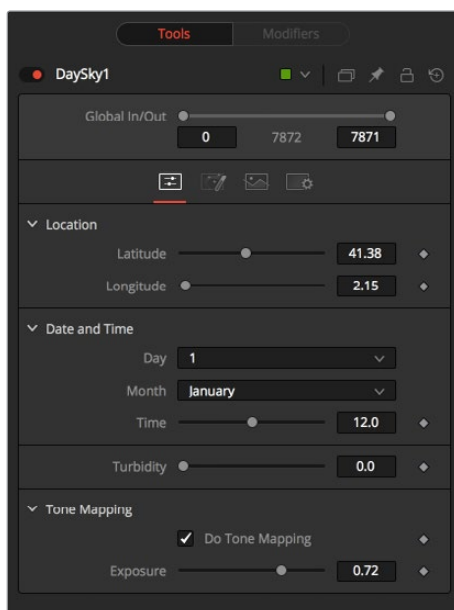
Day Sky [DS]



The DaySky 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>].

This 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.

Controls



Latitude, Longitude

Use these sliders to specify the Latitude and Longitude used to create the Day Sky simulation.

Day, Month, Time

Use these controls to specify the Day, Month, and Time for the DaySky simulation.

Turbidity

Turbidity causes light to be scattered and absorbed rather than 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


Deselect this checkbox to disable any tone mapping applied to the simulation. 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.

Generally, this option should only be deselected 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.

Advanced Tab



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 amount of backscatter light in the simulation.

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 on which frame that the clip starts and Global Out to specify on which frame this clip ends (inclusive) within the project's Global Range.

The node will not produce an image on frames outside of 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).

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 4 times 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.

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.

Fast Noise [FN]



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 and swirling fog to 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 shimmer, particle systems and dirtiness maps.

Controls



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.

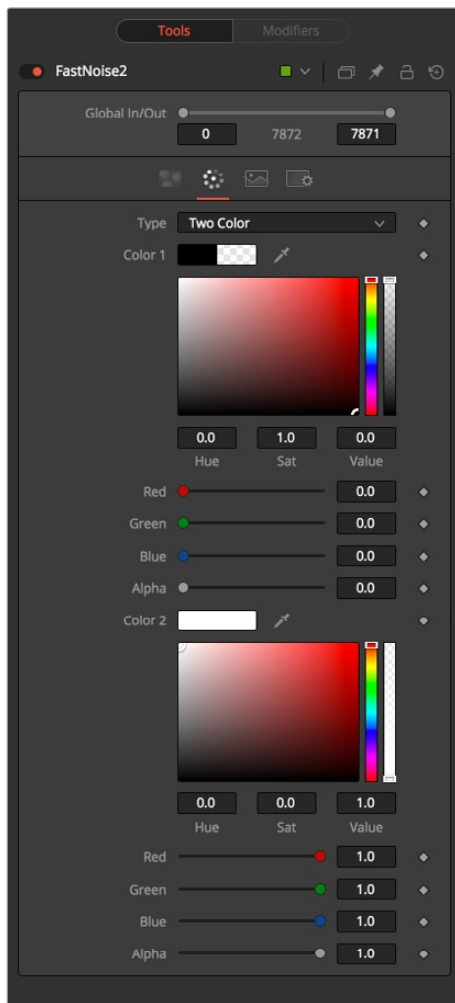
Seethe

Adjust this thumbwheel control to interpolate the noise map against a different noise map. This will cause a crawling shift in the noise, like 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.

Color Tab



Two Color

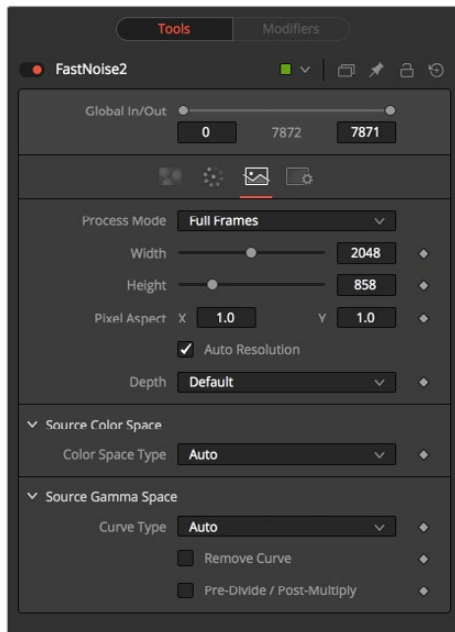
A simple two color gradient is used to color the noise map. The noise function will smoothly transition from the first color into the second.

Gradient

The Advanced Gradient control in Fusion is used to provide much more control over the color gradient used with the noise map.

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 on which frame that the clip starts and Global Out to specify on which frame this clip ends (inclusive) within the project's Global Range.

The node will not produce an image on frames outside of 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).

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 4 times 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.

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.

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

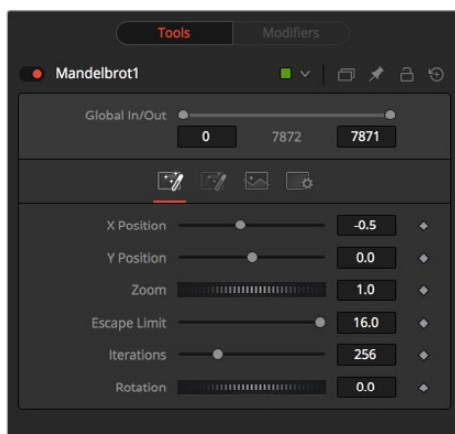
A mask connected to this input can be used to control the noise map completely, such as boosting it in certain areas, combining it with other textures, or if Detail is set to 0, replacing the Perlin Noise map altogether.

Mandelbrot [MAN]



This node creates an image pattern based on the Mandelbrot fractal theory set.

Generation Tab



Position X and Y

This designates 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.

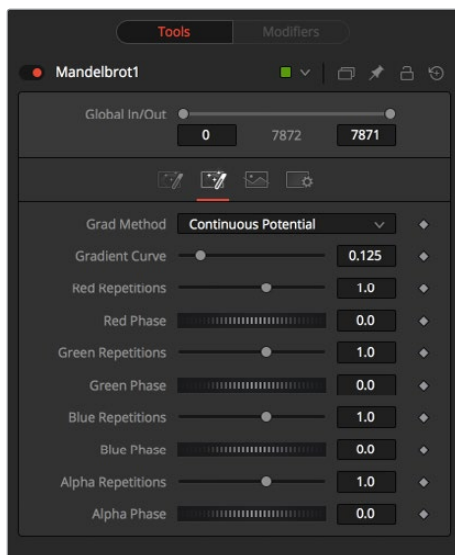
Iterations

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.

Color Tab



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

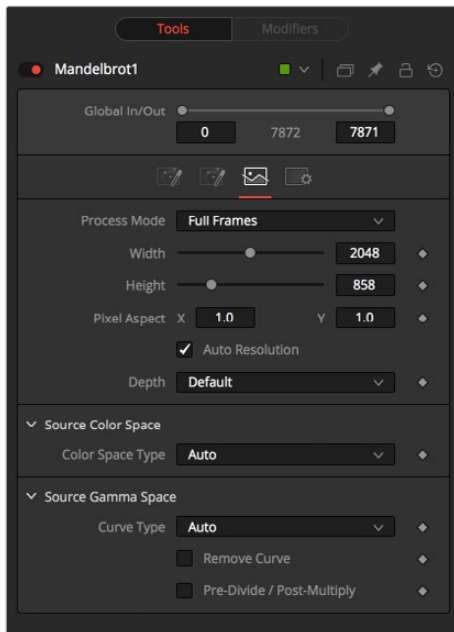
This affects the width of the gradation from the pattern to the background color.

R/G/B/A Phase/Repetitions

Set the color values of the pattern generators.

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 on which frame that the clip starts and Global Out to specify on which frame this clip ends (inclusive) within the project's Global Range.

The node will not produce an image on frames outside of 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).

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 4 times 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.

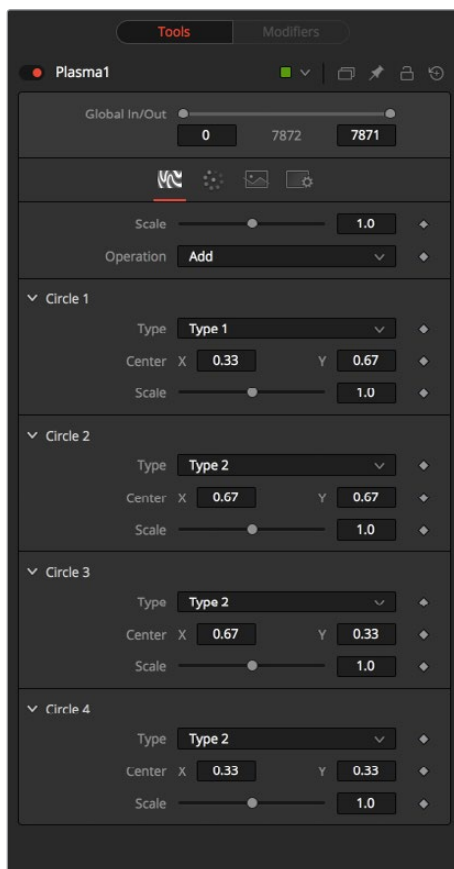
Right-click on the Width, Height or Pixel Aspect controls to display a menu listing the file formats defined in the preferences Frame Format tab. Selecting any of the listed options will set the width, height, and pixel aspect to the values for that format accordingly.

Plasma [PLAS]



This background generation node 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. See also the Fast Noise node.

Circles Tab



Scale

The Scale control is used to adjust the size of the pattern created.

Operation

The options in this menu determine the mathematical relationship between 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

Determine the size of the circle to be used for the pattern.

Color Tab



Phase

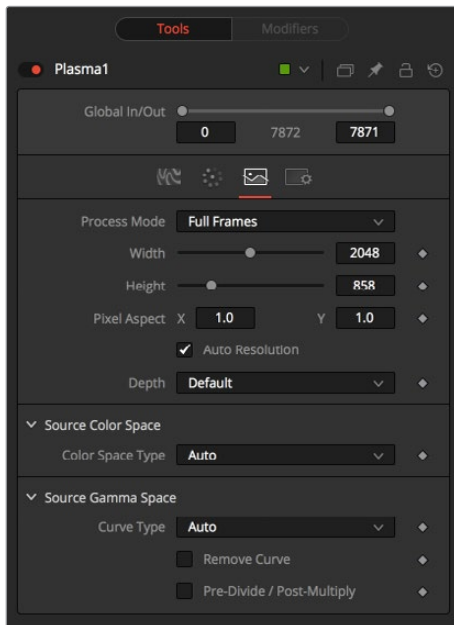
Phase changes the color phase of the entire image. When animated, this creates psychedelic color cycles.

R/G/B/A Phases

Changes the phase of the individual color channels and the alpha. When animated, this creates color cycling effects.

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 on which frame that the clip starts and Global Out to specify on which frame this clip ends (inclusive) within the project's Global Range.

The node will not produce an image on frames outside of 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).

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 4 times 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.

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.

Text Plus [TXT+]

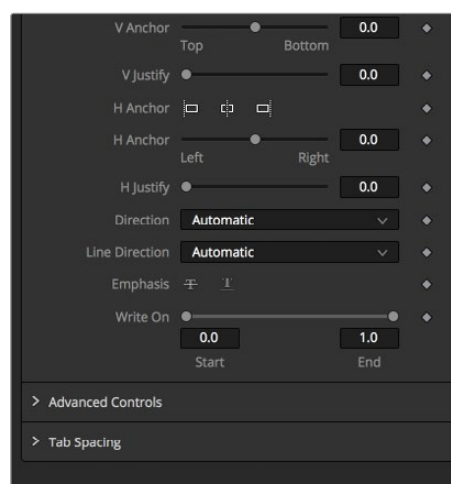
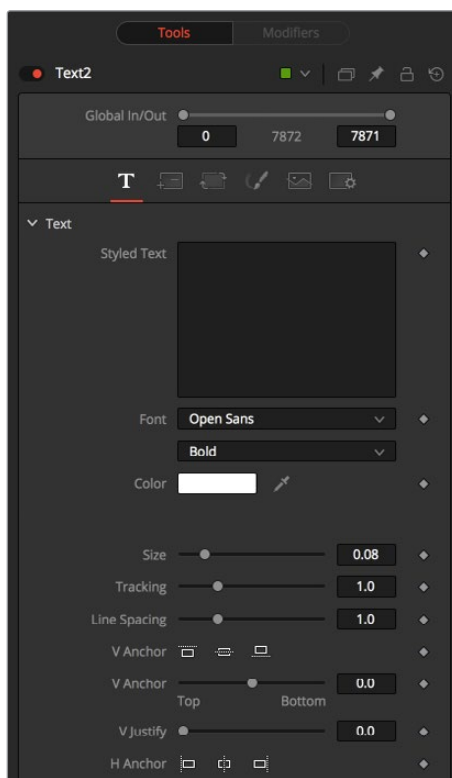


Fusion's Text node is an advanced character generator capable of 3D transformations, multiple styles, and several layers of shading. Text can be laid out to a user-defined frame, circle, or along a path.

Any True Type 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.

Text 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.

Font

The Font controls are used to select the font used by Fusion to create the text. Fusion has limited support for third-party font managers. Fonts managed by a third-party font manager may need to move the Mac OS X or Windows Fonts folder.

Underline and Strikeout

These checkboxes enable the addition of emphasis styles to the font used.

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 on the edit box displays a custom contextual menu.

The Styled Text contextual menu has the following options.

Animate

Use this command to set to a keyframe on the entered text and animate the content over time.

Character Level Styling

This command enables Character Level Styling, which will place a set of controls in the Modifiers tab. Use these controls to affect changes in the font, color, size and transformations applied to individual characters.

Comp Name

Comp Name puts the name of the composition in the edit box and is generally used 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.

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.

International Font Controls

This drop-down menu can be used to select a language specific to a subset of a font.

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-bottom, bottom-top, left-right, or right-left.

Force Monospaced

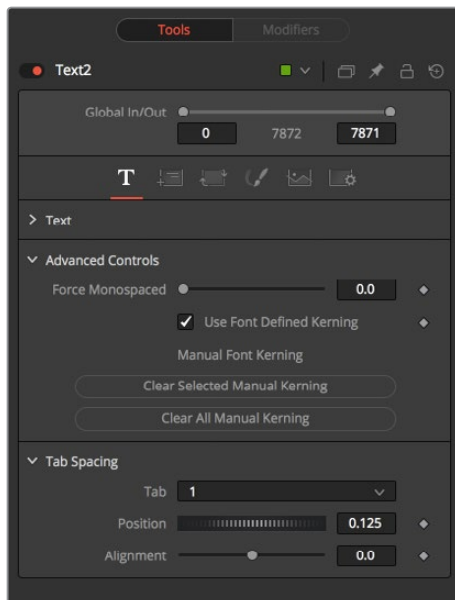
This slider control can be used to override the kerning (spacing between characters) that is defined in the font. Setting this slider to zero (the default value) will cause Fusion to rely entirely on the kerning defined with each character. A value of one will cause the spacing between characters to be completely even, or monospaced.

Do Font Defined Kerning

This enables kerning as specified in the true type font and is on by default.

Advanced Font Controls

Leave the Font Defined Glyphs, Reordering and Ligation checkbox selected unless instructed to do otherwise by support.

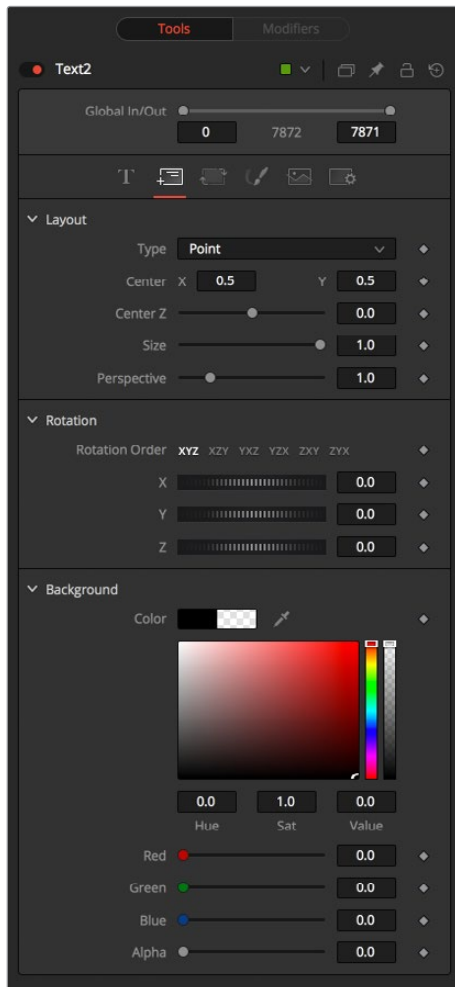


Manual Font Kerning/Placement

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's documentation 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 button array.



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 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.

Layout Rotation

Select this checkbox and another set of options appears to set the Rotation options of the text.

Rotation Order

This button array allows you to select the order in which 3D rotations are applied to the text.

Angle X, Y and Z

These Angle controls 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 only visible when the Layout mode is set to Frame. They are used to adjust the dimensions and aspect 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.

Fit Characters

This menu control is only visible 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 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

This label only appears 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.

The text generated by this node is normally rendered against black. This color picker control can be used to set a new background color.

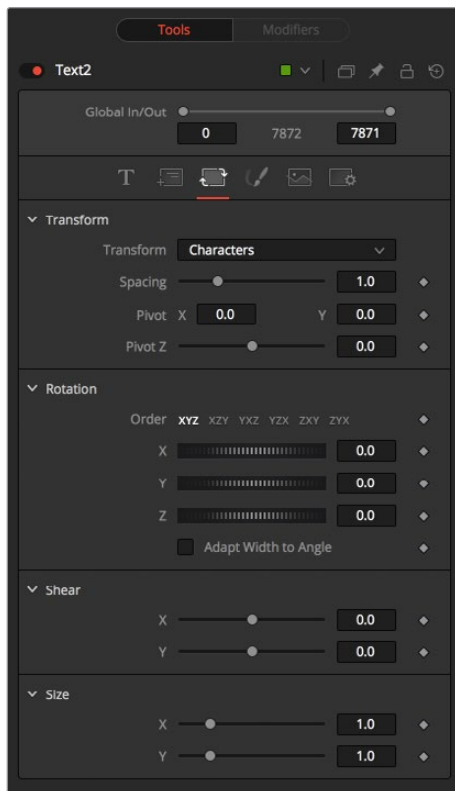
Right-Click Here for Shape Animation

Consult the Motion Paths chapter of this manual for details.

Horizontal and Vertical Alignment

Two identical sets of controls are used to control Vertical and Horizontal Alignment of the text. Use the first array of buttons to choose the alignment of the text. The slider beneath controls the justification.

Transform Tab



Select Transform

There are three buttons 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 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 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 farther along the axis (away from the viewer). Negative values will 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.

Angle 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

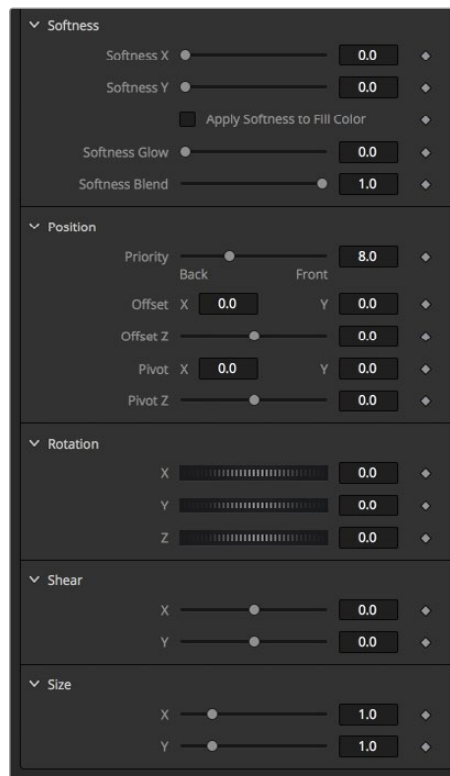
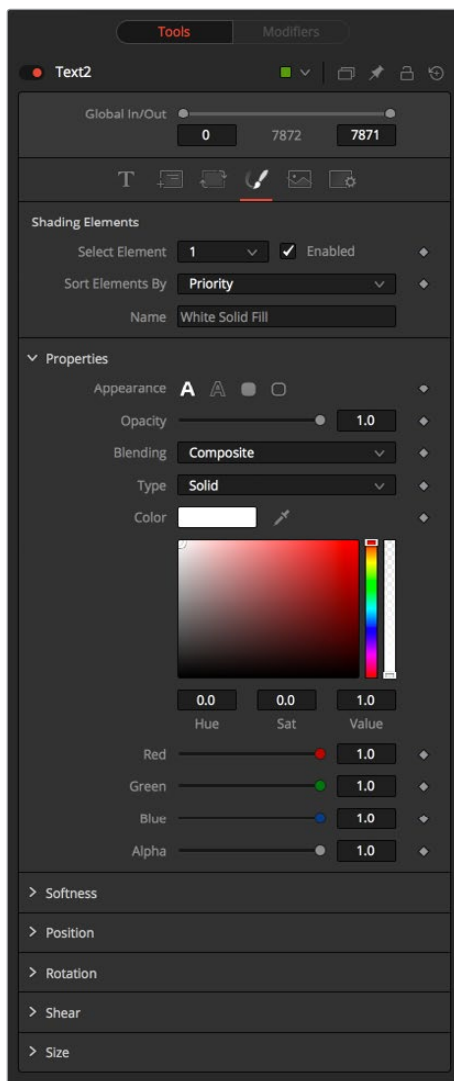
Adjust these sliders to modify the slanting of the text elements along the X- and Y-axis.

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. The Number menu is used to select the element affected by adjustments to the controls in this tab.



Name

This text label can be used to assign a more descriptive name to each shading element.

Enabled

Select this checkbox to enable or disable each layer of shading elements. Element 1 is enabled by default. The controls for a shading element will not be displayed unless this checkbox is selected.

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.

Priority Back/Front

This slider determines the layer's order for the shading elements, also known as the Z-order. 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.

Overlap

This menu is used to select how the renderer deals with an Overlap between two characters in the text.

Composite

This Overlap option will merge the shading over top of itself.

Solid

This option sets the pixels in the overlap region to pure opaque.

Transparent

This option sets the pixels in the overlap region to pure transparent.

Element Type (Buttons)

There are four options available from this menu, providing control over how the shading element is applied to the text. Different controls will appear below depending on the element 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.

Overlap

(All Types) Overlap is used to determine how the shading is handled when portions of the same shading element overlap. Setting this menu to transparent will cause the pixels' color and alpha channels to be set to 0 (transparent).

Thickness

(Outline only) Use this slider control to adjust 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 much 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) This menu provides options for how the corners of the outline are drawn. Options include Sharp, Rounded, and Beveled.

Line Style

(Outline only) This menu offers additional control over the style of the line. In addition to the default solid line, a variety of dash and dot patterns are available.

Shape

(Border Fill only) Shape creates a solid rectangular image around the character.

Shape

(Border Outline only) Creates a rectangular outline around each character.

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 only) This slider is used to round off the edges of the border.

Color Types

In addition to solid shading, it is also possible to 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 if it comes from an external image source. Different controls will be displayed below depending on the Color Type selected.

Solid Mode

When the Type menu is set to Solid mode, color selector controls are provided to select the color of the text.

Image Mode

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 revealed when this option is selected.

Color Image

(Image Mode only) This Color Image text box is used to 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.

Image Size

(Image Mode only) The Image Size menu is used to select how the image is mapped to the text.

Full Image

This option applies the entire image to the text.

Text

This option applies the image to fit the entire set of text.

Line

This option applies the image per line of text.

Word

This option applies the image per each word of text.

Character

This option 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.

Softness On Fill Color Too

Selecting this checkbox will cause blur (softness) to be applied to the shading element as well. The effect is best seen when applied to a shading element colored by 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.

Transform Controls

Selecting the Transform button in the Shading tab will display controls for performing transformations to the shading elements. These controls work in exactly the same way as their equivalents in the Alignment and Layout Tabs, with the addition of a Center Offset control.

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 center 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 farther away from the camera, while positive values will bring it closer to the camera.

Tabs Tab

The controls in the Tabs area are used to configure the horizontal screen positions of 8 separate tab stops. Any tab characters in the text will conform to these positions.

Because the Tab key is used by Fusion to advance to the next control, it is not possible to enter a tab directly into the Styled Text input. Enter a tab using one of the following methods.

Copy and Paste

Copy a tab from another document, such as Notes 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 0.0 to 1.0, where 0 is the far left side and 1 is the far right side of the frame. The position of the tab will be indicated in the Viewer by a thin vertical white line when the Text node is selected and the Tabs tab is open.

Alignment

Each tab can be set to 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. Small white boxes at the top of the tab lines indicate that there are tabs present in the node tree. Clicking within these boxes will toggle the alignment of the tab between the three states.

Rendering Tab

Image Shading Sampling

Use this button array 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 Shading Edges

Use this button array to choose how transformations applied to image shading elements are handled when they wrap off the text's edges.

Sort Shading Elements

This button selection determines the ordering of the shading elements. The default is By Priority, which organizes the shading elements back to front, according to the priority back/front slider in each shading element's controls.

Selecting the By Depth (Z-Position) option will re-organize these according to each element's Z-position, as set by the element's transformations.

Clip Characters Behind Camera

This checkbox determines whether characters that go beyond the plane of the camera will be clipped, or if they will reflect back toward the center of the image. This should normally be enabled for clipping, but some interesting effects can be produced when clipping is disabled.

Anti-Aliasing

This slider is used to increase or decrease the amount of anti-aliasing applied to the text. Higher values mean exponentially longer render times, while lower values reduce render times at the expense of quality.

Render to Flash File

Click this button to render the output of the Text node to an Adobe Flash file.

Library Tab

The Shading Library is used to store and retrieve the settings of a Text node for easy re-use.

Put

To add a new text element to the library, click on the Put button. All of the current settings for the node will be saved according to the name assigned. A thumbnail will appear in the Shading Library showing how that style looks on the letter A.

Get

Click on a shading element in the library, then click on the Get button to apply those settings to the current node. Get will replace all of the text as well as the style.

Shading

Clicking on Shading will replace the shading elements only, without affecting the text entered. Right-click in the Shading Library window for a list of display options.

Text+ Toolbar

When the Text node is selected, a toolbar will appear in the Viewer.



Allow Typing In Preview Window

Enable this button to type and edit text directly in the Viewer, click on the text to produce a cursor that can be positioned within the text. The cursor can be moved using the arrow keys. Type normally.

Allow Manual Kerning

This button enables Manual Kerning, which overrides the automatic kerning normally applied to text. A small red dot will appear beneath each character of text in the Viewer. Drag on the dot to

reposition a character. Also select multiple characters and move them together. Hold the Option or Alt key down while dragging to constrain motion to a single axis.

Use the arrow keys on the keyboard to make manual adjustments to the position of the selected characters. Hold the Command or Ctrl key down while pressing arrow keys to move the character in smaller increments. Shift will move the characters in larger increments.

To animate the position of each character, right-click on the control label Manual Font Kerning/Placement beneath the Text tab's Advanced Font 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

This button 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.

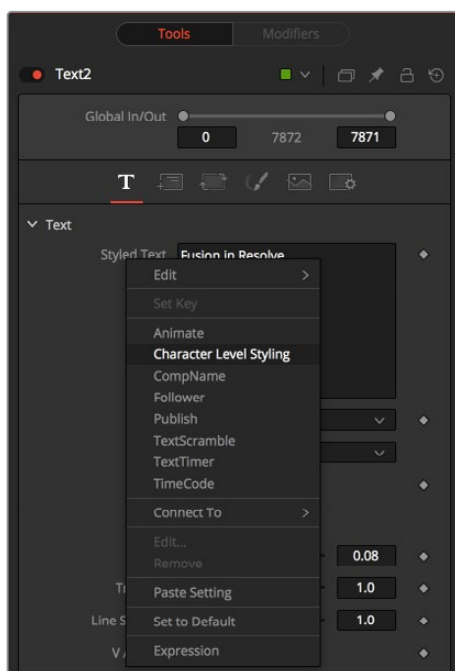
Text Outline Outside Frame Only

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 off screen and is no longer rendering a visible result.

Show Always Text Outline

This button draws an outline around the edges of text at all times, whether the text is visible within the frame or not.

Modifiers



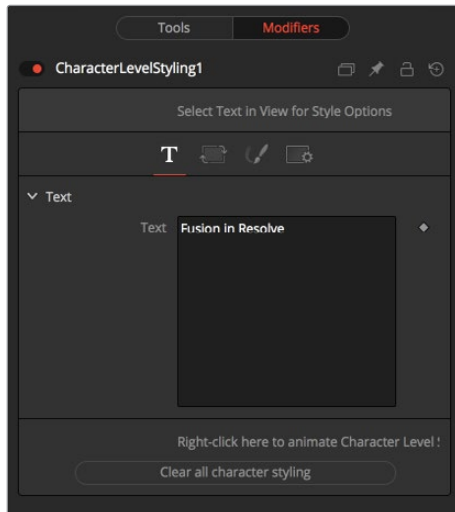
Character Level Styling

The Character Level Styling modifier only works on Text+ nodes. You can then select individual characters directly in your view and apply different text attributes to them. Think of larger capitals at the beginning of each line, different fonts in one word, colored highlighting, and more.

It can be applied by right-clicking into the Styled Text field of a Text+ node and selecting Character Level Styling.

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 means of copying the Text+, right-clicking on the Text 3D and choosing Paste Settings.

Text Tab



Text Controls, Alignment, Transform and Shading Tab

For details see the Text+ node documentation.

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.

Comp Name

The Comp Name only works on Text+ nodes. It 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 into the Styled text field of a Text+ node and selecting Comp Name.

Controls


This Modifier Has No Controls

For further details see the Text+ node documentation.

Follower

The Follower only works on Text and Text3D nodes. This node allows for a plethora of cool motion graphics effects. The basic idea is that you animate the parameters of a single character and the other characters will follow that animation with a delay. It can be applied by right-clicking into the Styled Text field of a Text node and selecting Follower.

Timing Tab



Range

Allows the user to select if all characters should be influenced or only a selected range. You can drag-select a range of characters directly on the screen.

Order

Determines in which order the Characters are influenced. Notice that empty spaces are counted as characters as well. Available options are:

- **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 effected characters can be specified by sliders.

Delay Type

Determines what sort of delay is applied to the animation. Available options are:

- **Between Each Character:** The more characters there are in your text, the longer the animation will take to the end.
- **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 Tab

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 needs to 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 only works on Text+ nodes. It scrambles the Text around, randomly replacing 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 TextScramble.

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. Text can be entered either here or in the Text+ node.

Animate On Time

When set, the characters will get scrambled randomly on every new frame. This switch has no effect when Randomness is set to 0.

Animate On Randomness

When set, the characters will get 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 set, the length of the single words will stay the same, though their characters get scrambled around.

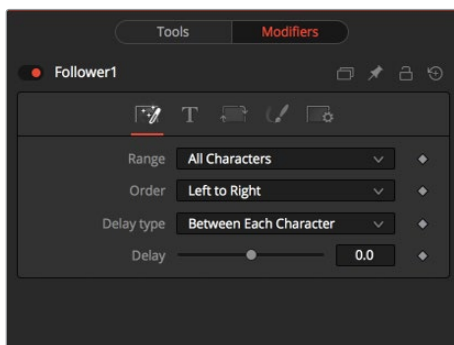
Substitute Chars

Defines which characters are used to scramble the text.

Text Timer

The Text Timer only works on Text+ nodes. It makes the Text+ node either a Countdown, a Timer or a Clock. This is quite useful for on screen real-time displays or to burn in the creation time of a frame into the picture.

It can be applied by right-clicking into the Styled text field of a Text+ node and selecting TextScramble.



Mode

Sets the mode the timer is working in. In Clock mode the current system time will be displayed.

Hrs, Mins, Secs (Switches)

Defines which parts of the clock should be shown on screen.

Hrs, Mins, Secs (Sliders)

Set 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 to automate burn-ins for daily renderings.

It can be applied by right-clicking into the Styled text field of a Text+ node and selecting Time Code.

**Hrs, Mins, Secs, Frms, Flds**

Activate or de-activate these options to customize the time code display to show hours, minutes, seconds, frames, and fields respectively. Activating Frames only will 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.

Chapter 34

DeepPixel Nodes

This chapter details DeepPixel nodes found in Fusion.

Contents

Ambient Occlusion [SSAO]	809
Depth Blurred	812
Fog	814
Shader	815
Texture	817

Ambient Occlusion [SSAO]



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. For this reason, 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 approximates expensive raytraced global illumination quickly. Being a post effect, it exposes similar aliasing issues like the Shader, Texture, and VolumeFog nodes. Hence, artifacts may appear in certain situations.

Usage

The AO node rarely works out of the box, but requires some tweaking. The setup process involves adjusting the KernelRadius and NumberOfSamples to get the desired affect.

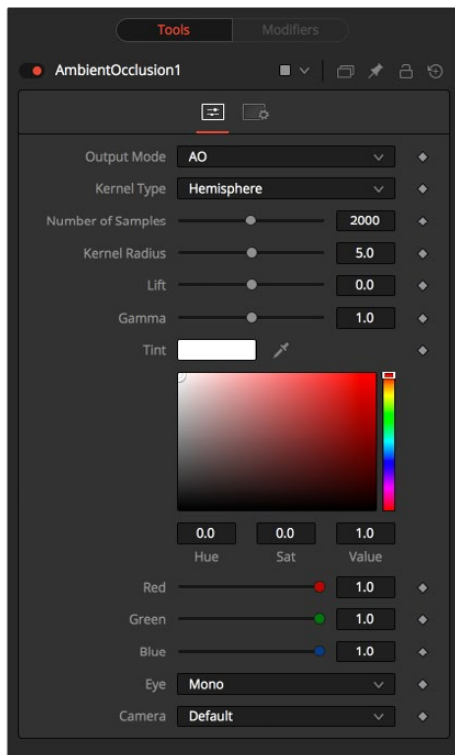
The KernelRadius depends on the natural 'scale' of the scene. Initially there might appear to be no AO at all. At this point either the KernelRadius is too small or too big and working values have to be found.

Inputs

- **Image (required):** Requires RGBA, Z-Depth and Normals.
- **Scene (required):** The Scene or 3D Camera the image was rendered with.

If any of these are not supplied, the node will fail with a descriptive error.

Controls



Output Mode

- **Color:** The incoming image with Ambient Occlusion applied
- **AO:** Outputs the pure Ambient Occlusion as a grayscale image.

White corresponds to regions in the beauty pass 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 for creative freedom to combine the passes in various ways.

Kernel Type

To determine the AO, rays are cast outward from a point on the surface being shaded outwards to a large enclosed sphere.

The number of unoccluded rays, that is those rays that reach the sphere, determines the AO factor.

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 will receive 100% ambient intensity, while other parts will be 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 will receive 50% ambient intensity, while other parts will be made darker or brighter.

Number of Samples

Increase the number of samples until artifacts in the AO pass disappear. Higher values can generate better results but also increase render time.

Kernel Radius

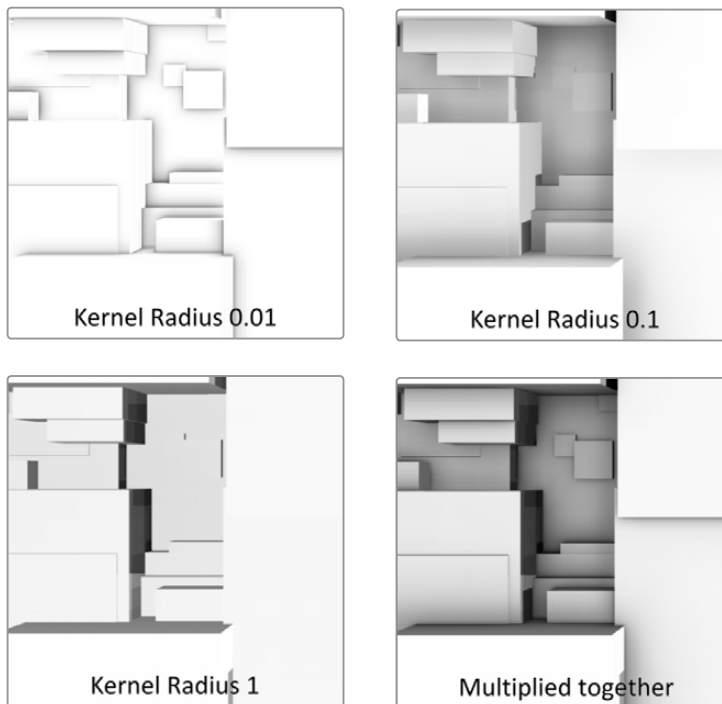
The Kernel Radius controls the size of the filter kernel in 3D space. For each pixel, it controls how far one searches in 3D space for occluders. Most likely, the Filter Kernel is finicky and may need to be adjusted manually for each individual scene.

If made too small, nearby occluders can be missed. If made too large, the quality of the AO will decrease and the number of samples needs to 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 needs to be large as well. With tiny Z values, a small kernel size like 0.1 should be sufficient.

Lift/Gamma/Tint

Used to adjust the AO for artistic effects.



TIP: Combining multiple AO passes with different kernel radii can produce better effects.

Known Issues

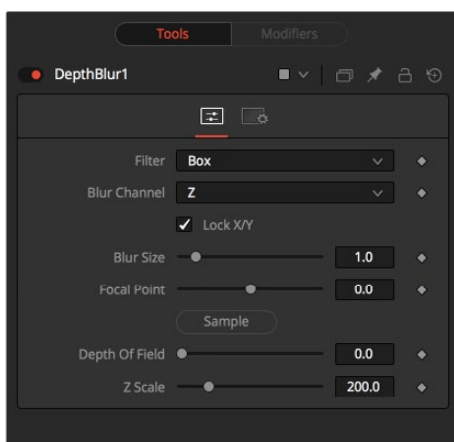
- **Transparency/Translucency:** AO is designed to work with opaque objects. There will be two kinds of problems: those with transparent receivers and those with transparent occluders. You can work around some of these problems by splitting out the transparent/translucent objects into separate scenes and only computing AO on the opaque objects.
- **Particles:** Do not use AO on particles, unless the particles are solid opaque geometry. This is just the transparency problem again. Antialiased edges (another form of transparency) will also cause problems with AO. There's not much you can do about this.
- **Supersampling:** See Antialiased edges. In order for this to work, the AO needs to be computed in the big image before it is downsampled to the final image.
- **View Dependence:** AO methods work in view space and the results are view dependent. This is a limitation of the technique itself. 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 UV renderer can be used to bake AO into the textures on models.

Depth Blurred



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 by means of the Blur Channel controls. If the Blur Image input is connected, channels from that image will be used to control the blur.

Controls



Filter

These buttons are used to select the filter used for applying the blur.

Box

This applies a 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 Channels

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 a node is connected to the node's Blur Image input on the node tree.

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 applied to the image.

Focal Point

This control is only visible when the Blur channel is set to use the Z-buffer.

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.

Drag the Pick button (the pointer changes to a dropper) over the displayed image and select a Z-buffer value. If the image does not contain a valid Z-buffer, no change will occur in the value.

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 was 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 it 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.

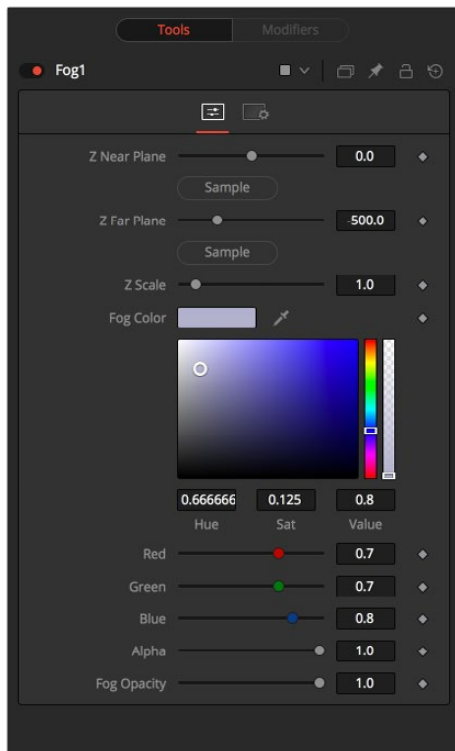
Fog



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.

The second image input on the Fog node can be used to provide 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 provided here.

Controls



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 it 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 of the fog.

Shader



The Shader node can control the lighting, reflection mapping and 3D shading of elements in a rendered image. The reflection map image (connected to the green input) can be projected onto all elements in the scene or to elements selected by the Object and Material ID channels in the common controls. Effect masking can also be used to limit the effect of this node.

The Shader node relies on the presence of the X, Y, and Z normal map channels in 3D rendered images. If these channels are not present, this node has no effect.

Light Tab



Ambient

Ambient controls the amount of Ambient color present in the scene or for the selected object. This is a base level of light, added to all pixels, even in completely shadowed areas.

Diffuse

This option controls the amount of 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 amount of 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 amount of Reflection contribution in the scene or for the selected object. High levels make objects appear mirrored, low levels overlay subtle reflections giving a polished effect. It has no effect if no reflection map is connected.

Reflection Type

Select from these three buttons to determine 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



Edit Diffuse and Specular

Toggle these controls on to enable editing of the Shader curves for the individual channels 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 a number of points to group select them. Right-clicking displays a menu with options for adjusting the spline curves.

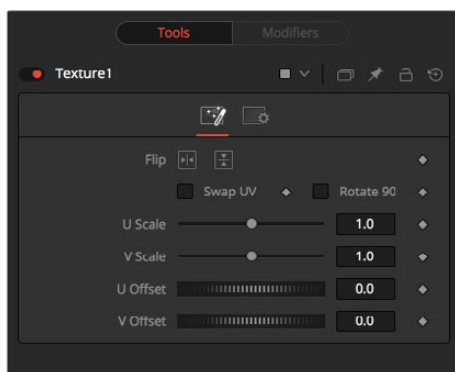
Texture



The Texture node can control the texture mapping of elements in a rendered image. The texture-map image (connected to the green input) can be wrapped around objects to replace the current texture. 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 will 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.

Texture Tab



Swap UV

When this checkbox is selected the U and V channels of the source image are swapped.

Flip Horizontal and Vertical

The texture-map image is flipped horizontally and/or vertically when this control is toggled on.

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.

Chapter 35


Effect Nodes

This chapter details the Effect nodes in Fusion.

Contents

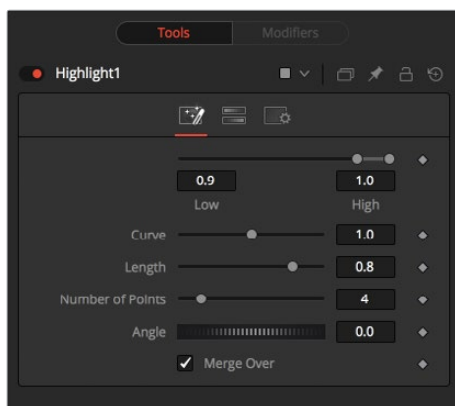
Highlight [HIL]	821
Hot Spot [HOT]	822
Pseudo Color [PSCL]	828
Shadow [SH]	829
Trails [TRLS]	831
TV [TV]	834

Highlight [HIL]



The Highlight filter creates star-shaped highlights in bright regions of the image, similar to a lens star filter effect.

Settings Tab



Low and High

This range control designates the range of Luminance values in the image that will generate highlights. Values less than the Low value will not receive highlights. Values above the High value will receive the full highlight effect.

Curve

The Curve value changes the drop off over the length of the highlight. Higher values will cause the brightness of the flares to drop off closer to the center of the highlight, whereas lower values will 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.

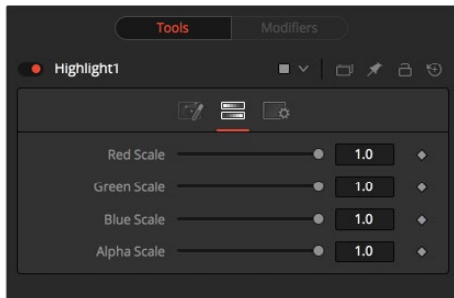
Angle

Use this control to rotate the highlights.

Merge Over

When this checkbox is on, it will overlay the effect on the original image. With the checkbox off, the output will be the highlights only. This is useful for downstream color correction of the highlights.

Color Scale Tab



Red, Green and Blue Scale

Moving the sliders of one or all of these channels down will change the falloff color of the highlight.


Alpha Scale

Moving the Alpha slider down will make highlight falloff more transparent.

Highlight Masks

The Highlight node offers an additional mask input called the Highlight Mask. This is a pre-mask that determines what pixels can be affected by the highlight before the highlight is applied. Unlike regular masks, it will not crop off highlights from source pixels when the highlight extends past the masks edges.

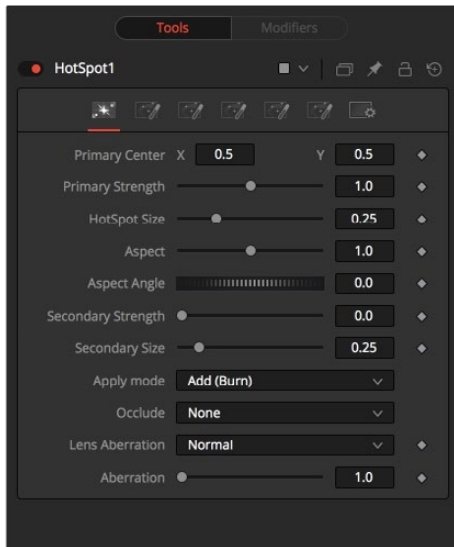
Hot Spot [HOT]



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 present 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 a bright star.

Hot Spot Tab



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.

Primary Strength

This control determines the brightness of the primary hotspot.

Hotspot Size

This control determines the diameter of the primary hotspot. 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 hotspot. Values above 1.0 will elongate the circle horizontally and values below 1.0 will elongate the circle vertically.

Aspect Angle

This control can be used to rotate the primary hotspot.

Secondary Strength

This control determines the strength, which is to say the brightness, of the secondary hotspot. The secondary hotspot is a reflection of the primary hotspot. It will always be positioned on the opposite side of the image from the primary hotspot.

Secondary Size

This determines the size of the secondary hotspot.

Apply Mode

This control determines how the hotspot 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

Use this button array to select which channel of the image connected to the Hotspot node's Occlusion input is used to provide the occlusion matte.

Occlusion happens when the source of the hotspot is blocked by something between it and the camera. When the source of a hotspot is occluded, the hotspot simply winks out.

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. The white pixels in the image will occlude the hotspot. Gray pixels will partially suppress the hotspot.

Lens Aberration

Aberration changes the shape and behavior of the primary and secondary hotspots.

In and Out Modes

Elongates the shape of the hotspot into a flare. The hotspot stretches toward the center of the image when in In mode, stretching toward the corners when in 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 will cause the effect to become more severe the closer the hotspot gets to the center. Flare Out causes the effect to increase, as the hotspot gets closer to the edges of the image.

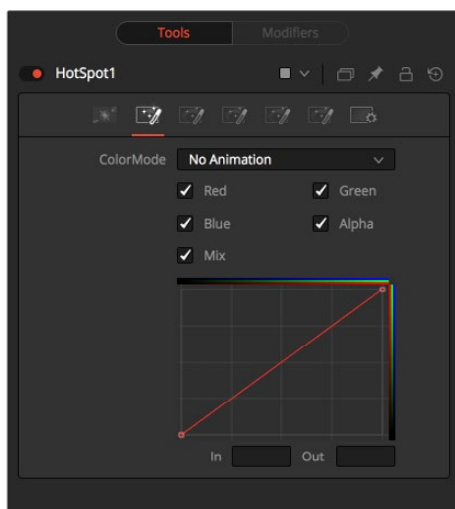
Lens

This mode emulates a round, ringed lens effect.

Aberration

The Aberration slider controls the overall strength of the lens aberration effect.

Color Tab



Color Channel and Mix

When selected, these checkbox controls enable the editing of the chosen splines in the LUT below.

Red, Green, Blue and Alpha Splines

The Red, Green, Blue, and Alpha splines are used to adjust the color of the spotlight along the radius of the hotspot.

The vertical axis represents the intensity or strength of the color channel, from a value of 0 at the bottom to 1 at the top. The horizontal axis represents the position along the radius of the hotspot, from the outside edge on the left to the inside on the right.

So the default curve of the node indicates that the red, green, blue, and alpha channels all have a linear falloff from the outside edge of the curve to the inside edge.

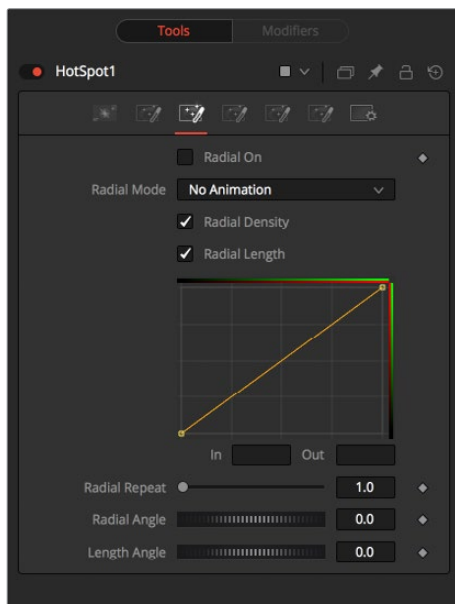
Mix Spline

The Mix spline is used to determine the influence of the controls that the Radial tab will have along the radius of the hotspot. 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 hotspot to blend with the color hotspot. A value of 0 is all radial hotspot, while a value of 1.0 is all color hotspot.

NOTE: Right-clicking in the LUT will display a contextual menu with options related to modifying spline curves.

A complete description of LUT editor controls and options can be found in Chapter 82, “LUT Nodes.”

Radial Tab



Radial On

When selected, this control enables the Radial splines. Otherwise, the radial matte created by the splines is not applied to the hotspot, and the Mix spline in the color controls will have no effect on the hotspot.

Radial Length and Radial Density Splines

The key to these splines is realizing that the LUT Editor's horizontal axis represents a position around the circumference of the hotspot. A value of 0.0 is 0 degrees and 1.0 is 360 degrees. With that in mind, the length determines the radius of the light making up the hotspot along the circumference, and the density represents how bright the light is along the circumference.

Radial Repeat

This control will repeat 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 again between 180 and 360.

Length Angle

This control will rotate the effect of the Radial Length spline around the circumference of the hotspot.

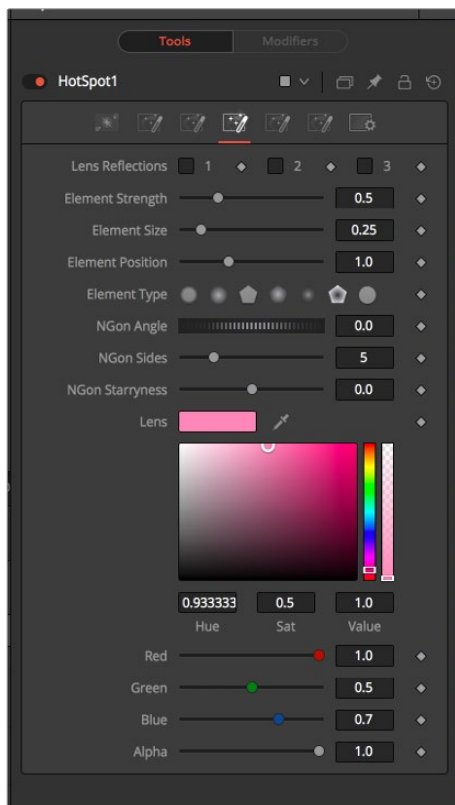
Density Angle

This control will rotate the effect of the Radial Density spline around the circumference of the hotspot.

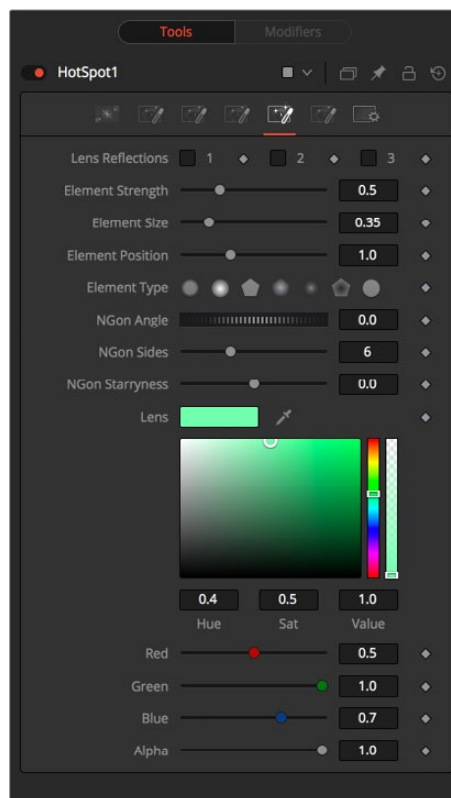
NOTE: Right-clicking in the spline area will display a contextual menu containing options related to modifying spline curves.

A complete description of LUT Editor controls and options can be found in Chapter 82, "LUT Nodes."


L1, L2 and L3 Tab



L1 Tab



L2 Tab



L3Tab

Lens Reflect 1-3

When selected, the reflection caused by the element is enabled.

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 hotspot position and the center of the image.

Element Type

Use this array 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 amount 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

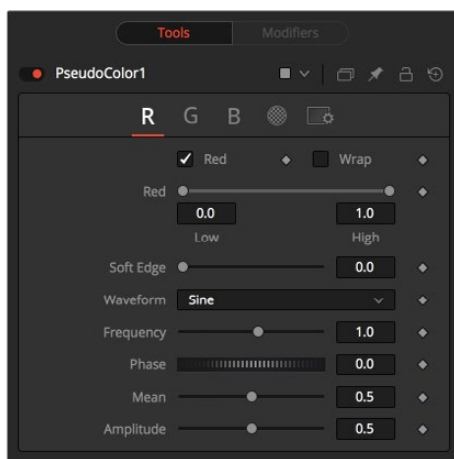
These controls determine the color of the lens that affects the colors of the reflections. To choose a lens color, pick one from a displayed image, or enter RGBA values using the sliders or input boxes.

Pseudo Color [PSCL]



The Pseudo Color node provides the ability to produce variations of color based on waveforms generated by the node's controls. Static or animated variances of the original image can be produced. The node's controls are separated into four identical tabs, one for each of the color channels.

R/G/B/A Tabs



Color Check Box

When selected, the Pseudo Color node will affect this color channel.

Wrap

When selected, waveform values that exceed allowable parameter values will be 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 will increase the number of occurrences of the variances.

Phase

This modifies the Phase of the waveform. Animating this control will produce color cycling effects.

Mean

This determines the level or Mean value of the waveform selected. Higher values will increase the overall brightness of the channel until the allowed maximum is reached.

Amplitude

Amplitude increases or decreases the overall power of the waveform.


Shadow [SH]



Shadow is a versatile node used in the creation of shadows that are based on the alpha channel information from an image. It is also possible to use an additional image to control the shadow's apparent depth.

The Shadow node is designed to cast simple 2D shadows. Use a Spotlight node and an Image Plane node for full 3D shadow casting.

Controls



Shadow Offset

This control sets the apparent distance of the shadowed object from the background. Adjusting the position of the Shadow Offset crosshair in the views is the quickest way to create simple drop shadows.

Softness

Softness controls how blurry the shadow's edges appear. The further the background is from the object, the fuzzier it becomes.

Shadow Color

Use this control to select the color of the shadow. The most realistic shadows are usually the ones that are 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 will vary the apparent distance of the light between infinity (1.0) and being at 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 Z Map input will be used to create the shadows depth map. Selections exist for the RGB and A, Luminance, and Z-buffer channels.


Output

The Output image can either contain the image with shadow applied or the shadow only.

Changing the setting in the output drop-down list controls which mode will be used. This 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.

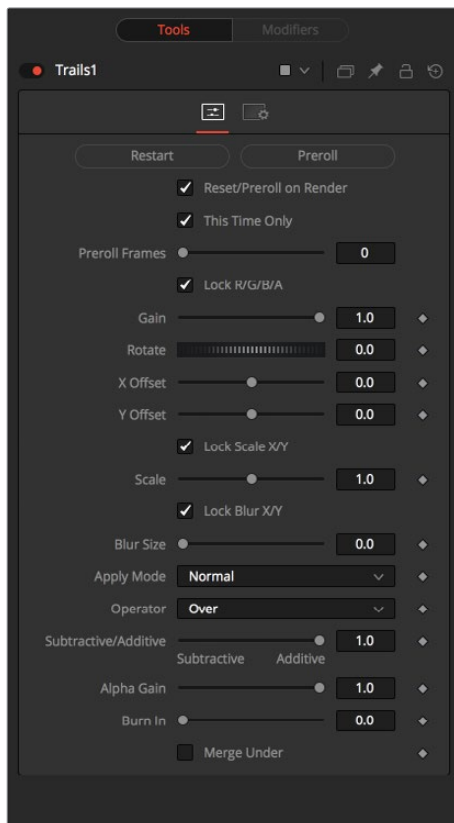
For example, Shadow Alpha controls the shadow's degree of transparency.

Trails [TRLS]



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.

Controls



Restart

This control clears the image buffer and displays a clean frame, without any of the ghosting effect.

Preroll

This will make the Trails node pre-render the effect by the number of frames on the slider.

Reset/Pre-Roll on Render

When this checkbox is enabled, the Trails node will reset itself when a preview or final render is initiated. It will Pre-roll the designated number of frames.

This Time Only

Selecting this checkbox will make the pre-roll use this current frame only and not the previous frames.

Number of Pre-roll Frames

This determines the number of frames to pre-roll.

Lock RGBA

When selected, this checkbox allow the Gain of the color channels to be controlled independently. This allows for tinting of the Trails effect.

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.

Lock Blur X/Y

When selected, this checkbox allows the blurring of the image buffer to be controlled separately for each axis.

Gain

The Gain control affects the overall intensity and brightness of the image in the buffer. Lower values in this parameter will create a much shorter, fainter trail, whereas higher values will 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.

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.

Scale

The Scale control resizes the image in the buffer before the current frame is merged into the effect.

Blur Size

The Blur Size control applies a blur to the image in the buffer before the current frame is merged into the effect.

Apply Mode

This menu is used to determine the method used by the Trails node when merging one sample over another. The methods in this menu are documented more completely in the Merge node's documentation in Chapter 72, "Composite Nodes."

Operator Mode

This menu is used to select to determine how the foreground and background are combined to produce a result. The menu will only be visible when the node's Apply mode is set to Normal.

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. Note that some modes not listed in the Operator dropdown (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.

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.

$x = 1, y = 1 - [\text{foreground alpha}]$

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.

$x = [\text{background alpha}], y = 0$

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.

$x = 1 - [\text{background alpha}], y = 0$

ATop

ATop places the foreground over the background only where the background has a matte.

$x = [\text{background alpha}], y = 1 - [\text{foreground alpha}]$

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.

$x = 1 - [\text{background alpha}], y = 1 - [\text{foreground alpha}]$


Subtractive/Additive, Alpha Gain, Burn In

For details on these controls and their effect, consult the documentation for the Merge node in Chapter 72, "Composite Nodes."

Merge Under

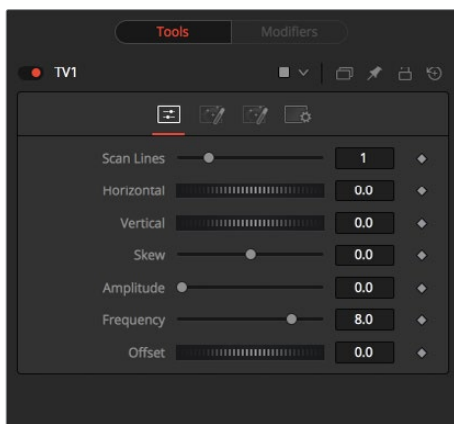
This merges the current image under the generated trail, rather than the usual, over top operation.

TV [TV]



The TV node is a simple node designed to mimic some of the typical flaws seen in analog television broadcasts and screens.

Controls



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) will drop every second line. A value of 2 shows one line, 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 will increase the intensity of the deformation. Use frequency 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.

Noise Tab



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 will be static. Change the value over time to cause the static to change from frame to frame.

Roll Bar Tab



Bar Strength

At the default value of 0 no bar is drawn. The higher the value, the darker the area covered by the bar will become.

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.

Chapter 36


Film Nodes

This chapter details the Film nodes available in Fusion.

Contents

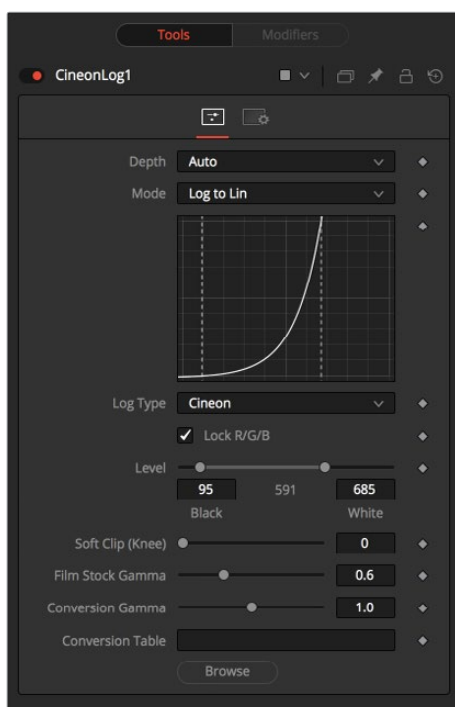
Cineon Log [LOG]	838
Film Grain [FGR]	840
Grain [GRN]	843
Light Trim [LT]	846
Remove Noise [RN]	846

Cineon Log [LOG]



The Cineon Log node is used to convert image data from logarithmic to linear. Use this node if the log-lin conversion was bypassed in a Cineon Loader to return processing to Linear.

Controls



Mode

These buttons offer two options, one for converting log images to linear and one for converting linear images to logarithmic.

Lock RGB

When selected, the settings in this tab will affect all of the color channels equally.

De-select this control to convert the red, green and blue channels of the image using separate settings for each channel.

Black Level and White Level

Use this control to set the black point and white point in the log image before converting. Pixels with values in log space below the Black Level will become out of range values below 0.0. Pixels with values above the White Level will become out of range values above 1.0 after conversion.

When processing in floating point color space, the out of range values stored in the Cineon format are preserved through the conversion. In 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 will immediately cause 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 Look Up Table can be created with specific conversion values. The ASCII LUT file can be loaded using the File Folder Icon button.

Black Rolloff

Since a mathematical $\log()$ operation on a value of zero or lower will result in invalid values, Fusion will clip 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 very small 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 the 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 will crush most of the image's range into very very small values. This is followed by a Custom node (described below), 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 will lift 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 falloff from the native Fusion process



The falloff from the ramped clipping process

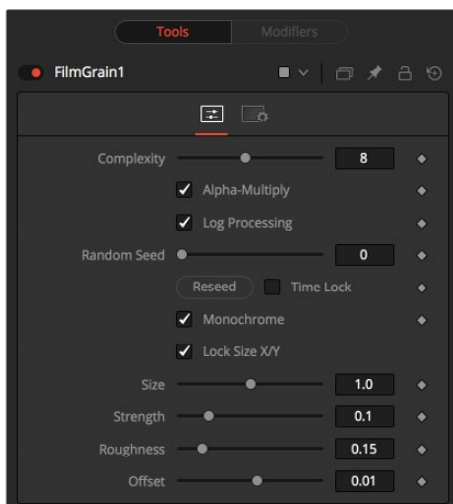
Film Grain [FGR]



The Film Grain node takes a new approach to grain, which should be more closely aligned with the grain profiles of modern film stocks. This provides you with more control over the final appearance of the grain.

NOTE: The Film Grain node does not replace the original Grain node, which is still provided to allow older compositions to load and render, but in almost every case, use of the new Film Grain node is now encouraged.

Controls



Complexity

The complexity of the Grain 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 complex results, without the apparent regularity often perceivable in digitally-produced grain.

Alpha Multiply

When the Alpha Multiply checkbox is enabled, the Film Grain node will multiply its results by the source images 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 semi-transparent 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 will ensure that the strength of the grain is accurate.

Log Processing

When this checkbox is enabled (default), the grain applied to the image will have its intensity applied non-linearly to match the grain profile of most film. Roughly speaking, the intensity of the grain will increase exponentially from black to white. When this checkbox is disabled the grain will be 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 very little 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 will appear completely 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 response to light of the red, green, and blue layers of the film also cause each color component of the film to present a different grain profile, typically with the blue channel presenting the largest amount of grain.

As a result, the most important control in the new 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 will closely mimic the results of preceding the old grain node with a Linear > Log conversion and following with a Log > Linear conversion immediately after.

Seed

The 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 on the Randomize button to randomly select a new seed value, or adjust the slider to manually select a new seed value.

Monochrome

When the Monochrome checkbox is enabled (default) the grain will be 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, so that changing the resolution of the image does not impact the relative appearance of the grain. The default Grain Size of 1.0 will produce 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 will produce an output value of $p \pm \text{strength}$. 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 over simplification, especially when the complexity exceeds 1. Enabling the Log Processing checkbox will also cause that variation to be affected such that there will be 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 you will observe that the grain produced has a very even luminance variation across the whole image. Increase the roughness to 1.0 and you will 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 0.1 to receive grain as if its value was actually 0.2.


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



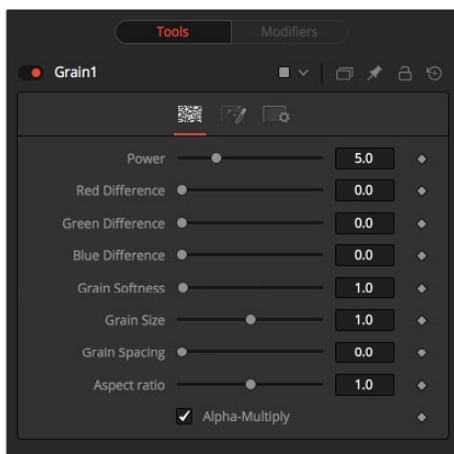
When Log processing is off, the Grain is applied evenly to the entire image, as shown here

Grain [GRN]



The Grain node offers comprehensive film grain emulation. This is useful for introducing simulated grain into a video or computer-generated image and matching existing grain in a given scene.

Grain Tab



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 will multiply the image by the alpha, clearing the black areas of any grain effect.

Spread Tab



Edit RGB Checkboxes

Separate red, green, and blue to enable the custom curves for each channel. More grain would appear in the blue channel than the red, and the green channel would receive the least. This curve mimics usual film responses. Right-clicking in the spline area will display a contextual menu containing options related to modifying spline curves. A complete description of the LUT Editor control and its options can be found in Chapter 82, "LUT Nodes."

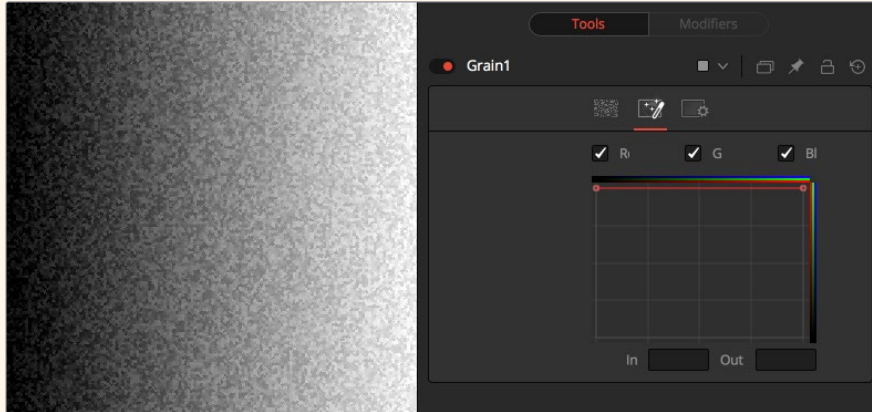
In and Out

This control gives direct editing of points on the curve by setting In/Out point values.

Examples

Default Spread

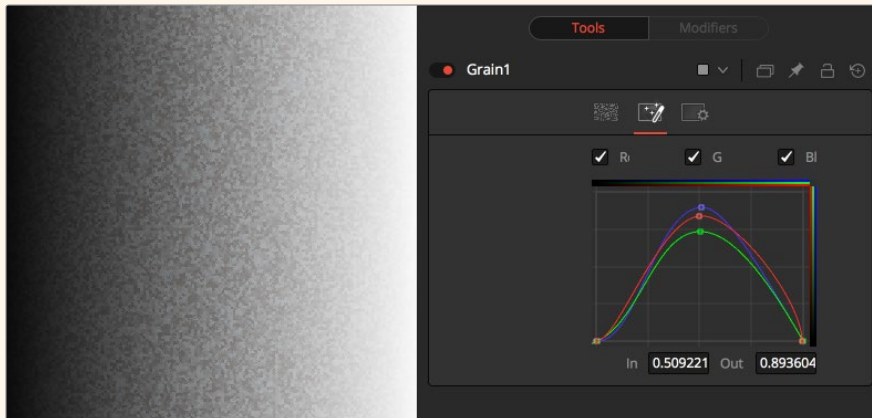
In the default setting, the Grain is applied evenly to the entire image as shown here. However, film often shows a different amount of grain in the blacks, mids and whites.



Default Spread

Bell Shape Spread


Setting a bell shape is often a good starting point to create more realistic looking grain. Here we have a non-uniform distribution with different amounts of grain in the red, green, and blue channel as well.



Bell-Shaped Spread


In both examples, the Grain's power has been over-exaggerated to show the effect a bit better.

Light Trim [LT]



This node emulates film scanner light trims. By design, this node works best with logarithmic data, such as the images stored by the Cineon file system. When logarithmic data is provided, the Light Trim node can be used to increase or decrease the apparent exposure level of the image.

Controls




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, de-select this checkbox.

Trim

This slider shifts the color in film, optical printing and lab printing points. 8 points equal one stop of exposure.

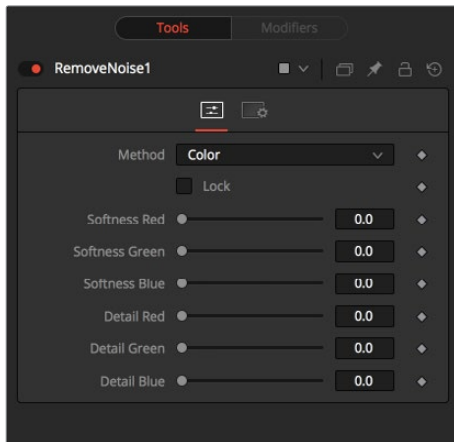
Remove Noise [RN]



The Remove Noise node provides simple noise management. The basic principle of its operation is the node blurs the image channels, 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.

Controls / Color



Method

Use these buttons to choose whether the node processes color using Color or Chroma methods. This will also give you a different set of control sliders.

Lock

This checkbox will link 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.

Controls / Chroma



Method

The same principle applies as with the RGB Controls, but in this case the Luminance and Chrominance channels are influenced.

Chapter 37


Filter Nodes

This chapter details the Filter node available in Fusion.

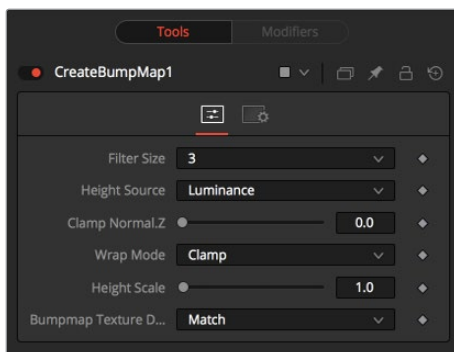
Contents

CreateBumpMap	850
Custom Filter Node	851
ErodeDilate Node	853
Filter Node	854
Rank Filter Node	856

CreateBumpMap



The Create Bumpmap node converts a grayscale (height map) image into a bump map. Check the notes for the naming conventions used in Fusion. Since the resulting bump vector information is represented as RGB, it's possible to modify them using all the image processing nodes in Fusion. Use the Create BumpMap node for applying to a material.



Input Port

CreateBumpMap.Input (white): Receives the RGBA channels from an image output for the bump calculation.

Controls

This tab contains all parameters for the node.

Filter Size

The process of generating the bump information is basically a Custom Filter. This multi-button control sets the filter size.

Height Source

Set the channel from where to extract the grayscale information from.

Clamp Normal.Z

Clips the lower values of the blue channel in the resulting bump texture.

Filter Wrap Mode

Basically wraps the image at the borders, so the filter produces correct result when using seamless tileable textures.

Height Scale

Changes the contrast of the resulting values in the bump map. Increasing this value yields in a more visible bump map.

Bump Map Texture Depth

Optionally converts the resulting bump texture into the desired bit depth.

Notes on Bump Maps

This tab contains all parameters for the node.

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).

Custom Filter Node



The Custom 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 context menu.

A Kernel filter is an array (or grid) of either 3 x 3, 5 x 5, or 7 x 7 values. 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 one multiplies the pixel's effect on the result. Negative values can also be entered, where the value of the pixel will be subtracted from the average. Only integer values can be entered; 0.5 is not valid.

For example, a filter with the values...

```
0 0 0
0 1 0
0 0 0
```

...will have zero effect from its neighboring pixels and the resulting image would be unchanged. A blurring effect would be...

```
1 1 1
1 1 1
1 1 1
```

...where the neighboring pixels are averaged with the center, resulting in a softening effect.

```
-5 0 0
0 1 0
0 0 5
```

This example will subtract five times the value from the top left and add five times the value from the lower right. If parts of the image that is processed are very smooth in color, the neighboring values will be very similar. In parts of the image where the pixels are different (e.g., an edge), the results will be different and tend to highlight or emboss edges in the image.

Using the values...

```
1 1 1
1 1 1
1 1 1
```

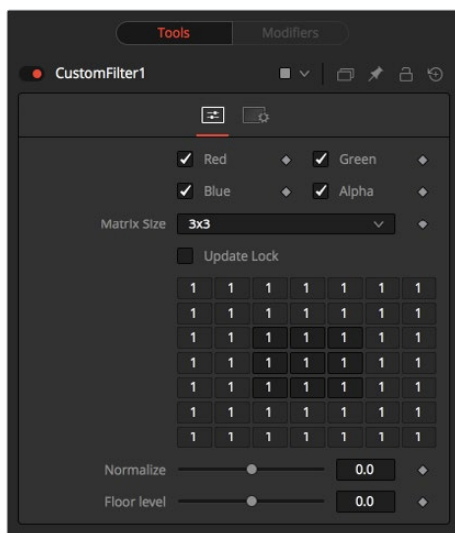
...in a filter and sliding Normalize to Positive will make the image go brighter or glow, simulating film over-exposure.

Using the values...

```
-1 0 0
0 0 0
0 0 1
```

...in a filter and sliding Floor Level to Positive will look like a Relief filter.

Controls



Color Channels (RGBA)

The color corrector defaults to operating on R, G, B, and A channels. Selective channel editing is possible by clicking the checkboxes beside each channel, making selected channels active or inactive.

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 will cause 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

Use this drop-down 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 will not render the filter. This is useful for setting up each value of the filter, then turning Update Lock off 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 will be analyzed. If a larger Matrix size is set, more of the text boxes in the grid will be enabled for input.


Normalize

This controls the amount of filter normalization that is applied to the result. Zero will give a normalized image. Positive values will brighten or raise the level of the filter result. Negative values will darken or lower the level.

Floor Level


This will add or subtract a minimum, or Floor Level, to the result of the filtered image. Zero will not add anything to the image. Positive values will add to the filtered image and negative values will subtract from the image.

ErodeDilate Node



This simple node erodes or dilates the image, depending on whether the Amount slider is set to a negative or positive value.

Controls



Color Channels (RGBA)

The filter defaults to operating on R, G, B, and A channels. Selective channel filtering is possible by clicking the checkboxes beside each channel to make them active or inactive. 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 will cause 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.

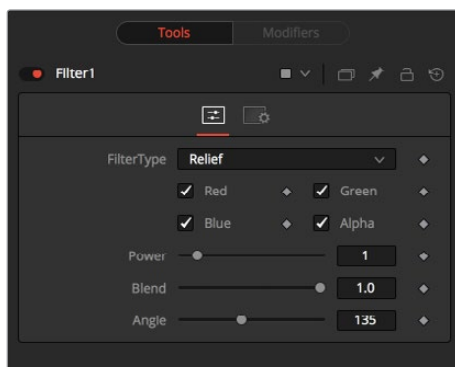
A positive value for Amount causes the image to dilate, similar to the effect of overexposing a camera. Regions of high luminance and brightness grow, eating away at the darker regions of the image. Both techniques will eradicate fine detail in the image and tend to posterize fine gradients.

Filter Node



The Filter node contains several pre-set filters, enabling a variety of effects, from radically changing the look of an image to adding subtle randomly-generated film grain.

Controls



Filter Type

This drop-down control provides a selection of filter types. The options are displayed below.

Relief

This appears to press the image into metal, such as an image on a coin. The image will appear to be bumped and overlaid on gray.

Emboss Over

Embosses the image over 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 amazing 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 mid-range). 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 defaults to operating on R, G, B, and A channels. Selective channel filtering is possible by clicking the checkboxes beside each channel to make them active or inactive.

Power

Values range from 1 to 10. It proportionately increases the amount by which the selected filter affects the image. This does not apply to Laplacian filter type.

Angle

This control has a range from 0 to 315 degrees and changes the effect in increments of 45 degrees. This only applies to the Relief and Emboss filters.

Median

Depending on which Filter Type is selected, the control may appear. It varies the Median filter's effect. A value of 0.5 will produce the true median result, as it will find the middle values. A value of 0.0 will find the minimums and 1.0 will find the maximums. Applies to the Median filter only.

Seed

This control is only visible 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 will always produce the same result given the same seed value.

Animated

This control is only visible 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, remove the selection from this checkbox.

Rank Filter Node



The Rank Filter is a very simple node. This filter examines nearby pixels, sorts the pixels by value, and then replaces the color of all the pixels examined by the color of the pixel with the specified rank.

For example, a Rank filter with a size of 1 will sample 9 adjacent pixels for their values. The following shows some example values, sorted by value.


0.0
0.4
0.4
0.5
0.5
0.7
0.7
0.7

A Blur filter with the same size would average these values together to produce a value of 3.9.

A Rank filter lets you choose a value from the list to select. For example, a rank of four would select the fourth item in the list, producing a value of 0.5

A Median filter is simply a linear Rank filter that selects the median, or middle value, from the list of sorted values.

Controls



Size

This control determines the Size in pixels of the area sampled by the filter. A value of 1 will sample only the pixels immediately adjacent to the center pixel, while larger values sample from a larger area.

Small values are excellent for removing salt and pepper style noise, while larger values produce an effect similar to water color paintings.

Rank

The Rank determines which value from the sampled values will be chosen, where 0 is the lowest value and 1 is the brightest.

Chapter 38


Flow Nodes

This chapter details the Sticky Note and Underlay nodes available in Fusion.

Contents

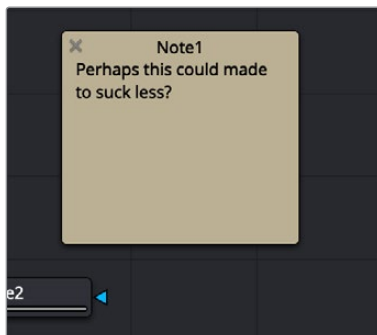
Sticky Note [NTE]	860
Underlay [UND]	861

Sticky Note [NTE]



Sticky Notes are a useful way of providing notes, comments and history for a specific area of a comp. By changing their size and color, they can provide unobtrusive comments or crucial notices, as required. Sticky Notes make a good complement to the Comments tab.

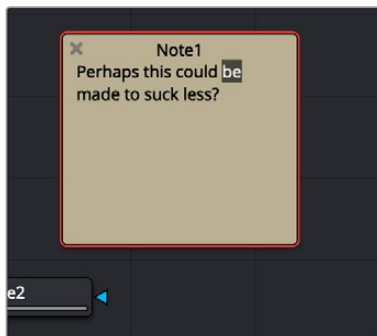
Usage



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. The new Sticky Note will be created with its title bar centered on the last click position. Alternately, 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.


Like Groups, Sticky Notes are created in a smaller, collapsed form. They can be expanded by double-clicking anywhere on them, or by clicking the icon in the top right corner. Once expanded, they can be resized by dragging on any side or corner. To collapse them again, click the icon in the top right corner once more.

Sticky Notes can be renamed, deleted, copied, and pasted and have their node color and text color changed, using the Node Editor’s contextual menu like any other node. It is also possible to lock them to prevent editing.



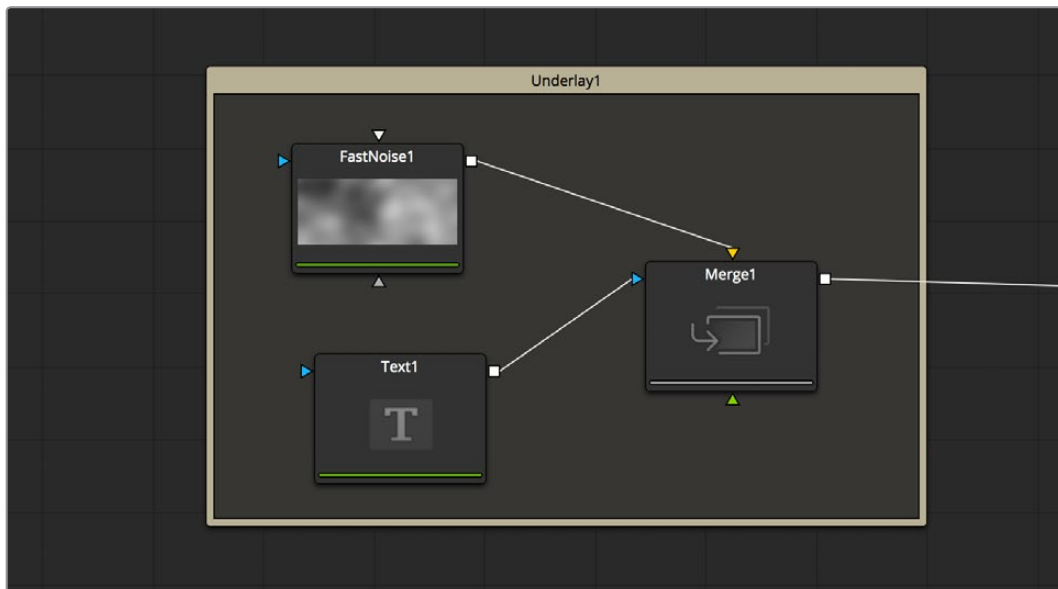
To edit the text in a Sticky Note, first expand it by double-clicking anywhere on the note, then click below its title bar. If the note is not locked, the text will become editable.

Underlay [UND]



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



As with conventional nodes, an Underlay can be added to a comp by selecting it from the Nodes menu, in the Node category, or from the Node Editor's Add Node contextual menu. The new Underlay will be created with its title bar centered on the last click position.

Underlays can be resized by dragging on any side or corner. This will not affect any nodes.

Underlays can also be used as simple selection groups. Activating an Underlay, by clicking on 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.

Chapter 39

Flow Organizational Nodes

This chapter details the Group, Macro, and Pipe Router nodes, which are designed to help organize your compositions, making the flow easier to see and understand.

Contents

Groups [NTE]	864
Macro [-/-]	865
Pipe Router [-/-]	867

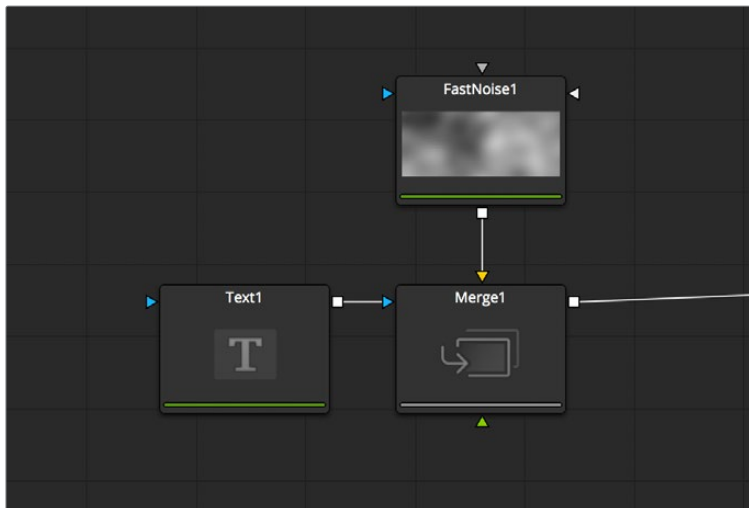
Groups [NTE]



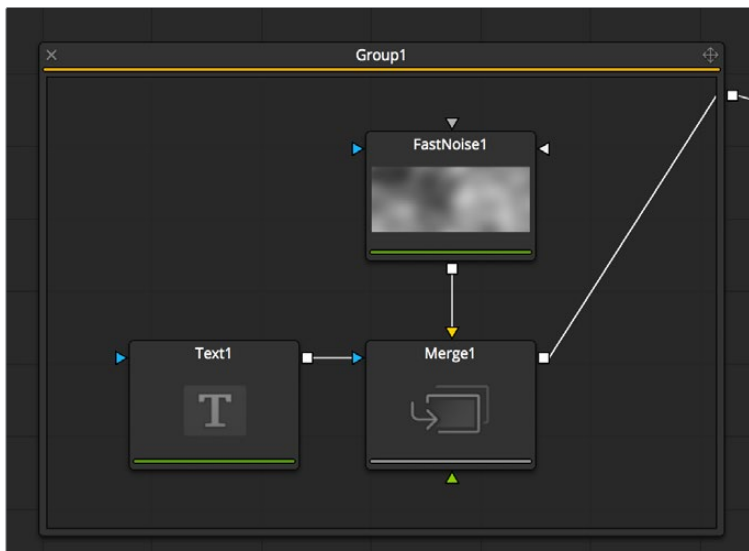
Groups can be used to neatly organize your comps, especially by putting together multiple nodes.

Usage

- To group nodes, select them and either press Command-G or Ctrl-G or right-click and select Group from the contextual menu. Press Command-E or Ctrl-E to either Expand or Collapse a selected Group.



- Right-click and select Ungroup from the contextual menu to ungroup all nodes.



- When opened, Groups will hover over existing elements, allowing editing of the enclosed nodes.

Macro [-/-]

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 for building your own customized nodes.

Usage



Creating a Macro

To create a Macro, select the nodes intended for the macro, right-click on any of them and select Macro > Create Macro from the contextual menu.

Macro Editor

The Macro Editor then allows you to specify and rename the controls that shall be exposed in the final macro.

In this example we only expose the Threshold and Gain sliders of a Softglow node, which is sandwiched between two ChannelBooleans.

After setting up the Macro to your liking, type in a name in the Macro Name field and select File > Save.



To add the Macro to your node tree, right-click anywhere on the node tree and select Macro > [NameOfYourMacro] from the contextual menu.

The Final Macro

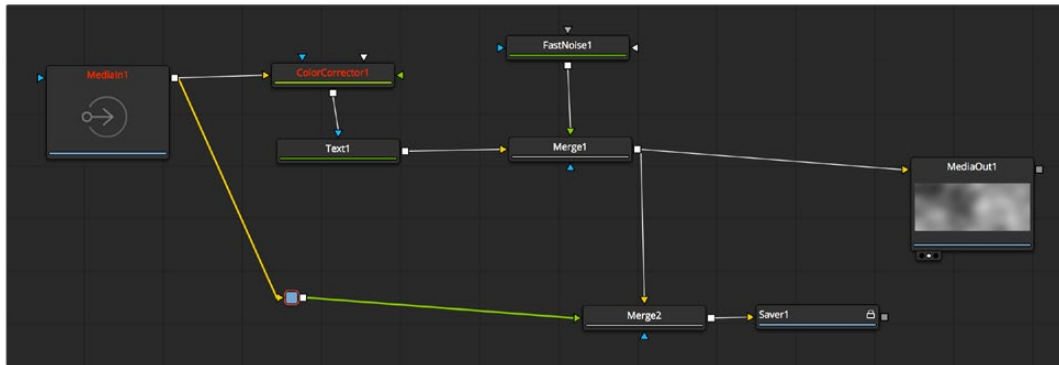
The final Macro looks and behaves just like any other node in Fusion.

As another example, you could take a single ChannelBoolean, set it to Add mode, and make it into a macro exposing no controls at all, thus creating the equivalent of a Add Node like the one that can be found in programs like Shake.

Pipe Router [-/-]

Pipe Routers can be used to neatly organize your comps by creating diversions in your pipes to make them better visible and help you understand the flow more easily. Pipe Routers do not have any influence on render times.


Usage



Pipe Router

To insert a Router in your pipe, Option or Alt left-click on the pipe and place the Router anywhere on the node tree.

Though Routers have no actual controls, they still can be used to add comments to your comp.



An example comment in a PipeRouter node

Chapter 40

Fuses

This chapter introduces Fuses, scriptable plug-ins that can be used within Fusion.

Contents

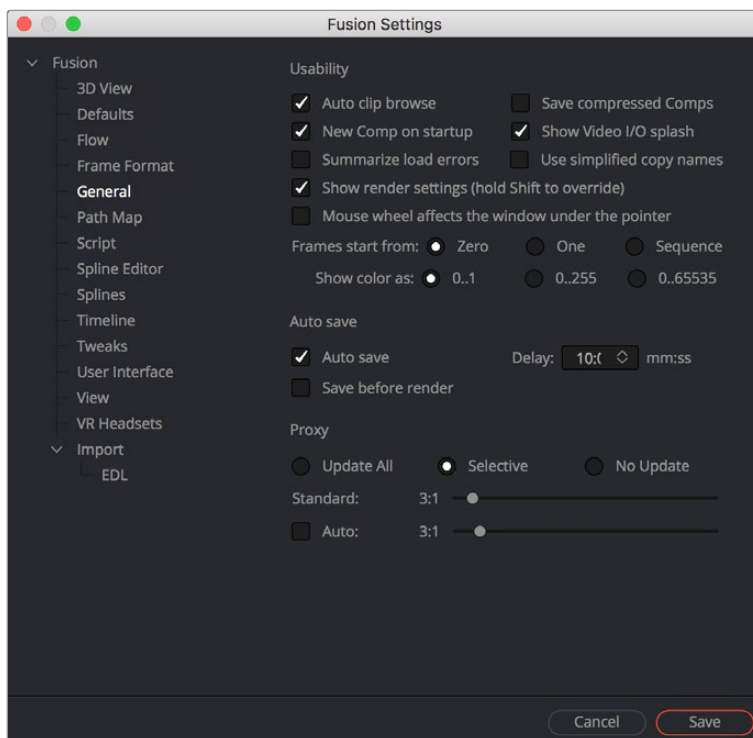
Fuses [FUS]	870
-------------	-----

Fuses [FUS]



Fuses are scriptable plug-ins. They can be edited within Fusion by pressing the Edit button, and will compile on-the-fly when clicking on Reload.

Fuses are a great way to prototype plug-ins or to deal with experimental code. Please refer to the SDK for further information.



Chapter 41

I/O Nodes

This chapter details the Loader and Saver nodes within Fusion.

Contents

Loader [LD]	873
Saver [SV]	883

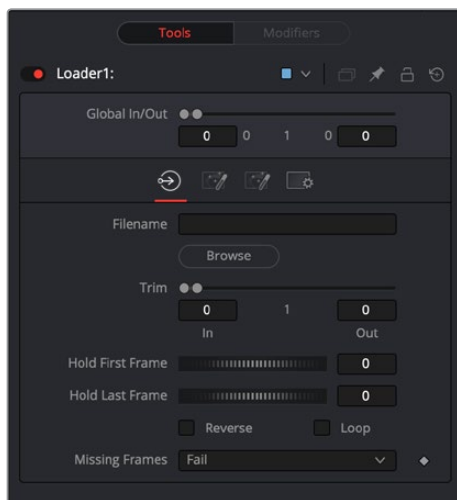
Loader [LD]



The Loader tool is used to select and load footage from the hard drive or network storage.

The Loader is responsible for trimming, looping, and extending the footage, as well as setting the field order, pixel aspect and color depth. Loader is arguably the most important tool in Fusion.

File Tab



Filename

Clicking on the yellow Folder button will display a standard Fusion file browser. The path to the footage can also be typed directly using the text box provided.

The text box supports filename completion. As the name of a directory or file is typed in the text box, Fusion will display a popup that lists possible matches. Use the arrow keys to select the correct match and complete the path.

When a Loader is added to the flow, a file dialog will appear automatically to allow selection of a clip from the hard drives. To select footage later, select Cancel. The Loader will still be added to the flow. Disable the automatic display of the file browser by disabling Auto Clip Browse in the Global > General Preferences.

File Sequences

It is common practice to use file sequences to identify an image that is part of a series. If the last part of a file's name is a number (not counting file extension), Fusion will automatically scan the directory looking for files that match the sequence.

For example, the following filenames would be valid sequences.

image.0001.tga

image.0002.tga

image.0003.tga

image151.exr

image152.exr

image153.exr

image1.tif

image2.tif

image3.tif

The following would not be considered a sequence, since the last characters are not numeric.

shot.1.fg.tga

shot.2.fg.tga

shot.3.fg.tga

It is not necessary to select the first file in the sequence; Fusion will search the entire folder for files matching the sequence in the selected filename.

Also, Fusion will determine 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.tga

image.0100.tga

Fusion will see this as a file sequence with 100 frames, not a file 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 contextual 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 only want to load a single frame out of a sequence, like, for example, a photograph out of a folder containing many other files as well. By default, Fusion will detect those as a sequence, but if you hold Shift while dragging the file from Explorer to the Flow Editor, Fusion will only take that specific file and disregard any sequencing.

Also, if you use the Preview controls on the bottom right of the screen, you can Command-click or Ctrl-click on the slider once you reached the frame you want to use as a still, and Fusion will set up the loader accordingly.

Proxy Filename

The Proxy Filename control only appears once the filename control points to a valid clip. This 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.

For example, create a 1/4 scale version of a Cineon film sequence to use as a file proxy. Whenever the proxy mode of the flow is enabled, the smaller resolution proxy clip will be loaded from disk and all processing will be performed at the lower resolution, significantly improving render times.

This is particularly useful when working with large film plates stored on a remote files 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 exactly the same format as the main files. In the case of formats with options, such as Cineon, DPX and OpenEXR, the proxies will use the same format options as the primary.

Trim In and Out

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 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 value of 95 would stop loading frames after the 95th.

Right-click on the range control to see three options appear in a submenu

Autodetect Clip Length

This rescans the clip to see if frames have been added or removed since the last time the clip was scanned (when it was loaded).

Set Clip Start Frame

Selecting this will display a dialog box to set the number of the actual start frame of the clip on disk. Use this when the first frames of the file sequence are not yet available to set the clip's properties manually.

Set Clip Length

This sets the actual clip length, overriding the scanned length. This is useful if the entire clip has not yet been rendered or captured. Used in combination with Set Clip Start Frame, it can be used to define a clip before the clip is available.

For example, imagine there is only a single animatic frame available for a scene that is being precomped. The animation department has yet to render the entire layer. Use the Set Clip Start and Set Clip Length to define the clip anyway, and the composition will load the correct frames when they become available. If the Missing Frames control is set to Wait, then rendering will pause until the frame is available, instead of failing.

Hold First and 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 played last.

Loop

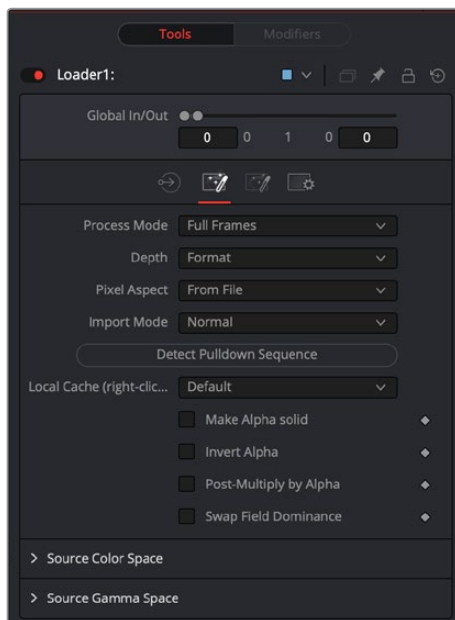
Select this checkbox to loop the footage until the end of the project. Any lengthening of the clip using Hold First/Last Frame or shortening using Trim In/Out is included in the looped clip.

Missing Frames

The Missing Frames drop-down menu provides options for selecting how the Loader will behave when an expected frame is missing from the footage, or is unable to load for any reason.

- **Fail:** The Loader will not output any image unless a frame becomes available. Rendering will abort.
- **Hold Previous:** Output the last valid frame 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:** Output a black frame until a valid frame becomes available again.
- **Wait:** Fusion will wait for the frame to become available, checking every few seconds. Useful for rendering a flow simultaneously with a 3D render. All rendering will cease until the frame appears.

Import Tab



Process Mode

Use this menu to select the Fields Processing mode used by Fusion when loading the image. The default option is determined by the Has Fields checkbox control in the Frame Format preferences, and the default height as well. Available options are Full frames, NTSC fields, PAL/HD fields, PAL/HD fields (reversed) and NTSC fields (reversed). The two Reversed options load fields in the opposite order, and will thus result in the fields being swapped both in time order and in vertical order as well. Use the Reverse Dominance checkbox (described in the Import tab below) to swap fields in time only.

Depth

The buttons in this array are used to select the Color Depth used by Fusion 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, TGA files will automatically process at 8-bit because the TGA file format does not store color depths greater than 8. Cineon files will load at Float, etc. If the color depth of the format is undetermined, the default depth defined in the flow's Frame Format preferences is used.

Formats that support multiple color depths will be 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 in which the image will be processed. For a more complete discussion of color depth and how it affects composites, consult the Frame Formats chapter of this manual.

Pixel Aspect

This button array is used to determine the image's pixel aspect ratio.

From File

The loader will conform 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 will be ignored. The pixel aspect set in the composition's Frame Format preferences will be used instead.
- **Custom:** Select this option to override the preferences and set a pixel aspect for the clip manually. Selecting this button will cause the X/Y Pixel Aspect control to appear. For a more complete discussion of pixel aspect and how it affects composites, consult the Frame Formats chapter of this manual.

Custom Pixel Aspect

This control is only visible when the Custom Pixel Aspect method is used. Enter the desired 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 24fps footage into 30fps. 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 will remove existing 3:2 pull-down applied to the image sequence, converting from 30fps to 24fps.
- **2:3 Pull Down:** The footage will have pull-down applied, converting 24fps footage to 30fps by creating 5 frames out of every four. The Process mode of a loader set to Pull Down should always be Full Frames.

First Frame

This control is used to determine which frame of the 3:2 sequence is used as the first frame of the loaded clip. It only appears if the Pull-up or Pull-down options are selected from the Import Mode menu.

Detect Pulldown Sequence

Pressing this button will cause Fusion to attempt to automatically detect and set the Pull-up sequence of the footage. It only appears if Pull-up or Pull-down is selected from the Import Mode menu. If it succeeds in detecting the order, the First Frame control will automatically be set to the correct value.

When checked, the original alpha channel of the clip will be cleared and set to solid white (completely opaque).

Invert Alpha

When checked, the original alpha channel of the clip will be inverted. This may also be used in conjunction with Make Alpha Solid to set the alpha to pure black (completely transparent).

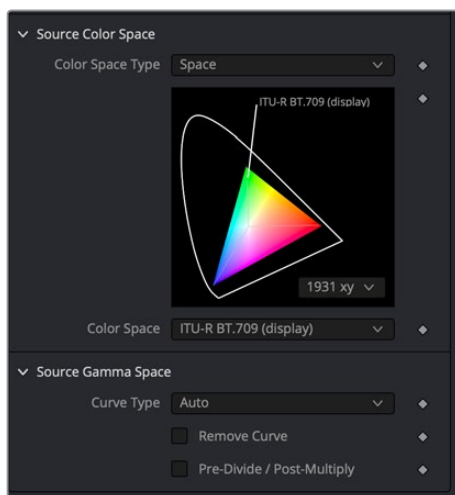
Post-multiply by Alpha

Selecting this option will cause 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 this control is selected, the field order (dominance) of the image will be swapped, so that the order in time that the fields appear in is reversed. Unlike the Process Mode control, this is done without spatially swapping the scanlines in the image.

Source Color Space



Sets the Color Space of the footage to help achieve a linear workflow.

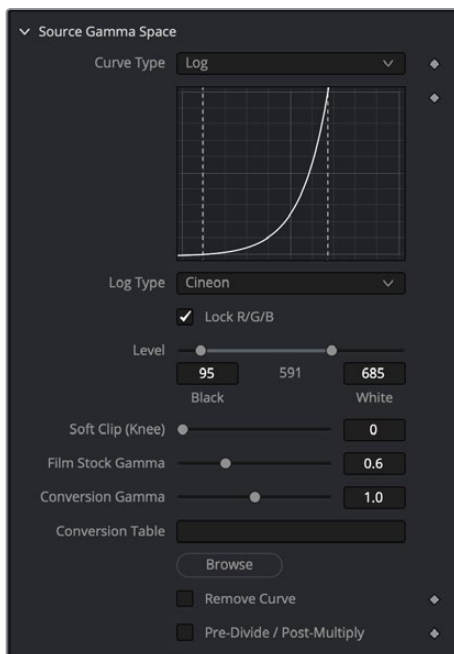
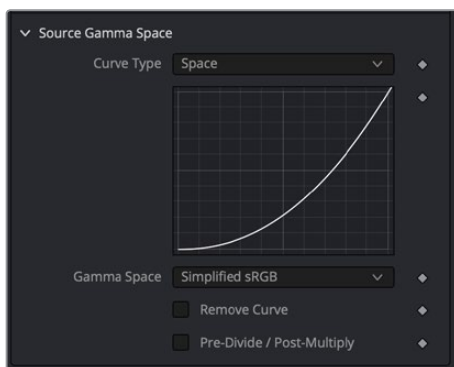
Unlike the Gamut tool, this doesn't perform any actual color space conversion, but rather adds the source space data into the metadata, if that metadata doesn't already exist. The metadata can then be used downstream by a Gamut tool with the From Image option, or in a Saver, if explicit output spaces 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

Determines the Gamma Space of the footage and gives 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 the user to set the Gamma Space from a variety of options.
- **Log:** Brings up the Log/Lin settings, similar to the Cineon Tool.

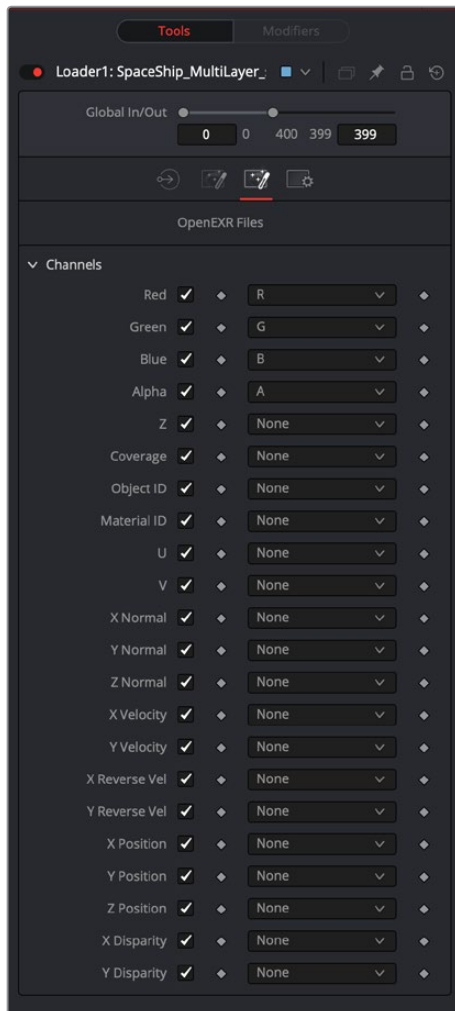


Remove Curve

Depending on the selected Gamma Space 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 information, options and settings relative to loading the image format. Not all file formats have options. Notably, the Cineon, DPX, PSD, OMF, OpenEXR and QuickTime formats all provide additional options when loaded. See Appendix A, File Formats, for a description of all supported formats.

Options for JPG and DPX are shown for reference on the right.

OMF Format

An OMF file is not just a media format; it can contain information about edits, multiple sequences, and even multiple versions of the same shot. If multiple clips exist within the OMF file, the format options may be used to select the desired clip. To import all clips from an OMF file, complete with edit and combine information, use the Import OMF option from Fusion's File menu to create a entire composition from the contents of the OMF.

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 individually from a PSD file, with appropriate blend modes to combine them, use the Import PSD option from Fusion's File menu to create a entire composition from the contents of the PSD file.

See the chapter Film in Fusion and the chapter for the Cineon Log tool for a detailed discussion of the options found in the Cineon and DPX file formats.

OpenEXR Format

Industrial Light and Magic developed the OpenEXR format. Its initial design was to provide a compact and flexible format to support storage of high dynamic range images (float). The format has been expanded to support a variety of extra channels and metadata.

The format options for OpenEXR files provide a mechanism for mapping any non-RGBA channels to the channels supported natively in Fusion. Enter the name of a channel 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 <http://www.openexr.com/>.

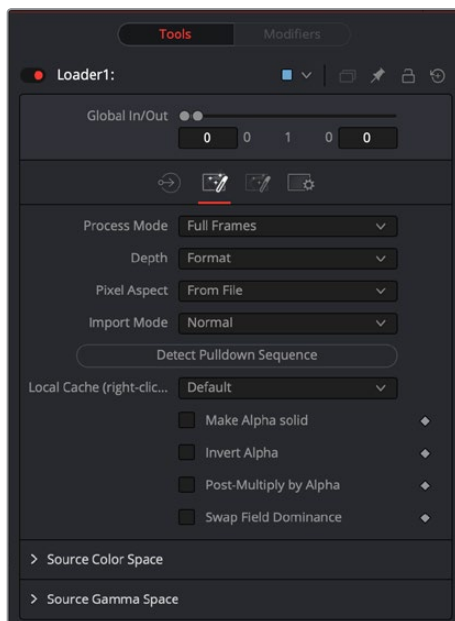
QuickTime

QuickTime files can potentially contain multiple tracks. Use the format options to select one of the tracks.

Image File Lists

One of the more interesting formats supported by Fusion is the Image File List format. An IFL is actually a text file containing a list of files that should be loaded. Use IFL files to specify a sequence of arbitrarily named files, or files from different directories.

Common Controls



Global In and Out

Use this control to specify the position of this tool within the project. Use Global In to specify the frame on which that the clip starts and Global Out to specify the frame on which this clip ends within the project's Global Range. The tool will 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 amount of available frames in the clip, Fusion will automatically trim 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 amount of available frames in the clip, Fusion will automatically lengthen 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.

To slip 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 control.

The Magic Comp Variable

As you can see in the screenshots, the pathnames in this example start with **Comp:**.

The Comp variable in Fusion works for Loaders and Savers and helps you to keep your work organized. **Comp:** stands for the folder your actual composition is stored in.

So as long as all your source footage is stored in subfolders of your Comp folder, Fusion will find 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 will 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 MediaIn 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 MediaIn node would then be:

Comp:\..\Footage\ Greenscreen\0815Green_0000.dpx

Observe how the two dots .. set the directory to go up one folder. Pretty much the same like CD .. in a command shell window.

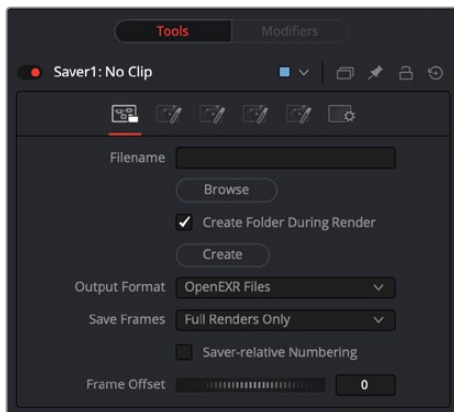
Saver [SV]



The Saver is responsible for writing the results of a composition to disk. It can be inserted into a composition at any point. The Saver tool can also be used to add scratch track audio to your flow, which can be heard during interactive playback.

A composition can contain any number of Saver tools, and Savers can be placed at any point in the flow.

File Tab



Filename

Use the Filename dialog to select the name and path of the rendered image output. Click on the yellow folder icon to display a file browser to select a folder and filename to be used for output.

Sequence numbering is automatically added to the filename when rendered. For example, if `c:\renders\image.tga` is entered as the filename and 30 frames of output is rendered, the files will automatically be numbered as `image0000.tga`, `image0001.tga`, `image0003.tga`...and so on. Four digit padding is automatically used for numbers lower than 10000.

Specify the number of digits to use for padding the sequence number by explicitly entering the digits into the filename.

For example, `image000000.tga` would apply 6 digit padding to the numeric sequence, `image.001.tga` would use 3 digit padding and `image1.tga` would use none.

Saving to an Image File List

It is possible to save to an IFL file if more control is required over the filenames written to disk than specified here. To create an IFL, create a text file with the extension .ifl.

Each line of the text file specifies the name for one frame of the output.

For example, to render using roman numerals instead of decimals for the file sequence numbering would create the following files:

```
image.i.tga  
image.ii.tga  
image.iii.tga  
image.iv.tga  
image.v.tga
```

The render will fail if there are not enough lines to match the number of frames rendered. In the above example, the render would fail after the fifth frame.

Format options cannot be specified when using image file lists to save to disk. The defaults will be used instead.

Output Format

Use this control to select the image format to be saved. 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.

For a list of file formats supported for saving by Fusion, see Appendix I.

Save Frames

This control selects between two modes. Full Renders Only

Images are only 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 will save each frame to disk as it is processed interactively.

When used correctly, this feature can completely eliminate the need to perform a final render after rotoscoping. It can cause tremendous confusion when used in conjunction with a flow 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 will not automatically be re-rendered. Either step through each frame again or perform a final render to make certain that the result is correct.

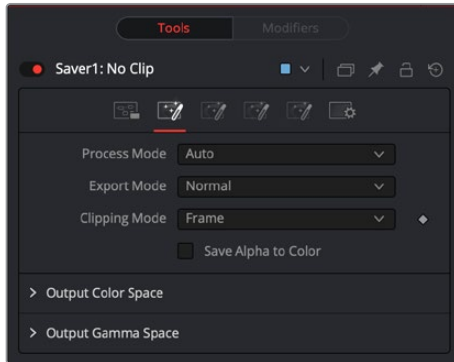
Saver-relative numbering

Normally, Fusion will use the render range of a composition to determine the numeric sequence used when rendering a file sequence to disk. Enable this checkbox to reveal the Sequence Start Frame control to adjust the starting number in the sequence to a custom value.

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 will normally be numbered 0001-0030. If the Sequence Start Frame is set to 100, the rendered output would be numbered from 100-131.

Export Tab



Process Mode

The Process Mode menu is used to set the method of field processing used by Fusion when the flow is rendered.

The default option is Auto. This will render the flow using the field ordering set in the Process mode of the most background Loader or Creator tool in the flow.

To force a flow to render individual fields or full frames regardless of the settings of other tools in the flow, change this menu's setting to one of the other listed options.

For example, if the Saver's process mode is set to NTSC fields, a Loader set to process full frames will still load a full frame from disk, but animation applied to that layer will be applied with field level accuracy. In all cases, footage will be re-interlaced, if necessary, and saved as frames on disk.

See the Frame Formats chapter for additional details on frame and fields processing in Fusion.

Export Mode

This menu is used to apply a SMPTE standard 3:2 pulldown to the footage as it is saved to disk, converting the footage from 24fps to 30 fps.

First Frame

First Frame determines the cadence of the 3:2 pulldown by choosing what frames are combined from the 24fps source to create the first frame in the 30fps pulldown result. Normally, this should be left to AA. When saving a clip that originally had pulldown that was removed to apply effects and it needs to be reinserted into an existing edit, the cadence of the original clip may need to be matched.

Clipping Mode

These radio buttons define how the edges of image should be treated. This can also be called source image clipping.

They default to Frame, which will provide the same behavior as previous versions of Fusion. Since this option will clip to the parts of the image visible within its visible dimensions, it will break any infinite-workspace behavior.

If the upstream DoD is smaller than the frame, the remaining area in the frame will be treated as black/transparent. None does not perform any source image clipping at all.

This means that any data that would normally be needed outside the upstream DoD will be treated as black/transparent. Be aware that this might create humongous images which can consume a huge amount of disk space. So you should use this option only if really needed.

For more information about ROI, DoD and Infinite Workspace, please see the dedicated chapter.

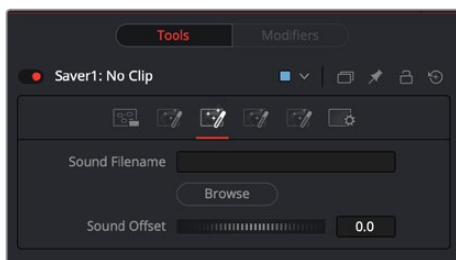
Save Alpha to Color

When selected, this control will cause the alpha channel to be saved into the color channels as a grayscale image. This will completely overwrite any existing color information.

Audio Tab

The Audio functionality is included in Fusion 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 will load the entire audio file into memory for efficient display of the waveform in the Timeline.

The audio track will be included in the saved image if a QuickTime or AVI file format is selected. Fusion currently supports playback of WAV audio.



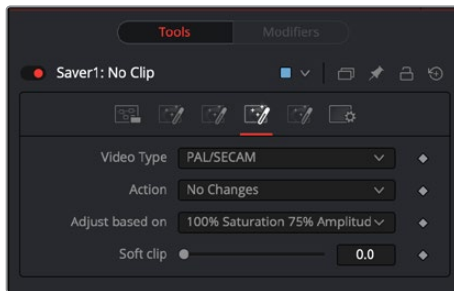
Audio Browse

Select the optional audio scratch track to be used. Select the *.WAV file of choice, then expand the Saver segment in the Timeline layout 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 create an offset in time between the audio and images of the flow.

Legal Tab



Video Type

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 will treat illegal colors in the image.

Adjust to Legal

This will cause 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.

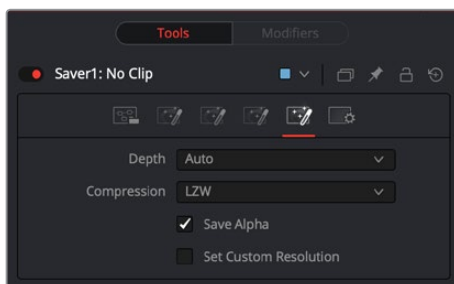
No Changes

This causes the images to be saved unaffected.

Adjust Based On

This menu is used to choose whether Fusion will legalize 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%.

Format Tab



The Format tab contains information, options and settings specific to the image format being saved. The controls for a TGA sequence will be entirely different from the ones displayed when a DPX file is saved.

TGA and DPX are displayed on the right for reference.

When the Saver is set to DPX, it's important to understand the reason for the "Data is Linear" option. When saving log data into a DPX, and not using the Saver'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 will think the data is linear, and will not perform any log-lin conversion.

The Magic Comp Variable

As you can see in the screenshots, the pathnames in this example start with **Comp:**.

The Comp variable in Fusion works for Loaders and Savers and helps you to keep your work organized. **Comp:** stands for the folder your actual composition is stored in.

So as long as all your source footage is stored in subfolders of your Comp folder, Fusion will find 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 will 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 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 would then be: **Comp:\..\Footage\Greenscreen\0815Green_0000.dpx**

Observe how the two dots .. set the directory to go up one folder. Pretty much the same like **CD ..** in a command shell window.

Chapter 42

LUT Nodes

This chapter details the LUT nodes available in Fusion.

Contents

File LUT [FLU]	891
LUT Cube Analyzer [LCA]	892
LUT Cube Apply [LCP]	893
LUT Cube Creator [LCC]	894

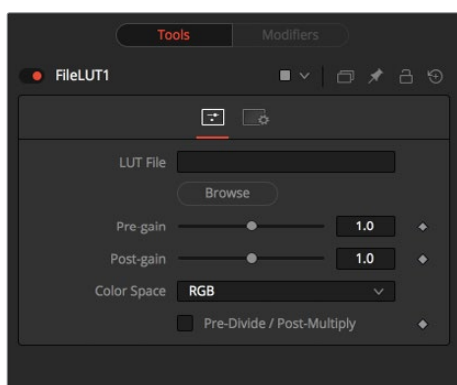
File LUT [FLU]



The File LUT node applies a Look up 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 the system or network.

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 very 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, simply by changing the contents of the LUT. This can be useful when the same LUT-based color correction is applied in many different compositions.

Controls



LUT File

Use this control to select the path to the file describing the LUT. Currently, this node supports LUTs exported from Fusion in .LUT and .ALUT formats, Shake's LUT format, and a number of 3D LUT formats as well. The node will fail with an error message on the Console if it is unable to find or load the specified file.

Color Space


Use this control 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 will cause the image pixel values to be divided by the alpha values prior to 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.

LUT Cube Analyzer [LCA]



The LUT Cube Analyzer takes an image created 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 image into the node would result in an unaltered or 1:1 LUT file.

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 will create a LUT that exactly resembles your color pipeline.

Usage

Connect the output of any node modifying an image that was created with the LUT Cube Creator to the input of the Analyzer, view the Analyzer, select the desired output format, specify a filename, and press Write File to create the 3D LUT.

Controls



Type

Select the desired output format of the 3D LUT.


Filename

Specify a path and name the LUT file is written to.

Write File

Press this button to generate the 3D LUT file based on the settings above.

LUT Cube Apply [LCP]



The LUT Cube Apply 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.

Usage


Connect any image that is meant to be modified according to the LUT to the Background input of the node. Connect a LUT Cube image to the Foreground input of the node.

LUT Cube Creator [LCC]



The LUT Cube Creator creates an image for further use with the LUT Cube Analyzer or LUT Cube Apply. 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.

Controls



Type

- **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.

Size

Determines the resolution of the color cube.

NOTE: Higher resolutions yield more accurate results but are also more memory and computational extensive.

Common values for color cubes are 17x17x17 or 33x33x33, or in a mathematical fashion n^2+1 .

Chapter 43

Mask Nodes

This chapter details the Mask nodes available in Fusion.

Contents

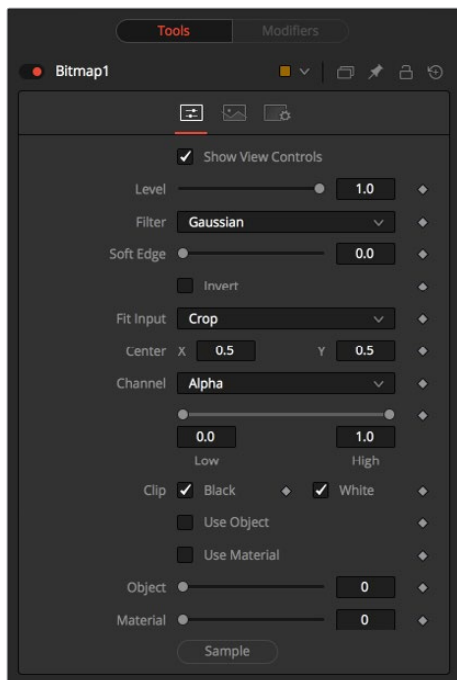
Common Mask Controls [CMC]	897
Bitmap Mask [BMP]	901
B-Spline Mask [BSP]	904
Ellipse Mask [ELP]	906
Mask Paint [PNM]	908
Polygon Mask [PLY]	909
Ranges Mask [RNG]	911
Rectangle Mask [REC]	912
Triangle Mask [TRI]	914
Wand Mask [WND]	915

Common Mask Controls [CMC]

Common Controls can be found in every Mask node. This chapter is therefore valid for all Mask nodes.

Common Mask Controls

Although each mask has its own set of controls unique to that mask type, several of the controls shown are common for all types of masks. The controls listed here are generally found on all masks.



Show View Controls

Use the Show View Controls checkbox to disable the display of the Mask controls in the Viewer. Polylines, centers, angles, and other controls will not be displayed, even when the node is selected.

Level

The Level control designates the transparency level of the pixels in the Mask channel. When the value is 1.0, the effect mask is completely opaque (unless it has a soft edge). Lower values will 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 will lower 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 will lower the values of all of the pixels in the Mask channel, regardless of the fact that the Rectangle mask beneath it is still opaque.

Filter

This control selects the filtering algorithm to be used when applying Soft Edge to the mask.

Box

This is the fastest method, but at reduced quality. This is best suited for very small 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 to let you control the quality. At 1 and 2 passes, results are identical to Box and Bartlett, respectively. At 4 and above, results are usually as good as Gaussian, in less time and with no edge “ringing.”

Gaussian

The default filter, this 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 extremely slight edge “ringing” on floating-point pixels.

Soft Edge

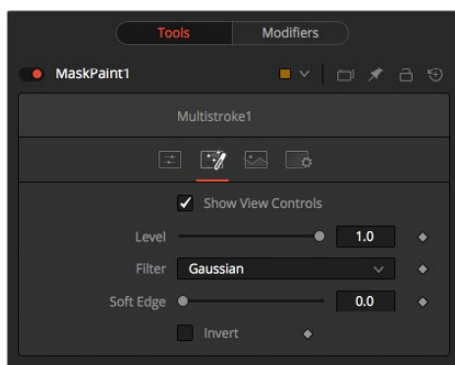
Use the Soft Edge slider to blur (feather) the edges of the mask, using the selected filter. Higher values will cause the edge to fade off well beyond the boundaries of the mask. A value of 0.0 will create a crisp, well-defined edge.

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, an outline of the mask shape is drawn and the thickness of the outline is handled with this control.

Paint Mode

Although each mask has its own set of controls unique to that mask type, several of the controls shown are common for all types of masks. The controls listed here are generally found on all masks.



Merge

Merge is the default for all masks. The new mask is merged together with the input mask.

Add

The mask’s values are added to the input mask’s values.

Subtract

In the intersecting areas, the new mask values are subtracted from the input mask's values.

Minimum

The input mask's values are compared to the new mask, and the lowest (minimum) value is taken.

Maximum

The input mask's values are compared to the new mask, and the highest (maximum) value is taken.

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. This differs from the Invert Paint mode in that it affects all pixels, regardless of whether they are covered by the new mask or not.

Solid

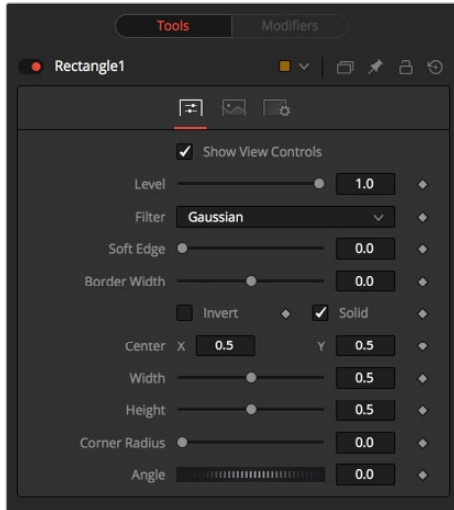
When the Solid checkbox is enabled, all areas completely enclosed by the mask will be filled solid white. Turning off the checkbox will treat the mask as an outline, with the width of the outline determined by the Border Width control. This checkbox is enabled by default.

Center

All masks have a center, with the exception of the common controls found on Creator nodes, such as Background and Fast Noise.

Process Mode

Use this menu control to select the Fields Processing mode used by Fusion to render changes to the mask. 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 mask created will be locked to values defined in the composition's Frame Format preferences. If the Frame Format preferences change, the resolution of the mask produced 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 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 display monitor) and an aspect of 0.91 would create a slightly rectangular pixel (like an NTSC monitor).

Depth

The Depth button array 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 will set the width, height, and pixel aspect to the values for that format.

Bitmap Mask [BMP]



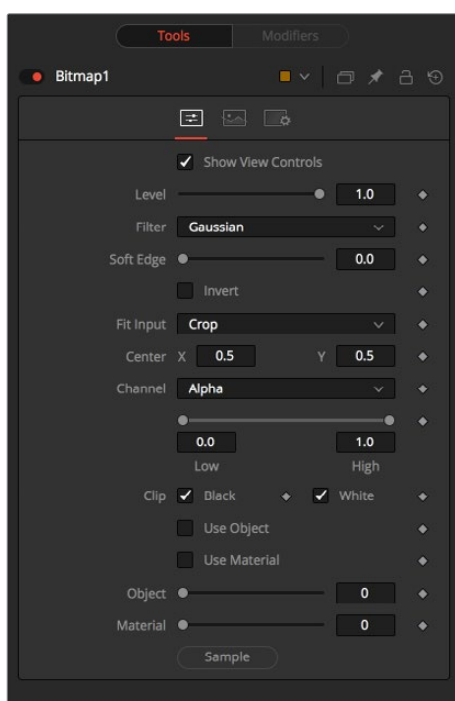
The Bitmap Mask allows images from the node tree to act as masks for nodes and effects. Bitmap masks can be based on values from any of the color, alpha, hue, saturation, luminance, and auxiliary coverage channels of the image. Nodes can also be masked based on the Object or Material ID of a 3D rendered image (provided those channels were included when the file was rendered).

The output of any node can be connected directly to another node's Effect Mask input. The Bitmap Mask node is not required for many common tasks. If the output is connected directly, the Common Controls tab for the masked node will display a control 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.

Many of the controls found in this Mask node are common to all Mask nodes. These controls are documented at Common Mask Controls.

Controls



Level, Filter, Soft Edge and Border Width

Please refer to the Common Mask Controls.

Fit Input


This control defines how the image source is treated if it does not fit the dimensions of the generated mask.

Imagine a 720*576 image source being used to generate a 1920x1080 mask.

In the following section we refer to image source as being the 720x576 image (yellow) and mask as being the mask that is actually generated (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 bigger 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 dimension 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.



Other Controls

Channel

Use this control to select 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 value of the low control will clip pixels below the specified value to black (0.0). Decreasing the high value will force pixels higher than the specified value to white (1.0).


Use Object/Use Material

This control has no affect unless the input image contains a Material or Object ID channel. When toggled on, the selected Object ID and/or Material ID is used to create a mask based on the selected object or material. When toggled off, the regular color channels will generate the mask.

Image Tab

Please refer to the Common Mask Controls.

B-Spline Mask [BSP]

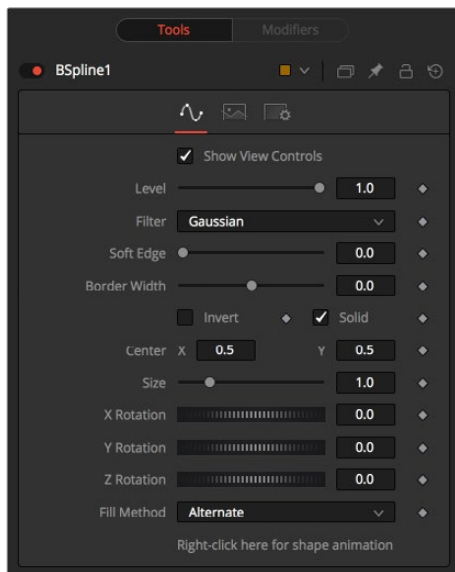


A B-spline Mask is identical to a Polygon Mask in all respects except one. Where Polygon masks use Bezier splines, this Mask node uses B-splines. Where Bezier splines employ a main point and two handles to manage the smoothing of the spline segment, a B-spline requires only a single point. This means that a B-spline shape requires far fewer control points to create a nicely smoothed shape.

The smoothness of a B-spline is determined by the tension of the control points. To adjust the tension of a B-spline's control points, select the point, 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.

Many of the controls found in this Mask node are common to all Mask nodes. These controls are documented at Common Mask Controls.

Controls



Level, Filter, Soft Edge and Border Width

Please refer to the Common Mask Controls.

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 effect mask along any axis.

Fill Method

The Fill Method drop-down 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 together.

Adding Points

Adding Points to a B-spline effect mask is relatively simple. Immediately after creating the mask there are no points, but the mask will be in Click Append mode. Simply 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.




When a B-spline mask is added to a node, a toolbar will appear in the view with buttons that offer easy access to the modes and nodes. Hold the mouse pointer over any button in the toolbar to display a tooltip that describes that button's function.

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.

Image Tab

Please refer to the Common Mask Controls.



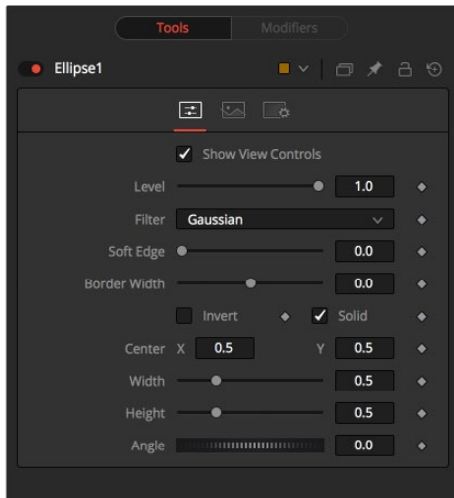
Ellipse Mask [ELP]



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.

Many of the controls found in this Mask node are common to all Mask nodes. These controls are documented at Common Mask Controls.

Controls



Level, Filter, Soft Edge and Border Width

Please refer to the Common Mask Controls.

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 view using the pointer. Any changes will be reflected on 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 on this control.


To change the mask's size without affecting the aspect ratio, drag the on screen 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 in the input boxes 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.

Image Tab

Please refer to the Common Mask Controls.



Mask Paint [PNM]

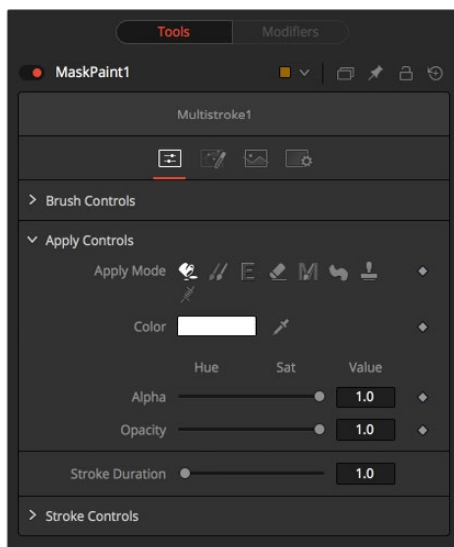


The Mask Paint node allows direct painting on mask images, using the mouse pointer as if it was a paintbrush. In addition to regular paint strokes, it is possible to apply basic primitive shapes and polyline style strokes.

Each stroke can have a duration that lasts for the entire project, a single frame or field, or an arbitrary number of fields. The strokes can have independent durations in the Timeline for easy manipulation of time. Alternatively, Multistrokes is a faster but non-editable way for doing many mask cleanup paint tasks.

Controls

As the Paint Mask node is fundamentally identical to the Paint node, see the Paint node and the Paint and Rotoscoping chapter for more details on the many options and capabilities available. The only difference is that, as Paint Mask operates on single-channel mask images, there is no Channel Selector control and all color controls have only a single Alpha value.



Mask Tab

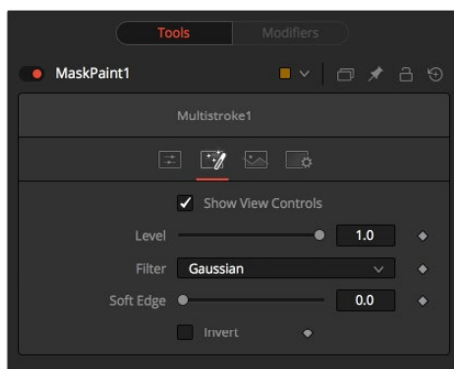
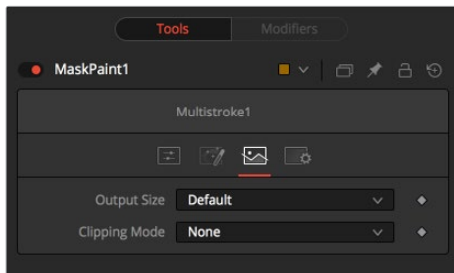



Image Tab

Please refer to the Common Mask Controls.



Polygon Mask [PLY]



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.

Controls



Level, Filter, Soft Edge and Border Width

Please refer to the Common Mask Controls.

Size

Use the Size control to adjust the scale of the Polygon 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 effect mask along any axis.

Fill Method

The Fill Method drop-down menu offers two different techniques for dealing with overlapping regions of a polyline. If overlapping polyline segments in a mask are causing undesirable holes in the mask, try switching the setting of this control from Alternate to Non Zero Winding.

Right Click Here for Shape Animation


By default, all polyline 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 masks together.

Adding Points

Adding Points to a polygonal effect mask is relatively simple. Immediately after creating the mask there are no points, but the mask will be in Click Append mode. Simply 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.




When a Polygon mask is added to a node, a toolbar will appear in the view with buttons that offer easy access to modes and nodes. Hold the mouse pointer over any button in the toolbar to display a tooltip that describes that button's function.


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.

Image Tab

Please refer to the Common Mask Controls.

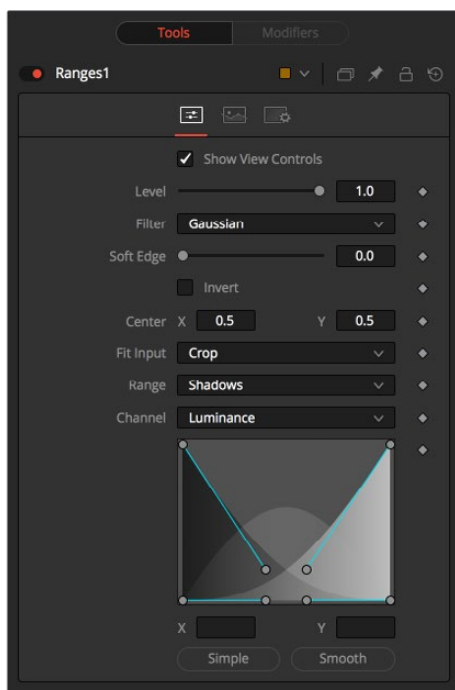


Ranges Mask [RNG]



Similar to Bitmap Mask, the Range 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, Range allows spline-based selection of low, mid and high ranges, akin to Color Corrector.

Controls



Level, Filter, Soft Edge and Border Width

Please refer to the Common Mask Controls.

Shadows/Midtones/Highlights

These buttons are used to select which range will be 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 will 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.

Spline Display

The extent of the ranges is selected by manipulating the spline handles. There are four spline points, each with one Bezier 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 Bezier 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. In other words, after low and high masks have been applied, midtones is everything else.


The X and Y text controls below the Spline Display can be used to enter precise positions for the selected Bezier point or handle.

Presets


This sets the splines to two commonly-used configurations. Simple gives a straightforward linear-weighted selection, while Smooth uses a more natural falloff.

Image Tab

Please refer to the Common Mask Controls.

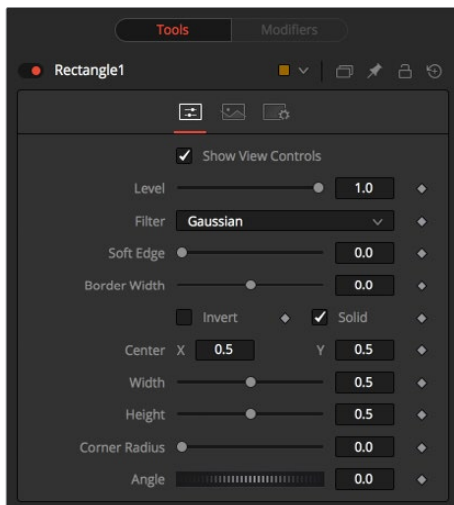


Rectangle Mask [REC]



The Rectangle Mask creates a simple square or rectangular effect mask. Many of the controls found in this Mask node are common to all Mask nodes. These controls are documented at Common Mask Controls.

Controls



Level, Filter, Soft Edge and Border Width

Please refer to the Common Mask Controls.

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 will apply 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. 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.

Image Tab

Please refer to the Common Mask Controls.



Triangle Mask [TRI]

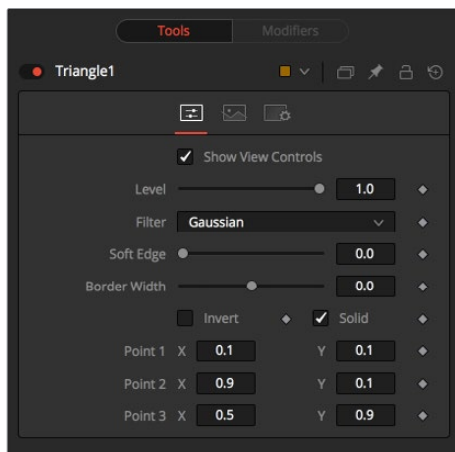


The Triangle Mask is unique in that it has no Center, Size or Angle control. Unlike most other types of masks, all three points of the triangle may be attached to a tracker or motion path.

Complex effect masking is possible using trackers and other nodes' paths to manipulate the triangle shape.

Many of the controls found in this Mask node are common to all Mask nodes. These controls are documented at Common Mask Controls.

Controls



Level, Filter, Soft Edge and Border Width


Please refer to the Common Mask Controls.

Point 1, Point 2, Point 3

These controls indicate the position of the three corners of the triangle. Each point can be published, connected to other controls, animated with a path, or attached to trackers. To perform any of these tasks, right-click on the Position control in the mask controls, or directly on the point in the Viewer, and select the appropriate option from the contextual menu.

Image Tab

Please refer to the Common Mask Controls.



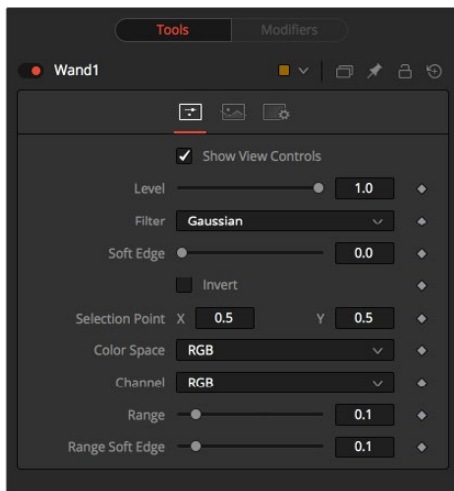
Wand Mask [WND]



The Wand Mask provides the ability to mask an image based on a magic wand-style selection, similar to the magic wand tools found in traditional 2D paint applications. As with a Bitmap mask, any image in the composition can be used as a source for the mask. Generally, the default is most useful, where the source image is the input of the node to which the mask is applied.

NOTE: When adding a Wand mask to a node, a crosshair will appear in the Viewers. This crosshair should be positioned in the image to select the color used to create the Wand mask. The mask itself is created by examining the color of the pixel beneath the selection point and adding that color to the mask. The mask then expands to examine the pixels surrounding the selection point. If the surrounding pixels are the same color, they are also added to the mask. The mask stops expanding when no connecting pixels fall within the color range of the mask. The node to be used as the image source for the mask should be connected to the Source (orange) input on the node tree. As with other masks, the Effect Mask (blue) input is used to combine the wand's result with other Mask nodes. Many of the controls found in this Mask node are common to all Mask nodes. These controls are documented at Common Mask Controls.

Controls



Level, Filter, Soft Edge and Border Width

Please refer to the Common Mask Controls.

Selection Point

The Selection Point is a pair of X and Y coordinates that determine from 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 from an individual channel only.

The exact labels of the buttons will depend on the color space selected for the Wand Mask operation. If the color space is RGB, the options will be R, G, or B. If YUV is the color space, the options will be Y, U, or V.

Range


The Range slider controls the range of colors around the source color that will be included in the mask. If the value is left at 0.0, only pixels of exactly the same color as the source will be considered part of the mask. The higher the value, the more that similar colors in the source will be 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 will be treated as 100% within the mask. If the soft range is set to 0.0, no other pixels will be considered for the mask. Increasing the soft range will increase the number of colors close to, but not quite within, the range that will be included in the mask. These pixels will be semi-transparent in the mask.

Image Tab

Please refer to the Common Mask Controls.



Chapter 44


Matte Nodes

This chapter details the Matte nodes available in Fusion.

Contents

Alpha Divide [ADV]	919
Alpha Multiply [AMI]	919
Chroma Keyer [CKY]	919
Clean Plate	923
Delta Keyer	924
Difference Keyer [DfK]	932
Luma Keyer [LKY]	934
Matte Control [MAT]	935
Ultra Keyer [UKY]	939


Alpha Divide [ADV]



As the name gives away, the Alpha Divide's sole purpose is to divide an incoming image by its Alpha Channel.

This node has no controls


Alpha Multiply [AMI]



As the name gives away, the AlphaMultiply's sole purpose is to multiply an incoming image with its Alpha Channel.

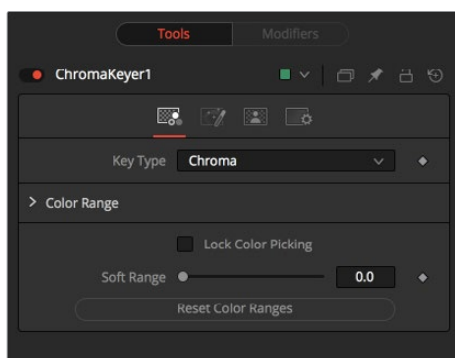
This node has no controls

Chroma Keyer [CKY]



The Chroma Keyer node creates an alpha channel (matte) for an image by removing selected colors from the scene. Unlike the Ultrakeyer, which has specific optimizations for keying from blue and green colors, the chroma keyer works equally well with any color.

Chroma Key Tab



Key Type

This determines the type of selection to be used for the matte creation.

Chroma

Chroma causes a matte to be created based on the RGB values of the selected color range.

Color

This causes a matte to be created based on the hue of the selected color range.

Color Range

These range controls update automatically to represent the current color selection. Colors are selected by selecting the Chroma Keyer node in the node tree, then dragging in 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 generally selecting colors in the displays is all that is required.

Lock Color Picking

When this checkbox is selected, Fusion will prevent accidental growing of the selected range by selecting more colors from the view. It is a good idea to select this checkbox once the color selection has been made for the matte. All other controls in the node remain editable.

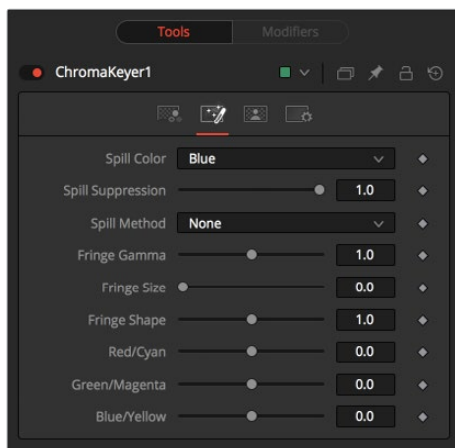
Soft Range

This control softens the selected color range to include additional colors into the matte.

Reset Color Ranges

Clicking on this button resets the Chroma Keyer's range controls, discarding all color selections. All other sliders and controls maintain their values.

Image Tab



Spill Color

Use these buttons to select the color used as the base for all spill suppression techniques.

Spill Suppression

Spill is generally caused by the transmission of the color of the background 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 fringe of the foreground element.

Spill suppression attempts to remove color from the fringe. The process used is optimized for either blue or green screens; you select which color is used as the base from the control above.

When this slider is set to 0 no spill suppression is applied to the image.

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, 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 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.

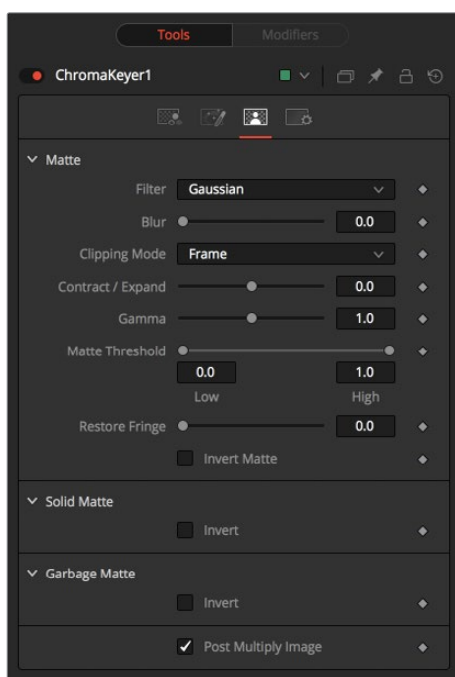
Fringe Shape

Fringe Shape forces the fringe to be pressed toward the external edge of the image or pulled toward the inner edge of the fringe. Its effect is most noticeable while the Fringe Size sliders 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 semi-transparent pixels that still contain color from the original background to match the new background.

Matte Tab



Matte Blur

Matte Blur 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 applied to the matte.

Matte Contract/Expand

This slider shrinks or grows the semi-transparent 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 only affects semi-transparent areas, it will have no effect on a hard edge's matte.

Matte Gamma

Matte Gamma raises or lowers the values of the matte in the semi-transparent 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 only affects semi-transparent areas, it will have no effect on a hard edge's matte.

Matte 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.

This control is often used to reject salt and pepper noise in the matte.

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.

Garbage Matte Mode

Garbage Mattes are Mask nodes or images connected to the Garbage Matte input on the node's tile. 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.

Make Transparent

Select this button to make the Garbage matte transparent.

Make Solid

Select this button to make the Garbage matte solid.

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 pre-multiplied for purposes of merging it with other images. Use the Subtractive option of the Merge node instead of the Additive option.

For more information, see the Merge node documentation.

Clean Plate



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 so that later keying can key fine detail without choking or clipping the matte.



Method

Color Uses a difference method to replace the color; choose by Click-Drag clicking to the image.

Ranges uses a chroma range method to separate the background color.

Matte 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. This control is often used to reject salt and pepper noise in the matte.

Erode

Will decrease the size of the screen area.

Crop

Will trim in from the edges of the image.

Grow Edges

Will expand the color of the edges of the subject.


Fill

Will fill in remaining holes with color from the surrounding screen color.

Time Mode

- **Sequence:** will generate a new clean plate every frame.
- **Hold Frame:** will hold the clean plate at a single frame.

Mask Tab




Garbage Mask

Garbage mask can be applied to clear areas before growing edges or filling remaining holes.

Invert

Invert will use the transparent parts of the mask to clear the image.

Delta Keyer



The Delta Keyer is a classic color difference keyer, with many features and controls for tuning the matte and separating the subject from blue or green screen.

It contains several keying systems; the Key tab is the master difference keyer, Pre Matte is a pre clean plate to smooth out screen color. Tuning, Fringe, and Matte finish the key process.

How to Key


Use the Pick button on the background color to select the blue or green screen color from the image. Hold Option (Alt) while click dragging the pick and it will pick the color from the up stream image, making the key not flicker.

The Pre Matte is a clean plate generator that smooths out the color of the screen and is instigated by box selecting areas of the screen color. Tweaking the Erode will expand the pre matte so it does not clip into the subject of the image.

Inputs

- Delta keyer has the image input.
- Garbage Matte
- Clean Plate
- Effect mask
- Solid Matte

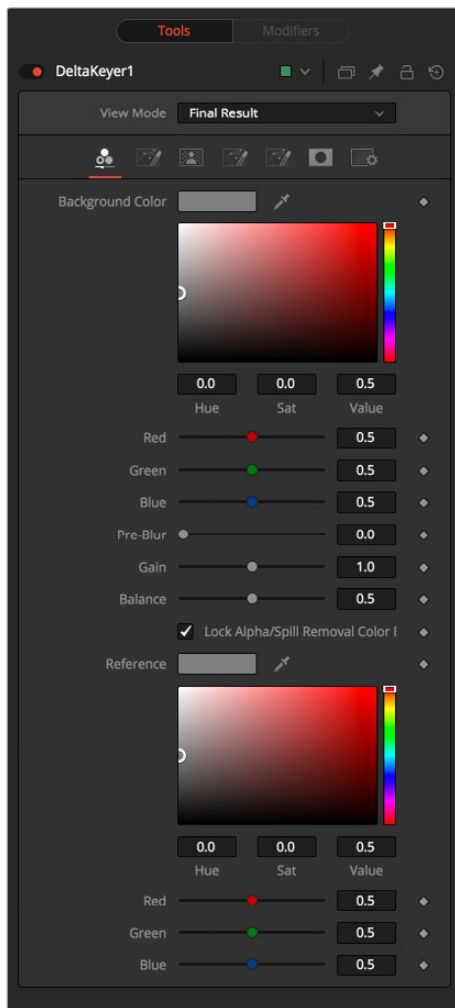
View Mode



At the top of the controls is View Mode, the default is to show the final result. This can be changed to see the output of the various intermediate stages of the Key process.

- **Pre Matte:** Will display the output of the Pre Matte key.
- **Matte:** Will display the alpha of the key before being combined with Solid and Garbage masks.
- **Tuning Ranges:** Will display the Shadow Midtone and Highlight range of the image. Shadows are in the red channel, Midtones in the green channel and Highlights 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:** Is the original source image color channels combined with the final matte. This can be combined with further DeltaKeyer tools.
- **Final Result:** Is the final keyed image with spill suppression, ready to merge onto a scene.

Key Tab



Background Color

This is the color of the blue or green screen, the keying color. This will be turned black with no alpha. Use the Pick button on the background color to select the blue or green screen color from the image. Hold Option (Alt) while click dragging the pick and it will pick the color from the upstream image, making the key not flicker.

Pre-Blur

Applies a blur before generating the alpha. This can help with certain types of noise and edge enhancements and artifacts in the source image.

Gain

Increases the influence of the background color. This will cause areas of background color to become more transparent.

Balance

The key is performed by comparing the differences between the dominant channel determined by the background color and the other two channels, with balance determining the proportions of the other two channels. A value of 0 will use the minimum of the other two channels, where a value of 1 will use the maximum. A value of 0.5 will use half of each.

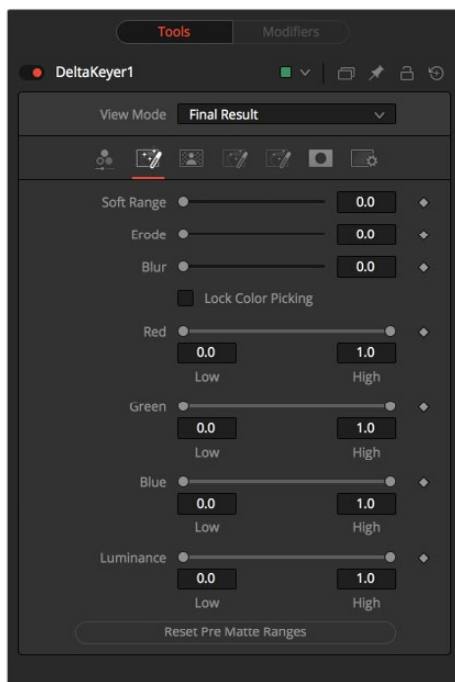
Lock Alpha/Spill Removal Color Balance Reference

Unlocking this allows a different color references to be used when generating the alpha and when determining how much of the background color to subtract from the image.

Color Balance Reference

Can be used to adjust for lighting or white balance that might be reducing background color purity and saturation. A correction can be applied based on the reference of a neutral colored object when generating the key alpha and determining the amount of background color subtraction, without altering the background color that is subtracted.

Pre Matte



Soft Range

The soft range will extend the range of selected color and roll off the screen color.

Erode

Erode will contract the edge of the pre matte, so the edge detail will not be clipped.

Blur

This will soften the edges of the pre matte.

PreMatte Range

These range controls update automatically to represent the current color selection. Generally, the Reveal control does not have to be opened to display these controls. Colors are selected by selecting the Ultra Keyer node's tile in the flow and dragging in 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 generally selecting colors in the displays is all that is required.

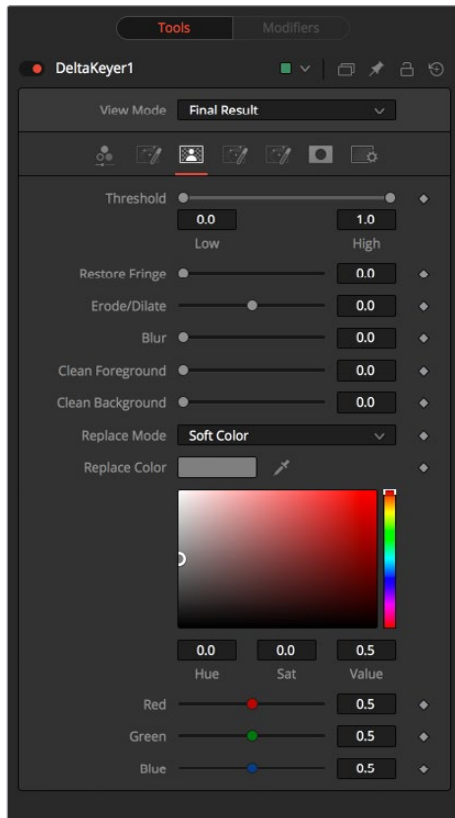
Lock Color Picking

When this checkbox is selected, Fusion will prevent accidental growing of the selected range by selecting more colors from the view. It is a good idea to select this checkbox once the color selection is made for the matte. All other controls in the node remain editable.

Reset Pre Matte Ranges

This discards all color selection by resetting the ranges but maintains all other slider and control values.

Matte Tab



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 to get a key, the edge of the subject where you have hair will get clipped out, Restore Fringe will bring back that edge while keeping the matte solid.

Erode/Dilate

Expands or contracts the matte.

Blur

Softens the matte.

Clean Foreground

Fills slightly transparent areas of the matte.

Clean Background

Clips the bottom dark range of the matte.

Replace Mode

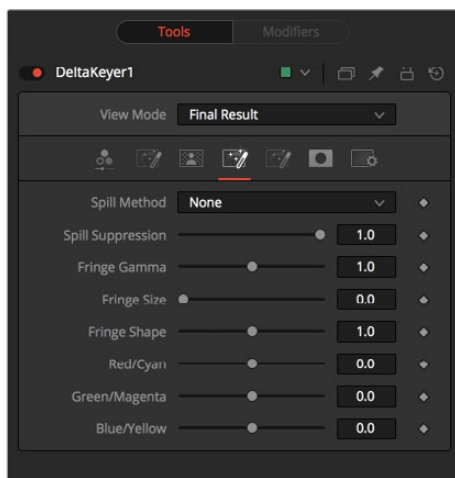
Determines how matte adjustments will cause color to be restored to the image.

- **None:** Is no color replacement. Matte processing will not affect the color.
- **Source:** Is the original color from the unkeyed image.
- **Hard Color:** Is solid color.
- **Soft Color:** Is 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.

Fringe Tab



Spill Suppression

Spill is generally caused by the transmission of the color of the background through the semi transparent 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 fringe of the foreground element.

Spill suppression attempts to remove color from the fringe. The process used is optimized for either blue or green screens; you select which color is used as the base from the control above.

When this slider is set to 0, no spill suppression is applied to the image.

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, 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.

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

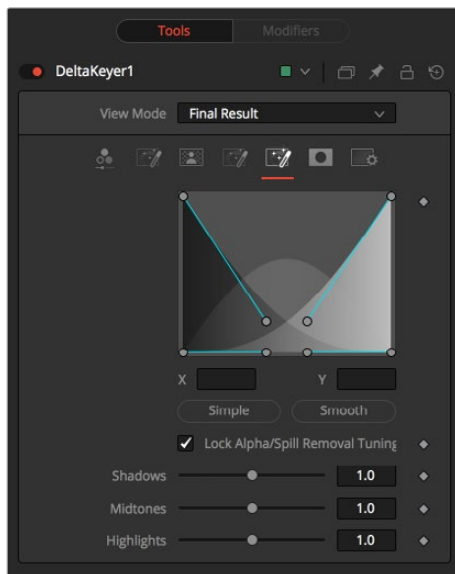
Fringe Shape forces the fringe to be pressed toward the external edge of the image or pulled toward the inner edge of the fringe. Its effect is most noticeable while the Fringe Size value is large.

Cyan/Red, Magenta/Green and Yellow/Blue

Use these three controls to color correct the fringe of the image.

This is useful for correcting semi-transparent pixels that still contain color from the original background to match the new background.

Tuning Tab



Range controls

This defines how much color range is in the Dark Shadows, Midtones, and Highlight bright areas of the image. This spline controls allow for easy adjusting of the tonal ranges of each Shadow and Highlight tonal map.

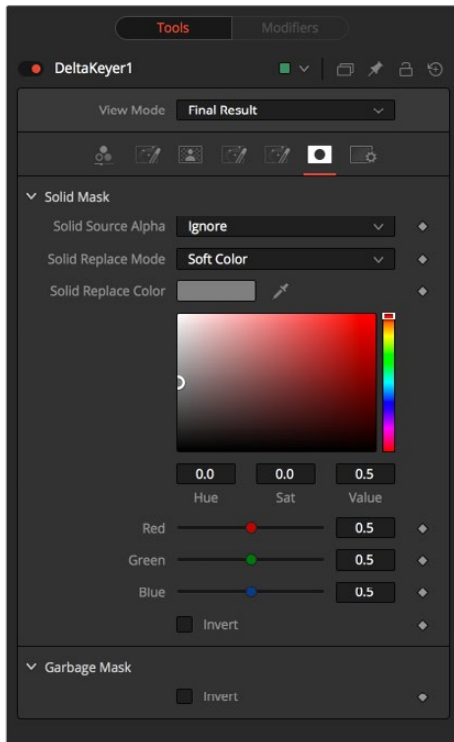
Preset Simple will set the range to be linear. Preset Smooth sets a smooth tonal gradient for the ranges.

Lock Alpha/Spill Removal Tuning

Unlocking this allows a different tuning to be 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.

Mask Tab



Solid Mask

Solid Source Alpha

Used to combine the existing alpha from the source image into the solid mask.

- **Ignore:** Does not combine the alpha from the source image.
- **Add:** Solid areas of the source image alpha will be made solid in the solid mask.
- **Subtract:** Transparent areas of the source image alpha will be made transparent in the solid mask.

Solid Replace Mode

This determines how the solid mask will cause color to be restored to the image.

- **None:** Is no color replacement. The solid mask will not affect the color.
- **Source:** Is the original color from the unkeyed image.
- **Hard Color:** Is solid color.
- **Soft Color:** Is 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:** will invert the solid mask, before it is combined with the source alpha.

Garbage Mask

- **Invert:** Normally solid areas of the Garbage mask will clear the image. When inverted it's the transparent areas of the mask that will clear the image.

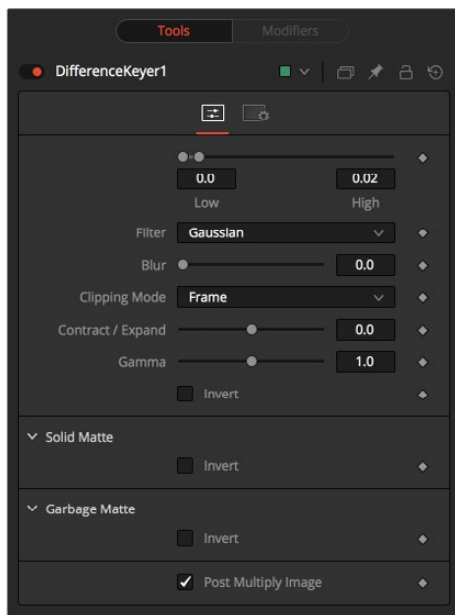
Difference Keyer [DfK]



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.

The Difference Keyer node is very sensitive and, although the process sounds reasonable at first glance, subtle variations in the position of the camera from shot to shot usually make it difficult to pull a highly-detailed alpha channel using this method. 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 will always be 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.

Controls



Threshold High and Low

This slider works by defining a range of difference values between the images to create a matte. A difference below the lower threshold becomes black in the matte. Any difference above the upper threshold becomes white (solid) in the matte. The difference values in the range in between create a gray scale matte.

Matte 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.

Matte Contrast

The Matte Contrast slider changes the Look up Table curve of the matte's luminance values. This creates a soft cropping of the matte at the low end of the slider and a hard edge expansion of the matte at higher slider values.

Matte Gamma

Matte Gamma raises or lowers the values of the matte in the semi-transparent areas. Higher values cause the gray areas to be more opaque and lower values cause the gray areas to be more transparent. Completely 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.

Garbage Matte Mode

Garbage Mattes are Mask nodes or images connected to the Garbage Matte input on the node's tile. 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 node. 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.

Make Transparent

Select this button to make the Garbage matte transparent.

Make Solid

Select this button to make the Garbage matte solid.

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 pre-multiplied for purposes of merging it with other images. Use the Subtractive option of the Merge tool instead of the Additive option.

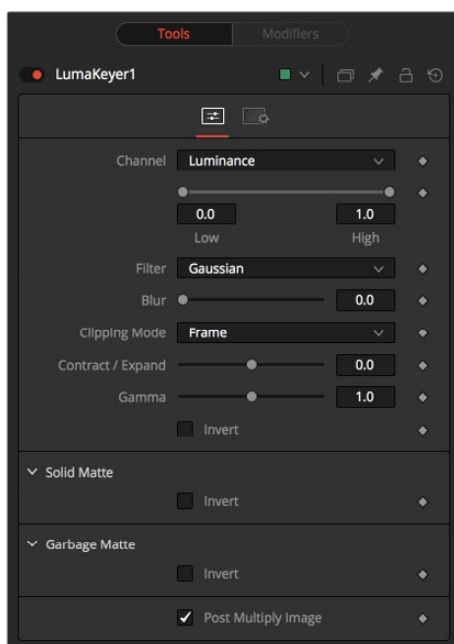
For more information, see the Merge Nodes documentation.

Luma Keyer [LKY]



The Luma Keyer node uses the overall luminance of an image to create an alpha channel. When this node was first created it was used exclusively on the luminance channel of the image, but it has since grown to allow pulling mattes from virtually any channel Fusion understands. In some respects, it would now be more accurate to call this node a Channel Keyer.

Controls



Channel

Use this drop-down list to select the color channel used for creating the matte. Select from the Red, Green, Blue, Alpha, Hue, Luminance, Saturation, and Depth (Z-buffer) channels.

Threshold High and Low

This slider works by defining a range of luminance values in the image to create a matte. Any value below the lower threshold becomes black in the matte. Any value above the upper threshold becomes white (solid) in the matte. All values within the range create the gray scale matte.

Matte Blur

Matte Blur 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 applied to the matte.

Matte Contrast

The Contrast slider changes the Look up Table curve of the matte's luminance values. This creates a soft cropping of the matte at the low end of the slider and a hard edge expansion of the matte at higher slider values.

Matte Gamma

Matte Gamma raises or lowers the values of the matte in the semi-transparent areas. Higher values cause the gray areas to be more opaque and lower values cause the gray areas to be more transparent. Completely black or white regions of the matte remain unaffected.

- **Invert:** When toggled on, the matte is inverted, causing all transparent areas to be opaque and all opaque areas to be transparent.

Garbage Matte Mode

Garbage Mattes are Mask nodes or images connected to the Garbage Matte input on the node's tile. 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 node. A Matte Control tool is often used after a Keyer node to add a Garbage matte with the opposite effect of the matte applied to the keyer.

Make Transparent

Select this button to make the Garbage matte transparent.

Make Solid

Select this button to make the Garbage matte solid.

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 pre-multiplied for purposes of merging it with other images. Use the Subtractive option of the Merge node instead of the Additive option.

For more information, see the Merge Nodes documentation.


Matte Control [MAT]



Keyer nodes are generally used to create an alpha channel on an image that does not already have one. The Matte Control node is used to manipulate an existing alpha channel or to create one by hand via rotoscoping.

The Matte Control node also has a Foreground Image input. Use this node to copy a color channel or alpha channel from the foreground to the background, or to combine alpha channels from the two images.

Controls



Matte Combine

The Matte Control node can combine alpha or color channels from an image in the foreground input with the background image. 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 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

Selection of 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.

Matte 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 applied to the matte.

Matte Contract/Expand

This shrinks or grows the matte to exclude some of the keyed image or include some of its surrounding area. Values above 0.0 expand the matte and values below 0.0 contract it.

Matte Gamma

This raises or lowers the values of the matte in the semi-transparent 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.

Matte 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.

- **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.

Garbage Matte

Garbage Mattes are Mask nodes or images connected to the Garbage Matte input on the node's tile. 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 node. 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.

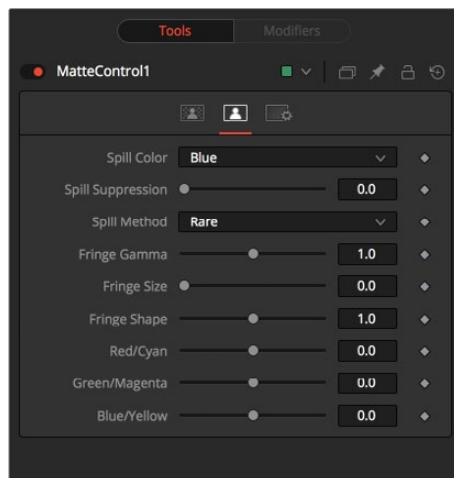
- **Make Transparent:** Select this button to make the Garbage matte transparent.
- **Make Solid:** Select this button to make the Garbage matte solid.
- **Post Multiply Image:** Select this option to cause the keyer to multiply the color channels of the image against the alpha channel it creates for the image. This option is usually enabled and is on by default.

Deselect this checkbox and the image can no longer be considered pre-multiplied for purposes of merging it with other images. Use the Subtractive option of the Merge node instead of the Additive option.

- **Multiply Garbage Matte:** When selected, the values of the image's Red, Green, and Blue channels will be multiplied against the Garbage Matte input.

For more information, see the Merge nodes documentation.

Spill Tab



Spill Color

Use these buttons to select the color used as the base for all spill suppression techniques.

Spill Suppression

Spill is generally caused by the transmission of the color of the background through the semi-transparent 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 fringe of the foreground element.

Spill suppression attempts to remove color from the fringe. The process used is optimized for either blue or green screens; you select which color is used as the base from the control above.

When this slider is set to 0, no spill suppression is applied to the image.

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, 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.

Fringe Gamma

This control can be used to adjust the brightness of the fringe or halo that surrounds the keyed image.

Cyan/Red, Magenta/Green and Yellow/Blue

Use these three controls to color correct the fringe of the image. This is useful for correcting semi-transparent pixels that still contain color from the original background to match the new background.


Fringe Size

This expands and contracts the size of the fringe or halo surrounding the keyed image.

Fringe Shape

Fringe Shape forces the fringe to be pressed toward the external edge of the image or pulled toward the inner edge of the fringe. Its effect is most noticeable while the Fringe Size sliders value is large.

Ultra Keyer [UKY]



The Ultra Keyer node has 2 keyers built in to it, a Prematte keyer acts as garbage matte creator and the color difference keyer that will extract fine detail and transparency. This is optimized to extract mattes from images using blue screen or green screen backgrounds.

How to Key

Use the Pick Button on the Background Color to select the blue or green screen color from the image. Hold Option (Alt) while click dragging the pick and it will pick the color from the upstream image, making the key not flicker.

The Pre Matte garbage keyer is instigated by box selecting areas of the screen color, and tweaking the Pre Matte Size will expand the Garbage matte so it does not clip into the subject of the image.

Pre Matte Tab



Background Color

This 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 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 will iteratively merge the pre-keyed image over either a blue or green background before processing it further.

In certain cases 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.

- **PreMatte Range:** These range controls update automatically to represent the current color selection. Generally, the Reveal control does not have to be opened to display these controls. Colors are selected by selecting the Ultra Keyer node's tile in the node tree and dragging in 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 generally selecting colors in the displays is all that is required.
- **Lock Color Picking:** When this checkbox is selected, Fusion will prevent accidental growing of the selected range by selecting more colors from the view. It is a good idea to select this checkbox once the color selection is made for the matte. All other controls in the node remain editable.

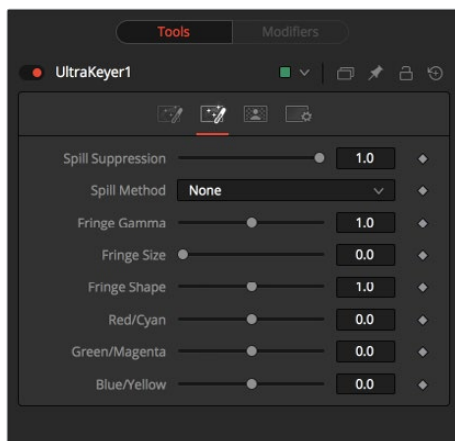
Pre Matte Size

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 semi-transparent areas of the subject. This generally will also 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.

Image Tab



Spill Suppression

Spill is generally caused by the transmission of the color of the background 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 fringe of the foreground element.

Spill suppression attempts to remove color from the fringe. The process used is optimized for either blue or green screens; you select which color is used as the base from the control above.

When this slider is set to 0, no spill suppression is applied to the image.

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, 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.

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.

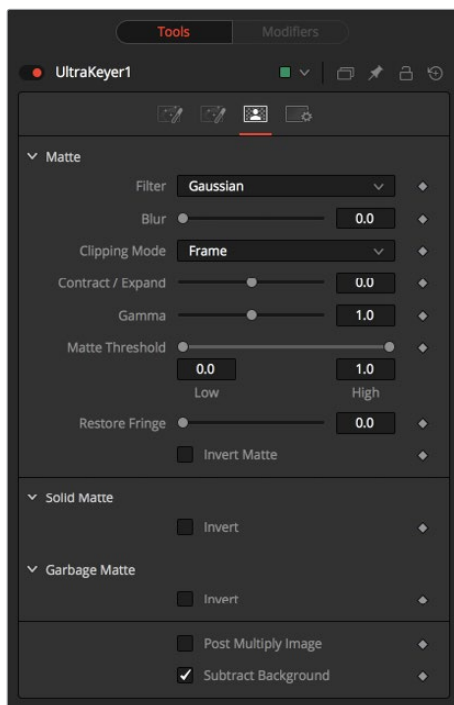
Fringe Shape

Fringe Shape forces the fringe to be pressed toward the external edge of the image or pulled toward the inner edge of the fringe. Its effect is most noticeable while the Fringe Size sliders 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 semi-transparent pixels that still contain color from the original background to match the new background.

Matte Tab



Matte Blur

Matte Blur 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 applied to the matte.

Matte Contract/Expand

This slider shrinks or grows the semi-transparent 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 only affects semi-transparent areas, it will have no effect on a hard edge's matte.

Matte Gamma

Matte Gamma raises or lowers the values of the matte in the semi-transparent 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 only affects semi-transparent areas, it will have no effect on a hard edge's matte.

Matte 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.

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 to get a key, the edge of the subject where you have hair will get clipped out, Restore Fringe will bring 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.

Make Solid

Select this button to make the Garbage matte solid.

Make Transparent

Select this button to make the Garbage matte transparent.

Garbage Matte

Garbage Mattes are Mask nodes or images connected to the Garbage Matte input on the node's tile. 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.

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 pre-multiplied for purposes of merging it with other images. Use the Subtractive option of the Merge node instead of the Additive option.

For more information, see the Merge nodes documentation.

Chapter 45


Metadata Nodes

This chapter details the Metadata nodes available in Fusion.

Contents


Copy Metadata [META]	946
Set Metadata [SMETA]	947
Set Timecode [TCMETA]	948

Copy Metadata [META]



Copy Metadata combines, replaces or clears the metadata in your image.

Controls



Operation

The drop-down defines how the metadata of foreground and background input are treated.

Imagine having metadata in the background image looking like this:

```
FrameRate = 24  
Record = Scratched  
TimeCode = 00:00:08:15  
Hovercraft_Filling = Eels
```

and the foreground's metadata looking like this:

```
Hovercraft_Filling = Oysters
```

Merge (Replace Duplicates)

All values will be merged, but values with duplicate names will be taken from the foreground input.

The output looks like this:

```
FrameRate = 24  
Record = Scratched  
TimeCode = 00:00:08:15  
Hovercraft_Filling = Oysters
```


Merge (Preserve Duplicates)

All values will be merged, but values with duplicate names will be taken from the background input.

The output looks like this:

```
FrameRate = 24  
Record = Scratched  
TimeCode = 00:00:08:15  
Hovercraft_Filling = Eels
```

Replace


```
Hovercraft_Filling = Oysters
```

The entire metadata in the background will be replaced by the ones in the foreground. The output looks like this:

Clear


All metadata will be discarded.

Set Metadata [SMETA]



Set Metadata allows the user to create new Name = Value pairs in the metadata.

Controls




Field Name

The name of the Metadata Value. Do not use spaces.

Field Value


The value assigned to the name above.

Set Timecode [TCMETA]



Set Timecode inserts dynamic timecode values into the metadata table based on the FPS settings.

Controls



FPS

You can choose from a variety of Frame Per Second settings here. 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 can look like this:

```
TimeCode:    00:00:08:15  
Frames:      207
```

Chapter 46

Miscellaneous Nodes

This chapter details miscellaneous nodes within Fusion.

Contents

Auto Domain [ADOD]	951
Change Depth [CD]	952
Custom [CT]	953
Fields [FLD]	962
Run Command [RUN]	964
Set Domain [DOD]	967
Time Speed [TSPD]	969
Time Stretcher [TST]	970

Auto Domain [ADOD]




The Auto Domain node automatically sets the image's domain of definition based on bounds of the input image's background Canvas color. It does not change the image's physical dimensions. This node can be used to speed up compositions by optimizing the DoD of images based on their content rather than their dimensions.

For example, a CG character rarely takes up the entire frame of an image. The Auto Domain node would set the DoD to a rectangular region encompassing the portion of the scene actually containing the character. The DoD is updated on each frame to accommodate changes, such as a character walking closer to the camera.

See the Set Canvas Color node for more information about the Canvas color.

Controls



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).

Change Depth [CD]




The Change Depth node has one simple use, to change the bits per color channel used to process a node. The single control for this node is Depth, which contains five buttons. Select Keep to leave the color depth as is, or either 8-bit, 16-bit or Float to change to the selected color depth.

This node is often used after color correcting Cineon files, converting 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.

Controls



Depth

Keep doesn't do anything to the image but rather keeps the input depth. The other options change the bit depth of the image to the respective value.

Dither

When down converting from higher bit depth it might be useful to add Error Diffusion or Additive Noise to camouflage artifacts that result from problematic (high contrast) areas.

Custom [CT]

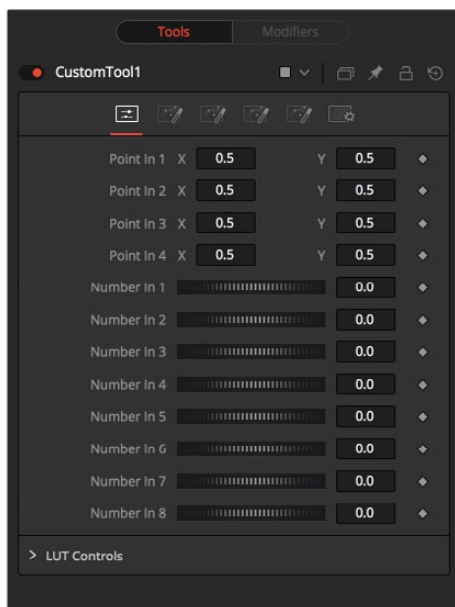


The Custom node is quite likely the most complex, and the most powerful, node in Fusion. Any user moderately experienced with scripting, or C++ programming, should find the structure and terminology used by the Custom node to be familiar.

The Custom node is used to create custom expressions and filters to modify an image. In addition to providing three image inputs, the Custom 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.

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.

Controls



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

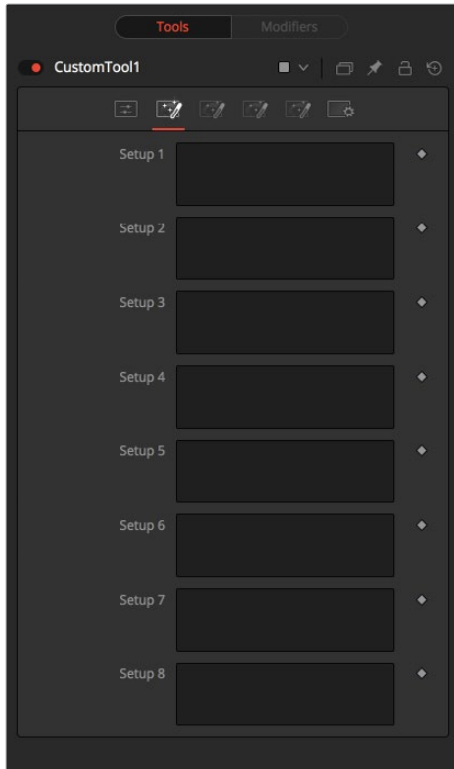
The values of these controls are available to expressions entered in the Setup, Intermediate, and Channels tabs as variables `n1`, `n2`, `n3`, ..., `n8`. They are normal slider controls and can be animated or connected to modifiers exactly as any other node might.

LUT in 1-4

The Custom 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 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 will still see the values as `n1`, `n2`, ..., `n8`.

Custom Setup Tab

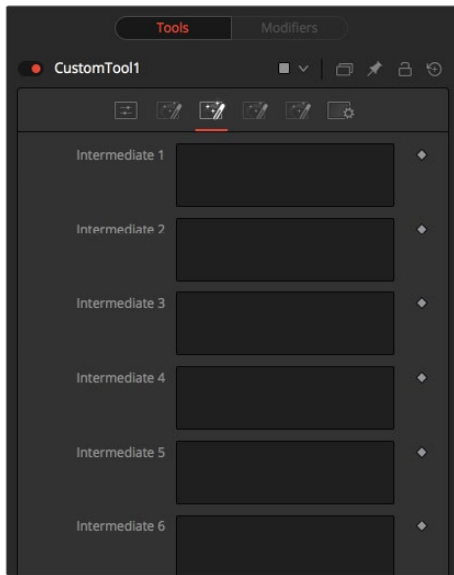


Setup 1-4

Up to four separate expressions can be calculated in the Setup tab of the Custom noder. 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 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`, and so on. Allowable values include constants, variables like `n1`.. `n8`, `time`, `W` and `H`, and so on, and functions like `sin()` or `getr1d()`.

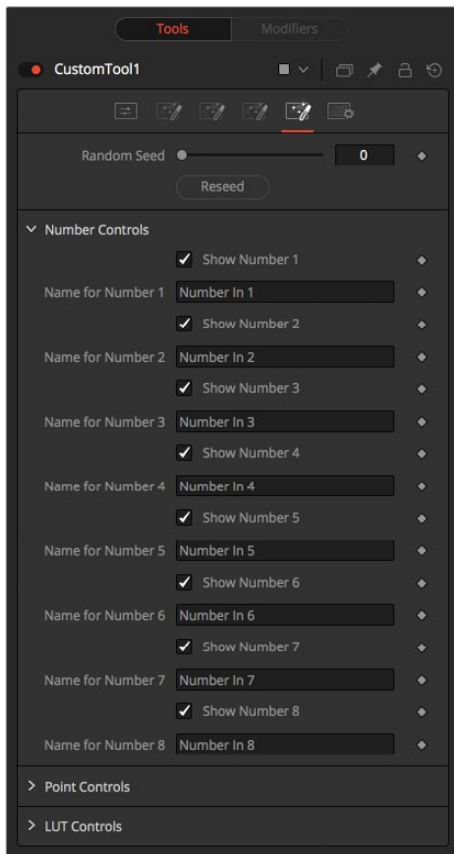
Custom Inter Tab



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.

Custom 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 nodes are required with different random results for each.

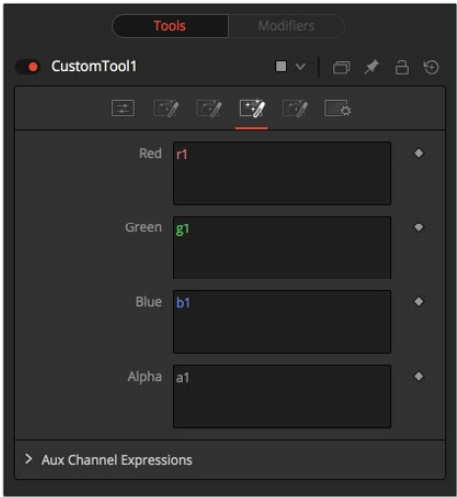
Number Controls

There are eight sets of Number controls, corresponding to the eight Number In sliders in the Controls tab. Untick 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. Untick 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.

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, and the result is used to create 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 will be 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 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

Value Variables	
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)
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 copy/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 other locations than the current one in an image.

In the above description, [ch] is a letter representing the channel to access, and [#] 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
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

Mathematical Expressions	
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
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
xly	1.0 if either x or y (or both) are not 0.0, otherwise 0.0
xlly	1.0 if either x or y (or both) are not 0.0, otherwise 0.0

Example

The following examples are intended to help you understand the various components of the Custom node.

Rotation

To rotate an image, we need the standard equations for 2D rotation:

$$x' = x * \cos(\theta) - y * \sin(\theta) \quad y' = x * \sin(\theta) + y * \cos(\theta)$$

Using the n1 slider for the angle theta, and a sample function, we get (for the red channel):

```
getrlb(x * cos(n1) - y * sin(n1), x * sin(n1) + y * cos(n1))
```

This will calculate the current pixel's (x,y) position rotated around the origin at (0,0) (the bottom-left corner), and then fetch 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 afterwards:

```
getrlb((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 getrlb() 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:

```
getrlb(i1, i2)  
getglb(i1, i2)  
getblb(i1, i2)  
getalb(i1, i2)
```

With the Intermediate expressions substituted in, we only have to do all the adds, subtracts and multiplies once per pixel, instead of four times per pixel. 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 image aspect at all. It is left as an exercise to the reader to include this (sorry). Another improvement could be to allow rotation around different points than the center.

Filtering

Our second example duplicates the functionality of a 3x3 Custom Filter node set to averages the current pixel together with the eight pixels surrounding it. To duplicate it with a Custom node, add a Custom 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 will be 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).

```
(getrlw(x-s1, y-s2) + getrlw(x, y-s2) + getrlw(x+s1, y-s2) + getrlw(x+s1, y) + getrlw(x-s1, y) + r1 +getrlw(x-s1, y+s2) + getrlw(x, y+s2) + getrlw(x+s1, y+s2)) / 9
```

This expression adds the nine pixels above the current pixel together by calling the `getrlw()` 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 exact 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.

```
(getglw(x-s1, y-s2) + getglw(x, y-s2) + getglw(x+s1, y-s2) + getglw(x+s1, y) + getglw(x-s1, y) + g1 +getglw(x-s1, y+s2) + getglw(x, y+s2) + getglw(x+s1, y+s2)) / 9
```

```
(getblw(x-s1, y-s2) + getblw(x, y-s2) + getblw(x+s1, y-s2) + getblw(x+s1, y) + getblw(x-s1, y) + b1 +getblw(x-s1, y+s2) + getblw(x, y+s2) + getblw(x+s1, y+s2)) / 9
```

```
(getalw(x-s1, y-s2) + getalw(x, y-s2) + getalw(x+s1, y-s2) + getalw(x+s1, y) + getalw(x-s1, y) + a1 + getalw(x-s1, y+s2) + getalw(x, y+s2) + getalw(x+s1, y+s2)) / 9
```

It is time to view the results. Add a Background node set to 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 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 node we created, but the flexibility of the Custom 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.

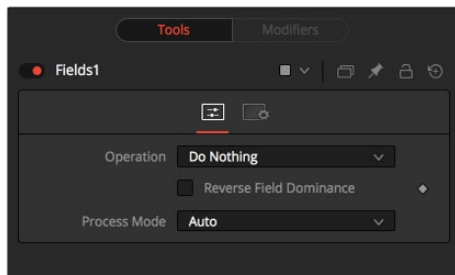
This is just one example; the possibilities of the Custom node are limitless.

Fields [FLD]



The Fields node is a robust multi-purpose utility offering several functions related to interlaced video frames. It interpolates separate video fields into video frames, or separates video frames into individual fields. It 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.

Controls



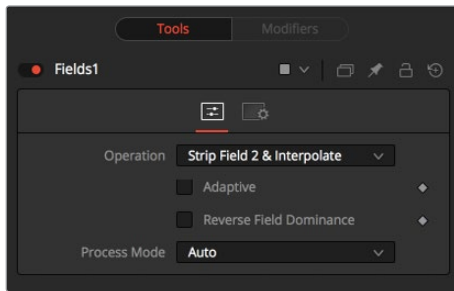
Operation

This control is used to select the type of operation the node will perform. See below for a detailed explanation.

Process Mode

This control is used to select the field's format used for the output image. See below for a detailed explanation.

Operation



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 will be combined to form half of the amount of double height frames. If supplied with two image streams, single frames from each stream will be combined to form double height images.

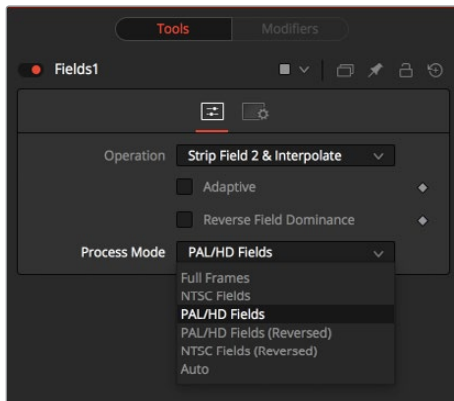
De-Interlace

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



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 is used if the input types are mixed.

Run Command [RUN]

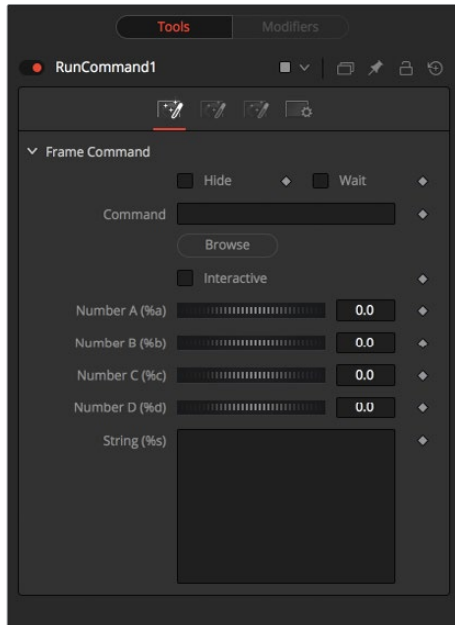


The RunCommand node is used to execute an external command or batch file at certain points during a render. Choose to run a command once at the start, or at the end of a render, or have the command execute once for each frame.

An image input is not required for this node to operate. However, if RunCommand is connected to a node's output, the command will only be launched after that 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-0 result, the node will also fail.

RunCommand can be used to net render other command line applications using the Fusion render manager, as well as a host of other useful functions.

Controls



Frame Command

The first file browser in the node is used to specify the path and parameters for the command to be run after each frame is rendered. Select the Hide checkbox to prevent the application or script from displaying a window when it is executed.

Hide

Select this checkbox to suppress the display of any window or dialog started by the command.

Wait

Enable this checkbox to cause the node to Wait for the remote application or tool to exit before continuing. If this checkbox is cleared, the system will continue rendering without waiting for the external application.

Number A (%B) 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**: Substituted with the text from the large text entry field

If you want to add zero padding 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.

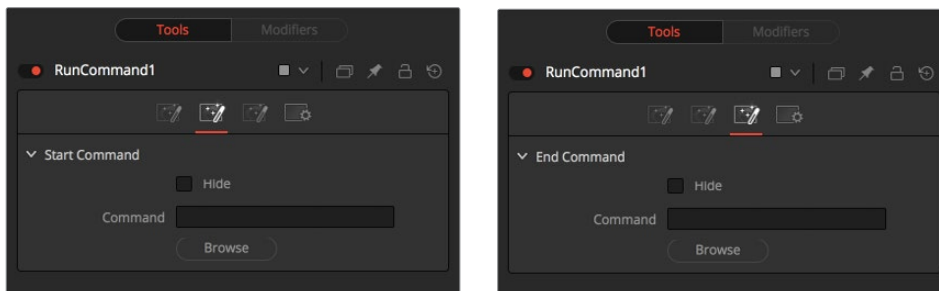
Process Priority

The Process Priority buttons provide options for selecting the priority at which the launched process runs. This determines how much processor time the launched process receives compared to other applications.

Interactive

This checkbox determines whether the launched application should run interactively, allowing user input.

RunCommand 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.tga, test0001.tga, etc. You may have to modify the example to match.

Add a RunCommand node after the Saver, to ensure the Saver has finished saving first. Now enter the following text into the RunCommand node's Frame Command text box:

```
C\copytest.bat D\test%04f.tga 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 RunCommand node could be used to FTP the files to a remote drive or Abekas device on the network, to print out each frame as it is rendered, or to execute a custom image-processing tool.

The RunCommand 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.

Set Domain [DOD]

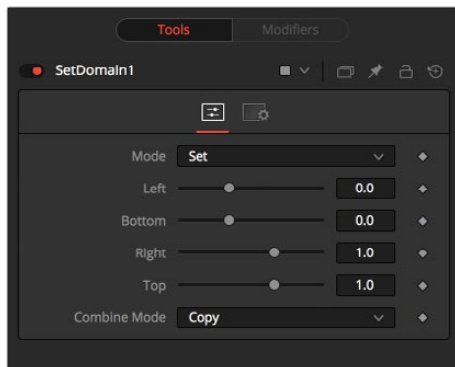


Set Domain is used to adjust or set the active area of an image, or in other words the area of the image considered to have valid data.

It does not change the image's physical dimensions. Anything outside the DoD will not be processed by downstream nodes, thus speeding up rendering of computational 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.

Controls/Set Mode



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. 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.

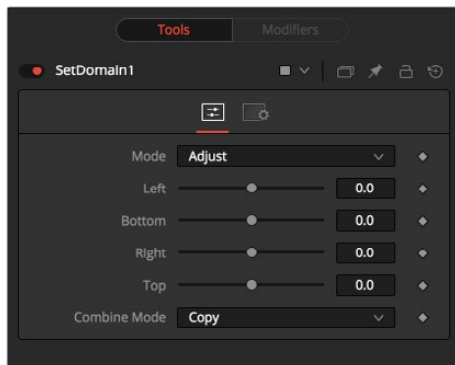
1 represents the top border of the image, 0 represents the bottom border. The slider defaults to 1 (Top Border).

External Inputs

- **SetDomain.Input:** [orange, required] This input must be connected to the output of a node that produces a 2D image.
- **SetDomain.Foreground:** [green, optional] This input expects 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.


Set Mode defaults to the full extent of the visible image.

Controls/Adjust Mode



In Adjust mode, basically the same operations can be carried out like in Set Mode. All Sliders default to 0, though, marking their respective full extent of the image. Positive values shrink the DoD while negative values expand the DoD to include more data.

Time Speed [TSPD]



Time Speed 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 manually upstream using an OpticalFlow node or by loading the forward/reverse vector channels from disk.

TimeSpeed does not interpolate the aux channels, but rather destroys them. In particular, the Vector/BackVector channels are consumed and destroyed after computation.

Add an OpticalFlow after the FlowSpeed if you want to generate flow vectors for the retimed footage.

Controls



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% is 1.0, 50% is 0.5 and 10% is 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 will offset time back and positive numbers will advance.

Interpolate Between Frames

When checked, frames before and after the current frame will be interpolated to create new frames. This usually offers smoother and cleaner results. When cleared, no interpolation will take place.

Sample Spread

This slider controls the strength of the interpolated frames on the current frame. A value of 0.5 causes 50% of the frame before and 50% of the frame ahead of the current frame to be blended with 0% of the current frame.

Depth Ordering

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 will produce small or slow vectors. The car will produce larger or faster vectors.

The Depth Ordering in this case is Fastest Ontop, since the car will draw over the background.

In a shot where the camera pans to follow the car, the background will have faster vectors, and the car will have slower vectors, so the Depth ordering method would be Slowest Ontop.

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 only use clamp edges 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 only have one checkbox enabled, you can use a larger softness at around 0.03.

Source Frame and Warp Direction

This control set allows for the choosing of 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:** Will take the previous frame and use the Forward vector to interpolate the new frame.
- **Next Forward:** Will take the next frame in the sequence and use the Forward vector to interpolate the new frame.
- **Prev Backward:** Will take the previous frame and use the Back Forward vector to interpolate the new frame.
- **Next Backward:** Will take the next frame in the sequence and use the Back vector to interpolate the new frame.

Time Stretcher [TST]




The Time Stretcher node is similar to the Time Speed node, but it permits the speed of the clip to be animated over the course of the effect. 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 non-linearly. To apply steady time changes such as frame rate changes, use a TimeSpeed instead. When operating in Flow mode, Optical Flow data is required.

This node does not generate optical flow directly; you have to create it manually upstream using an OpticalFlow node or by loading the forward/reverse vector channels from disk.

FlowStretcher does not interpolate the aux channels but rather destroys them. In particular, the Vector/BackVector channels are consumed/destroyed. Add an OpticalFlow after the FlowStretcher if you want to generate flow vectors for the retimed footage.

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, the Source Time control already contains a Bezier spline with a single keyframe set to 0.0. The position of the keyframe is determined by the current time when the node is added to the node tree.

(The Source Time spline may not be immediately visible until Edit is selected from the Source Time's contextual menu, or Display all Splines from the Spline Window's contextual menu.)

Interpolate Between Frames

This toggles whether the Time Stretcher will interpolate between the next and previous frames.

Sample Spread

This determines the strength of the interpolated frames on the current frame. A value of 0.5 causes 50% of the frame before and 50% of the frame ahead of the current frame to be blended with 0% of the current frame. A value of 0.25 would blend 25% of the previous and next frames with 50% of the current frame. Set this control over 0.25 only in extreme cases.

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 only use clamp edges 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 only have one checkbox enabled, you can use a larger softness at around 0.03.

Source Frame and Warp Direction

This control set allows for the choosing of 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:** Will take the previous frame and use the Forward vector to interpolate the new frame.
- **Next Forward:** Will take the next frame in sequence and use the Forward vector to interpolate the new frame.
- **Prev Backward:** Will take the previous frame and use the Back Forward vector to interpolate the new frame.
- **Next Backward:** Will take the next frame in sequence and use the Back vector to interpolate the new frame.

Depth Ordering

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 will produce small or slow vectors. The car will produce larger or faster vectors. The Depth ordering in this case is Fastest Ontop, since the car will draw over the background.

In a shot where the camera pans to follow the car, the background will have faster vectors, and the car will have slower vectors, so the Depth ordering method would be Slowest Ontop.

Example

Make certain 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 is to be displayed in during the project.

To shrink a 100-frame sequence to 25 frames, follow these steps:

- Change the Current Time to frame 0.
- Change the Source Time control to 0.0.
- Advance to frame 24.
- Change the Source Time to 99.
- Check that the spline result is linear.
- Fusion will render 25 frames by interpolating down the 100 frames to a length of 25.
- Hold the last frame for 30 frames, then play the clip backward at regular speed. Continue the example from above and follow the steps below.
- Advance to frame 129.
- Right-click on the Source Time control and select Set Key from the menu.
- Advance to frame 229 (129 + 100).
- Set the Source time to 0.0.

Chapter 47

Optical Flow

This chapter details the Optical Flow nodes available in Fusion.

Contents

Optical Flow [OF]	975
Repair Frame [REP]	977
Smooth Motion [SM]	979
Tween [TW]	980

Optical Flow [OF]



This node computes the Optical Flow between the frames of the input sequence. The optical flow can be thought of as a per pixel displacement vector which matches up features between two frames.

The computed optical flow is stored within the Vector and BackVector aux channels of the output image. At its highest quality settings, the Optical Flow node can be slow to process. If you find optical flow is too slow, you should consider tweaking the quality settings, using a disk cache, or pre-rendering it out into OpenEXR files.

There are quite a few quality settings to tweak, many with small or diminishing returns; depending on the settings, there is easily a variation of 10x in rendering time. As a hint to those interested in reducing process time, try experimenting with the Proxy, Number of Iterations, and Number of Warps sliders and changing the filtering to Bilinear.

OpticalFlow can only work with the frames you allow. If you trim an upstream Loader to a sub frame range, OpticalFlow cannot see beyond the subframe range, even if there are additional frames on disk that it could use.

If the footage input flickers on a frame-by-frame basis, it is a good idea to deflicker the footage beforehand.

Currently, OpticalFlow will have to render twice for a downstream Time Stretcher or Time Speed to evaluate. This is because Time Speed needs A. FwdVec and B. BackVec to work, but OpticalFlow generates A. BackVec and A. FwdVec when it processes.

When pre-generating Optical Flow vectors, consider adding a SmoothMotion node afterward with smoothing for forward/ backward vectors enabled.

Inputs/Outputs

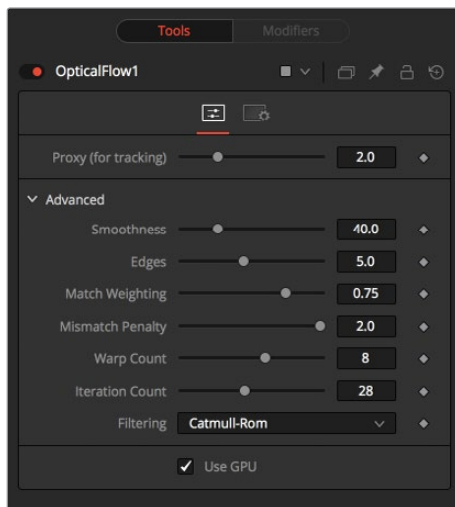
Input

This is the sequence of images for which you want to compute optical flow.

Output

This is the sequence of images with the optical flow stored in its Vector and BackVector channels.

Controls



Proxy (for tracking)

The input images are re-sized down by the proxy scale and tracked to produce the optical flow. This option is purely to speed up calculation of 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 speed up and a proxy scale of 3 will give a 9x speed up.

Advanced

The Advanced Control section has parameter controls to tune the OpticalFlow vector 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 will not be needed.

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 smoothing based upon 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 Loose, the optical flow becomes smoother and tends to overshoot edges. When it is set Tight, details from the color images start to slip into the optical flow, which is not desirable, and edges in the flow more tightly align with the edges in the color images. As a rough guideline, if you are using the disparity to produce a Z-channel for post effects like DoF, then prefer it tight, and if you are using the disparity to do interpolation, you might want to be looser.

In general, if it is too tight, there can be problems with streaked out edges when the optical flow is used for interpolation.

Match Weight

This controls how the matching of neighborhoods in the subsequent image is done. When set to Match Color, large structural color features are matched. When set to Match Edges, small sharp variations in the color are matched. Typically, a good value for this slider is in the [0.7, 0.9] range although on some shots, values closer to 0.0 work well. 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 do 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 gives the 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.

Number of Warps

Turning this option down makes the optical flow computations faster. In particular, the computational time depends linearly upon this option. To understand what this option does, you need to 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 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 what the optical flow is doing at the same time.


Number of Iterations

Turning this option down makes the computations faster. In particular, the computational time depends linearly upon this option. Just like adjusting Number of Warps, at some point adjusting this option higher will give 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 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.

Repair Frame [REP]



RepairFrame is used for replacing damaged or missing frames or portions of frames with scratches or other temporally transient artifacts.

You can use the Mask input to limit the repairs to certain areas. If your footage varies in color from frame to frame, sometimes the repair can be noticeable because, to fill in the hole, RepairFrame needs to pull color values from adjacent frames. Consider some kind of deflickering, color correction, and/or using a soft edged mask to help reduce these kinds of artifacts.

RepairFrame replaces parts of a frame by examining its two neighboring frames and thus requires three sequential frames as opposed to TimeStretcher/TimeSpeed, which work between two sequential frames to create a new in-between frame.

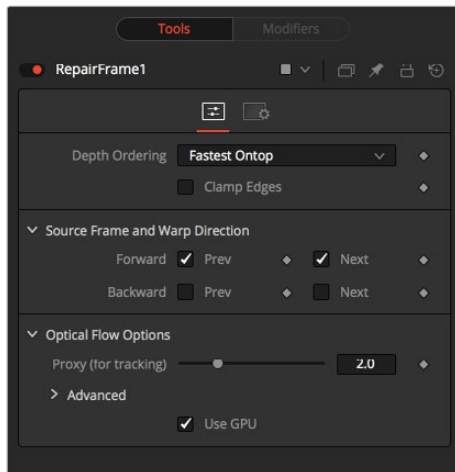
Unlike TimeStretcher and TimeSpeed, RepairFrame does not require, nor need, input Vector/BackVectors in order to work. An upstream OpticalFlow is not required.

RepairFrame internally will compute the optical flow it needs. This can make it slow to process.

RepairFrame will not pass through, but rather destroy, any aux channels after the computation is done.

See the Optical Flow node for control and setting information.

Controls



Depth Ordering

The Depth Ordering is used to determine which parts of the image should be rendered on top. This is best explained by an example.

In a locked off camera shot where a car is moving through the frame, the background does not move, so it will produce small or slow vectors, while the car will produce larger or faster vectors.

The depth ordering in this case is Fastest Ontop since the car will draw over the background.

In a shot where the camera pans to follow the car, the background will have faster vectors and the car will have slower vectors, so the Depth ordering method would be Slowest Ontop.

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 only use clamp edges 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 only have one checkbox enabled you can use a larger softness at around 0.03.

Source Frame and Warp Direction

This control set allows for the choosing of 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:** Will take the previous frame and use the Forward vector to interpolate the new frame.
- **Next Forward:** Will take the next frame in the sequence and use the Forward vector to interpolate the new frame.
- **Prev Backward:** Will take the previous frame and use the Back Forward vector to interpolate the new frame.
- **Next Backward:** Will take the next frame in the sequence and use the Back vector to interpolate the new frame.
- **Advanced:** Please see the Advanced Controls chapter of Optical flow.

Smooth Motion [SM]



This node takes an image sequence and smooths it using optical flow to look at neighboring frames.

It is important that the input sequence has pre-computed Vector and BackVector channels contained in it, otherwise this tool will print error messages.

Check on the channels you want to temporally smooth. If a channel selected for smoothing is not present, SmoothMotion will not fail nor will it print any error messages.

SmoothMotion was initially designed for smoothing of the Disparity channel, where it helps reduce temporal edge/fringing artifacts.


It can also be used to smooth the Vector and BackVector channels, however do be aware that in various situations this can make the interpolated results worse, especially if there are conflicting motions or objects in the shot that move around erratically, jitter or bounce rapidly.

One thing you can try is using two or more SmoothMotion nodes in sequence to get additional smoothing. With one SmoothMotion node, the previous, current and next frames are examined for a total of 3; with 2 SmoothMotion nodes, 5 frames are examined, and with 3 SmoothMotion nodes, 7 frames are examined.

Another thing you can try is using 2 SmoothMotion nodes, but in the first node enable the smoothing of the Vector and BackVector channels, and in the second SmoothMotion enable 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.

Controls



Channel

SmoothMotion can be applied to more than just the RGBA channels. It can also do all of the other aux channel groups in Fusion.

Tween [TW]



Tween constructs an in between frame by interpolating between two frames using the optical flow. Tween is nearly identical in functionality to TimeSpeed and TimeStretcher. 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 BackVector aux channels stored in the images, and it must manually generate the optical flow. There is no need to add an upstream OpticalFlow node. The generated optical flow is thrown away and not stored back into the output frames.

Since optical flow is based upon 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 will destroy any input aux channels. See the Optical Flow node for control and setting information.

Inputs/Outputs

InputA

This is an image to interpolate from.

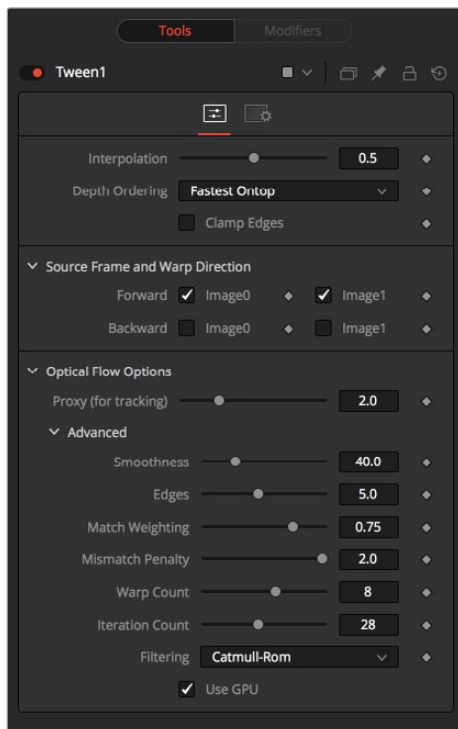
InputB

This is an image to interpolate to.

Output

This is the interpolated image.

Controls



Interpolation Parameter

This option determines where the frame we are interpolating is, relative to the two source frames A and B. An Interpolation Parameter of 0.0 will give frame A back, a parameter of 1.0 will give frame B back, and a parameter of 0.5 will give a result that is halfway between A and B.

Depth Ordering

The Depth ordering is used to determine which parts of the image should be rendered on top. This is best explained by an example.

In a locked off camera shot where a car is moving through the frame, the background does not move, so will produce small or slow vectors. The car will produce larger or faster vectors.

The Depth Ordering in this case is Fastest Ontop since the car will draw over the background.

In a shot where the camera pans to follow the car, the background will have faster vectors and the car will have slower vectors, so the Depth Ordering method would be Slowest Ontop.

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 only use clamp edges 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 only have one checkbox enabled you can use a larger softness at around 0.03.

Source Frame and Warp Direction

This control set allows for the choosing of 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:** Will take the previous frame and use the Forward vector to interpolate the new frame.
- **Next Forward:** Will take the next frame in the sequence and use the Forward vector to interpolate the new frame.
- **Prev Backward:** Will take the previous frame and use the Back Forward vector to interpolate the new frame.
- **Next Backward:** Will take the next frame in the sequence and use the Back vector to interpolate the new frame.

Chapter 48


Paint Node

This chapter details the Paint node available in Fusion.

Contents

Paint	985
Hot Keys	989
Modifiers	990
Circle	990
CloneMultistroke	991
Copy Ellipse	994
Copy Polyline	995
Copy Rectangle	997
Fill	998
Multistroke	999
Paint Group	1003
Polyline Stroke	1004
Stroke	1007

Paint



Paint is an extremely flexible, stroke-based system for making changes directly to a series of images. Use the Paint node for wire and rig removal, image cloning, or to rapidly create custom masks and mattes.

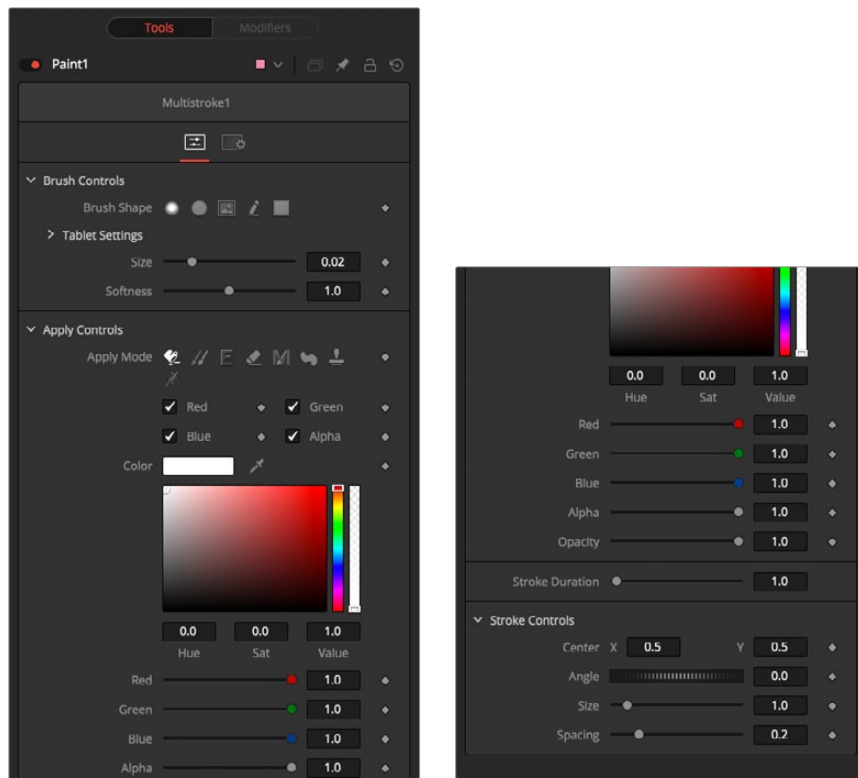
Fusion's paint can even be used to create new images and artistic elements from scratch.

Each Paint node is made up of a series of brush strokes. These strokes are vector shapes created directly on a view. The type of brush used, the size of the stroke and the effect of the stroke on the image are all user-controllable. A wide range of apply modes and brush types are available.

Brush strokes can be made into 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 re-ordered, deleted and modified with virtually infinite flexibility.

Controls



Not all of the controls described here appear in all modes. Certain controls are only useful in a specific Paint mode and are hidden when they are not applicable. Additionally, several of the controls are considered to be self-explanatory; the purpose of a Center control, Angle or Size control should be relatively straightforward to determine.

To reduce complexity, these controls are not all described.

Color Space

The Color Space array of buttons is only visible when the current mode is set to Fill. These are used to select the color space when sampling colors around the fill center for inclusion in the fill range.

R, G, B and Alpha

When selected, these checkboxes reflect which color channel is being painted. For example, with R, G, and B off, and Alpha on, painting will occur on the Alpha channel.

Brush Controls

Brush Shape

- **Soft Brush:** The Soft Brush type is a circular brush with soft edges. Modify the size of the brush in the Viewer by holding the Command or Ctrl key down while dragging the mouse.
- **Circular Brush:** A Circular Brush is a brush shape with hard edges. Resize this brush interactively.
- **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. See “Creating Custom Brushes” later in this chapter.
- **Single Pixel Brush:** The Single Pixel Brush is perfect for fine detail work, creating a brush exactly one pixel in size. No anti-aliasing is applied to the single pixel brush.
- **Square Brush:** A Square Brush is a brush shape with hard edges.

Vary Size

- **Constant:** The brush will be a constant size over the stroke.
- **With Pressure:** The stroke size will vary with the actual applied pressure.
- **With Velocity:** The stroke size will vary with the speed of painting. The faster the stroke, the thinner it is.

Vary Opacity

- **Constant:** The Constant brush will be a constant transparency over the entire stroke.
- **With Pressure:** The stroke transparency will vary with the applied pressure.
- **With Velocity:** The stroke transparency will vary 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 sources brush images.

Node

The image source is derived from the output of a node on the node tree. Drag the node into the Source node input to set the source.

Clip

The image source is derived from an image or sequence on disk. Any file supported by Fusion's Loader can be used. Locate the file using the filename Clip browser that appears to set the clip used as a source.

Brush

Images stored in the Fusion > Brushes directory are used as a brush for the Paint node. Select the brush from the menu that appears.

Apply Controls

Apply Mode


- **Color:** The Color Apply Mode paints simple colored strokes on the screen. When used in conjunction with an image brush, it can also be used to tint the brush.
- **Clone:** The Clone Apply Mode copies portions of one image into another image, or to clones from the same image using adjustable positions and time offsets. 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 with out 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 brush stroke as a guide.
- **Stamp:** Stamp 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 drawing them in toward the stroke.

Stroke Controls

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 straight line along the underlying vector shape that composes a stroke or polyline stroke). Increasing the value of this slider increases the density of the stroke, whereas decreasing the value is likely to cause the stroke to assume the appearance of a dotted line.




Stroke Animation:

The Stroke Animation menu control provides several pre-built animation effects that can be applied to a paint stroke. This menu only appears for Vector strokes.


- **All Frames:** This default displays the stroke for All Frames of the project where a valid target image is available to 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 will be added to the paint stroke that precisely duplicates the timing of the paint stroke's creation. The stroke will be 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 mode 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 will perform 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 mode will apply a Write On and then a Write Off animation mode to the stroke. Trail: Selecting the Trail mode will cause 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 segment of the stroke that follows the stroke as if it were a path. As with the Write On and Off effects, this will start at the frame that is current when the animation mode is selected. The timing of the animation can be adjusted manually using the Spline or Timeline Editors.
- **Duration:** Duration sets the duration of each stroke in frames. This control is only present for Multistrokes 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 will have a duration in the Timeline that can be trimmed independently from one stroke to the next. The duration can be set to 0.5, which will allow 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. Drag the low value upward to give the impression that the stroke is being erased, or drag the high value from 0.0 to 1.0 to give the impression that the stroke is being drawn 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, Write Off modes of the Stroke Animation menu.
- **Make Editable:** This button only appears for Vector strokes. Clicking on 'Make Editable' turns the current stroke into a polyline spline so that the shape can be adjusted or animated.

Working with Paint Strokes



- **Strokes Menu:** To create a new paint stroke, you can select the type of stroke you want from the menu that is displayed in the view. When a paint stroke is selected or edited, a menu is displayed in the view to select different editing options.

Types of Paint Strokes



- **Circle:** Creates a circular shape with animatable control over radius and center.
- **Clone Multistroke:** Like the Multistroke described further below, but specifically meant to clone elements from one image to the other.
- **Copy Polyline:** A Polyline area with animatable offset to clone elements from one image to the other.
- **Copy Rectangle:** A rectangular shape area with animatable offset to clone elements from one image to the other.
- **Fill:** Fills an area of the image based on adjustable color values.
- **Multistroke:** Perfect for those 100-strokes-per-frame retouching paint jobs like removing tracking markers. Much faster than the Stroke, but not editable later on.
- **Paint Group:** Allows easy grouping of multiple strokes with full control over center and size.
- **Polyline Stroke:** A fully editable stroke based on animatable polylines. Can be connected to existing polylines like Masks or Animation Paths.
- **Rectangle:** Creates a rectangular area.
- **Stroke:** The “standard” Stroke. Fully animatable and editable. Might become slow if hundreds of strokes are used in an image; for huge amounts of strokes it is better to use MultiStroke.

Hot Keys

Hot key control over Paint makes it interactive to adjust painting styles and color without having to navigate menus.

While painting:

- Hold Ctrl while left-dragging to change brush size.
- Hold Alt while clicking to pick color.

While cloning:

- Alt+click to set the clone source position. Strokes start cloning from here.
- 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:

- Painting will “rub through” pixels onto the destination.
- Arrow keys will change the clone source position; you can also drag crosshair and angle control or Size slider.
- Alt+left or right will change the clone source angle.
- Alt+up or down will change the clone source size.
- Shift + Ctrl can be used with the above for greater or lesser adjustments [&] (Left and right square brackets) will change the clone source Time Offset (this requires a specific Clone Source node to be set in the Source Node field).

Copy Rect/Ellipse:

Hold Shift while dragging out the source to constrain the shape.

With single stroke selected (not multi or polyline strokes):

Press X or Y to flip the stroke.

Paint Groups:

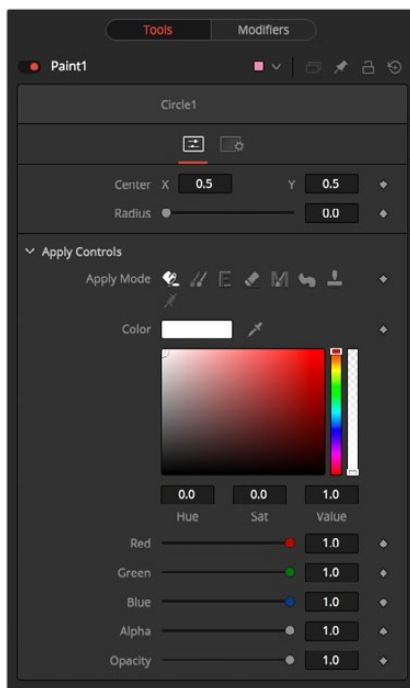
Ctrl+drag to change the position of a group's crosshair, without changing the position of the group.

Modifiers

Circle

The Circle only works on Paint nodes. Creates a circular shape with animatable control over radius and center. It can be applied by clicking on Circle in the Paint node's Stroke menu.

Circle Controls

**Center**

The Center of the circle.

Radius

The Radius of the circle.

Apply Mode

For details see the Apply Controls of the Paint node.

Color

The Color of the circle.

Opacity

The Opacity of the circle.

CloneMultistroke

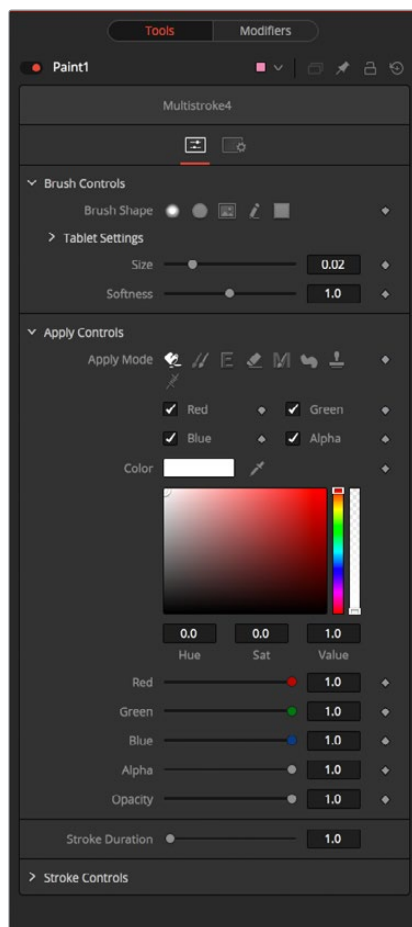
The CloneMultistroke is pretty much like the normal Multistroke, but specifically meant to clone areas from one image to the other. In addition to the Multistroke, it has a Source Node input, described later in this manual.

- Perfect for those 100-strokes-per-frame retouching paint jobs like removing tracking markers.
- Much faster than the Stroke, but not editable later on.
- It can be applied by clicking on the CloneMultistroke button in the Paint node's Stroke menu.

CloneMultistroke Controls

Not all of the controls described here appear in all modes. Certain controls are only useful in a specific paint mode and are hidden when they are not applicable. Additionally, several of the controls are considered to be self-explanatory; the purpose of a Center control, Angle or Size control should be relatively straightforward to determine.

To reduce complexity, these controls are not all described.



Brush Shape

- **Soft Brush:** The Soft Brush type is a circular brush with soft edges. Modify the size of the brush in the Viewer by holding the Command or Ctrl key down while dragging the mouse.
- **Circular Brush:** A Circular Brush is a brush shape with hard edges. Resize this brush interactively.
- **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. See Creating Custom Brushes later in this chapter.
- **Single Pixel Brush:** The Single Pixel Brush is perfect for fine detail work, creating a brush exactly one pixel in size. No anti-aliasing is applied to the single pixel brush.
- **Square Brush:** A Square Brush is a brush shape with hard edges.

Vary Size

- **Constant:** The brush will be a constant size over the stroke.
- **With Pressure:** The stroke size will vary with the actual applied pressure.
- **With Velocity:** The stroke size will vary with the speed of painting. The faster the stroke, the thinner it is.

Vary Opacity

- **Constant:** The Constant brush will be a constant transparency over the entire stroke.
- **With Pressure:** The stroke transparency will vary with the applied pressure.
- **With Velocity:** The stroke transparency will vary 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 sources brush images.

- **Node:** The image source is derived from the output of a node on the node tree. Drag the node into the Source Node input to set the source.
- **Clip:** The image source is derived from an image or sequence on disk. Any file supported by Fusion's Loader can be used. Locate the file using the filename Clip browser that appears to set the clip used as a source.
- **Brush:** Images stored in the Fusion > Brushes directory are used as a brush for the Paint node. Select the brush from the menu that appears.

Apply Controls

Apply Mode

- **Color:** The Color apply mode paints simple colored strokes on the screen. When used in conjunction with an image brush, it can also be used to tint the brush.
- **Clone:** The Clone apply mode copies portions of one image into another image, or clones from the same image using adjustable positions and time offsets. 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 brush stroke as a guide.
- **Stamp:** Stamp 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 drawing them in toward the stroke.
- **Source Node:** Shows which node's image output is used to clone from.


Stroke Controls

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 click-dragging the mouse pointer.

Spacing

The Spacing slider determines the distance between dabs (samples used to draw a straight line along the underlying vector shape that composes a stroke or polyline stroke). Increasing the value of this slider increases the density of the stroke, whereas decreasing the value is likely to cause the stroke to assume the appearance of a dotted line.



Stroke Animation

The Stroke Animation menu control provides several pre-built animation effects that can be applied to a paint stroke. This menu only appears for Vector strokes.

- **All Frames:** This default displays the stroke for All Frames of the project where a valid target image is available to 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 will be added to the paint stroke that precisely duplicates the timing of the paint stroke's creation. The stroke will be 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 mode 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 will perform 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 mode will apply a Write On and then a Write Off animation mode to the stroke.

- **Trail:** Selecting the Trail mode will cause 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 segment of the stroke that follows the stroke as if it were a path. As with the Write On and Off effects, this will start at the frame that is current when the animation mode is selected. The timing of the animation can be adjusted manually using the Spline or Timeline Editors.

Duration

Duration sets the duration of each stroke in frames. This control is only present for Multistrokes or when the stroke animation mode is set to Limited Duration. It is most commonly employed for frame-by-frame rotoscoping through a scene.

NOTE: Each Vector stroke applied to a scene will have a duration in the Timeline that can be trimmed independently from one stroke to the next. The duration can be set to 0.5, which will allow 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. Drag the low value upward to give the impression that the stroke is being erased, or drag the high value from 0.0 to 1.0 to give the impression that the stroke is being drawn 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, Write Off modes of the stroke animation menu.

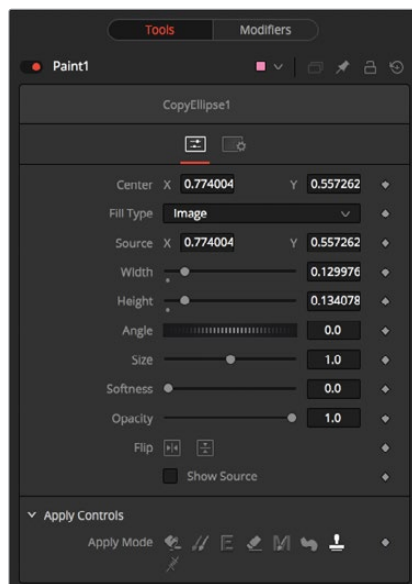
Make Editable

This button only appears for Vector strokes. Clicking on 'Make Editable' turns the current stroke into a polyline spline so that the shape can be adjusted or animated.

Copy Ellipse

The Copy Ellipse only works on Paint nodes. It creates an elliptical shape to clone elements from one area of the image to the other. It can be applied by clicking on the Copy Ellipse button in the Paint node's Stroke menu.

Copy Ellipse Controls



Center X Y

The Center of the Ellipse. Move this control to determine where the content will be copied to.

Image

The Source will be used to copy content to the destination.

Fill

The Fill color will be used to create a plain fill on the destination.

Source Center X Y

The Center of the Source Ellipse. Move this control to determine where the content will be copied from. Activate Show Source to see the onscreen controls.

Width/Height Sliders

The Width and Height of the ellipse. This can also be modified with the onscreen controls.

Angle

The Rotation of the ellipse. This can also be modified with the onscreen controls.

Size

The overall Size of the ellipse. This is a multiplier to the Width and Height settings.

Softness

The Softness of the edge of the ellipse.

Opacity

The overall opacity of the output.

Flip Horiz/Flip Vert

The area inside the ellipse will be flipped horizontally and/or vertically.

Show Source

Show the onscreen controls to adjust the position of the Source Ellipse.

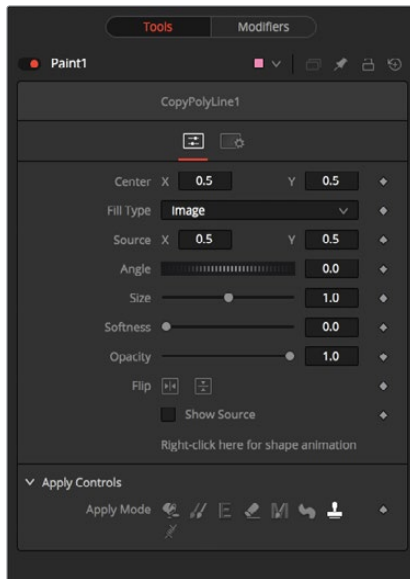
Apply Mode

For more details see the Apply Controls of the Paint node section.

Copy Polyline

The Copy Polyline only works on Paint nodes. It creates a user-definable Polyline to clone elements from one area of the image to the other. See the Polylines and Rotoscoping section of this manual. It can be applied by clicking on the Copy Polyline button in the Paint node's Stroke menu.

Copy Polyline Controls



Center X Y

The Center of the Polyline. Move this control to determine where the content will be copied to.

Image

The Source will be used to copy content to the Destination.

Fill

The Fill color will be used to create a plain fill on the Destination.

Source Center X Y

The Center of the Source Polyline. Move this control to determine where the content will be copied from. Activate Show Source to see the onscreen controls.

Width/Height Sliders

The Width and Height of the Polyline. This can also be modified with the onscreen controls.

Angle

The Rotation of the Polyline. This can also be modified with the onscreen controls.

Size

The overall Size of the Polyline. This is a multiplier to the Width and Height settings.

Softness

The Softness of the edge of the Polyline.

Opacity

The overall Opacity of the output.

Flip Horiz/Flip Vert

The area inside the Polyline will be flipped horizontally and/or vertically.

Show Source

Show the onscreen controls to adjust the position of the source Polyline.

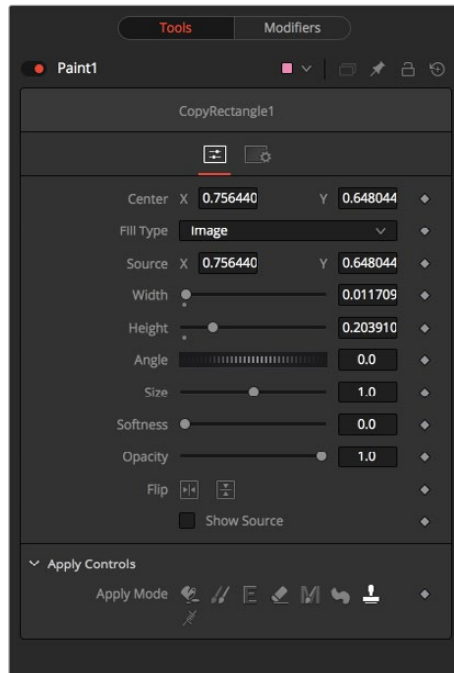
Apply Mode

For more details see the Apply Controls of the Paint node section.

Copy Rectangle

The Copy Rectangle only works on Paint nodes. It creates a user-definable Rectangle to clone elements from one area of the image to the other. It can be applied by clicking on the Copy Rectangle button in the Paint node's Stroke menu.

Copy Rectangle Controls



Center X Y

The Center of the Rectangle. Move this control to determine where the content will be copied to.

Image

The Source will be used to copy content to the Destination.

Fill

The Fill color will be used to create a plain fill on the Destination.

Source Center X Y

The Center of the Source Rectangle. Move this control to determine where the content will be copied from. Activate Show Source to see the onscreen controls.

Width/Height Sliders

The Width and Height of the Rectangle. This can also be modified with the onscreen controls.

Angle

The Rotation of the Rectangle. This can also be modified with the onscreen controls.

Size

The overall Size of the Rectangle. This is a multiplier to the Width and Height settings.

Softness

The Softness of the edge of the Rectangle.

Opacity

The overall Opacity of the output.

Flip Horiz/Flip Vert

The area inside the Rectangle will be flipped horizontally and/or vertically.

Show Source

Show the onscreen controls to adjust the position of the source Rectangle.

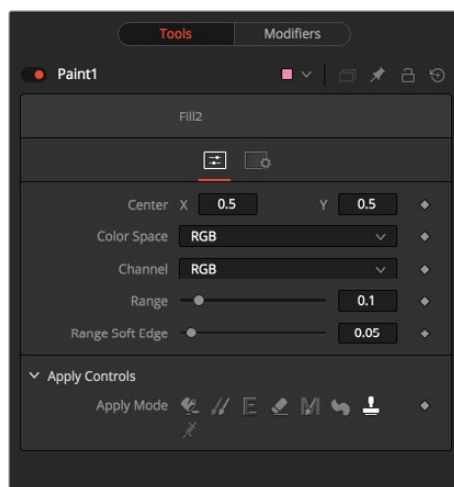
Apply Mode

For more details see the Apply Controls of the Paint node section.

Fill

The Fill only works on Paint nodes. It fills an area of the image with a user-definable color. From the way it works, it can be compared with the Wand mask. It can be applied by clicking on the Fill button in the Paint node's Stroke menu.

Fill Controls



Selection Point

The Selection Point is a pair of X and Y coordinates that determine from where in the source image the fill 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 fill 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 from an individual channel only.

The exact labels of the buttons will depend on the Color Space selected for the fill operation. If the color space is RGB, the options will be R, G, or B. If YUV is the color space, the options will be Y, U, or V.

Range

The Range slider controls the range of colors around the source color that will be included in the fill. If the value is left at 0.0, only pixels of exactly the same color as the source will be considered part of the fill. The higher the value, the more similar colors in the source will be considered to be wholly part of the fill.

Range Soft Edge

The Range Soft Edge determines the falloff range of the colors selected. Any pixel within the range defined above will be treated as 100% for the fill. If the soft range is set to 0.0, no other pixels will be considered for the fill. Increasing the soft range will increase the number of colors close to, but not quite within, the range that will be included in the fill. These pixels will be semi-transparent in the fill.

Apply Mode

For more details see the Apply Controls of the Paint node section.

Multistroke

The Multistroke is the standard stroke in the Paint node. Perfect for those 100-strokes-per-frame retouching paint jobs like removing tracking-markers. Much faster than the Stroke, but not editable later on.

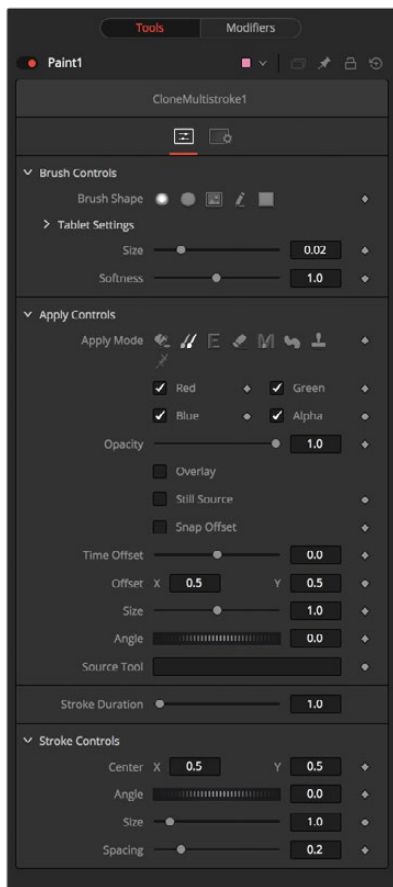
While Multistrokes aren't directly editable, they can be grouped with the PaintGroup modifier, then tracked, moved and rotated by animating the PaintGroup instead. Alternatively, the Edit Multistrokes and Combine Strokes node scripts can be used to convert multistrokes into ordinary, editable strokes, and vice versa.

It can be applied by clicking on the Multi Stroke button in the Paint node's Stroke menu.

Multistroke Controls

Not all of the controls described here appear in all modes. Certain controls are only useful in a specific paint mode and are hidden when they are not applicable. Additionally, several of the controls are considered to be self-explanatory the purpose of a Center control, Angle or Size control should be relatively straightforward to determine.

To reduce complexity, these controls are not all described.



Brush Controls

Brush Shape

- **Soft Brush:** The Soft Brush type is a circular brush with soft edges. Modify the size of the brush in the Viewer by holding Command or Ctrl down while dragging the mouse.
- **Circular Brush:** A Circular Brush is a brush shape with hard edges. Resize this brush interactively.
- **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. See “Creating Custom Brushes” later in this chapter.
- **Single Pixel Brush:** The Single Pixel Brush is perfect for fine detail work, creating a brush exactly one pixel in size. No anti-aliasing is applied to the single pixel brush.
- **Square Brush:** A Square Brush is a brush shape with hard edges.

Vary Size

- **Constant:** The brush will be a constant size over the stroke.
- **With Pressure:** The stroke size will vary with the actual applied pressure.
- **With Velocity:** The stroke size will vary with the speed of painting. The faster the stroke, the thinner it is.

Vary Opacity

- **Constant:** The Constant brush will be a constant transparency over the entire stroke.
- **With Pressure:** The stroke transparency will vary with the applied pressure.
- **With Velocity:** The stroke transparency will vary 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 sources brush images.

- **Node:** The image source is derived from the output of a node on the node tree. Drag the node into the Source Node input to set the source.
- **Clip:** The image source is derived from an image or sequence on disk. Any file supported by Fusion's Loader can be used. Locate the file using the filename Clip browser that appears to set the clip used as a source.
- **Brush:** Images stored in the Fusion > Brushes directory are used as a brush for the Paint node. Select the brush from the menu that appears.

Apply Controls

Apply Mode

- **Color:** The Color apply mode paints simple colored strokes on the screen. When used in conjunction with an Image brush, it can also be used to tint the brush.
- **Clone:** The Clone apply mode copies portions of one image into another image, or to clones from the same image using adjustable positions and time offsets. 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 brush stroke as a guide.
- **Stamp:** Stamp 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 drawing them in toward the stroke.

Stroke Controls

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 click-dragging the mouse pointer.

Spacing

The Spacing slider determines the distance between dabs (samples used to draw a straight line along the underlying vector shape that composes a stroke or polyline stroke). Increasing the value of this slider increases the density of the stroke, whereas decreasing the value is likely to cause the stroke to assume the appearance of a dotted line.



Stroke Animation

The Stroke Animation menu control provides several pre-built animation effects that can be applied to a paint stroke. This menu only appears for Vector strokes.

- **All Frames:** This default displays the stroke for All Frames of the project where a valid target image is available to 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 will be added to the paint stroke that precisely duplicates the timing of the paint stroke's creation. The stroke will be 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 mode 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 will perform 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 mode will apply a Write On and then a Write Off animation mode to the stroke.
- **Trail:** Selecting the Trail mode will cause 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 segment of the stroke that follows the stroke as if it were a path. As with the Write On and Off effects, this will start at the frame that is current when the Animation mode is selected. The timing of the animation can be adjusted manually using the Spline or Timeline Editors.
- **Duration:** Duration sets the duration of each stroke in frames. This control is only present for Multistrokes or when the Stroke Animation mode is set to Limited Duration. It is most commonly employed for frame-by-frame rotoscoping through a scene.

NOTE: Each Vector stroke applied to a scene will have a duration in the Timeline that can be trimmed independently from one stroke to the next. The duration can be set to 0.5, which will allow 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. Drag the low value upward to give the impression that the stroke is being erased, or drag the high value from 0.0 to 1.0 to give the impression that the stroke is being drawn 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/Write Off modes of the Stroke animation menu.
- **Make Editable:** This button only appears for Vector strokes. Clicking on 'Make Editable' turns the current stroke into a polyline spline so that the shape can be adjusted or animated.

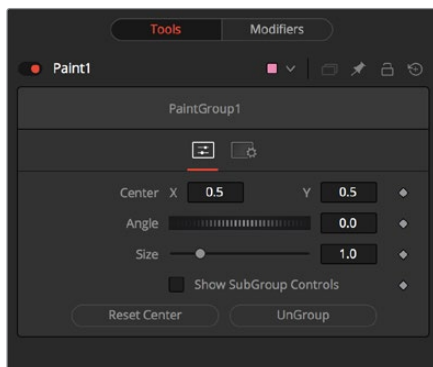
Paint Group

The Paint Group only works on Paint strokes and objects. It groups all currently selected paint strokes together, making it easier to handle them all together.

Multistroke modifiers can be grouped as well, which includes all the multistrokes contained within that modifier. By animating the Paint Group, it is thus possible to move, track and rotate a collection of multistrokes.

It can be applied by clicking on the Paint Group button in the Paint node's Stroke menu.

Paint Group Controls



Center X Y

The Center of the Group. Modify this control to move the group around.

Angle

The entire group can be rotated using this control.

Size

Scales the entire group.

Show Subgroup Controls

Displays the controls of the individual strokes in the group and allows the user to modify them.

Reset Center

Puts the Center back to the position where it was when the group was created.

Ungroup

Disassembles the group back into individual strokes and deletes the Group modifier. For temporary opening of the group, use the Show Subgroup controls switch instead.

Polyline Stroke

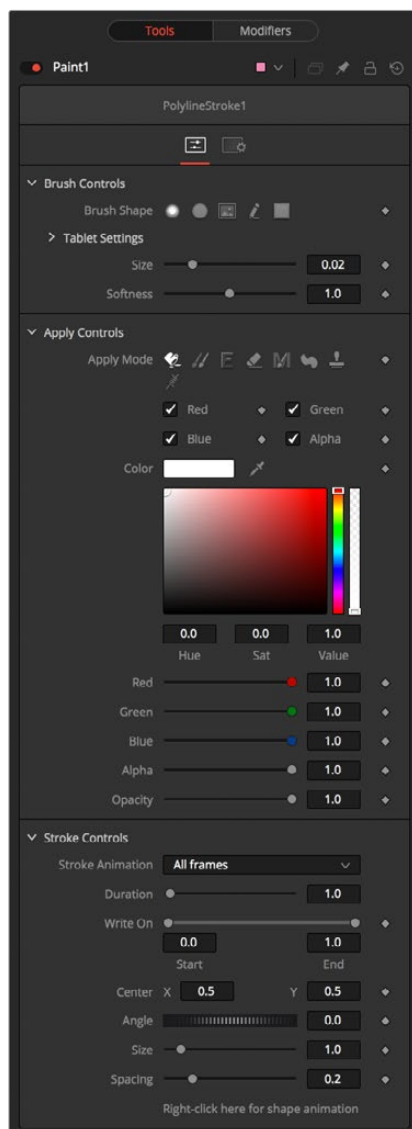
Fusion offers a Polyline Stroke mode. This provides the ability to create and manipulate a stroke in the same way that a path or mask might be created. To add a polyline stroke to the node tree, select the Polyline button from the Paint node's Stroke toolbar and click in the view to add the first point. Continue clicking to add additional points to the polyline.

Notice that the usual Polyline toolbar buttons will appear in the Viewer toolbar. Polyline strokes are created in Click Append mode by default, but they can also be created in Draw Append mode.

Polyline Stroke Controls

Not all of the controls described here appear in all modes. Certain controls are only useful in a specific paint mode and are hidden when they are not applicable. Additionally, several of the controls are considered to be self-explanatory; the purpose of a center control, angle or size control should be relatively straightforward to determine.

To reduce complexity, these controls are not all described.



Brush Controls

Brush Shape

- **Soft Brush:** The Soft Brush type is a circular brush with soft edges. Modify the size of the brush in the Viewer by holding Command or Ctrl down while dragging the mouse.
- **Circular Brush:** A Circular Brush is a brush shape with hard edges. Resize this brush interactively.
- **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. See “Creating Custom Brushes” later in this chapter.
- **Single Pixel Brush:** The Single Pixel Brush is perfect for fine detail work, creating a brush exactly one pixel in size. No anti-aliasing is applied to the single pixel brush.
- **Square Brush:** A Square Brush is a brush shape with hard edges.

Vary Size

- **Constant:** The brush will be a constant size over the stroke.
- **With Pressure:** The stroke size will vary with the actual applied pressure.
- **With Velocity:** The stroke size will vary with the speed of painting. The faster the stroke, the thinner it is.

Vary Opacity

- **Constant:** The Constant brush will be a constant transparency over the entire stroke.
- **With Pressure:** The stroke transparency will vary with the applied pressure.
- **With Velocity:** The stroke transparency will vary 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 sources brush images.

- **Node:** The image source is derived from the output of a node on the node tree. Drag the node into the Source Node input to set the source.
- **Clip:** The image source is derived from an image or sequence on disk. Any file supported by Fusion’s Loader can be used. Locate the file using the filename Clip browser that appears to set the clip used as a source.
- **Brush:** Images stored in the Fusion > Brushes directory are used as a brush for the Paint node. Select the brush from a the menu that appears.

Apply Controls

Apply Mode

- **Color:** The Color apply mode paints simple colored strokes on the screen. When used in conjunction with an image brush, it can also be used to tint the brush.
- **Clone:** The Clone apply mode copies portions of one image into another image, or to clones from the same image using adjustable positions and time offsets. 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 brush stroke as a guide.
- **Stamp:** Stamp 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 drawing them in toward the stroke.
- **Source Node:** Shows which node's image output is used to clone from.


Stroke Controls

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 click-dragging the mouse pointer.

Spacing

The Spacing slider determines the distance between dabs (samples used to draw a straight line along the underlying vector shape that composes a stroke or polyline stroke). Increasing the value of this slider increases the density of the stroke, whereas decreasing the value is likely to cause the stroke to assume the appearance of a dotted line.



Stroke Animation

The Stroke Animation menu control provides several pre-built animation effects that can be applied to a paint stroke. This menu only appears for Vector strokes.

- **All Frames:** This default displays the stroke for All Frames of the project where a valid target image is available to 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 will be added to the paint stroke that precisely duplicates the timing of the paint stroke's creation. The stroke will be 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 mode 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 will perform 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 mode will apply a Write On and then a Write Off animation mode to the stroke.
- **Trail:** Selecting the Trail mode will cause 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 segment of the stroke that follows the stroke as if it were a path. As with the Write On and Off effects, this will start at the frame that is current when the Animation mode is selected. The timing of the animation can be adjusted manually using the Spline or Timeline Editors.

Duration

Duration sets the duration of each stroke in frames. This control is only present for Multistrokes or when the stroke animation mode is set to Limited Duration. It is most commonly employed for frame-by-frame rotoscoping through a scene.

NOTE: Each Vector stroke applied to a scene will have a duration in the Timeline that can be trimmed independently from one stroke to the next.

The duration can be set to 0.5, which will allow 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. Drag the low value upward to give the impression that the stroke is being erased, or drag the high value from 0.0 to 1.0 to give the impression that the stroke is being drawn 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/Write Off modes of the stroke animation menu.

Right Click Here for Shape Animation

To animate the Polyline shape like Polyline Masks, right click on this control and choose animate.

Stroke

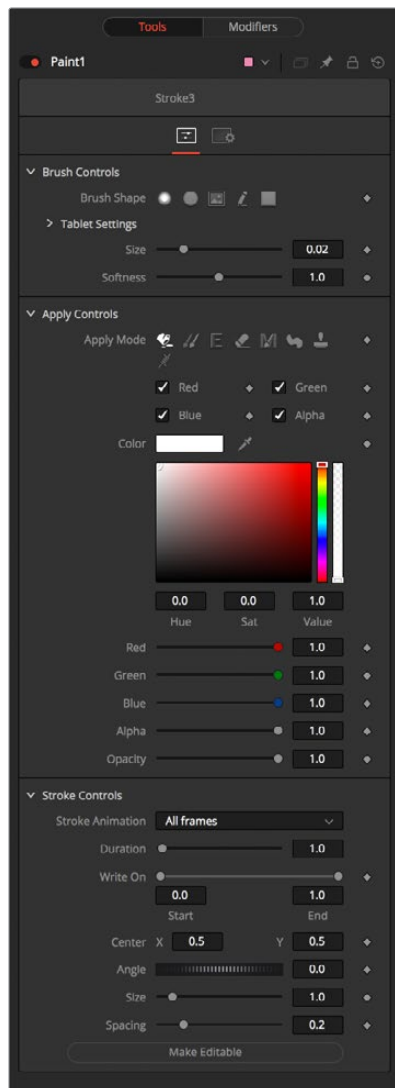
When the Paint node is first added to the node tree, the brush type is set to a medium sized soft, circular brush with a white Color apply mode.

Release the mouse button to end the creation of the stroke. To draw another stroke, simply click again and continue. A new stroke will be added to the image with the same settings as the last applied stroke. When the painting is complete, choose the Select button in the Paint toolbar to avoid accidentally adding new strokes.

While in Brush or Polyline modes, the controls displayed in the Node tab affect the next stroke created, not the strokes that are already created. To adjust the settings of an existing stroke, switch to Select mode using the Paint toolbar and select the stroke.

To add a stroke to the node tree, select the Polyline Button from the Paint Node's Stroke toolbar, place the mouse or pen over the image, click-hold the left mouse button and start drawing.

Stroke Controls



Not all of the controls described here appear in all modes. Certain controls are only useful in a specific paint mode and are hidden when they are not applicable. Additionally, several of the controls are considered to be self-explanatory; the purpose of a center control, angle or size control should be relatively straightforward to determine. To reduce complexity, these controls are not all described.

Brush Controls

Brush Shape

- **Soft Brush:** The Soft Brush type is a circular brush with soft edges. Modify the size of the brush in the Viewer by holding Command or Ctrl down while dragging the mouse.
- **Circular Brush:** A Circular Brush is a brush shape with hard edges. Resize this brush interactively.
- **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. See “Creating Custom Brushes” later in this chapter.
- **Single Pixel Brush:** The Single Pixel Brush is perfect for fine detail work, creating a brush exactly one pixel in size. No anti-aliasing is applied to the single pixel brush.
- **Square Brush:** A Square Brush is a brush shape with hard edges.

Vary Size

- **Constant:** The brush will be a constant size over the stroke.
- **With Pressure:** The stroke size will vary with the actual applied pressure.
- **With Velocity:** The stroke size will vary with the speed of painting. The faster the stroke, the thinner it is.

Vary Opacity

- **Constant:** The Constant brush will be a constant transparency over the entire stroke.
- **With Pressure:** The stroke transparency will vary with the applied pressure.
- **With Velocity:** The stroke transparency will vary 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 sources brush images.

- **Node:** The image source is derived from the output of a node on the node tree. Drag the node into the Source Node input to set the source.
- **Clip:** The image source is derived from an image or sequence on disk. Any file supported by Fusion's Loader can be used. Locate the file using the filename Clip browser that appears to set the clip used as a source.
- **Brush:** Images stored in the Fusion > Brushes directory are used as a brush for the Paint node. Select the brush from the menu that appears.

Apply Controls

Apply Mode

- **Color:** The Color apply mode paints simple colored strokes on the screen. When used in conjunction with an image brush, it can also be used to tint the brush.
- **Clone:** The Clone apply mode copies portions of one image into another image, or to clones from the same image using adjustable positions and time offsets. 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 brush stroke as a guide.
- **Stamp:** Stamp 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 drawing them in toward the stroke.
- **Source Node:** Shows which node's image output is used to clone from.

Stroke Controls

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 click-dragging the mouse pointer.

Spacing

The Spacing slider determines the distance between dabs (samples used to draw a straight line along the underlying vector shape that composes a stroke or polyline stroke). Increasing the value of this slider increases the density of the stroke, whereas decreasing the value is likely to cause the stroke to assume the appearance of a dotted line.



Stroke Animation

The Stroke Animation menu control provides several pre-built animation effects that can be applied to a paint stroke. This menu only appears for Vector strokes.

- **All Frames:** This default displays the stroke for All Frames of the project where a valid target image is available to 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 will be added to the paint stroke that precisely duplicates the timing of the paint stroke's creation. The stroke will be 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 mode 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 will perform 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 mode will apply a Write On and then a Write Off animation mode to the stroke.
- **Trail:** Selecting the Trail mode will cause 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 segment of the stroke that follows the stroke as if it were a path. As with the Write On and Off effects, this will start at the frame that is current when the Animation mode is selected. The timing of the animation can be adjusted manually using the Spline or Timeline Editors.

Duration

Duration sets the duration of each stroke in frames. This control is only present for Multistrokes or when the stroke Animation mode is set to Limited Duration. It is most commonly employed for frame-by-frame rotoscoping through a scene.

NOTE: Each Vector stroke applied to a scene will have a duration in the Timeline that can be trimmed independently from one stroke to the next. The duration can be set to 0.5, which will allow 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. Drag the low value upward to give the impression that the stroke is being erased, or drag the high value from 0.0 to 1.0 to give the impression that the stroke is being drawn 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, Write Off modes of the stroke animation menu.

Right Click Here for Shape Animation

To animate the shape, like Polyline Masks, right-click on this control and choose animate.

Chapter 49

Particle Nodes

This chapter details the Particle nodes available in Fusion.

Contents

pAvoid [PAV]	1014
pBounce [PBN]	1016
pChangeStyle	1018
pCustom	1020
pCustomForce	1022
pDirectionalForce	1025
pEmitter	1027
pFlock	1031
pFriction	1033
pGradientForce	1034
pImageEmitter	1036
pKill	1040
pMerge	1041
pPointForce	1041
pRender [PRN]	1043
pSpawn [PSP]	1048
pTangentForce [PTF]	1049
pTurbulence [PTR]	1051
pVortex [PVT]	1053
Particle Common Controls	1055

pAvoid [PAV]




The pAvoid node is used to create a region or area within the image that affected particles will attempt to avoid entering and/or crossing.

It has two primary controls, one that determines the distance from the region a particle should be before it begins to move away from the region, and another to determine 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 will not overcome its momentum and the particle will cross that region anyway.

Controls



Randomize

The 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 on the Randomize button to randomly select a new seed value, or adjust the slider to manually select a new seed value.

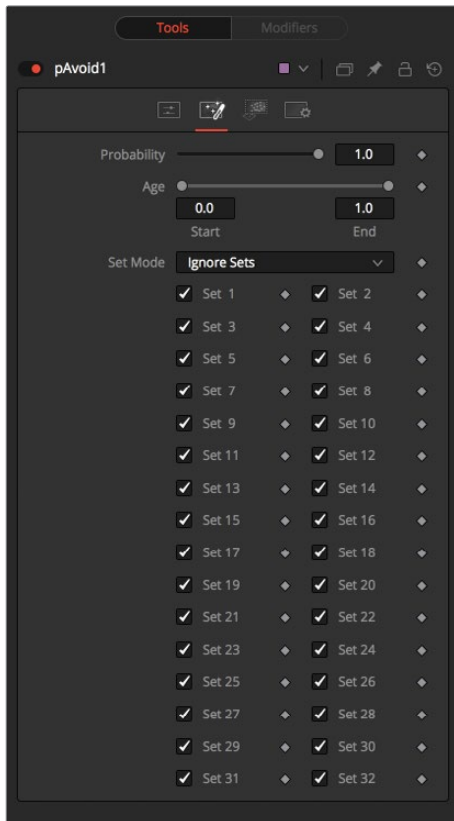
Distance

Determines the distance from the region a particle should be before it begins to move away from the region.

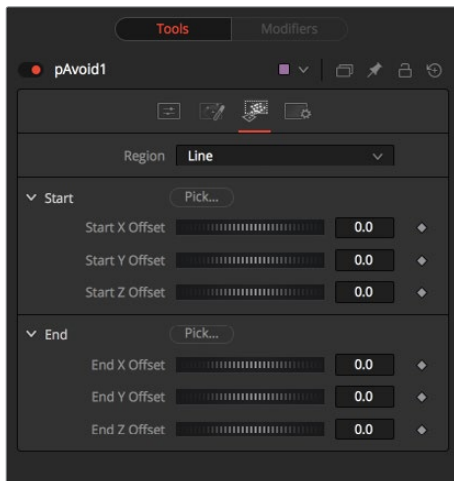
Strength

Determines how strongly the particle moves away from the region. Negative values will make the particles move toward the region instead.

Conditions Tab



Regions Tab



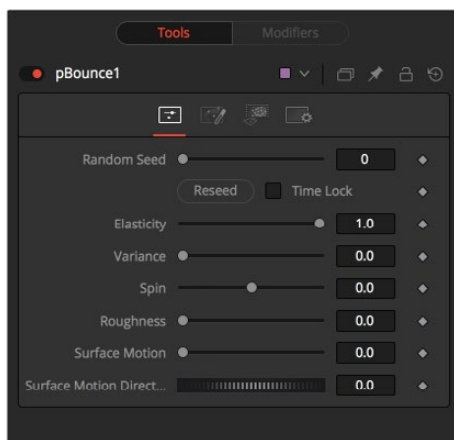
See Particle Common Controls in this chapter.

pBounce [PBN]



The pBounce node is used to create a region from which affected particles will bounce away when they come into contact with the region. The pBounce node has three main controls, as described below.

Controls



Randomize

The 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 on 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.

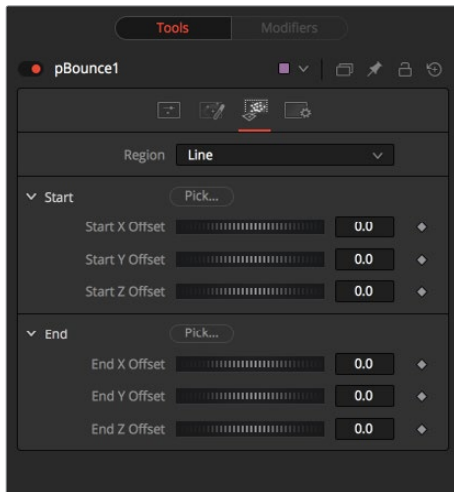
Surface Motion Direction

This thumbwheel control sets the angle relative to the bounce surface.

Conditions Tab




Regions Tab



See Particle Common Controls in this chapter.

pChangeStyle



The pChangeStyle node provides a mechanism for changing the appearance or style of particles that interact with a defined region. The primary controls in this node perfectly mirror those found in the Style tab of the pEmitter node. Particles that intersect or enter the region defined for this node will change as described by this node.

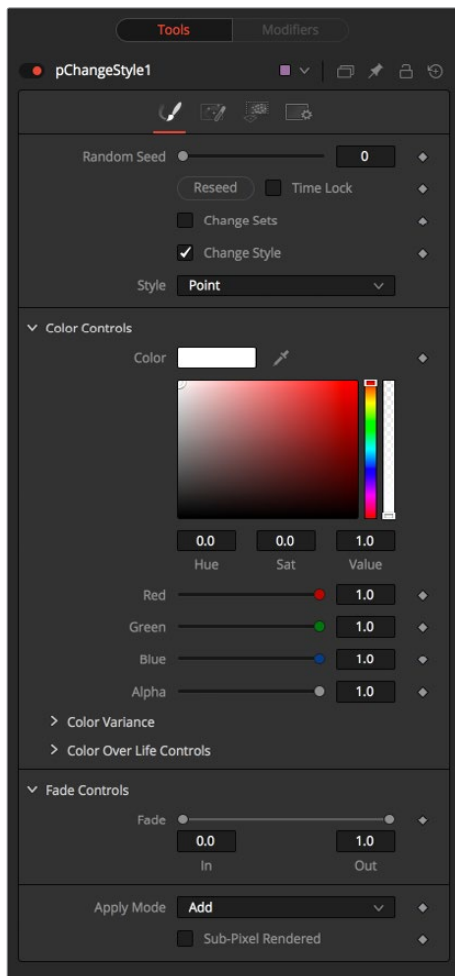
With the possible exception of the pCustom node, this is the only node that modifies the appearance of a particle, rather than its motion. It is often used to cause the appearance of particles changing in response to some event, like striking a barrier.

When using the pChangeStyle node in this fashion, it would be natural to assume that the node should be placed after the node causing the event. As an example, consider the creation of a particle system that appears to change its style after bouncing off a pBounce using a line region. In this case, the pChangeStyle node also uses a line region, positioned identically to the one in the pBounce node. Placing the pBounce before the pChangeStyle in the node tree causes the particles to bounce off the region before the pChangeStyle gets an opportunity to calculate its effect on the particle.

The result is that the particle is no longer intersecting with the pChangeStyle node's region, and so the style never changes.

As a rule, to create a change in style that appears to be caused by a physical event created by another modifier node in the node tree, the pChangeStyle node must be placed before that node for the effect to work properly.

Style Tab



Randomize

The 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 on the Randomize button to randomly select a new seed value, or adjust the slider to manually select a new seed value.

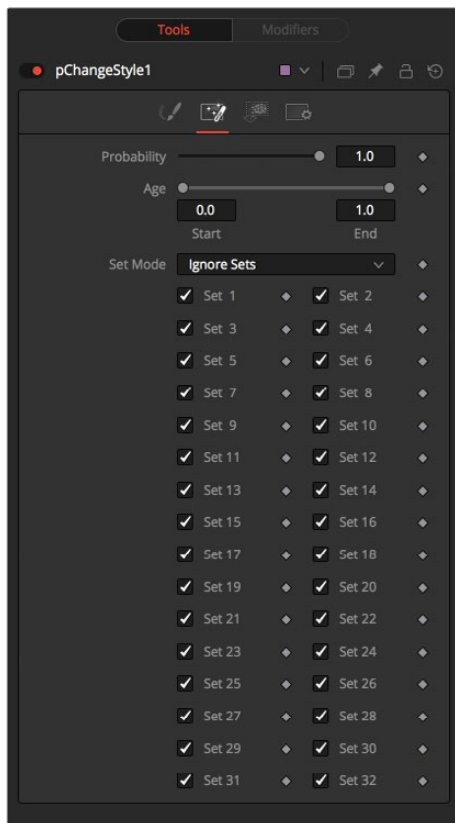
Change Sets

This option allows the user to change the particle's Set to become influenced by other forces than the original particle. See Chapter 70, "Particle Systems," to learn more about Sets.


Style

This option allows the user to change the particle's Style and thus the look. See Chapter 70, "Particle Systems," to learn more about Styles.

Conditions Tab




Regions Tab



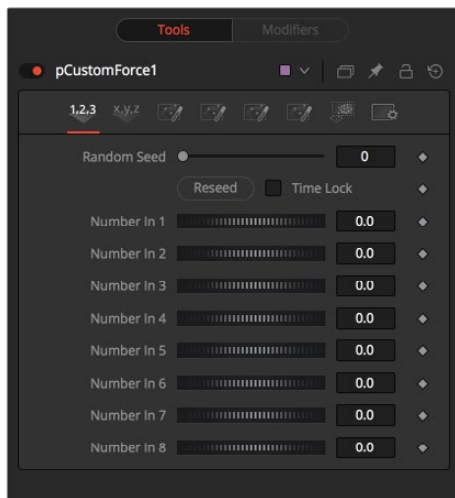
See Particle Common Controls in this chapter.

pCustom



The pCustom node is used to create custom expressions that affect the properties of particles. This node is almost identical to the Custom node, except that the calculations affect properties of the particles rather than the properties of a pixel.

Numbers Tab



Conditions and Regions Tab

See Particle Common Controls in this chapter.

The following particle properties are exposed to the pCustom control:

Particle Properties 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, rgnni, 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

Particle Properties exposed to the pCustom control	
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

Additional information on the custom class of nodes can be found in documentation for the Custom node.

All of the operators, functions and conditional statements described for that node apply to pCustom as well, including Pixel-read functions for the two image inputs (e.g., get1w(x,y), getz2b(x,y), and so on).

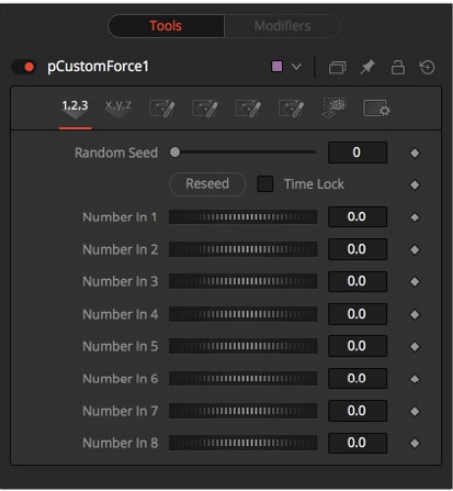
pCustomForce




The Particle Custom Force node allows you to change the forces applied to a particle system or subset. This node is quite likely the most complex and the most powerful node in Fusion. Any user moderately experienced with scripting or C++ programming, 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.

Conditions Tab

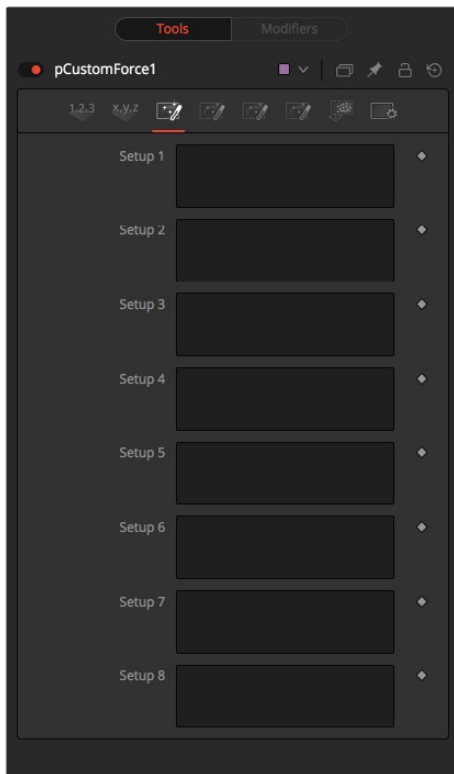


Regions Tab

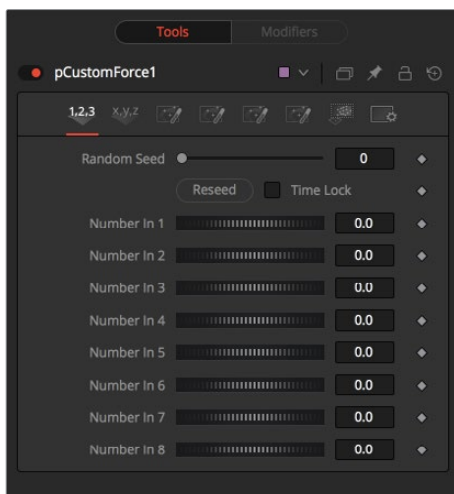


See Particle Common Controls in this chapter.


Setup Tab



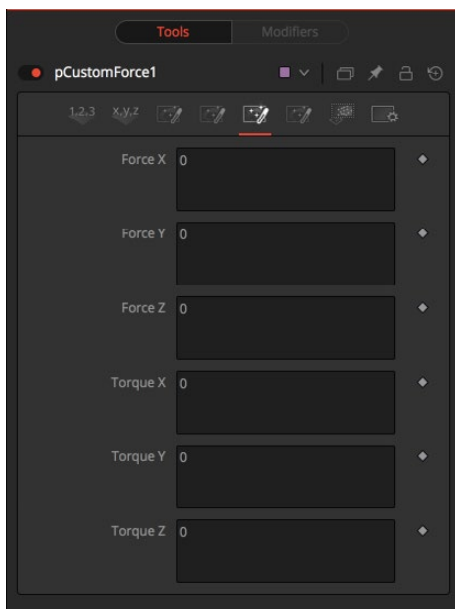
Numbers Tab



Inter Tab




Force Tab



Positions Tab




pDirectionalForce



This node applies a uni-directional force that pulls the affected particles in a specified direction. Its primary controls affect the strength of the force, and the angle of the forces pull along the X, Y, and Z axis.

As 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.

Controls



Randomize

The 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 on 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.

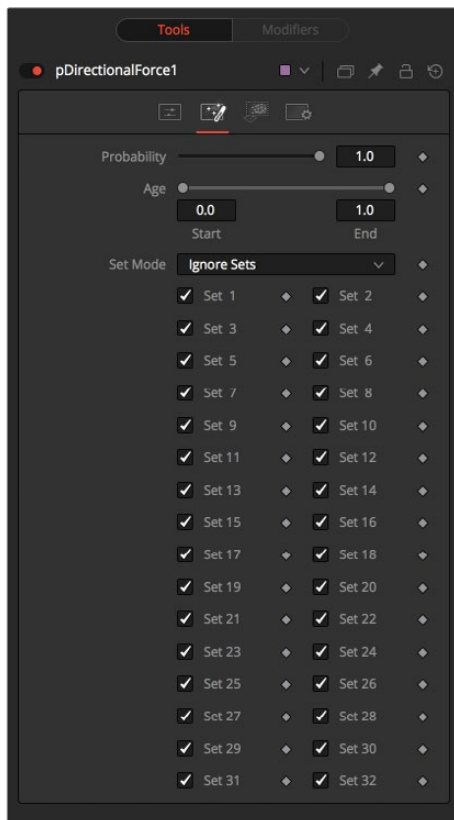
Direction

Determines the direction in X/Y Space.


Direction Z

Determines the direction in Z Space.

Conditions Tab



Regions Tab



See Particle Common Controls in this chapter.

pEmitter

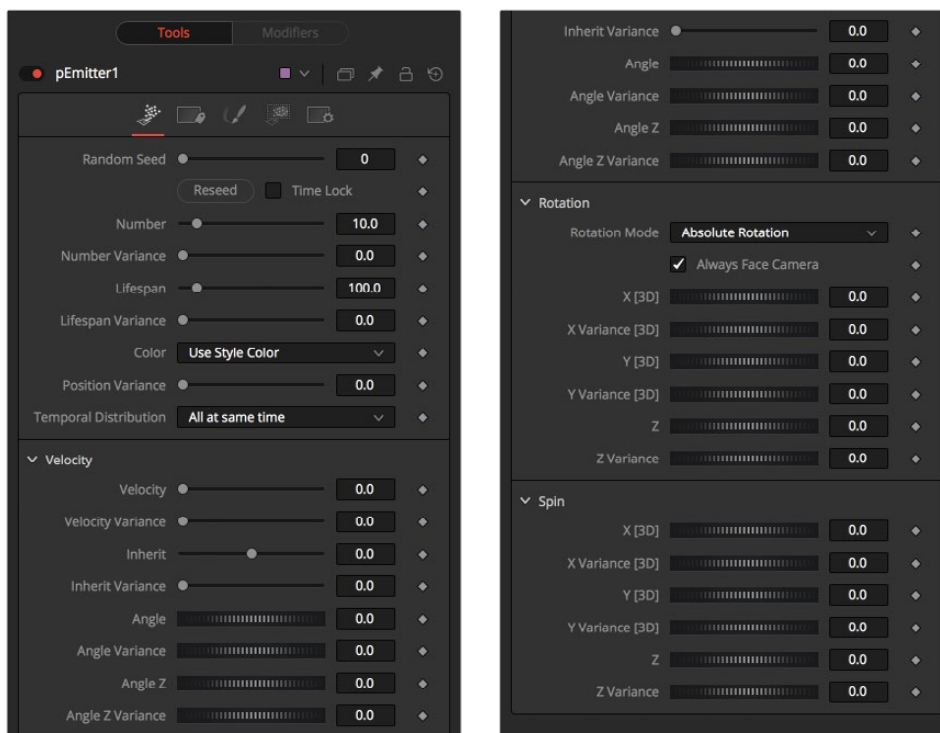


The pEmitter node is the main source of Particles (pImageEmitter is another) and will usually be the first node used in any new particle system. This node contains controls for setting the initial position, orientation and motion of the particles, as well as controls for the visual style of each particle.

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 on a viewer. To view the output of a particle system, add a pRender node after the pEmitter.

Controls

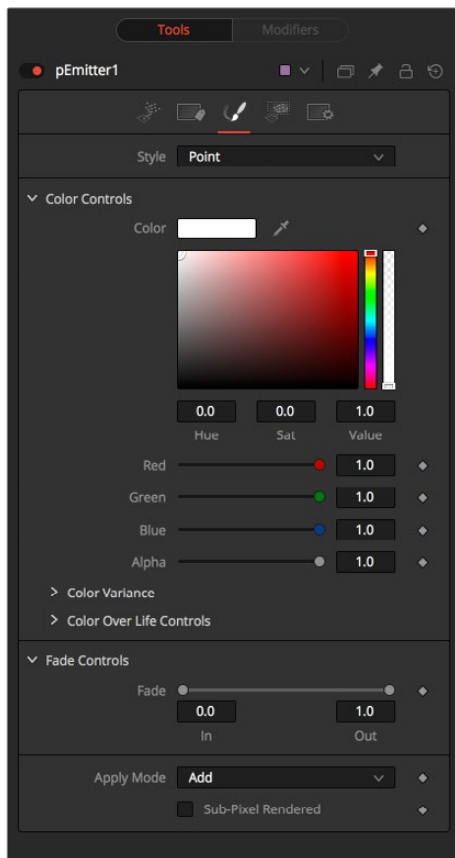
This tab contains settings that affect the physics of the particles emitted by the node. These settings do not directly affect the appearance of the particles. They modify behavior like velocity, spin, quantity and lifespan instead.



Randomize and Random Seed

The Random Seed slider is used to seed all of 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 on the Randomize button to automatically set a randomly chosen value for the Random Seed.

Style Tab



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 Source

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.

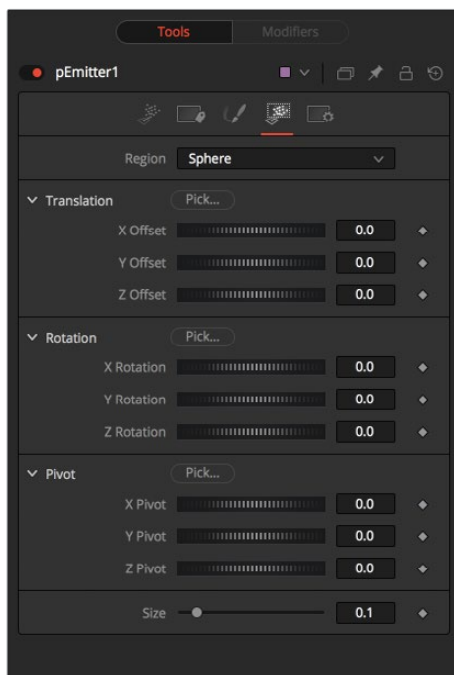
The alternate setting is Use Color From Region, which overrides the color settings from the Style tab and uses the color of the underlying bitmap region.

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.

Regions Tab



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.

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 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.

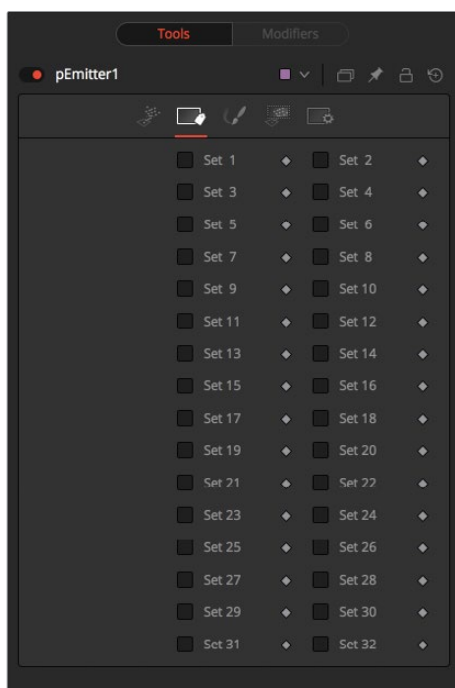
Spin XYZ and Spin Variance

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. They modify behavior like velocity, spin, quantity and lifespan instead.



pFlock

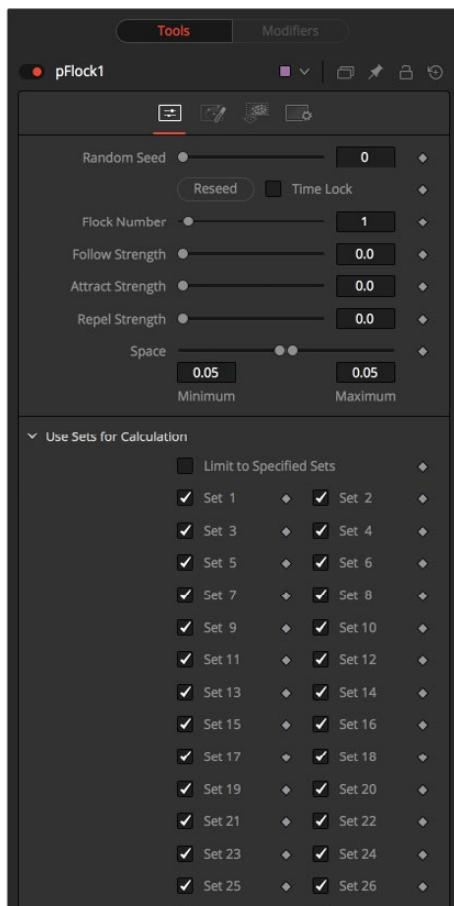


Flocking is a mechanism that can be used to simulate the behavior of organic systems, such as a flock of birds or a colony of ants. Its use can make an otherwise mindless particle system appear to be motivated, or acting under the direction of intelligence.

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.

Controls



Randomize

The 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 on 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 to be in the particle system and the larger the groups of particles will appear to be.

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.

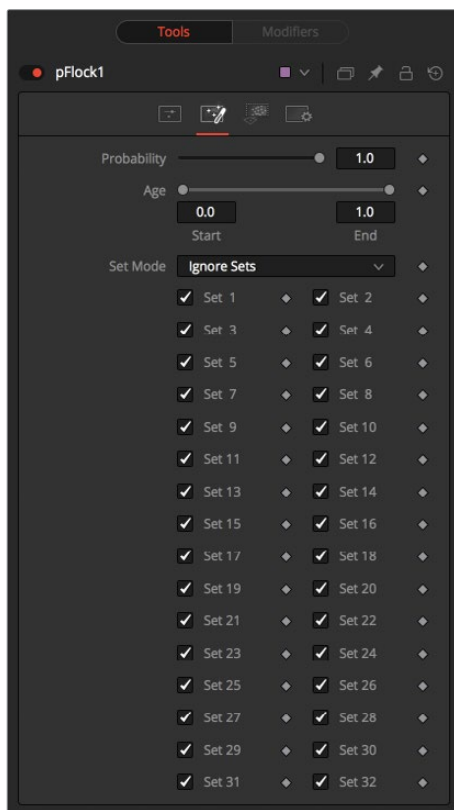
Repel Strength

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.

Conditions Tab




Regions Tab




See Particle Common Controls in this chapter.

pFriction



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.

Controls



Randomize

The 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 on 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.

Conditions Tab



Regions Tab



See Particle Common Controls in this chapter.


pGradientForce



The Gradient Force control accepts two inputs, one from a particle system and one from a bitmap image. 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).

This node can be used to give particles the appearance of moving downhill, or of following the contour of a provided shape.

Controls



Randomize

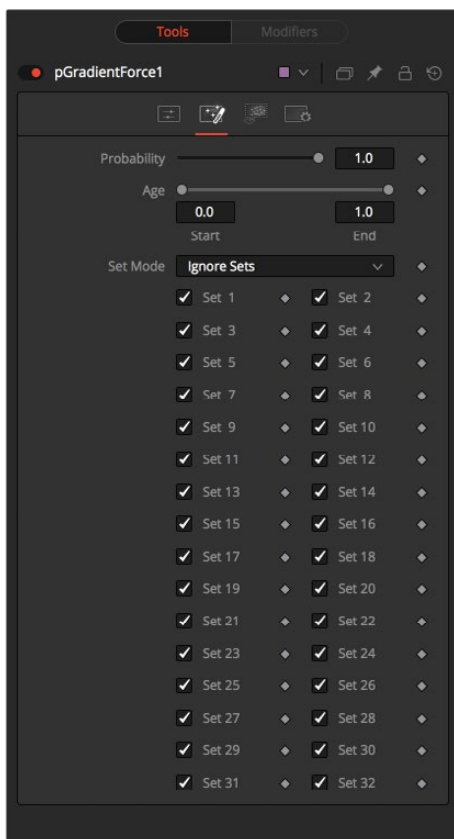
The Seed slider and Randomize button are presented whenever a Fusion node relies on a random result.

Two nodes with the same seed values will produce the same random results. Click on 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, it 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).

Conditions Tab




Regions Tab



See Particle Common Controls in this chapter.

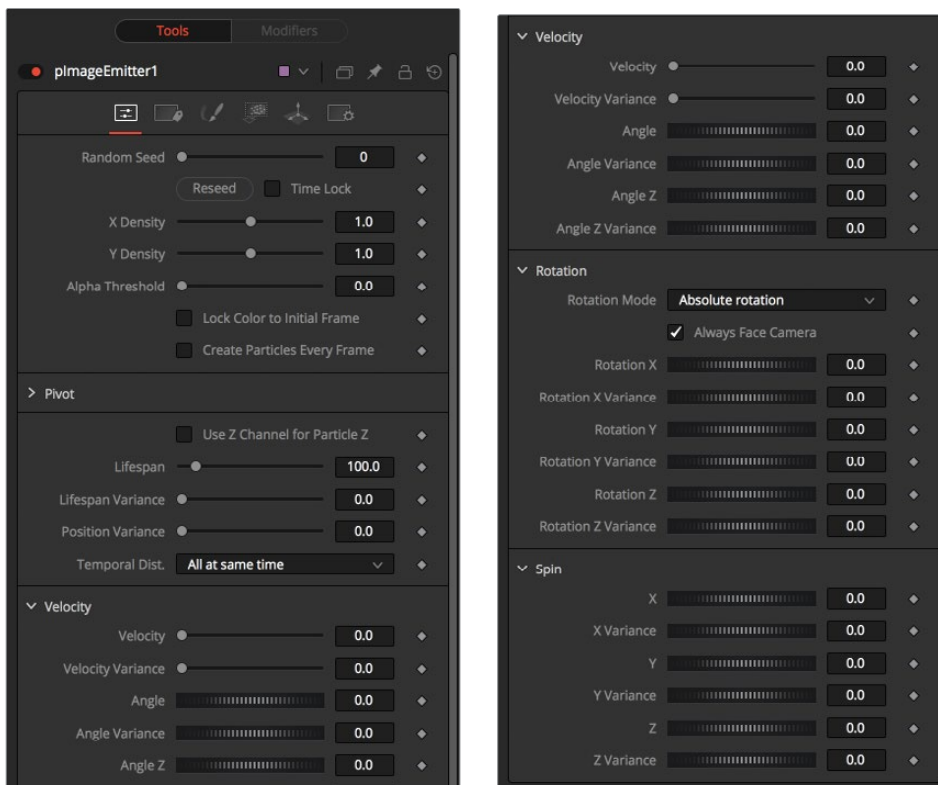
plmageEmitter



The plmageEmitter node takes an input image and treats each pixel of the image as if it were a particle. The main differences between the plmageEmitter 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.

Controls

The great majority of controls in this node are identical to those found in the pEmitter, and those controls are documented there. Below are the descriptions of the controls unique to the plmageEmitter 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 semi-transparent 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 has to be before it will generate a particle. Note that the default threshold of 0.0 will create particles for every pixel, regardless of alpha, though many may be transparent and invisible.

Lock Particle Color to Initial Frame

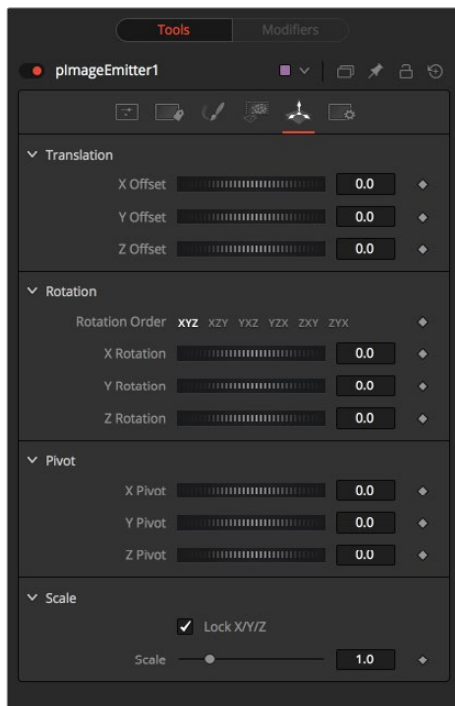
Select this checkbox to force the particles to keep the color with which they were born throughout the life of the particle. If this is off, and the input image changes on successive frames, the particles will also change color to match the image. This allows video playback on a grid of particles.

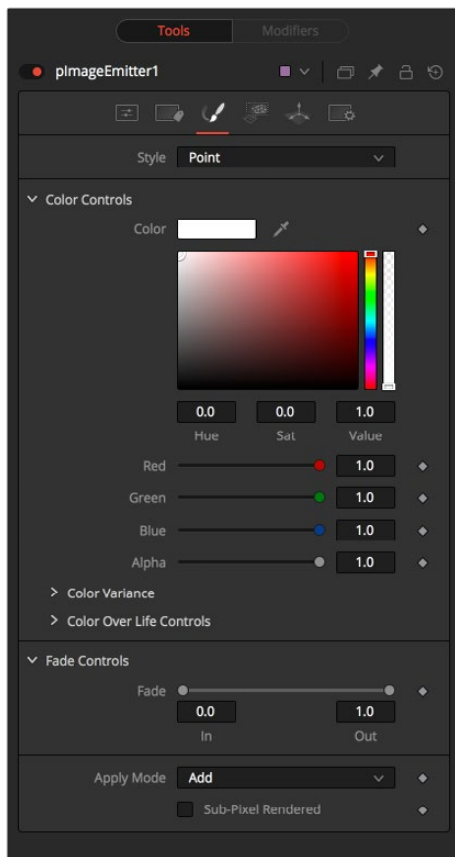
Create Particles Every Frame

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, 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

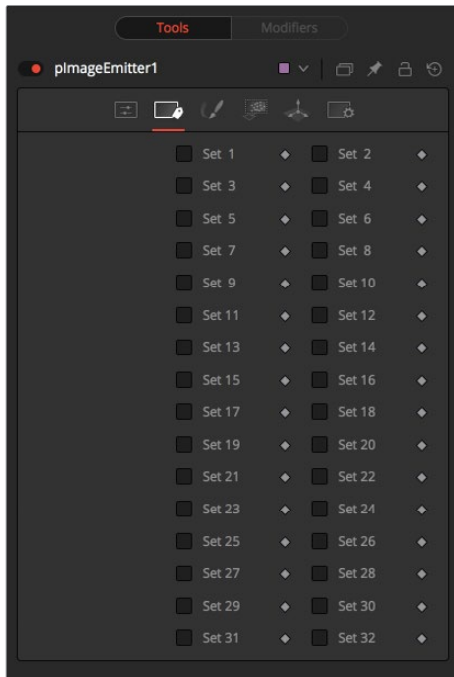
If the input image used to generate the particles has a Z depth channel, that channel can be used to determine the initial position of the particle in Z space. This can have an interesting hollow shell effect when used in conjunction with camera rotation in the pRender node.

Regions Tab



See Particle Common Controls in this chapter.

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 plImageEmitter (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.

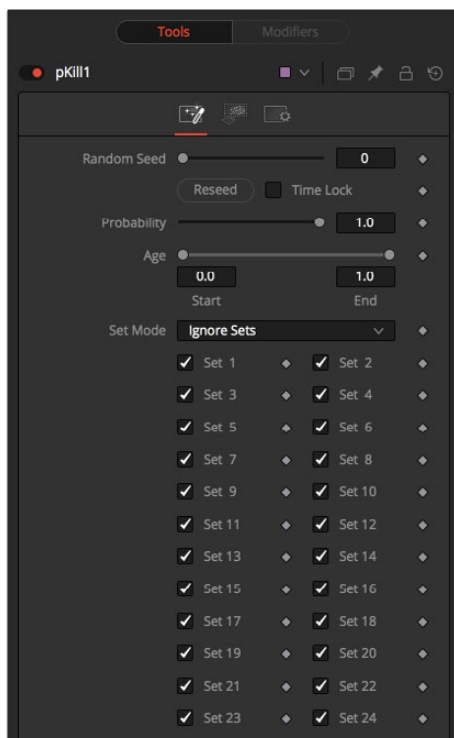
pKill




The great majority of controls in this node are identical to those found in the pEmitter, and those controls are documented there. Below are the descriptions of the controls unique to the pImageEmitter node.

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 which fall within a certain region, age, set, or by reducing the probability of the node applying to a given particle.

Conditions Tab



Regions Tab



See Particle Common Controls in this chapter.


pMerge



This node has no controls whatsoever. It serves to combine particles from two streams. Any nodes downstream of the Particle Merge node will treat the two streams as one.

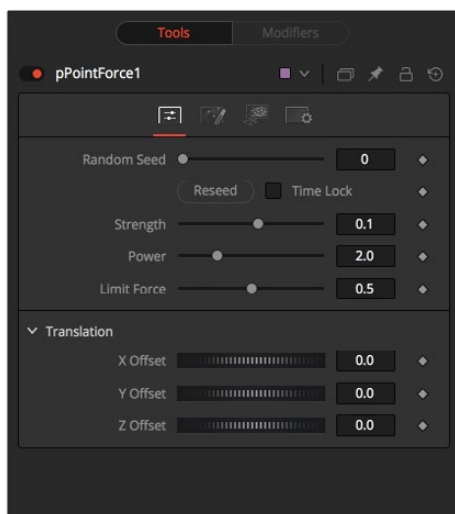
The combined particles will preserve any sets assigned to them when they were created, making it possible for nodes downstream of the pMerge to isolate specific particles when necessary.

pPointForce



This node applies a force to the particles that emanates from a single point in 3D space. The pPointForce can either attract or repel particles within its sphere of influence. There are four controls specific to the pPointForce node.

Controls



Randomize

The 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 on 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 there to be no falloff of strength. Higher values will impose an ever-sharper falloff in strength of the force with distance.

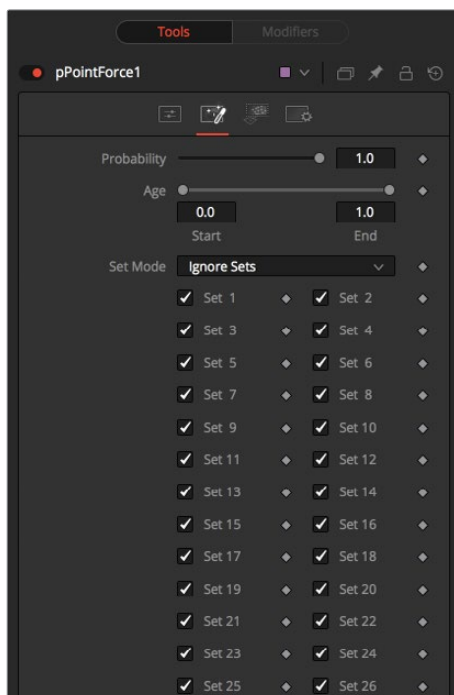
Limit Force

The Limit Force control is used to counter-balance potential problems with temporal sub-sampling. Because the position of a particle is only sampled 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.

Conditions Tab




Regions Tab



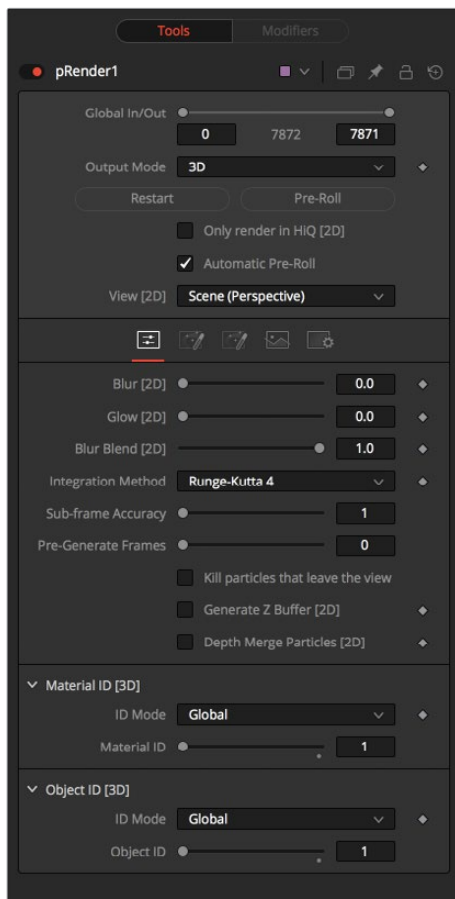
See Particle Common Controls in this chapter.

pRender [PRN]



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.

Controls



Output Mode (2D/3D)

While the pRender defaults to 3D output, it can be made to directly render a 2D image instead, as with DF4. 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 display view's context menu.

In 3D mode, the only controls in the pRender node that have any affect 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 connection of a 3D 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 Display Views

When the pRender node is selected in a node tree, all of the onscreen controls from Particle nodes connected to it are presented in the display views. 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 affect 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 will cause the particle system to recalculate, starting from the beginning of the render range up to the current frame. It will not render the image produced. It will only calculate 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 display view, 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 views, to prevent distracting visual disruptions.

About Pre-Roll

Pre-Roll is necessary because the state of a particle system is completely 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, add a pEmitter and pRender node to the composition. View the pEmitter in one of the display views. Right-click in the display view and select Views > 2D Viewer from the context menu.

Set the Velocity of the particles to 0.1. Place the pEmitter on the left edge of the screen and set the Current Frame to 0. Set a Render Range from 0–100 and press the Play button. Observe how the particle system behaves.

Stop the playback and return the current time to frame 0. Make sure that the Automatic Pre-Roll option is off in the pRender node. Now try jumping from frame 0 to 10, then to frame 60 and 90. Do not pass through the frames in between. Use the current time edit control or click in the

ruler directly to jump straight to the frame.

See 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 Pre-Render 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 over-ridden when the Hi-Q checkbox is de-selected, 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 display view 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.

Output Mode

Blur, Glow and Blur Blend

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 together. The end result is no different than adding a Blur after the pRender node in the node tree. For 2D particles only.

Sub Frame Calculation Accuracy

This determines the number of sub-samples taken between frames when calculating the particle system. Higher values will increase the accuracy of the calculation but also increase the amount of time taken 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. Pre-Generate Frames would be set 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 will automatically destroy any particles that leave the visible boundaries of the image. This can help to speed render times. Particles destroyed in this fashion will never return, regardless of any external forces acting upon them.

Generate Z Buffer

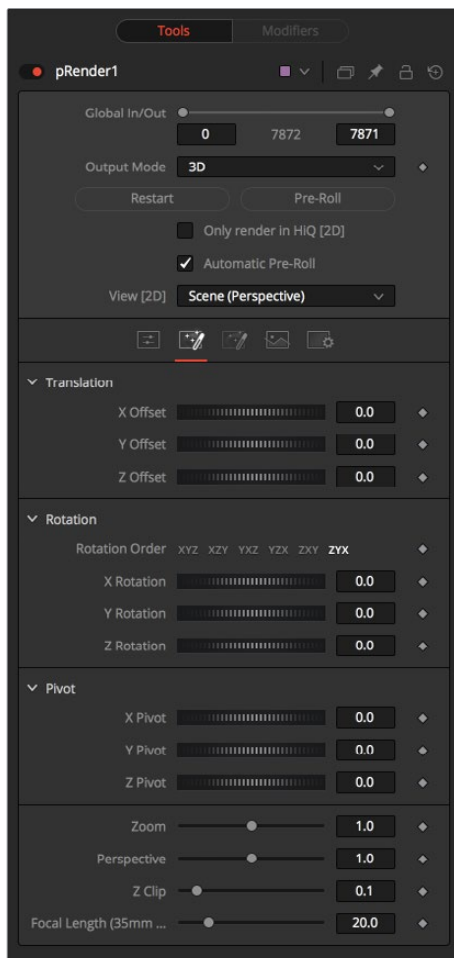
Selecting this checkbox will cause the pRender node to produce a Z Buffer channel in the image. The depth of each particle is represented in the Z Buffer. This channel can then be used for additional depth operations like Depth Blur, Depth Fog, and Downstream Z Merging.

Enabling this option is likely to increase the render times for the particle system dramatically.

Depth Merge Particles

Enabling this option will cause the particles to be merged together using Depth Merge techniques, rather than layer-based techniques.

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

None of the controls in this tab have any effect on 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.


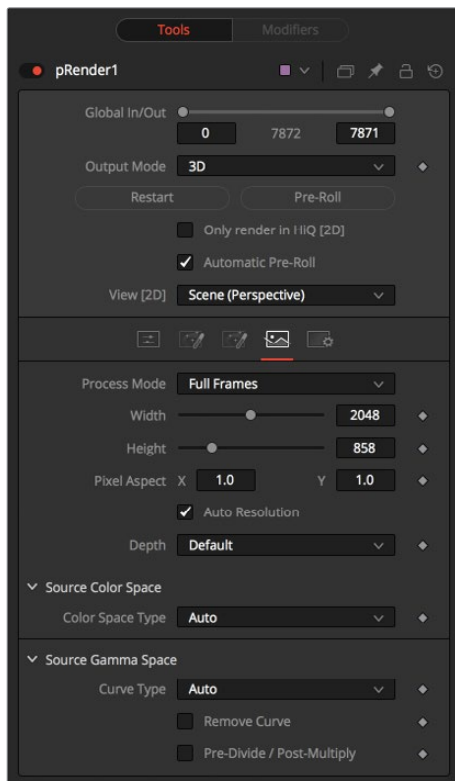


Image Tab

The controls in the Image tab of this node determine the width, height and aspect of the image. Refer to the Creator Common Controls page for details on the use of these controls.




Motion Blur

As with other 2D nodes in Fusion, Motion Blur is enabled from within the Fusion 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]

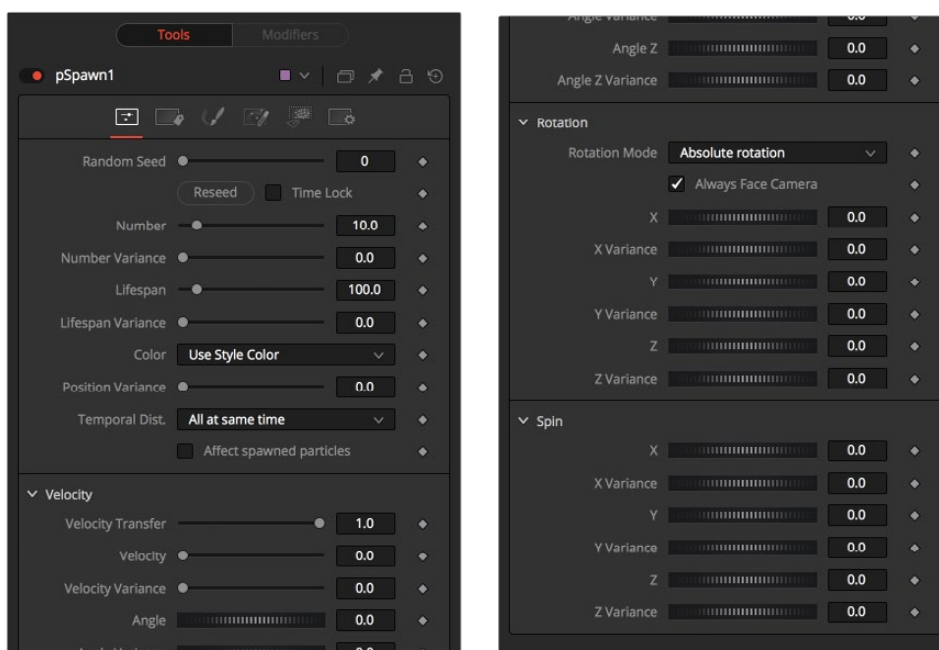


The pSpawn node makes each effected particle act as an emitter that can produce one or more particles of its own. The original particle continues until the end of its own lifespan, and each of the particles it emits becomes wholly independent with a lifespan and properties of its own.

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.

Controls

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.




Affect Spawned Particles

Selecting this checkbox causes particles created by spawning to also become affected by the Spawn 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.

Regions Tab



See Particle Common Controls in this chapter.

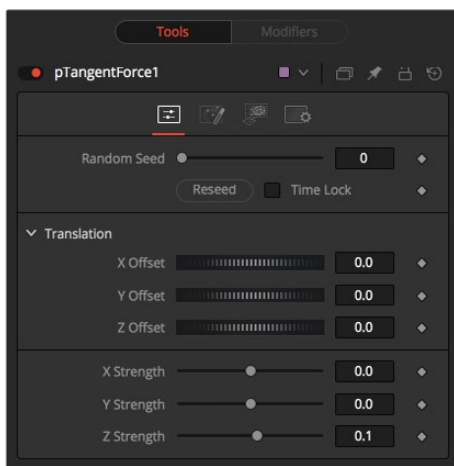
pTangentForce [PTF]



This node is used to apply a tangential force to the particles, a force that is applied perpendicularly to the vector between the pTangentForce's region and the particle it is affecting.

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.

Controls



Randomize

The 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 on the Randomize button to randomly select a new seed value, or adjust the slider to manually select a new seed value.

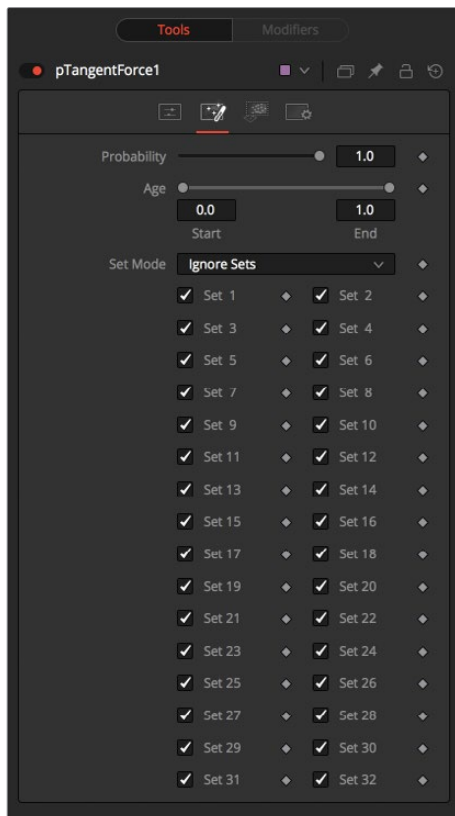
X, Y, Z Center Position

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.

Conditions Tab




Regions Tab



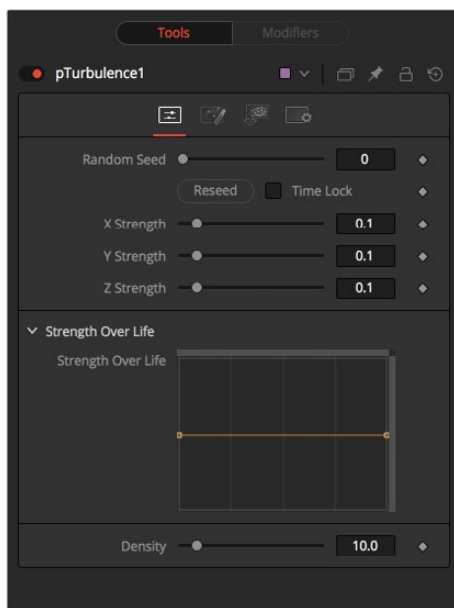
See Particle Common Controls in this chapter.

pTurbulence [PTR]



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.

Controls



Randomize

The 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 on 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.

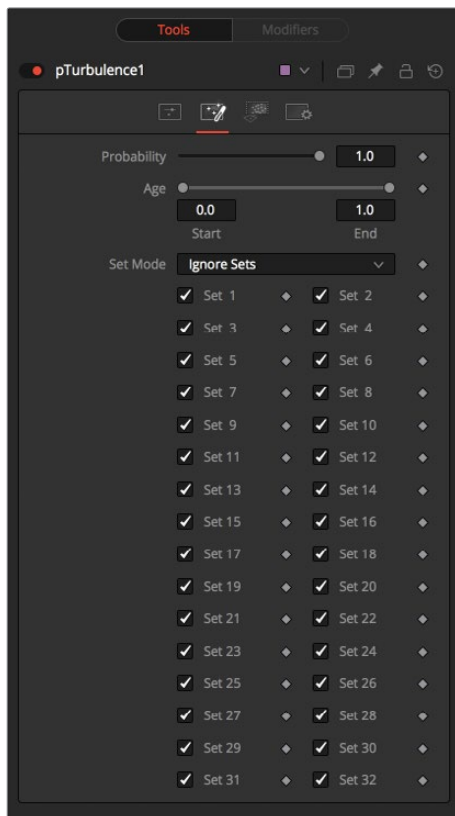
Strength Over Life

This LUT 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 of the turbulence. Higher values will produce finer variations in the turbulence produced.

Conditions Tab



Regions Tab



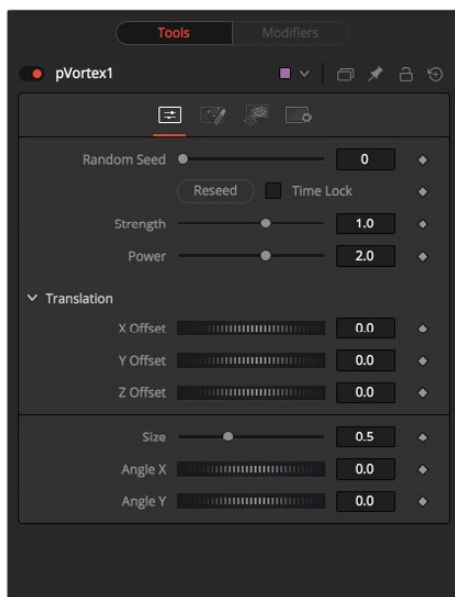
See Particle Common Controls in this chapter.

pVortex [PVT]



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.

Controls



Randomize

The 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 on 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.

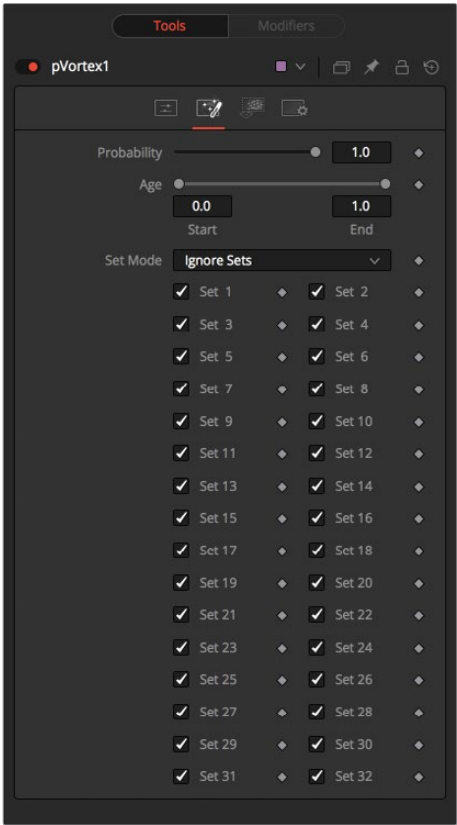
Size

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 axis.

Conditions Tab



Regions Tab

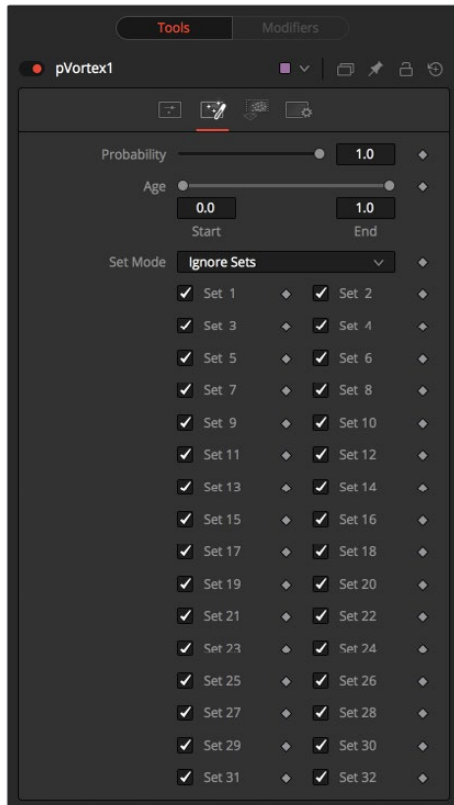


See Particle Common Controls in this chapter.

Particle Common Controls

The Common Controls can be found in every Particle node inside Fusion.

Conditions Tab



Probability

This control determines the Probability or percentage chance that the node will affect any given particle.

The default value of 1.0 will affect 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 affect a particle with a lifespan of 100 frames.

Set Mode and Sets

This menu control determines if the Particle node's effect will be applied to all particles, limited only to particles from a specified set(s), or applied to all particles except the selected set(s).

Sets are assigned by the nodes that create particles the pEmitter, plmage Emitter, the pChangeStyle and the pSpawn nodes.

Regions Tab



Region

Regions are used to restrict the node's effect to a geometric region or plane, and to determine the area where particles are created in a pEmitter node. There are seven types of regions, each with its own 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.

Bezier

Bezier mode uses a user-created polyline to determine the region where particles are created. The Bezier mode works in both 2D and 3D modes, however the Bezier polyline region can only be created in 2D.

To animate the shape of the polyline over time or to connect this polyline to another polyline, right-click on the Polyline label at the bottom of the control and select the appropriate option from the context 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 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. As with Bezier, this region type works in 3D, but the line itself can only be created and adjusted in 2D.

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 deeper 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 (3D)

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.


Start/End X,Y,Z Offset

(Line mode only) These controls define the Start and End point of the line in 3D Space.

Mesh Regions

Region Type

The Region Type control allows the artist to choose whether the region will include the inner volume, or just the surface. For example, when a mesh region is used with a pEmitter, this control would determine whether the particles are emitted from the surface or the full volume.




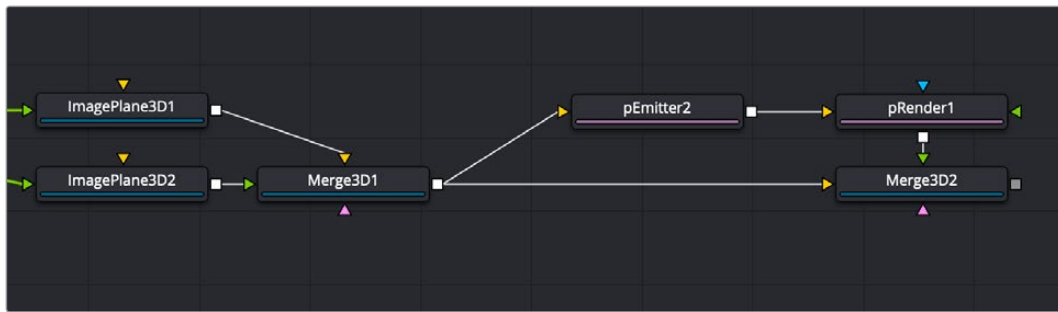
Winding Rule and Winding Ray Direction

The Winding Rule and Winding Ray Direction controls determine how the mesh region will handle particle creation with meshes that are not completely closed, as is common in many meshes imported from external applications. This scenario is common with imported mesh geometry, and even geometry which 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 shows 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, ObjectID

Selecting this checkbox allows the artist to use the slider found below the checkbox to filter the scene geometry so that only geometry with a matching ObjectID property is used as part of the region.

This menu control determines if the Particle node's effect will be applied to all particles, limited only to particles from a specified set(s), or applied to all particles except the selected set(s).

Sets are assigned by the nodes that actually create particles. These are the pEmitter, plmage Emitter, the pChangeStyle and the pSpawn nodes.

pStyle

The Styles tab occurs in the Particle Emitter, Particle Spawn, Particle ChangeStyle and Particle Image Emitter. In the Style tab, the type and look of the particles is determined.

Style Tab

The Style tab provides controls that affect the appearance of the particles, allowing the look and feel of the particles to be determined and animated over time.

Style Type

The Style Menu control provides access to the various types of particles supported by the Particle Suite. Each style has its own specific controls, as well as controls it will share with other styles.

Point Style

This option produces particles exactly 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 controls with the Point style. Additional controls specific to Point Cluster style are Number of Points and Number Variance.

Color

This standard Color Control selects 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 wider range of colors.

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 extremely useful for fire-type effects (for example, the flame may start out 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 context 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.

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 exactly 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 further adjust 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 ($\text{velocity} * \text{size} + \text{velocity} = \text{newsize}$). 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 is increased or decreased according to its depth (position in **Z** space). 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 LUT 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 of the features available to a standard LUT 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 view.

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 exactly as documented in the standard Merge node. The Burn-In control will cause the particles to overexpose, or "blow out," when they are combined together.

None of the Merge controls will have any effect on a 3D particle system.

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 and Blur Variance

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 LUT 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 view.

Z Blur (DoF) 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.

Apply Mode

This control only applies to 2D particles. 3D particle systems are not affected.

Add

Overlapping particles are combined by adding the color values of each particle together.

Merge

Overlapping particles are merged together.

Style Bitmap

(Bitmap Style Only) 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 Only) This drop-down list 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.

Animate

(Bitmap Style Only) This list determines how the animation of the bitmap source is applied to newly-created particles. It can be hard to visualize the impact of this control until it has been experimented with. A good way to experiment is to load the node tree in the Fusion > Examples directory called pTextParticleAge.comp and try out the three settings. View the Text node in the display view and the pRender node another view, then step through the frames using the [and] keys.

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 exactly 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 Only) This control allows the bitmap source frame to be Offset in time from the current frame.

Time Scale

(Bitmap Style Only) This control can be used to scale 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 Only) This control is used to apply a correction to the overall Gain of the image that is used as the bitmap. Higher values produce a brighter image, whereas lower values reduce both the brightness and the transparency of the image.

Noise

(Blob Style Only) Increasing this control's value will introduce grain-type noise to the blobby particle.

Fade

(Line Style Only) 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.

Sub Pixel Rendered

(Point and Point Cluster Style Only) This checkbox determines if 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 Style Only) The value of this control determines how many points are in each Point Cluster.

Chapter 50

Position Nodes

This chapter details the Position nodes available in Fusion.

Contents

Volume Fog [VLF]	1066
Volume Mask [VLM]	1072
Z to WorldPos [Z2W]	1075
WPP Concept	1076

Volume Fog [VLF]



The Volume Fog node is used to create sophisticated volumetric fog on images containing XYZ Position channels.

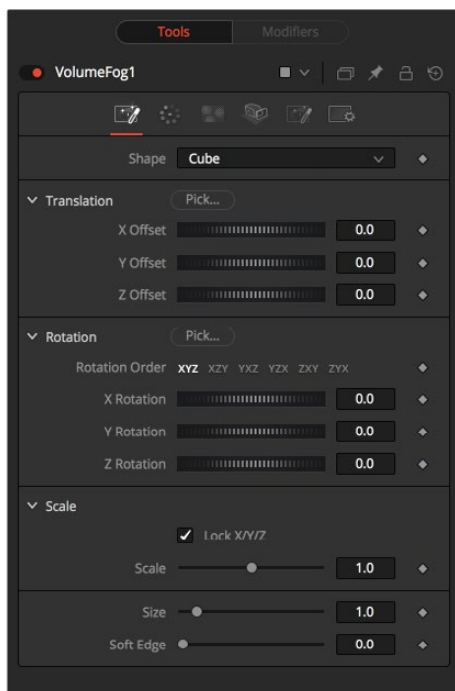
As opposed to 3D-rendered volumetric Fog, it works on 2D images and delivers much faster results and interactive feedback when setting up the fog. See the “WPP Concept” section for further explanation on how this technology works and to learn about the required imagery.

External Inputs

The following inputs appear on the node’s tile in the Node Editor.

- **VolumeFog.Image:** [orange, required] This input expects an image containing a World Position Pass in the XYZ Position channels.
- **VolumeFog.FogImage:** [green, optional] For creating volumetric fog with varying depth and extent; an image sequence can be connected here.
- **VolumeFog.EffectMask:** [blue, optional] The standard Effect Mask input as found on other Fusion nodes.
- **VolumeFog.SceneInput:** [pink, optional] This input expects a 3D scene containing a 3D Camera.

Shape Tab



Shape

Toggle between a basic spherical or rectangular volume to be placed in your image. These Volumes can then be further refined by means of the MaskImage or MaskImage sequence.

Translation Pick

Left-click and hold this button to pick 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

Left-click and hold this button to pick rotational values from any 3D Scene or 2D image containing those values, like an XYZ-Normal-Pass, to re-orient the 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 Volume around its center.

X, Y, Z Scale

Scale the Volume in any direction from its center to further refine the overall Size value specified below.


Size

The overall Size of the volume created.

Soft Edge

Controls how much the Volume is faded toward the center from its perimeter to achieve a softer look.

Color Tab



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 this value, the more images from the connected MaskImage sequence will be used to form the Depth of the Volume.

You can, for example, use a FastNoise with a high Seethe Rate to create such a sequence of images. Be careful with the resolution of the images though. 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 256MB 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 MaskImage sequence.

Make sure that both GlobalIn and GlobalOut as well as the valid range of your source node fall within the range of FirstSliceTime + Z Slices.

Color

Allows you to modify the color of the Fog generated. This will multiply over any color provided by the connected MaskImage.

Gain

Increases or decreases intensity of the fog created. 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 if the fog is composed onto the image in additive or subtractive mode, leading to brighter or dimmer appearance of the fog.

Fog Only

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.

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 allows the user to choose the correct Camera needed to evaluate the Volume. Instead of connecting a Camera, position values can be provided manually or by connecting the XYZ values to other controls.

Translation Pick

Left-click and hold this button to pick 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.

Light Tab

To utilize this feature you need to have actual lights in your 3D Scene. Connect that scene, including Camera and Lights, to the 3D input of the node.

Do Lighting

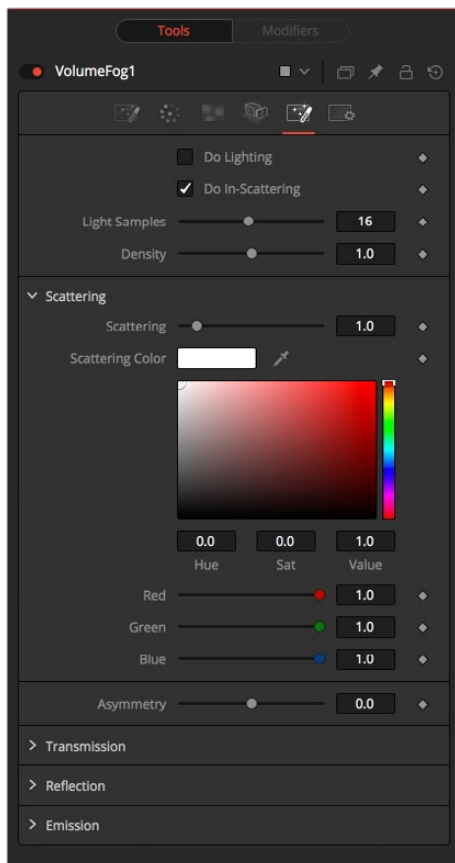
Enables or disables lighting calculations. Keep in mind that when not using GPU (i.e., rendering on the CPU), these calculations may become a bit slow.

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 higher 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 on 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. Light that doesn't get scattered out will tend toward this color. It is a multiplier, though, so if you have red light, but blue transmission, you won't see any blue.

Reflection


Changes the intensity of the light that is scattered out. It 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 was 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, or adding energy/light back into the calculation. In fact, 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 actually happen in the real world.

Examples

In these examples we are looking at a volume from the outside. On the right you see how the VolumeFog 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. On the right the actual light in the scene has been moved, which also influences the look of the volume.

Volume Mask [VLM]



The Volume Mask node is used to create volumetric masks from images containing XYZ Position channels.

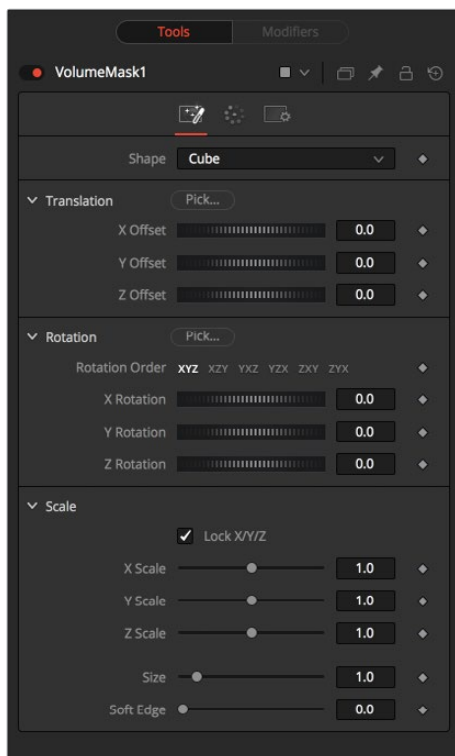
This can, for example, be used to isolate objects for color correction without the need to track or roto the scene. See the “WPP Concept” section for further explanation on how this technology works and to learn about the required imagery.

External Inputs

The following inputs appear on the node tile in the Node Editor.

- **VolumeFog.Image:** [orange, required] This input expects an image containing a World Position Pass in the XYZ Position channels.
- **VolumeFog.MaskImage:** [green, optional] For refining the mask, an image can be connected here.
- **VolumeFog.EffectMask:** [blue, optional] The standard Effect Mask input as found on other Fusion nodes.
- **VolumeFog.SceneInput:** [pink, optional] This input expects a 3D scene input containing a 3D Camera.

Shape Tab



Shape

Toggle between a basic spherical or rectangular Mask to be placed in your image. The Mask can be further refined by means of the MaskImage.

Translation Pick

Left-click and hold this button to pick XYZ coordinates from any 3D scene or 2D image containing XYZ values, like a rendered World Pass, to position the center of the volumetric 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

Left-click and hold this button to pick rotational values from any 3D scene or 2D image containing those values, like an XYZ-Normal-Pass, to re-orient the Mask.

When picking from a 2D image, like an XYZ-Normal-pass, make sure it's rendered in 32-bit float and WorldSpace 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.

Color Tab



Color

Allows user to modify the color of the Mask generated. This will add to any color provided by the connected MaskImage.

Subtractive/Additive Slider


Similar to the Merge node, this value controls if the Mask is composed onto the image in Additive or Subtractive mode, leading to 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.

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 allows the user 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

Left-click and hold this button to pick 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 32bit 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.

Z to WorldPos [Z2W]



The Z to WorldPos 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 VolumeMask and Volume Fog could be applied to real-world scenes. The quality of the resulting WPP depends largely 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.

External Inputs

The following inputs appear on the node tile in the Node Editor.

- **ZtoWorld.Image:** [orange, required] This input expects an image containing a World Position Pass or a Z-depth pass, depending on the desired operation.
- **ZtoWorld.EffectMask:** [blue, optional] The standard Effect Mask input as found on other Fusion nodes.
- **ZtoWorld.SceneInput:** [pink, required] This input expects a 3D scene input containing a 3D Camera.

Controls



Mode

Toggle 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 allows the user to choose the correct Camera needed to evaluate the image.

WPP Concept


The Position nodes in Fusion offer a completely new way of working with Masks and Volumetrics for footage containing XYZ Position channels. ZtoWorld offers the option to create those channels out of a Z-channel and 3D Camera information. Just for the sake of this overview we will refer to the World Position Pass as WPP.

What is a WPP?

The WPP reflects each pixel's XYZ position in the original scene as RGB color values.


So if the face the pixel was derived from in the original scene sits at 0/0/0, the resulting pixel will have an RGB value of 0/0/0 and thus will be black. If said face sits at 1/0/0 in the original scene, the resulting pixel will be fully red. Of course if the coordinates of our face are something like -60/75/123, the pixel will hold those values as well.

As you can see here, 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. This image shows a 3D Rendering of a scene with its center sitting at 0/0/0 in 3D Space and the according WPP. For better visibility the WPP has been normalized.



Different Coordinate Spaces


WPPs can be rendered in different Coordinate Spaces. These include World Space, Eye Space and Object Space. The image below depicts how those different spaces look, though the nodes in Fusion require the WPP to be rendered 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 ZtoWorld node, Volume Mask and Volume Fog will still 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, greatly improves 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 said 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 will be 0/0/0; 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 “with Sphere” below.



Chapter 51

Stereo Nodes

This chapter details the Stereo nodes available in Fusion.

Contents

Anaglyph [ANA]	1080
Combiner [COM]	1083
Disparity [DIS]	1084
Disparity To Z [D2Z]	1087
Global Align [GA]	1090
New Eye [NE]	1091
Splitter [SPL]	1094
Stereo Align [SA]	1094
Z To Disparity [Z2D]	1098


Anaglyph [ANA]



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.


Controls

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. Based on one of the methods described below, the separate images are then combined to create a stereoscopic output




Color Type Menu

The ColorType 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 turns out to be the most commonly used.



Method Menu

In addition to the color used for encoding the image, you can also choose five different methods: Monochrome, Half-color, Color, Optimized and Dubois.




Monochrome


The left eye contains the luminance of the left image, and the right eye contains the luminance of the right image (e.g., the luminance of the left image is placed in the output red channel, and the luminance of the right image is placed in the output green and blue channels).

Half-Color

The left eye contains the luminance of the left image, and the right eye contains the color channels from the right image that match the glasses color for that eye (e.g., the luminance of the left image is placed in the output red channel, and the green and blue channels of the right image are placed in the output green and blue channels).




Monochrome




Half-Color

Color

The left eye contains the color channels from the left image that match the glasses color for that eye, and the right eye contains the color channels from the right image that match the glasses color for that eye (e.g., the red channel of the left image is placed in the output red channel, and the green and blue channels of the right image are placed in the output green and blue channels).



Color




Optimized

Optimized

Used with red/cyan glasses, for example, the resulting brightness of what shows through the left eye will be 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's 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 a 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 (e.g., $1.05\times$ the green and $0.45\times$ the blue channels of the left image is placed in the output red 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.).

Dubois

Images with fairly saturated colors can produce retinal rivalry with Half-color, Color and Optimized methods because the color is only visible in 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, and also tends to reduce ghosting produced when one eye 'leaks' into the other eye. The particular calculated matrix we use 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 the user to easily swap the left and right eye input.

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.


Combiner [COM]



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.

Controls

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.



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.

Swap Eyes

Allows the user to easily swap the left and right eye input.

Disparity [DIS]



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 upon matching regions in the left eye to regions in the right eye by comparing colors and gradients of colors, it is important that colors in the two eyes are 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 will confuse the disparity tracking (and also cause problems if you are using ColorCorrector's histogram match ability to do the color matching).

In Stack mode, left and right outputs will output the same image. If the left and right images have a global vertical offset bigger than a few pixels from each other, it can help the disparity tracking algorithm if you approximately 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 will become a combined Disparity and Lens Distortion map. This can have advantages and disadvantages.

One disadvantage is 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/outputs

Left Input

Connect either the left eye image or the stacked image.

Right Input

Connect the right eye image. This input is invisible unless Stack Mode is set to Separate.

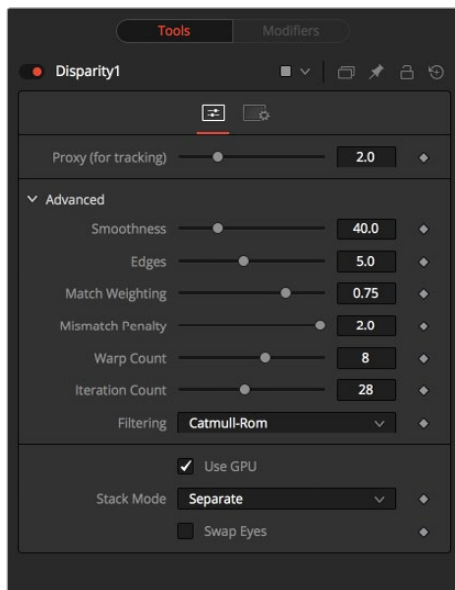
Left Output

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 only visible if Stack Mode is set to Separate.

Controls



Proxy (for tracking)

The input images are re-sized 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 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 will give a 4x speed up and a proxy scale of 3 will give a 9x speed up. 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, which if present in large quantities can completely obliterate any gains moving from 2:1 to 1:1 proxy and in some situations even make things worse (in some sense you can think of the Proxy setting as acting as a simplistic low-pass filter for removing noise/grain).

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 need to be connected.

Swap Eyes

With Stacked Mode, image stereo pairs, left and right images can be swapped.

Advanced

The Advanced Control 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 will not be needed.

Smoothness

This controls the Smoothness of the disparity. Higher smoothness helps deal with noise, while lower smoothness brings out more detail.

Edges

This slider is another control for smoothness but applies smoothing based upon 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 Loose, the disparity becomes smoother and tends to overshoot edges. When it is set to Tight, 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, then you can experiment with making it tighter, but if you are using the disparity to do interpolation, you might want to keep it looser.

In general, if it is too tight, there can be problems with streaked out edges when the disparity is used for interpolation.

Match Weight

This controls how the matching of neighborhoods in the left to neighborhoods in the right image is done. When set to Match Color, large structural color features are matched. When set to Match Edges, small sharp variations in the color are matched. Typically, a good value for this slider is in the [0.7, 0.9] range, although on some shots values closer to 0.0 work well. 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 do a color match on 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 gives the 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.

Number of Warps

Turning this option down 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.


Number of Iterations

Turning this option down makes the disparity computations faster. In particular, the computational time depends linearly upon on this option. Just like adjusting Number of Warps, at some point adjusting this option higher will give 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 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.

Disparity To Z [D2Z]



Description

DisparityToZ 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 Camera3D or a tracked stereo camera is connected into DisparityToZ, 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.

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/outputs

Left Input

Connect either the left eye image or the stack image.

Right Input

Connect the right eye image. This input is invisible unless Stack Mode is set to Separate.

Stereocamera

An external stereo camera tool.


Left Output

Holds the left eye image with a new Z channel or a stacked mode image with a new Z channel.

Right Output

Holds the right eye image with a new Z channel. This output is only visible if Stack Mode is Separate.

Controls



In addition to outputting Z values in the Z channel, this option 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 to see what the Z channel looks like. Note that Dimension's Z values are negative, getting more negative the further you are from the camera, so you will need view Normalization On to see them correctly.

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

This option refines the depth map based upon the RGB channels. The refinement causes edges in the flow to align more closely to 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

This option only processes in HiQ.

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.

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 you just want any Z-buffer and are not that particular about the exact details on how it is offset and scaled, or if there is no camera available, the Artistic option might be helpful.



External Mode

An input will appear on the node tree to connect an existing stereo Camera3D. This can either be a single stereo Camera3D (i.e., its eye separation is set to non-zero), or a pair of (tracked) Camera3Ds connected together via the Camera3D > Stereo > RightCamera 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 still useful. To reconstruct the Disparity > Z Curve, pick (D, Z) values for a point in the foreground and a point in the background

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), you can always create a bogus Camera3D, connect it into the DisparityToZ > Camera input, and then fiddle with the Camera3D’s controls.

Foreground Disparity (pick from left eye)

This is the disparity for the closest foreground object. It will get mapped to 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)

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 $\text{depth} = \text{constant}/\text{disparity}$. When set to Linear, the curve behaves like $\text{depth} = \text{constant} * \text{disparity}$. 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 not to, prefer Hyperbolic, as it is more physically accurate, while Linear does not correspond to nature and is purely for artistic effect.

Global Align [GA]

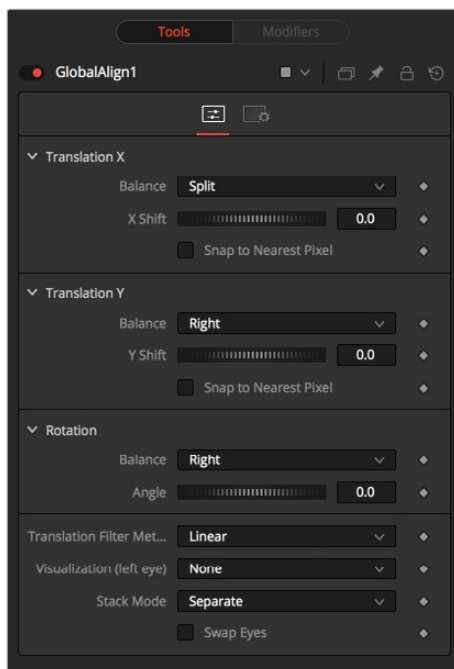


As opposed to Stereo Align, this node does not utilize optical flow at all. It's meant as a fast and convenient way to do simple stereo alignment for both X and Y as well as rotation.

Global Align comes in handy at the beginning of the node chain to visually correct major differences between 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 to generate more accurate results.

Controls



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 translation slider, 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

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

A drop-down to choose a 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 need to be connected.

Swap Eyes

With Stacked Mode, image stereo pairs, left and right images can be swapped.

New Eye [NE]



This option constructs a New Eye by interpolating between two existing eyes 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.

NewEye does not interpolate the aux channels, but rather destroys them. In particular, the Disparity channels are consumed/destroyed. Add another Disparity node after the NewEye if you want to generate Disparity for the realigned footage.

Inputs/outputs

Left Input

Connect either the left eye image or the stacked image.

Right Input

Connect the right eye image. This input is invisible unless Stack Mode is set to Separate.

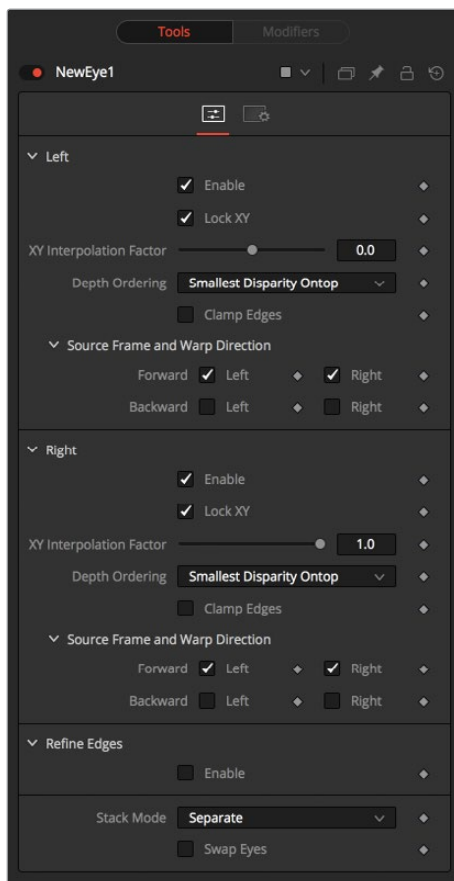
Left Output

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 only visible if Stack Mode is set to Separate.

Controls



Enable

When an eye is enabled, NewEye will replace it 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.

Interpolation

Interpolation determines where the frame we are interpolating is relative to the two source frames: Left and Right. An interpolation parameter of -1.0 will give frame Left back and a parameter of 1.0 will give frame Right back. A parameter of 0.0 will give 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 should be drawn on top.

- **Largest Disparity Ontop:** The larger disparity values will be drawn on top in the overlapping image sections.
- **Smallest Disparity Ontop:** 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 only use clamp edges 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 only have 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 is 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:** This will take the Left frame and use the Forward Disparity to interpolate the new frame.
- **Right Forward:** This will take the Right frame and use the Forward Disparity to interpolate the new frame.
- **Left Backward:** This will take the Left frame and use the Back Disparity to interpolate the new frame.
- **Right Backward:** This will take the Right frame and use the Back Disparity to interpolate the new frame.

Splitter [SPL]



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.

Controls



None

No operation will take place. The output image on both outputs is identical to the input image.

Horiz

The tool 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 the user to easily swap the left and right eye output.

Stereo Align [SA]



This extremely versatile node for fixing Stereo issues can be used for doing any of the following or combinations of:

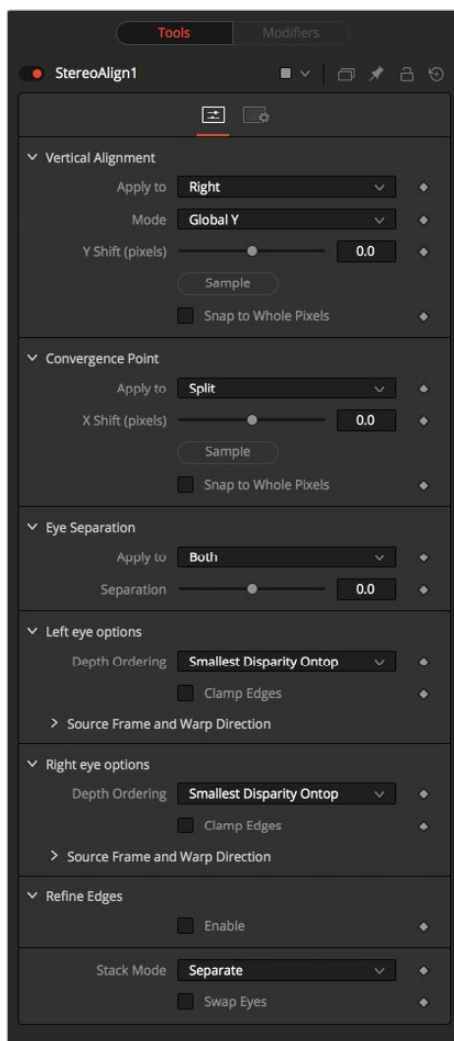
- Vertical Alignment of one eye to the other
- Changing the convergence
- Changing the eye separation

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.

Changing the eye separation can cause a lot of 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 have to fill the holes in manually. For now, this node will modify just the RGBA channels.

StereoAlign does not interpolate the aux channels, but rather 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.

Controls



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 note that you have to be careful about lens distortion because, even if two cameras are perfectly vertically aligned, they will still have vertical disparities due to lens distortion. It is probably a good idea to remove the lens distortion before computing the Disparity since, by doing a vertical alignment of the right eye, you are in effect removing the Y-component of the lens distortion in the right eye and it will look wrong later when you try to redistort it.

- **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 only available in Global Mode. You can either adjust the Y-shift manually to get a match or drag the pick button, which picks from the Disparity channel of the left eye. Also remember that, if you use this node to modify disparity, you can't pick from the node's output.

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.

Convergence

Convergence is just a global X-translation of L/R images.

Eyes

This determines which eyes are affected by convergence. In most cases, this will be set to both. If you set the eyes to Both/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 now rather than keeping one eye as a pure reference master.

X-shift

X-shift can be picked 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 apart depending on their distance from the camera.

This has the same effect as Eye Separation option in Camera3D.

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.

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.

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.

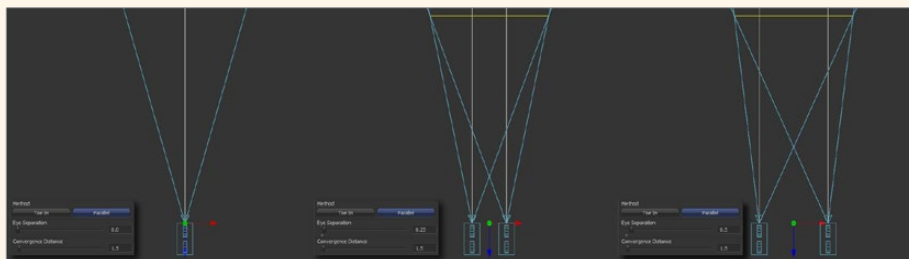
Because of these artifacts, it is a good idea to only use clamp edges to correct small gaps around the edges of an interpolated frame.

Softness

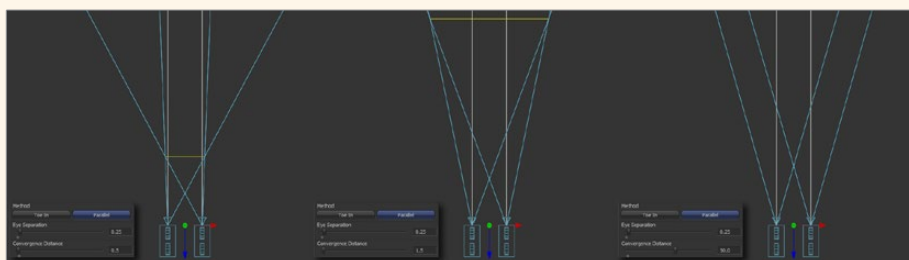
Helps to reduce the stretchy artifacts that might be introduced by Clamp Edges.

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 only have one checkbox enabled you can use a larger softness at around 0.03.

Example



Different settings for Eye Separation



and example settings for Convergence

Z To Disparity [Z2D]



ZToDisparity 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 generator node.

Inputs/outputs

Left

This is the left image or stack input.

Right

This is the right image.

Stereocamera

A stereo perspective camera may be either a Camera3D with eye separation or a tracked L/R Camera3D.

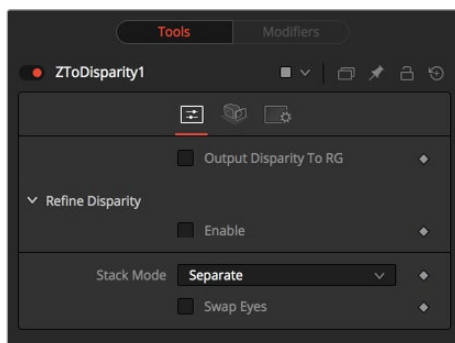
Left Output

This is the left image or stack output containing new Disparity channel.

Right Output

This is the right image that contains a new Disparity channel. This output will be hidden in Stack Mode.

Controls



Output Disparity To Rgb

In addition to outputting disparity values into the Disparity channel, this option causes ZToDisparity to also output the disparity values into the color channels as {x, y, 0, 1}.

When enabled, 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 upon 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.

Camera Tab

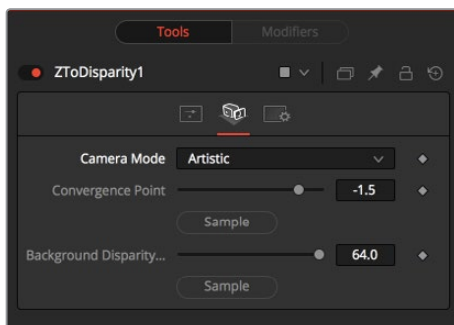
If you need correct real-world disparity values because you are trying to match some effect to an existing scene, you should use the External Camera options to get precise disparity values back. If you just want any disparity and do not particular care about the exact details on how it is offset and scaled, or if there is no camera available, then the Artistic option might be helpful.

External



An input will appear on the node tree to connect an existing stereo Camera3D, and use the Camera settings to determine the Disparity settings.

Artistic



If you do not have a camera, you can adjust these controls to produce an “artistic” Disparity channel whose values will not be physically correct but 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 Camera3D. At this distance, objects in the left and right eyes are at exactly the same position (i.e., have zero disparity).

Objects closer appear to pop out of the screen, and objects further appear behind the screen.

Background Disparity (pick 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.

Chapter 52

Tracker Nodes

This chapter details the Tracker nodes available in Fusion.

Contents

Tracker [TRA]	1102
Tracker Onscreen Controls	1102
Inspector Controls – Trackers Tab	1103
Inspector Controls – Operation Tab	1108
Options Tab	1111
Planar Tracker Node	1112
What the Planar Tracker Saves	1112
The Planar Node's Main Inputs	1112
A Typical Planar Track Workflow	1113
The Controls Tab in Track Mode	1113
Steady Mode	1117
Corner Pin Mode	1117
Stabilize Mode	1118
Options Tab	1120
Planar Transform Node	1121
Camera Tracker	1122
The Track Tab	1123
Camera Tab	1125
Solve	1127
Export	1131
Options	1134

Tracker [TRA]

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 Drip). Additionally, trackers can be used to stabilize an image, or to apply de-stabilization to one image based on the motion of another.

Please also refer to Chapter 57, “Using the Tracker Node” for more information.

Tracker Onscreen Controls

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 a tracker is selected in the node tree.



- The onscreen control is indicated by a red rectangle with a handle in the top left corner. This rectangle indicates the position of the pattern in the image. Every pixel within the rectangle is considered to be part of the pattern used for tracking. Resize the pattern by clicking and dragging on the rectangle's borders.



- Whenever the mouse moves over the pattern rectangle, a second rectangle with a dashed outline appears. This represents the search area, which determines how far away from the current pattern the Tracker looks on the next frame. This 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 will take to track, so try not to make the search area larger than necessary.



- Dragging on the handle repositions the pattern. While repositioning the pattern, a thumbnail window with an enlarged view of the pattern is displayed in the views, to assist with precise positioning of the pattern. This thumbnail disappears when the mouse button is released. The magnification ratio can be adjusted in the Options tab. If the selected Tracker has a custom name, the name of that Tracker will be displayed as a label at the bottom right of the search area rectangle.



There is no limit to the number of trackers that can be used in one Node Editor, or in the number of connections that can be made to a tracker. This chapter serves as a reference for the various controls in the Tracker, but we strongly suggest you read the more general information about using the Tracker in Chapter 57, “Using the Tracker Node.”

The Tracker can be employed in two forms: as a node in the Node Editor, or as a modifier attached to a control. When the Tracker is used as a node in the Node Editor, the image tracked comes from the input to the Tracker node. There is no limit to the number of patterns that can be tracked by a single Tracker node.

When the Tracker is used as a modifier, its controls appear in the Modifier tab for any node with a control connected to that modifier. Only one pattern can be tracked by a Tracker Modifier, but the image source can come from anywhere in the composition. This technique is generally used when tracking a quick position for a control.

Inspector Controls – Trackers Tab

The Trackers tab contains all of the controls you need for creating and positioning trackers, and for customizing and initiating tracking operations. A set of offset controls may be used after tracking has been accomplished to improve the alignment of an image being transformed by a track with the subject being tracked.

Track Buttons

There are four buttons used to initiate tracking and one in the middle used to stop a track in progress. These buttons can track the current pattern forward or backward in time. Hold the mouse pointer over each button for a moment, and a tooltip with the name of the button will appear.

The buttons operate as follows:

- **Track Reverse:** Clicking on this button will cause all Active trackers to begin tracking their patterns, 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 on this button will cause all Active trackers to begin tracking their patterns, starting at the current frame and moving backward through time until the beginning of the render range.
- **Stop Tracking:** Clicking on this button will Stop the tracking process immediately. This can also be achieved by pressing ESC. This button will only be active when tracking is in process.
- **Track Forward From Current Time:** Clicking on this button will cause all Active trackers to begin tracking their patterns, starting at the current frame and moving forward through time until the end of the render range.
- **Track Forward:** Clicking on this button will cause all Active trackers to begin tracking their patterns, 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 should be generated.

Frames Per Path Point

The value of this slider determines how often the Tracker sets a Keyframe on the Path. The normal default is 1, which sets a Keyframe on the tracked path at every frame.

Increasing the value of this slider will cause the tracked path to be less accurate. This may be desirable if the track is returning fluctuating results, but under normal circumstances this control should be left at its default value.

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 tracked pattern fields. For fielded footage tracked in Field mode, you will get better results setting this slider to a value of 2, which will result in one keyframe per frame of your footage.

Adaptive Mode

Fusion is capable of re-acquiring the tracked pattern, as needed, to help with complex tracks. This button array determines what mode of Adaptive tracking is employed.

If selected, Fusion searches only for the pattern originally selected in each single frame.

- **Every Frame:** If selected, Fusion re-acquires the pattern every frame. This helps the Tracker compensate for gradual changes in profile and lighting over time.
- **Best Match:** If selected, the Tracker will compare the pattern acquired at each frame and compare it to the original selected pattern. If the variation between the two patterns exceeds the threshold amount defined by the Match Tolerance control, Fusion will not re-acquire 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

The two buttons in this button array determine how the Tracker behaves when re-positioning a pattern. These controls are used when switching a path from one pattern to another, which happens when a pattern leaves the frame, or changes so significantly it can no longer be tracked.

- **Pattern Center:** When Pattern Center is the Active mode, the tracked path continues from the center of the new path. This is appropriate when replacing an existing path entirely, but when trying to append to a path using a new pattern, this will cause a discontinuity.
- **Track Center (append):** When Track Center (append) is selected, 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 technique will work best if the new pattern is located close to the position of the old pattern to avoid any problems with parallax or lens distortion. This mode can also be used to virtually continue the tracking of patterns that move out of the frame or get obstructed by other objects.

Tracker List Controls

Each tracker you create to use is managed in the Tracker List, with it's accompanying controls.

Add/Delete Tracker

Use these to add or delete Trackers from your Tracker List.

Tracker List


The Tracker List shows the names of all of the patterns created on this Tracker. It is also used to add new Trackers.

The Tracker node is capable of hosting a virtually unlimited number of Tracker patterns.

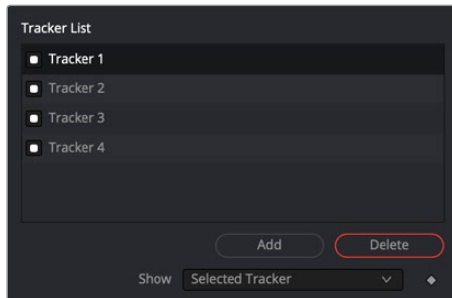
- Click on the Add button immediately above the list to add a new Tracker pattern.
- Each pattern appears in the list by name, next to a small checkbox. Clicking on the name of the pattern will select that pattern.
- The controls below the list will change to affect that pattern only. Click once on a selected pattern to rename the pattern to something more descriptive.
- Clicking on the checkbox changes the state of the Tracker.

Tracker States

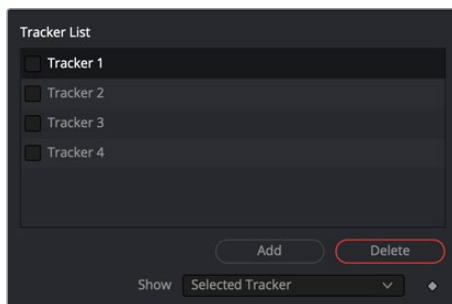
- **Enabled (Black Checkbox):** An Enabled pattern will re-track each time the Track is initiated. Its path data is available for use by other nodes, and the data is available for Stabilization and Corner Positioning.



- **Suspended (Gray Checkbox):** 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.



Show

These two buttons determine what controls are displayed in the node controls. They have no effect on the operation of the Tracker; they only affect the interface.


- **Selected Tracker Details:** When Selected Tracker Details is selected, 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 Trackers is displayed simultaneously below the Tracker List.

Left Pattern Display

The Pattern Display has two image windows next to each other, 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 are moved for a pattern, the display in the left most window will update to show the pattern. As the pattern is moved, the vertical bars immediately to the right of the image display indicate the clarity and contrast of the image channels.

The channel, or channels, with the best clarity are automatically selected for tracking. These channels have a white background in the vertical bar representing that channel. The automatic tracking can be allowed to stand, or you can override the automatic selection and choose the channel used for tracking by disabling the button labeled with a question mark, and selecting the button beneath the channel to track.



Under normal circumstances, the channel selected is indicated 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 only represented in full color if all three channels are selected for tracking.

Override this behavior by selecting the Show Full Color button beneath the Pattern Display, instead of the Show Selected Channel button.

As Fusion looks for the channel with the highest contrast automatically, you might end up tracking the blue channel, especially on scanned film material. Unfortunately the blue channel on most film stock contains the most grain as well, which naturally leads to unclear tracks. Before tracking it's always a good idea to zoom into 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 clear until the first time the selected pattern is actually tracked. 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 is accumulated 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 originally 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 that Fusion has detected extreme variations in the current pattern, and is no longer certain of its accuracy.

After tracking, the pattern display will show a small Flipbook of the track for that pattern, overlaid with a frame number to help identify problem frames for the Track.

Tracker Sizes

Each tracker has a set of sizing parameters that let you adjust the Pattern and Search box on-screen controls.

- **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 re-acquire 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 path from a Tracker pattern, right-click on this control and select Remove Path from the contextual menu.

X and Y Offset

You will often need to track the position of an object in the scene, but that object does not provide for a very reliable pattern. The Offsets permit the tracking of something close to the intended object instead. Use these Offsets to adjust the reported position of the Tracker so that the data is reported for the intended pattern instead of the actual one.

The Offset can also be adjusted directly in the view by activating the positioner's Offsets onscreen control in the Viewer toolbar.



Inspector Controls – Operation Tab

While the Trackers tab controls let you customize how the Tracker node analyzes motion to create motion paths, the Operation tab lets you use this analyzed motion data, stored within each tracker, to perform 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 in 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.

Operation Buttons

These four buttons select the exact function performed by the Tracker. The remaining controls in this tab fine-tune the result of the Operation.

- **None:** The Tracker performs no additional operation on the image beyond simply locating and tracking the chosen pattern. This is the default mode, used to create a path that will then drive another control in Fusion.
- **Match Move:** This mode is used when stabilizing an image, or when matching the position, rotation and scaling of the layer for match moving another element into the scene. Stabilizing requires a minimum of one Tracker to determine position, and two or more to determine scaling and rotation in the sequence.
 - **Corner Positioning:** The Corner Positioning mode is used to track the four corners of a rectangular object and replace the contents with a new image. A minimum of four Trackers is required for this mode, and if there are not enough Trackers, new ones will be created until the total equals four.
 - **Perspective Positioning:** This mode is similar to the Corner Positioner, but rather than replacing the contents of the rectangle, the four tracked corners are mapped to the four corners of the image. This is generally used to remove Perspective from an image. Like the Corner Positioner, this mode requires four Trackers, which will be created if there are not already that many.

Additional Controls For Match Move, Corner Positioning, and Perspective Positioning

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 Foreground input of the Tracker. This array of buttons 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 an image sequence.
- **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 Pop-Up Menus

This drop-down menu provides a variety of options to determine how the two layers should be combined. The options in this menu are explained in more detail in the reference chapter for the Merge node.

Subtractive - Additive Slider

This determines whether the Foreground layer is placed over the Background in using Additive or Subtractive merging techniques. This control is explained in more detail in the documentation for the Merge node.

Filter Method

Determines which filter to use to handle image transforms made using the Tracker node.

Edges (Match Move)

This button array only appears if the Operation mode is set to Match Move. The various options select 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

The Position, Rotation and Scaling checkboxes only appear 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

This checkbox only appears 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 only appears in the Corner Positioning mode. There are two options in the button array.
- **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, and the Classic mode is included pretty much for compatibility reasons only.

Perspective

The preferred setting for this control is True Perspective.

Corner Selector

When the operation of the Tracker is set to either Corner or Perspective Positioner modes, these four drop-down menus appear. They are used to select which trackers map to each of the four corners of the rectangle used by these modes. This is useful when a Tracker has more than four patterns selected, and you must choose which the positioners use.

Rotate Clockwise and Counter-Clockwise Buttons (Corner or Perspective Positioner)

These controls only appear when the operation of the Tracker is set to either Corner or Perspective Positioner 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.

Axis Type (Stabilize Setting)

Under virtually all circumstances, the Axis for any stabilization should be the average position of all Trackers on that frame, however, the rare occasion arises when the Tracker's axis must be elsewhere.

This array of buttons allows for the selection of an axis for the Stabilization based on the position of a single Tracker, or a manual position.

Reference (Stabilize Setting)

The Reference controls are used to set the "snapshot frame" for Stabilization. When you stabilize an image, there must be a position that is considered correct, to which all subsequent movement is detected and corrected.

Match Move Settings

These settings determine how tracking data is correlated with the reference pattern for making transforms.

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.

X/Y Paths

By default the Tracker applies a Displacement Path to the tracked points. To apply an XY Path to the tracked points go to Preferences > Globals > Splines.

Options Tab

The Options tab lets you customize the look of on-screen controls in the Viewer.

Show Pattern Names in Preview

This option defines if the Tracker's name will be displayed in the view. Switch it off to see just the Pattern Rectangle instead.

Show Enlarged Pattern on Dragging

This option defines if there is an enlarged view on positioning the Pattern Rectangle or not.

Enlargement Scale

The Zoom factor that is used on positioning the Pattern Rectangle when the above option is on. 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 (similarly for the 2nd, 3rd and so on)

Planar Tracker Node

The Planar Tracker node is designed to solve a match moving problem that commonly comes up during post-production. Live action footage containing a planar surface such as a license plate, a road sign, or a brick wall and need to replace the numbers in the license plate, a city's name in the road sign, or place a billboard poster on the blank brick wall. The problem is that the camera is moving in the shot so the license plate, road sign, or brick wall undergo continuous changes in perspective. The artist cannot simply merge a new license plate over the existing background without accounting for the perspective distortions. A time intensive way to solve this problem would be to use Fusion's Corner Pin node and manually keyframe the four corners to match the perspective distortions. Planar Tracker automates this keyframing process, to track the perspective distortions of a planar surface in a background plate over time and then re-apply those same perspective distortions to a different foreground.

Part of using 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. Planar Tracker is a useful time saving node in the artist's toolbox and, while it may track most shots, it is not a 100% solution.

What the Planar Tracker Saves

While Planar Tracker does save the final resulting track in the composition on disk, it does not save temporary tracking information such as the individual point trackers (compare with Camera Tracker which does save the individual point trackers). Some consequences of this are:

- The point trackers no longer appear in the view 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 tracking within one Fusion session.
- The size of composition files is kept reasonable (in some situations Planar Tracker can produce 100s of megabytes of temporary tracking data).
- Saving and loading of compositions is faster and more interactive.

The Planar Node's Main Inputs

The Planar Tracker has the following inputs:

- **Background:** Contains the planar surface to be tracked.
- **Corner Pin 1:** An 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:** 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:** Used to mask the output from node.

A Typical Planar Track Workflow

The standard tracking workflow with Planar Tracker is:

- 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 MediaIn node whose footage has a planar surface that can be tracked over time and view the Planar Tracker node in a view.
- 3 Select a reference frame:** Move to a frame where the planar surface to be tracked is unoccluded and set this as the 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 Timeline, adjust the render range to match the range of frames where the planar surface is visible.
- 6 Adjust track options:** Frequently changed options are Tracker, Motion Type, and Track Channel.
- 7 Mask out occluders:** If moving objects partially cover up the planar surface, you may wish to connect in an occlusion mask to Planar Tracker. When using the Hybrid tracker, providing a mask to deal with occluding objects is pretty much mandatory, while with the Point tracker it is recommended to try tracking without a mask.
- 8 Track:** Go back to the reference frame. Press the Track To End button and wait for the track to complete. Go back to the reference frame. Press the Track To Start button and wait for the track to complete. Note that the tracks in the view are not selectable or deletable like in 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:** The Steady, Corner Pin, and Stabilize operation modes use the tracking data produced in Track mode. At this point, you may also choose to export a Planar Transform node that can be used to mirror the tracked perspective distortion onto masks.

The Controls Tab in Track Mode

The contents of the Planar Tracker.

Operation Mode

Lets you choose what you want to do with the PlanarTracker 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 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.

The last three modes (Steady, Corner Pin, Stabilize) use the tracking data produced in Track mode. Note that 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 taken from. It is also the time from which tracking begins. The reference time 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 needs to be carefully chosen to give the best possible quality pattern to track from.

You choose a reference time by moving the playhead to an appropriate frame, and then you click 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 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 maximal size (e.g., when tracking an approaching road sign, it is good to pick a later frame where it is 400x200 pixels big rather than 80x40 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.

After you've drawn a pattern, a set of Pattern parameters let you transform and invert the resulting polygon, if necessary.

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 small patterns like Fusion's trusty Tracker node. The point tracker possesses the ability to automatically create its own 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 view, 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 view.

There is no best tracker. They each have their own advantages and disadvantages:

- **Artist Effort (occlusion masks):** the Point tracker will automatically create its own 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 of 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 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 & 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 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 simpler motion model. This can happen (for example) when a lot of track points are clustered in one side of the region to be tracked or when tracking a small region where Planar Tracker does not have a lot of pixels to work with.

Output

Controls what is output from the Planar Tracker node while in the Track operation mode.

- **Background:** Outputs the input image unchanged.
- **Background - Preprocessed:** 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 current frame to the previous frame.
- **Stop Tracking:** Stops any ongoing tracking operations.
- **Step tracker to next frame:** Tracks from 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 which 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 4x4 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 information currently encoded in the Track spline is shared with the Planar Transform node so that it can replicate the planar distortions tracked by the Planar Tracker node.

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 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 to do 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 time.

Invert Steady Transform

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 2nd set to invert the first 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 degradation (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 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 view.
- 3 **Connect in the texture:** In the Node Editor view, connect the output of the Medialn 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 view 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 and 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 view.

Corner Pin 1 Input Group

- **Each corner pin has a group of related inputs:**
- **Enable:** controls the visibility of the corner pin in the view.
- **Show Grid:** shows a grid over the corner pin. This can be useful when positioning the corners.
- **Merge Options:** control 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 smoothing out the tracked transforms over neighboring frames. Note that Stabilize Mode only smooths out motions, while Steady Mode tries to completely “lock off” all motion.

One thing to be aware of is that Planar Tracker stabilizes based upon the motion of the pattern so it is important 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 a shot has a truck moving down the road and the camera is on a vehicle also moving alongside the truck, and the tracking pattern is picked to be the side of the truck, Planar Tracker will end up smoothing the combined motion of both the truck and the vehicle the camera is mounted on. In some cases, this may not be desired and 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 only the motion of the camera being smoothed.

One unavoidable side effect of the stabilization process is that transparent edges are introduced along the edges of the image. These edges come about 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 be sometimes useful to film at a higher resolution (or lower zoom).

The stabilization workflow with Planar Tracker is:

- 1 **Track:** For the pattern, select a roughly planar region in the shot 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, 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. Playback 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 4x4 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 are:

- **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, Planar Tracker will examine all of 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, Planar Tracker will examine all of the frames and pick the smallest possible zoom factor that removes all of 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

Provides access to a spline whose keyframes contain 4x4 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.

Darken Image

Darkens the image while in Track mode in order to better see the in view controls and tracks. 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 the changing of 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 view. If the pattern is moving very fast, the tracks can look like spaghetti in the view. 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.

Planar Transform Node

The Planar Transform node applies perspective distortions computed by the analysis done by a Planar Tracker node to any input mask. The Planar Transform node can be used to reduce the amount of artist time spent 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 has been 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 needs to be done.

Depending on how well the Planar Tracker followed the object, this can result in a substantial time savings in the amount of tedious rotoscoping that needs to be done. The key to using this technique is recognizing situations where the Planar Tracker will do passably well tracking 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 **Evaluate:** Check the quality of the resulting track. If it does not follow the object passably well then give up and do the rotowork the old fashioned way.
- 3 **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 paste into another composition as desired.
- 4 **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.
- 5 **Refine:** Scrub the timeline to see how well the polygon follows the object. Make tweaks to the polyline on frames where it is off. It is possible to add new points to further refine the polygon.

Controls Tab

The Planar Transform node has very few controls. It's designed to apply the analyzed Planar Tracking data as a match move,

Reference Time

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 4x4 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.

Camera Tracker

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. The Camera Tracker will also create a point cloud in 3D space which can be used to align objects and other 3D models to the live action scene.

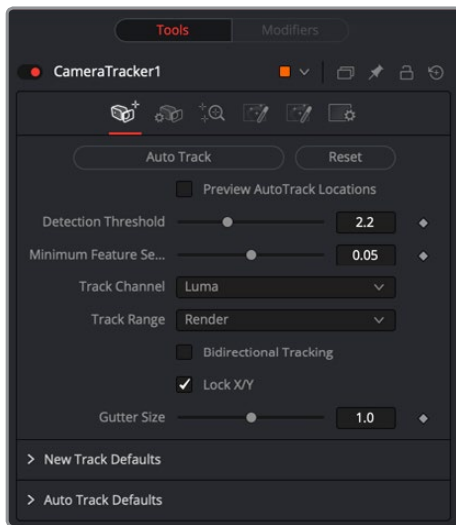
For more information about how to use the Camera Tracker, see Chapter 69, "3D Camera Tracking."



For more information on how the Camera Tracker workflow is designed to operate, see Chapter 69, "3D Camera Tracking"

The 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 - this includes the tracking data and the solve data (camera motion path and point cloud). To only delete 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. 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.

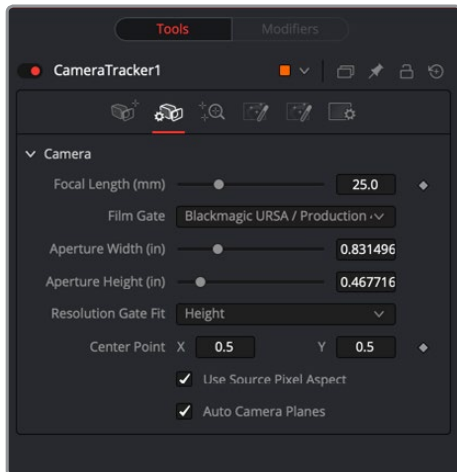
- **Tracker:** Internally, all the Trackers use the Optical Flow Tracker to follow features over time and then further refine the tracks with the trusted Fusion Tracker or Planar Tracker. The Planar Tracker method allows the pattern to warp over time by various types of transforms to find the best fit. 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 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 three top checkboxes define which parameters the solver can adjust to make the solve work.

Refine Focal Length

This will allow the solver to adjust the focal length of the lens to match the tracking points.

Refine Centre Point

Normally ticked off, Camera lenses are normally 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.

Camera

The camera group has 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 a guess if the Refine Focal Length option is ticked.

Film Gate

Choose a film gate preset from the dropdown menu or manually enter the film back size in the Aperture Width and Aperture Height inputs. Note, these values are in inches.

Aperture Width

In the event that the camera used to shoot the scene is not in the preset pulldown, manually enter the aperture width (inches).

Aperture Height

In the event that the camera used to shoot the scene is not in the preset pulldown, manually enter the aperture height (inches).

Resolution Gate Fit

This defines how the image fits the sensor size. Often film sensors are of a size 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 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, or Anamorphic cinema scope is 2:1 aspect.

Auto Camera Planes

When this is 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 anomalously fling points off really deep into the scene and consequently the image plane ends up being pushed really far out, making the resulting scene unwieldy to work with in the 3D views. In these cases, use this option to override this default behavior (or delete the offending tracks).

Lens

When solving for the camera's motion path, Camera Tracker internally creates and then uses a simulated lens 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:

- 1 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).
- 2 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 Parameters

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:

- **None:** Do not do any lens curvature simulations. This should be picked if there is a very low distortion lens or the lens distortion has already been removed from the source footage in a preprocessing step.
- **Radial:** Model only radial lens curvature. This causes the low and high order distortion values to be solved for.
- **Radial & Tangential:** Model and solve for both radial and tangential distortion.

Lower Order Radial Distortion

Determines the quadratic lens curvature.

Higher Order Radial Distortion

Determines the quartic lens curvature.

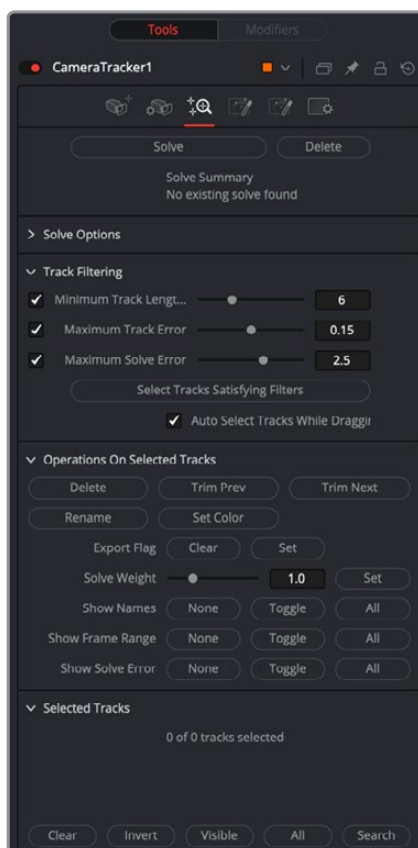
Tangential Distortion X/Y

Determines skew distortion.

Solve

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 cleanup 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.



Average Solve Error

Once the camera has been solved, the Average Solve Error will display here. This number is a good indicator whether the camera solve was successful or not. 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.

Solve Extent

Displays the frame range for which the camera's motion will be reconstructed. The Solver will reconstruct any frames that have tracks on them.

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 of the tracking data.

Clean Tracks By Filter

Clicking this button selects tracks based on the 'Track Filtering' options. If the 'Auto Delete Tracks By Filter' checkbox is selected, 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 check box, 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 check box, 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 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 (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.

Track Filtering

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, giving a stronger set of tracks to solve from. Each filter can be individually ticked on or off.

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.

Minimum 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).

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 view.

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.

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	Will remove the tracks from the set. When there are bad tracks, the simplest and easiest option is to simply delete them.
Trim Previous	Will cut the tracked frames from the current frame to the start of the track. Sometimes it can be more useful to trim a track than deleting 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	Will cut the tracked frames from the current frame to the end of the track.
Rename	Will replace the current auto generated name with a new name.
Set Color	Will allow 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 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 is still has a reconstructed 3D locator. Setting a tracks' 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	Will clear/hide the selected tracks.
Toggle	Will swap the selected tracks and unselect sets.
All	Will select all tracks.
Show Names	Will display the track name, by default these are a number.
Show Frame Range	Will display the start and end frame of a track.
Show Solve Error	Will display the amount of solve error each selected track has.

Selected Tracks

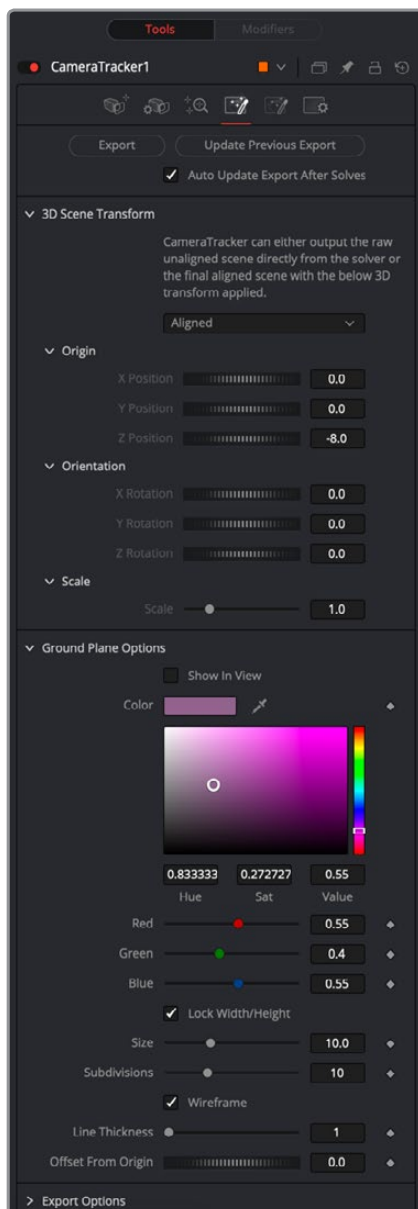
This area displays the properties of a track point or group of points. It has options to:

- **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.

NOTE: Also select tracks directly in the 2D view using the mouse or in the 3D view 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 actually used for compositing.



Export

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 view, 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. An example of where 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 this button to “overwrite” the existing nodes in place.

Automatically Update Previous Export After Solves

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

This defines the origin and alignment of the virtual camera, point cloud, and ground plane. 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 it is looking down the -Z axis. In this default orientation, the physical ground plane in the footage will often not match up with the virtual ground plane in the 3D view and the 3D Scene Transform provides a mechanism to correct that. Note that 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. Also note that if these options need changing, it is important to manually hit “Update Previous Export” to see the results in the exported nodes.

Unaligned, Aligned

The Unaligned button allows the origin and ground plane settings to be adjusted, either manually or by using a selected set of locators from the point cloud. When in unaligned mode, a 3D transform control will be shown in the 3D view which can be manually manipulated to adjust the origin.

Once alignment has been complete, the section is locked by switching to the Aligned button. From the Coordinate System, either output the Raw Unaligned scene where the 3D grid runs directly through the solved camera, or an Aligned scene where the 3D grid can be adjusted to the plane of the user’s choice.

Set From Selection

Takes selected 3D points from the point cloud and aligns the ground plane to fit those points. This can be adjusted individually for position, rotation, and scale. To set a ground plane, select a number of Point Cloud points that are on the ground. To get the best result, try to select as many points as possible belonging to the ground and having a wide separation. Under the Origin pull down, press Set from Selection. This will reposition the Ground Plane. Under the Orientation pull down, press Set from Selection. This will alter the orientation of the Ground plane to the selected points. Finally, press the Aligned button, which repositions the 3D system so the 3D grid is aligned to the Ground Plane.

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 view rather than trying to box select locators in the 3D view.

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	Will set 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 view.
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

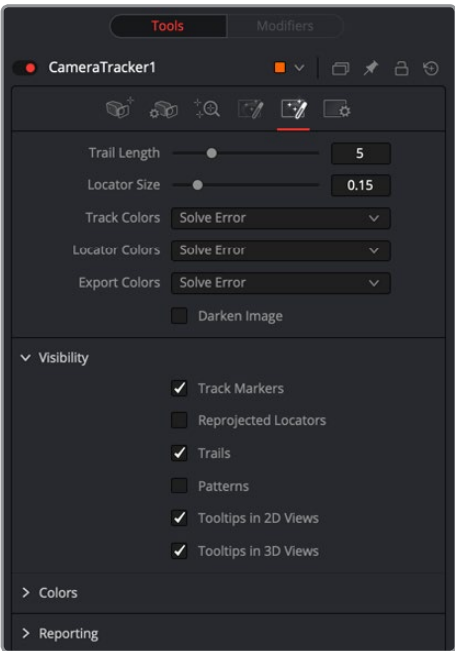
Provides a checkbox list of what will be exported as nodes when the Export button is pressed. These options are Camera, Point Cloud, Ground Plane, Renderer, Lens Distortion, and Enable Image Plane in the camera.

Previous Export

When the Update Previous Export button is clicked, the previously exported nodes listed here are updated with any new data generated (this included the camera path and attributes, the point cloud, and the renderer).

Options

The Options tab lets you customize the CameraTracker onscreen controls so you can work most effectively with the scene material you have.



Trail Length

Will display trail lines of the tracks overlaid on the view. The amount of frames forward and back from the current frame is set by length.

Location Size

In the 3D View, 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	are the onscreen tracks in the 2D view.
Locator Colors	are the Point Cloud locators in the 3D view.
Export Colors	are colors of the locators that get exported within the Point Cloud node.

Darken Image

Dim the brightness of the image in views to better view the overlaid tracks. This affects both the 2D and 3D view.

Visibility

Toggles which overlays will be displayed in the 2D and 3D views. The options are Tracker Markers, Trails, Tooltips in the 2D View, Tooltips in the 3D View, 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 view 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 feature films, Camera Match moving is often farmed out to experts who have a lot of experience with the process of tracking and solving these types of shots. There is rarely a shot where automatic, press a couple of buttons and the process is complete, can be done. It does take understanding of the whole process and what is important to get a good solved track.

The Camera Tracker has to solve for hundreds of thousands of unknown variables, which is a complex task. To create an accurate solve, it is important to get to a precise set of good tracks that exists for a long time. False or bad tracks will skew the result. This section explains how to clean up false tracks and other theory to get a good solve.

Workflow

Track -> Solve -> Refine Filters -> Solve -> Cleanup tracks -> Solve -> Cleanup from point cloud -> Solve -> Repeat.

The Solve is a repeated process to get a good result. Initially, there will be a lot of tracks, and not all are good, so a process of filtering and cleaning up unwanted tracks to get to the best set is a process. At the end of each clean up stage, pressing Solve will give you a result for Solve Error. This needs to be below 1 for it to be good for use, the lower the better. Refining the tracks will result in a better solve.

False Tracks

False tracks are tracks that move or are incorrect and are caused by a number of conditions, such as moving people or 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 Moire effects can cause the track to creep. Reflections in glass on buildings are warped, trees in the wind, 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 important; longer makes the solve fitting better. Bi-directional tracking 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

There are two seed frames that are used in the solve process, the algorithm chooses two frames that are as far apart in time yet share the same tracks, hence longer tracks makes the greater 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 of the sequence.

There is an option that auto detects Seed frames, which can take some time. In the workflow of solving, auto detect is a good idea. When refining the Trackers and resolving, it should be set to Manual mode and use the previous solve's seed frames, which are displayed in the Solve dialog window.

Refine Filters

After the first solve, all of 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 window will indicate how many tracks are affected by the filter. 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 on screen, and deleted by pressing the delete key. This process takes an experienced eye to spot tracks that go bad. Solve again.

View the 3D scene in perspective, 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. Solve again.

Repeat process until the Solve error is below 1.0 before exporting.

Selecting Points for Ground plane

Selecting tracks or points in the point cloud will allow for alignment and fitting of the ground plane.

Chapter 53

Transform Nodes

This chapter details the Transform nodes available in Fusion.

Contents

Camera Shake [CSH]	1139
Crop [CRP]	1142
DVE [DVE]	1144
Letterbox [LBX]	1145
Resize [RSZ]	1148
Scale [SCL]	1150
Transform [XF]	1153

Camera Shake [CSH]

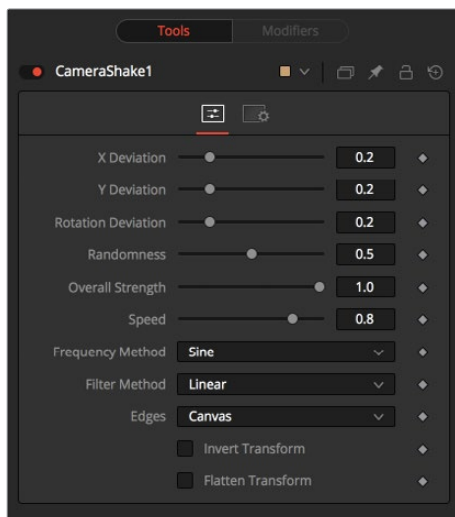


This node can simulate a variety of Camera Shake style motions from organic to mechanical. It is not the same as the Shake Modifier, which generates random number values for node controls.

See the Connections and Modifiers chapter for more information about the Shake Modifier.

The Camera Shake node concatenates its result with adjacent transformation nodes.

Controls



Deviation X and Y

These controls determine the amount of shake that is applied to the image along the horizontal (X) and vertical (Y) axis. 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 will generate a much more mechanical looking motion than a Sine.

Edges

This determines how the Edges of the image are treated.

Black

This causes the edges that are revealed by the shake to be 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.

Invert Transform

Select this control to Invert any position, rotation or scaling transformation. This option might be useful to 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.

Filter Modes

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 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. 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.

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, Bessel

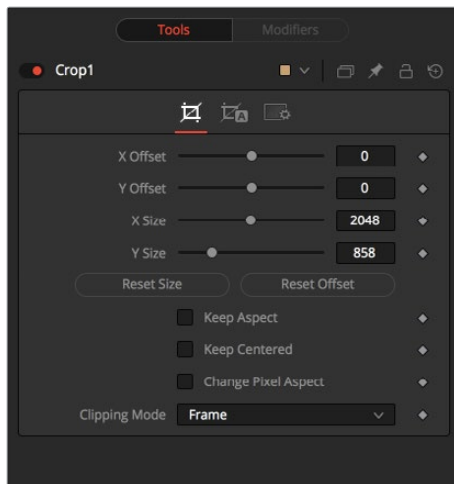
Crop [CRP]



The Crop node can be used to cut out a portion of an image or to offset the image into a larger image area. This node actually changes the resolution of the image and should not be animated under normal circumstances.

It is possible to crop an image in the view by viewing the upstream node and selecting a region with the mouse. To do so, first activate Allow Box Selection by clicking on the icon next to your view, then dragging a rectangle to perform the operation.

Controls



Offset X and Y

These controls position the top left corner of the Cropping window. Values larger than the actual resolution of the cropped image cause it to disappear off the edges of the output image. Values below 0 will push the input image toward the bottom right of the result. The values of these controls are measured in pixels.

Size X and Y

Use this control 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.

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 pixel aspect that the image is considered to have.

Clipping Mode

This option sets the mode used to handle the edges of the image 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 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 which would normally be outside the upstream DoD will be treated as black/transparent.

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

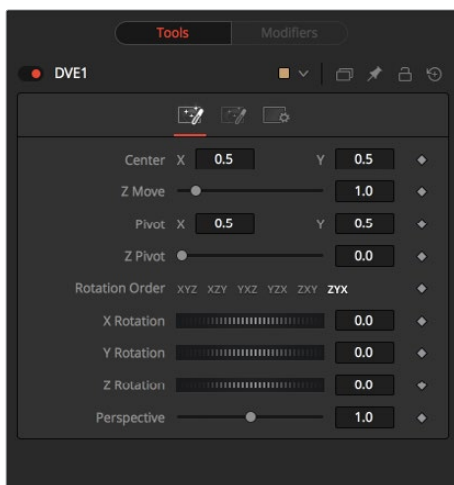
This evaluates the image and attempts to determine the background color. It then crops each side of the image to the first pixel that is not this color.

DVE [DVE]



The DVE node is a 3D-image transformation similar to nodes found in an online edit suite. The node encompasses image rotations, perspective changes and Z moves. The axis can be defined for all transformations.

Controls



Pivot X, Y and Z

Position 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 on screen. 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.

Unlike regular effect masking, the masking process occurs before the transformation. To use this feature, connect a mask to the DVE node's DVE Mask input. Alternatively, select the DVE node, right-click on the view and select DVE Mask from the contextual menu. All of 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 image over the input image to be either Subtractive or Additive. For more information on Subtractive and Additive merges, see the Merge node.

Letterbox [LBX]



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 adapt film resolution images to NTSC or PAL 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.

Controls



Formats Contextual Menu

Place the pointer on the Aspect X or Y control and press the right mouse button to display a menu with available frame formats. Select any one of the choices from the menu to set the Height, Width, and Aspect controls automatically, or enter the required information manually.

Width and Height

The values of these controls determine the size of the output image as measured in pixels.

Pixel Aspect X and Y

The 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 Modes

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 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. 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. 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



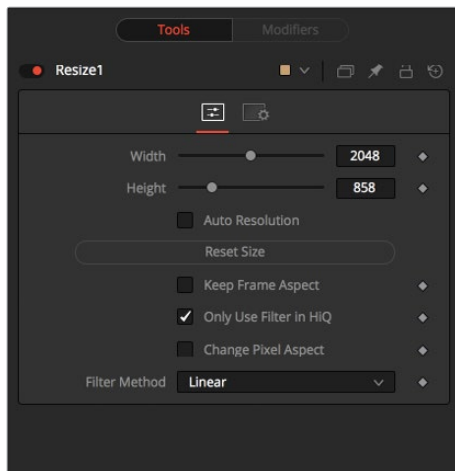
Different Resize Filters. From left to right: Nearest Neighbor, Box, Linear, Quadratic, Cubic, Catmull-Rom, Gaussian, Mitchell, Lanczos, Sinc, Bessel

Resize [RSZ]



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).

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.

Keep Frame Aspect

When toggled on, the Resize node maintains the aspect of the original image, preserving the original ratio between width and height.

Reset Size

Reset the image dimensions to the original size of the image.

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 pixel aspect that the image is considered to have.

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 will often depend on the amount of scaling and on the contents of the image itself.

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 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. 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. 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.


NOTE: Because this node changes the physical resolution of the image, we do not normally advise animating the controls.

Example



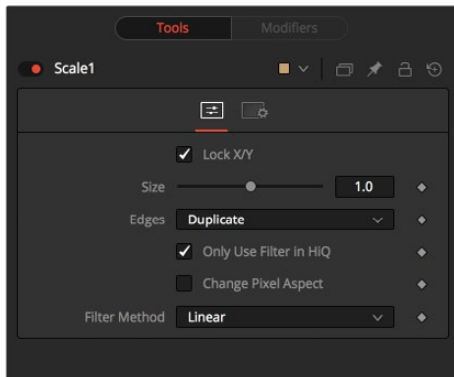
Different Resize Filters. From left to right: Nearest Neighbor, Box, Linear, Quadratic, Cubic, Catmull-Rom, Gaussian, Mitchell, Lanczos, Sinc, Bessel

Scale [SCL]



The Scale node is almost identical to the Resize node, except that Resize uses exact dimensions where the Scale node uses relative dimensions to describe the change to the source image's resolution.

Controls



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 X/Y

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 effect 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.

Change Pixel Aspect

Enable this checkbox to reveal a Pixel Aspect control that can be used to change the Pixel Aspect that the image is considered to have.

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.

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 will often depend on the amount of scaling and on the contents of the image itself.

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 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. 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.

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.

NOTE: Because this node changes the physical resolution of the image, we do not normally advise animating the controls.

Example



Different Resize Filters. From left to right: Nearest Neighbor, Box, Linear, Quadratic, Cubic, Catmull-Rom, Gaussian, Mitchell, Lanczos, Sinc, Besse

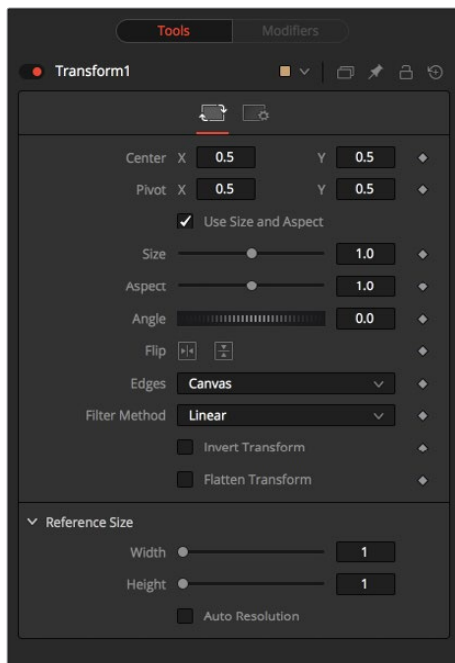
Transform [XF]



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. For more information on concatenation, consult the Transformations chapter.

Controls



Flip Horizontally and Vertically

Toggle this control on to flip the image along the X- or Y-axis.

Edges

Edges determines how the edges of the image will be treated.

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.

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 Size, or 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 counter-clockwise direction. Decreasing the Angle rotates the image in a clockwise direction.

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 re-introducing 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. See the Transformations chapter earlier in this manual for details on concatenated transformation.

Reference Size

The controls under the Reference Size reveal 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 by 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 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 only visible in the value shown for Transform 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.

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.

Filter Method

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 results.

Chapter 54

Warp Nodes

This chapter details the Warp nodes available in Fusion.

Contents

Coordinate Space [CDS]	1158
Corner Positioner [CPN]	1159
Dent [DNT]	1160
Displace [DSP]	1161
Drip [DRP]	1162
Grid Warp [GRD]	1164
Lens Distort [LENS]	1169
Perspective Positioner [PPN]	1171
Vector Distortion [DST]	1172
Vortex [VTX]	1173

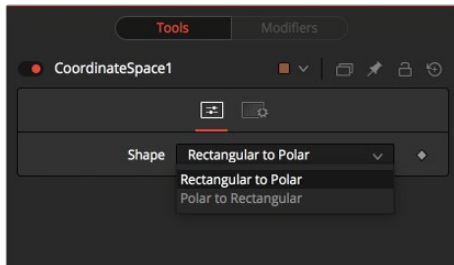
Coordinate Space [CDS]



The Coordinate Space node changes the coordinate space of the image at its input from Rectangular to Polar or from Polar to Rectangular.

Controls

Select between Rectangular to Polar and Polar to Rectangular. Consider the following example to demonstrate the two coordinate spaces.




Example

To demonstrate a basic tunnel effect that can be achieved with this node:

- Create some text and animate it to move along a path from the top of the frame to the bottom.
- Connect the output of the Text+ node to a Coordinate Space node.
- Select Polar to Rectangular from the Shape menu.

As the text moves from top to bottom along the original path, it will appear to move from an infinite distance in the Coordinate Space node. It may be necessary to flip the text using the Transform (Xf) 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.

Corner Positioner [CPN]



The Corner Positioner can be used to interactively position the four corners of an image. 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.

Controls



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 will be 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. Useful when the corners are attached to Trackers with patterns that may not be positioned exactly where they are needed.

Dent [DNT]



The Dent function creates a circular deformation of an image similar to a Fish Eye Lens effect, with the choice of six different Dent filters.

Controls



Type

Select the type of Dent filter to use from this menu. All parameters for Dent are animatable.

Dent 1

This creates a bulge dent.

Kaleidoscope

This creates a dent, mirrors, 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

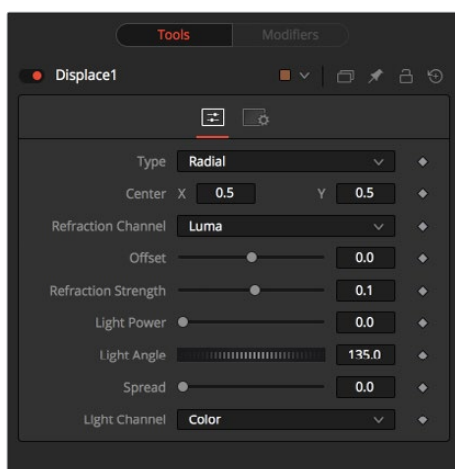
This changes the overall Strength of the Dent.

Displace [DSP]



This 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.

Controls



Type

The Type buttons can 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 from.

Refraction Channel

This array of buttons 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 spread out.

Light Channel

Select the channel from the refraction image that will be used as the simulated light source. Select from Color, Red, Green, Blue, Alpha, or Luminance channels.


NOTE: The Radial mode pushes pixels inwards or outwards 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.

Drip [DRP]



The Drip function 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.

Controls



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 randomly dispersed noise that distorts your image. 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 like 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 Drip.

Frequency

This changes the number of ripples emanating from the center of the Drip effect. A value of 0.0 indicates that there will be no ripples. Move the slider up to a value of 100, to correspond with the density of desired ripples.

Phase

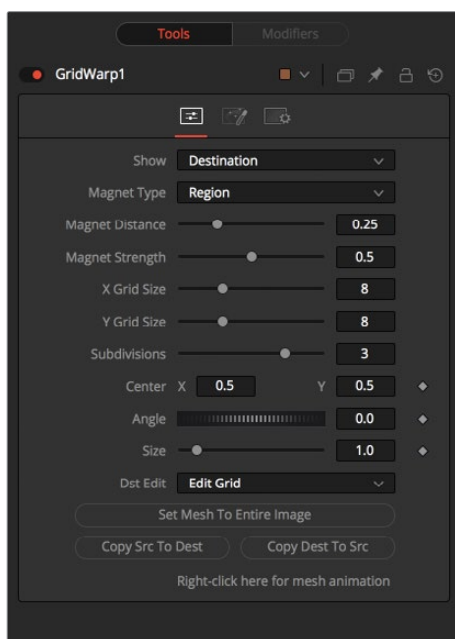
This controls the offset of the frequencies from the center. Animate the Phase value to make the ripple emanate from the center of the effect.

Grid Warp [GRD]



The Grid Warp node is a simple 2D deformation grid with flexible vertices. The image will be deformed so that the Source grid matches the Destination grid.

Controls



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 will be highlighted to indicate that is the active grid at the moment.

All of the 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 will be 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, when the left mouse button is clicked, will move. New points that enter the region during the move will be ignored. Choosing this option will expose 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 when the left mouse button is clicked will move. New points that enter the region during the move will be affected as well. Choosing this option will expose 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 affect for the magnet is, as in the size of the circle. Drag on the grid and any vertex within the range of the slider will move.

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 will increase or decrease the falloff of the magnet cursor's effect. At a setting of 0.0, the magnetic cursor has no effect, and vertices will 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 will reset the entire grid. Set the X and Y grid sizes to the appropriate resolution before making serious 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 there are, the smoother the deformation is likely to be (and the slower it will be to render the deformation).

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 on a face, 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 extremely useful for creating grids around organic shapes. When this mode is enabled, all onscreen controls will 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 will be 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 only visible when the Edit Line mode is enabled. The Point Tolerance slider determines how much tessellation the grid will apply to closely match the density of points in the spline. The lower this value, the fewer vertices there will be in the resulting grid and the more uniform the grid will appear. Higher values will start applying denser grids with variations to account for regions in the spline that require more detail.

Oversize Amount

This control is only visible 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 only visible when the Edit Line mode is enabled. The Snap Distance slider dictates how strongly the drawn spline will attract surrounding vertices. If a vertex is close enough to a spline's edge, the vertex will move to line up with the spline. The higher the value, the farther the reach of the spline.

Right-Click Here for Mesh Animation

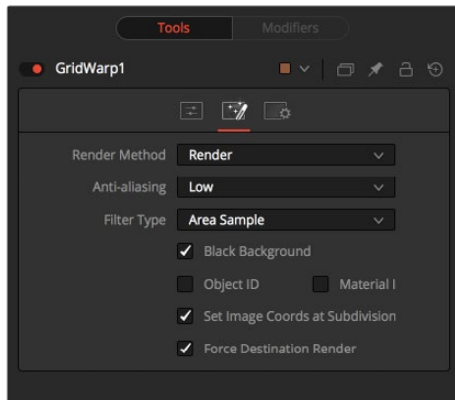
The grids are static by default. Right-clicking on the Right-Click Here for Mesh Animation label will provide a contextual menu with options for animating the grid or connecting it to another grid in the composition.

The grid uses a Polychange spline. Any adjustment to any of the control points will add or modify the keyframe for all points to the Polychange spline.

Right-Click Here for Shape Animation

This label only appears in the Edit Line mode. Right-clicking on the Right-Click Here for Shape Animation label will reveal a contextual menu that can be used to animate the shaping polyline or to connect it to other polylines in the composition.

Render Tab



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 is a checkbox in Wireframe Render mode.

It is otherwise 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 almost never for a final render.

Filter Type

The Area Sample control is only visible when the Render Method is not set to Wireframe. Checked by default, disabling this checkbox prevents the grid from calculating area samples for each vertex in the grid. Area sampling vastly improves render quality at the cost of speed.

Wireframe Width

This control only appears when the Render Method is set to Wireframe. It determines the width of the lines that make up the wireframe.

Anti-Aliased

This control only appears when the Render Method is set to Wireframe. Use this checkbox to enable/disable antialiasing 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 will be set to black or if they will be 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

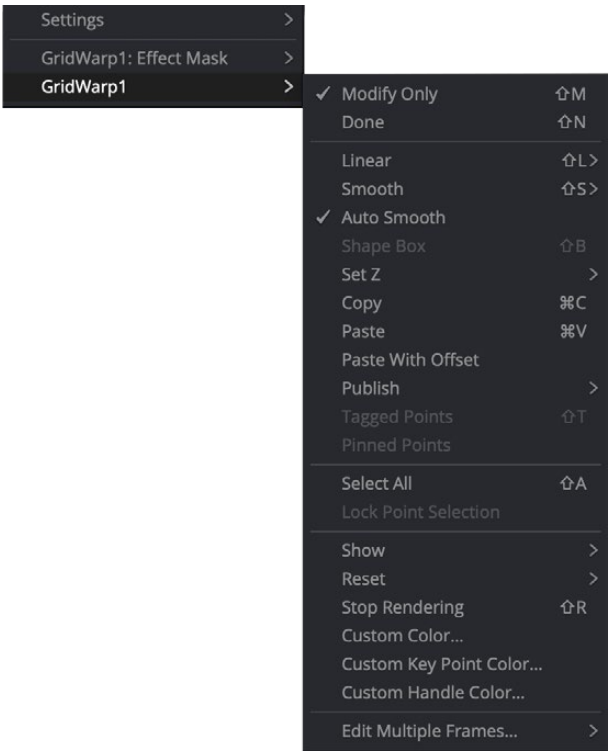
Checkbox that defaults to On.

Force Destination Render

Checkbox that defaults to On.

Right Click Menu

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 all 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 to apply or remove smoothing from selected vertices.

Auto Smooth Points

When Auto Smooth Points is enabled, the vertices in the grid will automatically be smoothed whenever they are moved. This is generally on 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 will Select 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), Bezier handles and subdivisions in the Viewers.

Reset Selected Points

Reset Selected Points (vertices) to their default positions.

Reset All Points

This will Reset All Points (vertices) in the mesh to their default positions.

Stop Rendering


This option will enable Stop 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.

Screen Controls

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 above in the contextual menu.



Lens Distort [LENS]




This node can be used to remove or add lens distortion in an image. The amount and type of lens distortion in an image depends on a variety of factors such as the actual lens that has been used, the quality and amount of lensgroups in that lens, adjustments like shift lenses and much more.

One reason to remove lens distortion is to comp an undistorted layer as, for example, a 3D rendering, on top of a distorted layer as, for example, a real world camera shoot. When combining the layers without removing the lens distortion, unwanted effects like straight lines not matching up on foreground and background will occur. The resulting composite will not look believable.

In a typical workflow one would apply the LensDistort in Undistort mode to the original layer, add the 3D elements, do all other compositing work and finally apply LensDistort again with exactly the same settings, but this time in Distort mode to get the original look and distortion back into the image.

Controls



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 black pixels would in that case result in 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 will be discarded.

Output Distortion Map

Outputs the location of pixels as a warped screen-coordinate map.

Camera Settings

The options known from the Camera3D are duplicated here. They can either be set manually or connected to an already existing Camera3D.

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 will be noticeable differences to higher settings.

Supersampling Mode [HiQ]

The type of sample done for each supersample. Nearest will lead to a crisper but more aliased image. Bilinear will give 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 the 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 a software like 3DEqualizer which analyzes the footage and delivers a dataset that can be imported into the LensDistort node right away.

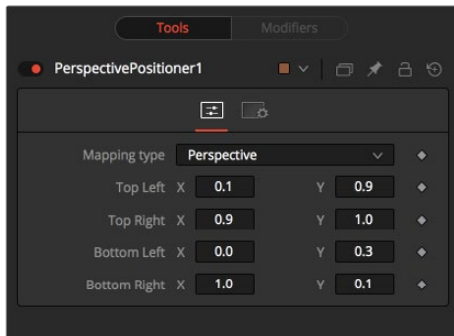
And finally, one could try to manually eyeball the amount of lens distortion using the control sliders. To do that, one could either look out for horizontal or vertical lines in the footage that are supposed to be straight and straighten them out using the controls, or shoot a full frame checkerboard pattern on set as a reference.

Perspective Positioner [PPN]



The Perspective Positioner is the complementary node to the Corner Positioner node. By positioning corner points on an image and moving them, it is possible to remove the perspective from an image. This function can also be used to wobble and warp the image by animating the points over time.


Controls



Corners X and Y

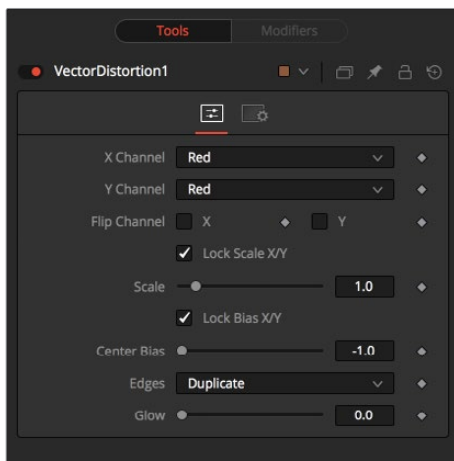
There are the four control points of the Perspective Positioner. Interactively drag these in the views to position each corner of the image.

Vector Distortion [DST]



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 a channel from a second image or reference image.

Controls



X Channel and Y Channel

Use these buttons to select which channel of the reference image will be used to distort the X and Y channels. If no Distort reference image is connected to the second (green) input of the node, then channels from the main input are used instead.

Flip X and Flip Y Channel

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 multiplication 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

The Edges buttons are used to set how the node deals with pixels that reach the edge of the screen.

Glow


Use this slider to add a glow to the result of the vector distortion.

Vortex [VTX]



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.

Controls



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

Change the area affected by the Vortex in the display window by dragging the circumference of the effect or by using 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.

Chapter 55

Modifiers

This chapter details the Modifiers available in Fusion.

Contents

Bezier Spline	1177
B-Spline	1178
Calculation	1178
Cubic Spline	1181
Expression	1181
Expression Syntax Formulas	1184
Expression Syntax Operators	1185
From Image	1186
GradientColorModifier	1188
MIDI Extractor	1189
More About MIDI	1192
General	1193
Natural Cubic Spline	1193
Offset	1194
Path	1197
Perturb	1198
Probe	1200
Publish	1202
Shake	1202
Track	1204
Vector	1205
XY Path	1206


Bezier Spline

The BezierSpline is one of the animation modifiers in Fusion and normally is applied to numerical values rather than point values. It is applied by default each time you right-click on a numerical control and select Animate.

It can also be applied by right-clicking on a numerical control and selecting Modify with > BezierSpline.

Usage


Being an animation spline, this modifier has no actual Controls tab. However, its effect can be seen and influenced in the Spline Editor. The Bezier Spline offers individual control over each point's smoothness by means of handles. The smoothness can be applied by multiple ways.



- To make the keys smooth, select them and press Shift-S. The handles can then be used to further modify the behavior of in and out.
- To make the keys linear, select them and press Shift-L. These operations can also be performed using the contextual menu.
- Select the keyframe(s), right click and select Smooth or Linear. The menu also allows the user to apply smoothing using a Savitzky-Golay filter. Select the keyframe(s), right-click and select Smooth Points -Y Dialog.

Ease In/Out Can Also Be Modified by Using the Control Slider

Select the keyframe you want to modify, right-click and select Ease In/Out... from the contextual menu. Then use the sliders to individually control the Ease In/Out numerically.



B-Spline

The B-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 on a numerical control and selecting Modify with > B-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. 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 left-clicking on the keyframe to select it, holding down W as well as the left mouse button, and moving the mouse to the left to lower the tension and to the right to increase the tension. This also can be done with multiple selected keyframes simultaneously.

Calculation

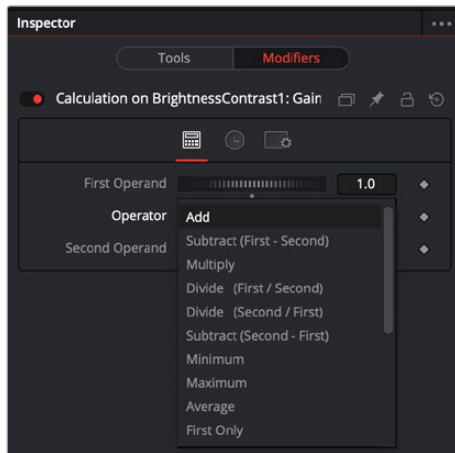
Calculations are used to create indirect connections between value controls. A calculation can perform a mathematical expression based on two operands, where each operand can be connected to another control or set manually by the user.

Additionally, the Calculation control can access the values of a connected control at times other than the current time, using Time offsets and Time Scale controls built into the Calculation Modifier.

The most common use for a calculation is when two controls need to be connected to each other, but the range or scope of the value's output by one control is inappropriate for the other control.

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 to a Calculation than it is to do so with an Expression.

Calc Tab



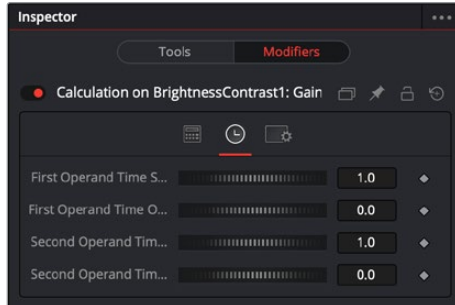
First and Second Operand

These sliders are either connected to published or animated controls from other nodes in the composition 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 on the down arrow opens the menu with the following options:

Time Tab



First and Second Operand Time Scale

These sliders cause the frame used to read the values of the Operands specified in the Calc tab to be scaled. A value of 1 returns the value of the operand at frame x when the composition is set to frame x. For example, if the first operand is animated with a value of 1 to 10 from frame 0 to 10, then a scale of 0.5 would cause the calculation to return a value of 5 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 in inverse proportion to the size of the text.

- Create a new composition, and set the Global Start and Render Start to 1. Set the Global End and Render End to 100.
- Add a Text Plus node to the composition.
- Enter a small amount of text in the StyledText input.
- Make sure the current frame is 0
- Set the Size parameter of the Text node to a value of 0.06 at frame 0.
- Right-click on the Size slider and animate the slider by setting a key.
- Advance to frame 100 and set the value of the Size control to 0.50.
- Add a Blur node immediately after the Text node. Connect the Blur input to the Text node's output.
- View the Blur node in one of the Viewers.
- We want the blur to decrease in strength as the text gets bigger. The controls cannot be directly connected together because the values of the Text Size control are getting bigger instead of smaller.
- Right-click on the Blur size and select Modify With Calculation from the contextual menu. A new set of controls will appear in the Modifiers tab while the Blur node is selected.
- Switch to the Modifier tab (F11)
- Right-click on the First Operand slider. Select Connect To > Text 1 > Size from the contextual menu. This connection isn't very useful, though; 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.
- Switch to the Time tab of the modifier and set the First Operand Time Scale to -1.0. Normally, the First Operand will get the value of the control it is connected to from the same frame as the current time. So at frame 10, the first operand will be 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 Nodes tab of the Node Control Area (F9). Press play (spacebar) and watch how the value of the Blur Size relates to the value of the Text Size.

Cubic Spline

The 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 on 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.



Expression

Adding an Expression modifier to a control will add the ability to control and manipulate either position or value controls 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.

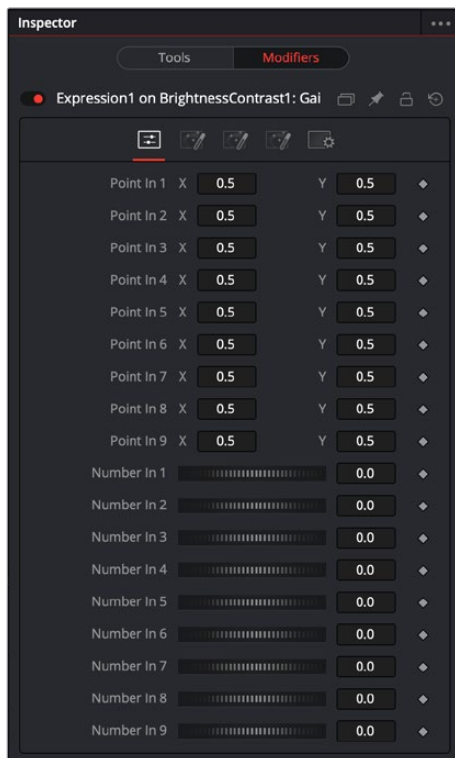
An Expression will accept up to nine value inputs and nine position inputs that are used as part of a user-defined mathematical expression to output a value.

To add an Expression to a control, right-click on the control and choose Modify With > Expression from the contextual menu. The type of value that will be returned by the Expression entirely depends on the type of control it is modifying.

When used with a value control (like a slider), the Expression in the Number Out tab will be evaluated to create the result. When used to modify a positional control (like a node's center), the Point Out tab will control the result.

The Modifiers view contains the controls for the Expression modifier. Its controls are described below.


Controls



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.

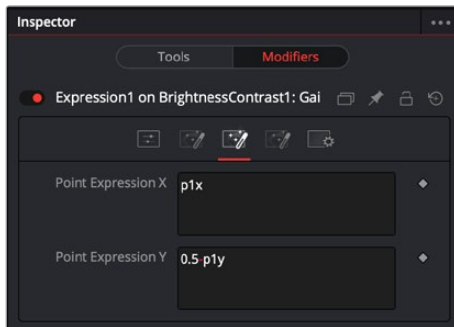
These values can be set manually by the user, connected to other parameters, animated and even connected to other Expressions or Calculations.

Number Out Tab



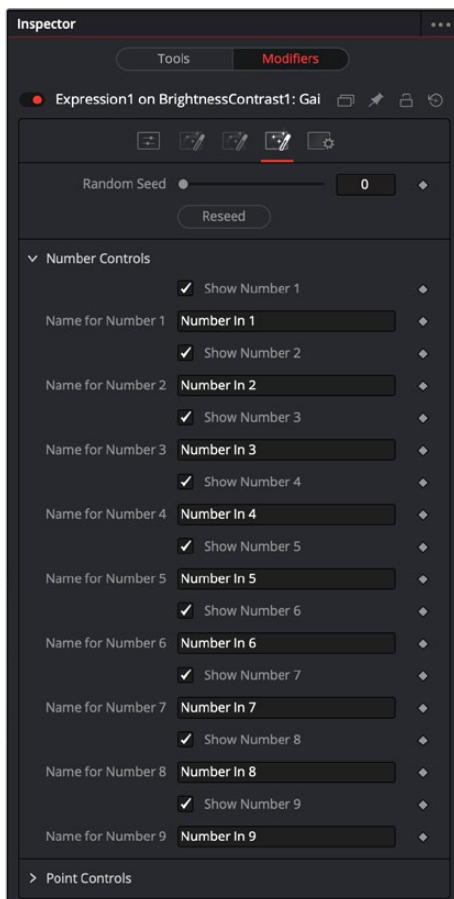
This enables a mathematical formula to be entered, which can access the values from both the Number In and the Point In tabs and output a value used to modify the control to which the Expression was applied. See below for the syntax to use in this field.

Point Out Tab



Each one of the text boxes in this tab can contain a mathematical formula that can access the values from both the Number In and the Point In tabs and output a value used to modify the control to 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.

Config Tab



A good expression can be re-used over and over again. As a result, it can be useful to provide more descriptive names for each input and to hide the ones that are unused. The Config Tab of the Expressions modifier is used to customize visibility and name for each of the nine point and number inputs.

Random Seed

The Random Seed control is used to seed 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 will always be 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 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 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.

<code>dist(x1, y1, x2, y2)</code>	The distance between point x1,y2 and x2,y2.
<code>dist3d(x1,y1,z1,x2,y2,z2)</code>	The distance between 3D points x1,y2,z1 and x2,y2,z2
<code>noise(x)</code>	A smoothly varying Perlin noise value based on x
<code>noise2(x, y)</code>	A smoothly varying Perlin noise value based on x and y
<code>noise3(x, y, z)</code>	A smoothly varying Perlin noise value based on x, y and z
<code>if(c, x, y)</code>	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.

<code>x + y</code>	x plus y.
<code>x - y</code>	x minus y.
<code>x < y</code>	1.0 if x is less than y, otherwise 0.0.
<code>x > y</code>	1.0 if x is greater than y, otherwise 0.0.
<code>!x</code>	1.0 if x = 0, otherwise 0.0.
<code>-x</code>	(0.0 - x).
<code>+x</code>	(0.0 + x) (effectively does nothing).
<code>x ^ y</code>	x raised to the power of y.
<code>x y</code>	x multiplied by y.
<code>x / y</code>	x divided by y.
<code>x % y</code>	x modulo y, (remainder of (x divided by y)).
<code>x <= y</code>	1.0 if x is less than or equal to y, otherwise 0.0.
<code>x >= y</code>	1.0 if x is greater than or equal to y, otherwise 0.0.
<code>x = y</code>	1.0 if x is exactly equal to y, otherwise 0.0.
<code>x == y</code>	1.0 if x is exactly equal to y, otherwise 0.0 (identical to above).
<code>x <> y</code>	1.0 if x is not equal to y, otherwise 0.0.
<code>x != y</code>	1.0 if x is not equal to y, otherwise 0.0 (identical to above).
<code>x & y</code>	1.0 if both x and y are not 0.0, otherwise 0.0.
<code>x && y</code>	1.0 if both x and y are not 0.0, otherwise 0.0 (identical to above).
<code>x y</code>	1.0 if either x or y (or both) are not 0.0, otherwise 0.0.
<code>x y</code>	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 Bezier 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 have a look at the results. (Try this one with motion blur.)

From Image

The From Image 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.

It can be applied by right-clicking on a Gradient control and selecting From Image.

Controls

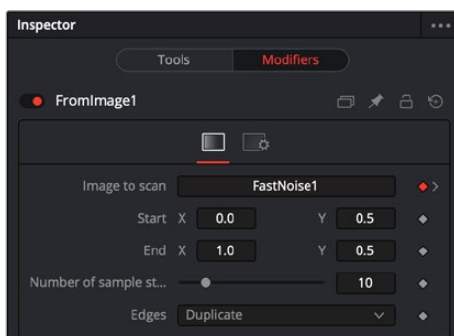


Image to Scan

Drop the node you want to be examined into this box.

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 view.

Number of Sample Steps

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 points you will see 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 will stay intact and can then be fine tuned by hand.

Edges

Edges determines how the edges of the image will be 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

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.

GradientColorModifier

The Gradient Color modifier allows the user to control the value of a parameter. A gradient with customized values is mapped into a specific time range to control the value. If both 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 onto a control and selecting Modify with > GradientColorModifier.

Controls



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 own 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 respectively and shifting the gradient to the left.

- **Once:** When using the Gradient Offset control to shift the gradient, the border colors will keep their values. Shifting the default gradient to the left will result in a white border on the left, shifting it to the right will result in a black border on the right.
- **Repeat:** When using the Gradient Offset control to shift the gradient, the border colors will be wrapped around. Shifting the default gradient to the left will result in a sharp jump from white to black, shifting it to the right will result 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 will result in the edge fading from white back to black, shifting it to the right will result 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.

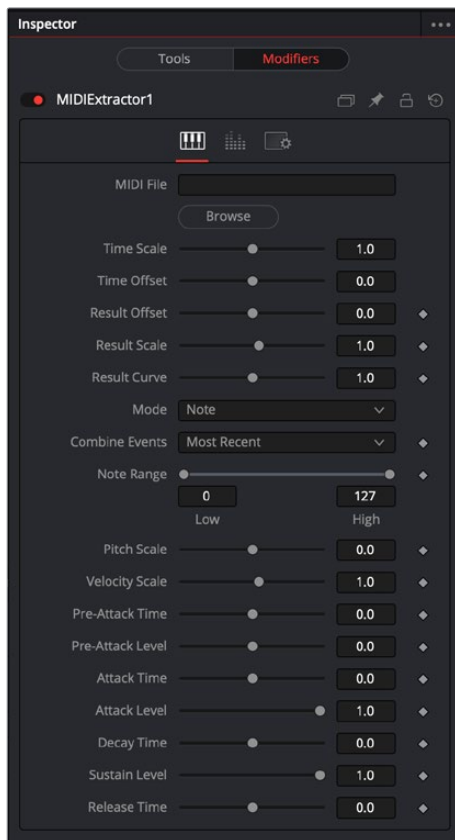
MIDI Extractor

The MIDI Extractor Modifier provides the ability to modify the value of a control using the values stored in a MIDI file. Using the modifier relies on some knowledge of MIDI, which is outside of the scope of this manual to describe in detail.

The value produced by the modifier is extracted from the MIDI event selected in the Mode menu. Each mode can be trimmed so 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 on a control and selecting Modify with > Midi Extractor.

Controls



MIDI File

This browser control is used to specify the MIDI file that will be 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 will play back the MIDI events at normal speed, 2.0 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 part way 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 will 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 result will fall linearly between 0.1 and 1.0 (for example, a velocity 127 note will generate 1.0, where 63 will generate approximately 0.5).

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 actually use any specific messages. It bases its event timing on the tempo map contained in the MIDI file.

Combine Events

This menu selects what will happen 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 will occur when in Beat mode. This is in Quarter notes so a value of 1.0 will give a beat every quarter.

Note Range Note and Poly Aftertouch Modes Only

This defines what range of notes will cause 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 will cause 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 will cause 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

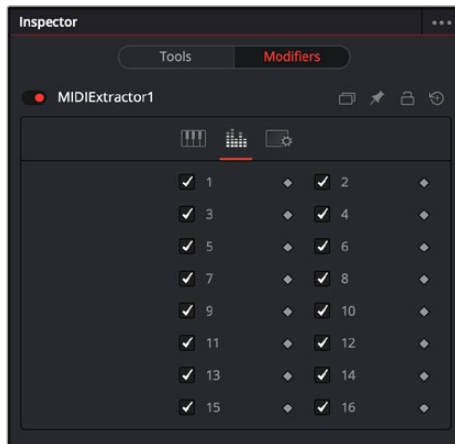
This specifies the MIDI controller number from which to extract events.

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 will start 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 on the beats.

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.

Channels Tab



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.

More About MIDI

A single MIDI interface allows 16 channels. Typically, these will be assigned to different instruments within a device or different devices.

Usually, MIDI data is 7-bits, ranging from 0..127. This is represented as a value between 0..1 to be more consistent with the way data is handled throughout the rest of Fusion.

There are quite a number of 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, how fast the key was 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, how fast the key was 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 things 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 whatever) during a note. It is 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).

General

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 value, this control has a range between -1 and 1, and a correspondingly finer degree of resolution.

For a resource on how MIDI works, have a look at <http://www.harmony-central.com/MIDI/Doc/doc.html>.

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 on a numerical control and selecting Modify with > Natural Cubic Spline.

NOTE: Unlike other spline types, Cubic splines have no control handles and attempt to automatically provide a smooth curve through the keypoints.

Usage

Being an animation spline, this modifier has no actual Controls tab, however, its effect can be seen and influenced in the Spline Editor.



Offset

Offsets are useful for creating constant or animated variances between values, relating to various controls, paths and points. There are three types of offsets available in Fusion:

- Offset Distance
- Offset Angle
- 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 of the offsets use the same set of controls, which behave differently depending on the offset type used. These controls are described below.

Offset Distance

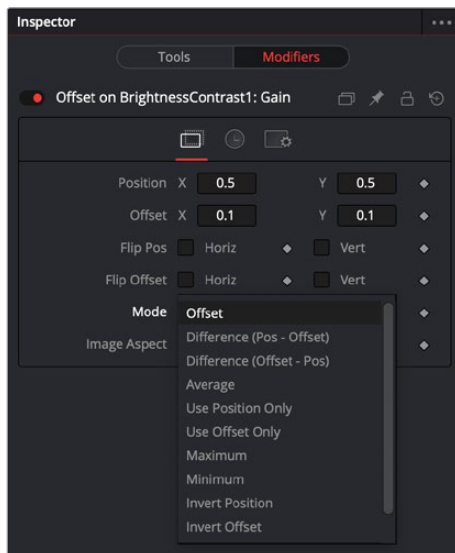
The Offset Distance modifier outputs a value that is based on the distance between two positional controls. This modifier is the first one discussed that is capable of outputting a value based on a mathematical expression applied to a position.

Offset Position

The Offset Position modifier outputs a position (X and Y coordinate) 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 on a control and selecting Modify with > Offset.

Offset Tab



Position X and Y

The first position value is used by the Position to generate the calculation.

Offset X and Y

The first position value is used by the Offset to generate the calculation.

Flip Position Horizontal and Vertical

When these controls are selected, the Position will be mirrored along the vertical or horizontal axis of the image.

Flip Offset Horizontal and Vertical

When these controls are selected, the Offset position will be mirrored along the vertical or horizontal axis of the image.

Mode

Select an option from the Mode menu to choose the mathematical operation performed by the offset control.

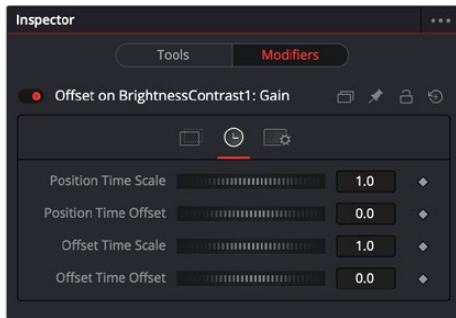
Available options are:

- 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.

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 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 Offset at the Time Offset specified.

Example

- Let's continue the Text banking example at the beginning of this chapter to illustrate one potential use of offsets.
- Select and view the Merge node in this node tree. Right-click on the Merge Size control and select Modify With > Offset Distance from the contextual menu. This will add two new crosshairs and Offset controls in the Modifier tab.
- The size of the text will now be determined by the distance or offset between the two new cross hairs. These points are animatable and can be connected to other controls.
- Connect the position value of the offset to the existing path by right-clicking on the Position control and selecting Connect To > Path on Merge 1 Center Value. Manually place the Offset crosshair at the bottom of the screen between the two path points.
- Now, the text should shrink near the ends of the path (when the distance between the offset and the path is at its minimum) and grow at its ends (where the distance between the offset and the path is at its maximum).

Path

The Path 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 a Coordinate control using a Path, right-click on the control and select Path from the contextual menu.

Controls



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 control along its path, represented as a value between 0.0 and 1.0. Displacement splines are used to control the speed of a control's movement along its 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.

- 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 will have an associated point on the Displacement spline.

Unlocked points will not have a corresponding point on the Displacement spline.

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.

Right Click Here for Shape Animation

It is possible to animate the shape of the path as well, or to connect it to other path controls like Polyline Masks or Paint Strokes.

Switching Default Paths

Change the Default Path type used when animating a Coordinate control to Path (if this is the preferred type of animation). Open the Global preferences and look under the Default tab for the drop-down list in the default animate section labeled Point With. Change this from the current value to Path. The next time Animate is selected from a Coordinate control's contextual menu, a Path will be used.

Perturb

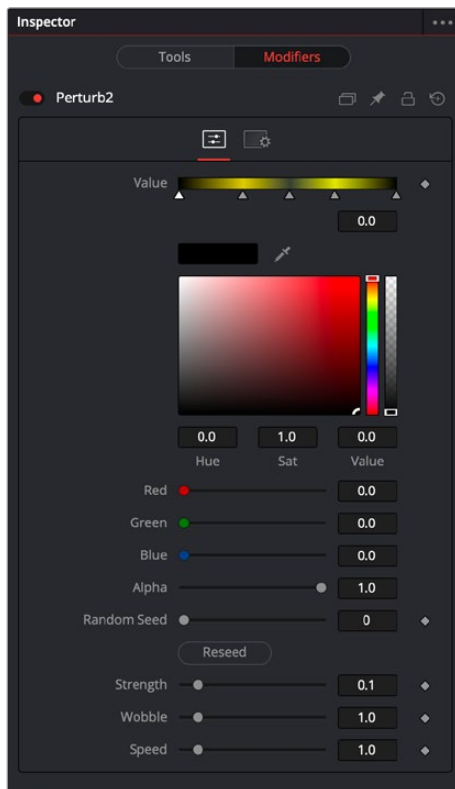
The Perturb modifier generates smoothly varying random values for several input types, based on Perlin noise. It can be used to add jitter, shake or wobble to any animatable control, even if the control is already animated. Its results are similar to those of the Shake modifier, though it uses a different set of controls that may be more suitable to your needs. Unlike other random modifiers, the Perturb modifier can also be applied to polylines, shapes, grid meshes and even color gradients.

For example, to add camera shake to an existing path, right-click on the crosshair and choose Insert > Perturb, then adjust the Strength down to suit. Alternatively, right-click on the path's Shape Animation control, and apply the wobble to the path's polyline itself (this works best if the polyline has many points, for example, if it has been tracked, or hand-drawn with the Draw Append pencil node). A third option is to insert the modifier onto the Displacement control, which will cause the motion along the path to jitter back and forth without actually leaving the line of the path.

NOTE: Perturb can only add jitter it cannot smooth out existing animation curves.

Controls





Value

The actual contents of this control will depend on what type of control the modifier was applied to. If the Perturb modifier was added to a basic Slider control, the Value will be a slider. If it was added to a Gradient control, then a Gradient control will be displayed here. Use the control to set the default, or center value, for the Perturb modifier to work on.

Jaggedness

(Polylines, Meshes only) This allows you to increase the amount of variation along the length of the polyline or mesh, rather than over time. Increasing this will give a squigglier polyline or more tangled mesh, independent of its movement.

Phase

(Polylines, 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, will produce two completely different results. Click on 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 will speed 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 pixel color or luminosity at a certain position or rectangular region of an image. Think of driving values by the presence or absence of an image or by color values of an image, using colors to create X Y positions on other controls, or measuring graded LUTs to compare values.

It can be applied by right-clicking on a control and selecting Modify with > Probe.

Probe Tab

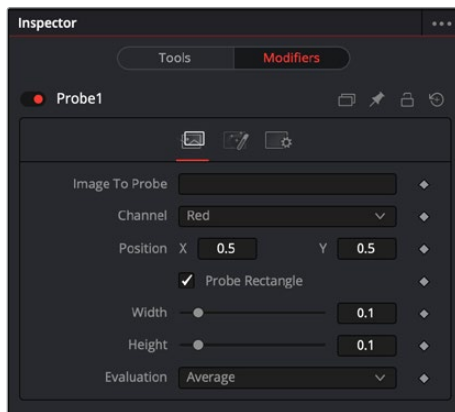


Image to Probe

Drop the node of the image you want to probe into this field.

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 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 only samples 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.

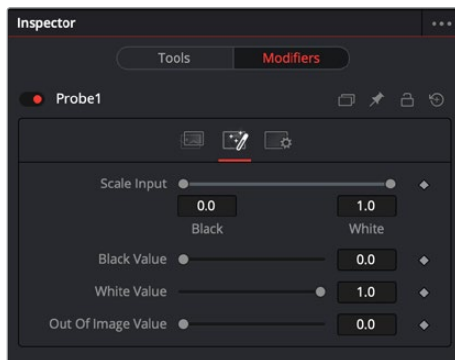
Evaluation

Sets how the pixels inside the rectangle are computed to generate the output value.

Options are:

- 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).

Value Tab



Scale Input Black White

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.

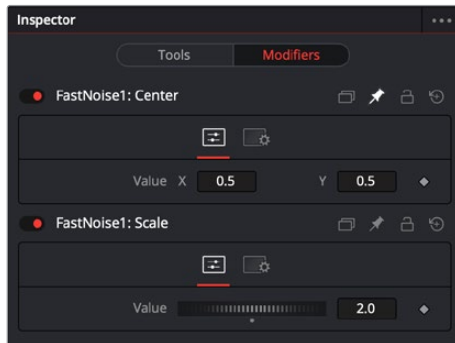
Out of Image Value

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 will not be generated before the entire rectangle is outside the frame boundaries of the image to be probed.

Publish

Only controls that are animated can be available from the Connect To menu option. A non-animated control (static) must be published before it can be connected. Animated controls are automatically published, whereas static controls have to be manually published.

To publish a control, right-click on the control and select Publish from the contextual menu.



Published Value

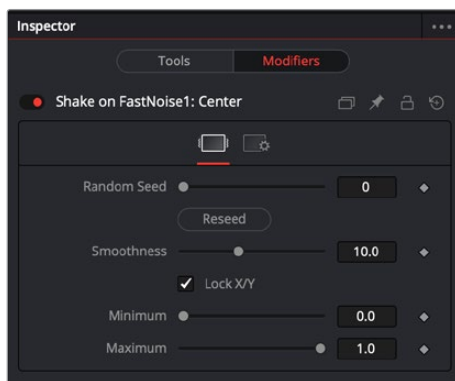
This is obviously dependent on which value is published from which node.

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 control, select Modify With > Shake from the Control's contextual menu. The Shake modifier uses the following controls to achieve its effect. It can be applied by right-clicking on a control and selecting Modify with > Shake.

Shake 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 will always produce 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 is set, the smoother the motion will appear. A value of zero will generate completely random results, with no smoothing applied.

Lock X and Y

This checkbox is used to Unlock the X- and Y-axis, revealing independent slider controls for each axis.

Min and Max

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 low to 0.0 and the high to 1.0. To restrict the motion to a smaller shake in the bottom right corner of the image, set the Min to 0.70 and the Max to 0.90.

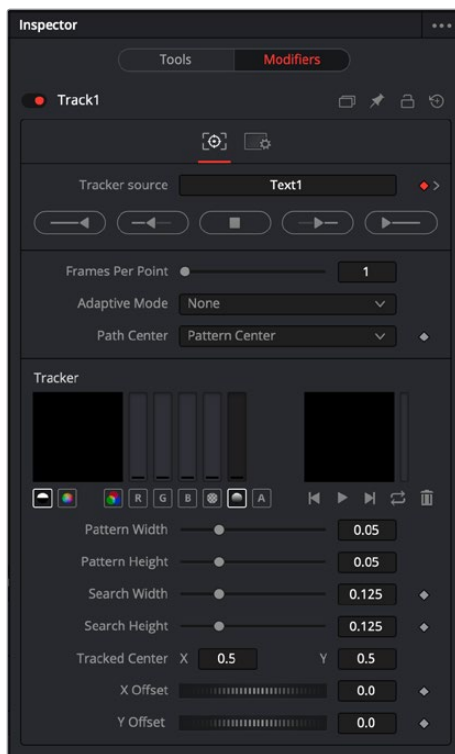
Example

- Create a new comp and add a Text node.
- Type something in the Text node and view it in the large display.
- Add a shake modifier by selecting Modify With > Shake Position from the contextual menu for the Text Center control.
- Switch to the Modifier tab and set the smoothing to 5.0.
- Set the Low to 0.1 and the High to 0.9.
- Go to frame 0 and add a Bezier spline to both the Low and the High controls.
- Advance to frame 90 and set the Low to 0.45 and the High to 0.55.
- Render out a preview and view the results.
- The text should start out by flying all over the screen and then tighten in toward the center of the screen.

Track

Another technique for adding a tracker directly to a control is to apply the Tracker as a modifier. Right-click on the center of the mask and select Ellipse x Center > Modify With > Tracker Position > Steady Position or Unsteady Position. This would add a modifier with a set of controls almost identical to those found in the Tracker node itself.

Tracker Tab



For an in-depth explanation of this node, refer to the Tracker chapter in this manual.

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.
- 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 will default to the output of the Loader). Set a different source image for the Tracker modifier by typing in the name of the node. Alternately, drag and drop the Source node from the Node Editor into the Text Box control.

Example

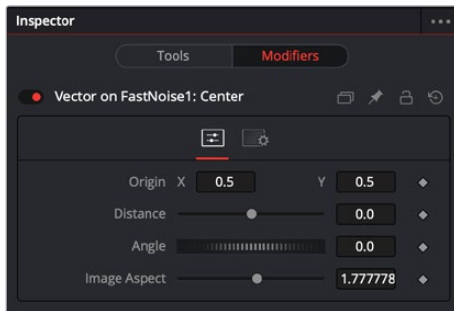
Imagine that you needed to track an actor's eyes so that an unearthly, alien glow could be applied to the eyes. Add a Glow node and then create an Ellipse mask in the shape of the eye. Right-click on the center of that mask and select modify with > Track. Track the actor's eye.

Vector

The Vector 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 on a control and selecting Modify with > Vector.

Controls



Origin

This control is used to represent the position from which the vector's distance and angle values originate.

Distance Slider Control

This control is used to determine the Distance of the vector from the origin.

Angle Thumbwheel Control

This control is used to determine the Angle of the vector relative to the origin.

Image Aspect Slider Control

This 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 Aspect value of 2. The default for this value is taken from the current Frame Format preferences using width/height. It may be required to modify this control to match the current image.

Example

- To illustrate, create a simple node tree consisting of a black background, a Text node and a Merge.
- Right-click on the center of the merge and select Modify With > Vector. This will add a Vector Control header under the small view. Expand it by clicking on the gold triangle. Drag the Distance control to distance the text from the vector origin.
- Drag the Angle thumbwheel to rotate the text around the vector origin.
- Add a path to the vector origin by right-clicking on the Origin control and selecting the Path option from the contextual menu. Verify that the current frame is set to frame 0 (zero) and drag the Vector Origin crosshair to the bottom left corner of the screen.
- Right-click on the Vector Angle thumbwheel and select Bezier Spline to animate this control.
- Set the Angle thumbwheel to a value of 10.
- Advance 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. This will cause the text to orbit around the path just created.

XY Path

The XY Path type uses a separate spline for the position along the X-axis, as it does for the position along the Y-axis.

To animate a coordinate control using an XY path, right-click on 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, 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 controls position will be keyframed on these splines. The advantage to the XY path is that it becomes very easy to work with motion along an individual axis.

Controls



X Y Z Values

These reflect the position of the animated control.

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 if or if not the actual path is displayed in the views.

Switching Default Paths

Change the Default Path type used when animating a coordinate control to XY path (if this is the preferred type of animation). Open the Global preferences and look under the default's tab for the drop-down list in the default animate section labeled Point With. 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 will be used instead of a Displacement path.

Chapter 56

VR Nodes

This chapter details Virtual Reality (VR) nodes available in Fusion.

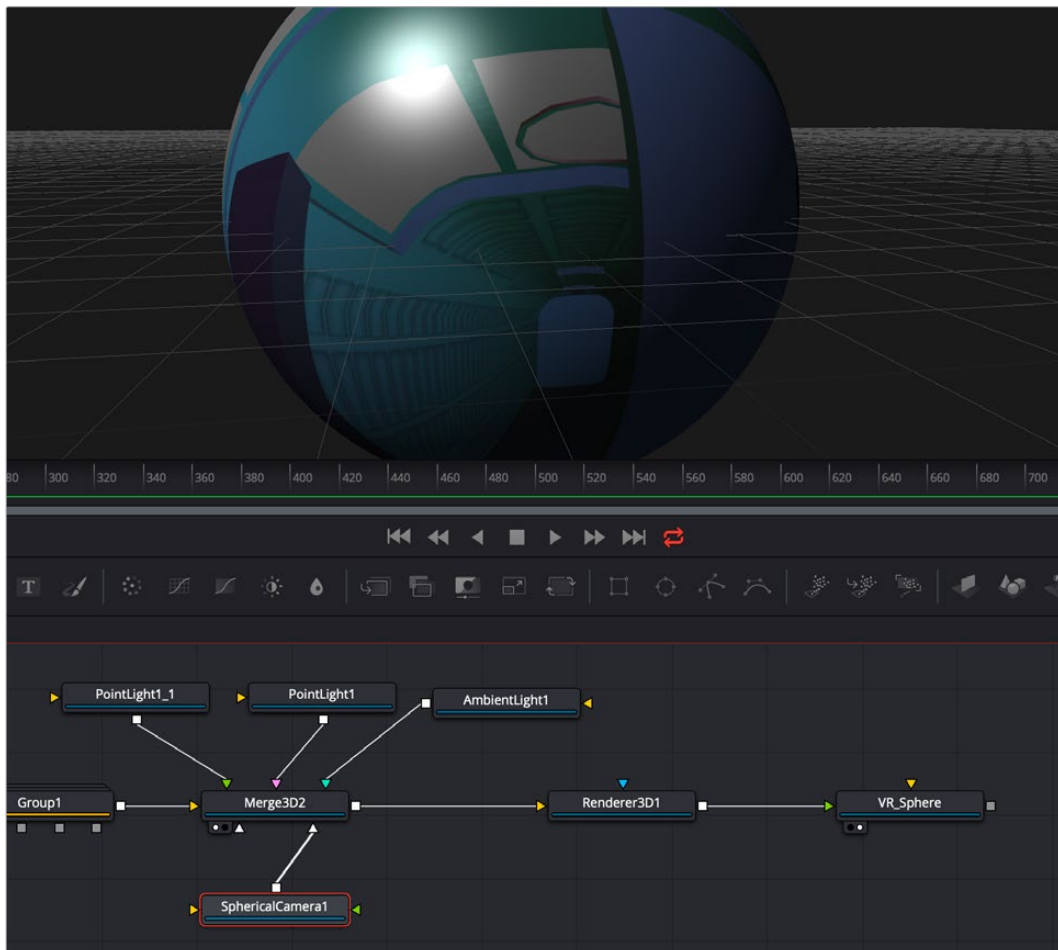
Contents

VR Nodes	1210
Spherical Stabilizer	1211
Spherical Camera	1212
LatLong Patcher	1214
PanoMap	1215

VR Nodes

360° (spherical) video, often described as Virtual Reality or VR, has been created and fixed in Fusion for over a decade. Dome productions, planetariums and other special-venue productions utilize the flexibility of Fusion and its 3D system to produce and deliver special content.

360° video is often represented with an equirectangular (lat-long) format, similar to how the globe is represented by a world map, with the poles at the top and bottom edges of the image and the forward view point at the center. Stereo VR can be created from two stacked lat-long images, one for each eye. VR video can also be created from 3D CGI scenes.




Fusion supports a number of common spherical image formats, and can easily convert between them.

- **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** 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** is a single 2:1 image in an equirectangular mapping.

Fusion 9 has built-in support for VR headsets like the Oculus Rift and HTC Vive. Fusion will display both spherical video and live 3D scenes from the comp directly to the headset.


Fusion's "Fix it in Post" tools for VR make it easy to do several important tasks that are common in these types of productions.

Spherical Stabilizer



VR live action is often shot with a handheld camera, so the footage can be shaky. The Spherical Stabilizer node will automatically identify and track visible features in the footage, then analyze 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. The node requires images in a spherical layout, which can be any of LatLong (2:1 equirectangular), Horizontal/Vertical Cross, or Horizontal/Vertical Strip.

Controls



Reject Outliers to Dominant Motion While Tracking

With this control ticked (the default setting), features that move contrary to the majority of other features will be ignored. This helps ignore the movement of subjects in the shot, preferring stable and consistent markers from the surrounding environment.

Track Controls

These initiate tracking and analysis of the shot. Note that the reference frame used for stabilization is set to the first frame tracked.

- Track Backwards from end frame starts tracking backwards from the end of the current render range.
- Track Backwards from current time starts tracking backwards 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 will cause the Track Controls to discard any previous tracking results and replace them with the newly-created track.
- Append will add 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).


Still – Smooth

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 in order to increase viewer comfort (Smooth mode, 1.0). This control allows either option, or anything in between.

Offset Rotation

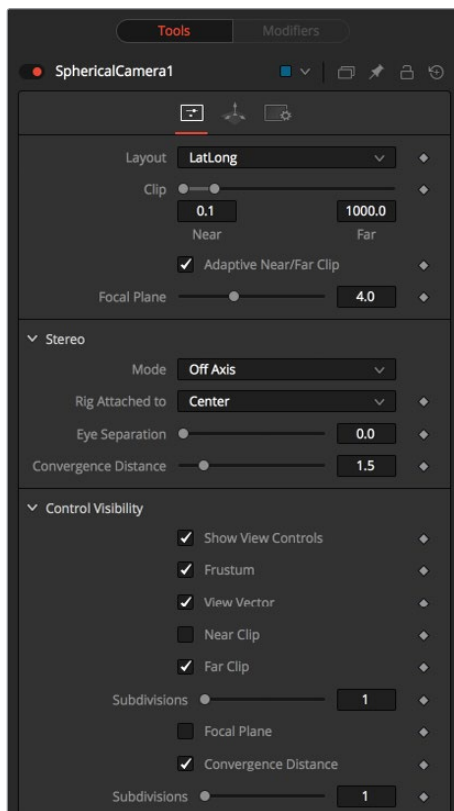
Often a shot is not completely level and needs the horizon to be realigned, or perhaps a desired pan should be reintroduced after fully stabilizing the shot. The Offset Rotation controls allow additional manual control of the Stabilizer's rotation of the footage, for pitch/tilt (X), pan/yaw (Y), and roll (Z), respectively. Rotation is always performed in that order, i.e., XYZ.

Spherical Camera



The Spherical Camera allows the 3D Renderer 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 is used as the size of each square cube face, so the resulting image may be a multiple of this size horizontally and vertically.

Controls



Layout

- **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** 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** is a single 2:1 image in equirectangular mapping.

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 which 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 object more than 20 units distant from the camera will be invisible to the camera. A near clipping plane of 0.1 means that any object closer than 0.1 units will also be invisible.

Adaptively Adjust Near/Far Clip

When selected, the Renderer will automatically adjust 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 only appears 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 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. 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

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 will be 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 on screen 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 Plane of Focus according to the respective 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.

LatLong Patcher



Equirectangular stitched images often need patches, paint work, or other VFX applied. The LatLong 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. This gives the ability to quickly pick a section of the spherical image to patch or paint, then apply it back to the original image. Note that matching rotations will be 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.

Controls



Mode

- **Extract** will pull a de-warped 90 degree square image from the equirectangular image.
- **Apply** will warp and merge 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.


Rotation

These will rotate the spherical image around each of the X, Y, and Z axes, offering independent control over pitch/tilt, pan/yaw, and roll, respectively.

Rotation Order


This chooses the ordering of the rotations around each axis. For example, XYZ will rotate first around the X axis (pitch/tilt), then around the Y axis (pan/yaw), then around the Z axis (roll). Any of the six possible orderings can be chosen.

PanoMap



Converts images from one spherical layout to another, such as cubemap to equirectangular formats. The node can also do rotations of the spherical images when converting.

Controls



From and To

- **Auto** will detect the incoming image layout from the metadata and image frame aspect.
- **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** 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** is a single 2:1 image in equirectangular mapping.

Rotation

These will rotate the spherical image around each of the X, Y, and Z axes, offering independent control over pitch/tilt, pan/yaw, and roll, respectively.

Rotation Order

This chooses the ordering of the rotations around each axis. For example, XYZ will rotate first around the X axis (pitch/tilt), then around the Y axis (pan/yaw), then around the Z axis (roll). Any of the six possible orderings can be chosen.