PART 6

Fusion

Chapter 44

Introduction to Compositing in Fusion

This introduction has been designed specifically to help users who are new to Fusion get started learning this exceptionally powerful environment for doing visual effects and motion graphics, now available right from within DaVinci Resolve.

The integration of Fusion into DaVinci Resolve is a significant engineering project so you can expect rapid changes to the user interface and feature set. That being the case, it's likely the user manual will be occasionally behind the engineering team, so there could be some UI screen shots and descriptions that are a little out of date. We will continue to update the manual at each release.

Contents

DaVinci Resolve, Now With Fusion Inside

The Fusion page is intended, eventually, to be a feature-complete integration of Blackmagic Design Fusion, a powerful 2D and 3D compositing application with over thirty years of evolution serving the film and broadcast industry, creating effects that have been seen in countless films and television series.

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

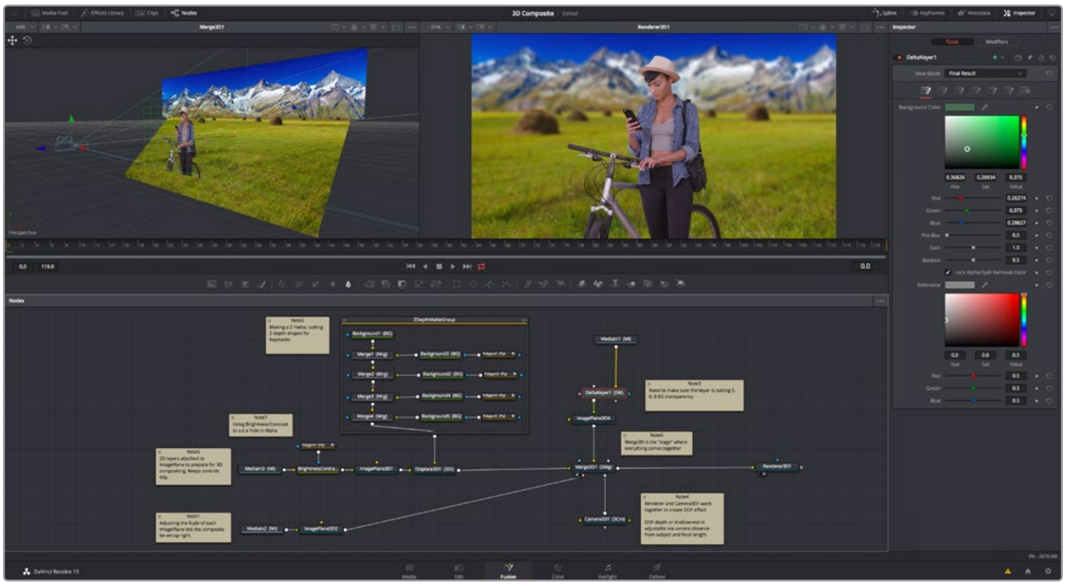

The Fusion page in DaVinci Resolve 15, showing Viewers, the Node Editor, and the Inspector

How Do I Use the Fusion Page?

At its simplest, you need only park the playhead over a clip you want to apply effects to, click the Fusion page button, and your clip is immediately available as a MediaIn node in the Fusion page, ready for you to add a variety of stylistic effects, paint out an unwanted blemish or feature, build a quick composite to add graphics or texture, or accomplish any other visual effect you can imagine, built from the Fusion page's toolkit of effects.

Alternately, you have the option of editing together all of the clips you want to use in your composition in the Edit page, superimposing and lining up every piece of media you'll need with the correct timing, before selecting them and creating a Fusion clip, which functions as a single item in the Edit page Timeline, but when seen in the Fusion page reveals each piece of media you've assembled as a fully built Fusion composition, ready for you to start adding nodes to customize for whatever effect you need to create.

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

How Do Fusion Effects Differ from Edit Page Effects?

While there are many effects you can create in the Edit page, the Fusion page's node-based interface has been designed to let you go deep into the minutiae of a composition to create sophisticated 2D and 3D effects with precise control and endless customization. If you like nodes for color correction, you'll love them for effects.

What Kinds of Effects Does the Fusion Page Offer?

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

3D Compositing

The Fusion page has powerful 3D nodes that include modeling text and simple geometry right inside the Fusion page. In Resolve Studio, this includes the ability to import 3D models in a variety of formats (that functionality has not yet been incorporated into DaVinci Resolve, but it's coming). 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.

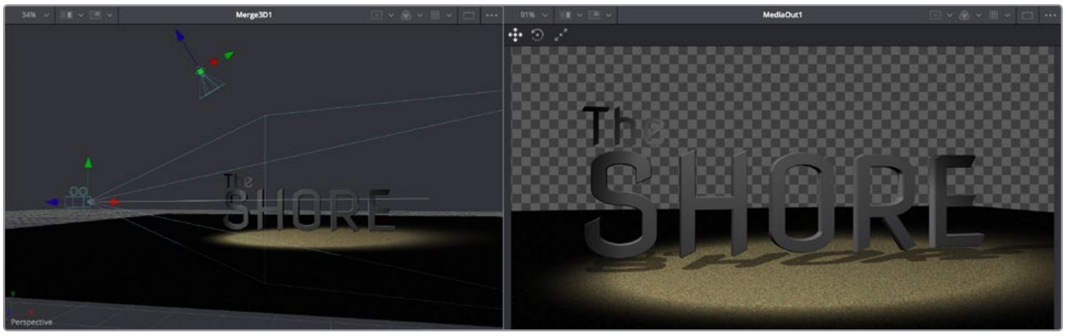

A 3D scene with textured 3D text, created entirely within the Fusion page

Particles

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

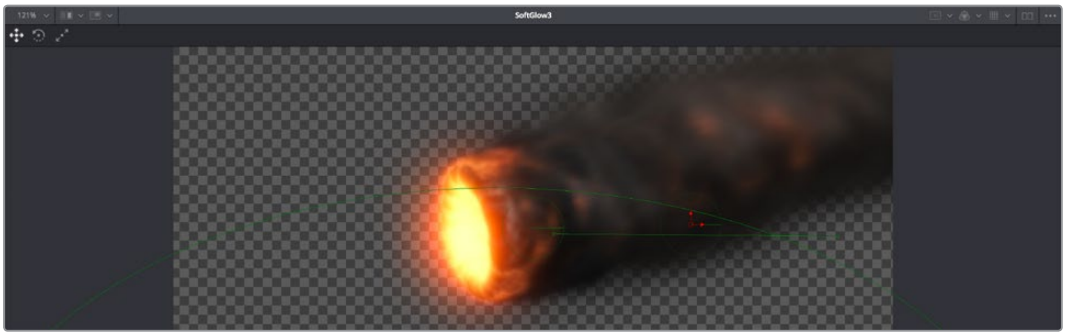

A 3D particle system, also created entirely within the Fusion page

Text

The Text tools in the Fusion page are exceptional, giving you layout and animation options that DaVinci Resolve has never had before, in both 2D and 3D. Furthermore, these Text tools have been incorporated into the Edit page as Fusion Titles, which are compositions saved as macros with published controls, right in Fusion, that expose those controls in the Edit page Inspector for easy customization and control, even if you're working with people who don't know Fusion.

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

And Lots More

The list goes on, with Stereo and VR adjustment nodes, Planar Tracking, Deep Pixel nodes for re-compositing rendered 3D scenes using Auxiliary Channel data, powerful Masking and Rotoscoping nodes, and Warping effects, the Fusion page 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.

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

Go on, give it a try, and remember that you have the chapters of this PDF to refer to, which includes Chapter 49, "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 45

Using the Fusion Page

This chapter provides an orientation on the user interface of the Fusion page, 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 Page User Interface

If you open everything up at once, the Fusion page is divided into four principal regions designed to help you make fast work of node-based compositing. The Media Pool and Effects Library share the area found at the left, the Viewer(s) are at the top, the Work Area is at the bottom, and the Inspector at the right. All of 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, all without leaving DaVinci Resolve.

The Fusion user interface shown completely

However, the Fusion page 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'll probably not 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 page UI, within which you can expose the three main panels used to construct compositions and edit animations in the Fusion page. These are the Node Editor, the Spline Editor, and the Keyframes Editor. By default, the Node Editor is the first thing you'll see, and the main area you'll be working within, but it can sit side-by-side with the Spline Editor and Keyframes Editor as necessary, 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 Keyframes Editor

Interface Toolbar

At the very top of the Fusion page is a toolbar with buttons that let you show and hide different parts of the Fusion page user interface. 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 bar with or without text labels.

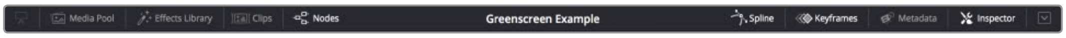

The UI toolbar of the Fusion page

These buttons are as follows, from left to right:

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

- **Spline:** Opens and closes the Spline Editor, where you can edit the curves that interpolate keyframe animation 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 the Fusion page, as well as to trim their In and Out points. A spreadsheet can be shown and hidden within which you can numerically edit keyframe values for selected effects.
- **Metadata:** Hides or shows the Metadata Editor, which lets you read and edit the available clip and project metadata associated with any piece of media within a composite.
- **Inspector:** Shows or hides the Inspector, which shows you all the editable parameters and controls that correspond to selected nodes in the Node Editor. You can show the parameters for multiple nodes at once, and even pin the parameters of nodes you need to continue editing so that they're displayed even if those nodes aren't selected.
- **Inspector Height 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 lets you have 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 DaVinci Resolve interface using the pointer, or use a keyboard shortcut to "select" a particular panel (such as in the Edit page), 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. A highlight appears at the top edge to show you which panel has focus so that you can keep track of which part of the current page is taking precedence, and you can switch focus as necessary to do what you need to do.

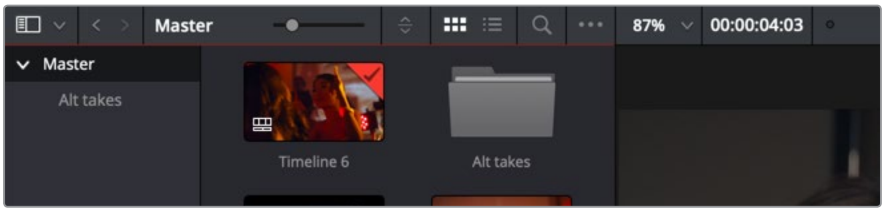

The focus indicator shown at the top edge of the Media Pool, shown next to a Viewer that doesn't have focus

Viewers

The Viewer area can be set to display either one or two viewers at the top of the Fusion page, 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 while seeing its effect on the overall composition in the other

Ordinarily, each viewer shows 2D nodes from your composition as a single image. However, when you're viewing a 3D node, you have the option to set that viewer to one of several 3D views, 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.

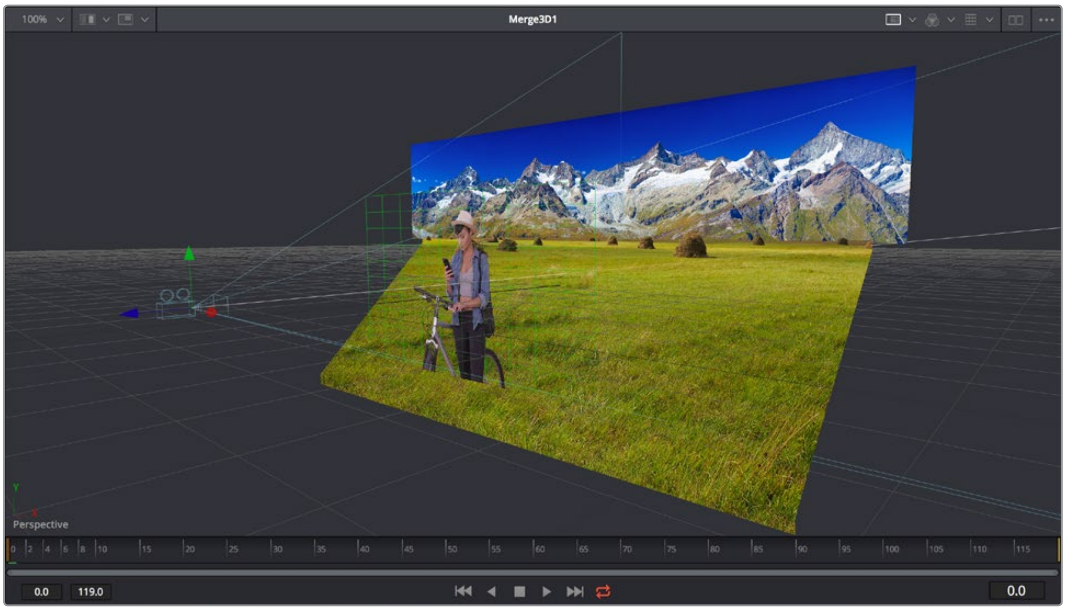

Loading a 3D node into a Viewer switches on a perspective view

TIP: In Perspective view, you can hold the Option key down and drag in the Viewer to pivot the view around the center of the world. All other methods of navigating viewers work the same.

Similarly to the Color page, the Viewers have a variety of capabilities you can use to compare and evaluate what you're looking at, except that there are many more options that are 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 Shift and Command down 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 the Command key down and use your pointer's scroll control to resize the Viewer.
- Hold the middle mouse button down, 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 the Option key down and drag to spin the stage around.

Loading Nodes Into Viewers

When you first open the Fusion page, the output of the current empty composition (the MediaOut1 node) is usually showing in Viewer2. If you're in Dual-viewer mode, Viewer1 remains empty until you assign a node to one of them.

To load specific nodes into specific viewers:

- Hover the pointer over a node, and click one of two buttons that appear at the bottom-left of the node.
- Click once to select a node, and press 1 (for the left Viewer) or 2 (for the right Viewer).
- Right-click a node and choose View On > None/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, 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 grading done in the Color page, since the Fusion page comes before the Color page in the DaVinci Resolve image processing pipeline. However, if you're working on clips that have been converted to linear color space for compositing, you may find it desirable to composite and make adjustments to the image relative to a normalized version of the image that appears close to what the final will be, and 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. The rest of this menu shows all LUTs from the /Library/Application Support/Blackmagic Design/DaVinci Resolve/LUT/VFX IO/ directory (on macOS) to use for viewing.
- **Option menu:** This menu contains various settings that pertain to the Viewer in the Fusion page.
  - **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.

Time Ruler and Transport Controls

The Time Ruler, located beneath the Viewer area, shows the frame range of the current clip or composition. However, the duration of this range depends on what's currently selected in the Timeline:

- If you've selected a clip, then the Time Ruler displays all source frames for that clip, and the current In and Out points for that clip define the "render range," or the range used in the Timeline and thus available in the composition by default. All frames outside of this range constitute the heads and tails of that clip that are unused in the edited Timeline.

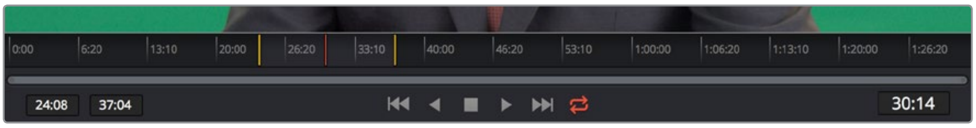

The Time Ruler displaying ranges for a clip in the Timeline via yellow marks (the playhead is red)

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

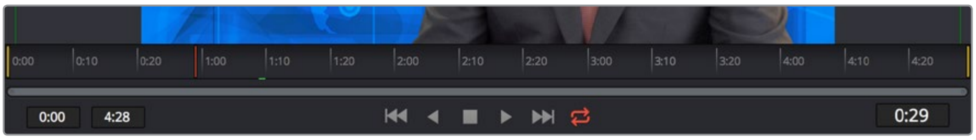

The Time Ruler displaying ranges for a Fusion clip in the Timeline

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

When setting ranges and entering frame numbers to move to a specific frame, numbers can be entered in 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 the Fusion page 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

The Time Ruler uses two different frame ranges, one for the duration of the entire clip or composition, and a Render range that currently determines either the duration of the current clip that appears within the Timeline, or the range of frames to cache in memory for previews.

Composition Start and End Range

The Composition Start and End range is simply the total duration of the current composition.

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.

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 the Command key down 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.

The Render Start and Render End time fields

Changing the Time Display Format

By default, all time fields and markers in the Fusion Page 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 > Fusion Settings.
2. When the Fusion settings dialog opens select the 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 looking at a clip with really long handles 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 six transport controls underneath the Time Ruler, including Composition first frame, Play Reverse, Stop, Play Forward, Composition last frame, and Loop.

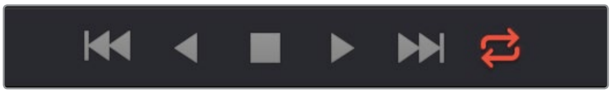

The Fusion page transport controls

Navigation Shortcuts

Many of the standard transport control keyboard shortcuts also work in the Fusion page, but some are specific to Fusion's particular needs.

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

- **Spacebar:** Toggles forward playback on and off.
- **JKL:** Basic JKL playback is supported, including J to play backward, K to stop, and L to play forward.
- **Back Arrow:** Moves 1 frame backward.
- **Forward Arrow:** Moves 1 frame forward.
- **Shift-Back Arrow:** Moves to the clip's Source End frame.
- **Shift-Forward Arrow:** Moves to the clip's Source Start frame.
- **Command-Back Arrow:** Jumps to the clip's In point.
- **Command-Forward Arrow:** Jumps to the clip's Out point.

Real Time Playback Not Guaranteed

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

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).

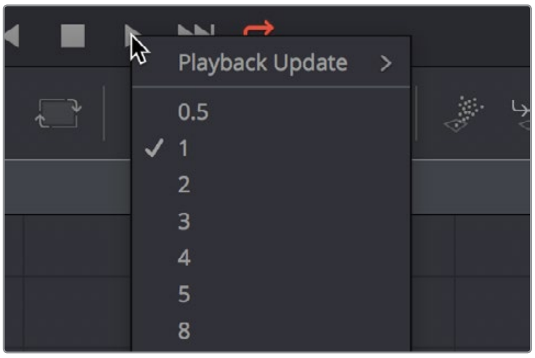

Right-click the Play Forward or Play Backward buttons to choose a frame increment in which to move the playhead

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.

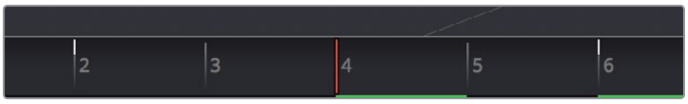

The Time Ruler displaying keyframe marks

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

Right-clicking anywhere in the transport control area lets you turn on and off Fusion page-specific quality controls, which lets you either enable high-quality playback at the expense of greater processing times, or enter various proxy modes that temporarily lower the display quality of your composition in order to speed processing as you work. Rendering for final output is always done at the highest quality, regardless of these settings.

High Quality

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

Motion Blur

Turning Motion Blur off 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.

Proxy

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

Auto Proxy

A draft mode to speed processing while you're building your composite. Turning on Auto Proxy reduces the resolution of the image while you click and drag 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 Fusion Settings, found in the Fusion menu.

Selective Updates (Available in Fusion Settings)

There are three options:

- **Update All:** Forces all of the nodes in the current node tree to render. This is primarily used when you want to update all of the thumbnails displayed in the Node Editor.
- **Selective:** (the default) Causes only nodes that directly contribute to the current image to be rendered, so named because only selective nodes are rendered.
- **No Update:** 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 the Fusion page 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 Fusion Memory Limits setting found in the Configuration panel of the System 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 the Fusion page.

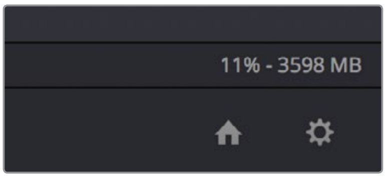

Percentage of the RAM cache that's been used at the bottom right of the Fusion page

Displaying Cached Frames

All frames that are 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 pay 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.

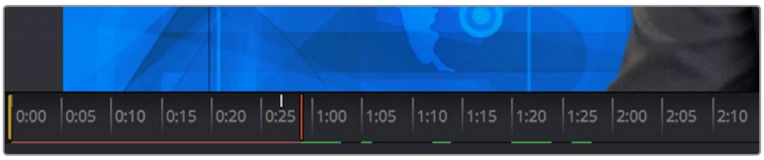

A red line indicates that cached frames from a different quality or proxy setting are being preserved

There's one exception to this, however. When you cache frames at the High Quality setting and you then turn High Quality off, 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 six sections that group commonly used nodes together. As you hover the pointer over any button, a tooltip shows you that node's name.

- **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 re-assigning 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, but even 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 the Fusion page, 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.

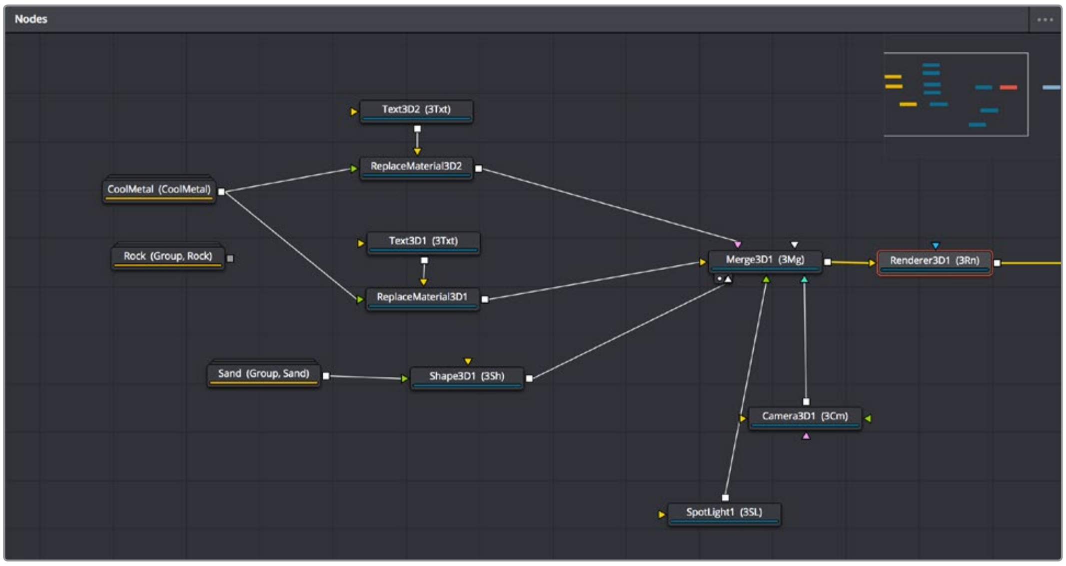

The Node Editor displaying a node tree creating a composition

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, but if there are no selected nodes or multiple selected nodes, then new nodes are added as disconnected from anything else.

Methods of adding nodes:

- 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.

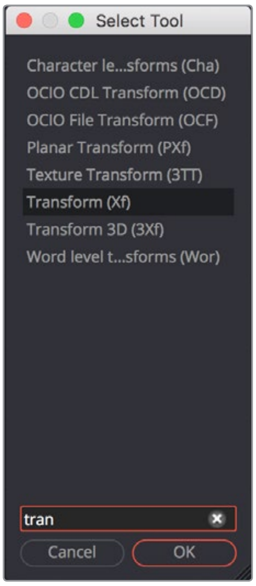

The Select Tool dialog lets you find any node quickly if you know its name

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").

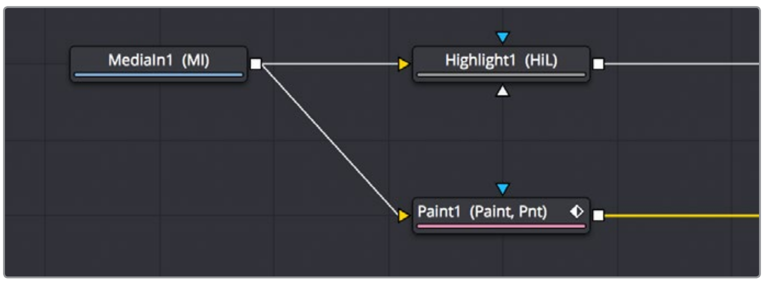

One node branching to two to split the image to two operations

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

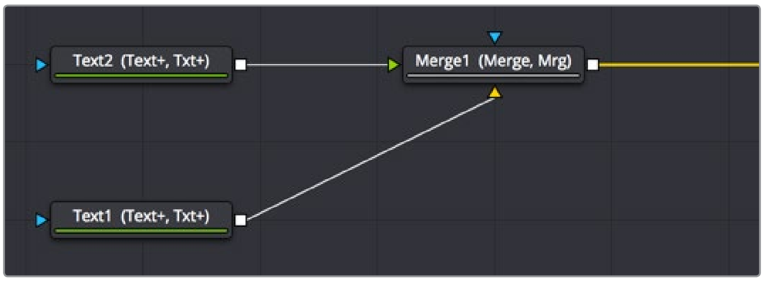

Two nodes being merged together into one to create a composite

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.

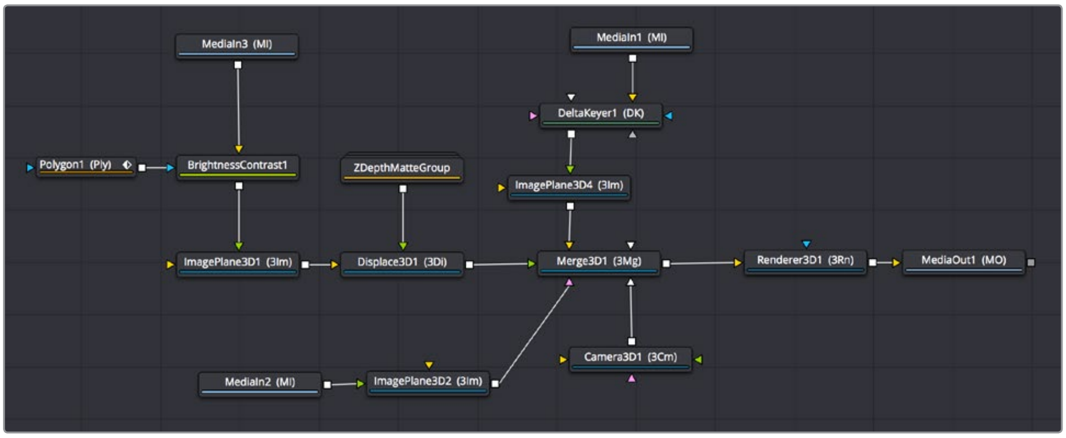

Nodes can be oriented in any direction; the input arrows let you follow the flow of image data

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-hand 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.

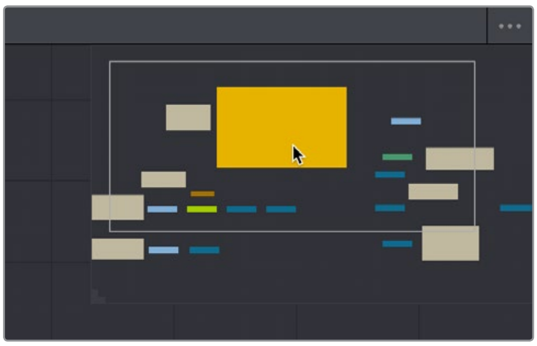

The Navigator pane for accessing offscreen parameters or tools

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 Shift and Command down and drag the Node Editor to pan.
- Press the Middle and Left buttons simultaneously and drag to resize the Node Editor.
- Hold the Command key down 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, the Fusion page 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 the Fusion page, immediately above the Resolve Page bar, shows you a variety of up-to-date information about things you're selecting and what's happening in the Fusion page. 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 appears in this area as you work.

The Status bar under the Node Editor showing you information about a node under the pointer

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.

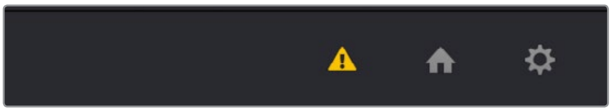

A notification that there's a message in the Console

Effects Library

The Effects Library on the Fusion page shows all of the nodes and effects that are available in the Fusion page, including effects that come with DaVinci Resolve and third-party OFX, if available. While the toolbar shows many of the most common nodes you'll be using in any composite, the Effects Library contains every single tool available in the Fusion page, 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.

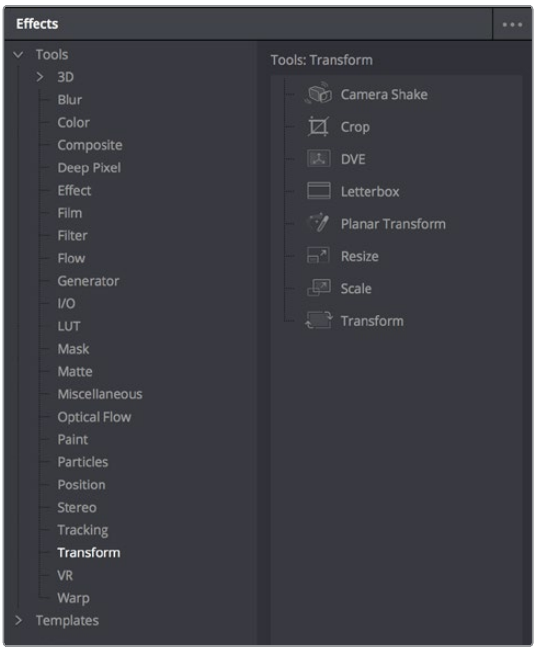

The Effects Library with Tools open

The hierarchical category browser of the Effects Library is divided into two sections. The Tools section contains every node that represents an elemental image processing operation in the Fusion page. The Templates section contains a variety of additional compositing functions, as well as libraries of content such as Lens Flares, Backgrounds, Generators, Particle Systems, Shaders (for texturing 3D objects) and other resources for use in your composites.

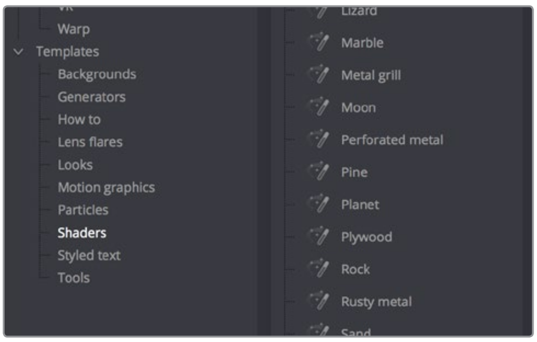

The Templates section of the Effects Library

Similar to the Media Pool, the Effects Library's bin list can be made full-height or half-height using a button at the far left of the UI toolbar. Additionally, an Option menu in the Effects Library gives you access to additional options and commands.

Inspector

The Inspector is a panel on the right side of the Fusion page 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 shows you the parameters of selected nodes. The Modifiers panel shows you different things for different nodes. For all nodes, it shows you the controls for Modifiers, or adjustable expressions, that you've added to specific parameters to automatically animate them in different ways. 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 lets you customize what kind of randomness is being added.

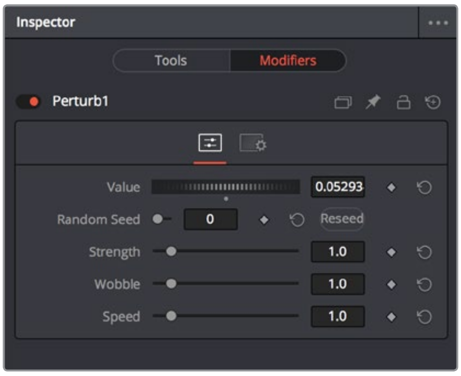

The Modifier panel showing a Perturb modifier

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.

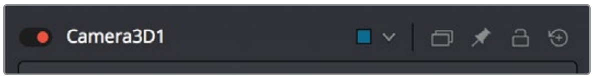

Common Inspector Controls

- **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 Fusion page Inspector is also capable of simultaneously displaying all parameters for multiple nodes you've selected in the Node Editor. Furthermore, a Pin button in the title bar of each node's parameters lets you "pin" that node's parameters into the Inspector so that they remain there even when that node is de-selected, 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.

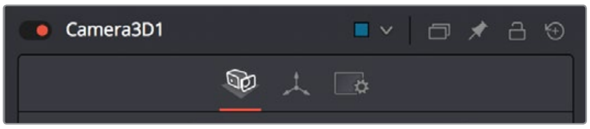

Nodes with several tabs' worth of parameters

Keyframes Editor

The Keyframes Editor displays each MediaIn 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 can use the Keyframes Editor to trim, extend, or slide MediaIn and effects nodes, or to adjust the timing of keyframes, which appear superimposed over each effect node unless you open them up into their own editable track.

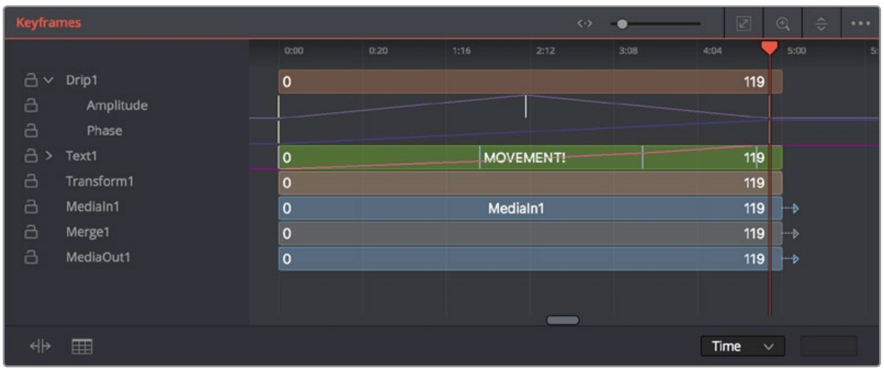

The Keyframes Editor is used to adjust the timing of clips, effects, and keyframes

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 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 tracks in various ways.
- An Option menu provides access to many other ways of filtering tracks and controlling visible options.

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

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

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

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

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

At the bottom-right, the Time/Toffset/Tscale pop-up menu and value fields let you numerically alter the position of selected keyframed either absolutely, relatively, or to a scale.

Adjusting Clip Timings

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

While much of this could be done in the Timeline prior to creating a Fusion clip that contains several MediaIn nodes, the Keyframes Editor also lets you adjust the timing of clips that you've added from directly within the Fusion page, as well as generators and 3D objects, which never originally appeared in the Edit page Timeline at all.

Adjusting Effect Timings

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

Adjusting Keyframe Timings

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

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

Keyframe tracks exposed

Keyframe Editing Essentials

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

Methods of adjusting keyframes:

- You can click on a single keyframe to select it.
- You can drag a bounding box over a series of keyframes to select them all.
- You can drag keyframes left and right to reposition them in time.
- You can right-click one or more selected keyframes and use 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, 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.

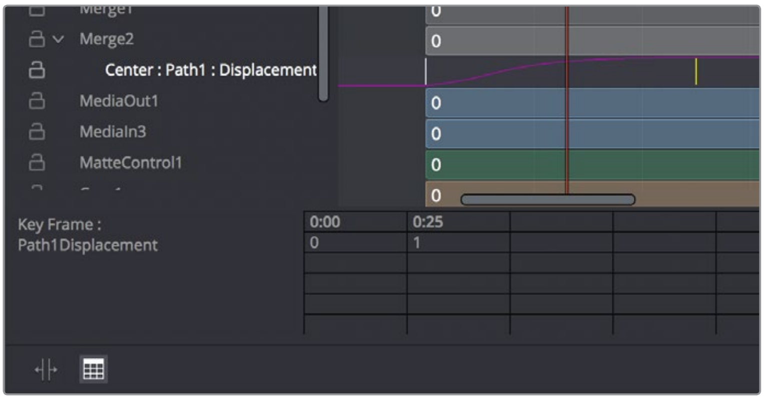

Editing keyframes in the Spreadsheet

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 and value 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.

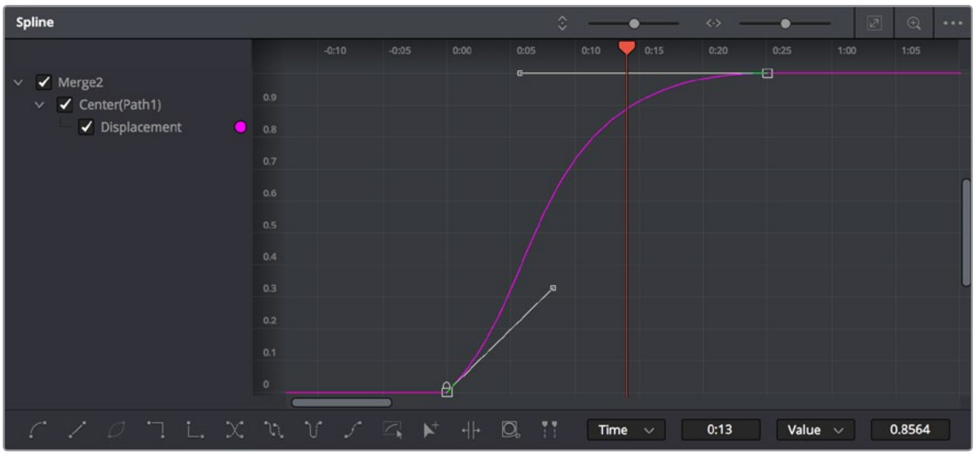

The Spline Editor is divided into the Zoom controls at top, Parameter list at left, and toolbar

Spline Editor Control Summary

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

- 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 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 layers in various ways.
- An Option menu provides access to many other ways of filtering layers 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.

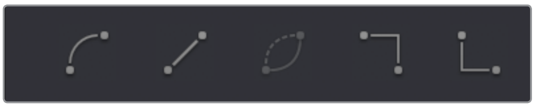

Control point interpolation controls

- **Smooth:** Creates automatically adjusted Bezier curves to create smoothly interpolating animation.
- **Flat:** Creates linear 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 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 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.

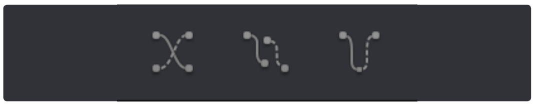

Spline Loop modes

- **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, click again to de-select it. 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), 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.

Thumbnail Timeline

Hidden by default, the Thumbnail timeline can be opened by clicking the Clips button in the UI toolbar, and appears underneath the Node Editor when it's open. The Thumbnail timeline shows you every clip in the current Timeline, giving you a way to navigate from one clip to another when working on multiple compositions in your project, and providing an interface for creating and switching among multiple versions of compositions, and resetting the current composition, when necessary.

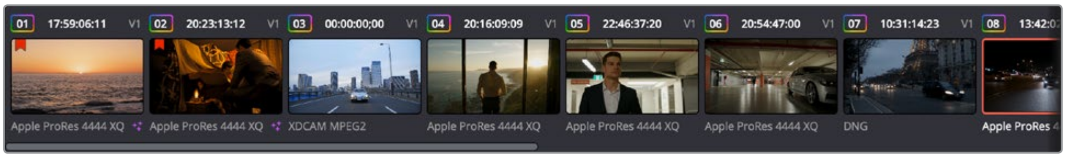

The Thumbnail timeline lets you navigate the Timeline and manage versions of compositions

Right-clicking on any thumbnail exposes a contextual menu

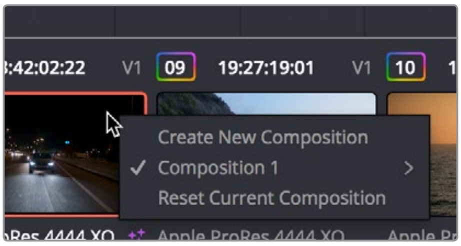

The contextual menu for the Thumbnail timeline

To open another clip:

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

To create and manage versions of compositions:

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

To reset the current composition:

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

To change how thumbnails are identified:

Double-click the area underneath any thumbnail to toggle among clip format, clip name, and a mystery that shall someday be solved by an intrepid team of adventurers embarking on a dangerous quest.

Media Pool

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

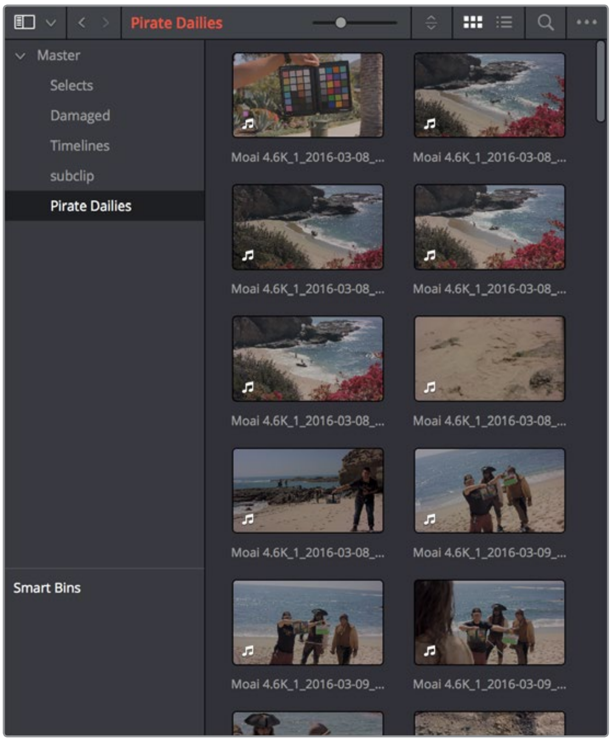

The Media Pool in Thumbnail mode showing video clips

TIP: If you drag one or more clips from the Media Pool onto a connection line between two nodes in the Node Editor so that the connection highlights in blue and then drop them, those clips will be automatically connected to that line via enough Merge nodes to connect them all.

The Bin List

The Bin list at the left, which can be opened and closed, shows a hierarchical list of all bins used for organizing your media as well as your timelines. By default, the Media Pool consists of a single bin, named "Master," but you can add more bins as necessary to organize timelines and clips by right-clicking anywhere in the empty area of the Media Pool and choosing Add Bin. You can rename any bin by double-clicking on its name and typing a new one, or by right-clicking a bin's name and choosing Rename Bin. The Bin list can be hidden or shown via the button at the upper left-hand corner of the Fusion page toolbar.

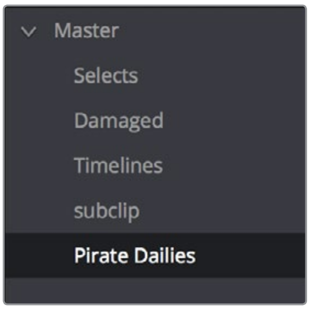

The Bin list

The browser area to the right shows the contents of the currently selected bin in the Bin list. Every clip you've added, every timeline you've created, and every AAF, XML, or EDL file you've imported appears here.

As elsewhere, the Media Pool can be displayed in either Icon or List view. In List view, you can sort the contents by any one of a subset of the total metadata that's available in the Metadata Editor of the Media page. Of particular interest to audio editors are columns for Name, Reel Name, different timecode streams, Description, Comments, Keyword, Shot, Scene, Take, Angle, Circled, Start KeyKode, Flags, and Usage.

For more information on using the myriad features of the Media Pool, see Chapter 11, "Adding and Organizing Media with the Media Pool." In the sections that follow, some key features of the Media Pool are summarized for your convenience.

Importing Media Into the Media Pool on the Fusion Page

While adding clips to the Media Pool in the Media page provides the most organizational flexibility and features, if you find yourself in the Fusion page and you need to quickly import a few clips for immediate use, you can do so in a couple of different ways.

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

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

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

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

For more information on importing media using the myriad features of the Media page, see Chapter 11, "Adding and Organizing Media with the Media Pool."

Bins, Power Bins, and Smart Bins

There are actually three kinds of bins in the Media Pool, and each appears in its own section of the Bin list. The Power Bin and Smart Bin areas of the Bin list can be shown or hidden using commands in the View menu (View > Show Smart Bins, View > Show Power Bins). Here are the differences between the different kinds of bins:

- **Bins:** Simple, manually populated bins. Drag and drop anything you like into a bin, and that's where it lives, until you decide to move it to another bin. Bins may be hierarchically organized, so you can create a Russian dolls nest of bins if you like. Creating new bins is as easy as right-clicking within the Bin list and choosing Add Bin from the contextual menu.
- **Power Bins:** Hidden by default. These are also manually populated bins, but these bins are shared among all of the projects in your current database, making them ideal for shared title generators, graphics movies and stills, sound effects library files, music files, and other media that you want to be able to quickly and easily access from any project. To create a new Power Bin, show the Power Bins area of the Bin list, then right-click within it and choose Add Bin.
- **Smart Bins:** These are procedurally populated bins, meaning that custom rules employing metadata are used to dynamically filter the contents of the Media Pool whenever you select a Smart Bin. Smart Bins are a a fast way of organizing the contents of projects for which you (or an assistant) has taken the time to add metadata to your clips using the Metadata Editor, adding Scene, Shot, and Take information, keywords, comments and description text, and myriad other pieces of information to make it faster to find what you're looking for when you need it. To create a new Smart Bin, show the Smart Bin area of the Bin list (if necessary), then right-click within it and choose Add Smart Bin. A dialog appears in which you can edit the name of that bin and the rules it uses to filter clips, and click Create Smart Bin.

Showing Bins in Separate Windows

If you right-click a bin in the Bin List, you can choose "Open As New Window" to open that bin into its own window. Each window is its own Media Pool, complete with its own Bin list, Power Bins and Smart Bins lists, and display controls.

This is most useful when you have two displays connected to your workstation, as you can drag these separate bins to the second display while DaVinci Resolve is in single screen mode. If you hide the Bin list, not only do you get more room for clips, but you also prevent accidentally switching bins if you really want to only view a particular bin's contents in that window. You can have as many additional Bin windows open as you care to, in addition to the main Media Pool that's docked in the primary window interface.

Filtering Bins Using Color Tags

If you're working on a project that has a lot of bins, you can apply color tags to identify particular bins with one of eight colors. Tagging bins is as easy as right-clicking any bin and choosing the color you want from the Color Tag submenu.

For example, you can identify the bins that have clips you're using most frequently with a red tag. A bin's color tag then appears as a colored background behind that bin's name.

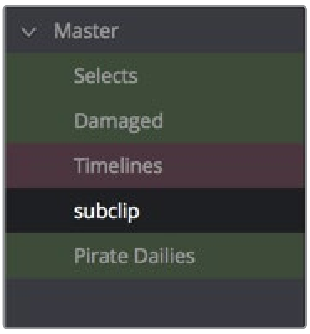

Using color tags to identify bins

Once you've tagged one or more Media Pool bins, you can use the Color Tag Filter pop-up menu (the pop-up control to the right of the Bin List button) to filter out all but a single color of bin.

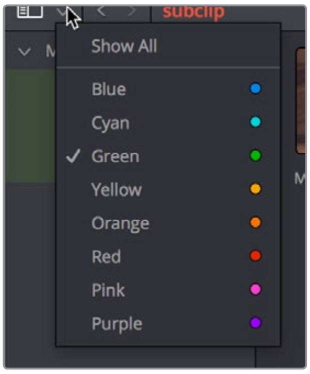

Using Color Tag filtering to isolate the blue bins

To go back to seeing all available bins, choose Show All from the Color Tag Filter pop-up.

Sorting the Bin List

The Bin list (and Smart Bin list) of the Media Pool can be sorted by bin Name, Date Created, or Date Modified, in either ascending or descending order. Simply right-click anywhere within the Bin list and choose the options you want from the Sort by submenu of the contextual menu.

You can also choose User Sort from the same contextual menu, which lets you manually drag all bins in the Bin list to be in whatever order you like. As you drag bins in this mode, an orange line indicates the new position that bin will occupy when dropped.

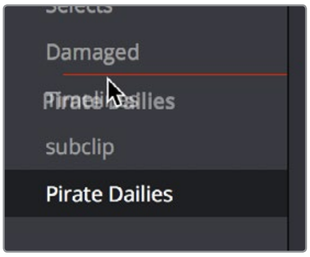

Dragging a bin to a new position in the Bin list in User Sort mode

If you use User Sort in the Bin list to rearrange your bins manually, you can switch back and forth between any of the other sorting methods (Name, Date Created, Date Modified) and User Sort and your manual User Sort order will be remembered, making it easy to use whatever method of bin sorting is most useful at the time, without losing your customized bin organization.

Searching for Content in the Media Pool

An optional Search field can be opened at the top of the Media Pool that lets you quickly find clips by name, partial name, or any of a wide variety of Media Pool metadata.

To search for a clip by name:

1. Select which bin or bins you want to search.
2. Click the magnifying glass button at the upper right-hand corner of the Media Pool.
3. Choose the particular column of information you want to search (or All Fields to search all columns) using the Filter by pop-up menu. Only selected bins will be searched.
4. Type your search string into the Search field that appears. A few letters should be enough to isolate only those clips that have that character string within their name. To show all clips again, click the cancel button at the right of the search field.

TIP: Smart Bins are essentially multi-criteria search operations that scope the entire project at once and are saved for future use.

Taking Advantage of the Media Pool's Usage Column

In List view, the Usage column does not automatically update to show how many times a particular clip has been used. However, you can manually update this metadata by right-clicking within the Media Pool and choosing Update Usage Data from the contextual menu that appears. Afterwards, each clip will display how many times it's been used in this column. Clips that have not been used yet display an x.

The Console

The Console, available by choosing Workspace > Console, is a window in which you can see the error, log, script, and input messages that may explain something the Fusion page 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.

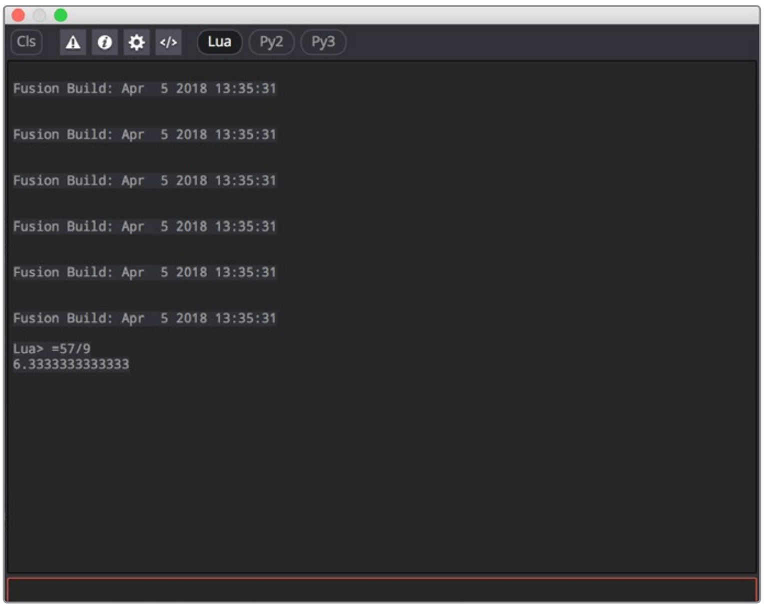

The Console window

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 three 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 do in the Console. More information on scripting will be forthcoming as it becomes available.

Customizing the Fusion Page

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

The Fusion Settings Window

The Fusion page has its own settings window, accessible by choosing Fusion > Fusion Settings. This window has a variety of options for customizing the Fusion experience, which will be documented in more detail at a later date.

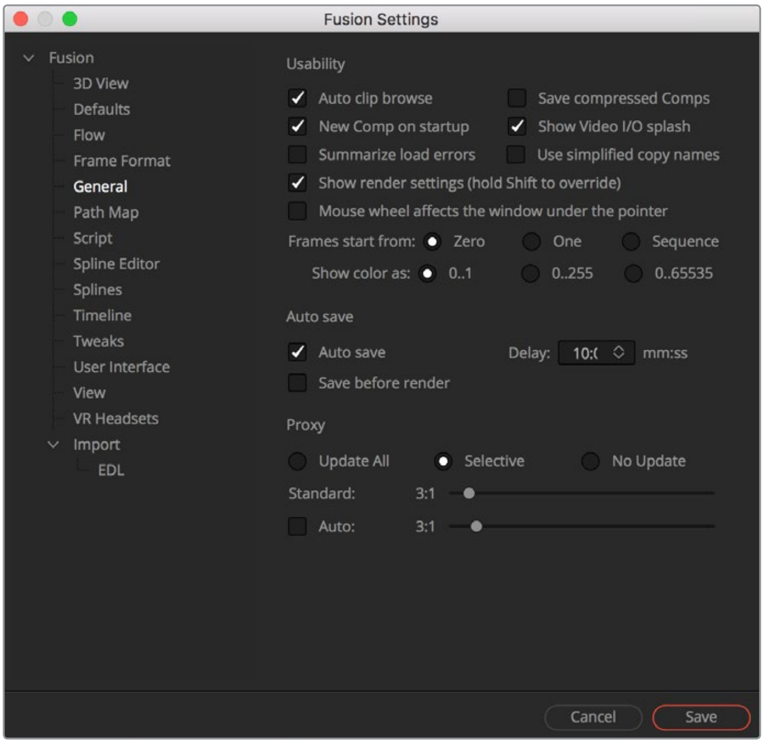

The Fusion Settings window set to the General panel

Showing and Hiding Panels

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

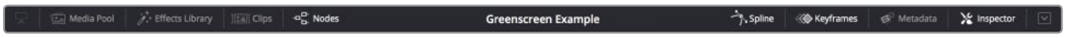

The UI toolbar of the Fusion page

Resizing Panels

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

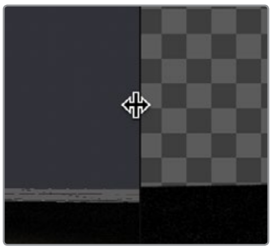

Dragging the edge between two viewers to resize it

Undo and Redo in DaVinci Resolve

No matter where you are in DaVinci Resolve, Undo and Redo commands let you back out of steps you've taken or commands you've executed and reapply them if you change your mind. DaVinci Resolve 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.

Because DaVinci Resolve integrates so much functionality in one application, there are three separate sets of undo "stacks" to help you manage your work.

- The Media, Edit and Fairlight pages share the same multiple-undo stack, which lets you backtrack out of changes made in the Media Pool, the Timeline, the Metadata Editor, and the Viewers.
- Each clip in the Fusion page has its own undo stack so that you can undo changes you make to the composition of each clip, independently.
- Each clip in the Color page has its own undo stack so that you can undo changes you make to grades in each clip, independently.

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 take advantage of this, there are three ways you can undo work to go to a previous state of your project, no matter what page you're in.

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

- Choose Edit > Undo (Command-Z) to undo the previous change.
- Choose Edit > Redo (Shift-Command-Z) to redo to the next change.
- On the DaVinci control panel, press the UNDO and REDO buttons on the T-bar panel.

TIP: If you have the DaVinci control panel, there is one other control that lets you control the undo stack more directly when using the trackballs, rings, and pots. Pressing RESTORE POINT manually adds a memory of the current state of the grade to the undo stack. Since discrete undo states are difficult to predict when you're making ongoing adjustments with the trackball and ring controls, pressing RESTORE POINT lets you set predictable states of the grade that you can fall back on.

You can also undo several steps at a time using the History submenu and window. At the time of this writing, this only works for multiple undo steps in the Media, Edit, and Fairlight pages.

To undo and redo using the History submenu:

1 Open the Edit > History submenu, which shows (up to) the last twenty things you've done.

2 Choose an item on the list to undo back to that point. The most recent thing you've done appears at the top of this list, and the change you've just made appears with a check next to it. Steps that have been undone but that can still be redone remain in this menu, so you can see what's possible. However, if you've undone several changes at once and then you make a new change, you cannot undo any more and those steps disappear from the menu.

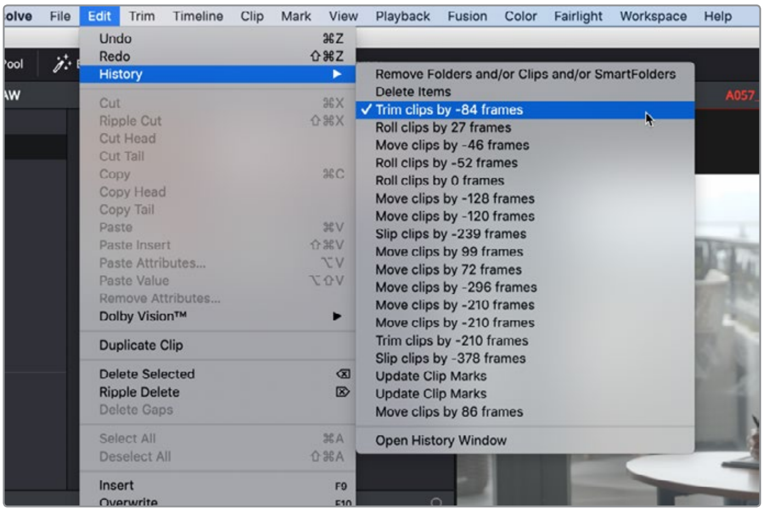

The History submenu, which lets you undo several steps at once

Once you've selected a step to undo to, the menu closes and the project updates to show you its current state.

To undo and redo using the Undo window:

1 Choose Edit > History > Open History Window.

2 When the History dialog appears, click an item on the list to undo back to that point. Unlike the menu, in this window the most recent thing you've done appears at the bottom of this list. Selecting a change here grays out changes that can still be redone, as the project updates to show you its current state.

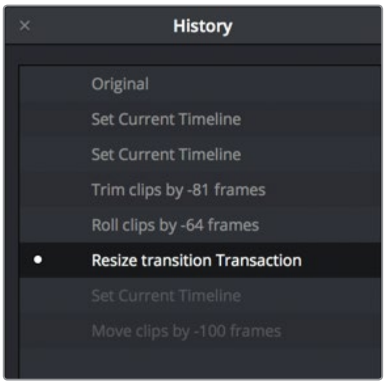

The Undo history window that lets you browse the entire available undo stack of the current page

3 When you're done, close the History window.

Chapter 46

Getting Clips into the Fusion Page

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

Contents

Getting Clips into the Fusion Page

Now that Fusion compositing is integrated into DaVinci Resolve, it's easy to get clips from your edit into the Fusion page to create any number of effects, prior to grading in the Color page. Depending on your needs, there are a few different ways that clips can find their way into the Fusion page.

Working on Single Clips in the Fusion Page

Each visible clip in a timeline appears in the Fusion page as a single MediaIn node connected to a MediaOut node. Clips that aren't visible because they're on lower tracks with fully opaque clips above them are ignored. These very simple default compositions are referred to unofficially in this manual as "single-clip compositions."

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

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

This initial node structure makes it easy to quickly use the Fusion page to create relatively simple effects that are better accomplished using the power of node-based compositing.

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

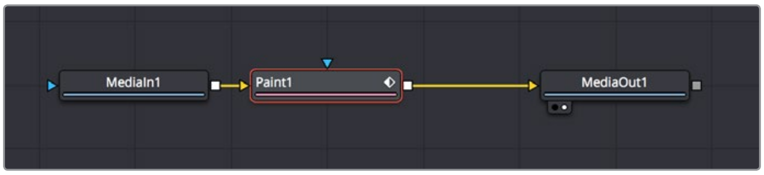

A simple paint effect applied to a shot with no camera motion

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

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

Turning One or More Clips into Fusion Clips

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

To create a Fusion clip:

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

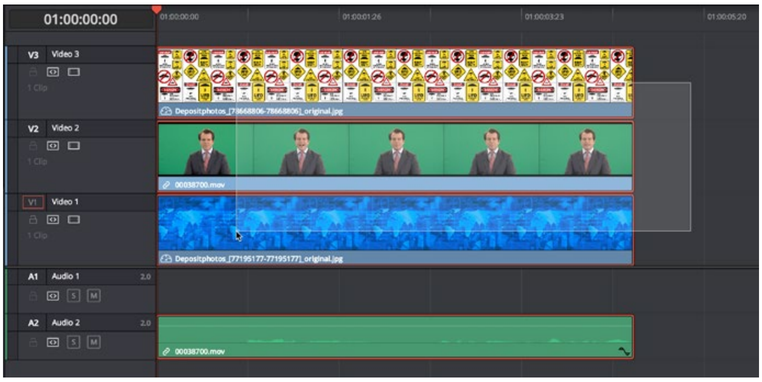

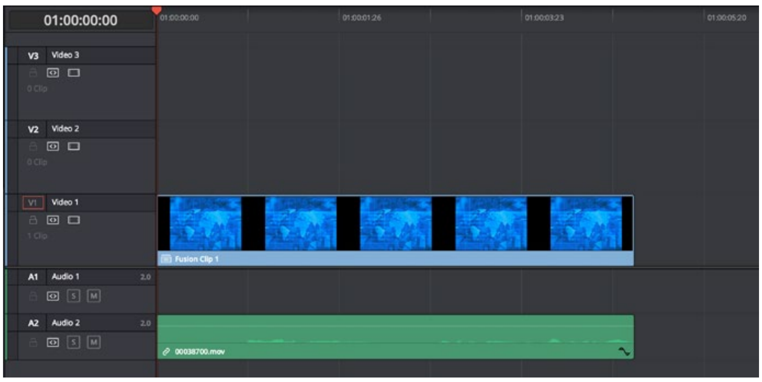

(Top) A stack of clips to use in a composite,
(Bottom) Turning that stack into a Fusion clip in the Edit page

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

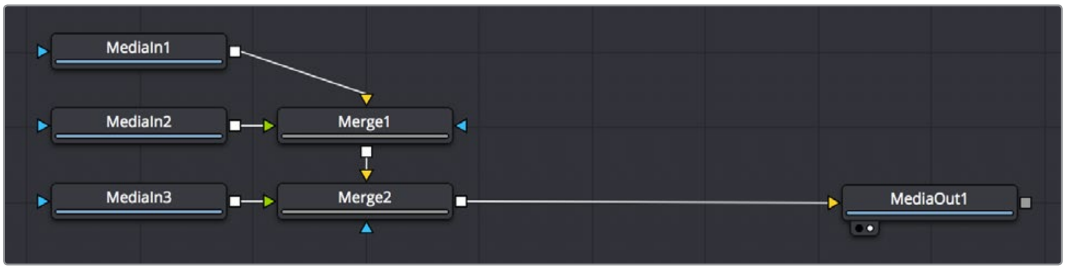

The initial node tree of the three clips we turned into a Fusion clip

Adding Fusion Composition Generators

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

To create a blank Fusion clip in the Edit page:

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

Adding More Clips to Fusion Compositions

No matter which of the above methods you use to move media into the Fusion page for compositing, you'll often find that even though you start out wanting to do something relatively simple, you end up needing to add more clips to create the effect that you really need.

Adding Clips From the Media Pool

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

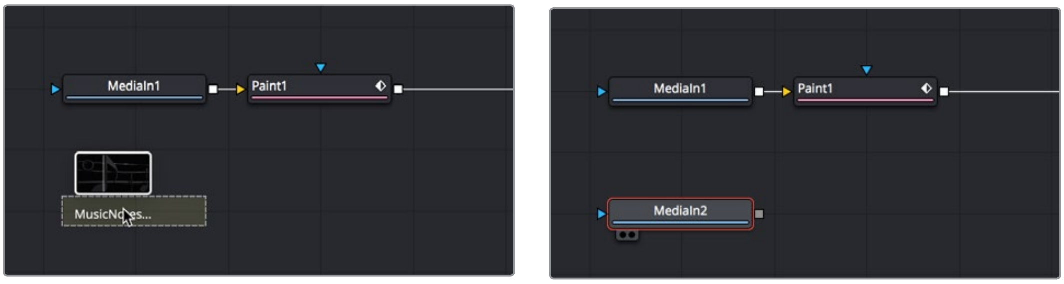

(Left) Dragging a clip from the Media Pool, (Right) Dropping it onto your composition

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

TIP: If you drag a clip from the Media Pool directly on top of a connection line between any two other nodes in the Node Editor, that clip will automatically be added as the foreground clip connected to a Blend node that composites the new clip over the top of whatever you had before.

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

Adding Clips From the File System

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

Using MediaIn Nodes

The MediaIn node is the foundation of every composition you create. This section goes into more detail about the controls that are available for adjusting MediaIn nodes.

MediaIn Node Inputs

MediaIn nodes have one "effects mask" input and one output. In the case of the effects mask input, connecting a mask node such as a Polygon or BSpline node automatically creates an alpha channel in the MediaIn node.

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

Inspector Properties of MediaIn Nodes

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

MediaIn Node Parameters for Clips in a Timeline

When you create a composition either for a single clip in the Timeline, or by turning multiple clips in the Timeline into a Fusion Clip, the MediaIn nodes for clips in the Timeline display fewer parameters than those imported directly from the Media Pool, because the timing of clips that have already been edited into a timeline is already set.

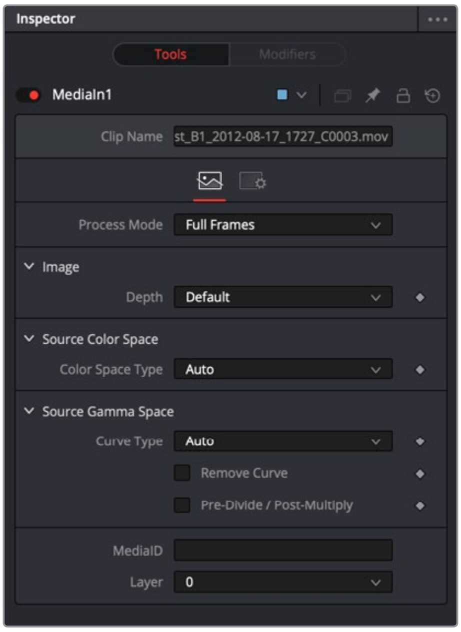

The Inspector parameters for a clip from the Timeline

Overall

- **Clip Name:** Displays the name of that clip.
- **Process Mode:** Lets you choose whether the clip represented by that node will be processed as Full Frames, or via one of the specified interlaced methods.

Image

- **Depth:** Lets you choose a bit depth for the image data output by this MediaIn node.

Source Color Space

- **Color Space Type:** Lets you choose a color space for the image data output by this MediaIn node. Auto uses the timeline color space, or whichever color space is assigned by Resolve Color Management (RCM) if it's enabled. Space lets you choose a specific setting from a Color Space pop-up menu, while a visual "horseshoe" graph lets you see a representation of the color space you've selected.

Source Gamma Space

- **Curve Type:** Lets you choose a gamma setting for the image data output by this MediaIn node. Auto uses the timeline gamma, or whichever gamma is assigned by Resolve Color Management (RCM) if it's enabled. Space lets you choose a specific setting from a Gamma Space pop-up menu, while a visual graph lets you see a representation of the gamma setting you've selected. Log lets you choose a specific log encoding profile from the Log Type pop-up menu, with a visual graph showing a representation of the log setting you've selected, and additional Lock RGB, Level, Soft Clip, Film Stock Gamma, Conversion Gamma, and Conversion table options presented to finesse the gamma output.
- **Remove Curve:** Depending on the selected gamma space or on the gamma space found in Auto mode, the associated gamma curve is removed from, or a log-lin conversion is performed on, the material, effectively converting it to a linear output space.
- **Pre-Divide/Post-Mujltiply:** Lets you convert "straight" alpha channels into pre-multiplied alpha channels, when necessary.

Other

- **MediaID:** An ID assigned by DaVinci Resolve for that clip.
- **Layer:** In the case of Compound clips, this identifies the layer of the clip.

Parameters for MediaIn Nodes of Clips From the Media Pool

When you add a clip from the Media Pool or from your file system directly into a composition, the resulting MediaIn node has more options than clips that were originally in the Timeline. This makes it easier to trim that clip's media, to hold the first or last frame for a longer duration than the original media, or to reverse or loop the clip to get more range for your composition.

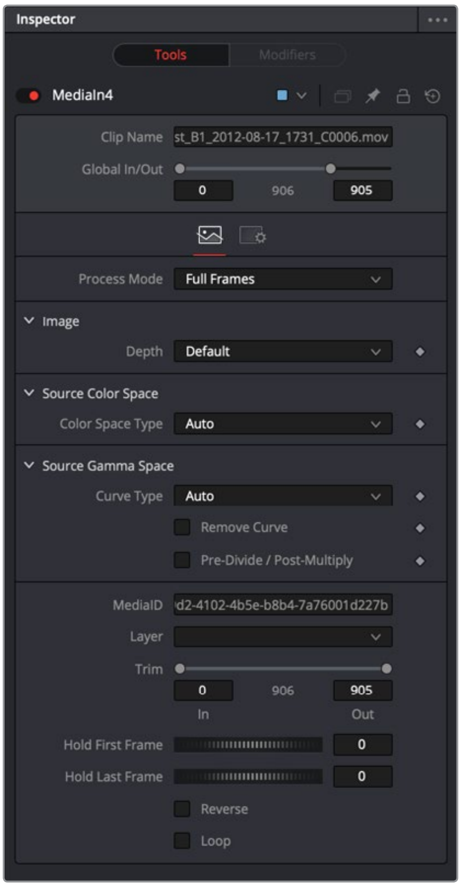

The Inspector parameters for a
clip imported from the Media Pool

Overall

- **Clip Name:** Displays the name of that clip.
- **Global In and Out:** Use this control to specify the position of this node 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 node 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.
- **Process Mode:** Lets you choose whether the clip represented by that node will be processed as Full Frames, or via one of the specified interlaced methods.

Image

- **Depth:** Lets you choose a bit depth for the image data output by this MediaIn node.

Source Color Space

- **Color Space Type:** Lets you choose a color space for the image data output by this MediaIn node. Auto uses the timeline color space, or whichever color space is assigned by Resolve Color Management (RCM) if it's enabled. Space lets you choose a specific setting from a Color Space pop-up menu, while a visual "horseshoe" graph lets you see a representation of the color space you've selected.

Source Gamma Space

- **Curve Type:** Lets you choose a gamma setting for the image data output by this MediaIn node. Auto uses the timeline gamma, or whichever gamma is assigned by Resolve Color Management (RCM) if it's enabled. Space lets you choose a specific setting from a Gamma Space pop-up menu, while a visual graph lets you see a representation of the gamma setting you've selected. Log lets you choose a specific log encoding profile from the Log Type pop-up menu, with a visual graph showing a representation of the log setting you've selected, and additional Lock RGB, Level, Soft Clip, Film Stock Gamma, Conversion Gamma, and Conversion table options presented to finesse the gamma output.
- **Remove Curve:** Depending on the selected gamma space or on the gamma space found in Auto mode, the associated gamma curve is removed from, or a log-lin conversion is performed on, the material, effectively converting it to a linear output space.
- **Pre-Divide/Post-Mujltiply:** Lets you convert "straight" alpha channels into pre-multiplied alpha channels, when necessary.

Other

- **MediaID:** An ID assigned by DaVinci Resolve for that clip.
- **Layer:** In the case of Compound clips, this identifies the layer of the clip.
- **Trim:** The Trim range control is used to trim frames from the start or end of a clip. Adjust the Trim In to remove frames from the start and 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.
- **Hold First Frame/Hold Last Frame:** The Hold First Frame and Hold Last Frame controls will hold the first or last frame of the clip for the specified amount of frames. Held frames are included in a loop if the footage is looped.
- **Reverse:** Select this checkbox to reverse the footage so that the last frame is played first and the first frame played last.
- **Loop:** Select this checkbox to loop the footage until the end of the project. Any lengthening of the clip using Hold First/Last Frame or shortening using Trim In/Out is included in the looped clip.

Using Loader and Saver Nodes

The Loader node lets you add a subset of supported media formats, such as OpenEXR, to a composition directly from the file system, not from the Media Pool. OpenEXR media, in particular, can contain high-quality floating-point image data, as well as alpha channel information and multiple layers of images and mattes; it's a good format for specialty compositing tasks where you need to import high-quality media for specific tasks that you don't need cluttering up the Media Pool.

To use a Loader node, simply create a Loader node anywhere within your composition that it's needed, and an import dialog will appear with which you can choose the OpenEXR file on disk that you want to load. Clicking Open attaches that media file to the Loader node you've created.

Loader Node Parameters

Loader nodes share parameters with MediaIn nodes, as described previously in this chapter.

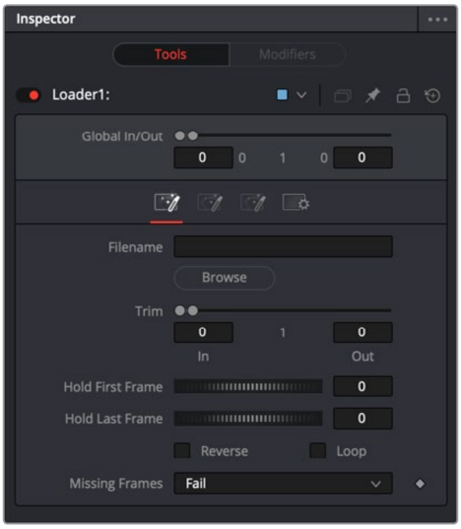

The Inspector parameters of a Loader node

Manually Caching Using Loader and Saver Nodes

The Loader node is also useful for compositing-specific workflows, such as rendering out specific branches of a node tree to OpenEXR via a Saver node, and then reimporting the result in place of the branch in order to improve the performance of your composition. Used this way, Loader and Saver nodes provide a manual workflow for caching but via media files that will never be purged unless you specifically delete them.

Outputting Images Using Saver Nodes

Saver nodes let you export OpenEXR media directly from the Fusion page. Most usefully, Saver nodes let you export a subset of the nodes in a composition that you need to export as a self-contained render. You can add as many Saver nodes as you like to whichever branches of your composition's node tree you need to output, in order to export multiple parts of a composition.

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

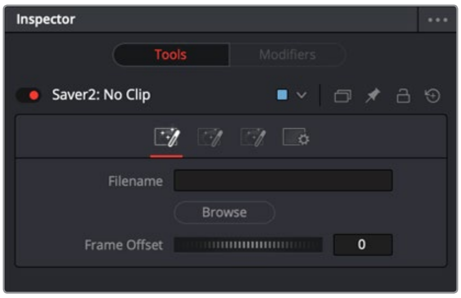

The Inspector parameters for a Saver node

To output one or more Saver nodes, choose Fusion > Render All Savers.

Chapter 47

Image Processing and Color Management

This chapter covers how the Fusion page fits into the overall image processing pipeline of DaVinci Resolve 15. It also discusses the value of doing compositing with clips using a linear gamma, and how to deal with color management in the Fusion page, so you can work with a linear gamma while previewing the image in the Viewer using the gamma of your choice.

Contents

Fusion's Place in the Resolve Image Processing Pipeline

In most workflows, clips are edited, then effects are applied to the edited clips, and the clips with these effects are graded, in pretty much that order. This is the order of operations that DaVinci Resolve follows, source clips edited into the Timeline in the Edit page flow into the Fusion page node tree, and image data from the Fusion page node tree flows into the Color page. DaVinci Resolve goes so far as to expose this via the order of the page buttons at the bottom of the screen, with the Edit page feeding the Fusion page, and the Fusion page feeding the Color page.

However, this isn't the whole story. As of DaVinci Resolve 15, the following sections describe which effects happen prior to the Fusion page, and which effects happen after the Fusion page.

Source Media into the Fusion Page

For ordinary clips, the MediaIn nodes in the Fusion page represent each clip's source media, as modified by whatever changes you've imposed on the source media via the Clip Attributes window, and whatever Edit page Transform and Cropping adjustments you've made to that clip.

Edit Page Plug-ins and the Fusion Page

If you add a ResolveFX or an OFX plug-in to a clip in the Edit page, and then you open the Fusion page, you won't see that plug-in taking effect. That's because these plug-ins are actually applied after the output of the Fusion page, but before the input of the Color page. If you open the Color page, you'll see the Edit page plug-in being applied to that clip, effectively as an operation prior to the grading adjustments and effects you apply in the Color page Node Editor.

With this in mind, the order of effects processing in the different pages of DaVinci Resolve can be described as follows:

Forcing Effects into the Fusion Page by Making Compound Clips

There is a way you can force clips with Edit page ResolveFX and OFX and Color page grades into the Fusion page, and that is to turn that clip into a compound clip. When Edit page effects and Color page grading are embedded within compound clips, MediaIn nodes corresponding to compound clips route the effected clip into the Fusion page.

Output from the Fusion Page to the Color Page

The composition output by the Fusion page's MediaOut node are propagated via the Color page's source input, with the sole exception that if you've added plug-ins to that clip in the Edit page, then the handoff from the Fusion page to the Color page is as follows:

What Viewers Show in Different Pages

Owing to the different needs of compositing artists, editors, and colorists, the Viewers show different states of the clip.

- **The Edit page Source Viewer:** Always shows the source media. If Resolve Color Management is enabled, then the Edit page Source Viewer will show the source media at the timeline color space and gamma.
- **The Edit page Timeline Viewer:** Shows clips with all Edit page effects, Color page grades, and Fusion page effects applied, so editors see the program within the context of all effects and grading.
- **The Fusion page Viewer:** Shows clips at the timeline color space and gamma, with no Edit page effects and no Color page grades.
- **The Color page Viewer:** Shows clips with all Edit page effects, Color page grades, and Fusion page effects applied.

Sizing and the Fusion Page

With the addition of the Fusion page, the order of sizing operations in DaVinci Resolve is a bit more complex. However, it's important to understand which sizing operations happen prior to the Fusion page, and which happen after, so you know which effects will alter the image that's input to the Fusion page, and which effects happen to the Fusion page's output. For example, lens correction, while not strictly sizing, is nonetheless an effect that will change how the image begins in your Fusion composition. However, stabilization is an effect that comes after the Fusion page, so it has no effect on the composition you're creating.

The order of sizing effects in the different pages of DaVinci Resolve can be described as follows:

Color Management in the Fusion Page

How the color of images input to the Fusion page by MediaIn nodes is handled depends on whether or not you have DaVinci YRGB Color Managed (RCM) or ACES enabled in the Color Science parameter of the Color Management panel of the Project Settings. Regardless of how you intend to handle your project's color, it's important to understand linear color, and why it's useful for compositing.

Converting to Linear in the Fusion Page

Because node operations in the Fusion page handle image data in very direct ways, you should ideally composite images that use a linear gamma, especially when you're combining images and effects involving bright highlights. This is because 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.

For example, you can apply filtering effects, such as a blur, to an image using 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.

Automatic Color Management in the Fusion Page

If you're using RCM or ACES color management, then all images coming into the Fusion page via MediaIn nodes are automatically converted to Linear, and all images leaving the Fusion page via MediaOut nodes are automatically converted back to the timeline color space. No user intervention is required to make this happen, and this sets you up to apply all necessary compositing and effects operations properly within a linear gamma, giving you the best results.

Manual Color Management in the Fusion Page

If you're not using Resolve Color Management (RCM) in your project, then images coming into the Fusion page via MediaIn nodes are always in the timeline gamma and. For some simple operations, this may be fine, but it's not always the ideal way to work.

The Fusion page has nodes you can use to convert the image out of each MediaIn node from the timeline gamma to linear gamma at the beginning of your composite, and then convert from linear back to the timeline gamma at the end of your composite, right before the MediaOut node feeds its result to the Color page.

- Each MediaIn node has source color space and source gamma space controls in the Inspector that let you apply a color space and gamma conversion without the need to add another node.
- The CineonLog node, found in the Film category of the Effects Library, lets you do a conversion from any of the formats in the Log Type pop-up menu to Linear, and vice versa. This is useful if your timeline gamma is a Log format.
- The FileLUT node, found in the LUT category of the Effects Library, lets you do a conversion using any LUT you want, giving you the option to manually load one of the LUTs that accompany DaVinci Resolve in the /Library/Application Support/Blackmagic Design/DaVinci Resolve/LUT/VFX IO/ directory (on macOS) to perform to/from linear gamma conversions.

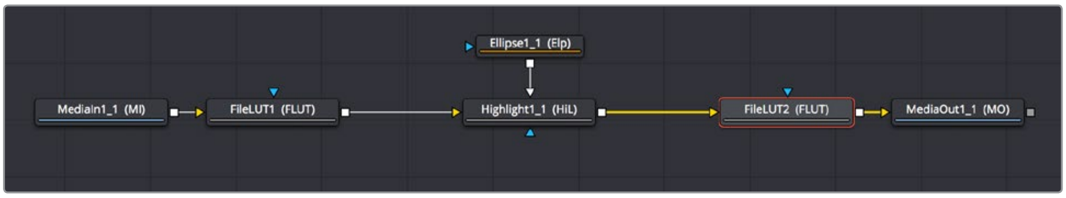

A node tree with "to linear" conversions at the beginning, and a "from linear" conversion at the end

NOTE: In the standalone version of Fusion, "Loader" nodes have color space and gamma conversion built-in when you expose their controls in the Inspector. However, this functionality has not yet been added to DaVinci Resolve 15.

Viewer Gamma and Color Space While Working in Linear

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

It would be obviously impossible to work if you couldn't see the image as it's supposed to look, in the final 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 of the many available color space and gamma conversions to view with. How this works depends, again, on whether you have RCM or ACES enabled, or not.

With RCM or ACES enabled, the only selectable option is Managed, which shows you the currently viewed node as it appears within the timeline color space. You can turn the LUT button off to view the linear image, but you can't change the color space being viewed.

If you're not using RCM or ACES, then you can choose whichever setting you want to use for viewing the image. The top five options (excluding Managed) let you choose Fusion controls, which can be customized via the Edit command at the bottom of this menu. The rest of this menu shows all LUTs from the /Library/Application Support/Blackmagic Design/DaVinci Resolve/LUT/VFX IO/ directory (on macOS) to use for viewing. So, if you're working in linear, you can choose VFX IO > Linear to Gamma 2.4 to see a "normalized" version of the composite you're working on.

The Viewer LUT button turned on, choosing a LUT to use from the pop-up menu

NOTE: By default, Viewers in Fusion show you the image prior to any grading done in the Color page, since the Fusion page comes before the Color page in the DaVinci Resolve image processing pipeline.

Chapter 48

Understanding Image Channels and Node Processing

This chapter seeks to demystify how the Fusion page handles image data, in the process showing you how different nodes need to be connected in order to get the results you expect. It also explains the mysteries of premultiplication, and presents a full explanation of how the Fusion page is capable of using and even generating auxiliary data.

Contents

Channels in the Fusion Page

If you're an old hand at compositing in Fusion, this chapter may be somewhat remedial. However, the Fusion page 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 the Fusion page.

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 the Fusion page supports. Incidentally, all image data in DaVinci Resolve, including the Fusion page, is 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 the Fusion page 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 the Fusion page 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 MediaIn 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.

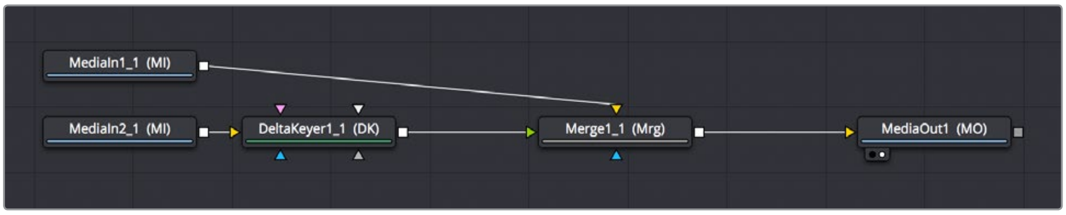

MediaIn2 node connected to a DeltaKeyer node, connected to a Merge node, which is connected to another MediaIn node to combine the two images using the alpha channel output by the DeltaKeyer

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

MediaIn 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 MediaIn node's output is connected to the input of a Highlight node to create a sparkly highlight effect.

MediaIn node connected to a Highlight node connected to MediaOut node

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.

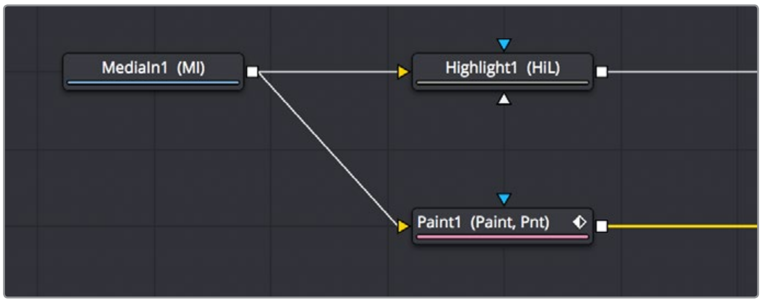

The MediaIn node's output is branched to the inputs of two other nodes

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 (a similar idea to using a KEY input to perform secondary corrections in the Color page). 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.

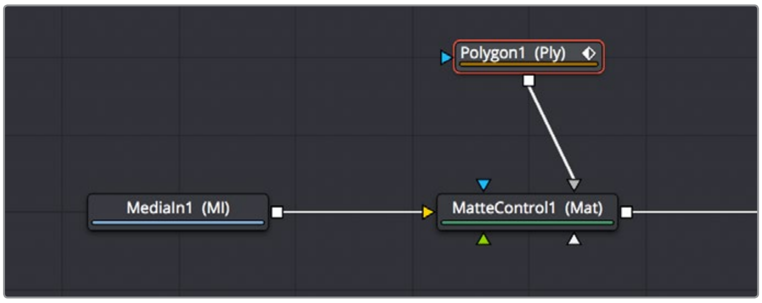

Polygon node connected to a MatteControl node for rotoscoping

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.

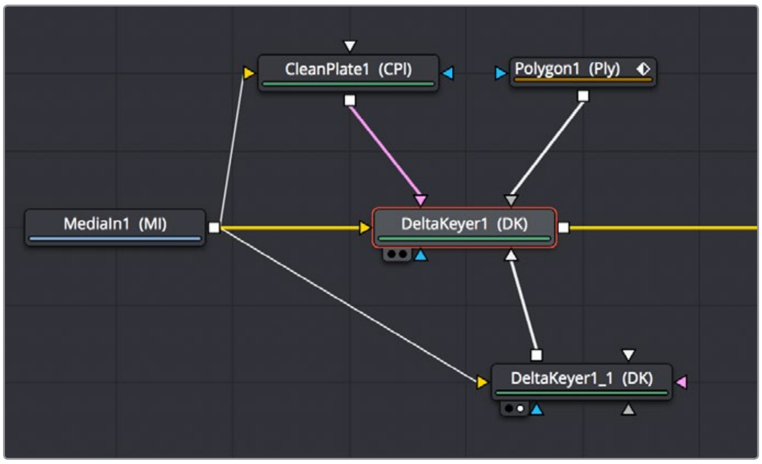

MediaIn node connected to the main "Input" input of a DeltaKeyer node;
other nodes connected to specific inputs for those particular nodes

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

(Left) The node input's tooltip in the Tooltip bar, (Right) The node tooltip in the Node Editor

Connecting to the Correct Input

When you're connecting nodes 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 commonly "Input" or "Background."

Side by side, dropping a connection on a node's
body to connect to that node's primary input

However, if you drop a connection line right on top of a specific input, then you'll connect to that input, so it's important to be mindful of where you drop connection lines as you wire up different node trees together.

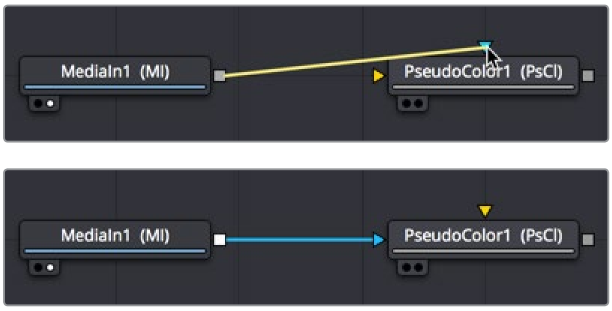

Side by side, dropping a connection on a specific node input, note how the inputs rearrange themselves afterwards to keep the node tree tidy-looking

TIP: If you hold the Option key down while you drag a connection line from one node onto another, and you keep the Option key held down while you release the pointer's button to drop the connection, a menu appears that lets you choose which specific input you want to connect to, by name.

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 Page 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 little 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, including the Merge node, the Matte Control node, and the Channel Booleans node as commonly used examples. To help you keep things straight, 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 Connect

Unlike the Color page, where each Corrector node is capable of performing nearly every kind of grading operation in combination for speed of grading, each of the many nodes in the Fusion page accomplish 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 MediaIn nodes and image processing operations that comprise one's composite.

Because each Fusion page node has a specific function, they're categorized by type to make it easier to keep track of which nodes require what types of image channels as input, and what image data you can expect each node to output. These general types are described here.

A node tree showing the main categories of node colors

Blue MediaIn Nodes and Green Generator Nodes

Blue MediaIn nodes add clips to a composite, and green Generator nodes create images. Both types of nodes output RGBA channels (depending on the source and generator), and may optionally output auxiliary channels for doing advanced compositing operations.

Because these are sources of images, both kinds of nodes can be attached to a wide variety of other nodes for effects creation besides just 2D nodes. For example, you can also connect MediaIn nodes to Image Plane 3D nodes for 3D compositing, or to pEmitter nodes set to "Bitmap" for creating different particle systems. Green Generator nodes can be similarly attached to many different kinds of nodes, for example attaching a FastNoise node to a Displace 3D node to impose undulating effects to 3D shapes.

2D Processing Nodes, Color Coded by Type

These encompass most 2D processing and compositing operations in DaVinci Resolve, 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, which 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 of the channels that are output from the upstream node to the downstream node. 2D nodes, which constitute most simple image processing operations in the Fusion page, propagate all channel data from node to node, including RGB, alpha, and auxiliary channels, regardless of whether or not that node actually uses or affects a particular channel.

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.

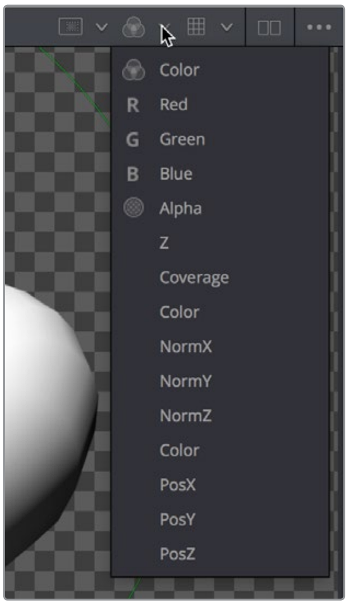

All channels that are available to the currently viewed node can be isolated via the Viewer's Color control

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.

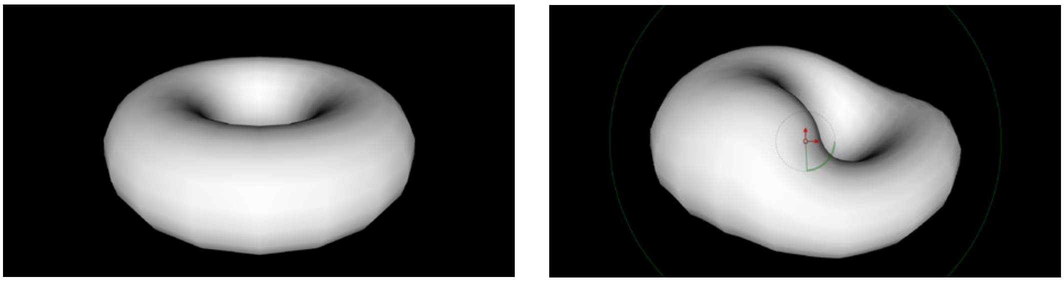

(Left) The Normal Z channel output by a rendered torus, (Right) The Normal Z channel after the output is connected to a Vortex node; note how this auxiliary channel warps along with the RGB and A channels

This is appropriate, because in most cases you want to make sure that all channels are transformed, warped, or adjusted together. You wouldn't want to shrink the image without also shrinking the alpha channel along with it, and the same is true for most other operations.

On the other hand, some nodes deliberately ignore specific channels, when it makes sense.

For example, the Color Corrector and Gamut nodes, both of which are designed to alter RGB data specifically, 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.

MediaIn, DeltaKeyer, Color Corrector, and Merge/MediaIn 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.

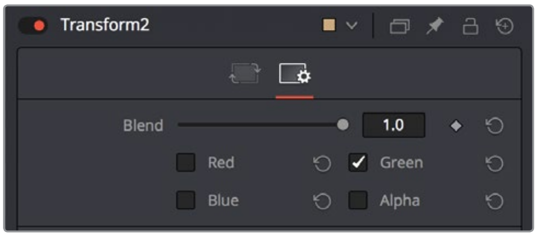

The channel limiting checkboxes in the Settings panel of a Transform node, so only the Green channel is affected

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

Skipping Channel Processing

Under the hood, most nodes actually process all channels first, but afterwards 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 whole reasons for compositing is to begin with an image that lacks an alpha channel, add one via keying or rotoscoping, and then composite that result against other images. While the methods for this are covered in detail in later chapters, here's an overview of how this is handled within the Fusion page.

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

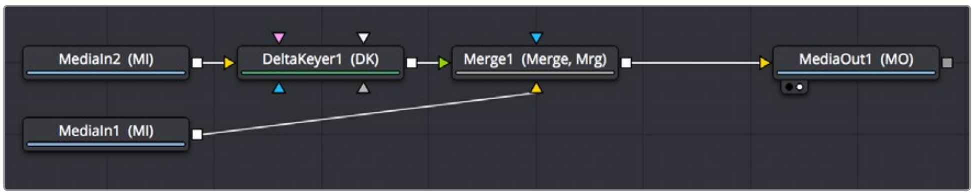

A simple node tree for keying; note that only one connection links the DeltaKeyer to the Merge node

In the case of rotoscoping using a Polygon node, you'll typically connect the image being rotoscoped to the background input of a MatteControl node, and a Polygon 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 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 rotoscoped subject over another image.

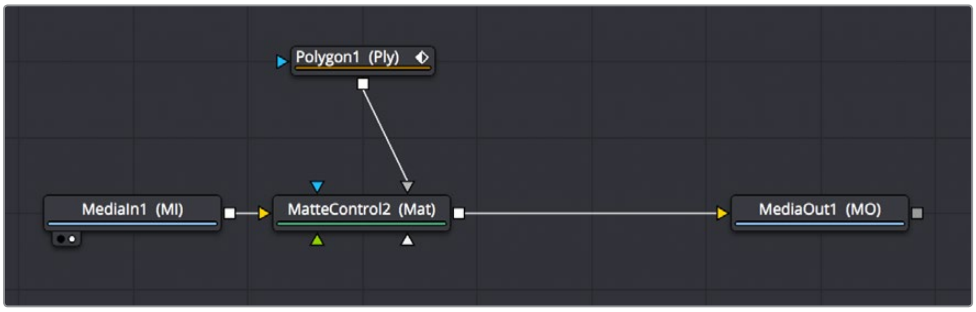

A simple rotoscoping node tree

In both cases, you can see how the Fusion page 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

As you've seen above, 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.

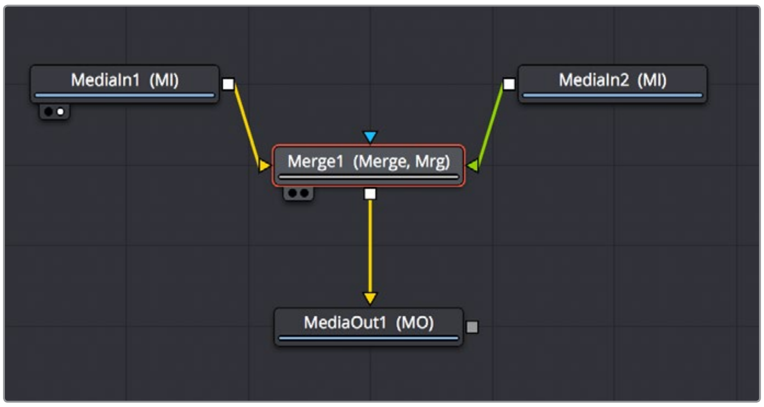

A simple Merge node composite

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 are 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:** Used to switch among or combine two sets of input channels in different ways, using a variety of simple pre-defined imaging math operations.
- **Channel Booleans:** Used to rearrange YRGB/auxiliary channels within a single input image, or among two input images, to create a single output image. If you only connect a single image to this node, it must be connected to the background input to make sure everything works.
- **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 the Fusion page, 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 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 pre-multiplied 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.

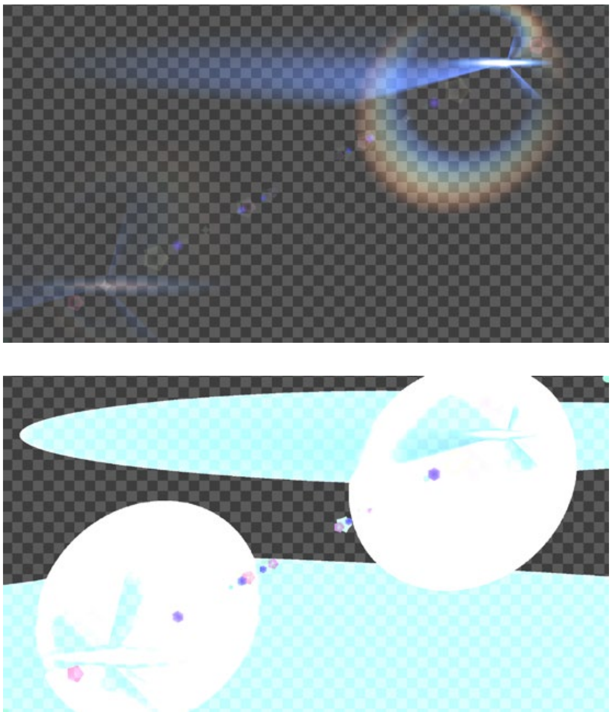

(Top) Premultiplied alpha image, (Bottom) Straight alpha image

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 two 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 MediaIn Nodes That Need It

When you select a MediaIn node, the Import panel in the Inspector have 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.

NOTE: This functionality was not yet available in DaVinci Resolve 15 as of this writing.

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.

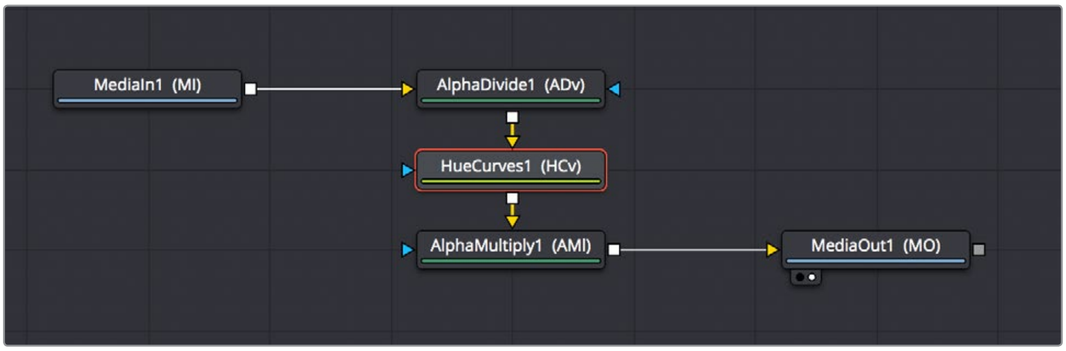

A node tree with explicit Alpha Divide and Alpha Multiply nodes

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 the Fusion page 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 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 the Fusion page, 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 of the auxiliary channels as well as other metadata used within Fusion.

Creating Auxiliary Channels in Fusion

The following nodes create auxiliary channels:

- **Renderer 3D:** Creates these channels in the same way as any other 3D application would, and you have the option of outputting every one of the auxiliary data channels that the Fusion page supports.
- **Optical Flow:** Generates Vector and Back Vector channels by analyzing pixels over consecutive frames to determine likely movements of features in the image.
- **Disparity:** Generates Disparity channels by comparing stereoscopic image pairs.

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.

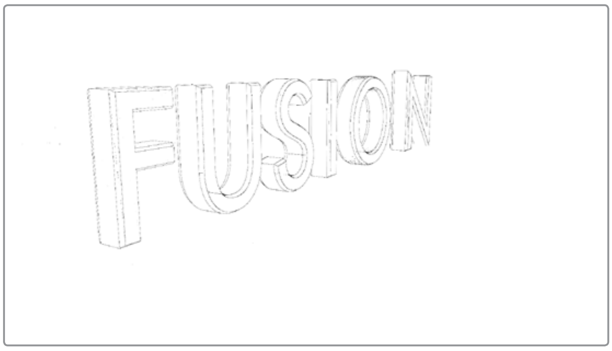

Z-Coverage channel

Background RGBA

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

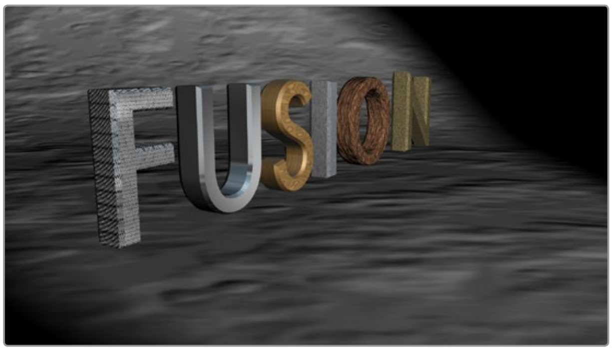

Background RGBA

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.

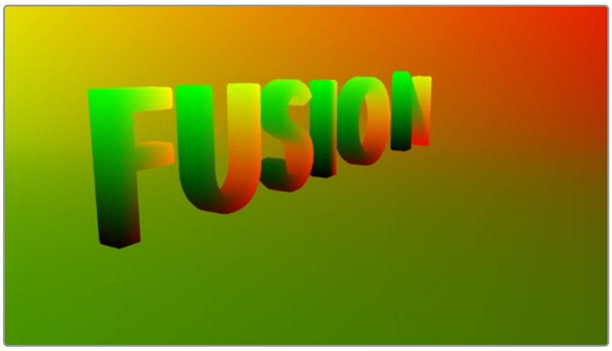

UV Texture

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.

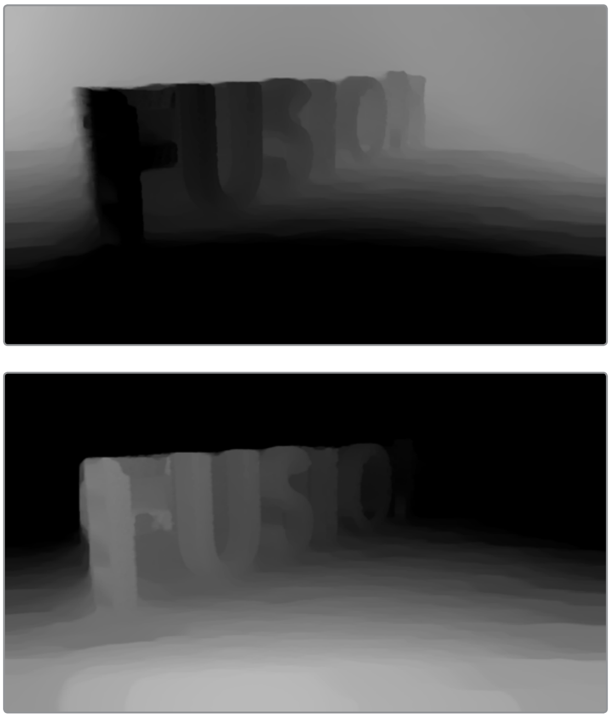

XY Disparity

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 no problem.

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.

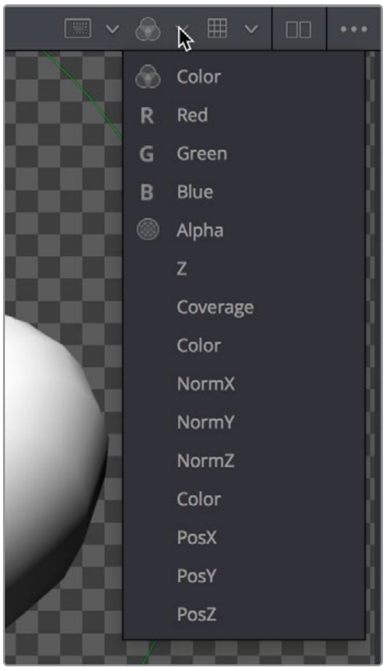

Selecting a channel from the Viewer's toolbar

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 closer to the camera and thickens in regions farther away from the camera. You use the Pick tool to select the Depth values from the image and to define the Near and Far planes of the fog's effect.

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 49

Learning to Composite in Fusion

This chapter is a grand tour of the basics of the Fusion page, walking you through the process of shepherding a clip from the Edit page to the Fusion page, 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 the Fusion page together to work efficiently.

Contents

What's a Composition?

A "composition" describes the collection of nodes that creates an effect in the Fusion page, just as a "grade" describes the collection of nodes that creates a color adjustment or look on the Color page. The relationship between the Edit page and the Fusion page, at a basic level, is similar to the relationship between the Edit page and the Color page. Every clip can have a grade applied to it in the Color page, and similarly every clip can have a composition applied to it in the Fusion page.

If you use the Fusion page to add effects or do any compositing at all, a badge appears on that clip in all timelines to show that clip has a composition applied to it.

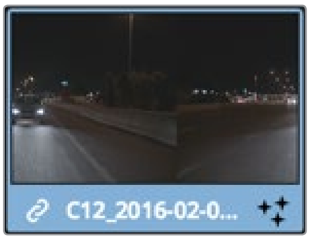

Clips with Fusion page compositions have a Fusion badge to the right of the name

Moving From the Edit to the Fusion Page

Whenever you want to create a composite using a clip from the Edit page, the simplest way to work is to just move the playhead so it intersects the desired clip in the Edit page timeline, make sure the clip you want to composite is the topmost clip of any superimposed stack of clips (whichever clip you see in the Timeline Viewer is the one you'll be compositing with), and then open the Fusion page.

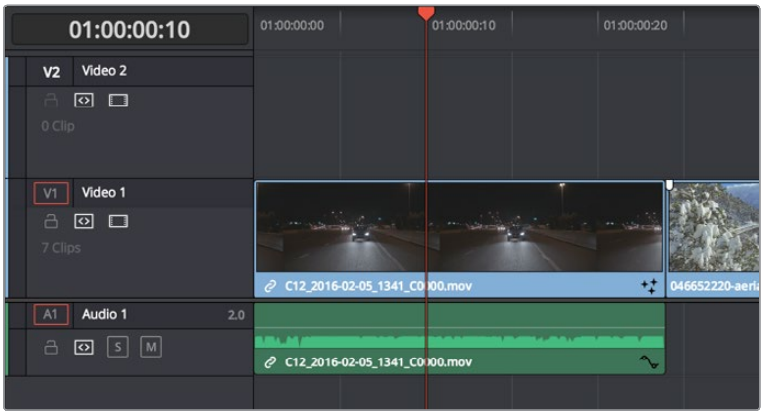

Positioning the playhead over a clip you want to use in a composition

In the Fusion page, you should see a single selected MediaIn1 node representing only the top-most clip you were parked on in the Edit page, and that image should be showing in the Viewer thanks to the MediaOut1 node automatically being loaded in the Viewer (the Viewer buttons visible underneath that node confirm this). Any clips that were underneath that clip are ignored when you work this way, because the idea is you're only doing a quick fix to the current clip at the position of the playhead.

How the Fusion page appears when you first open it while the playhead is on a new clip

The playhead should still be on the same frame you were parked on in the Edit page, except now it's on the equivalent source frame in the Time Ruler underneath the Viewer that represents that clip's media. Yellow markers indicate the range of the current clip that appears in the Timeline, while the source clip's handles extend to the left and right. Lastly, the selected MediaIn1 node displays its parameters in the Inspector to the right.

In the Node Editor at the bottom, the MediaIn1 node is connected to a MediaOut1 node. If this is all you see, there is no effect yet applied to this clip. It's only when you start adding nodes between MediaIn and MediaOut that you begin to assemble a composition.

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 "MediaIn," the actual names of nodes in the Fusion Node Editor have a number appended to them, to indicate which node is which when you have multiple instances of a particular type of node.

Applying and Masking Effects

Let's begin by looking as some very simple effects, and build up from there. Opening the Effects Library, then clicking the Disclosure control to the left of Tools, reveals a list of categories containing all the effects nodes that are 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.

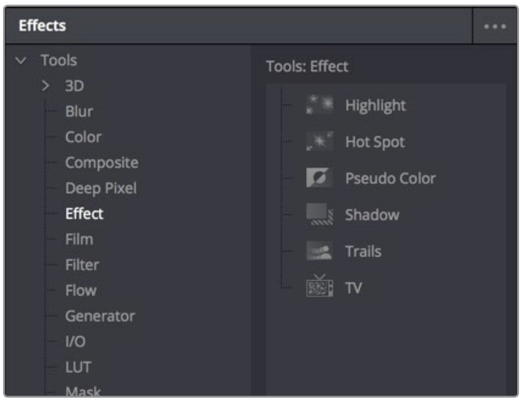

Browsing the Effect category to find the TV node

Adding a Node to the Tree

Assuming the MediaIn 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, and it immediately takes effect in the Viewer thanks to the fact that the MediaOut1 node is what's loaded in the Viewer, since that means that all nodes upstream of the MediaOut1 node will be processed and shown.

A new node added from the Effects Library

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.

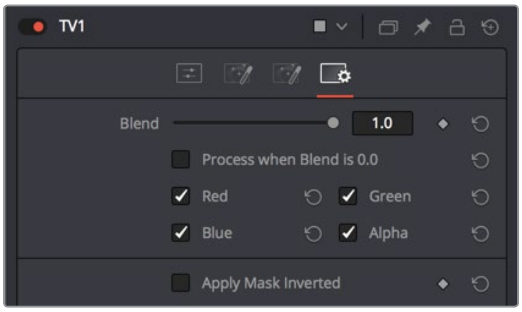

The Settings panel, which has channel limiting and mask handling controls that every node shares

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.

(Left) The original TV effect, (Right) Modifications to the TV effect to make the clip more solid

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 pizazz 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.

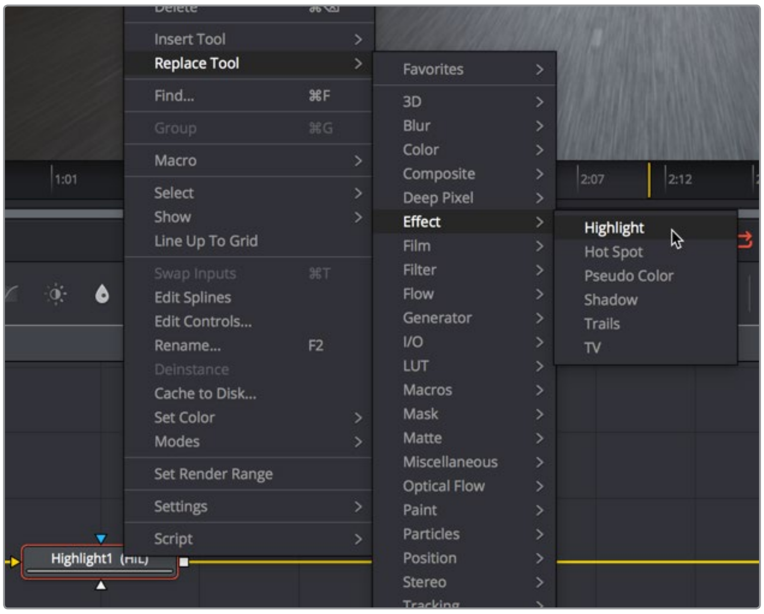

Right-clicking a node to use the contextual menu Replace Node submenu

It's time to use the Inspector controls to customize this effect, but first, let's take a look at how sliders in the Fusion page differ somewhat from sliders on other pages in DaVinci Resolve.

Adjusting Fusion Page Sliders

When you drag a slider in the Fusion page Inspector, in this case the “Number of Points” slider, a little dot appears underneath it. This dot indicates the position of the default value for that slider, and also serves as a reset button if you click it.

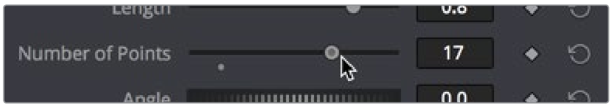

Adjusting a slider reveals a reset button underneath it

Each slider is limited to a different range of minimum and maximum values that is particular to the parameter you’re adjusting. In this case, the “Number of Points” slider maxes out at 24. However, you can 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.

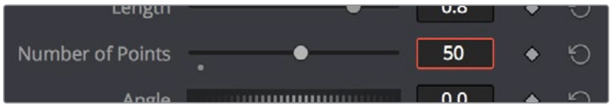

Entering a larger value to expand the range over which a slider will operate

Masking Node Effects

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

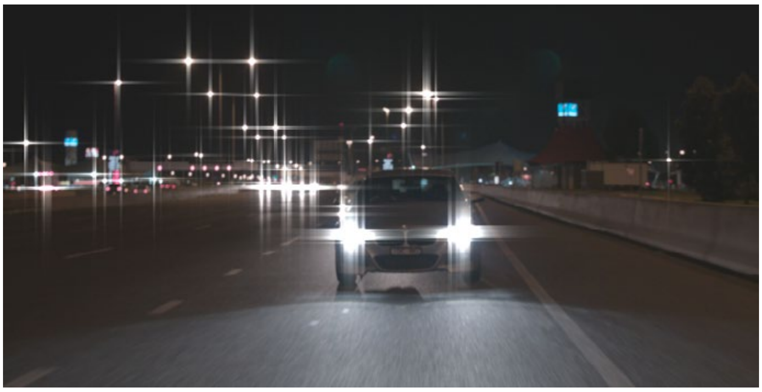

The Highlight effect with a raised Length value (zoomed in)

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 serves a similar function to the KEY input of nodes in the Color page; it lets you use a mask or matte to limit that node’s effect on the image, like a secondary adjustment in color correction. 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.

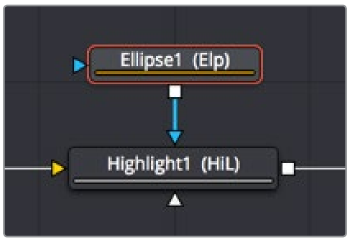

Automatically connecting an Ellipse node to the blue Effect Mask

Adjusting Mask Nodes

Masks, in the Fusion page, 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.

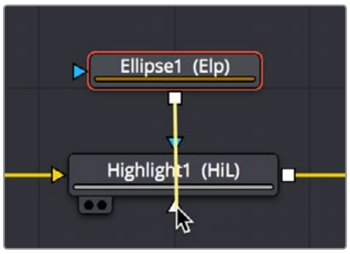

Dragging a connection from one node input to another

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 the Fusion page 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.

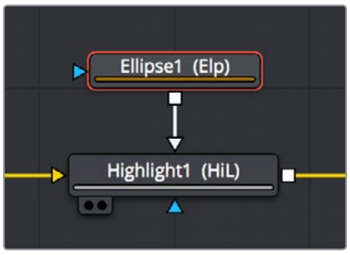

The Ellipse node now connected to the Highlight Mask input, which has moved to the top of the node to keep things neat

Now that the Ellipse1 node is connected to the Highlight Mask, the tight mask we've created just around the car headlights restrict 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.

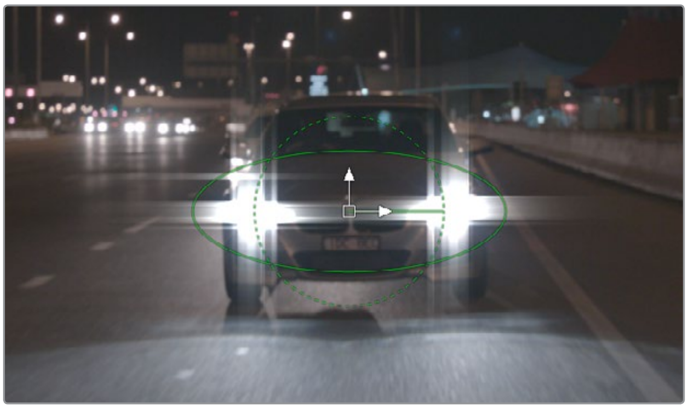

The Highlight effect is uncropped because the effect is being limited via the Highlight Mask input, rather than the Effect Mask input

Unlike nodes on the Color page that have largely the same inputs and outputs, nodes on the Fusion page may have any number of inputs that are 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.

Color Management

If you want to follow a professional compositing workflow of converting media gamma to linear as it comes into the Fusion page, and then back to the original timeline gamma prior to the MediaOut node, at the time of this writing, you must do this manually using either CineonLog or FileLUT nodes.

- The CineonLog node is good if you're converting from or to a log-encoded gamma.
- The FileLUT node is good if you're converting from or to a gamma of 2.4, since it can use the LUTs that are available in the /Library/Application Support/Blackmagic Design/DaVinci Resolve/LUT/VFX IO/ directory (on macOS).

What this looks like in the node tree is that each MediaIn node will have a CineonLog or FileLUT node attached to it doing a conversion to linear gamma, while the MediaOut node will have a CineonLog or FileLUT node attached just before it doing a conversion from linear gamma to the timeline gamma, whatever that happens to be. All other operations in the composition must be applied between these two conversions.

Converting all images to linear coming into the composite, and converting out of linear out of the composite, in this case using CineonLog nodes at the beginning and end

TIP: The CineonLog node can be found in the Film category of the Effects Library, while the FileLUT node is found in the LUT category.

In this example, the timeline gamma is BMD Film 4.6K, which is a log-encoded gamma, so the first CineonLog node does a Log to Lin conversion from this format. The second CineonLong node then does a Lin to Log conversion to this format, in order to move the image data out of the Fusion page in the way DaVinci Resolve expects.

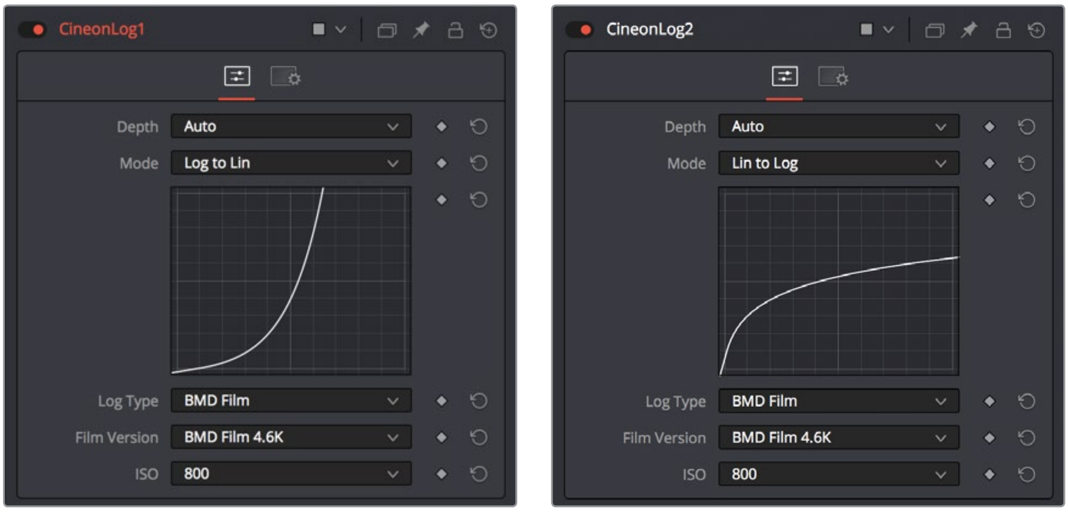

L(Left) The first CineonLog node does a Log to Lin conversion, (Right) the second CineonLog node does a Lin to Log conversion

While this is happening, you'll want to set your Viewer to a gamma setting that lets you see the image as it will look when the audience sees it (more or less). You can do this using the Viewer LUT button. Click to turn it on, and then choose a setting from the VFX IO submenu that you want to preview your composite with as you work, such as Linear to 2.4, corresponding to the BT.1886 standard for gamma used for SDR HD output.

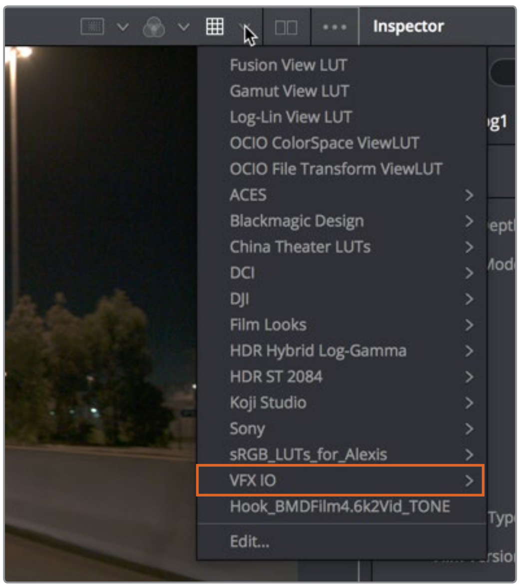

The Viewer LUT button (enabled) and pop-up menu

This is a bit of set-up, it's true, but it will provide superior compositing results, especially for compositions with filtering and lighting effects.

NOTE: Eventually, the MediaIn and MediaOut nodes will include built-in functionality for doing gamma and color space conversions, similarly to the standalone version of Fusion. Furthermore, the Fusion page will at some point in the future be governed by Resolve Color Management (RCM), similarly to the Edit and Color pages, so many of the examples in this chapter will omit these nodes from the node trees.

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

Clicking the Media Pool button in the UI toolbar opens the Media Pool, which if you're already familiar with DaVinci Resolve, is the same Media Pool that's now found in every page except for the Deliver page. The Media Pool shares the same area with the Effects Library, so if you have them both open at the same time, they'll be stacked one on top of another.

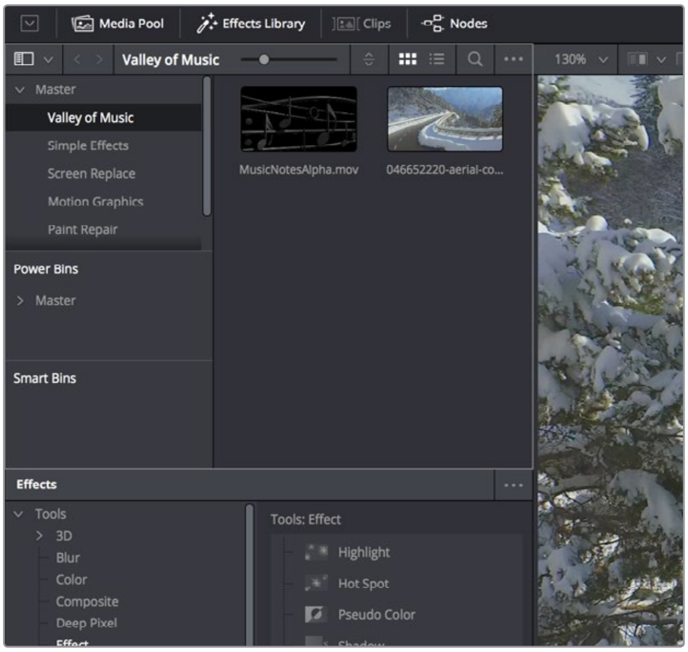

The Media Pool as seen in the Fusion page

If you drag a clip from the Media Pool to an empty area of the Node Editor, you'll add an unconnected MediaIn2 node (incremented to keep it unique) that you can then connect in any way you want.

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.

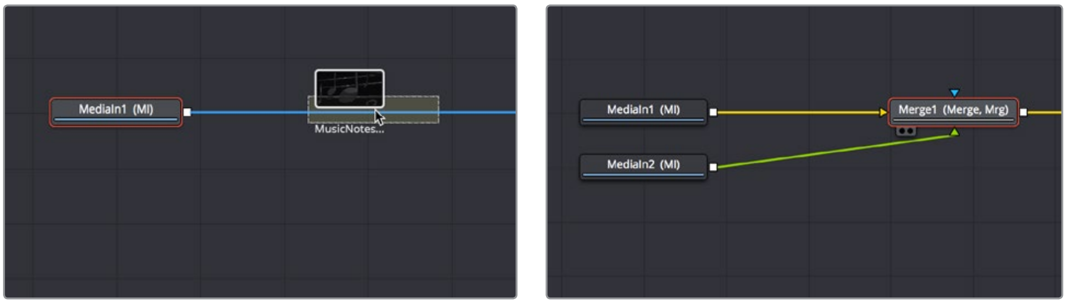

(Left) Dragging a node from the Media Pool onto a connection,
(Right) Dropping it to create a Merge node composite

The Fusion page Node Editor is filled with shortcuts like this to help you build your compositions more quickly. Here's one for if 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 orange, and foreground inputs are green.

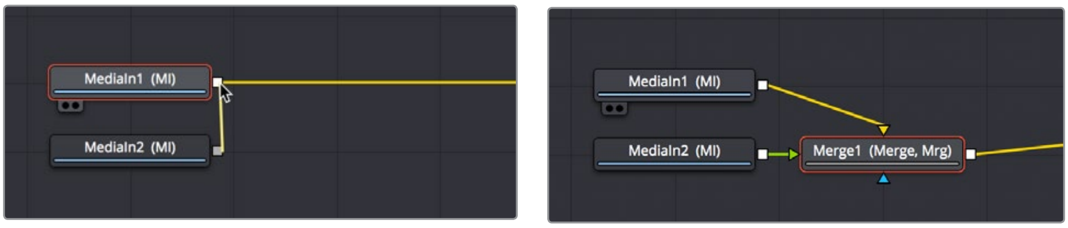

(Left) Dragging a connection from a disconnected node to another node's output, (Right) Dropping it to create a Merge node composite

Adding Clips to the Fusion Page From the File System

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

Adjusting the Timing of Clips Added From the Media Pool

Because the MediaIn2 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 MediaIn2 node's clip is being lined up with the very first frame of the MediaIn1 clip's handles, rather than the first frame of the actual composition range as seen by the yellow marks in the Time Ruler.

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

This a common issue with clips you add from the Media Pool in the Fusion page because those clips were never edited into the Timeline in the Edit page where they could be properly timed and trimmed relative to the other clips in the composition. Fortunately, you can slip clips and resize 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 MediaIn 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 Keyframe Editor are only intended to show you the timing of each composited clip of media.

In this case, we can see that the MediaIn2 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.

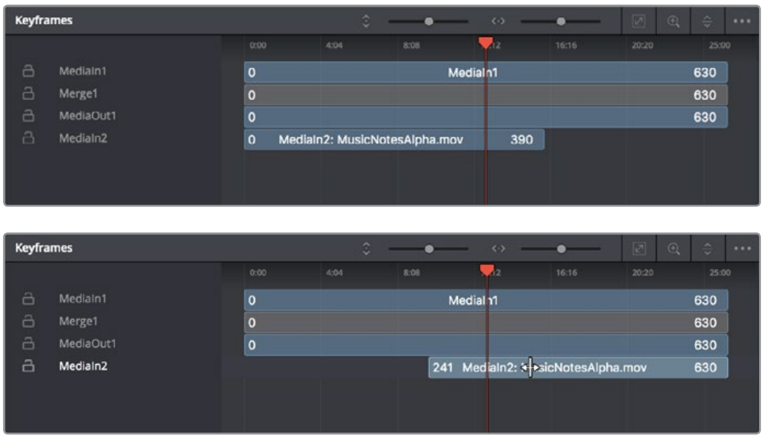

(Top) The original stack of layers, (Bottom) Sliding the MediaIn2 layer to line it up better with the MediaIn1 layer

As a result, the MediaIn2 clip lines up much better with the MediaIn1 clip.

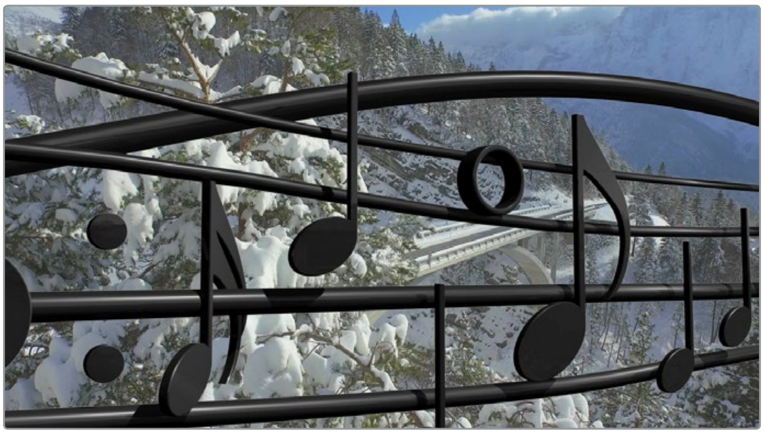

After sliding the MediaIn2 clip to improve its alignment with the other clip in the composite

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 pre-multiplied. 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.

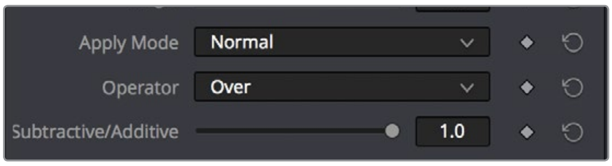

The Subtractive/Additive slider, which can be used to fix or improve fringing in composites

Drag the slider all the way to the left, to the Subtractive position, and the fringing disappears.


(Left) A clip with alpha exhibits fringing, (Right) Fixing fringing by dragging the Subtractive/Additive slider to the left

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 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, with the slider all the way to the right, 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, with the slider all the way to the left, 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.

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 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.

Composite 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 a composite mode in the Blend node to refine the screen replacement effect to incorporate real reflections from the scene.

Setting Up the Initial Composite

The base image in the MediaIn1 node is a clip that's been zoomed into in the Edit page. When you use the Transform, Cropping, or Lens Correction controls for a clip in the Edit page Inspector, those adjustments are passed along as the initial state of the image in the Fusion page, allowing for some prep work to be done in the Edit page, if necessary.

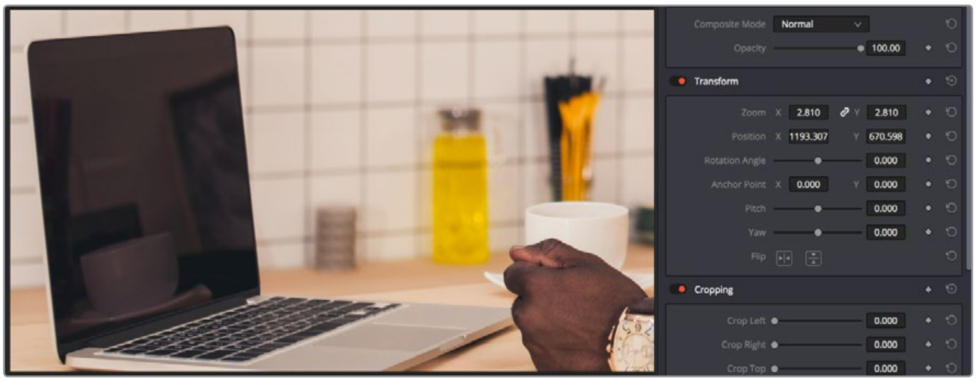

Adjusting the Edit sizing of a clip before moving it into the Fusion page for compositing

Because this particular example uses the Screen Composite mode to do a composite, we'll start by setting up some routine color management in the node tree, to illustrate how this should be handled.

Taking a Shortcut by Copying and Pasting Nodes

In the Fusion page, this first clip has been converted to and from linear gamma using a FileLUT node set to use the "Gamma 2.4 to Linear.cube" from the /Library/Application Support/ Blackmagic Design/DaVinci Resolve/LUT/VFX IO directory. However, after dragging and dropping the video image we need to insert into the screen onto a connection to automatically add it connected to a Merge1 node, we find we need to add another copy of the same FileLUT node after the new MediaIn2 node.

Happily, this is easy to do by selecting and copying the FileLUT1 node that's connected to the MediaIn1 node (Command-C), then selecting the MediaIn2 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.

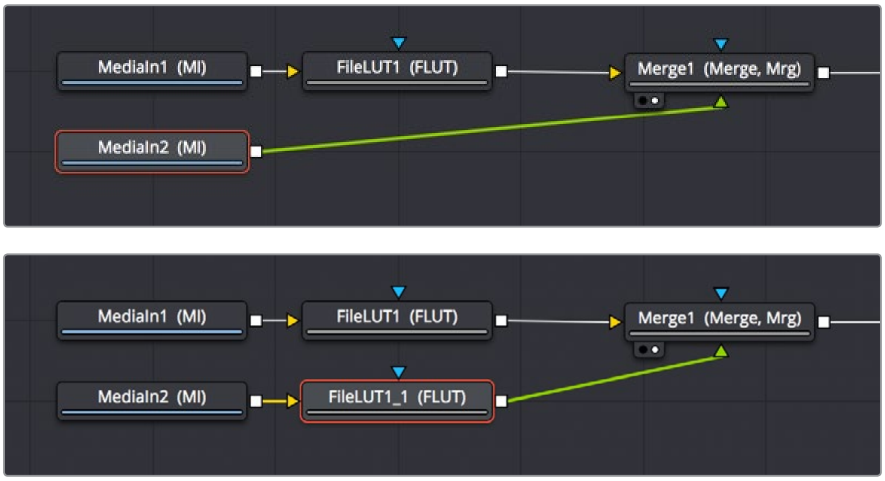

Copying a node from one part of the node tree and pasting to insert it into after a selected node

If we then select the Merge1 node, we can paste another instance of this FileLUT node to come just before the MediaOut1 node, setting its LUT File to the "Linear to Gamma 2.4.cube" LUT that's also found in the /Library/Application Support/Blackmagic Design/DaVinci Resolve/LUT/ VFX IO directory.

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 Which Node You See in the Viewer

Since we're doing gamma conversions to media coming into and going out of the Fusion page, it's no longer suitable to View the MediaOut node as we work, because the Viewer is currently only set up to convert the linear image data that's in-between the two sets of FileLUT nodes to something normal for you to look at (such as gamma 2.4). 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.

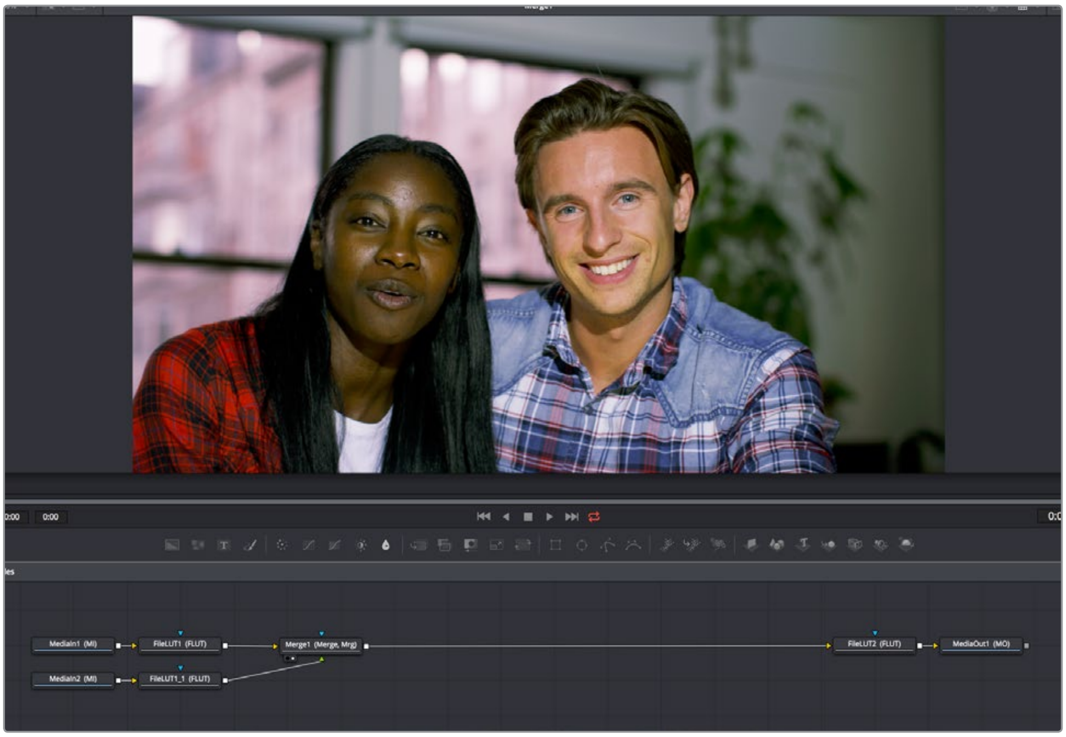

Loading a node from the middle of the node tree into the Viewer to see a specific node you're working on

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 MediaIn2 node, so we'll select the FileLUT_1 node that's attached to it.

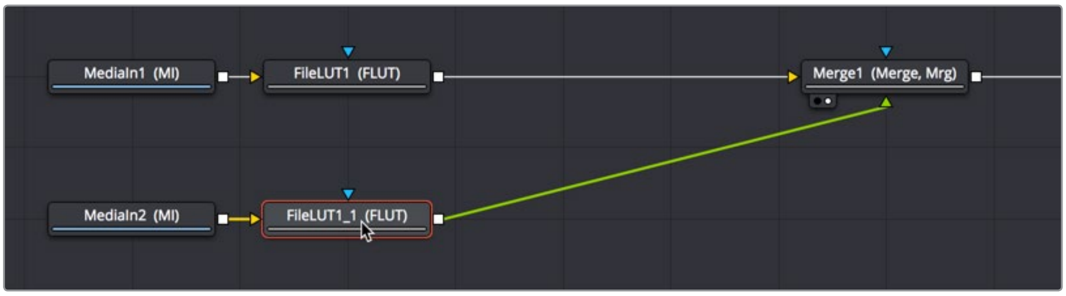

Selecting a node you want to add another node behind

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 the entire Fusion page. 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.

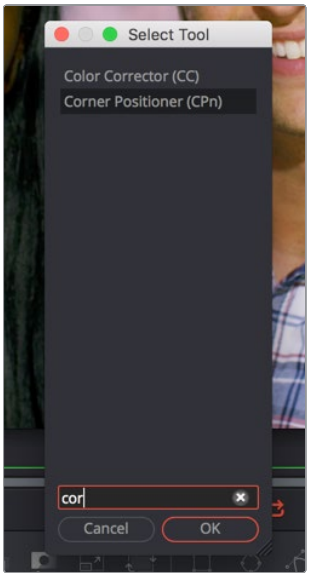

Pressing Shift-Spacebar opens the Select Tool dialog for quickly finding and adding new nodes

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.

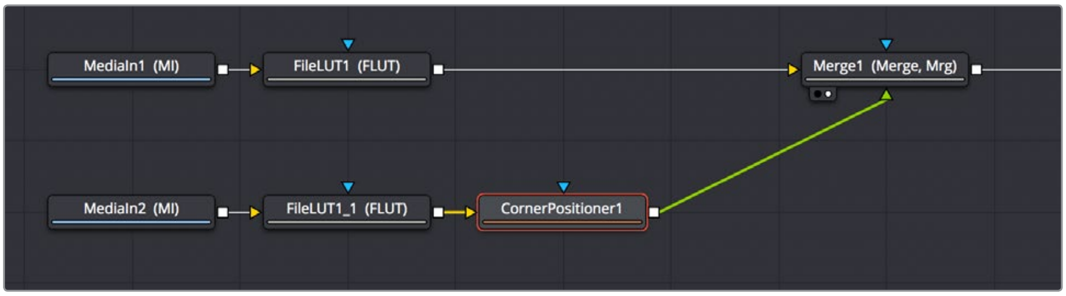

The CornerPositioner node added to cornerpin the foreground image prior to the Merge operation

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 to 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 he Viewer Option Menu. 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.

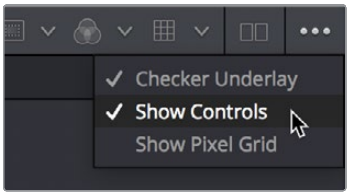

Show Controls in the Option menu toggles on-screen control visibility on and off

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 into or out of the Node Editor.
- Hold the Command key down and use your scroll wheel to zoom in and out of the Node Editor.

Using the Screen Composite Mode in the Merge Node

Once the foreground input image is fit to the 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 composite 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.

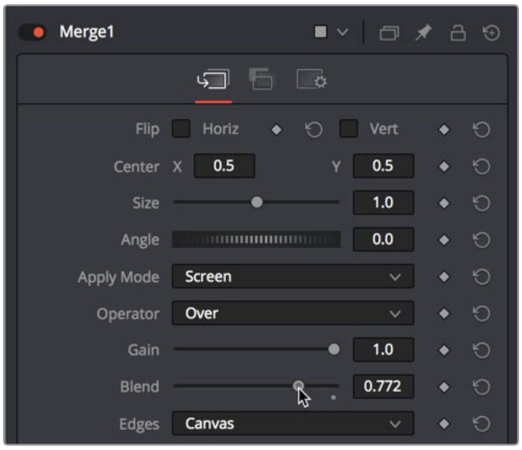

Adjusting the Apply Mode and Blend slider of the Merge node in the Inspector

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 which hide themselves when not needed or when a particular input isn't connected.

The Screen node is perfect for simulating reflections, and lowering Blend a bit lets you balance the coffee cup reflections from the display in the background with the image in the foreground. It's subtle, but helps sell the shot.

(Left) The original composite, (Right) The composite using the Screen apply mode

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 the Fusion page 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 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 HueCurves node from the toolbar. Selecting the CornerPositioner node we added, clicking the HueCurves button on the toolbar adds that node between the CornerPositioner and the Merge node.

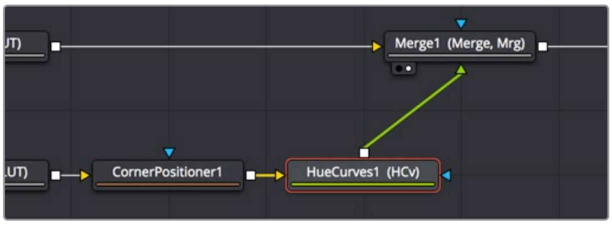

Adding the HueCurves node to make a correction to the foreground image

The HueCurves node exposes a curve control in the Inspector with options for adjusting 9 kinds of curves, each overlapping the others for simultaneous adjustment. 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 towards 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.

(Left) The uncorrected foreground, (Right) Using a HueCurves node to adjust the image for a better composite

Creating and Using Text

In this next example, we'll take a look at how to create a simple text object using the Text+ node. Then, we'll see how to use the text generator's alpha channel in another image to create a more complex composite.

Creating Text Using the Text+ Node

The Text+ node is the primary tool for creating 2D text in the Fusion page. This is also the new text generator that has become available in the Edit page, and because it's so ubiquitous, it appears in 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 MediaIn1 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

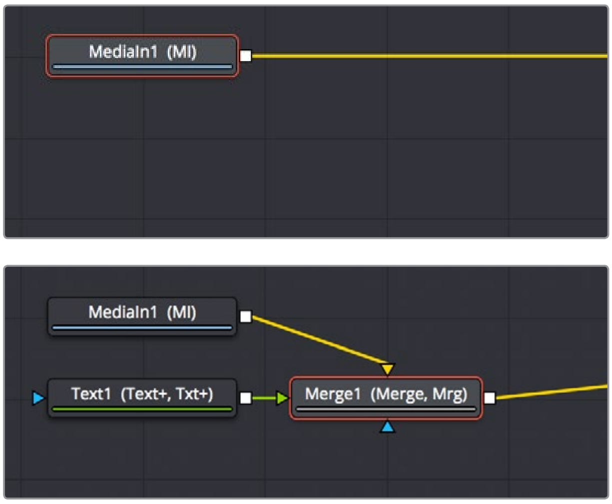

(Top) Selecting a node you want to append another node to, (Bottom) Clicking the Text+ button on the toolbar automatically creates a Merge composite with the text as the foreground input connection

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 that are 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).

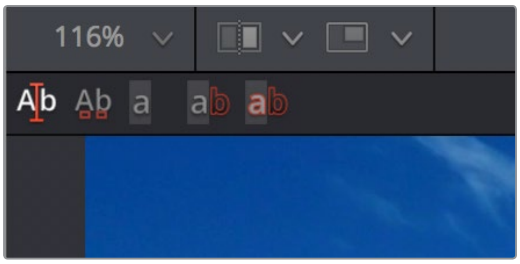

The Viewer toolbar for the Text node with tools for text entry, kerning, and outline controls

The text appears in the Viewer, superimposed against the background clip. Onscreen 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've got 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 the Command key down 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.

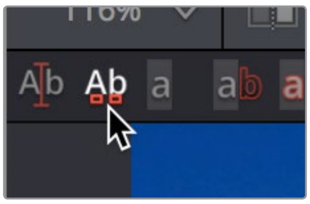

The Manual Kerning tool in the Viewer toolbar

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:

- 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.

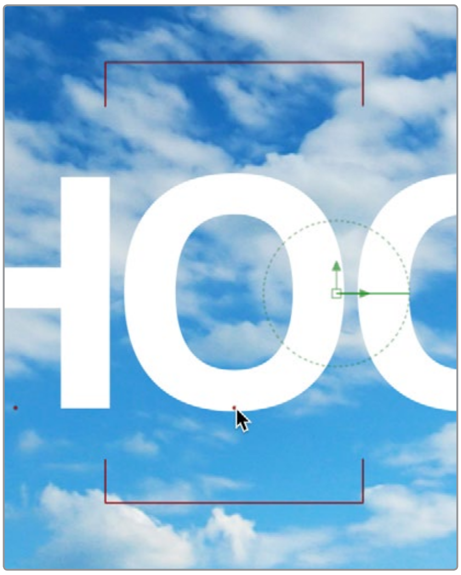

Option-dragging the little red dot revealed by the Manual Kerning tool to manually adjust kerning left or right

So there we go, we've now got a nice title, styled using the Viewer tools and Inspector controls on the Text panel. This looks good, but we've got 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're going to use the text as a matte to cut these letters out of another layer we'll be using for texture. So, first we'll drag another clip, of a chalkboard covered with math, from the Media Pool to the Node Editor as a disconnected MediaIn2 node.

Disconnecting and Reconnecting Nodes

Now we need to do a little rearranging. Moving the Merge1 node up, then clicking the last half of the connection from the Text1 node to the Merge foreground input to disconnect it.

(Top) Clicking the second half of a connection to disconnect it,
(Bottom) The result with the text node disconnected

Next, we'll drag a connection from the MediaIn2 node onto the Merge1 node's foreground input, so the entire Viewer becomes filled with the chalkboard (assuming we're still viewing the MediaOut node). At this point, we need to insert the Text1 node's image as an alpha channel into the MediaIn2 node's connection, and we can do that using a MatteControl node.

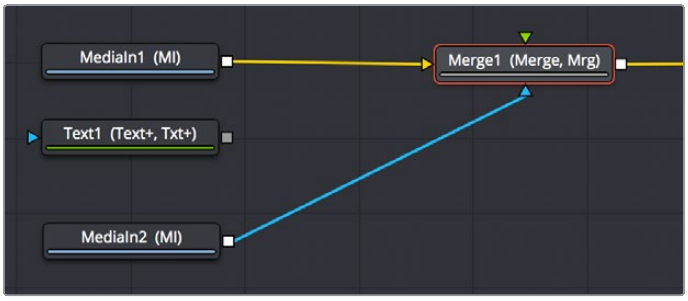

The updated composite, with two video images connected and the text node disconnected

Using Matte Control Nodes

Selecting the MediaIn2 node, we click the Matte Control button of the toolbar to add it between the MediaIn2 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 input of the MatteControl node, as seen in the following screenshot, or the MatteControl node won't work.

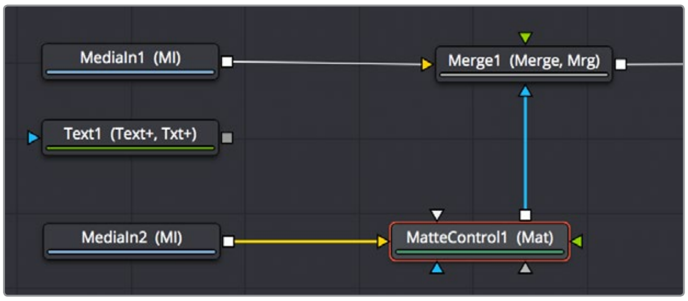

The second image properly connected to the Matte Control node's background input

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 or not you've connected the proper inputs.

One alternate method of connecting nodes together is to hold the Option key down 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.

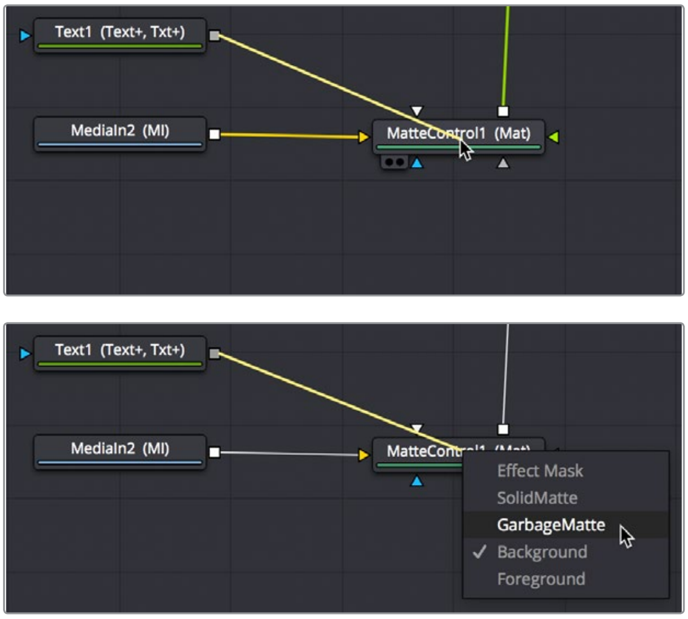

(Before/After) Option-dragging a node connection to drop onto another node exposes a node input menu

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.

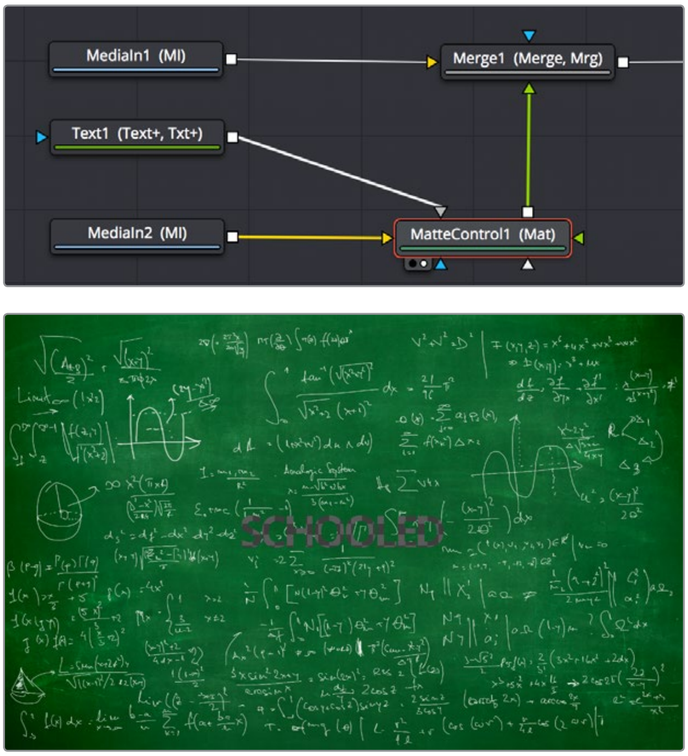

(Top) Connecting the Text node to the Matte Control node's garbage matte input,
(Bottom) The resulting hole punched in the image

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 garbage matte connected, a variety of controls are available to modify how the Garbage Matte input is applied to the image. Click Invert to create the effect we really want, which is text that is textured with the chalkboard image.

The alpha from the Text node punching a hole in another graphic

However, the new chalkboard layer is far bigger than the HD-sized elements we've been working with, so the alpha from the Text1 node is too small. This is easily fixed by setting the Fit pop-up menu to "Width," which automatically resizes the garbage matte to be as big as possible from edge to edge within the image.

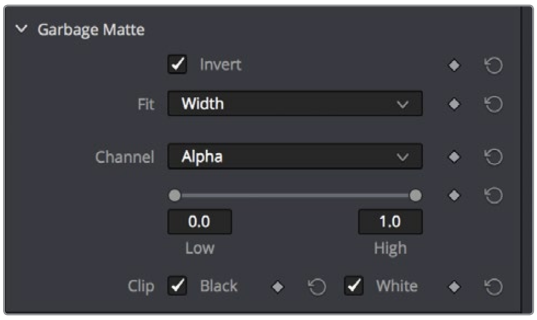

The Garbage Matte settings of the MatteControl node

Now, if we load the Merge1 node into the Viewer, we can see that the text effect is doing everything we want, but now the chalkboard text is too big.

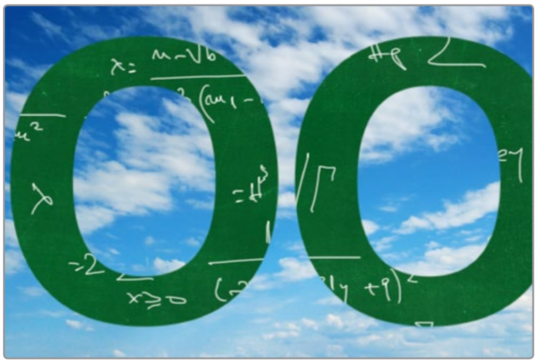

The final composite

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.

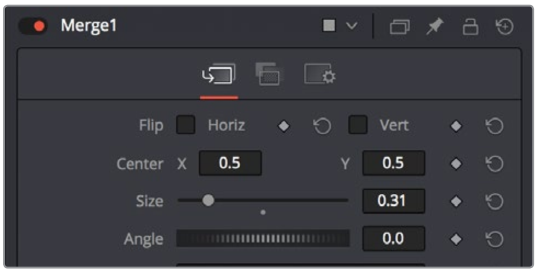

The Merge node transform controls that affect the foreground input's image

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 the Fusion page.

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 MediaIn node.

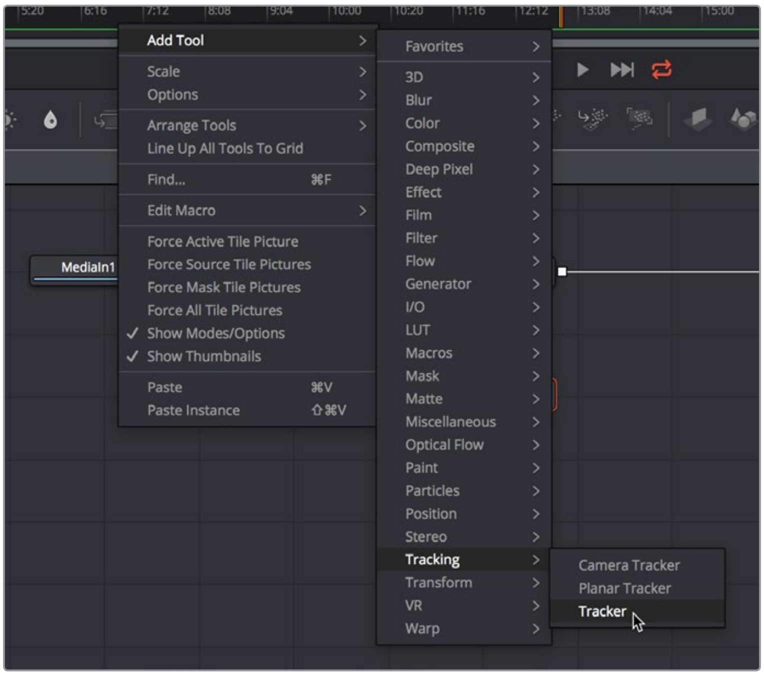

Creating a new node using the Node Editor contextual menu

Next, we'll drag a connection from the MediaIn1 node to the Tracker1 node to automatically connect the source clip to the Tracker1 background input. This branches the output from the MediaIn1 node to the Tracker node, so that the Tracker1 node processes the image separately from the rest of the node tree. This is not required, but it's a nice organizational way to see that the Tracker node is doing an analysis that must be referred to in another way than a "physical" connection.

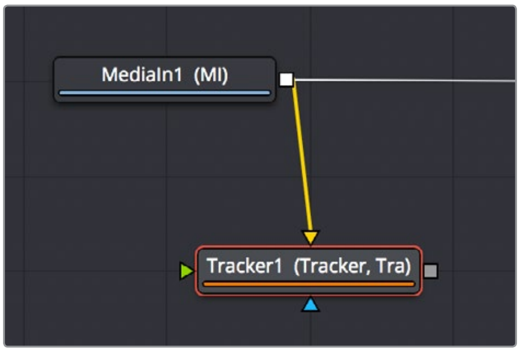

Branching a Tracker node to use to analyze an image

A Simple Tracking Workflow

The Tracker node is the simplest tracking operation the Fusion page has, and while there are several ways of using it, an extremely common workflow is to use the tracker node controls to analyze the motion of a subject in the frame with motion you want to follow, and then use the resulting motion path data by "connecting" it to the Center parameter of another node that's capable of transforming the image you want to match move.

Positioning the Tracker 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.

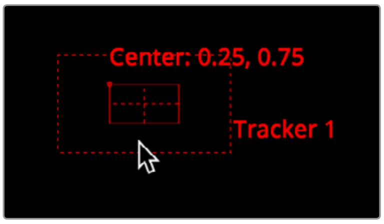

The on-screen controls of a selected tracker seen in isolation

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.

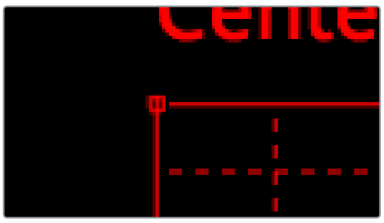

The handle for dragging the tracker boxes to move them around

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.

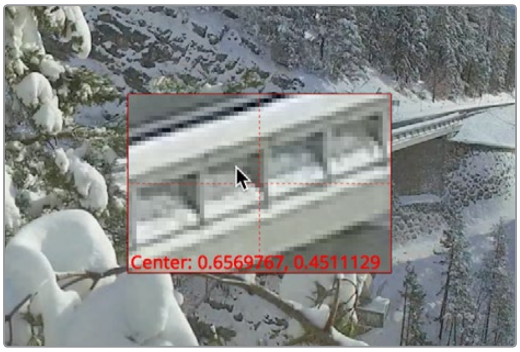

The zoomed-in preview that helps you position the pattern box as you drag it

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.

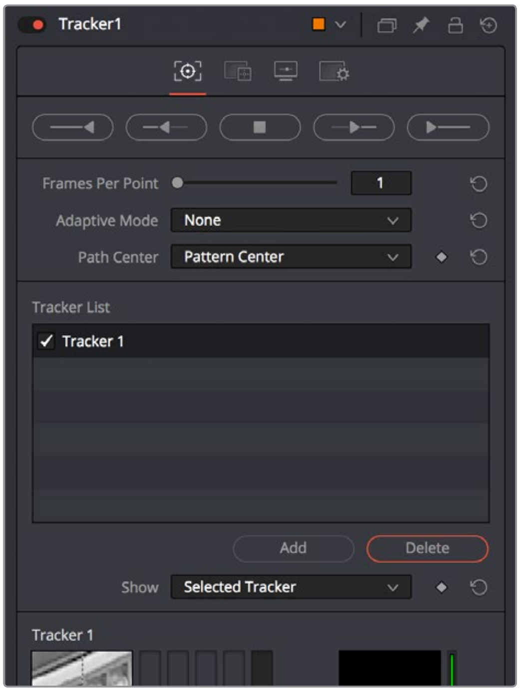

Tracker Inspector controls, with the tracking analysis buttons at top, the tracker options in the middle, and the Tracker List below

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 that much, but because the drone is flying in a circular pattern, the shape of the pattern area is changing as the clip plays. Fortunately, we can choose Every Frame from the Adaptive Mode pop-up, to instruct the tracker to update the pattern being matched at every frame of the analysis, to account for this.

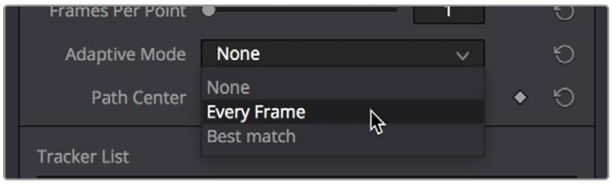

Changing the Adaptive Mode of the Tracker node to Every Frame to account for the camera's shift of perspective

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.

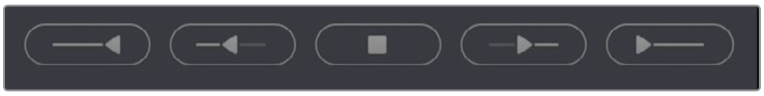

The analysis buttons, left to right, Track from Last Frame, Track Backward, Stop Tracking, Track Forward, Track from First Frame

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.

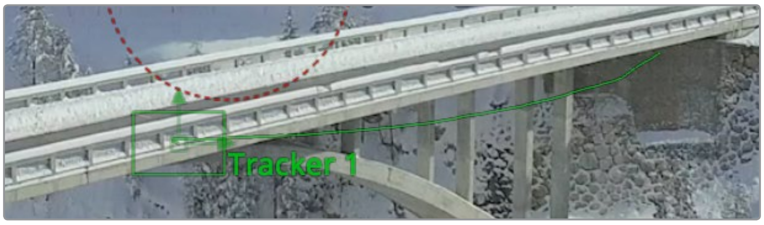

The analyzed motion path resulting from tracking a section of the bridge as the camera flies past

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 make that tracker easier to find in subsequent contextual menus, and lets you keep track of which trackers are following which subjects as you work on increasingly complex compositions.

Renaming a tracker to make it easier to find

Now it's time to connect the track we've just made to the text in order to start it in motion. 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.

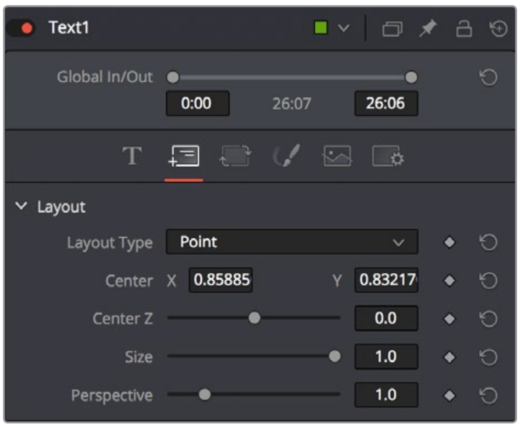

The Layout controls for a Text node, in the Layout panel

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.

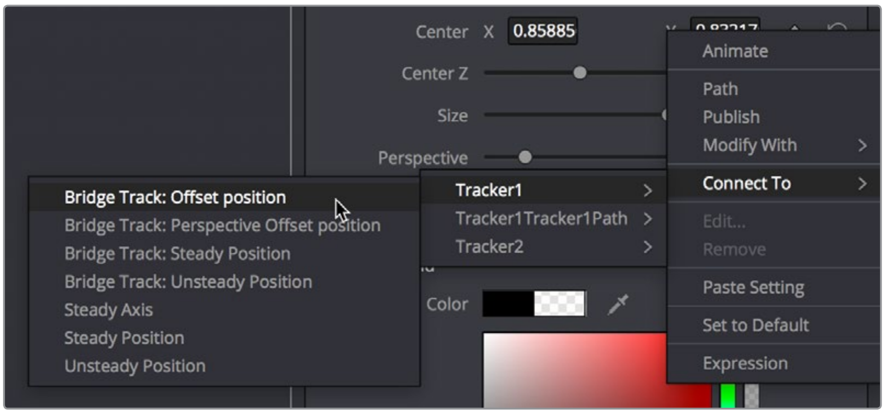

Connecting the Center X and Y parameter to the “Bridge Track: Offset position” motion path we analyzed

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.

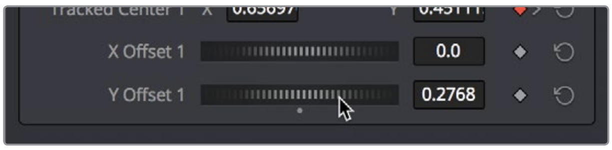

Using the X and Y Offset controls in the Tracker1 node to offset the text layer’s position from the tracked motion path

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 the Fusion page is to use the Planar Tracker, so we'll add the PlanarTracker node after the MediaIn1 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.

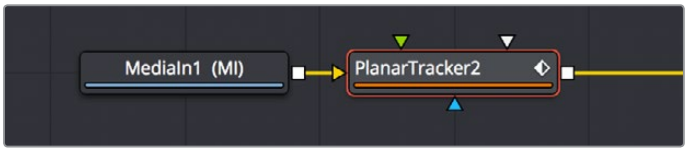

Adding a PlanarTracker node to analyze and steady the part of the image we want to paint on

With the PlanarTracker node selected, and either it or the MediaOut1 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 each on each corner of the man's forehead to create control points, clicking the first one we created to close the shape.

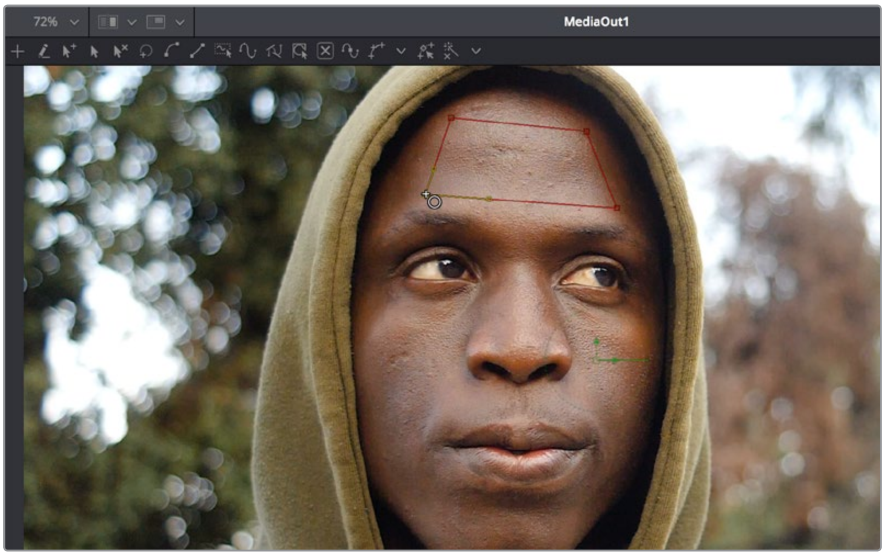

Drawing a shape over the man's forehead to prepare for Planar Tracking

Turning our attention to the Inspector, we can see that the PlanarTracker node has tracking transport controls that are 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.

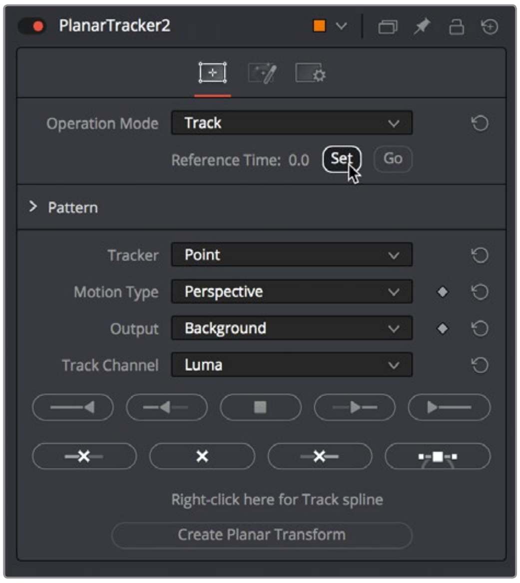

Setting a reference frame at the beginning of the range of frames we want to track

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 towards the end of the clip.

The Pattern controls let you set up how you want to handle the analysis. Of these controls, the Motion Type 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 that are available.

Once you initiate the track, a series of dots appear 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.

Clicking the Track from First Frame button to set the Planar Track in progress, green dots on the image and a green progress bar lets you know the track is happening

NOTE: If you click one of the Track buttons to begin tracking and nothing happens, or 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.

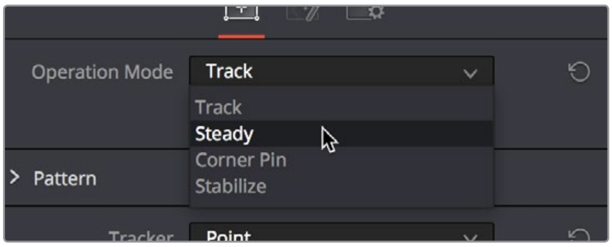

Setting the PlanarTracker node to Steady

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.

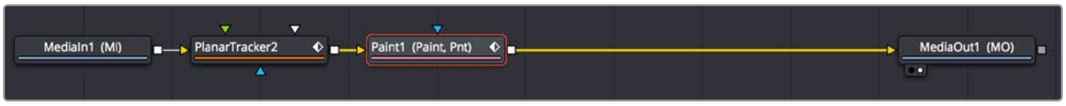

Adding a Paint node after the PlanarTracker to paint onto the steady surface

With the Paint node selected and the MediaOut1 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.

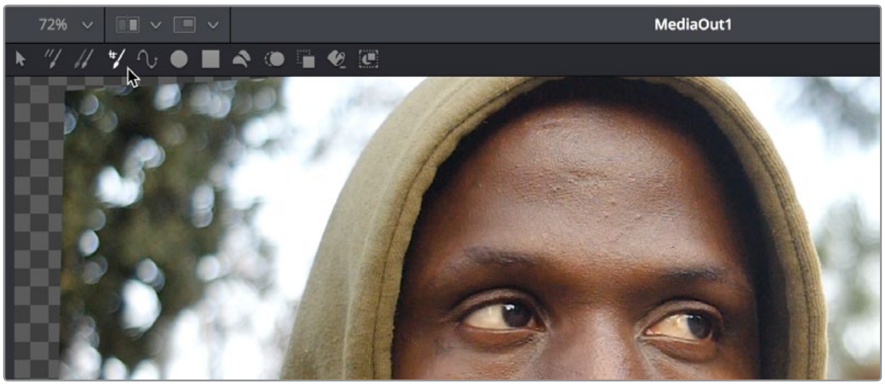

Choosing the Stroke tool from the Paint node's tools in the Viewer toolbar

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 be using 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 tool bar, 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 you want to hold the Option key down 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.

(Left) Setting an offset to sample for cloning, (Right) Dragging to draw a clone stroke

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.

(Top) Original image, (Bottom) After painting out two scars on the man's forehead with the Stroke tool set to Clone

TIP: You can adjust the size of the brush right in the Viewer, if necessary, by holding the Command key down 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 that you want to start making changes to strokes you've made, this is where they appear.

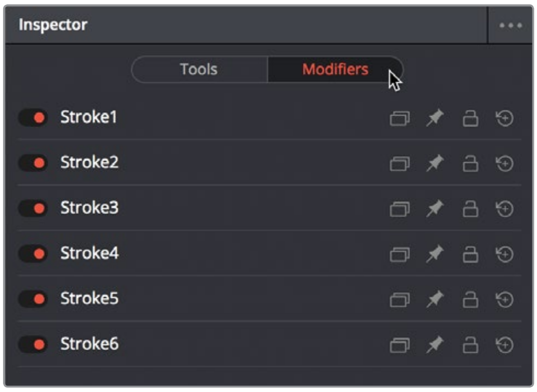

Each stroke made appears as an entry with controls in the Modifiers panel of the Inspector

Keep in mind that the last stoke 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 original PlanarTracker node.

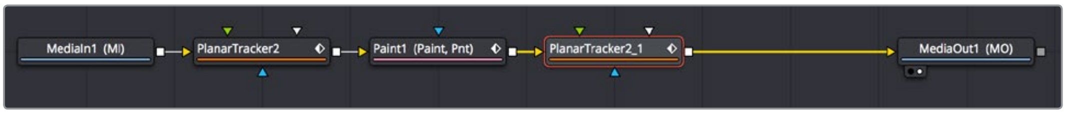

Pasting a second copy of the PlanarTracker node after the Paint node

With the second PlanarTracker node selected, we go into the Inspector and turn on the Invert Steady Transform checkbox, which 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.

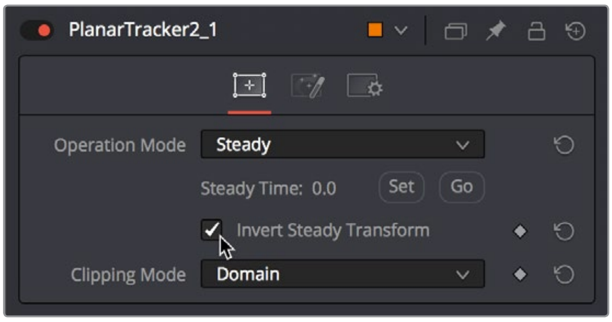

Turning on Invert Steady Transform to try
and put the image back to the way it was

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.

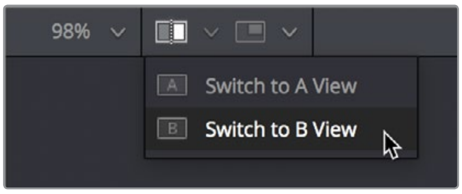

Opening the Split Wipe pop-up
menu in the Viewer

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, then switch back to A View and drag the MediaIn1 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).

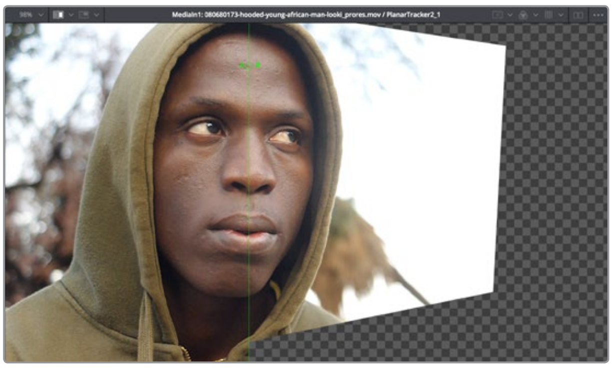

Comparing the "Invert Steady" version of the image
with the original image to see the forehead is the
same in both frames

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 Only Using 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.

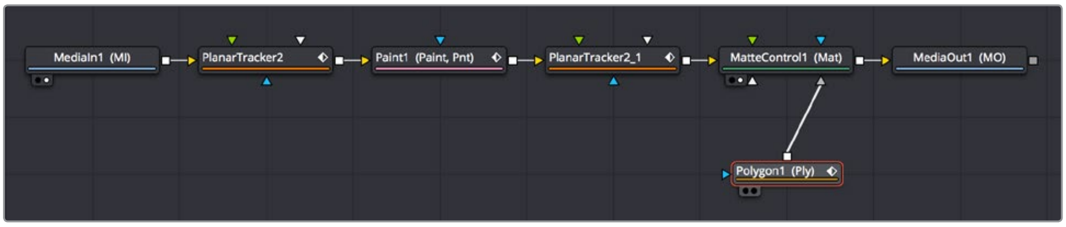

Adding a Polygon node, a MatteControl node, and a Merge node to composite the painted forehead on the original clip

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

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 MediaOut1 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 shape drawing in other spline-based environments, including the Color page:

- 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.

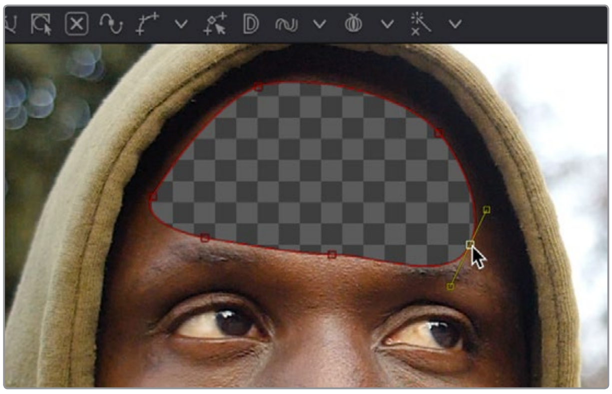

Drawing a shape to isolate the forehead gives an inverted result at first when using the Garbage Matte input of the MatteControl node to attach the Polygon to the MatteControl node

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.

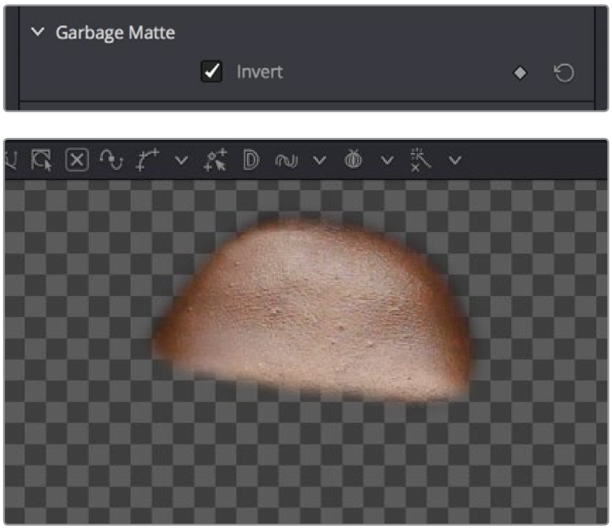

(Top) Inverting the Garbage Matte input, (Bottom) The resulting inverted mask inverting the forehead

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 MediaIn 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 MediaIn1 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, a fixed forehead and the background of the shot looks 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 MediaIn 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.

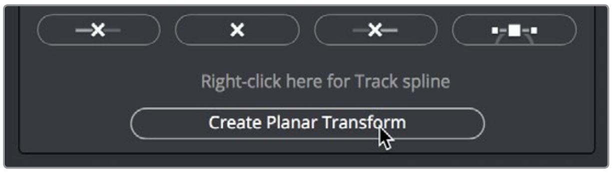

Creating a PlanarTransform node you can use to Match Move other images

We can insert this new node into the node tree to use it by holding the Shift key down and dragging it over the connection between the Polygon node and the MatteControl node, dropping it when the connection highlights.

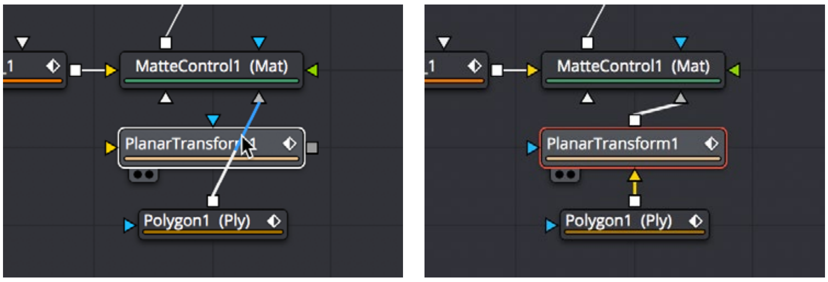

(Left) Inserting a PlanarTransform node by holding the Shift key down while dropping over a connection, (Right) After inserting the PlanarTransform mode

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!

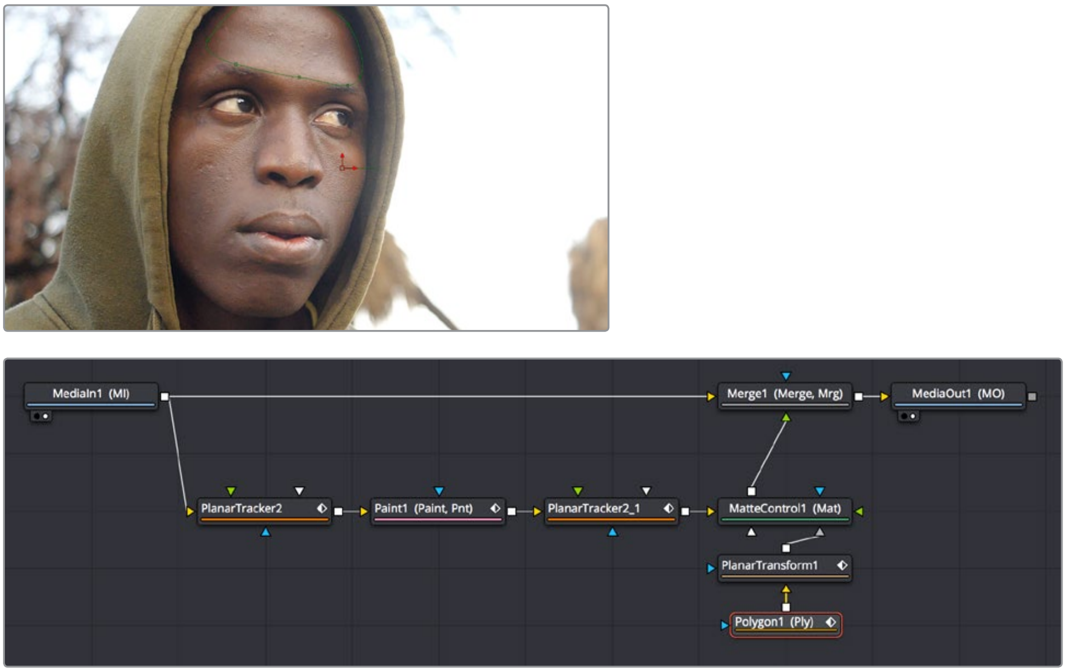

The final painted image, along with the final node tree

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 you can pre-organize the media you want to use in a composition in the Edit page, before creating a Fusion clip to bring all of it into the Fusion page in an organized way. Then, we'll do a simple composite using a greenscreen key and two other layers to create a news story.

Organizing Clips in the Edit Page to Create a Fusion Clip

For this example, we'll take a look at how you can organize multiple clips in the Edit page to use in the Fusion page by creating a "Fusion clip," which is effectively a special-purpose compound clip used specifically by the Fusion page. The next effect we need to create involves a greenscreen clip, a background graphic, and a foreground graphic. This is the kind of situation where superimposing all three layers on the Timeline to set up their order and timing can be the fastest way to set up the foundation of our composition.

With these clips edited together, we select all of them, right-click the selection, and choose New Fusion Clip from the contextual menu. This embeds them all within a single clip, which is easy to manage in the Edit page and keeps all the relevant media necessary for this composition within one handy object.

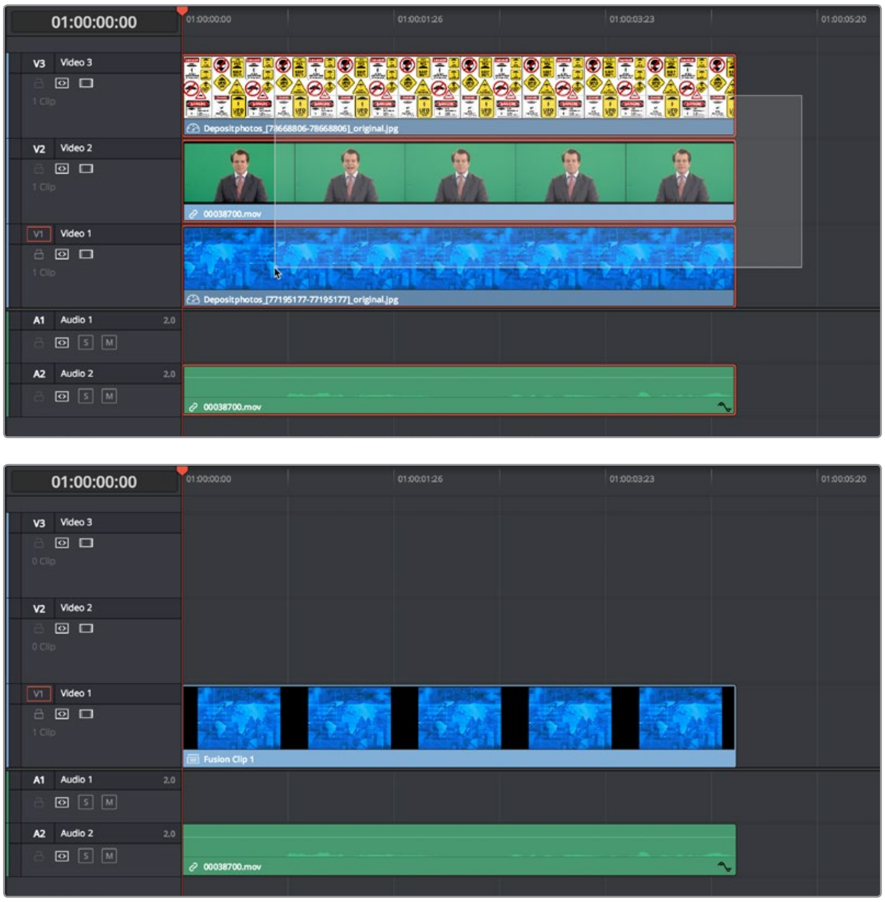

(Top) A stack of clips to use in a composite, (Bottom) Turning that stack into a Fusion clip in the Edit page

When we open the Fusion page to begin work, however, Fusion clips expose their contents in the Node Editor as a pre-built cascade of MediaIn nodes automatically connected by Merge nodes (one Merge node for each pair of clips) that take care of combining each layer of video the way they were in the Edit page Timeline.

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

With this node tree already assembled, we can focus our time on adding the nodes we'll need to each branch of this tree, rather than assembling everything from scratch.

Pulling a Greenscreen Key Using the Delta Keyer

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 be adding after the MediaIn nodes as we work.

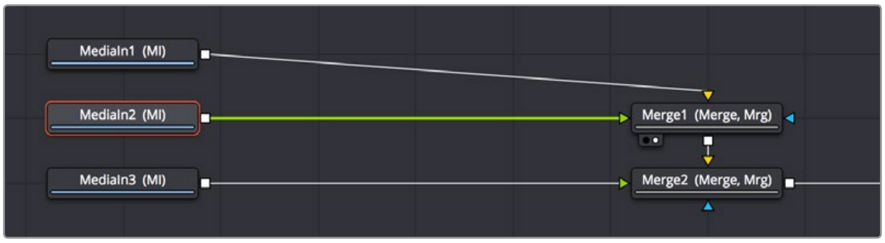

Creating space after the MediaIn nodes, and selecting the second one in preparation for adding a node

Selecting the MediaIn2 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.

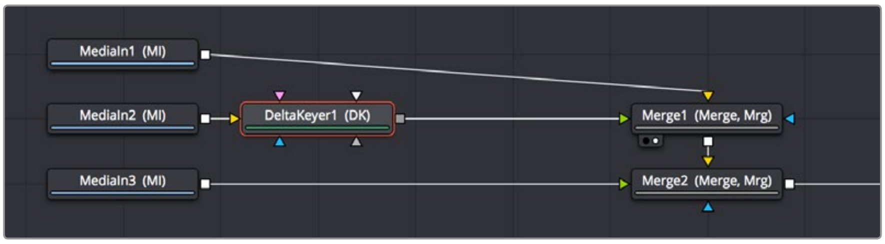

Adding a DeltaKeyer node inline after the MediaIn2 node

The DeltaKeyer node is a sophisticated keyer that is capable of impressive results combining different kinds of mattes and a clean-plate layer, but it can also be used very simply if the background that needs to be keyed is well lit. And once the DeltaKeyer creates a key, it embeds the resulting alpha channel in its output, so in this simple case, it's the only node we need to add. It's also worth noting that, although we're using the DeltaKeyer to key a green screen, it's not limited to 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 a shortcut to quickly sample the shade of green from the background of the image. The shortcut we'll use is a bit unorthodox, but it gives us the ability to preview how different areas of the background will key as we look for the right place to sample.

We hold the Option key down and click the eyedropper tool, and while continuing to hold the Option key down, we drag the pointer over the green of the background in the Viewer.

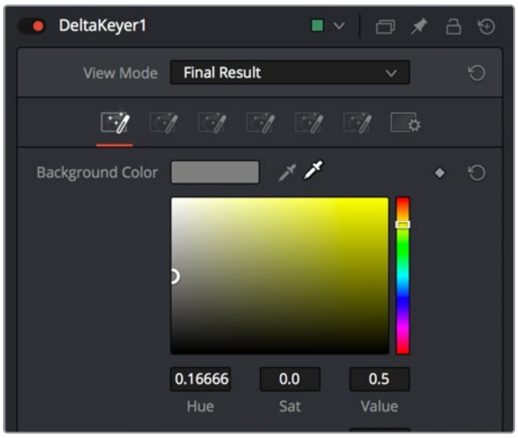

Option-clicking-and-dragging the eyedropper to the Viewer to sample the Background Color

As we drag in the Viewer, an analysis of the color picked up by the location of the eyedropper 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.

(Before) The original image, (After) Sampling the green screen using the eyedropper from the Inspector

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.

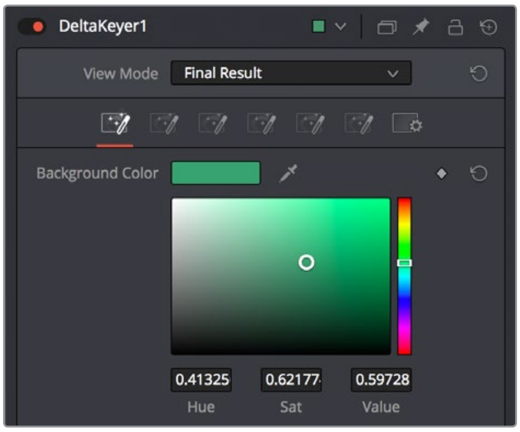

The DeltaKeyer Inspector updates with the sampled color

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 C 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 be using 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.

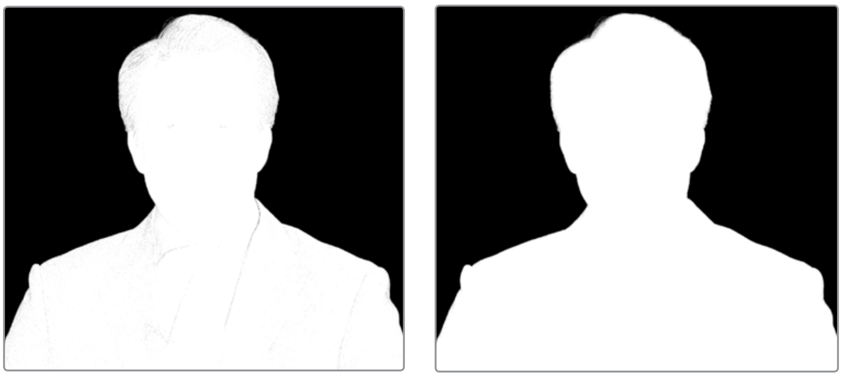

(Before) The original key, (After) The key after using the Clean Foreground slider

With this accomplished, we're good with the key, so we load the Merge1 node back into the Viewer, and press C to set the Color control of the Viewer back to RGB. We can see the graphic in the background, but right now it's too small to cover the whole frame, so we need to make another adjustment.

The final key is good, but now we need to work on the background

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 MediaIn1 node to take care of this. Selecting the MediaIn1 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 MediaIn1 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 onscreen 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's 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 layer we brought into the Fusion page, 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 MediaIn3 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.

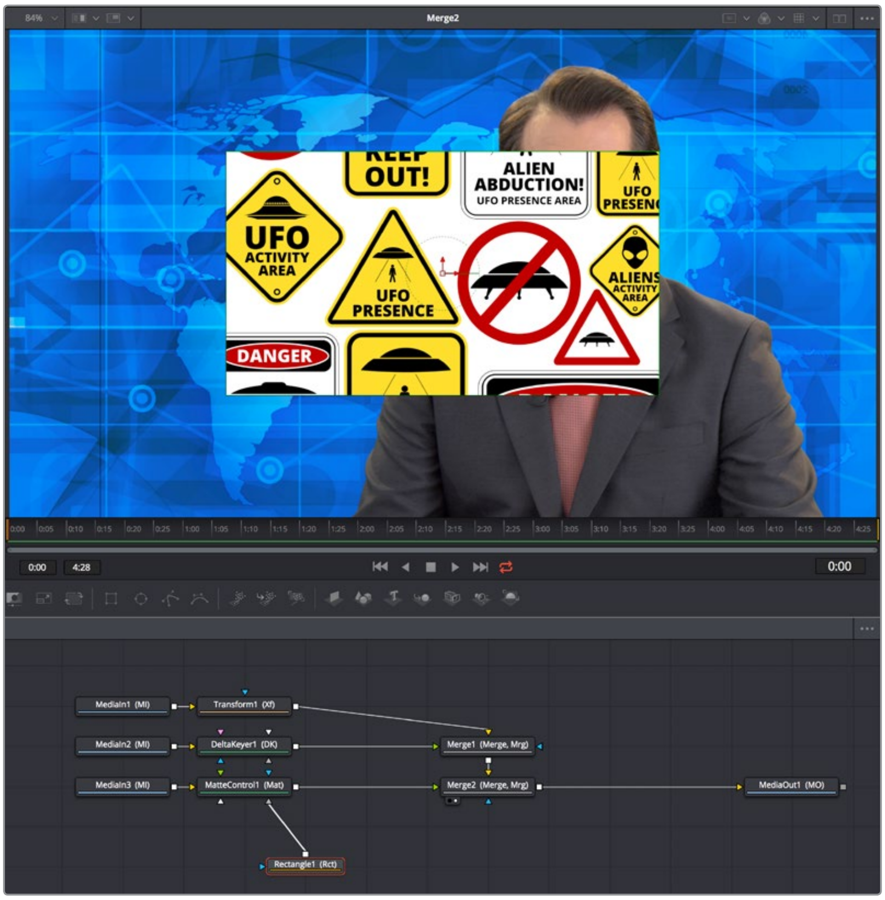

Masking the logo using a Rectangle mask connected to a MatteControl node

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.

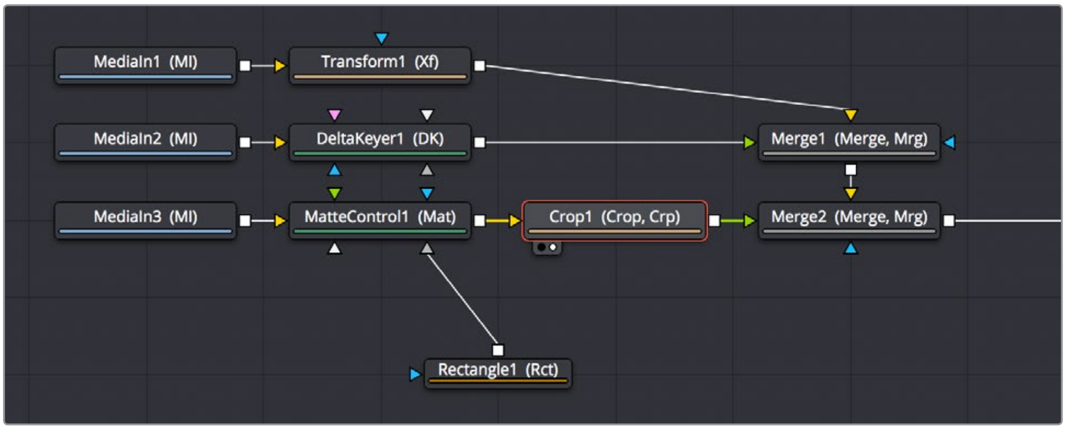

Adding a Crop node after masking the image to center the cropped logo on the frame

With the Crop node selected, we can click the crop tool in the Viewer toolbar

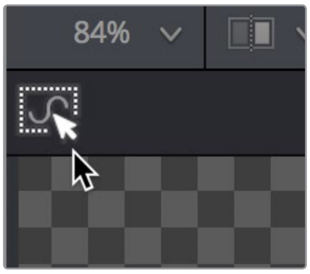

Selecting the crop tool in the Viewer toolbar

This lets us crop the image by dragging a bounding box around it.

(Left) Dragging a bounding box using the Crop tool, (Right) The cropped logo now centered on the frame

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 page functionality, but we've got 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 that parameter up 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.

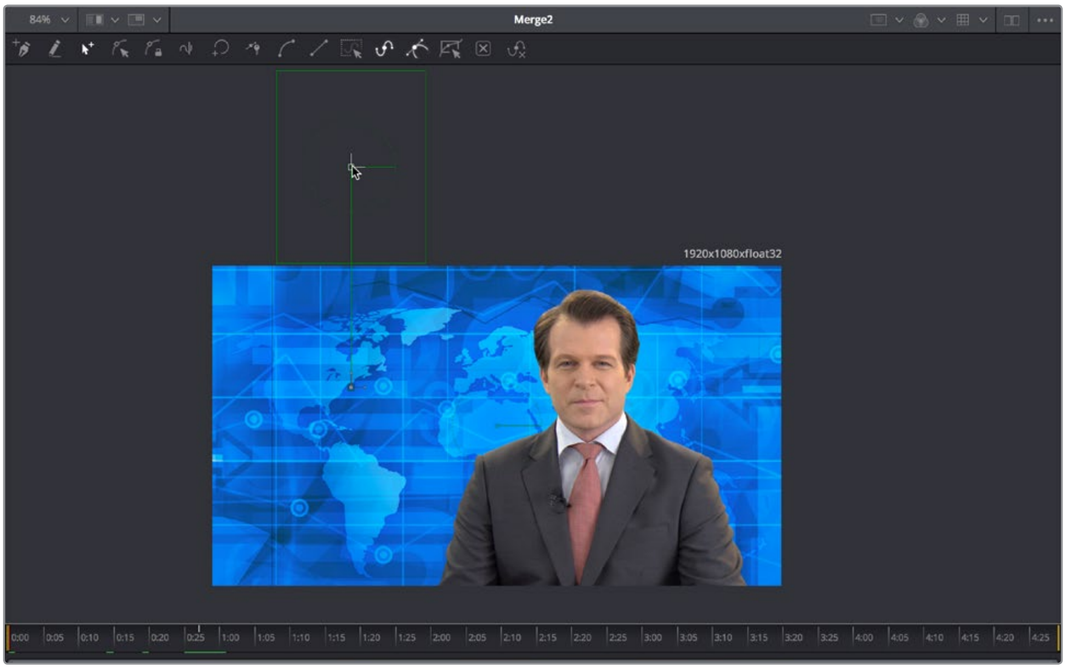

Moving an object in the Viewer to create animation via a motion path

At this point, if we play though 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.

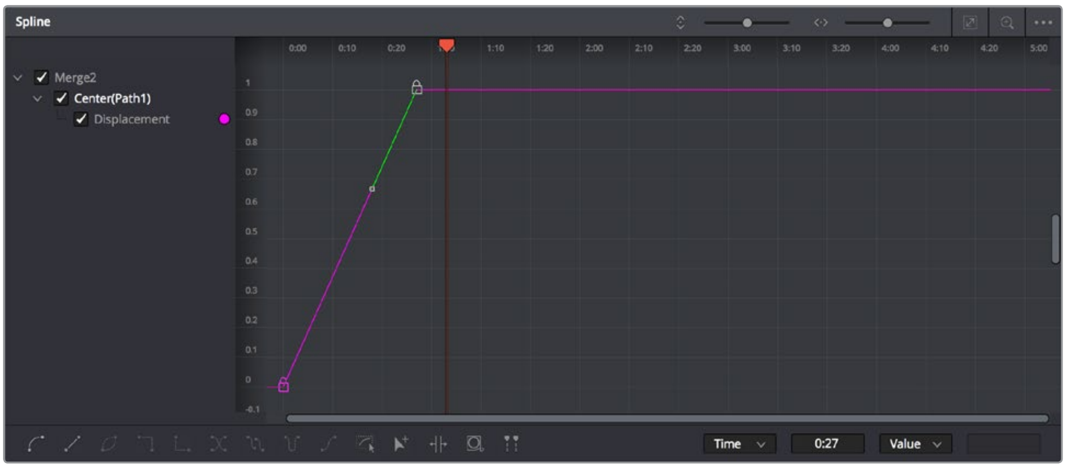

The Displacement curve from the animated Center parameter of the Merge2 node in the Spline Editor

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.

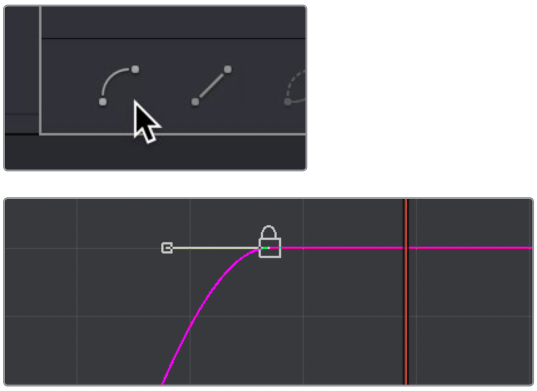

Clicking the Smooth button to turn the selected control point in the graph into a Bezier curve

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.

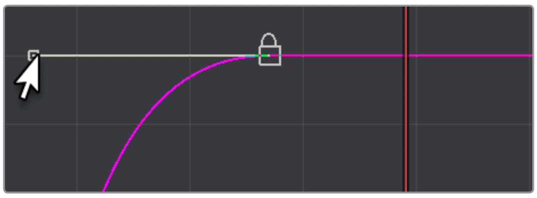

Editing the spline to create a steeper curve, making the logo coast more gradually to a stop

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 the Fusion page is capable of. However, this introduction should have given you a solid look at how to work in the Fusion page so that you can explore further on your own.

Have fun!

Chapter 50

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

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

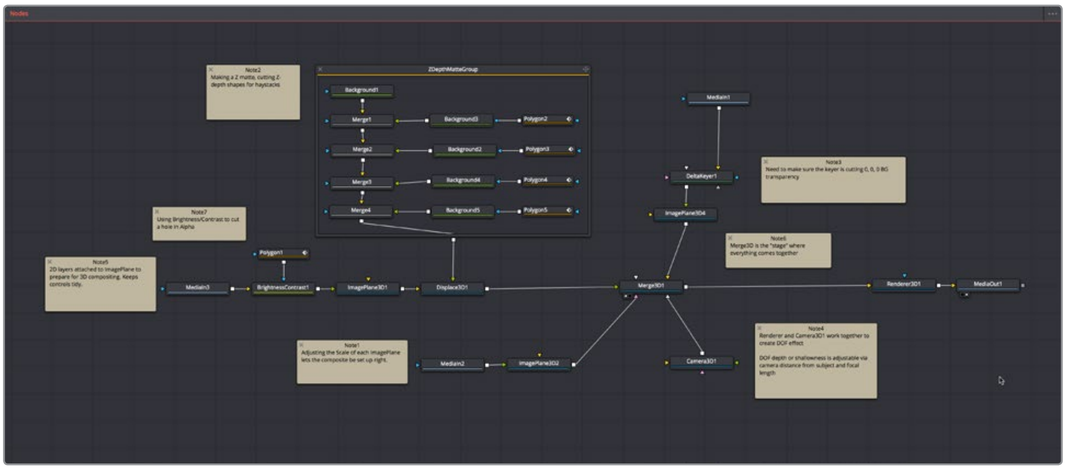

The Node Editor

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

To display the Node Editor:

Click the Nodes button on the UI toolbar.

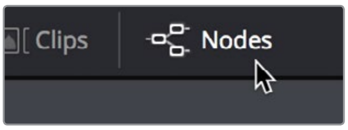

The Nodes button in the UI toolbar

Navigating Within the Node Editor

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

Methods of panning the Node Editor:

- Middle click and drag to pan around the Node Editor.
- Hold Shift and Command down, 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 the Command key down 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 that is not visible in the Node Editor becomes selected, either by using the Find command or by selecting a node's header in the Inspector, the Node Editor will automatically pan to display the node in the visible area.

Using the Node Navigator

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

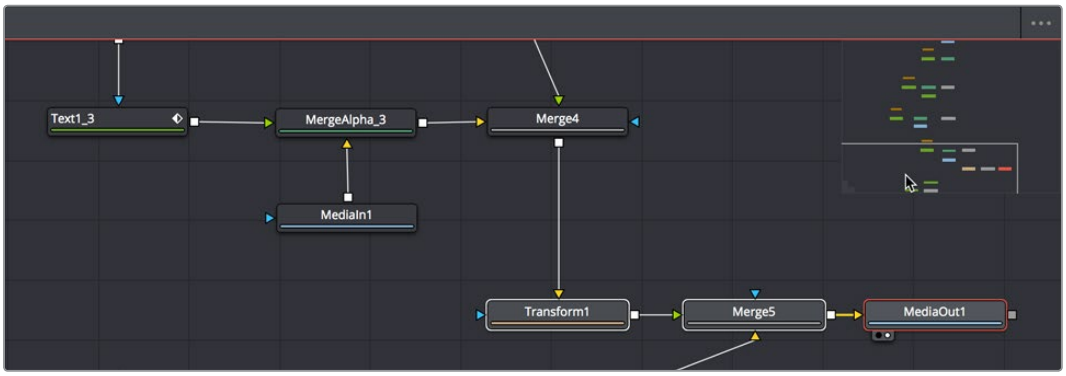

The Node Navigator

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

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

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

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

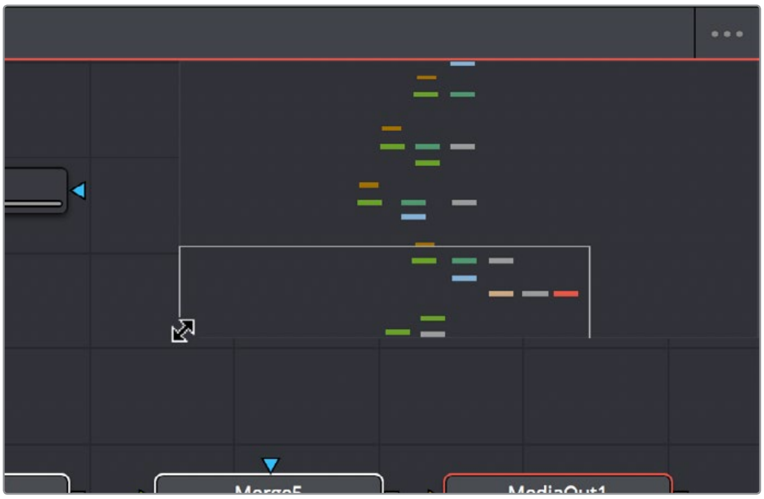

Drag the corner to resize the Navigator

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

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

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

Drag within the Node Navigator 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 Nodes That Are Compatible

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

The Fusion page toolbar

TIP: If you don't know which node a particular icon corresponds to, 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:** De-select all nodes in the Node Editor and then click a toolbar button.

Methods of adding nodes by dragging toolbar buttons:

- **To insert a new node into the node tree:** Drag a toolbar button into the Node Editor and onto the connection line between any two compatible nodes. When the connection highlights as the node is over it, drop the node and it'll be inserted.
- **To create a disconnected node:** Drag a toolbar button into an empty part of the Node Editor. Dragging a toolbar button into the Inspector also creates a disconnected node.
- **To insert a new node after a node loaded into a Viewer:** Drag a toolbar button onto a Viewer to insert a new node after whichever node is viewed, regardless of whether or not 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.

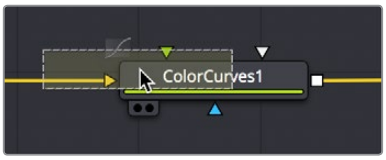

Dragging a node from the toolbar to replace an existing tool.

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 the Fusion page 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.

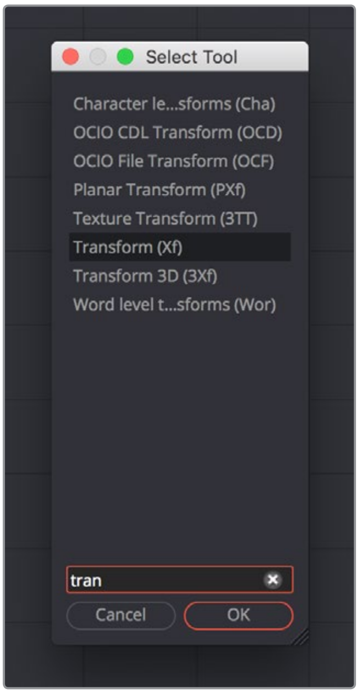

The Select Tool dialog lets you find any node quickly if you know its name

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

To open the Effects Library:

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

The Effects Library appears at the upper left corner of the Fusion page, 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.

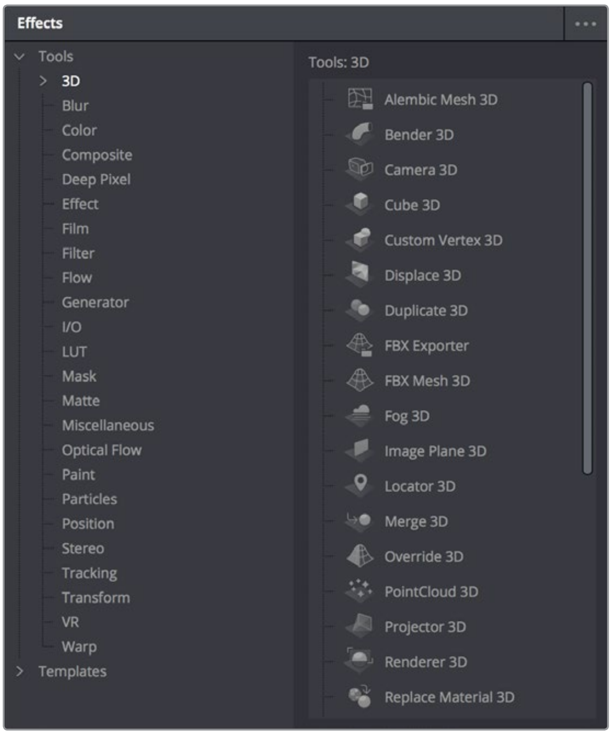

The Tools bin of the Effects Library exposing 3D nodes

By default, the category list shows two primary sets of effects, Tools and Templates, with disclosure controls to the left that hierarchically shows all sub-categories 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:** De-select 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 or not any nodes are selected.

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

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

Adding, Inserting, and Replacing Templates Using the Effects Library

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

Adding a Lens Flare effect

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

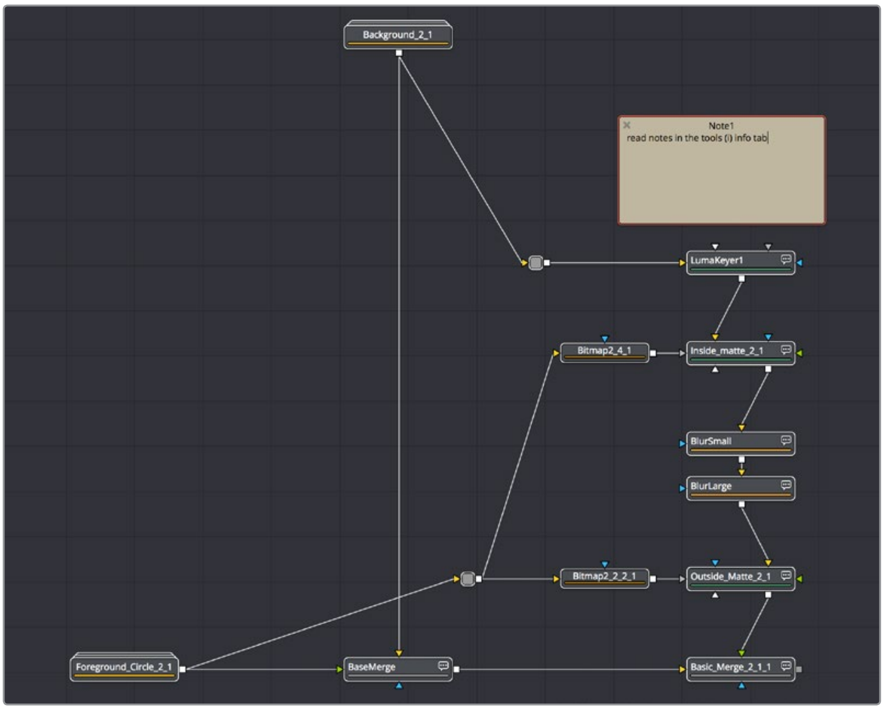

Adding a LightWrap effect from the "How to" bin of the Templates category of the Effects Library

Adding, Inserting, and Replacing Nodes Using the Contextual Menu

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

Methods of adding nodes using the contextual menu:

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

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

Deleting Nodes

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

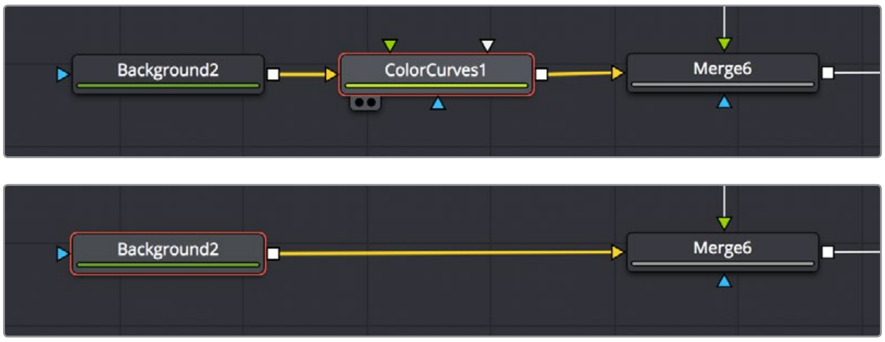

(Top) Before deleting a node from a node tree, (Bottom) After, upstream and downstream nodes have automatically reconnected

Disconnected Nodes

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

Selecting and De-Selecting 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 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.

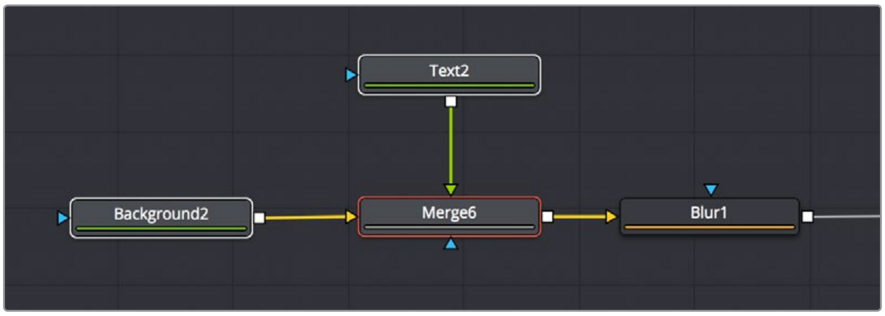

The active node is highlighted orange, while other selected nodes are highlighted white.

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 pretty much as you would expect.

Methods of de-selecting nodes:

- Click once in the background of the Node Editor to deselect all nodes.
- Command-click to de-select 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 MediaIn 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 the Fusion Page

When you first open the Fusion page, the output of the current empty composition (the MediaOut1 node) is usually showing in Viewer2. If you're in Dual-viewer mode, Viewer1 remains empty until you assign a node to one of them. 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 the Fusion page. The left and right indicators correspond to the left and right Viewers, regardless of whether both Viewers are visible, or just one.

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

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

- Clicking an indicator turns it white to show that node is currently loaded in the corresponding Viewer. Clicking it again turns the indicator black and removes it from the Viewer. Nodes only display View indicators if they're currently being viewed. If you want to view indicators, hovering the cursor over the node makes the indicators visible and available for clicking.
- You can also use keyboard short cuts 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 have the 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, 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 off if they're filled.
- Press ` (the tilde key) to empty both viewers.

Create/Play Preview

You can right-click a node, and choose an option from the Create/Preview Play On submenu of the contextual menu to render and play a preview of any node's output on one of the available Viewers.

TIP: Hold the Shift key when selecting the Viewer from the menu to bypass the Render dialog and to start creating the preview immediately using the default settings or the last settings used to create a preview.

Connecting and Disconnecting Nodes

Once you've started to add nodes to your composition, you need to connect them to perform their intended operations.

Node Basics

Each node displays small colored shapes around the edges. One or more arrows represent inputs, and the square represent the tool's processed output, of which there is always only one. Outputs are white if they're connected properly, gray if they're disconnected, or red to let you know that something's wrong and the node cannot process properly.

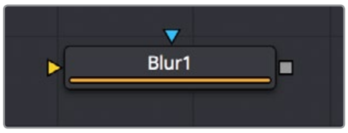

A Blur node with a Foreground Input, Mask Input, and Output

Each node takes as its input the output of the node before it. By connecting a Media In node's output to a Transform node, you move image data from the Media In 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.

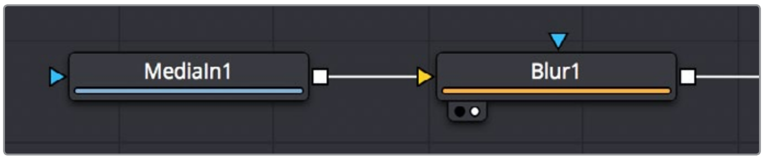

Two nodes connected together

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 order in which you drag node connections is not important; you can just as easily drag a connection from one node's input to another node's output and get the same results.

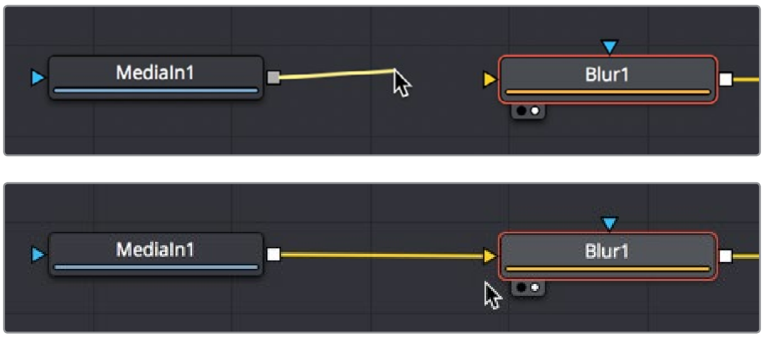

(Top) Before, (Bottom) After dragging a connection line and dropping it to connect two nodes

Dropping Connections On Top of Nodes

To make your life a bit easier, you can also drag a connection line and drop it directly on top of the 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 Media In node and dropped onto the body of a Blur1 node, and the background input is connected first.

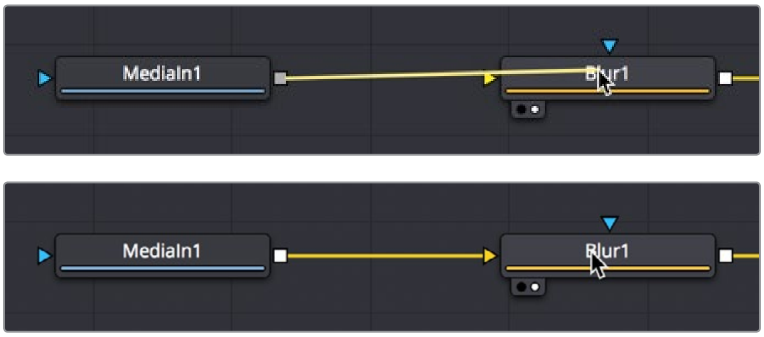

(Top) Before, (Bottom) After dragging a connection line and dropping it on top of a node

If you drop a connection on top of a node that already has the background input connected, then the second most important connection will be attached, which for multi-input nodes is the foreground input, and for other single-use nodes may be the Effects Mask input.

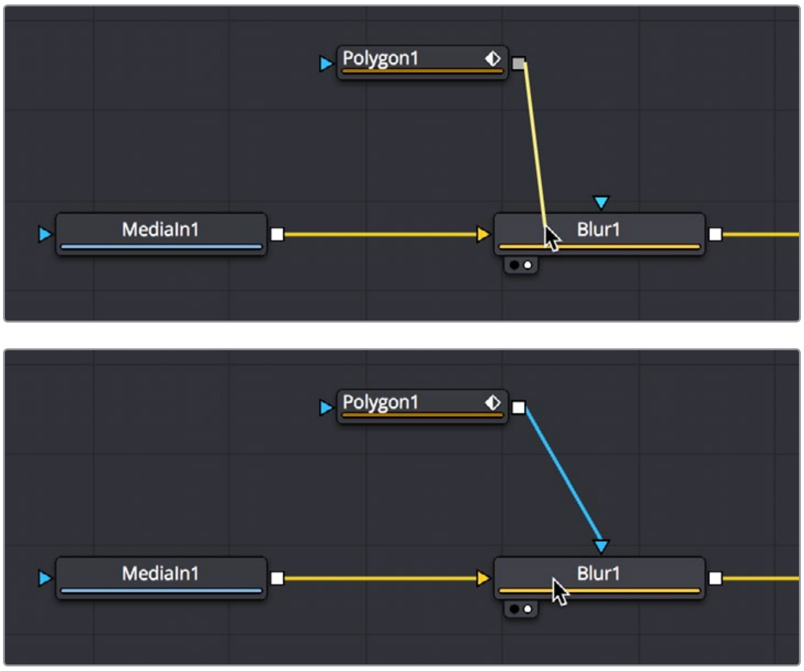

(Top) Before, (Bottom) After dragging a connection line and dropping it on top of a node that has the background input already connected

Some multi-input nodes are capable of adding inputs to accommodate many connections, such as the Merge3D node. These nodes simply add another input whenever you drop a connection onto them.

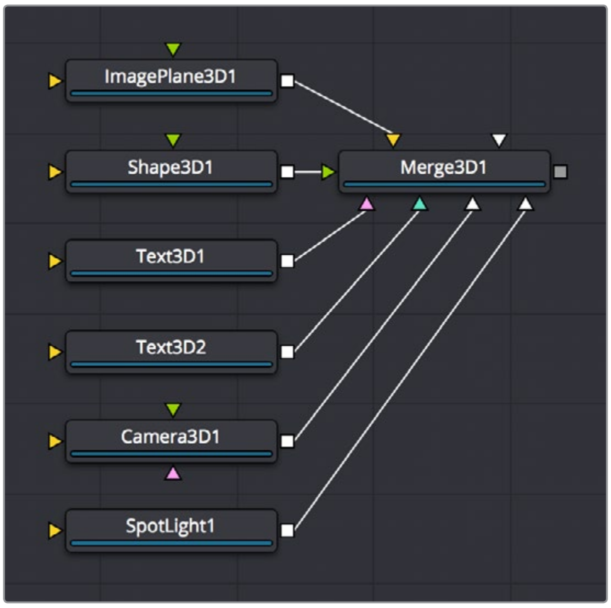

After dragging a connection line and dropping it on top of a Merge3D node

Attaching Connections to Specific Inputs

If you want to make sure you don't attach a connection to the default input of a node, then you need to drop it right on top of the specific node input you want to attach it to. If you can see the input's label in the tooltip bar, then you know you're correctly positioned to make a good connection.

However, there's an alternate method of connecting nodes together in instances where there are several inputs to choose from and you want to make sure you're choosing the correct one. Hold the Option key down 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.

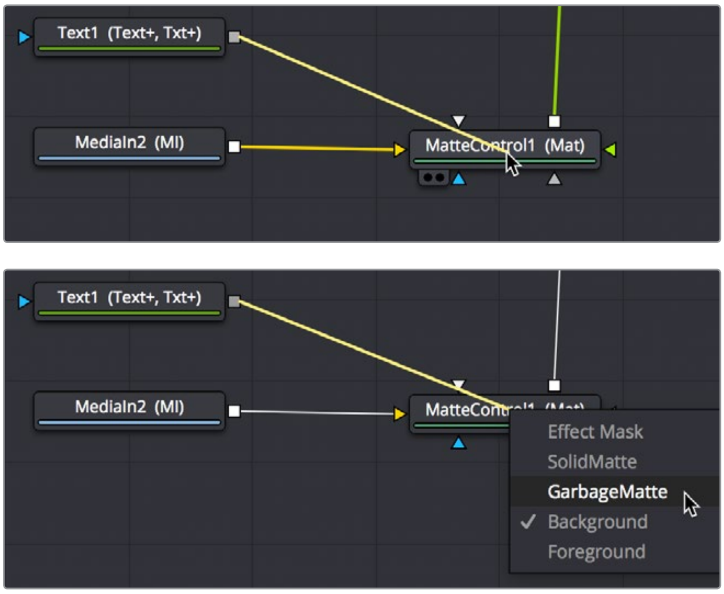

(Top/Bottom) Option-dragging a node connection to drop onto another node exposes a node input menu

Automatically and Manually Attaching Mask Nodes

Mask nodes, such as the Polygon, B-Spline, Ellipse, or Rectangle, have a different automatic behavior when you connect them to other nodes. If you drag a connection from a Mask node onto the body of another node, it will automatically connect itself to the default mask input, which is usually the effect mask input. The assumption is that you're using the mask to limit the node's effect, somehow. However, this isn't always the case, so you'll need to be careful of this behavior to make sure you're attaching your mask to the input that will actually create the effect you need.

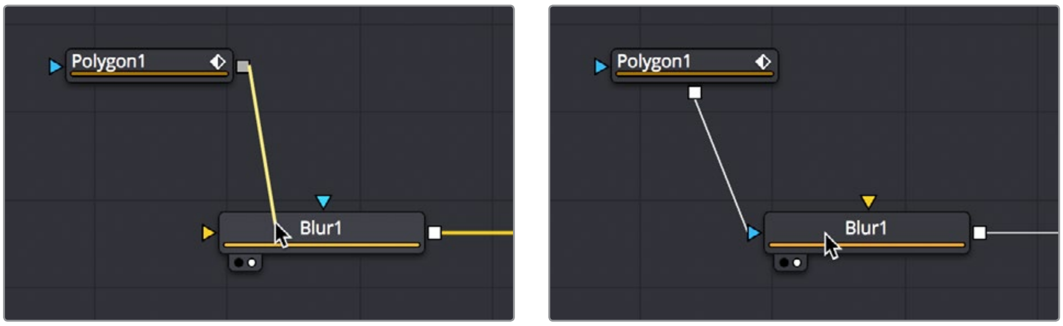

(Left) Before, (Right) After dragging a connection from a mask node and dropping it on top of a MatteControl node

Identifying Node Inputs

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

> **TIP:** Rather than remembering the different knot types press the right mouse button, hold Option and Alt and drag from the output of a node to the center of another tool. When you release the mouse a tooltip will appear allowing you to select the knot you want to connect to.

Node Order Matters

The order in which nodes are attached defines the order in which each image processing operation is applied to the image.

In the following example, a Media In 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 Media Out node, which sends the final image to the Color page for grading.

Adding a Defocus effect first, then the TV node second

As you can see above, connecting the Defocus node first, followed by the TV node, means that while the initial image is softened, the TV effect is sharp. However, if you reverse the order of these two nodes, then the TV effect distorts the image, but the Defocus node now blurs the overall result, so that the TV effect is just as soft as the image it's applied to. The explicit order of operations you apply makes a big difference.

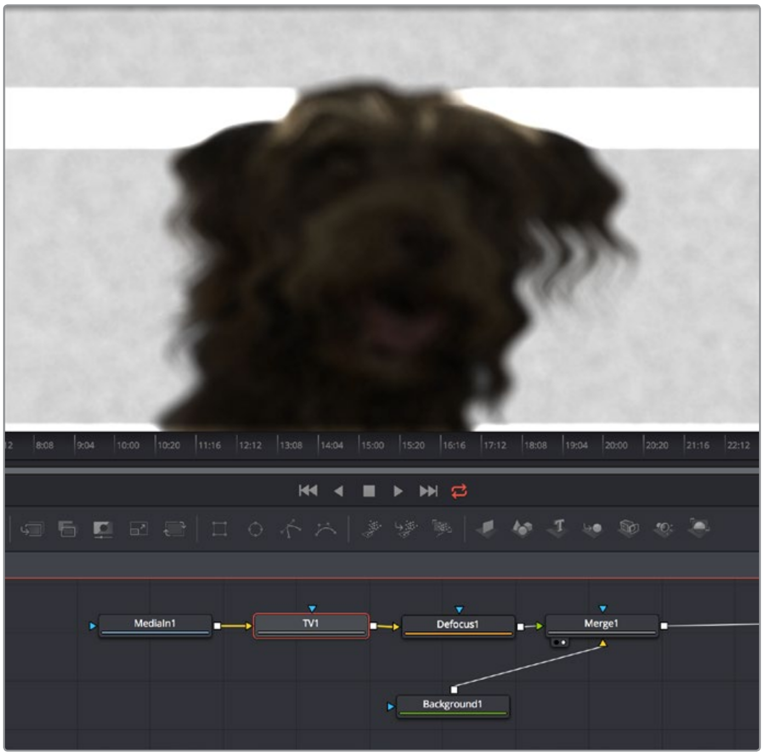

Adding a TV effect first, Defocus second

As you can see, the node tree that comprises each composition is a schematic of operations with tremendous flexibility. Additionally, the node tree structure facilitates compositing by giving you the ability to direct each node's output into separate branches, which can be independently processed and later recombined in many different ways, to create increasingly complex composites while eliminating the need to precompose, nest, or otherwise compound layers together, which would impair the legibility of your composition.

In the following example, several graphics layers are individually transformed and combined with a series of Merge nodes. The end 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.

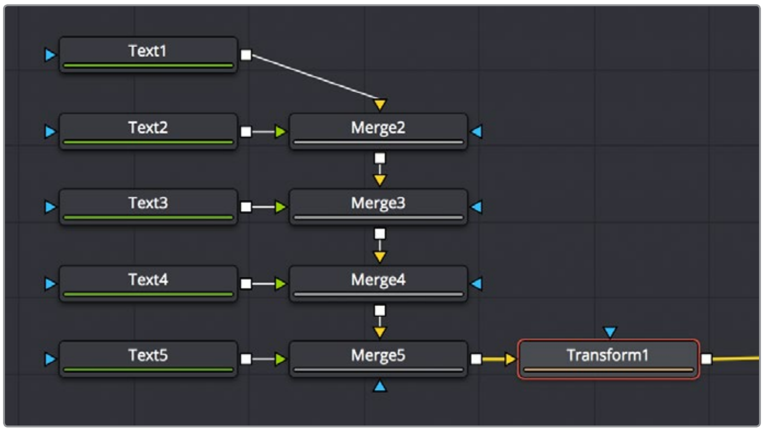

The output of five Text nodes being combined using
Merge nodes is modified by a single Transform node

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.

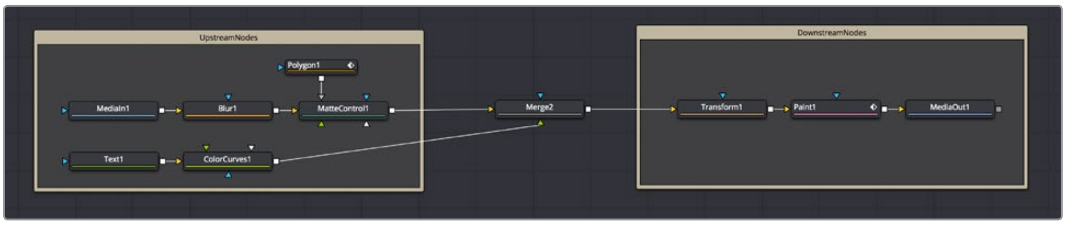

Tools upstream (at left) and downstream (at right) of the Merge node

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.

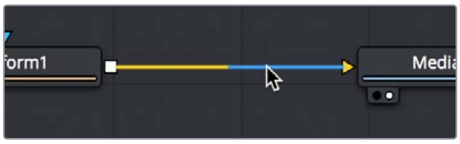

The two halves of a connection line that are revealed when you hover your pointer over it

By clicking and/or dragging these two halves, it's possible to quickly disconnect, reconnect, and overwrite node connections, which is essential to rearranging your node tree quickly and efficiently.

To disconnect two nodes, do one of the following:

- Click once on the input half of the connection between two nodes.
- Click on the input arrow to which a connection is attached, then drag to pull the connection away from the tool and drop it anywhere in an empty area of the Node Editor.

To overwrite a current connection:

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

To reconnect a connection from one node to another:

Drag the output or input half of a connection to disconnect it from one node, and drop the connection directly on another node's input or output.

Tracing Connections Through the Node Tree

Positioning the pointer over a node causes the connections attached to that node to become highlighted, which makes it easier to see which nodes it's attached to. Additionally, highlighted connections display the color of the input they're connected to, which makes it easy to see if they're connected to a foreground, a background, or a particular kind of mask.

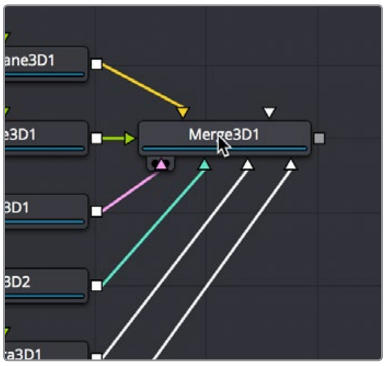

Hovering the pointer over a node highlights the color of all connections, telling you what kinds of inputs are connected

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

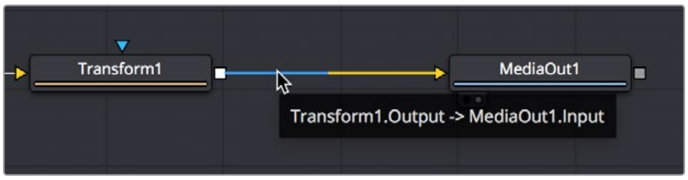

Hovering the pointer over a node highlights the connection between it and other nodes

Branching

Any single node input can only have one connection attached to it. However, a tool's output can be connected to inputs on as many nodes as you require. Splitting a node's output to inputs on multiple nodes is called branching. There are innumerable reasons why you might want to branch a node's output. A simple example is to process an image in several different ways before recombining these results later on in the node tree.

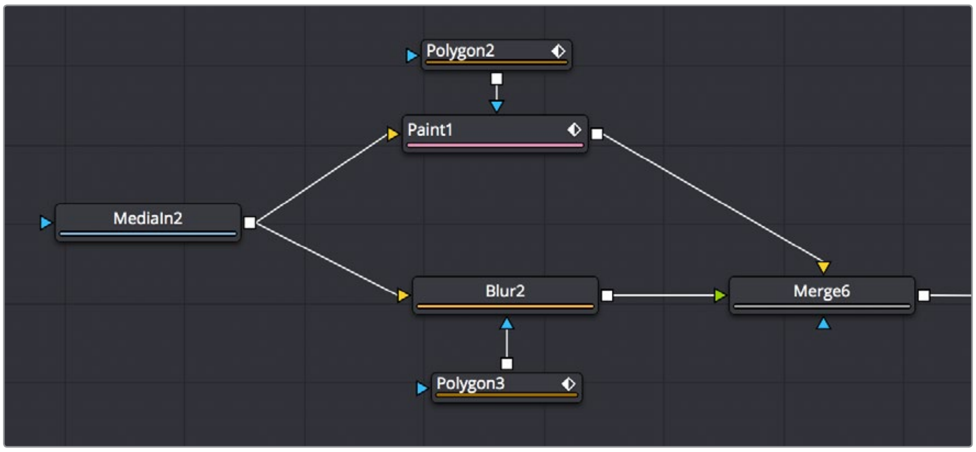

A MediaIn node branched to two node operations and then recombined using a Merge node

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.

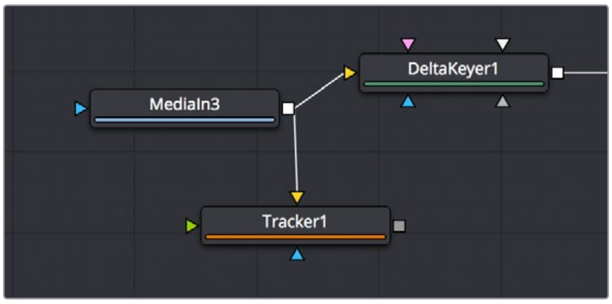

A MediaIn node branched to two different kinds of inputs, used separately

Connecting Merge Nodes

The Merge node is the primary tool available for compositing images together. Each Merge node is capable of combining two inputs to create a third, using standard compositing methods and composite modes. 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 Background input (such as a landscape), and a foreground image with an alpha channel to the Foreground input (such as a graphic with an alpha channel), the Merge node will combine them into a single image for further compositing.

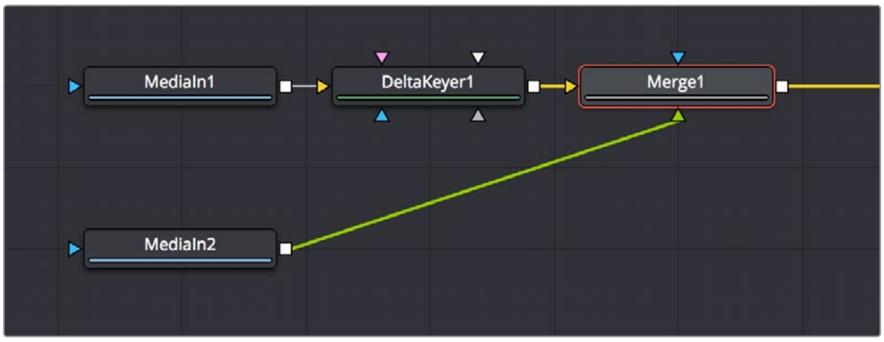

Two MediaIn nodes and a DeltaKeyer node attached to a Merge node creating a composite

Each Merge node has three inputs:

- **Background (green):** The default input. Whichever image is connected to this input defines the output resolution of the Merge node.
- **Foreground (orange):** 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, it's important to bear these inputs in mind when you connect to a Merge node. Of course, you can always drag a connection to a specific input to make sure you're connecting things the way you need. However, if you're in a hurry and you simply drag connections right on top of a Merge node:

- The first connection will be made to the background input.
- The second connection will be made to the foreground input.
- The third connection will be made to the effect mask input.

> **TIP:** When you add a Merge node after a selected node by clicking the Merge button on the toolbar, by clicking on the Merge icon in the Effects Library, or by right-clicking a node in the node tree and choosing Insert Tool > Composite > Merge from the contextual menu, the new Merge node is always added with the background connected to the upstream node coming before it.

Automatically Creating a Merge Node When Adding Nodes

There's a nice shortcut for connecting Merge nodes if you want to connect the incoming clip immediately to your node tree as the top layer of a composite, and that's to drag a clip from the Media Pool or a generator from the Effects Library right on top of any connection line.

When you drop the resulting node, this automatically creates a Merge node, the background input of which is connected to the next node to the left of the connection you dropped the clip onto, and the "foreground input" of which is connected to the new node that represents the clip or generator you've just added.

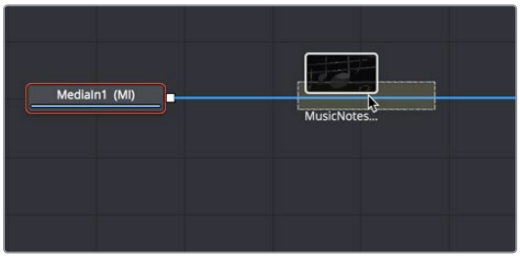

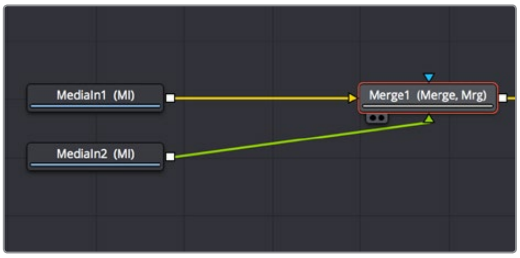

(Left) Dragging a node from the Media Pool onto a connection,
(Right) Dropping it to create a Merge node composite

Additionally, If you drag two or more nodes from the Media Pool into the Node Editor at the same time, Merge nodes will be automatically created to connect them all, making this a fast way to initially build a composite.

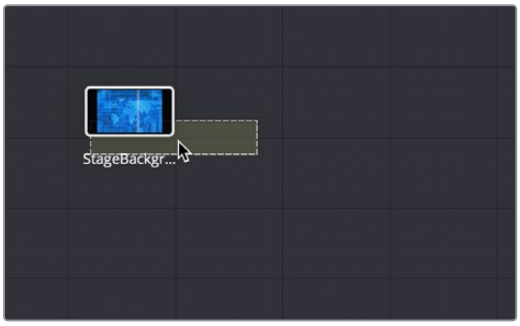

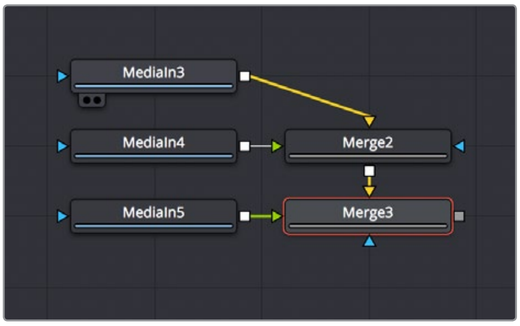

(Left) Dragging three nodes from the Media Pool to the Node Editor, (Right) Merge nodes are automatically created to connect them all

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.

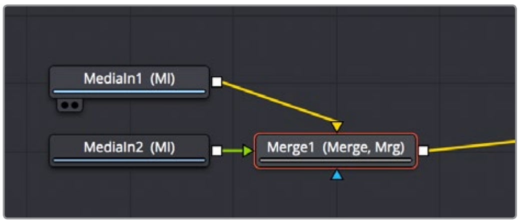

(Left) Dragging a connection from a disconnected node to another
node's output, (Right) Dropping it to create a Merge node composite

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.

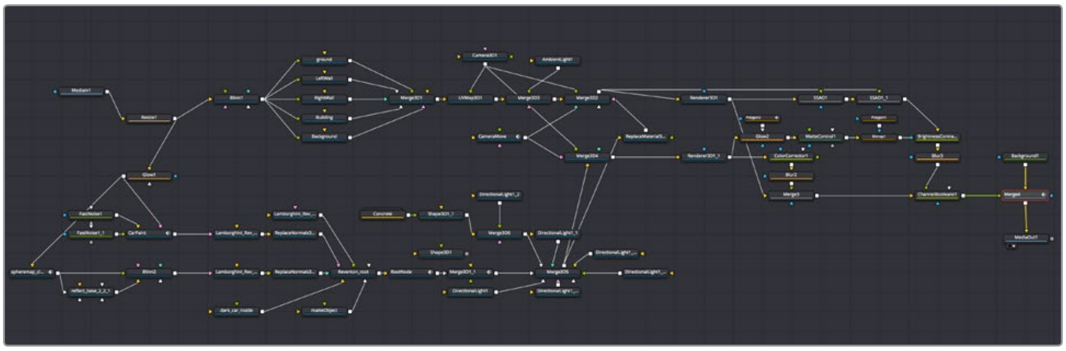

Linear connections between nodes

If you like, you can change how connections are drawn by enabling orthagonal connections, which automatically draws lines with right angles to avoid having connections overlap nodes.

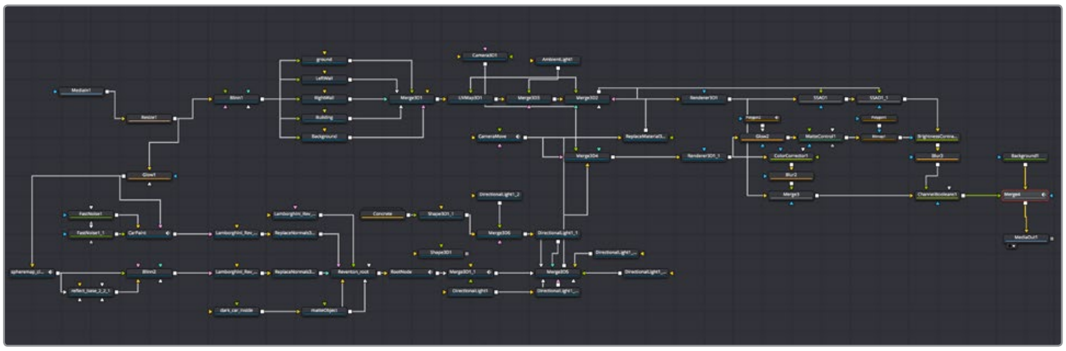

Optional orthogonal connections between nodes

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

To change how connections are drawn in the Node Editor:

- Right-click the Node Editor background and choose one of the following from the contextual menu:
  - Options > Direct Pipes
  - Options > Orthogonal Pipes

Using Routers to Reshape and Branch Connections

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

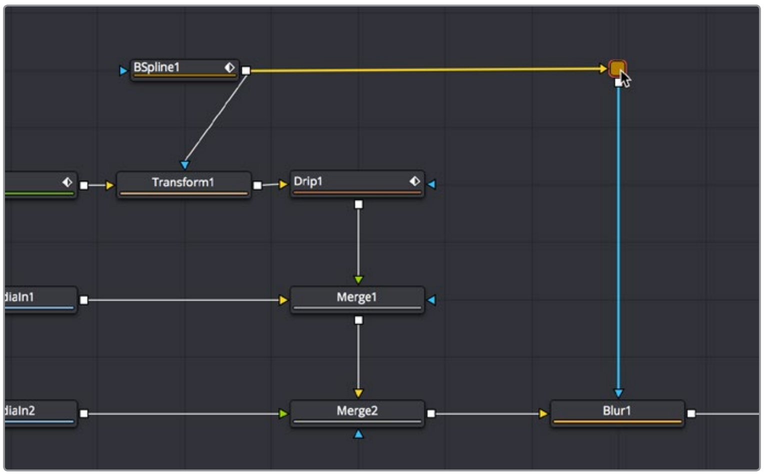

A router added to force a connection to be drawn at an angle

Routers are tiny nodes with a single input and an output, but with no parameters except for a comments field (available in the Inspector), which you can use to add notes about what's happening in that part of the composition.

Even more usefully, you can branch a router's output to multiple nodes, which makes routers even more useful for keeping node trees neat in situations where you want to branch the output of a node in one part of your node tree to other nodes that are all the way on the opposite end of that same node tree.

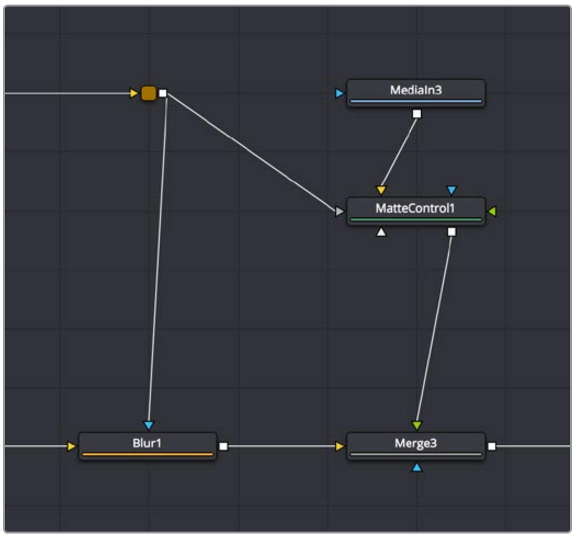

A router branching its output to multiple nodes

Methods of using routers:

- **To add a router to a connection:** Option-click anywhere on a connection.
- **To move a router:** Drag the router to a new location, and the connection will reshape itself as necessary.
- **To branch a router's output:** Drag a connection from the router output to the input of another node. You can branch a router's output as many times as you need to.
- **To remove a router:** Select any router and press the Delete key, or right-click a router and choose Delete from the contextual menu.

Swapping Node Inputs

For multiple-input nodes such as the Merge, Merge 3D, and Dissolve nodes, there's a quick method of swapping the Primary and Secondary inputs, such as the foreground and background inputs of a Merge tool, when you find you've accidentally connected them in the wrong order. If a node has more than two of its inputs connected, only the foreground and background inputs will be swapped.

To swap the primary inputs of a multi-input node, do one of the following:

- Select a node and press Command-T to reverse its inputs.
- Right-click a node and choose Swap Inputs from the contextual menu.

(Left) Before swapping node inputs, (Right) After swapping node inputs, connections don't move but the colors change

Because inputs in the Fusion page 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 (orange) 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 the Shift key down, 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 the Shift key down, 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.

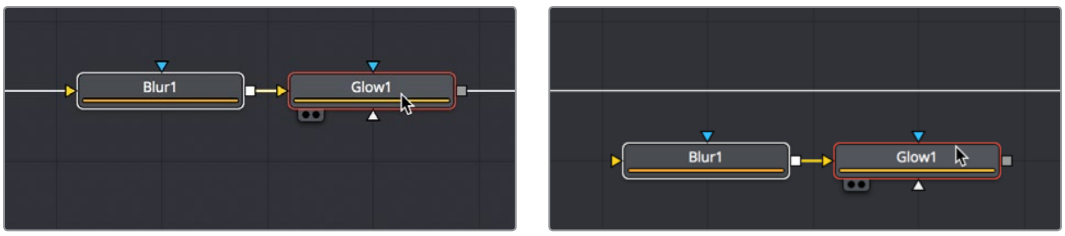

(Left) Before extracting a pair of nodes, (Right) After extracting a pair of nodes

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 the Shift key down 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.

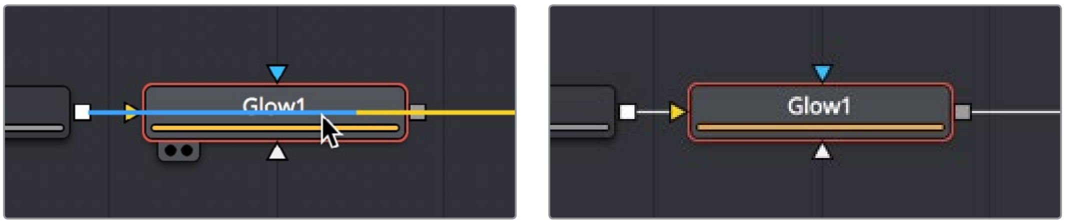

I(Left) Before inserting a node, (Right) After inserting a node

TIP: If you hold the Shift key down, you can extract a node and re-insert it somewhere else with a single drag.

Cut, Copy, and Paste Nodes

The standard operations of cut, copy, and paste are also available in the Node Editor. You can use them to temporarily remove nodes from the Node Editor, create duplicate nodes, or even copy the settings from one node and paste those settings into another node with compatible settings.

Cut, Copy, and Paste in the Node Editor

The standard commands all work, but with some special features specific to the Node Editor.

To copy one or more selected nodes, do one of the following:

- Right-click a node and choose Copy from the contextual menu.
- Choose Edit > Copy from the Edit menu (Command-C).

To cut one or more selected nodes, do one of the following:

- Right-click over the node and choose Cut from the contextual menu.
- Choose Edit > Cut from the Edit menu (Command-X).

When you paste into the Node Editor, you create a copy of the last node or nodes you've cut or copied. When pasting, there are a few different things you can do to control where pasted nodes appear.

To paste one or more selected nodes, do one of the following:

- **To paste 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, 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 Media In 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 the Fusion page. For example, if you copy the following set of three nodes:

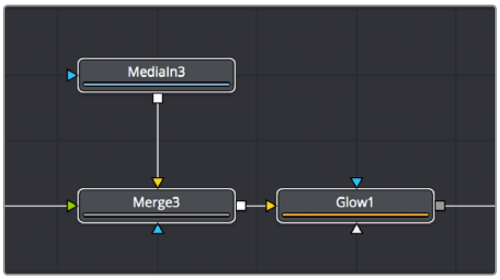

A set of three nodes being copied

And you then paste into a new text editing document, you get the following:

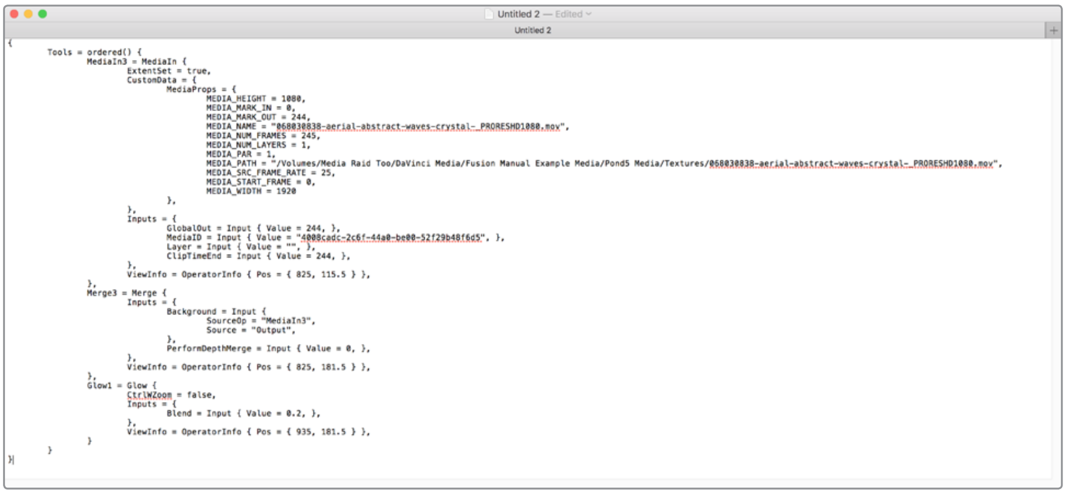

Untitled 2 — Edited

Untitled 2

```
{
	Tools = ordered() {
		MediaIn3 = MediaIn {
			ExtentSet = true,
			CustomData = {
				MediaProps = {
					MEDIA_HEIGHT = 1080,
					MEDIA_MARK_IN = 0,
					MEDIA_MARK_OUT = 244,
					MEDIA_NAME = "068030838-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/068030838-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 = "4008cadc-2c6f-44a0-be00-52f29b48f6d5", },
				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 the Fusion page 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. Alternately, you can right-click directly on a connection line, and choose Paste Instance from the contextual menu.

However you past an instance, the name of that instanced node takes the form "Instance_NameOfNode." If you paste multiple instances, each instance is numbered "Instance_NameOfNode_01."

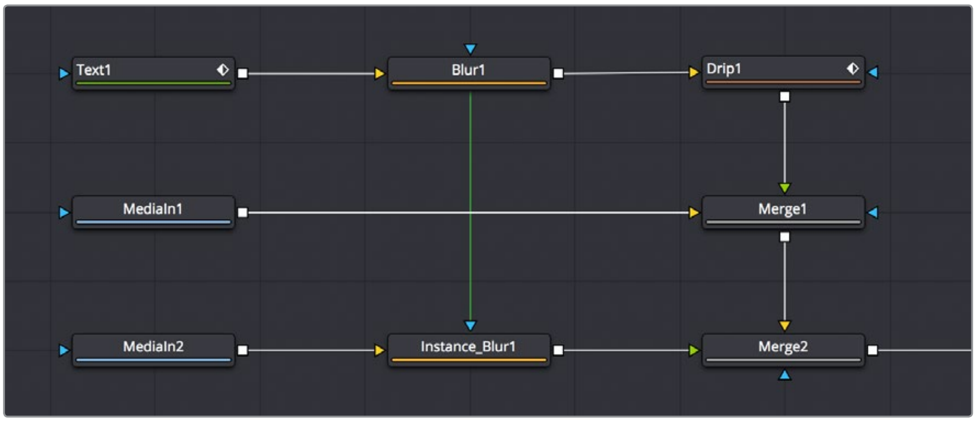

A green link line shows an instanced Blur node's relationship to the original Blur node it was copied from

When a node tree contains instanced nodes, a green line shows the link between the original node and its instances. You have the option to hide these green link lines to reduce visual clutter in the Node Editor.

To toggle the visibility of green instance link lines in the Node Editor:

1. Right click anywhere in the background of the Node Editor.
2. Choose Options > Show Instance Links from the contextual menu.

If you've been using an instance of a node and you later discover you need to use it to apply separate adjustments, you can "de-instance" the node.

To de-instance a node, making it independent:

1. Right-click an instanced node.
2. Choose Deinstance from the contextual menu. That node is now independent from the original node. Once you de-instance a node you cannot re-instance it, but you can undo the operation.

NOTE: If you've de-instanced a node and you cannot undo the operation because you've restarted DaVinci Resolve, you can only recreate an instance by copying the original and pasting an instance again.

De-Instancing and Re-Instancing Specific Parameters

By default, every parameter in an instanced node is linked to the original node, so that any change you make is rippled across. However, from time to time you'll find the need to independently adjust just one or two parameters while keeping the rest of that node's parameters linked. For this reason, instead of de-instancing the entire tool, you can de-instance individual parameters.

To de-instance a single parameter:

Right-click on a parameter's name or value in the Inspector, and choose Deinstance from the contextual menu.

If you've only de-instanced individual parameters, you can re-instance those parameters later on if you change your mind.

To re-instance a single parameter:

Right click on a parameter's name or value in the Inspector, and choose Reinstance from the contextual menu. That parameter immediately inherits the setting of the original node.

Keeping Node Trees Organized

Similar to working with files on your desktop, even the simplest of composites require you to do some amount of organization. In this section we'll look at some basic node operations, some of which you may already be familiar with just from using your computer's operating system or other applications.

Moving Nodes

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 the Fusion > Node Editor > Arrange To Grid or Arrange to Connected checkboxes on.

Commands to "Clean Up" a Node Tree

The grid in the background of the Node Editor can be used to align nodes, either by eye or automatically.

To "clean up" an unruly node tree:

Right-click in an empty section of the Node Editor, and choose "Line Up All Tools to Grid" from the contextual menu. All nodes in the Node Editor will move to align and center themselves along the nearest grid lines.

To "clean up" only one or more selected nodes:

Right-click one of the selected nodes and choose "Line Up to Grid" from the contextual menu. All selected nodes will move to align and center themselves along the nearest grid lines, while all unselected nodes will be left as they are.

Renaming Nodes

Each node that's created is automatically assigned a name (based on its function) and a number (based on how many of that node have been created already). For example, the first Blur node added to a composition will be called Blur1, the second will be Blur2, and so on. Although initially helpful, larger compositions may benefit from important nodes having more descriptive names to make it easier to identify what they're actually doing, or to make it easier to reference those nodes in expressions.

To rename a node:

1. Do one of the following:
   - Right-click a node and choose Rename from the contextual menu.
   - Select a node and press F2.
2. When the Rename dialog appears, type a new name, then click OK or press Return.

NOTE: If multiple nodes are selected, multiple dialogs will appear asking for a name for each tool.

Since the Fusion page can be scripted and use expressions, the names of nodes must adhere to a scriptable syntax. Only use alphanumeric characters (no special characters), and do not use any spaces. Also, you cannot start a node name with a number. If you accidentally create a name that doesn't exactly follow the guidelines, spaces 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. Alternately, 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.

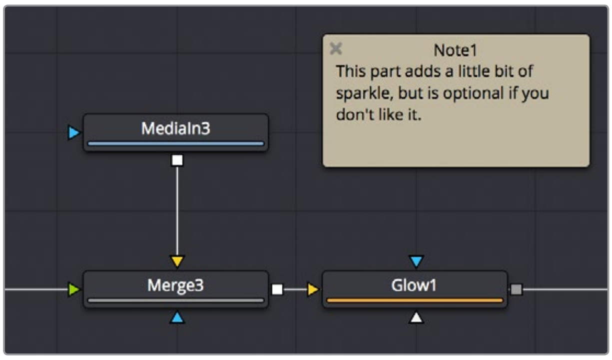

A Sticky Note in the Node Editor

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. 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.
- **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, then type a new name into the Rename dialog and click OK. Alternately, 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, then drag any of the edges or corners to make it larger or smaller.
- To minimize a Sticky Note: Click the close box at the upper left-hand corner of the Sticky Note and it collapses to a small tile.
- To delete a Sticky Note: 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.

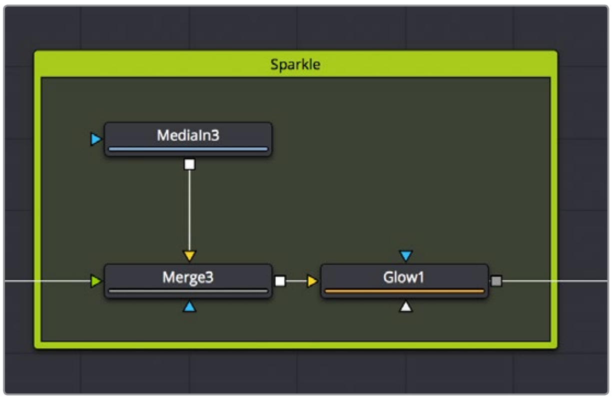

An Underlay in the Node Editor

Underlay Boxes can be named to identify the purpose of that collection of nodes, and they can be colored to be distinct from other Underlay Boxes or to adhere to some sort of color code for your compositions.

Methods of working with Underlay Boxes:

- To create an Underlay Box: Click somewhere in the Node Editor where you want the Underlay Box to appear. Then, press Shift-Spacebar, type "under," and press the Return key when the Underlay Box appears in the Select Tool window. Alternately, 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, 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, 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 of the nodes. Alternately, 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 de-selected, 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.
- **Method one:** With both the Underlay Box and nodes de-selected, drag a bounding box or Command-click to select all nodes in the box you want to remove, and drag them out.
- **Method two:** 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 source or an effect has been added to the Node Editor, it's represented by a node. By default, nodes are rectangular and thin, making it easier to fit reasonably complicated grades within a relatively small area. However, if you like, you can also display node thumbnails.

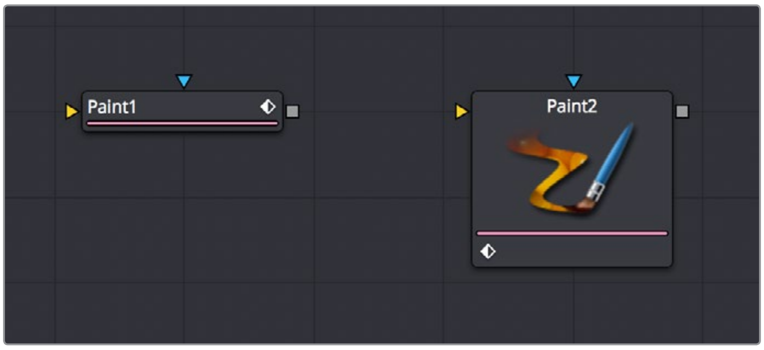

A node in the Node Editor shown both with and without a thumbnail

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 options for which nodes should display thumbnails in the contextual menu that appears when you right-click anywhere in the background of the Node Editor.

Force All Tile Pictures

This option shows thumbnails for every single node in the Node Editor. This can make simple node trees easier to read, but it'll make all node trees take up considerably more room.

NOTE: If Show Thumbnails is enabled, nodes may not update until the playhead is moved in the Time Ruler.

Force Active Tile Pictures

You may also choose to only show thumbnails for nodes that are currently selected, which can make it easier to see which node you're working on. When nodes become deselected, the thumbnails will be hidden again.

Force Source Tile Pictures

This enables thumbnails for all MediaIn 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 those nodes are selected or not.

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 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.

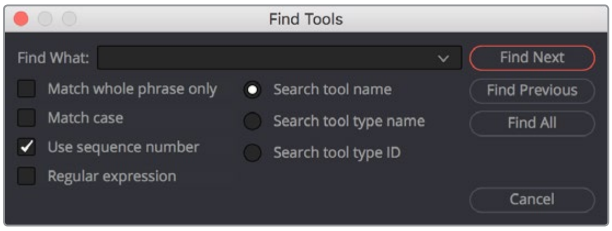

The Find dialog lets you quickly locate nodes wherever they are in the Node Editor

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 and select a downstream node matching the criteria.
   - Click Find Previous to try and select an upstream node matching the criteria.
   - Click Find All to try and 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 of the nodes based on the search term, and you can temporarily disable them by pressing the short cut for Bypass, Command-P.

Using Regular Expressions

If you need to do more complicated searches, you can turn on the Regular Expression checkbox, which lets you enter some simple expressions with which to create more complex find operations. Some useful examples of regular expressions that are valuable include the use of Character Sets.

Character Sets

Any characters typed between two brackets [ ] will be searched for. Here are some examples of character set searches that work in the Fusion page.

`[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 upper case T or a lower case t

`[aeiou]`

Finds: Every vowel

`[0-9]`

Finds: Every numeral

`[5-7]`

Finds: Every numeral from five to seven, 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 for facilities where multiple compositing artists use a common set of facility defaults, stored somewhere that's commonly accessible. The default paths are:

- For macOS systems, this path defaults to /UserName/Library/Application Support/Blackmagic Design/DaVinci Resolve/Fusion/Defaults.
- For Windows systems, this path defaults to C:\Users\<username>\AppData\Roaming\Blackmagic Design\DaVinci Resolve\Fusion\Defaults..
- For Linux systems, this path defaults to ~/.fusion/BlackmagicDesign/DaVinci Resolve/Fusion/Defaults..

If you browse this directory, the settings for each node are saved using a name taking the form INTERNALNAME_PUBLICNAME.settings, where INTERNALNAME is the internal name of the Fusion tool, and PUBLICNAME is the name of the Node that's derived from the internal Fusion tool. For example, the default setting for a Blur node would be called Blur_Blur.setting. This naming convention is partly to ensure that third-party 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, then click Open. Those settings are now applied to that node.

Adding Saved Settings From the File System

Saved settings in your File system can also be used to create new nodes by dragging the . setting file into the Node Editor from a standard file browser. Once dropped, that setting turns into a new node.

TIP: If you drop a setting directly onto a connection line, the new node will be inserted onto that connection.

Node Modes Including Disable and Lock

Right-clicking one or more nodes and opening the contextual menu reveals a series of commands in the Modes submenu, some with accompanying keyboard shortcuts, that let you set control visibility, disable, lock, update, and cache nodes.

- **Show Controls:** Sets whether or not that node reveals its parameters in the Inspector when its 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 re-enabled. While disabled, the last processed image for that node will be displayed as a freeze frame. One example of when this is useful is when you have a large or processor-intensive composition (such as a particularly intense particle system), and disabling this option temporarily will let you quickly make several quick parameter adjustments to different nodes without forcing you to wait for the node tree to re-render after every adjustment. Another example of when this is useful is when you want to quickly see the effect of animated downstream nodes while keeping upstream nodes that are too processor-intensive to play in real time from rendering additional frames.
- **Force Cache:** When enabled, this node's output for the current frame has an extremely high cache priority, essentially forcing it to stay cached in memory. Off by default.

Toggling any one of these node modes displays a badge within that node indicating its state.

Node Editor Options

Right-clicking in an empty area of the Node Editor will bring up the contextual menu and the Options submenu. The Options submenu contains several choices that can be used to customize how the Node Editor looks and behaves.

- **Pipes Always Visible:** Enabling this option causes a connection to cross over a node instead of beneath it, sometimes making it easier to follow the connection's path.
- **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.

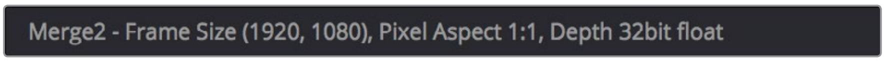

The Status bar located beneath the Node Editor

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 additional give you information about the Domain (Image and DoD), and the data range used by that clip.

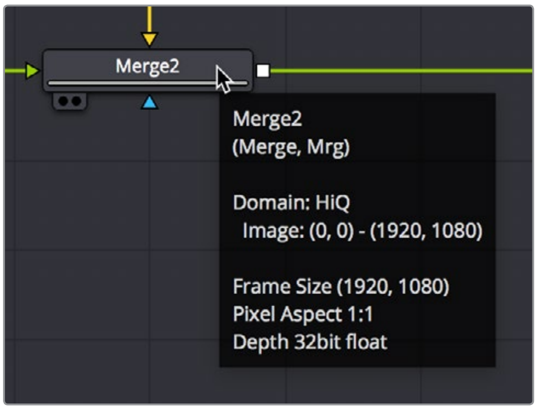

The floating tooltip showing node information that appears within the Node Editor

Chapter 51

Node Groups, Macros, and Fusion Templates

This chapter reveals how to use groups, macros, and templates in the Fusion page so working with complex effects becomes more organized, more efficient, and easier.

Contents

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 sub groups 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).

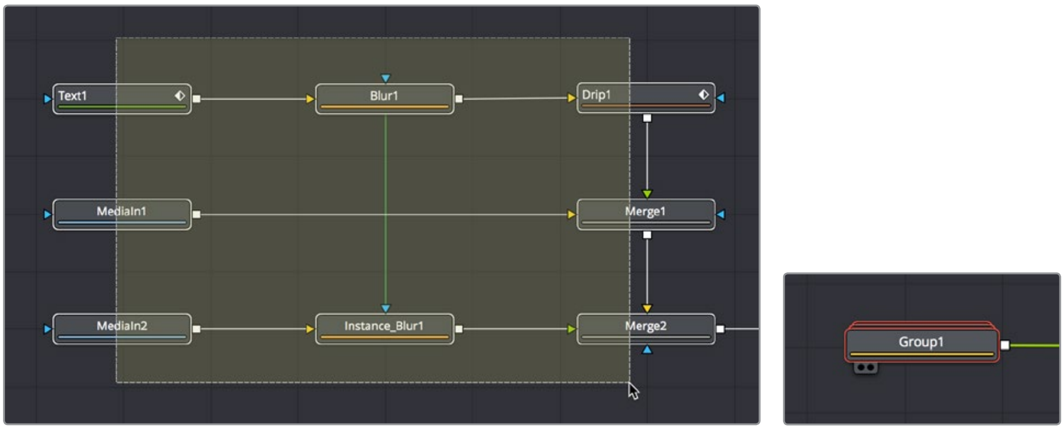

(Left) Several nodes selected in preparation for making a group, (Right) The resulting group

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 it by double-clicking it, or by selecting the group node and pressing Command-E.

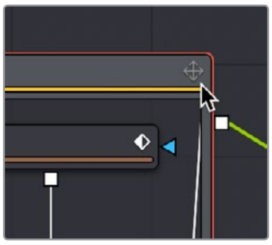

An open group window showing an expanded group

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 short cut.

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 this Position button on lets you size group nodes independently of the rest of the node tree.

Ungrouping Nodes

If you decide you no longer need a particular group, or you simply find it easier to have constant access to all the nodes in the group at once, you can decompose or "ungroup" the group without deleting the nodes within it to eliminate the group but keep the contents in the Node Editor.

To Ungroup nodes, do the following:

1. Right click on the group.
2. Choose Ungroup from the contextual menu. The nodes inside the group are placed back in the main node tree.

Saving and Reusing Groups

One of the best features of groups is that every group and its settings can be saved for later use in other shots or projects. Groups and their settings can be recalled in various ways.

A good example of when you might want to Save and Load 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 re-using 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.

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 re-use 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 /Users/UserName/Library/Application Support/Blackmagic Design/DaVinci Resolve/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 /Users/UserName/Library/Application Support/Blackmagic Design/DaVinci Resolve/Fusion/Macros/ directory.

Other Macro Examples

Macros can also be used as LUTs. Just copy the macro's .setting file to the LUTs: folder, and the macro will be selectable in the Viewers as a LUT. These LUT macros can be used for more than just a color adjustment; you could make a macro that does YUV 4:2:2 resampling, a resize, a sharpening filter, or just watermarking.

Creating Fusion Templates

The integration of Fusion into DaVinci Resolve has enabled the ability to use Fusion Titles in the Edit page. Fusion Titles are essentially generators that were created in the Fusion page and which can be edited into the Timeline of the Edit page as clips with custom controls. However, the really exciting thing about this is that you can create your own Fusion Title templates, from nearly any Fusion page composition you can build using Fusion-generated objects such as Text+ layers, Fusion generators, and even 3D geometry and 3D text with texture and lighting effects. 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 open.

Building a composition to turn into a template

Create a Macro

Macros are basically Fusion compositions that have been turned into self-contained nodes. Ordinarily, these nodes are used as building blocks inside of the Fusion page so that you can turn frequently-made compositing tricks that you use all the time into your own nodes. However, we can also use this macro functionality to build templates for the Edit page.

Having built your composition, select every single node you want to include in that template except for the MediaOut1 node.

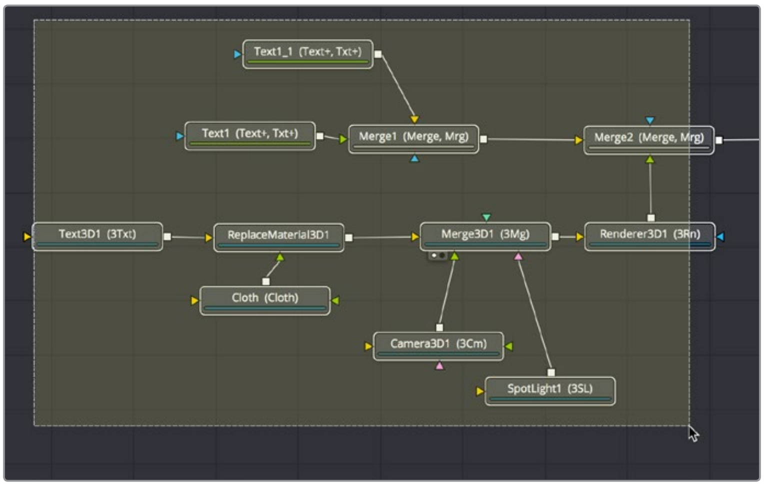

Selecting the nodes you want to turn into a template

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.

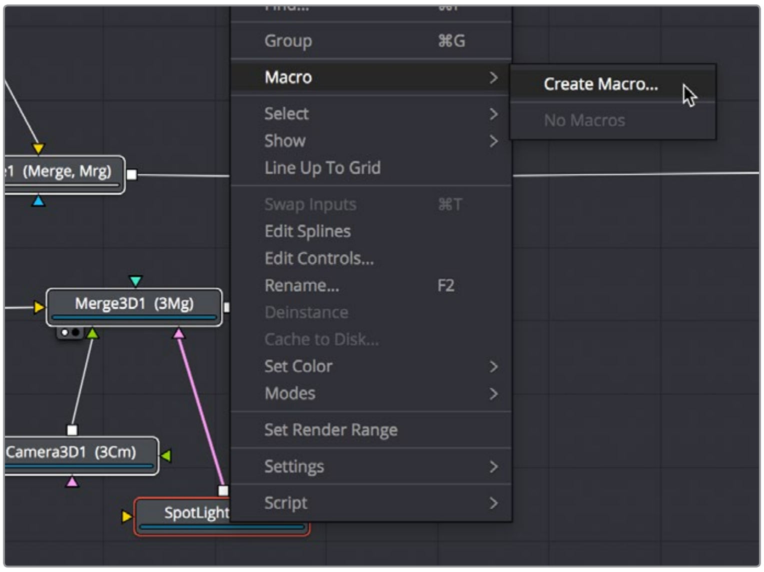

Creating a macro from the selected nodes

The Macro Editor window appears, filled to the brim with a hierarchical list of every parameter in the composition you've just selected.

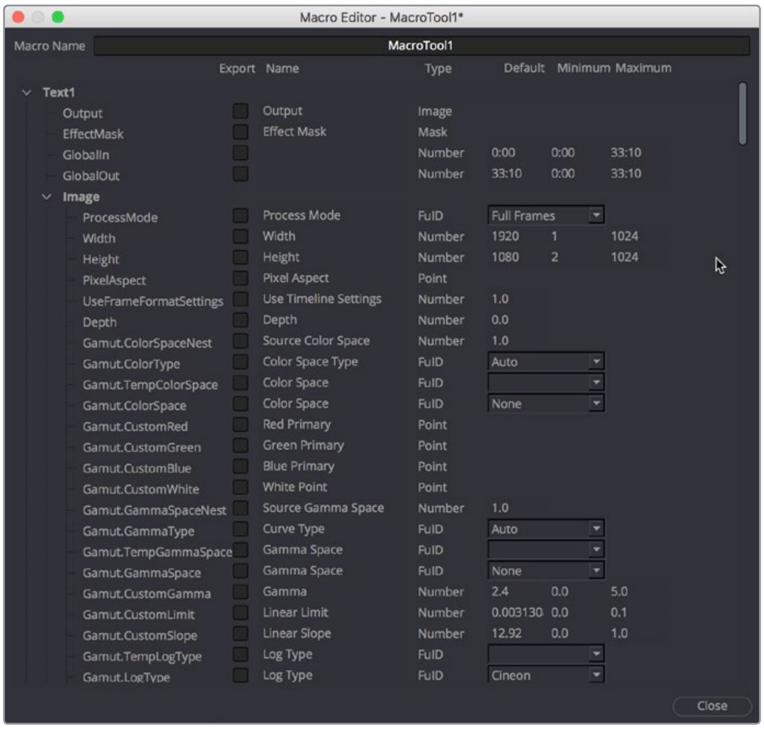

The Macro Editor populated with the parameters of all the nodes you selected

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

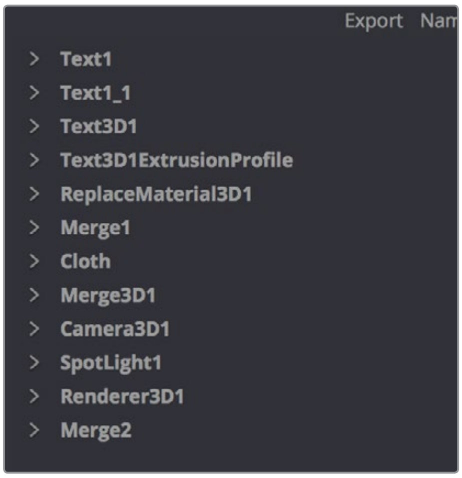

A simple list of all the nodes we've selected

Closing the top node's parameters reveals a simple list of all the nodes we've selected. The Macro Editor is designed to let you choose which parameters you want to expose as custom editable controls for that macro. Whichever controls you choose will appear in the Inspector whenever you select that macro, or the node or clip that macro will become.

So all we have to do now is to turn on the checkboxes of all the parameters we'd like to be able to customize. For this example, we'll check the Text3D node's Styled Text checkbox, the Cloth node's Diffuse Color, Green, and Blue checkboxes, and the SpotLight node's Z Rotation checkbox, so that only the middle word of the template is editable, but we can also change its color and tilt its lighting (making a "swing-on" effect possible).

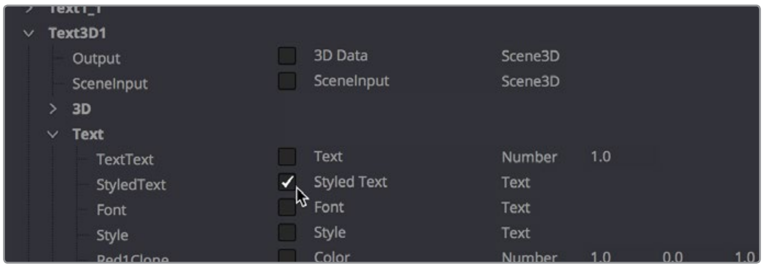

Turning on the checkboxes of parameters we'd like to edit when using this as a template

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. If we're using macOS, we navigate to the /Library/Application Support/Blackmagic Design/DaVinci Resolve/Fusion/Templates/Edit/Titles directory, enter a name in the field below, and click Save.

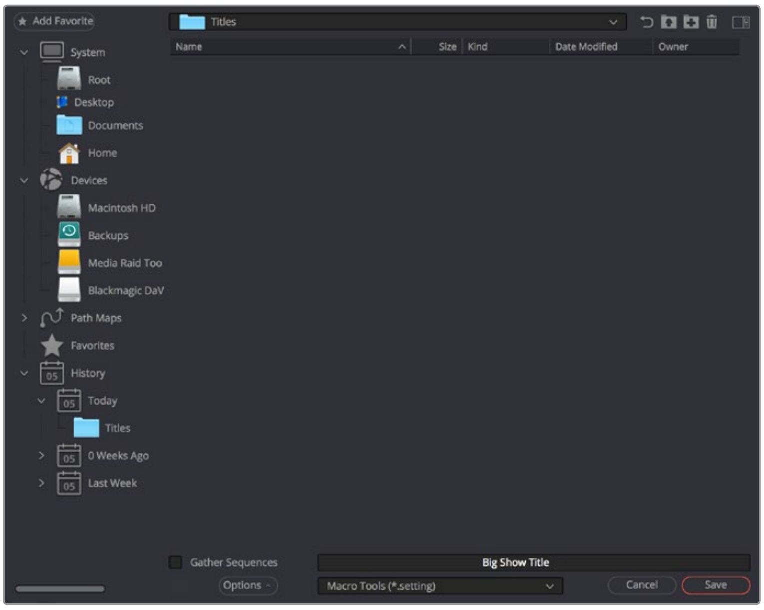

Choosing where to save and name the Macro

Restart DaVinci Resolve and Use Your New Template

After you've saved your macro, you'll need to quit and re-open DaVinci Resolve. When you open the Effects Library of the Edit page, you should see your new template inside the Titles category, ready to go in the Fusion Titles list.

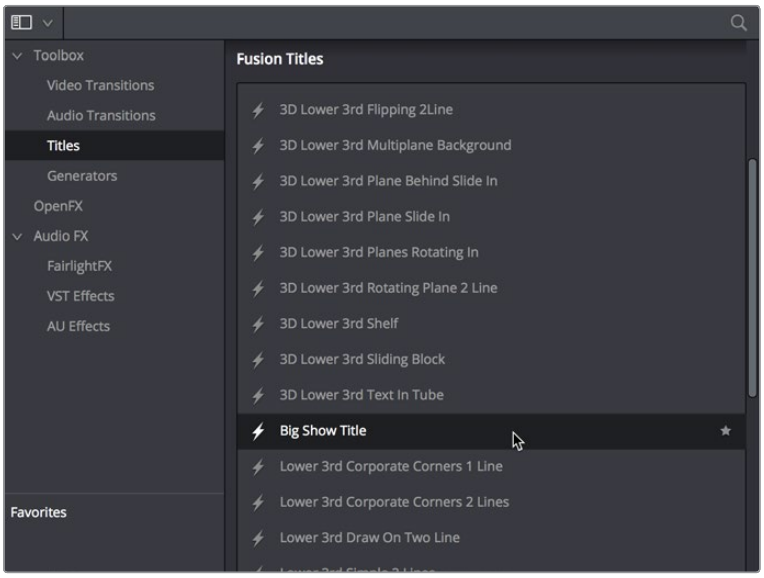

Custom titles appear in the Fusion Titles section of the Effects Library

Editing this template into the Timeline and opening the Inspector, we can see the parameters we enabled for editing, and we can use these to customize the template for our own purposes.

Customizing the template we made

And that's it!

Chapter 52

Using Viewers

This chapter covers working with viewers in the Fusion page, 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

Viewers in the Fusion page 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, (Left) a 3D viewer, (Right) a 2D viewer

Additionally, you can expose "subviews" including color inspectors, magnifiers, waveforms, histograms, and vectorscopes to help you analyze the image as you work.

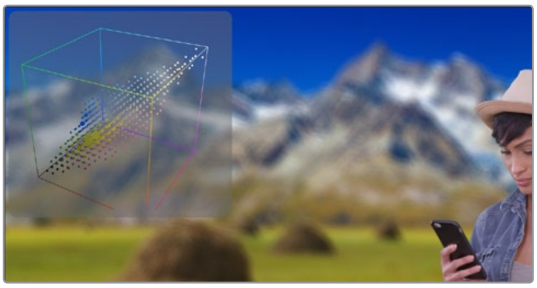

Viewer with a 3D Histogram subview at the upper left-hand corner

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

At the time of this writing, video output is not available from the Fusion page.

Loading Nodes Into Viewers

When you first open the Fusion page, the output of the current empty composition (the MediaOut1 node) is usually showing in Viewer2. If you're in Dual-viewer mode, Viewer1 remains empty until you assign a node to one of them.

To load specific nodes into specific viewers:

- Hover the pointer over a node, and click one of two buttons that appear at the bottom-left of the node.
- Click once to select a node, and press 1 (for the left Viewer) or 2 (for the right Viewer).
- Right-click a node and choose View On > None/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, 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, 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 scroll bar between the Viewers and the Work area below them.

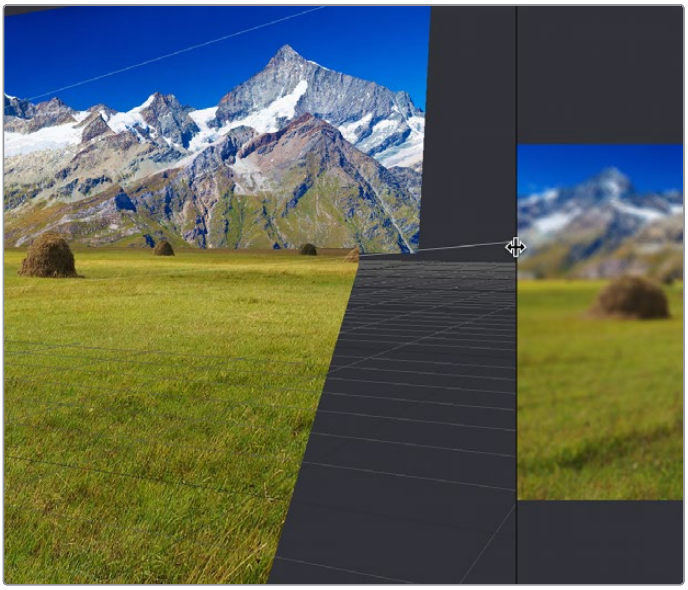

The Viewer Divider bar

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 Shift and Command down 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 the Command key down and use your pointer's scroll control to resize the Viewer.
- Hold the middle mouse button down, 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 the Option key down 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 the Option key down 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.

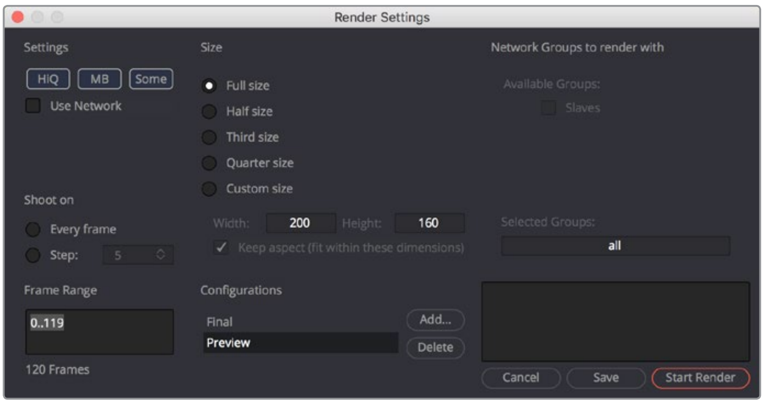

The Flipbook Preview Render dialog

3. When you've chosen the settings you want to use, click Start Render.

   The current frame range of the Fusion page 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, 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 the right mouse button down and drag left or right to scrub through frames.

To play back a flipbook using the keyboard, do one of the following:

- Press the Spacebar to start or stop playback.
- Hold Shift and press the Spacebar to play in reverse.

To scrub through a flipbook frame-by-frame using the keyboard, do one of the following:

- Press the Left or Right Arrow keys to move to the previous or next frame.
- Hold Shift and press the Left or Right Arrow keys to jump back or forward 10 frames.
- Press Command-Left Arrow to jump to the first frame.
- Press Command-Right Arrow to jump to the last frame.

TIP: The mouse and keyboard shortcuts work in 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 of the settings available for rendering flipbook previews to RAM.

Settings

The Settings section of the Preview Render dialog include 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 settings 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 that is 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 that are 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.

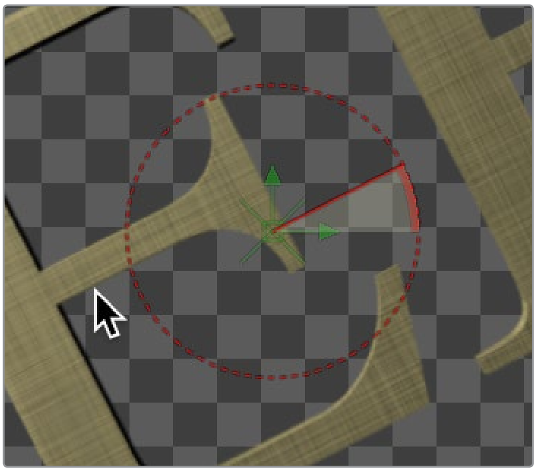

The Angle preview control

The controls that are 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.

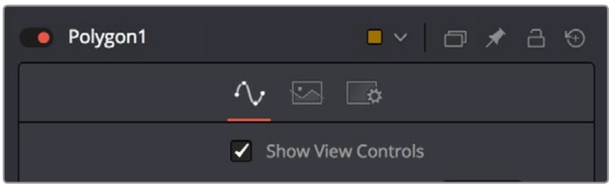

You can disable some nodes, like the Polygon node, on a per node basis

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.

- The Arrow keys can be used to adjust the position of an onscreen control along the vertical or horizontal axis by small steps.
- Holding the Command key down while using the 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

Viewer toolbars run 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.

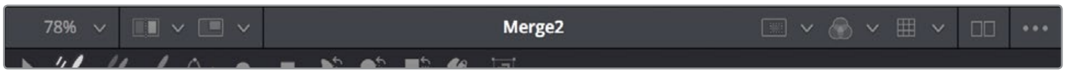

The Viewer toolbar

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.

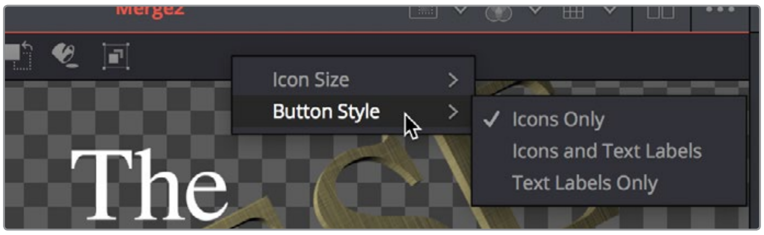

The contextual menu shown for the Node toolbar

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 two different images 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.

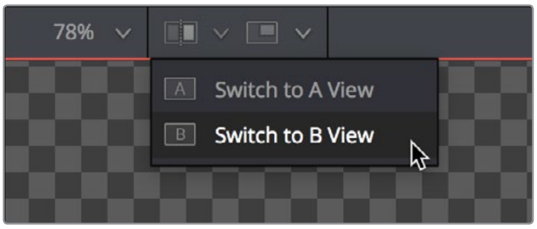

The Buffer menu lets you switch between buffers

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 the wipe off, click the Switch to Split Wipe View button again (or press /).

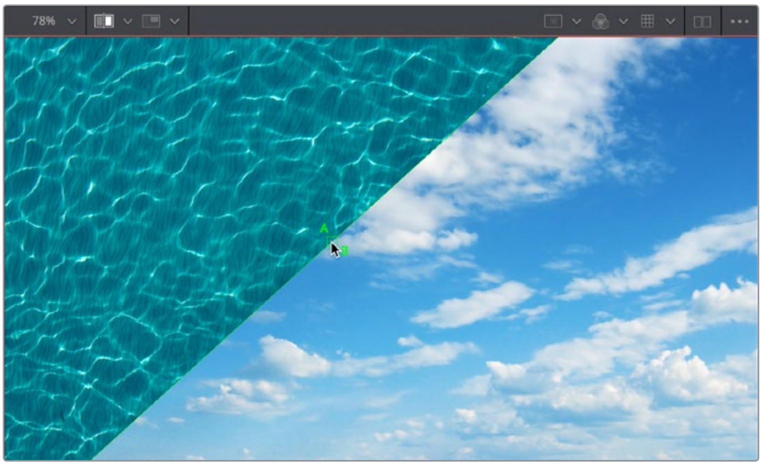

The wipe divider can be adjusted for comparing different areas of the A and B images

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 Command-Option down 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.

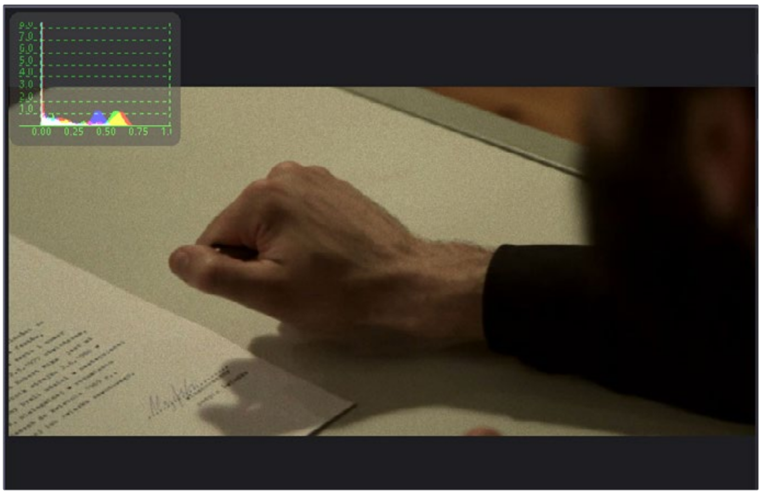

The Subview menu with the Histogram subview displayed

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 that are 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.

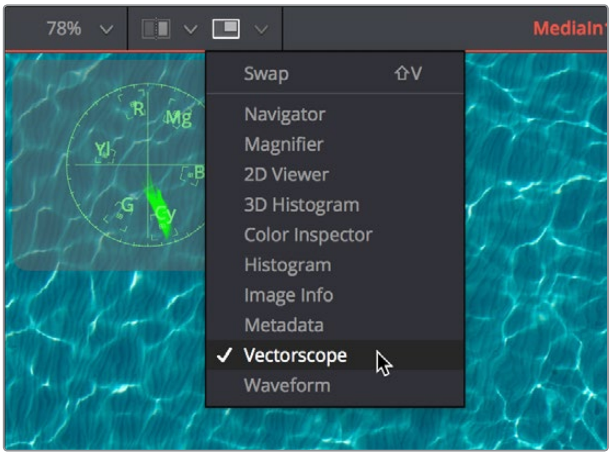

The Subview button in the Viewer toolbar

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.

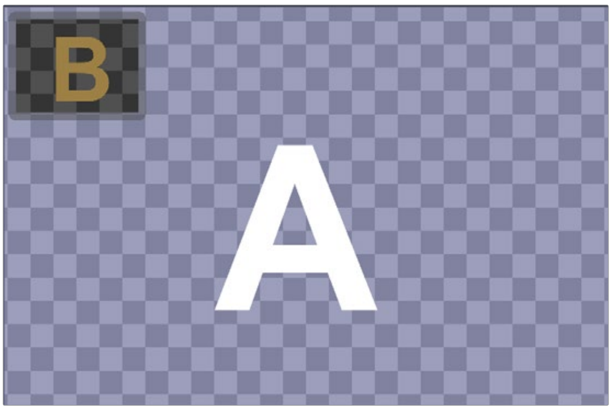

The Subview used as another viewer

3D Image Viewer

The 3D image viewer is available when viewing a node from the 3D category.

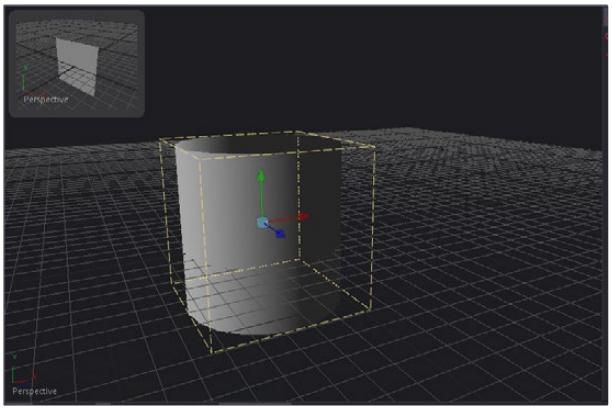

A 3D image viewer as a subview

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.

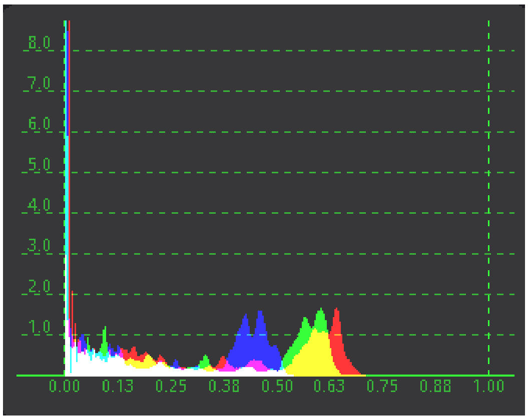

The Histogram viewer type for evaluating the contrast and color cast in an image

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.

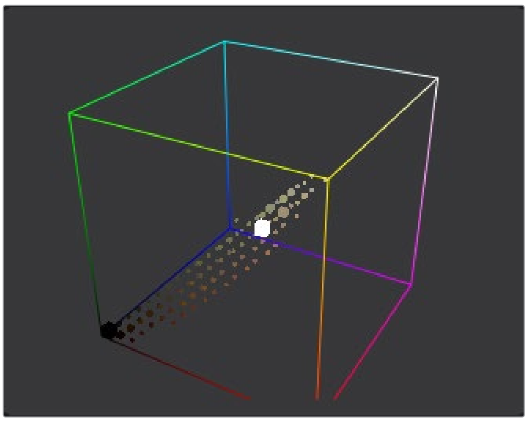

The 3D Histogram viewer type for evaluating out of range colors

To rotate within a 3D Histogram, do one of the following:

- Hold the Option key down and drag left or right using the middle mouse button.
- Hold the middle and right mouse buttons down 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.

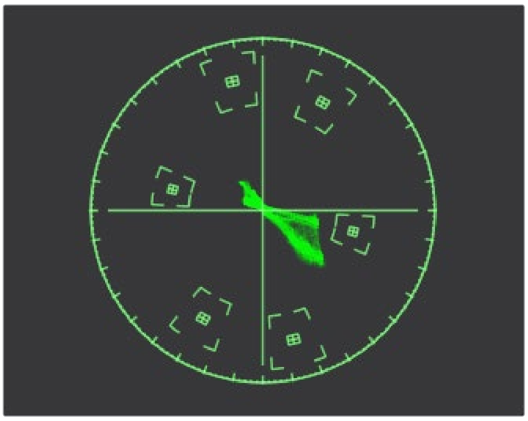

The Vectorscope viewer type for evaluating chrominance in an image

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.

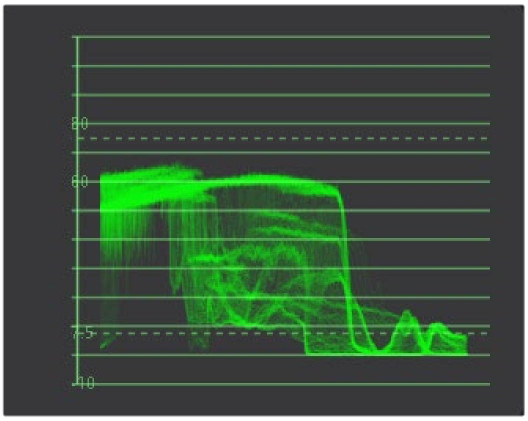

The Waveform viewer type for evaluating luminance in an image

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.

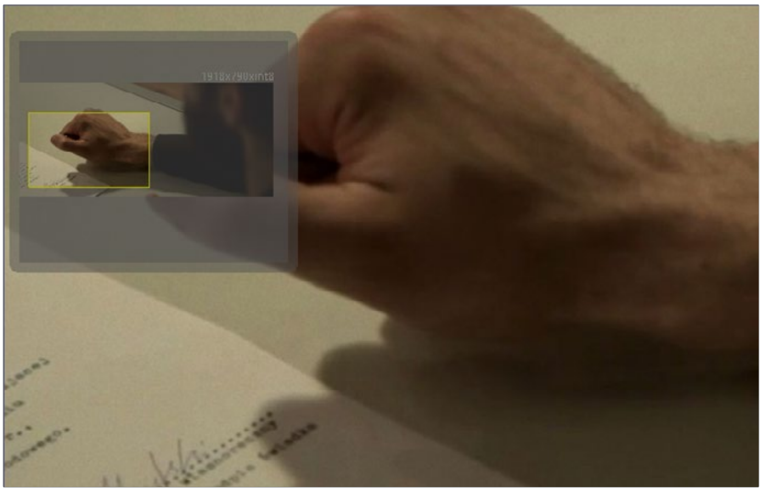

The Navigator subview for panning the image while zoomed in

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.

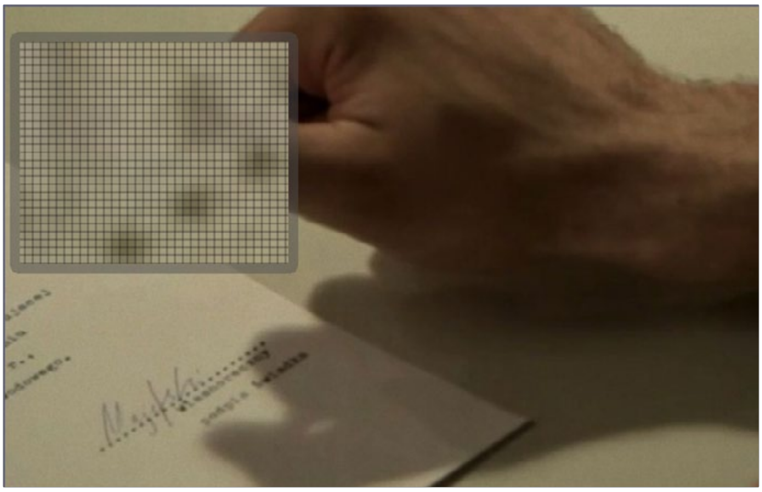

The Magnifier subview to view a zoomed in version of the image

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.

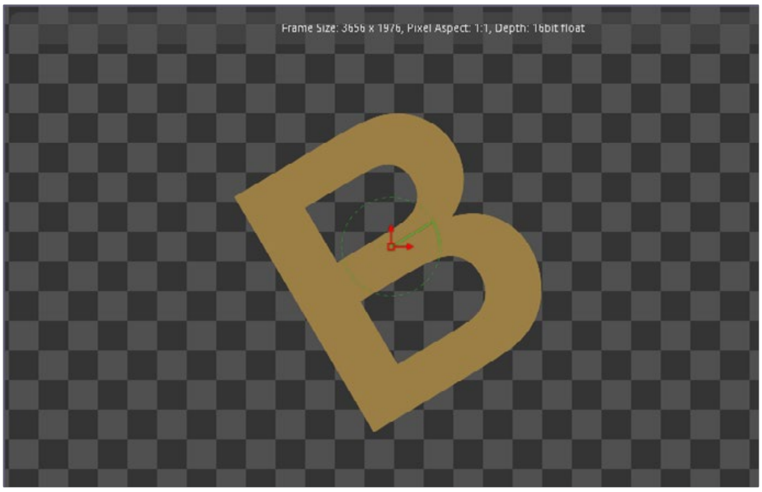

The Image Info subview for viewing size, pixel aspect, and color depth information

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.

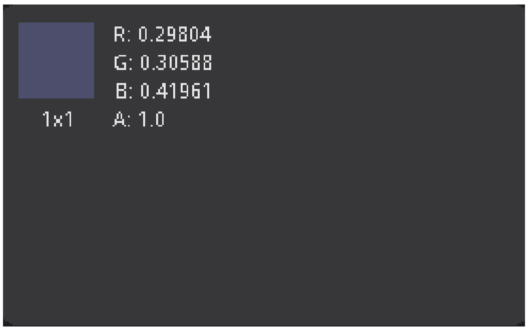

The Color Inspector subview for evaluating colors in an image

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.

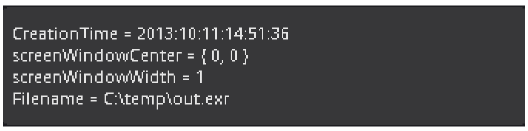

The Metadata subview for viewing embedded metadata

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.

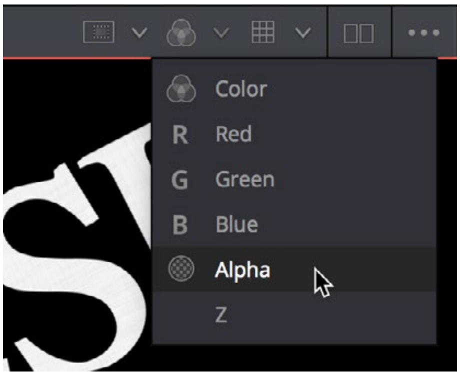

The Channel toolbar button for switching between color, alpha, and depth channels

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 of 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 than traditional 2D compositing. When a node from the 3D category or some particle systems is selected, a 3D Viewer is used to display the scene. The 3D Viewer shows a representation of a composite in a true GPU-accelerated 3D environment.

For more information on 3D controls, see Chapter 60, “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 forwards. 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 the Option key down and drag left or right using the middle mouse button.
- Hold the middle and right mouse buttons down 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 now 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.

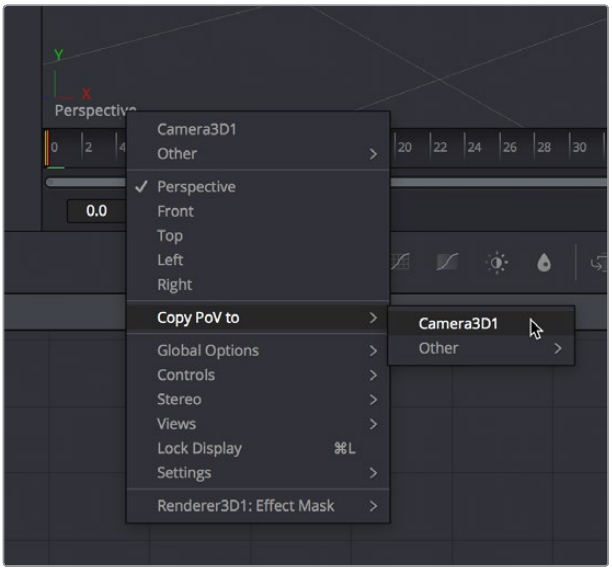

The Viewpoint label and Axis control in a 3D Viewer

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.

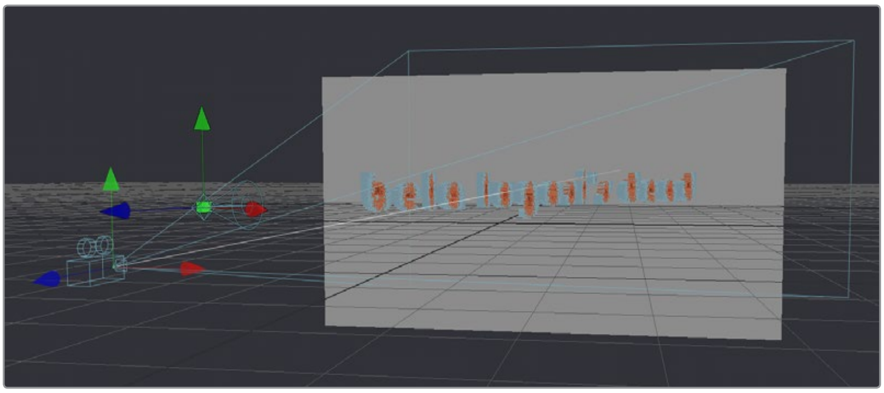

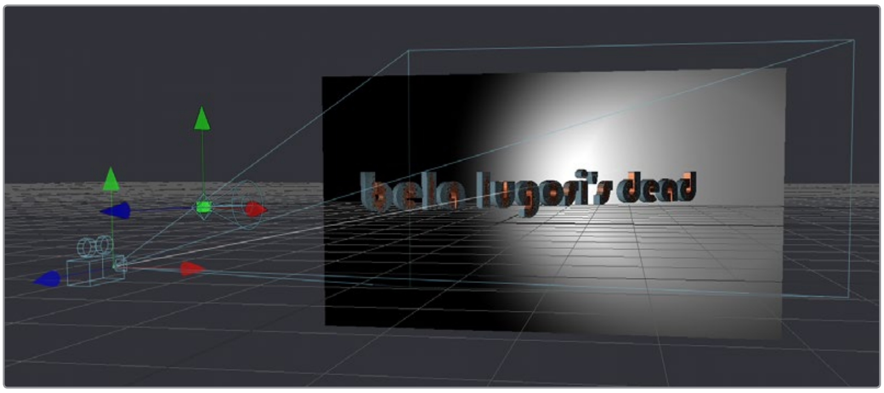

(Top) A 3D scene using default lights, (Bottom) The same scene with lighting turned on

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 have to 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.

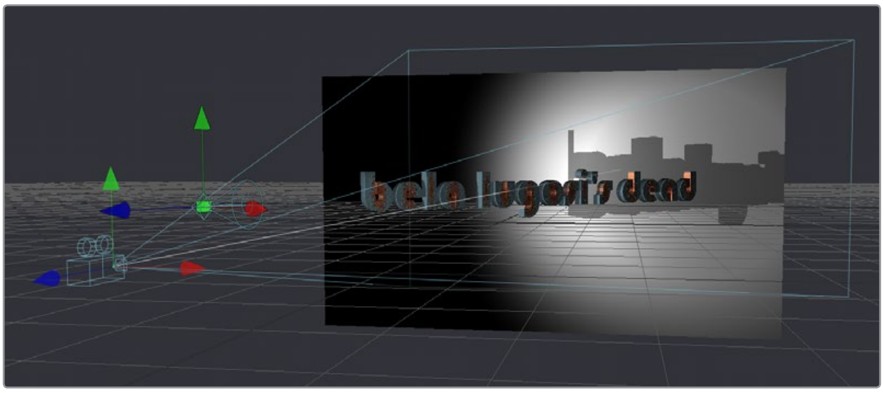

A 3D scene with shadows enabled along with the lights

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.

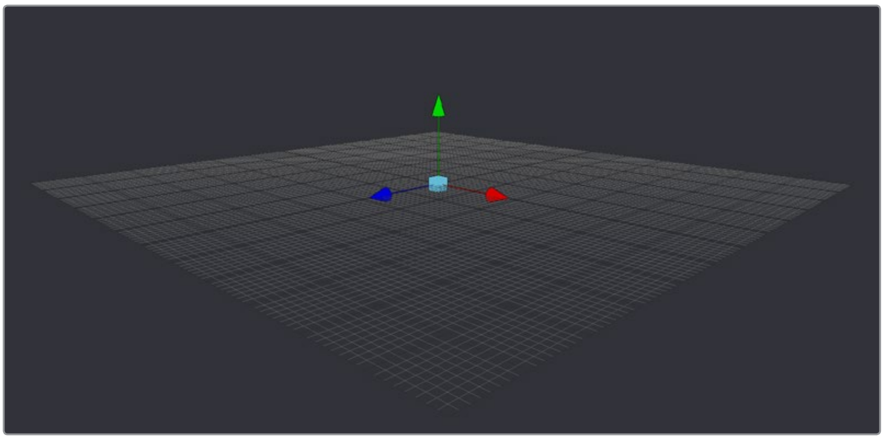

The default grid of the 3D Viewer grid with its origin at x=0, y=0 and z=0

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.

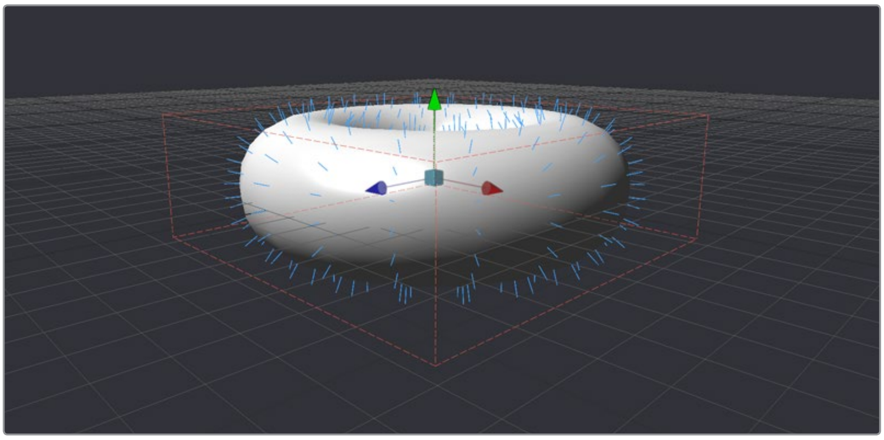

The Normals viewed in a 3D scene

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.

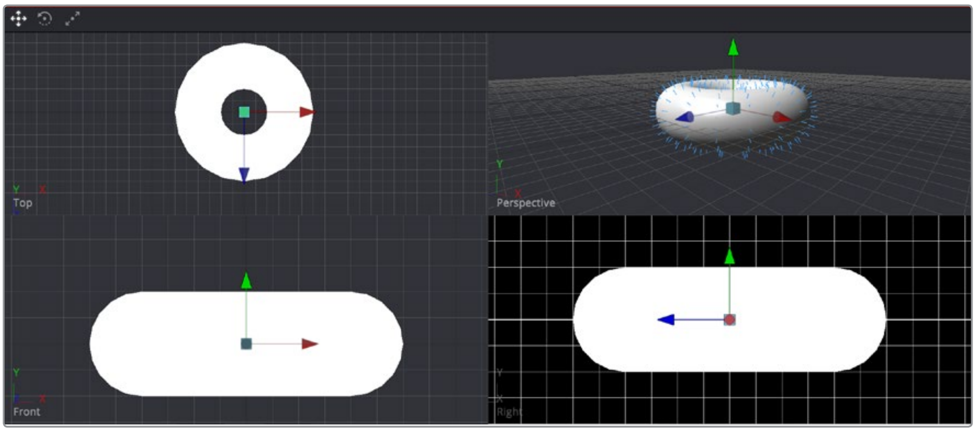

A Quad view of a 3D scene

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 and 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 that are used to help you compose elements within a boundary or along the center vertical and horizontal axis. 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, 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, 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, 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.

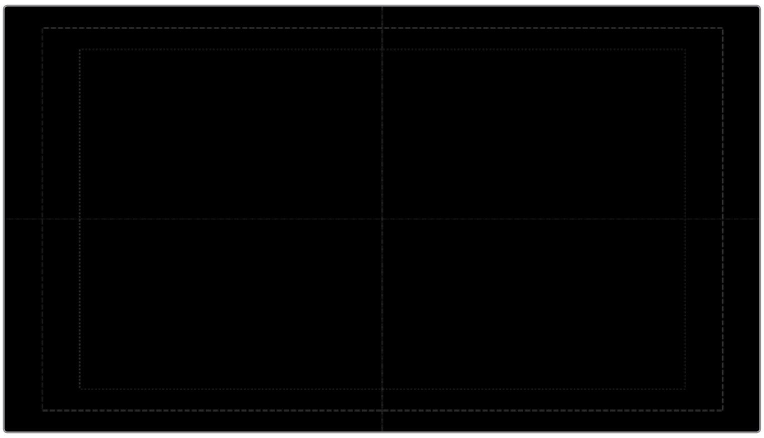

The Guides submenu in the Viewer's contextual menu

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.

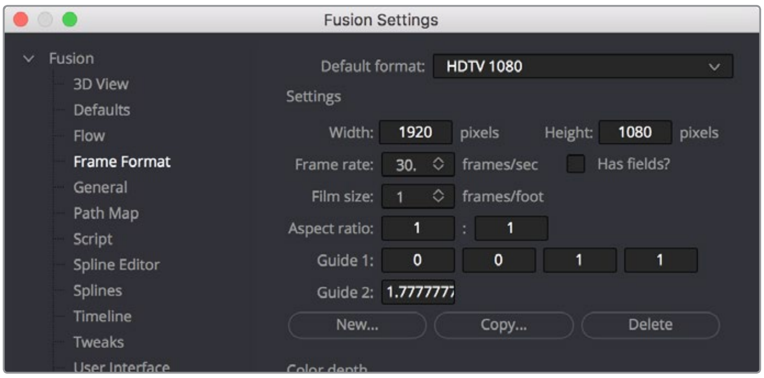

The Frame Format Guides settings

Domain of Definition and Region of Interest

As a compositing environment, the Fusion page uses the standard compositing conventions of Region of Interest (RoI) and Domain of Definition (DoD) to dramatically improve rendering.

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 the Fusion page 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, the Fusion page 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.

ShowDoDResult

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.

Clips from the Edit page timeline or Media Pool will typically also have the DoD default to the full image width of the source media. The exception is media stored in the .raw or OpenEXR format. These formats are capable of storing the data window of the image, and the Fusion page will apply this as the DoD when loading such an image and will write out the DoD if output directly.

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 (RoI)

The Region of Interest, frequently abbreviated to RoI, is a rectangular region similar to the Domain of Definition. However, unlike the DoD, which tells the node what pixels are actually present in the image, the RoI tells the node which pixels actually need to be rendered. When a node renders, it intersects the current RoI with the current DoD to determine what pixels should be affected.

Enabling RoI Controls

You can turn on the RoI controls to restrict rendering to a small region of the image to significantly improve performance when you're only working on a small part of a high-resolution or complex composition. For example, if you're using paint to clean up some holes in a matte on the floor of a composition with many, many high resolution layers, 3D, and Lighting operations, you can use the RoI controls to isolate the part of the floor you're working on, which makes caching that part of the composition much faster.

To enable the RoI controls, do one of the following:

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

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

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

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

Sometimes, it's faster to simply draw a rectangle where you want the RoI to be.

To quickly draw the RoI at the desired size:

1. Choose Set from the Viewer menu next to the RoI button, or right-click anywhere within the Viewer and choose Region > Set Region.
2. When the pointer turns into an RoI drawing cursor, drag within the Viewer to set a RoI rectangle.

Alternately, 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 the Fusion page to the rest of DaVinci Resolve. While the RoI is active, the Fusion page will only request rendering of the pixels inside the region when it displays an image in that Viewer. Flipbook previews that you create in that Viewer will also respect the current RoI. MediaOut 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, it can also reduce file I/O, since MediaIn nodes only load pixels from within the RoI, if one is specified. This does require that the file format used supports direct pixel access. Cineon, DPX and many uncompressed file formats support this feature, as well as OpenEXR and TIFF in limited cases.

Please note that changes to the viewed image size or color depth will cause the pixels outside the RoI to be reset to the image's canvas color. This also happens when switching in and out of Proxy mode, as well as during Proxy mode switching with Auto Proxy enabled. When the image size is maintained, so are the last rendered pixel values outside the RoI. This can be useful for comparing changes made within the RoI with a previous node state.

TIP: Image Overlay 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 2D viewer.

Image LUTs

Image LUTs can be applied to each viewer. In fact, you can even apply separate 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 of 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 grading in the Color page at the end of the node tree.

However, in more elaborate production pipelines, you may have the need to apply multiple LUTs consecutively.

Resolve Color Management

At the time of this writing, the Fusion page does not automatically interact in any way with Resolve Color Management (RCM). Images coming into the Fusion page via MediaIn nodes are in the Timeline gamma and color space. For some simple operations, this may be fine, but it's not always ideal.

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 most 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

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 DaVinci Resolve application support directory. On macOS, the /Users/UserName/Library/Application Support/Blackmagic Design/DaVinci Resolve/Fusion/LUTs/ directory lets that macro appear in the Viewer LUT menu.

For this to work, the macro must have one image input and one image output. Any controls exposed on the macro will be available when the Edit option is selected for the LUT. For more information about creating Macros, see Chapter 51, "Node Groups, Macros, and Fusion Templates."

Preset LUTs

All LUTs that are available to DaVinci Resolve in the Color page are also available to the Fusion page, which includes custom LUTs you've installed, as well as preset LUTs that come installed with DaVinci Resolve, such as the highly useful VFX IO category that includes a wide variety of miscellaneous to Linear and Linear to miscellaneous transforms. All of these LUTs appear, by category, in the Viewer LUT menu.

Fuse LUTs

Fuses are scriptable 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 information about fuses, see Chapter 51, "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 DaVinci Resolve, including custom LUTs you've installed yourself.

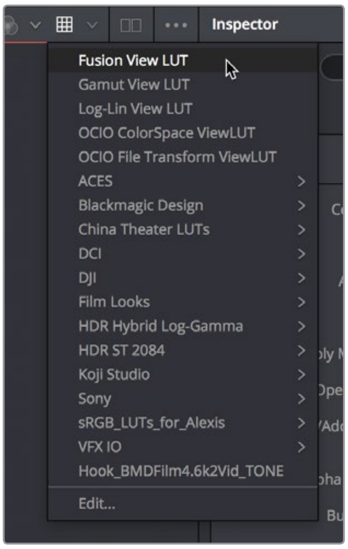

The Viewer LUT button and menu

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:

- Right-click anywhere within the Viewer, and choose Global Options > Buffer LUT > Enable.
- 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 color correction is 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, 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 only applied when viewing the alpha channel of an image, or when viewing masks.

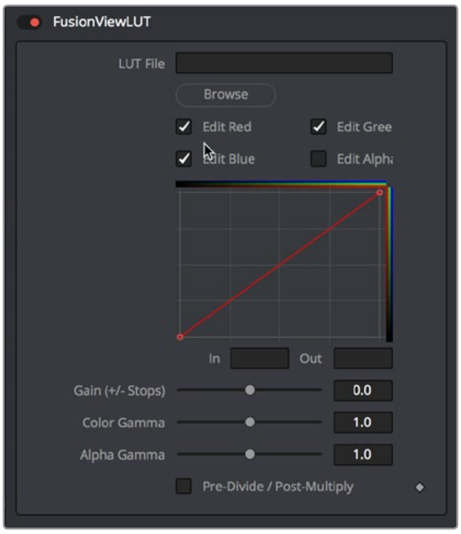

The LUT Editor for the default Fusion View LUT.

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.

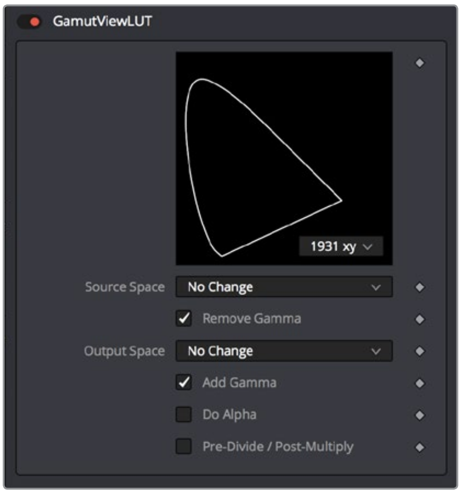

The Gamut View LUT Editor

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.

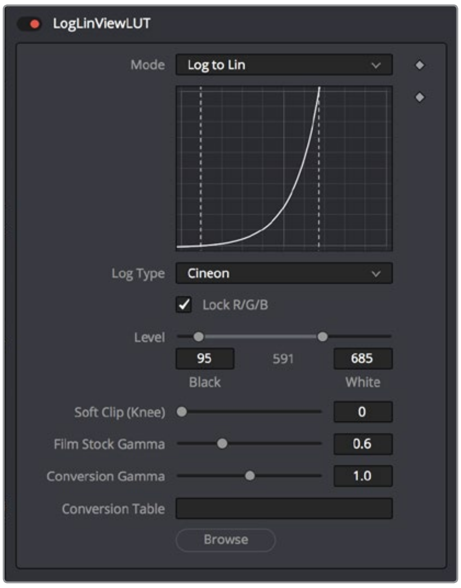

The Log-Lin LUT Editor

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 Cineon file in Log color space you may often apply 3 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 onscreen controls are drawn on top.

Applying Multiple LUTs

The Viewer contextual menu can be used to apply multiple image LUTs into a processing chain.

To apply an additional LUT, do the following:

1. Right-click anywhere within the Viewer.
2. From the Viewer's contextual menu, choose LUT -> Add New.
3. From the Add New submenu choose a LUT to add.

To remove a LUT other than the first LUT, do the following:

1. Right-click anywhere within the Viewer.
2. From the Viewer's contextual menu, choose LUT -> Delete.
3. From the Delete submenu choose a LUT to remove.

A complete stacked LUT configuration can be saved to and loaded from a .viewlut file, as described below.

Saving Custom LUTs

There are a variety of ways to create and use different Viewer LUTs in Fusion. You can save LUTs when you save Viewer settings, you can import LUTs that have been exported from Fusion or other applications and you can open any one of the various supported LUT file types. In addition, you can use the standard nodes in Fusion to create macros, which can then be saved and used as a LUT.

LUT Settings

The most straight forward 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, 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 popup 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 popup 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 .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.

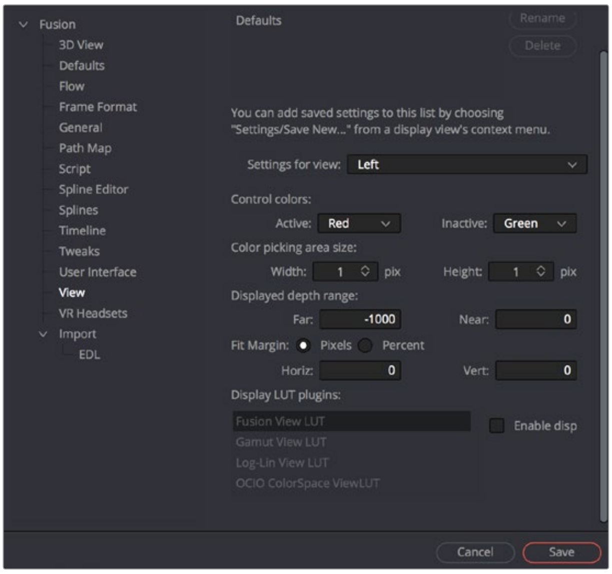

The LUT Default settings found in the View panel of the Fusion Settings window.

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 maybe 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 onscreen controls are not necessary or are getting in the way of evaluating the image, you can temporarily hide them using the Show Controls option. This option is toggled using Command-K.

Checker Underlay

The Checker Underlay shows a checkerboard beneath transparent pixels to make it easier to identify transparent areas. This is the default option for 2D viewers. Disabling this option replaces the checkerboard with black.

Pixel Grid

Enabling this option will show a light black grid that outlines the exact boundaries of pixels in the image when the image is scaled past a certain threshold. The default is Off.

Smooth Resize

The Smooth Resize option uses a smoother bilinear interpolated resizing method when zooming into an image in the Viewer. When smooth resize is disabled, scaling uses the nearest neighbor method and shows noticeable aliasing artifacts but is more useful for seeing the actual pixels of the viewed image when you zoom all the way down to a pixel level since there is no interpolation. This option is enabled by default and can be toggled by clicking on the SmR button in the Viewer toolbar.

Show Square Pixels

Depending on the frame format preferences and the type of footage loaded, many images may have pixels that are rectangular instead of square. Both the NTSC and PAL video standards, as well as some anamorphic film formats, use rectangular pixels. A computer monitor uses perfectly square pixels. To compensate for this, aspect correction is automatically performed when viewing non square pixels. This prevents non square pixel images from appearing squashed or stretched in the Viewer.

You can enable the Show Square Pixels option to override the aspect correction. Show Square Pixels can also be toggled on and off using the 1:1 button in the Viewer toolbar.

Gain/Gamma

Exposes or hides a simple pair of Gain and Gamma sliders that let you adjust the viewed image. Especially useful for “gamma slamming” a composite to see how well it holds up with a variety of gamma settings. Defaults to no change.

360° View

Sets the Fusion page Viewer to properly display spherical imagery in a variety of formats, selectable from this submenu. Disable toggles 360 viewing on or off, while Auto, LatLong, Vert Cross, Horiz Cross, Vert Strip, and Horiz Strip let you properly display different formats of 360° video.

Locking the Viewer (Command-L)

You can lock a viewer to prevent it from updating. The node that’s loaded into that viewer still processes and the new image is queued for display in the Viewer, but until you unlock it, the Viewer does not update. By default the Viewer is unlocked.

Additional Viewer Options

There are additional commands when you right-click anywhere within a viewer and choose from the generically named Options submenu.

Alpha Overlay

When you enable the alpha overlay, the Viewer will show the alpha channel overlaid on top of the color channels. This can be helpful when trying to see where one image stops and another begins in a composite. This option is disabled by default.

Overlay Color

When you turn the alpha overlay on, the default color is to show white for the area the alpha covers. There are times when white does not show clearly enough, depending on the colors in the image. You can change the color by choosing a color from the list of Overlay Color options.

Follow Active

Enabling the Follow Active option will cause the Viewer to always display the currently active node in the Node Editor. This option is disabled by default, so you can view a different node than what you control in the Control Panel.

Show Controls

When onscreen controls are not necessary or are getting in the way of evaluating the image, you can temporarily hide them using the Show Controls option. This option is toggled using Command-K.

Show Full Color Range

When working with floating point images, you will occasionally need to visualize the values that fall outside the normal luminance range. Enabling the Show Full Color Range option using the Toolbar button automatically normalize any image displayed in the Viewer. Normalization causes the brightest pixel in a color channel to be mapped to a value of 1.0 (white) and the darkest pixel to be mapped to a value of 0.0 (black). Midrange values are scaled appropriately to fit within that range. It is also useful when viewing Z-buffer or other auxiliary channels, which often use value ranges far different from those in the color channels.

Show Labels

Lets you toggle the display of the text that sometimes accompanies onscreen controls in the Viewer, without disabling the functions that are showing those overlays, and without hiding the onscreen controls themselves.

Status Bar Information

The status bar at the bottom of the Fusion window provide 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.

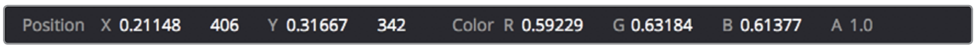

The status bar showing coordinates and color information

Chapter 53

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 controls that are available.

Contents

Overview of the Inspector

While the creation and connection of nodes in the Node Editor determine 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 the different methods for opening node parameters in the Inspector, to edit them in different ways according to the type of controls that are available.

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 Modifiers, which are automated expressions that you can attach to individual parameters to create animated effects, are placed and where Modifier parameters can be edited. Additionally, certain nodes such as the Paint node generate data such as Strokes, which are saved in the Modifiers panel.

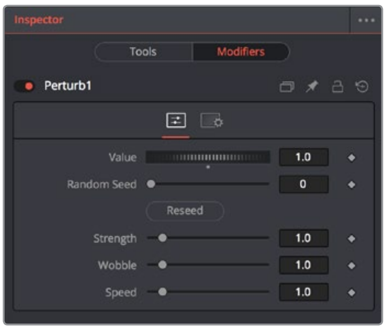

Modifiers displayed in the Modifiers panel

Customizing the Inspector

There are a variety of ways you can customize how the Inspector is presented.

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 the editing of parameters.

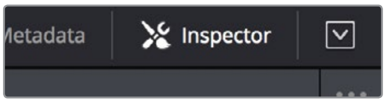

The Maximize button on the left side of the Inspector

In maximized height mode, the Inspector takes up up the entire right side of the UI, letting you see every control that a node has available, or creating enough room to see the parameters of two or three pinned nodes all at once. In half height mode, the top of the Inspector is aligned with the tops of the Viewers, expanding the horizontal space that's available for the Node Editor.

Inspector Display Preferences

By default, you only see selected nodes in the Inspector, and only the Active node is expanded to show its controls. You can change this behavior by choosing Fusion > Fusion Settings and opening the User Interface panel. In the User Interface preferences are checkboxes that manage the display of controls.

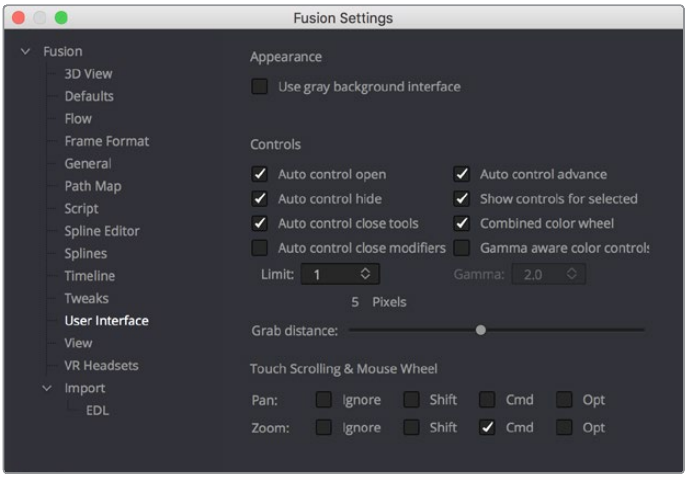

Control preferences in the User Interface category.

- **Auto Control Open:** When enabled (the default), whichever node is Active automatically open 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 that are highlighted white do not.

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 while closing 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, active, or not.

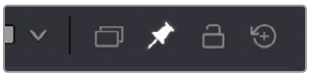

The Pin button of a node's control header in the Inspector

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 compositors who may be working on a composition in situations where you don't want to lock the node.

To Toggle the Inspector controls for a node or or off:

Right-click on the node in the Node Editor, or on the 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 the Fusion Settings.

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. Alternately, 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 codes 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 right-most 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 your project for future use.

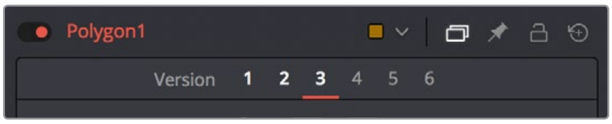

The Version bar, underneath a control header with versions enabled

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.

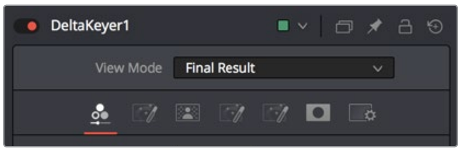

The parameter tabs of the DeltaKeyer node

The Settings Tab

Every node that comes with DaVinci Resolve 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.

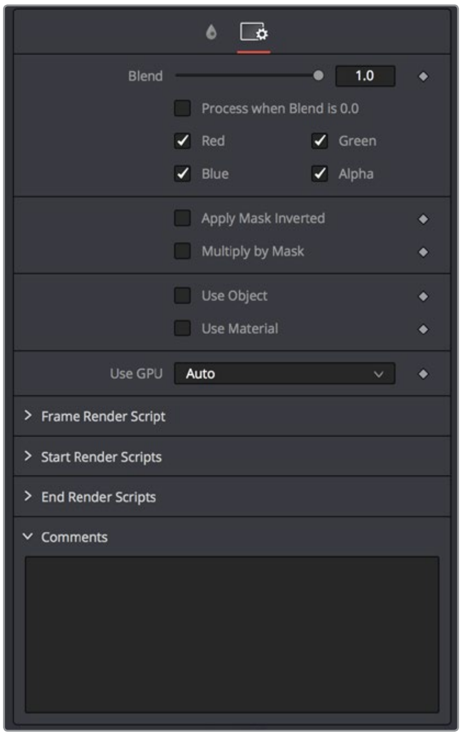

The Settings tab in the Inspector

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 MediaIn and Generator nodes. It is used to blend between the node's unaltered image input and the node's final processed output. When the blend value is 0.0, the outgoing image is identical to the incoming. 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 so only the green channel is affected

For example, if you wanted to use the Transform node to affect only the green channel of an image, you can turn off the Red, Blue, and Alpha checkboxes. As a result, the green channel is processed by this operation, and the red, blue, and alpha channels are copied straight from the node's input to the node's output, skipping that node's processing to remain unaffected.

Transforming only the green color channel of the image with a Transform effect

Skipping Channel Processing

Under the hood, most nodes actually process all channels first, but afterwards 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.

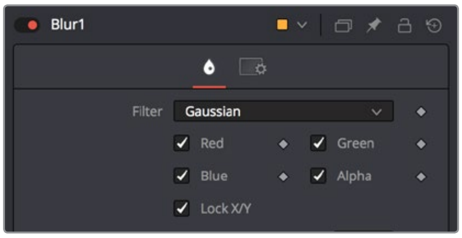

Channel checkboxes on the Controls tab of the Blur node indicates that disabled channels won't be processed at all, to save rendering time

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 check box only operates 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" 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 check box 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 Motion Blur on 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 information, see Chapter 73, "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.

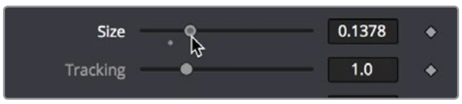

Hold Command or Ctrl while clicking in the gutter to move in smaller increments

Once you click directly on a slider handle you can make changes to its value using the Left and Right Arrow keys. The Command/Ctrl 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.

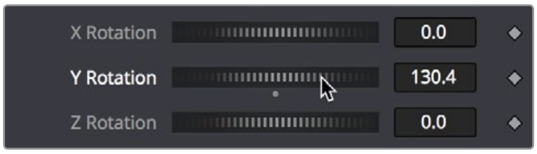

Thumbwheel controls for X,Y and Z rotation with arrows on either end for fine tune adjustments

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/Ctrl 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.

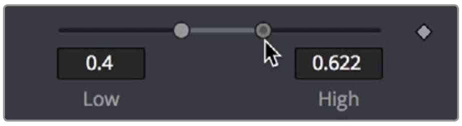

A Matte Threshold Range control

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.

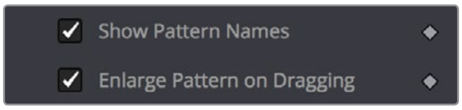

Checkboxes used to select options for tracking

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.

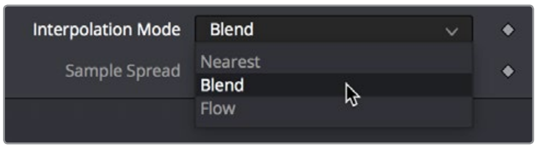

Pop-up menu in the TimeSpeed node

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.

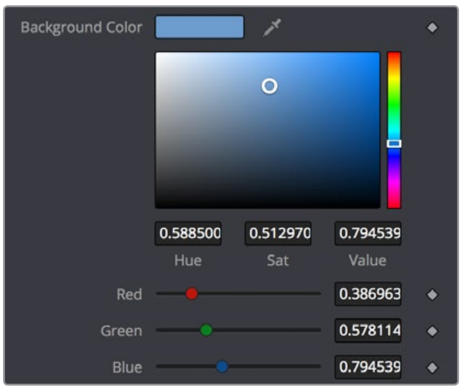

The color swatch with transparency preview

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.

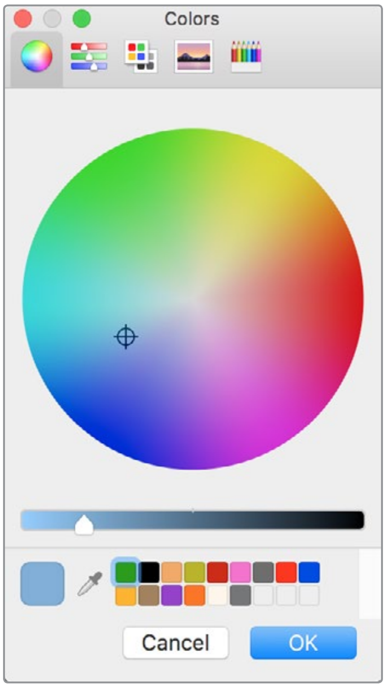

macOS Colors panel

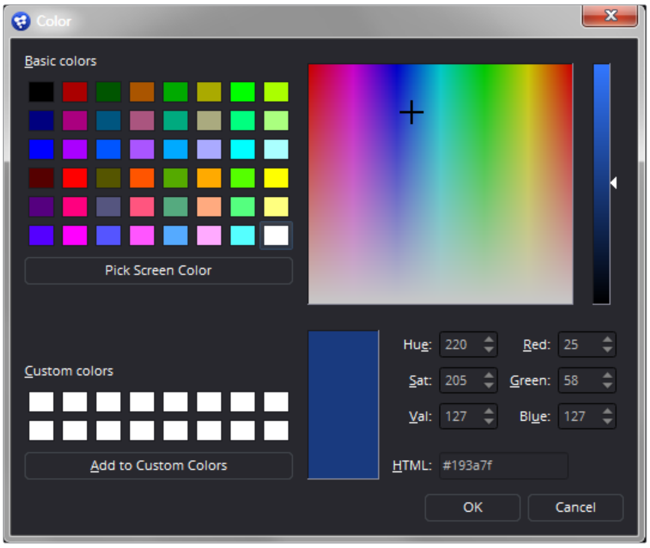

Windows Color dialog

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.

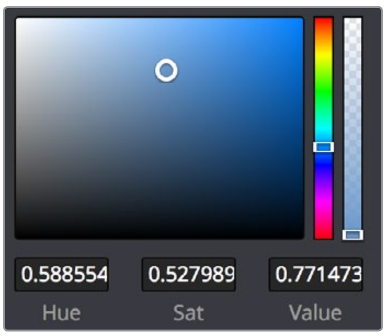

The color chooser in the Background node

Picking Colors from an Image

If you are trying to match the color from an image in the Viewer, you can press the cursor down over the Pick button, 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.

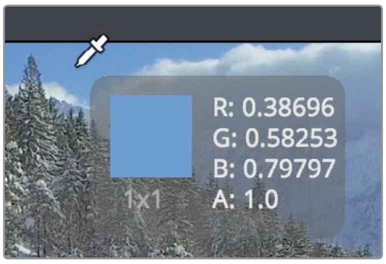

The eye dropper with color swatch

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, then holding Command (macOS) or Ctrl (Windows) and dragging out a rectangle for the sample side 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.

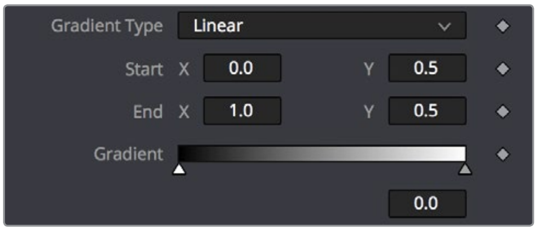

The default Gradient controls

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.

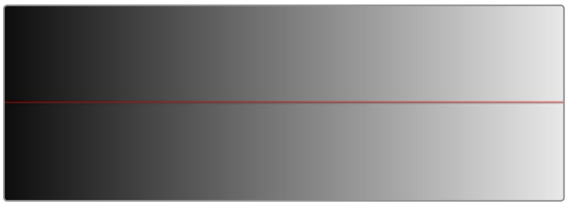

Linear gradient

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

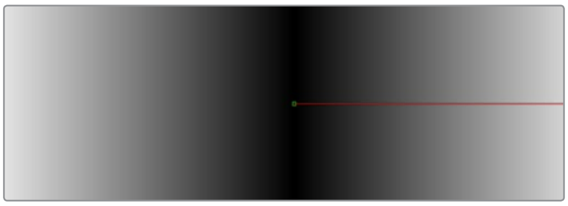

Reflect gradient

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

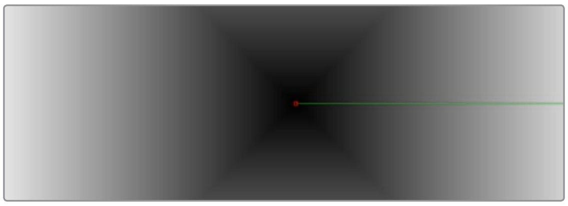

Square gradient

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

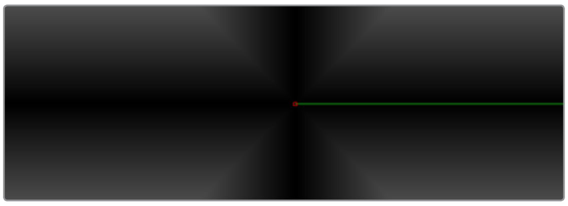

Cross gradient

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

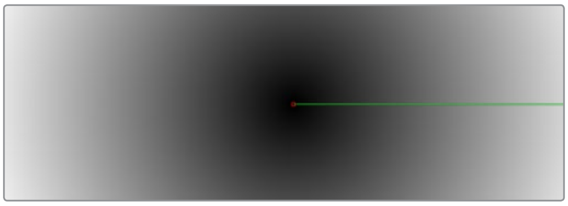

Radial gradient

Angle draws the gradient in a counter-clockwise sweep when the starting point is at the center of the image.

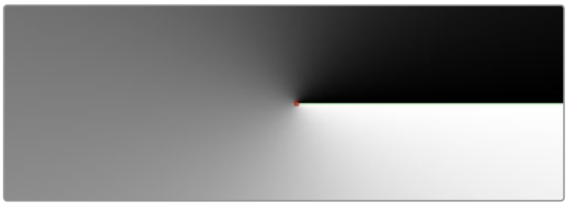

Angle gradient

Start and End Position

The Start and End Position controls have a set of X and Y edit boxes which 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, 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 (macOS) or Ctrl (Windows) 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 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

The Fusion page has keyframing controls that are similar to those found in other pages of DaVinci Resolve. 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; the list is endless. For convenience, a set of keyframing controls are available within the Inspector next to each keyframable parameter. These controls are:

- A gray Keyframe button to the right each keyframable parameter. Clicking this gray button creates a keyframe at the current position of the playhead, and turns the button orange.
- Whenever the playhead is sitting right on top 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 jump the playhead to those keyframes.

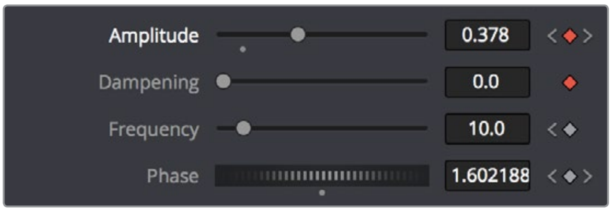

KeyframeControlsFusionPageInspector

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 information about keyframing in the Fusion page, see Chapter 54, "Keyframing in the Fusion Page."

Removing Animation From a Parameter

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).

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 time saver 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.

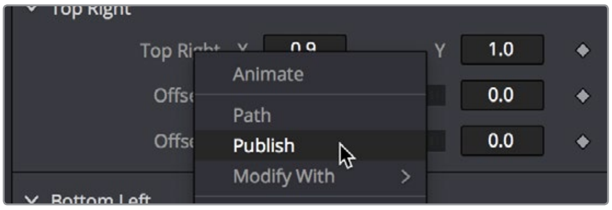

The Publish contextual menu

Connecting Parameters by Pickwhipping

You can also use simple expressions to link two parameters together. By using simple expressions and pickwhipping, 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 Pickwhip 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.

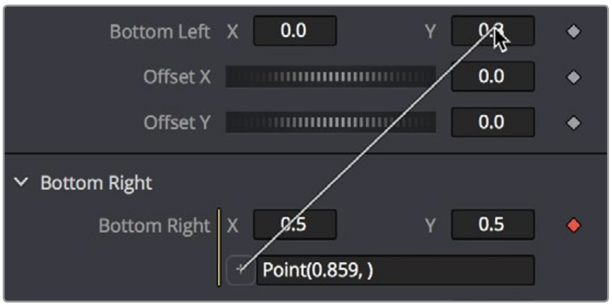

Pickwhipping one parameter to another

TIP: Disabling the Auto Control Close node's General preference, then selecting two nodes in the Node Editor will allow you to pickwhip two parameters from different nodes.

The Expression field can further be used to add mathematical formulas to the value received from the target parameter.

For more information on Pickwhipping and Expressions, see Chapter 56 "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 50, "Working in the Node Editor," for more information on these options.

Parameter Contextual Menus

The contextual menu for individual parameters is accessed by right-clicking over a 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 Add Controls dialog. Right-click the name of a node in the Inspector (in the header bar) and choose Edit Controls from the contextual menu. A new window will appear, titled Edit Control.

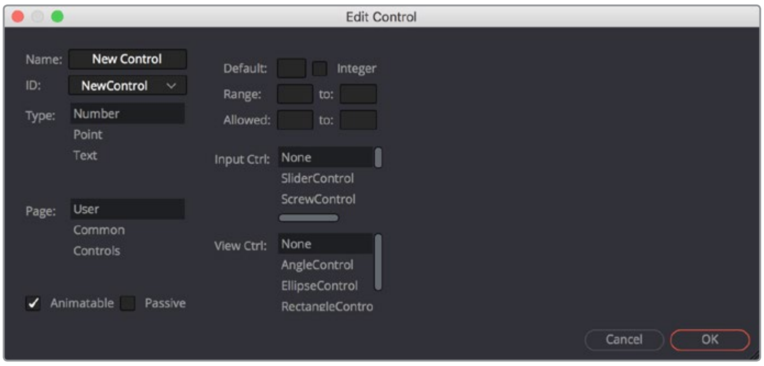

The Edit Control window

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.

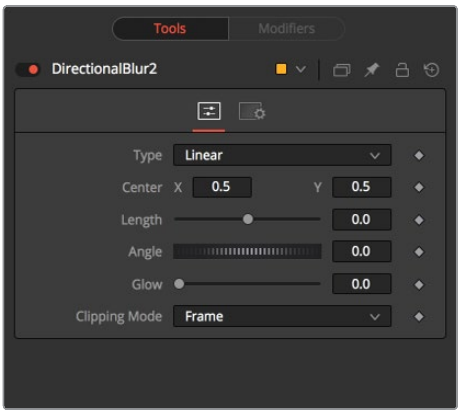

Default DirectionalBlur controls in the Inspector

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
```

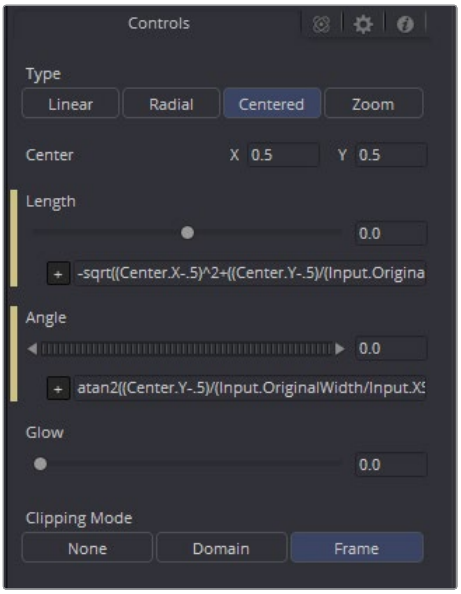

DirectionalBlur controlled by the Center's position

This functions fine, but the controls are confusing. The Center control doesn't work as the center anymore, 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 chose 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 4 entries to 2. 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.

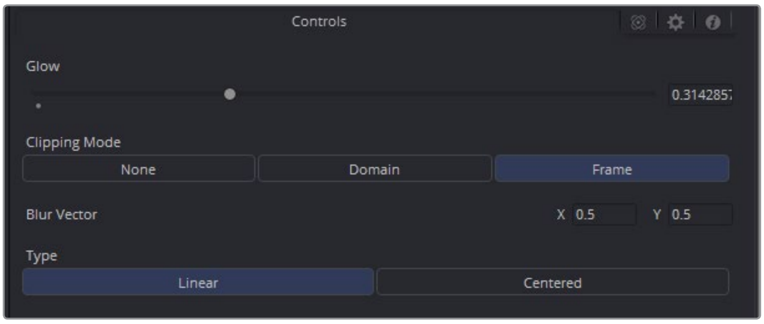

DirectionalBlurs with UserControls applied

Chapter 54

Keyframing in the Fusion Page

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

The Fusion page has keyframing controls that are similar to those found in other pages of DaVinci Resolve. 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; the list is endless.

For convenience, a set of keyframing controls are available within the Inspector next to each keyframable parameter. These controls are:

- A gray Keyframe button to the right each keyframable parameter. Clicking this gray button creates a keyframe at the current position of the playhead, and turns the button orange.
- Whenever the playhead is sitting right on top 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 jump 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 time saver 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 another 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 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

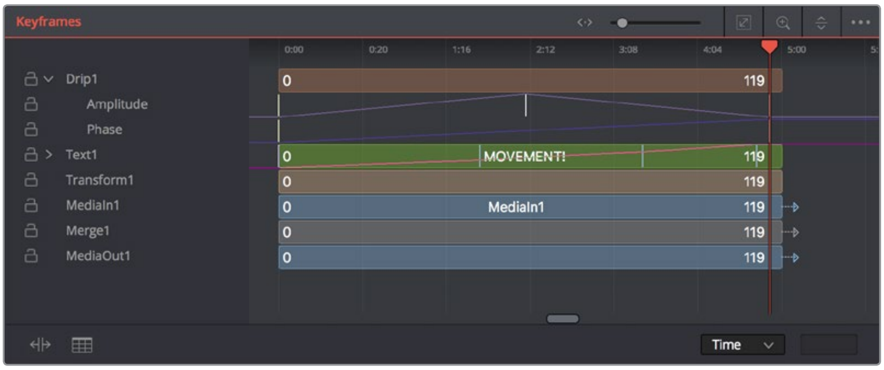

The Keyframe Editor

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).

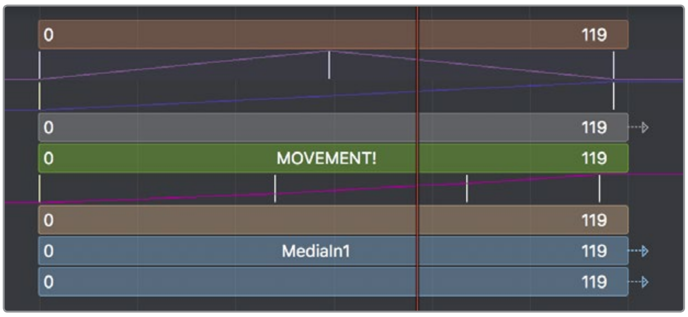

The Timeline tracks

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 DaVinci Resolve, the playhead is a vertical bar that runs through the Timeline view to indicate the position of the current frame or time. The Timeline 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 the Command-Option keys down 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.

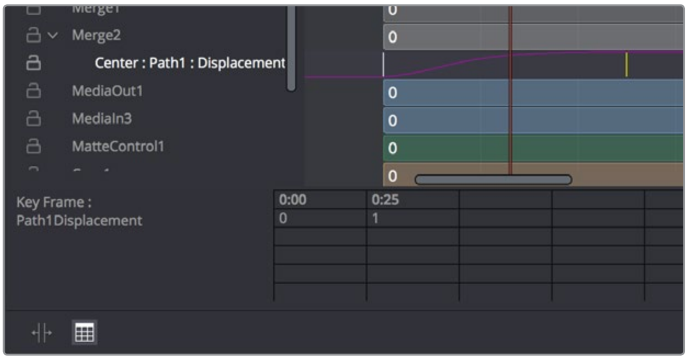

Editing keyframes in the Spreadsheet

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 controls let 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 that may have been added to a composition directly from the Media Pool.

To select a single segment in the Timeline, do one of the following:

- Click the node's name in the header.
- Click the node's segment in the Timeline.

To add another segment to the selection, do one of the following:

- Hold Command and click additional segments to select discontiguous selections.
- Select a segment, 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 de-select 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.

The Move cursor

Resizing Segments

Resizing segments has different effects on MediaIn and Effect nodes:

- Resizing a MediaIn node is similar to resizing clips in the Edit page, in that you're changing the in and out points of the range of media that clip makes available to your composition.
- Resizing 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 MediaIn node that represents a clip that's 100 frames long and is connected to a Defocus node that's animated from frames 80–100. There is little to no point in processing the defocus node between frames 0–79, so trimming the defocus segment to start at frame 80 in the Timeline will effectively prevent it from rendering and consuming either memory or processor time until needed.

To trim a segment in the Timeline, do the following:

Drag on either end of the node's segment in the Timeline.

The cursor changes to a vertical bar with a single arrow when the cursor is in the right location to trim.

The Trim cursor

Holding the First or Last Frame

If you want to hold a Loader's First or Last frame of a clip for a certain number of frames, also called a freeze frame, you can hold Command while you drag beyond the first or last of the segment in the Timeline.

Working with Keyframes in the Timeline

Keyframes can be drawn in one of two ways. When keyframe tracks are closed, they're drawn over the node's segment. Clicking on the disclosure icon to the left of the node's name in the track header expands the display so each keyframed parameter has its own track in the Timeline, enabling precise editing.

Furthermore, each keyframe track, whether open or closed, exposes a miniature curve overlay that provides a visual representation of the rise and fall of keyframed values. This little overlay isn't directly editable.

The Drip1 segment has its keyframe tracks exposed, while the Text1 segment has its keyframe tracks collapsed so they're displayed within the segment.

Drag and Drop Keyframe Editing

Here are pointer-based keyframe editing methods that will get you started.

Methods of selecting keyframes:

- Click a single keyframe to select it.
- Drag a bounding box over a series of keyframes to select them all.
- Command-click to select discontiguous keyframes.
- Shift-click the first and last of a range of keyframes to select a contiguous range.

Methods of adjusting keyframes:

- You can drag keyframes left and right to reposition them in time.
- You can right-click one or more selected keyframes and use contextual menu commands to change keyframe interpolation, copy/paste keyframes, or even create new keyframes.

Keyframe Editing Using the Time Editor

A 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, 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.

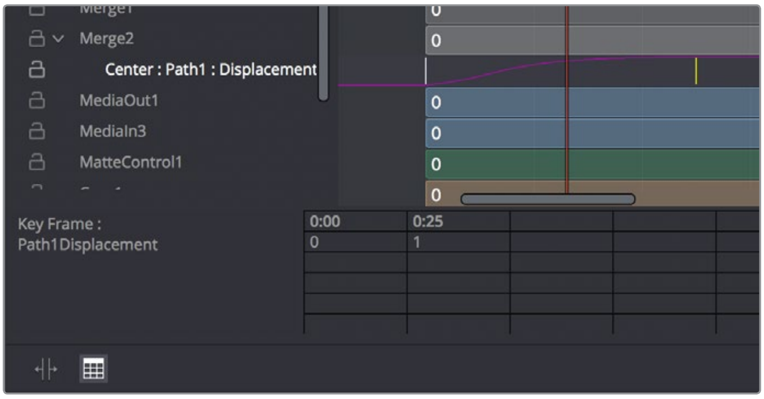

Editing keyframes in the Spreadsheet

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. 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

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.

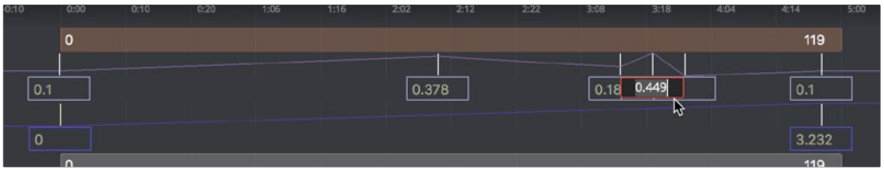

Keyframes with value displayed in the Timeline

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 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 display node layers that are under the playhead.
- If you've created custom filters, they appear here as well, in alphabetical order.

To go back to showing everything:

To disable the currently applied filter, choose Show All from the Keyframe Editor option menu. All layers should reappear.

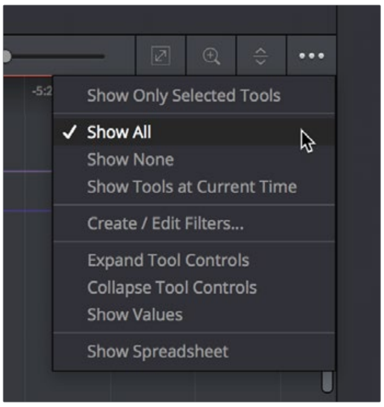

Choosing a Timeline filter

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.

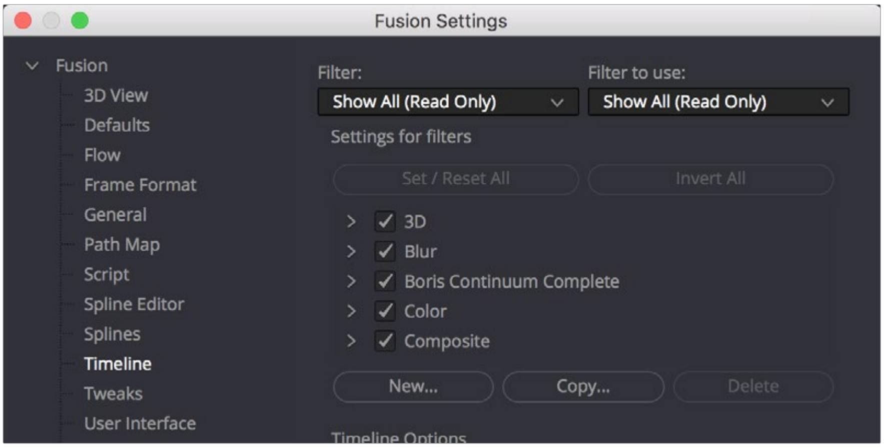

The Global Timeline preferences for enabling filters

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 right away turns all node categories off.
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 the "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 three 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.

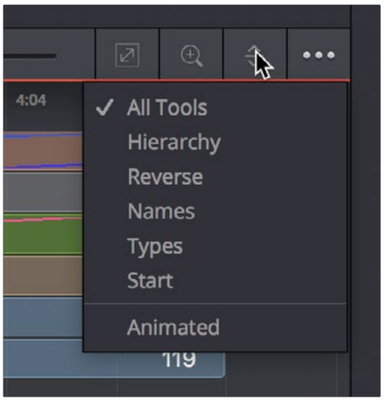

The Timeline Sort Order menu

- **All Tools:** An option that 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 similar to the markers found everywhere else in DaVinci Resolve, and they're 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 Time Ruler, where they appear as a small marker with a line extending vertically through the graph view.

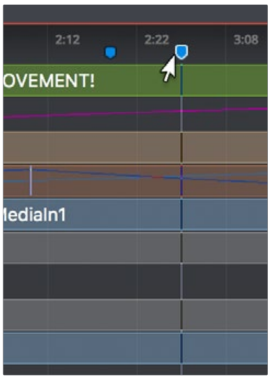

A guide being moved in the Timeline

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.

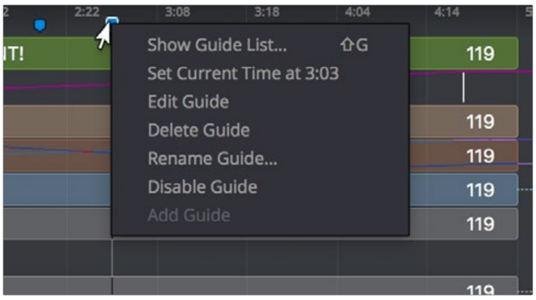

The Guide contextual menu is accessed by right-clicking on the top of the guide

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 and 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.

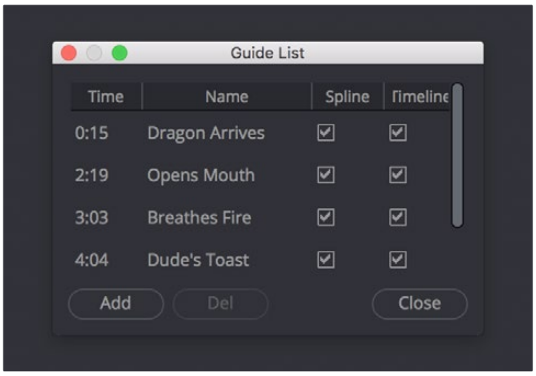

The Guide List dialog allows you to navigate through a composition using guides

There are a pair of checkboxes beside the names of each guide. One is for the Spline Editor and one is for the 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 spline keyframes and the start and end of segments as you drag in the Timeline, you can have them snap to a field, frame or to guides. Autosnap is an option that 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.

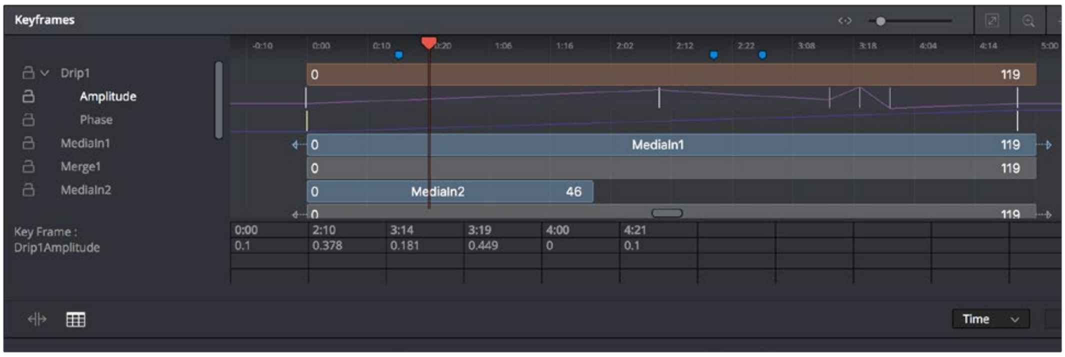

The Spreadsheet Editor showing editable data for six keyframes

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.

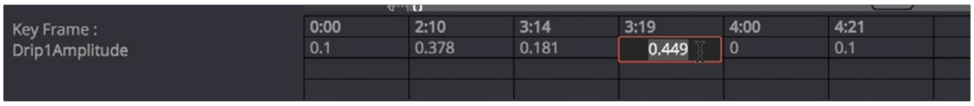

Clicking on the parameter's keyframe value allows you to change it

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 of 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 hight 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.

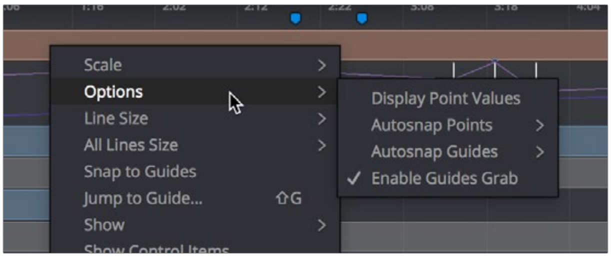

The Options submenu for changing Display Point Values

Here are the two options, compared.

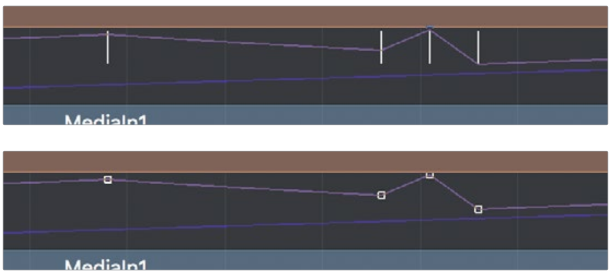

(Top) Keyframes displayed as bars, (Botom) Keyframes displayed as “Point Values”

Chapter 55

Animating with Motion Paths in the Fusion Page

Moving layers and 3D objects can have their animation edited and controlled using motion paths. This chapter discusses how motion paths can be used in the Fusion page.

Contents

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 onscreen with a crosshair or an X.

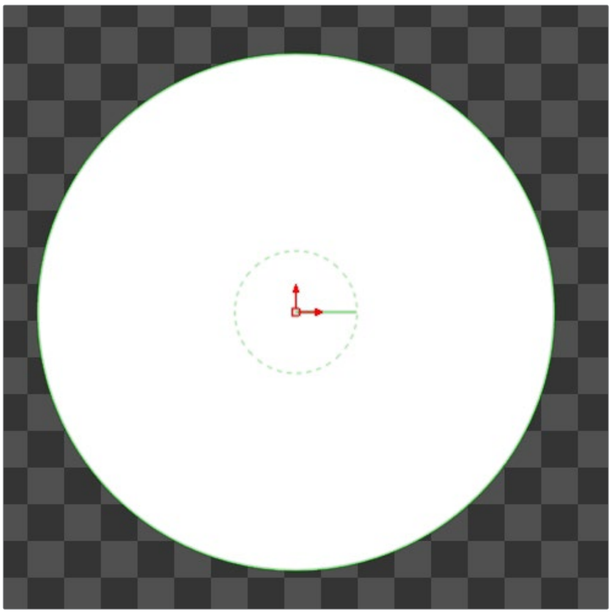

A Center Offset onscreen control for an Ellipse node

It's not possible to add a motion path to a one-dimentional 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 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.
- The Poly path is the default type for motion paths and most documentation in this chapter assumes that this type is used.
- 3D motion paths only pertain 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 on the onscreen control in the Viewer, and choose Path from the contextual menu for that control.

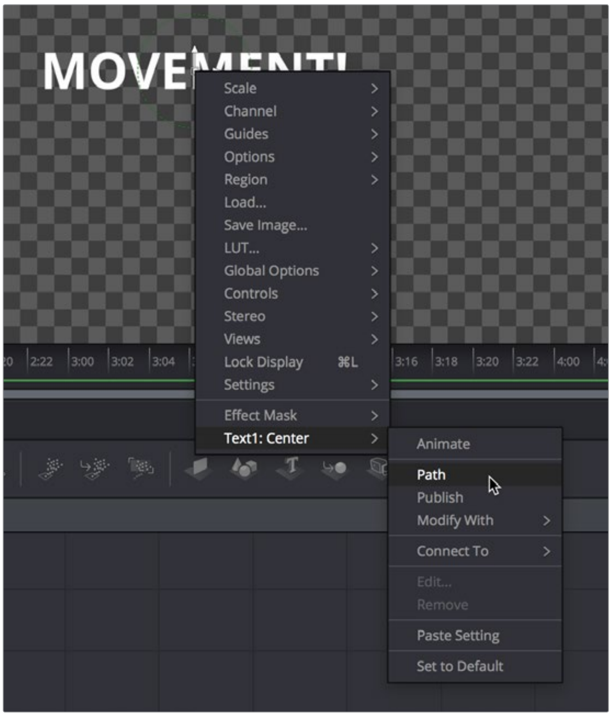

Initiating path keyframing for the center position of a text title

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 onscreen 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.

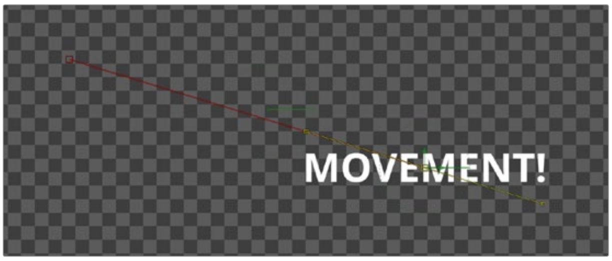

Two keyframes spaced several frames apart display a motion path showing the direction of animation

6. The position of the 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 coordinate control of the object you're animating, and choose "Remove Path1" from the submenu of the "NameOfObject: Center" submenu.

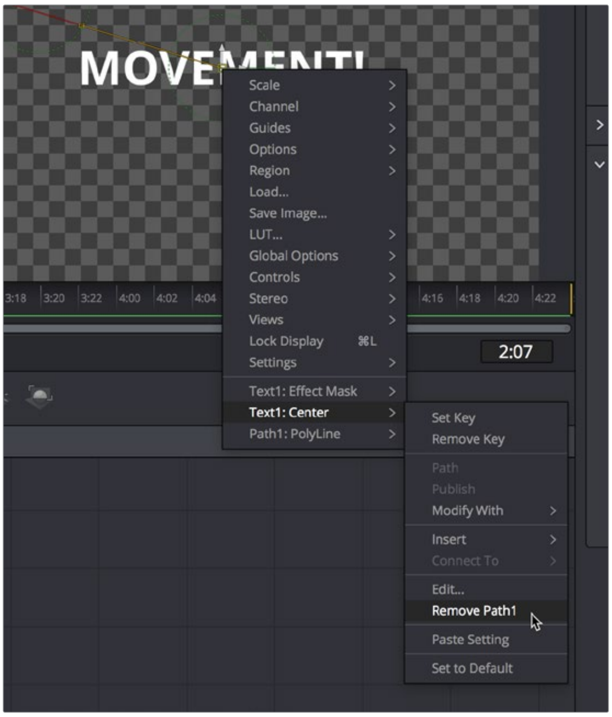

Removing an entire motion path at once

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.

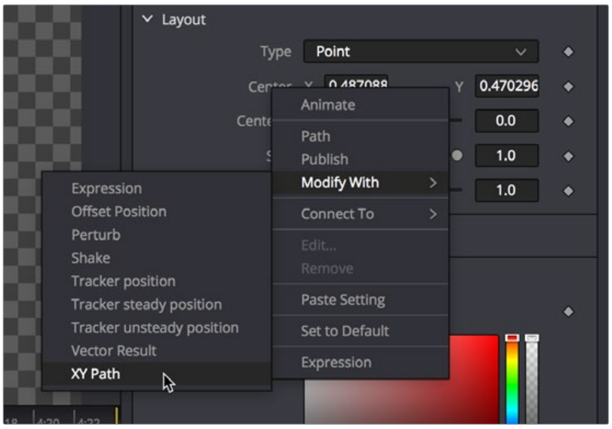

Adding the XY Path modifier to a Center parameter in the Inspector

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 set keyframes, and an editable motion path appears in the Viewer once you've set two or more keyframes to create animation.

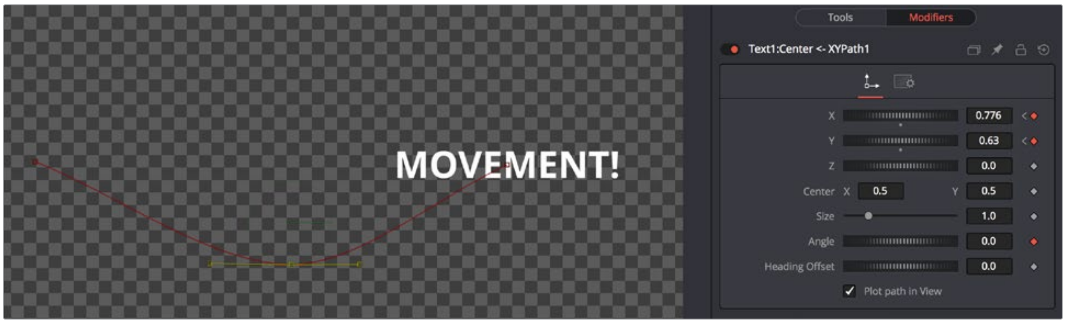

Using an XY Path modifier to animate a piece of text

Displacement Splines and Path Timing

Every Poly path has an associated Displacement Spline in the Spline Editor. The displacement spline represents the position of the animated control along its path, represented as a value between 0.0 and 1.0.

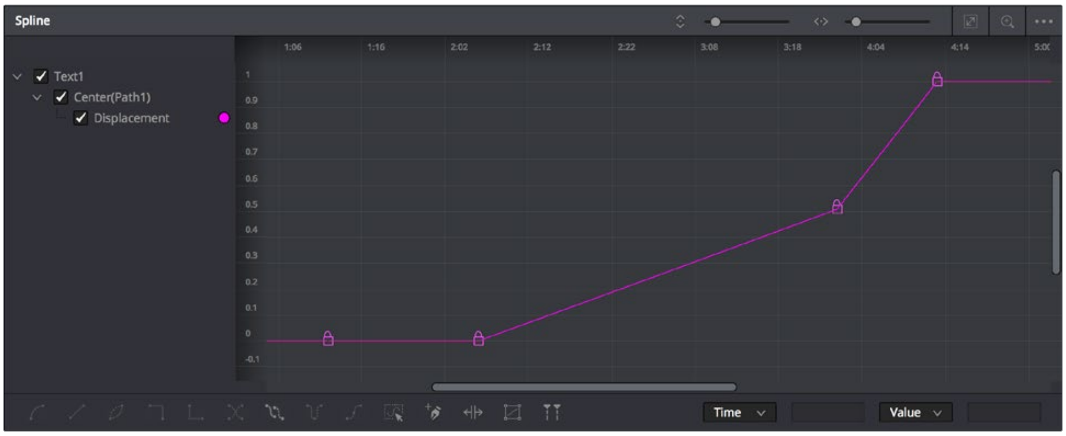

The Displacement curve of a Poly path

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 or not 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.

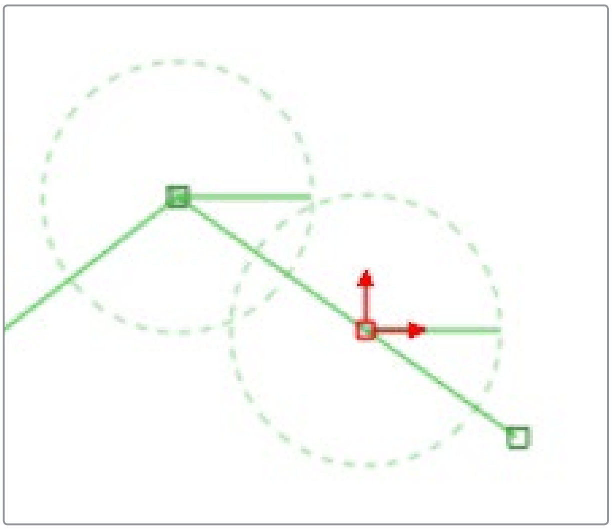

A locked point in the Viewer

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 mode. 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.

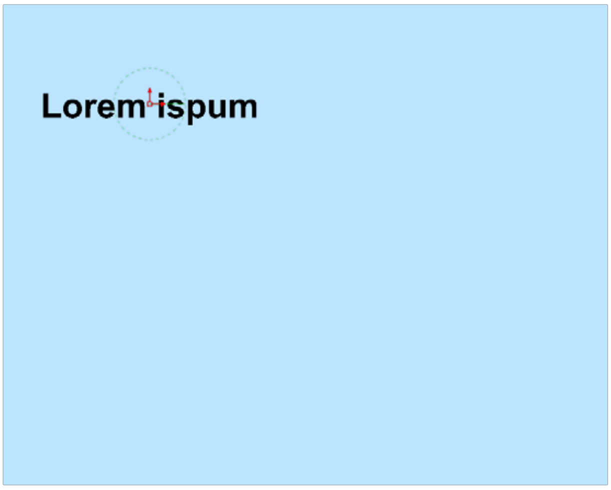

Text placed in the upper left corner of the frame

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.

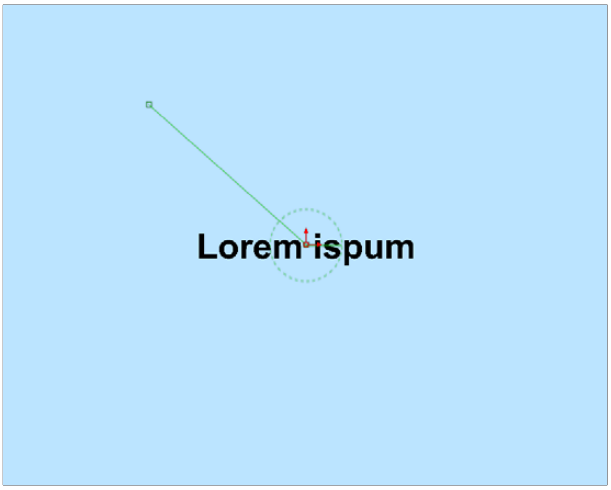

Moving the playhead and repositioning the text adds a locked point

   This sets the second locked point.

7. View the Spline Editor and display Path 1's: Displacement Spline.

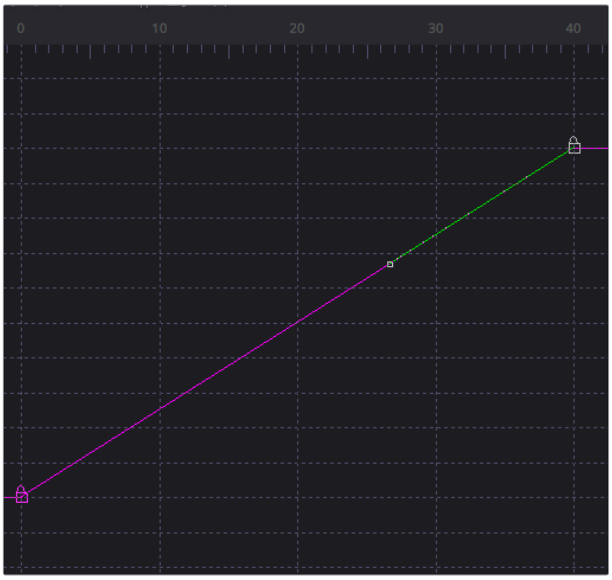

The path's displacement spline

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.

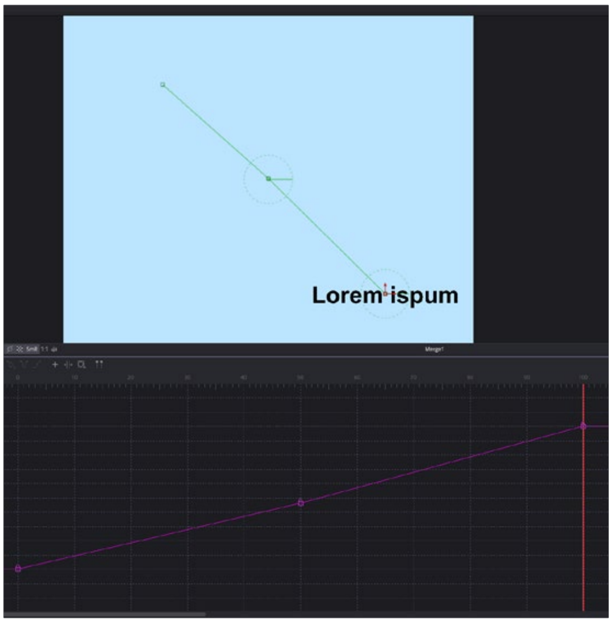

Position the playhead at 100 and drag to reposition the text, creating another locked point

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.

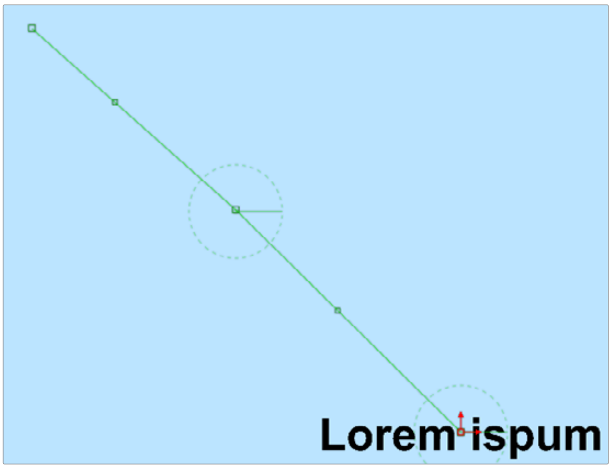

Add two new points to the path

13 Smooth these new points to create a curve in the path.

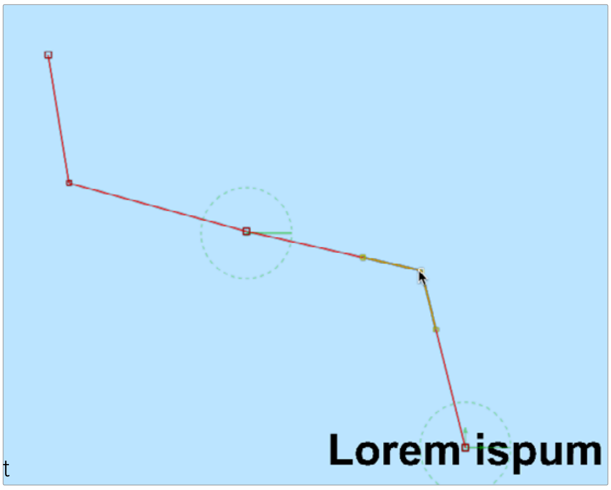

Adjust the path using the newly added points

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 Command (macOS) or Ctrl (Windows) while dragging it to frame 65. The point is copied at frame 65.

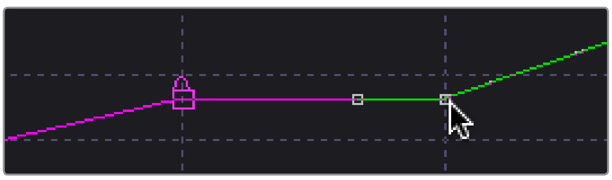

Create a pause in the motion by copying locked points

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 onscreen 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 onscreen control where you want it. Position the playhead again and move the onscreen control to its new position. The difference is that the control points are only there for spatial positioning. There 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 (macOS) or File > Preferences > Globals (Windows) 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.

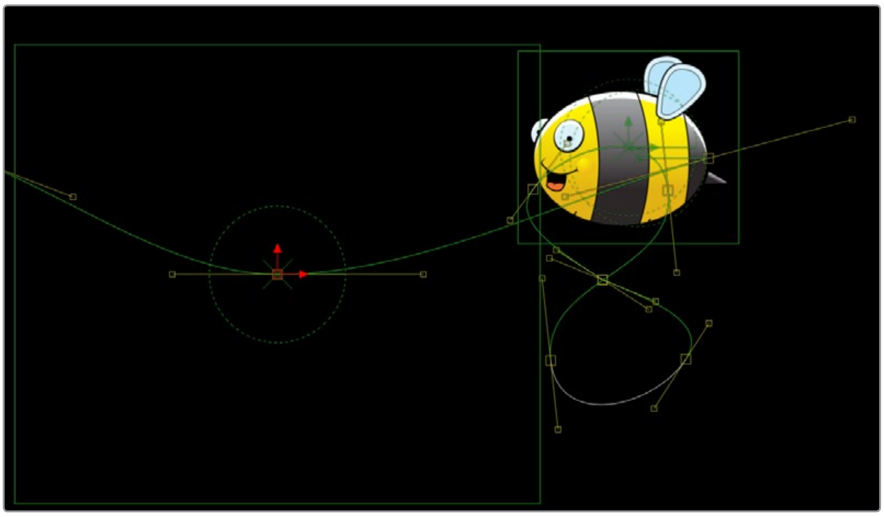

Two motion paths working together

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. Alternately, 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 to display its context menu and select Record from the contextual menu. This displays a submenu of available data that may be recorded.

Use the Record Time option in conjunction with the Draw Append mode to create complex motion paths that will recreate the motion precisely as the path is drawn.

The time used to record the animation may not suit the needs of a project precisely. Adjust the path's displacement spline in the Spline Editor to more correctly match the required motion.

Importing and Exporting Polylines

You can import and export polyline shapes into a common editable ASCII text file or its native format. These methods are used to save a particularly useful or generic mask or path for future use or for use in another application, such as Maya or LightWave. You can also import FXF, SSF or Nuke shape files.

Native Format

To save a polyline shape in Fusion's native ASCII format, 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, and 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, 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 56

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

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 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, points, each have different sets of modifiers that will work with them.

Knowing Which Parameters are Modified

The Fusion page 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 Modifier tab as an indication.

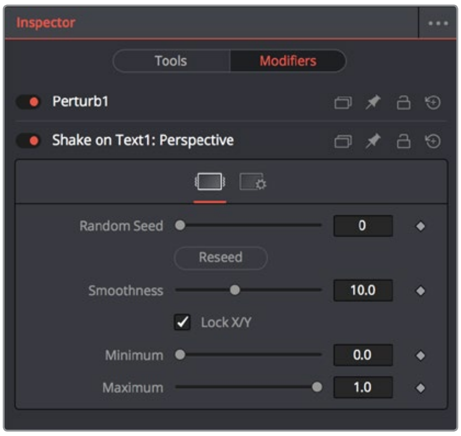

The Modifier tab

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 any other node in the Fusion page. 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. Insert submenu

For more information on all Modifiers available in the Fusion page, see Chapter 95, "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 uni-directional parameter connections, or a combination of both.

SimpleExpressions can be added from the parameter's contextual menu.

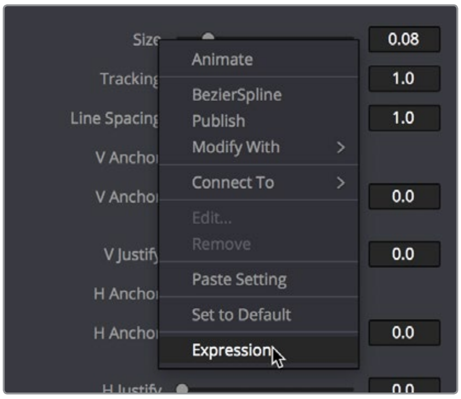

Right-clicking on a parameter to add an Expression from the contextual menu

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 page specific shorthand. Some examples of SimpleExpressions and their syntax:

```
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 5 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.

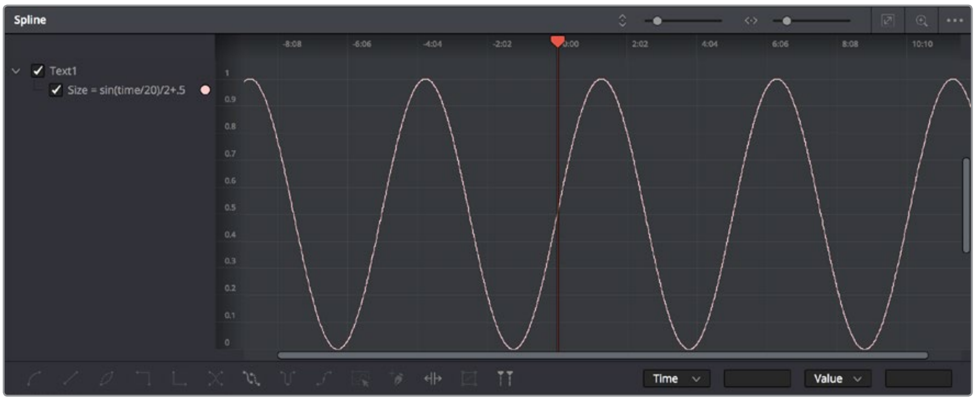

A sine wave in the Spline Editor, generated by the expression used for Text1: Size

```
iif(Merge1.Blend == 0, 0, 1)
```

This returns 0 if the Blend value is 0, and returns 1 if it is not. The iff() function is a shorthand conditional statement, if-then-else.

```
iif(Input.Metadata.ColorSpaceID == “sRGB”, 0, 1)
```

This returns 0 if the image connected to the current node’s Input is tagged with 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 page nodes is the main image input shown in the Node Editor as an orange triangle. Images have members that you can read, like 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”)..”\nfrom 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 file name, 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.

Pickwhipping

To the left of the SimpleExpression sits a button with a + on it. Click dragging the button onto another control, or "pickwhipping," will let you easily get the name of that control. Hovering over a tab while pickwhipping will open that tab.

Pickwhipping 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.

Chapter 57

Using the Tracker Node

This chapter shows the many capabilities of the Tracker node in the Fusion page, 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 on the Fusion Page

Tracking is one of the most powerful automation nodes available to a compositor, and the Fusion page 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

There are four primary applications for tracking data.

Stabilizing

You can use one or more trackers to remove all of 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 which 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 on the Fusion page 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.

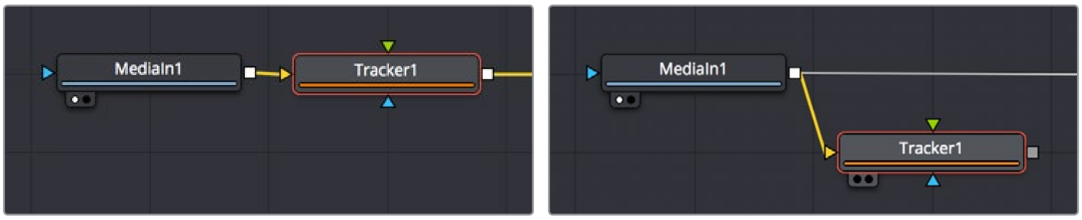

(Left) Tracker node connected serially so it can both track and transform the input image, (Right) Tracker connected as a branch to indicate it only contains tracking data, without using it directly

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 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:
Alternately, 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.

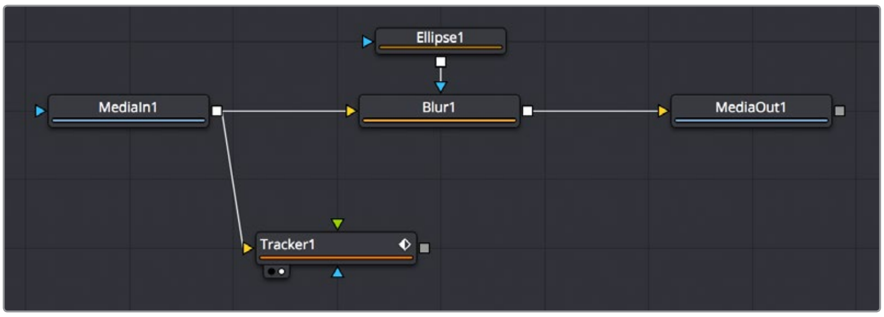

Blurring an actor's face by connecting tracking data to the center position of an Ellipse node

This is made easier by renaming the Tracker you created to something descriptive of what's being tracked.

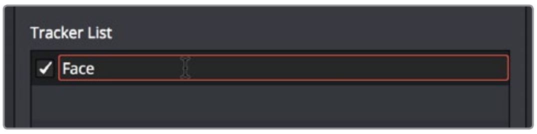

You can rename trackers in the Tracker list by double-clicking them and typing something descriptive

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 of 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.

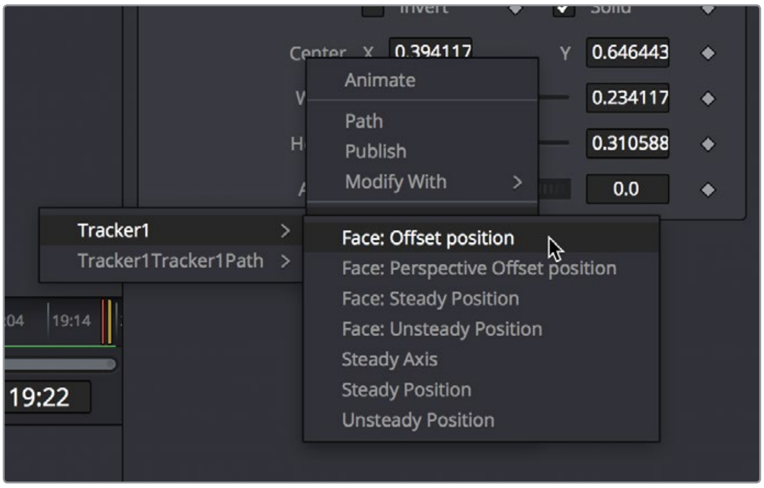

Right-clicking a parameter's label lets you connect tracking data to it to animate it

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.

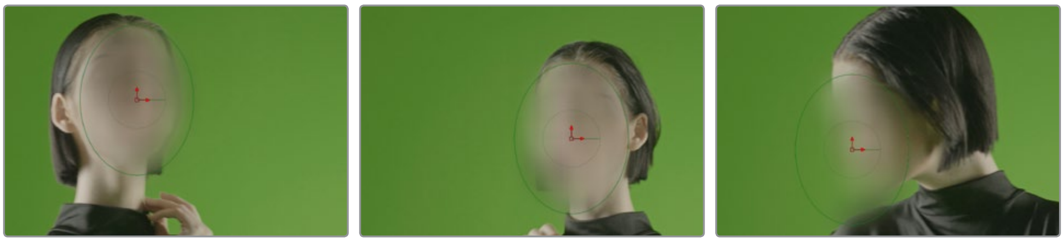

Connecting the center of the Ellipse node to the Face tracker animates the ellipse

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.

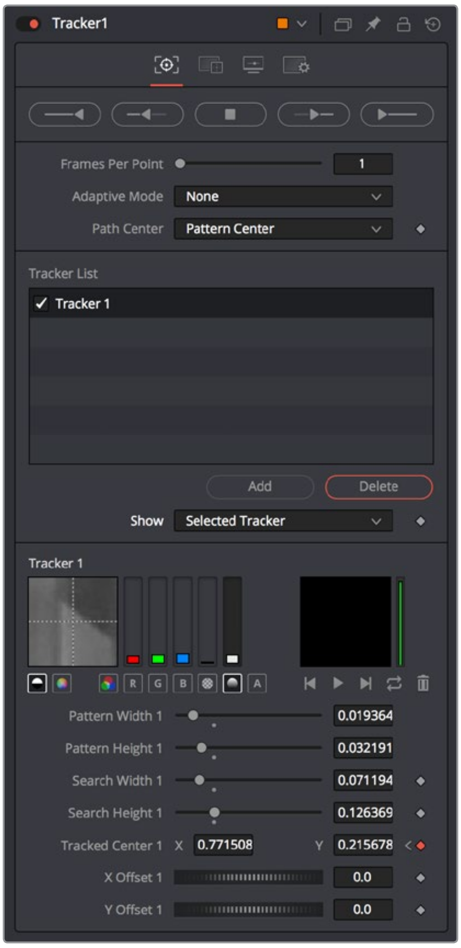

The Tracker Control tab

- **The Operations tab:** Where you decide how the tracking data is used.

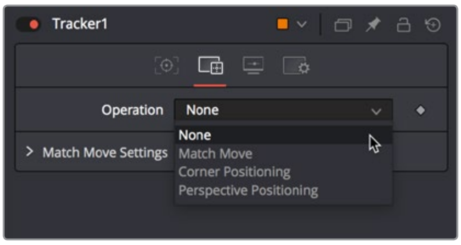

The Tracker Operations tab

- **The Display Options tab:** Where you can customize how the onscreen controls look in the Viewer.

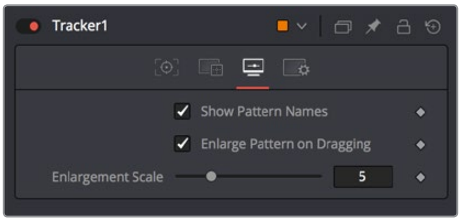

The Tracker Display Options tab

Motion Tracking Workflow In Depth

Analyzing motion using one or more trackers within the Tracker node is easy.

First: Connect a Tracker Node to the Image You Need to Track

Tracker nodes serve two purposes. They provide the nodes you need to analyze a subject you want to follow, and they serve as a container for the resulting track data. Regardless of whether or not you actually use the Tracker node itself to do anything with the tracking data, the background input (yellow) of a Tracker node must be connected to the output of the image you want to track 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.

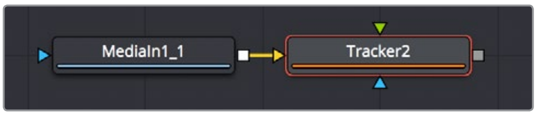

Connecting a MediaIn node's output to a Tracker node's background input

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 on the Add button immediately above the Tracker List control.

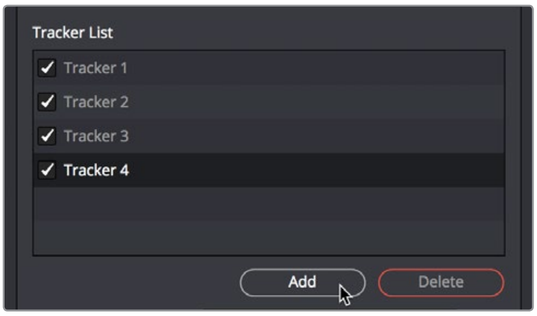

To add an additional tracker, click the Add button in the Inspector

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.

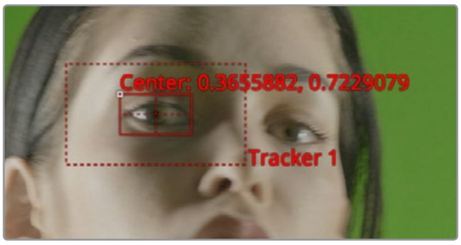

A pattern box positioned over an eye you want to track

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 has a dashed line and the inner 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.

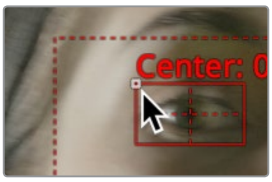

To move a pattern in the Viewer, drag it from the upper left corner

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.

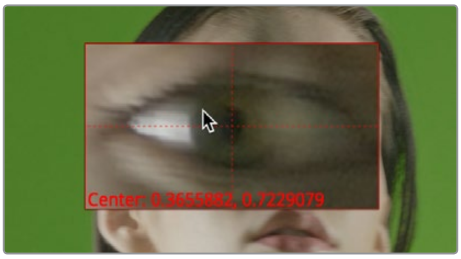

A zoomed view appears while you drag a pattern box to help you position it

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.

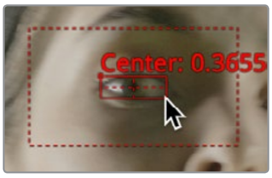

Resizing a pattern box to fit the eye

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 a 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.

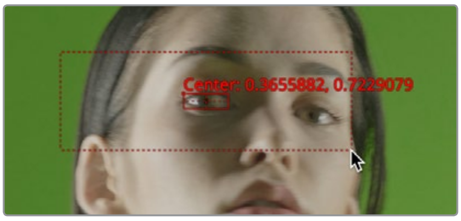

You can resize the search area by dragging the edges of the dotted outline

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 on 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.

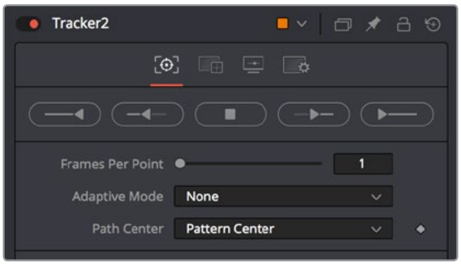

The tracking transport buttons and analysis parameters

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 anytime.

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.

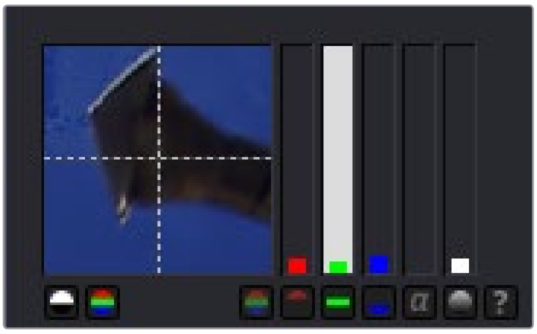

Channel bars indicate which channel is selected for tracking

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.

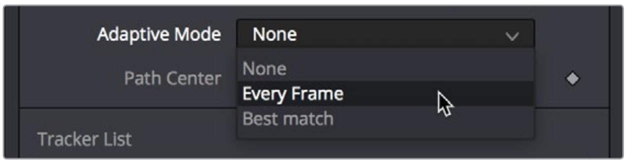

The Adaptive Mode options

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 re-acquired 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 re-acquire 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 take 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 non-linear motion in the tracked pattern.

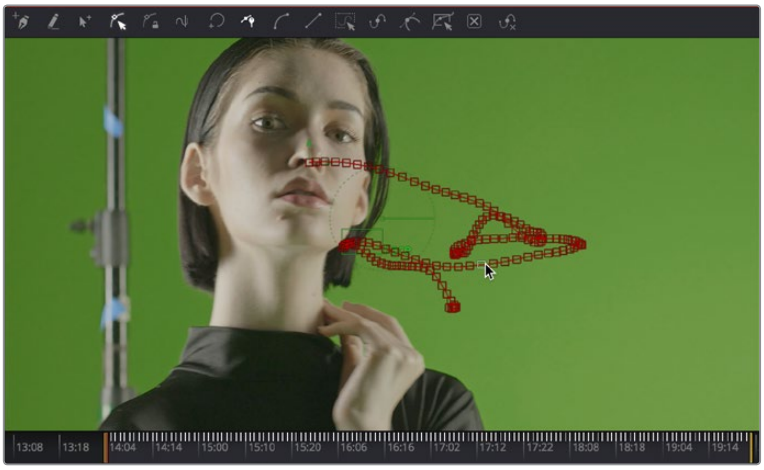

Tools for modifying tracker paths in the Node toolbar of the Viewer

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.

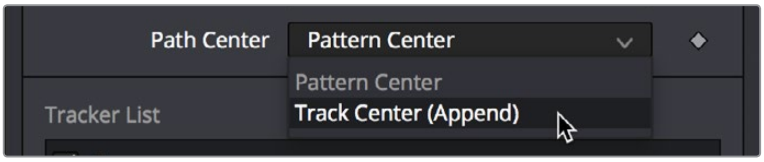

The Track Center (Append) mode pop-up menu

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.

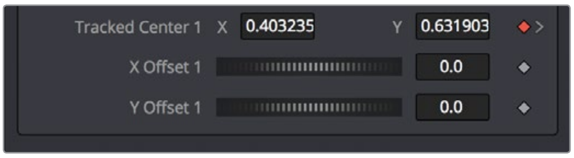

The Tracker Offset controls in the Inspector

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.

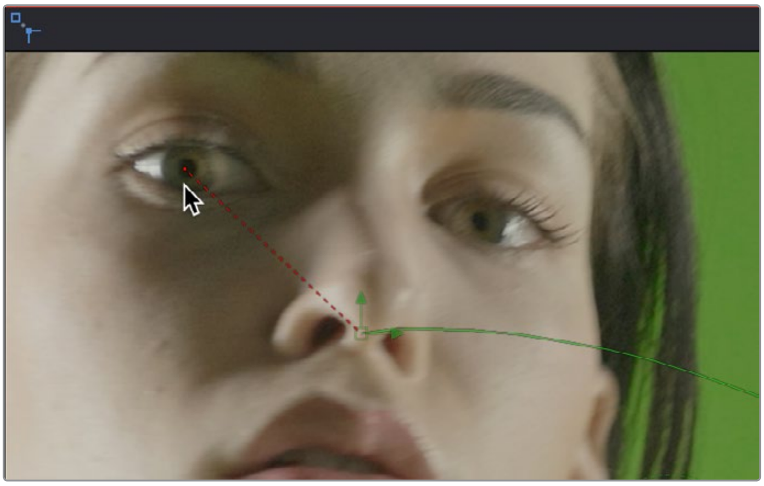

The Tracker Offset tool in the Node toolbar of the Viewer;
a track of the nose is being offset to the actor's eye

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.

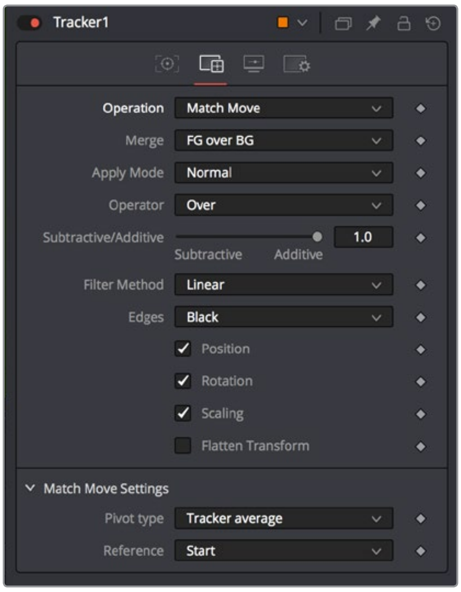

The Tracker Operation tab Match Move button

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.

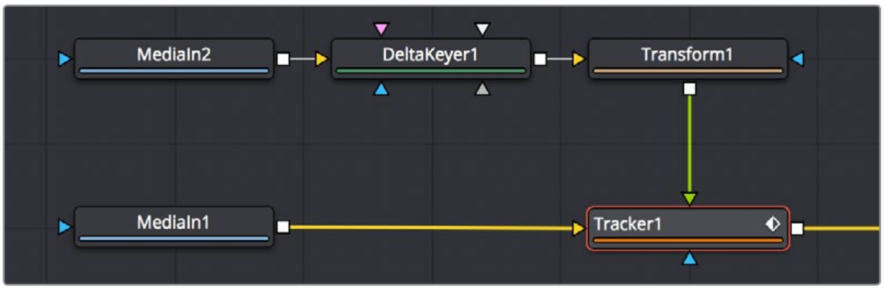

Connect a foreground image to the Tracker's foreground input

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, 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 actual 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 re-introduce 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 of 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 detailed information on the Tracking controls please see Chapter 92, "Tracker Nodes."

Chapter 58

Planar Tracking

This chapter provides an overview of how to use the Planar Tracker node, and how to use it to make match moves simple. For more information about the Planar Tracker node, see Chapter 92, “Tracker Nodes.”

Contents

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 re-applying those same perspective distortions to a different foreground.

TIP: Part of using the Planar Tracker is also knowing when to give up and fall back to using Fusion's Tracker node or to manual keyframing. Some shots are simply not trackable or the resulting track suffers from too much jitter or drift. The Planar Tracker is a time-saving node in the artist's toolbox and, while it can track most shots, it's not always a 100% solution.

Different Ways of Using the Planar Tracker Node

Like the other tracking nodes found in the Fusion page, 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.

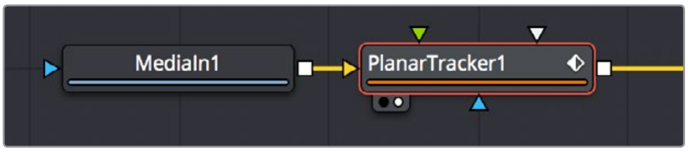

Connecting an image to the background input of a PlanarTracker node

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.

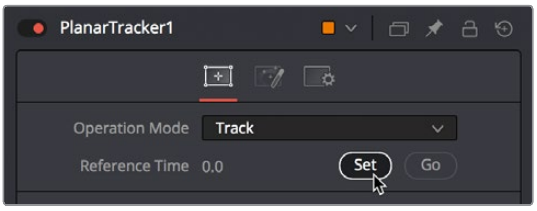

Clicking the Set button to set the reference frame to use for analysis

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, 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 lets you see how much detail is contributing to the track, and the direction of motion that's being analyzed.

During tracking, you can see Track Markers and
Trails to follow how well the track is going

6 Once the track is complete, play through the clip to visually inspect the track so you can evaluate how accurate it is. Does it stick to the surface? Switching to Steady mode can help here, as scrubbing through the clip in Steady mode will help you immediately see unwanted motion in the track.

7 Since we're doing a match move, click the Create Planar Transform button to export a Planar Transform node that will automatically transform either images or masks to follow the analyzed motion of the plane you tracked.

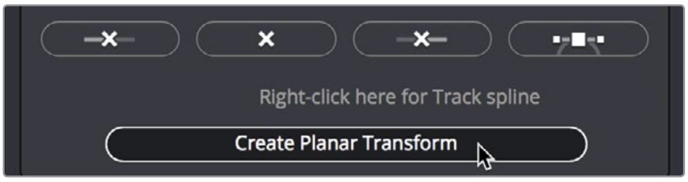

Clicking Create Planar Transform to create a node to use to transform other images or masks

In this case, the Planar Transform node will be inserted after a pair of Background and Paint nodes that are being used to put some irritatingly trendy tech jargon graffiti on the wall. The Planar Transform will automatically transform the Paint node's output connected to its background input to match the movement of the wall.

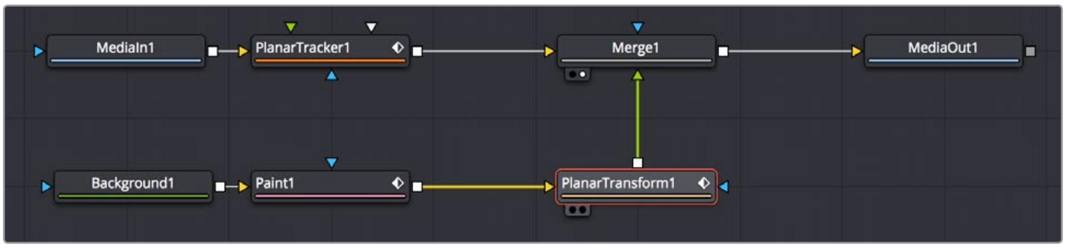

Adding the PlanarTransform node after a Paint node to match move it to the background image, combining it via a Merge node

The 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 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.

Chapter 59

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

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 information on the controls for each Mask node, see Part 7, "Fusion Page Effects."

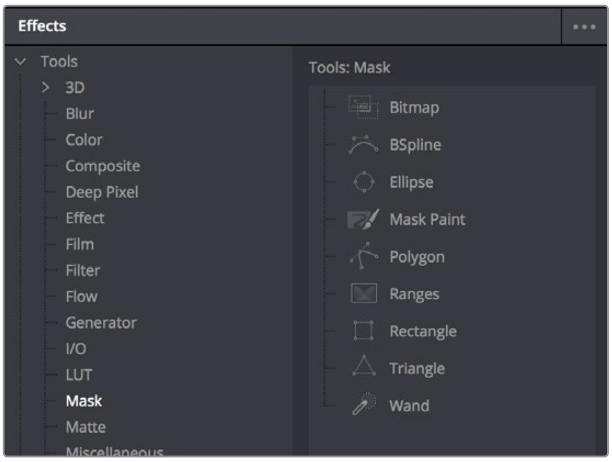

The available nodes in the Mask bin of the Effects Library

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 secondary color correction.

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 to the Color Corrector.

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 that are composed of control points and handles. Several points together are used to form the overall shape of a polyline.

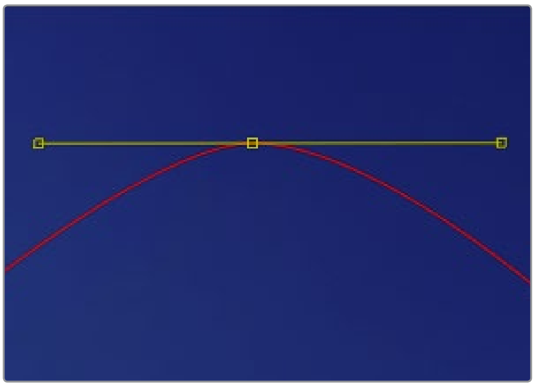

Bezier control point with direction handles extended to create a smooth curve

Each control point has a pair of handles that are 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.

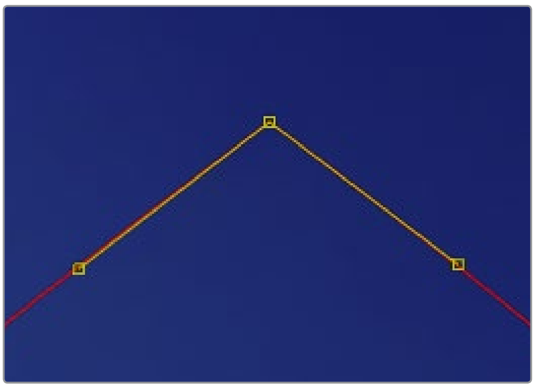

Bezier control point with direction handles aligned to create a linear segment

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.

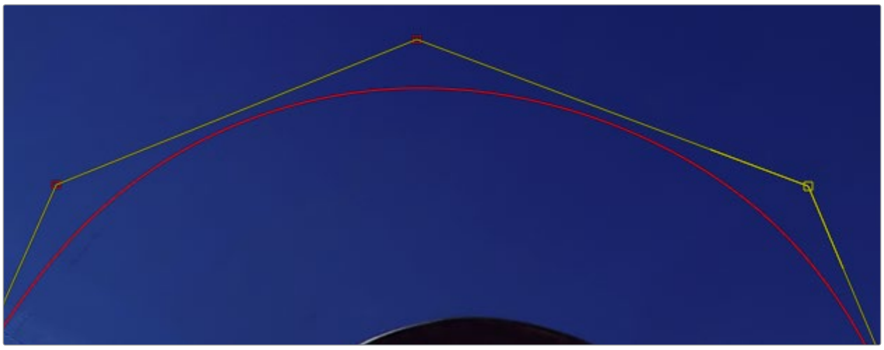

B-Splines excel at creating smooth curves

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.

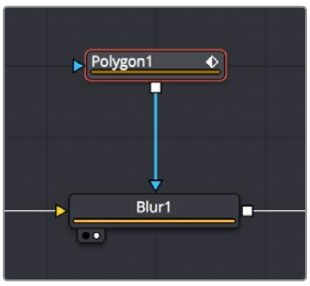

A Blur node with a Polygon node masking its effect

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.

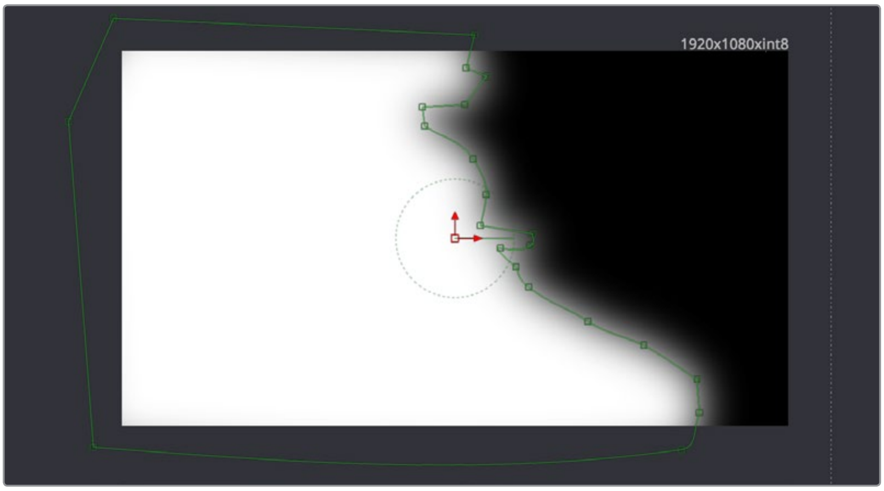

A Polygon node's mask seen in the Viewer

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.

The Fusion page 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.

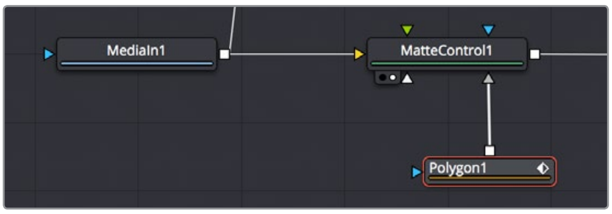

Feeding a Polygon node to a MatteControl node to do rotoscoping

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 MediaIn 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 load a MediaIn node directly into the Viewer while selecting the Mask node, you can draw a spline to rotoscope the image.

Rotoscoping a MediaIn node using a disconnected Polygon node

When you're finished rotoscoping, you simply connect the Polygon node's output to the MediaIn node's input, and an alpha channel is automatically added to that node.

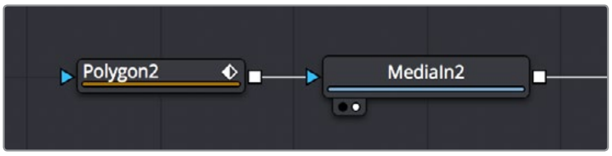

Connecting a Polygon node to a MediaIn node to use a spline as an alpha channel

TIP: If you connect a Mask node to a MediaIn node's effect input without any shapes drawn, that mask outputs full transparency, so the immediate result is that the image output by the MediaIn node becomes completely blank. This is why when you want to rotoscope by connecting a mask to the input of a MediaIn node, you need to work within a disconnected mask node first. Once the shape you're drawing has been closed, connect the Mask node to the MediaIn node'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.

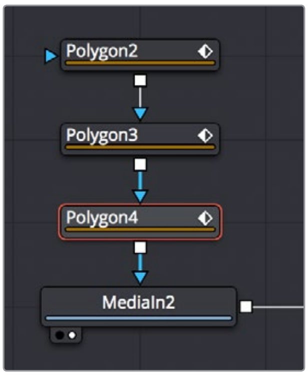

Combining multiple Polygon nodes one after the other in the node tree

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.

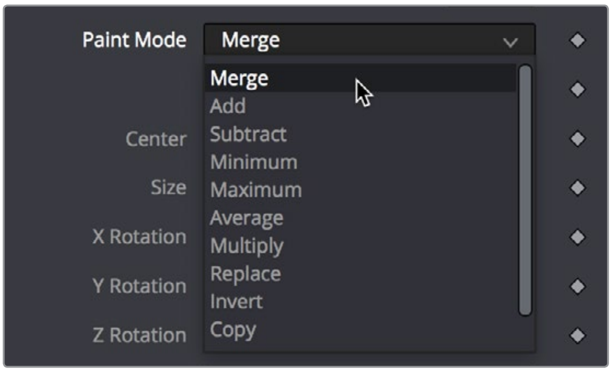

The Paint Mode parameter in the Polygon node Inspector parameters

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.

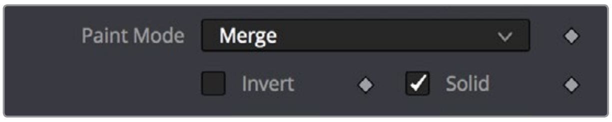

The Invert and Solid options

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 on the Fusion page has an Effect input (colored blue), which let you choose which parts of the image will or will not be affected by that node.

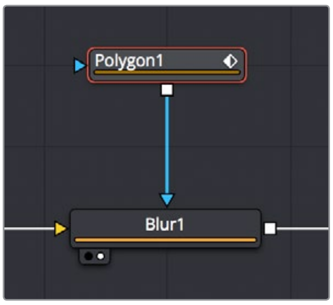

A Blur node with a Polygon node masking its effect

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 unaffected image data from the input over.

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 blue and green screen keys. In the following example, a lighting stand to the left is removed from the image via a B-Spline node's mask connected to the Garbage Matte input of the DeltaKeyer node.

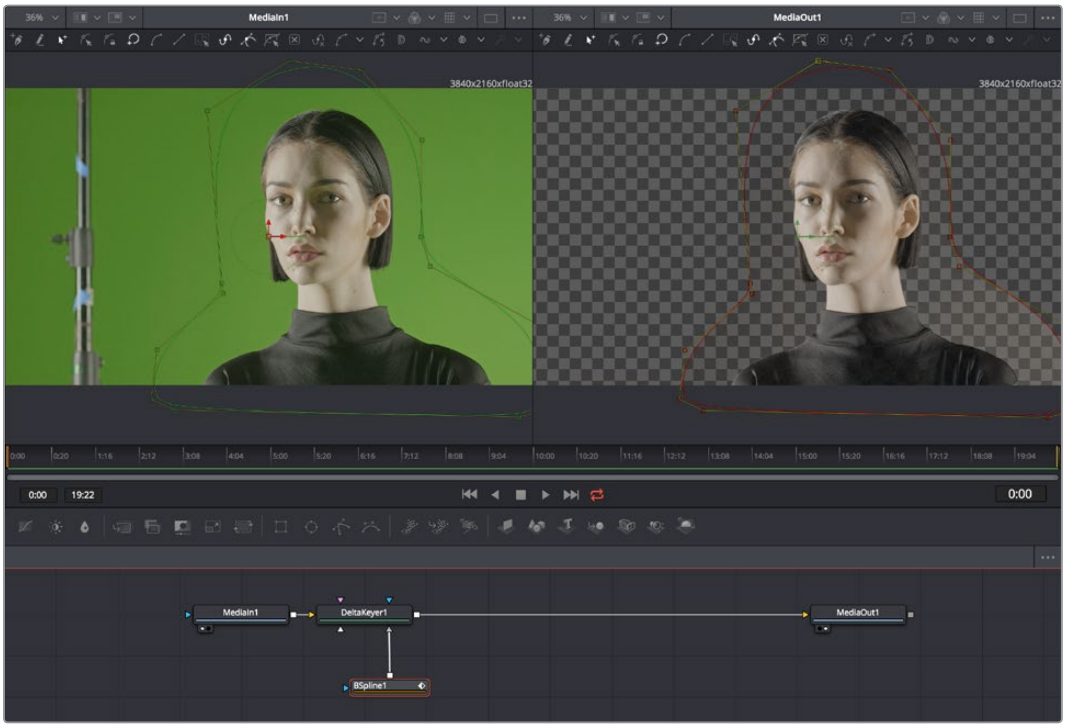

A B-Spline node is connected to the Garbage Matte input of a DeltaKeyer node to eliminate a light stand at the left of frame

TIP: You can quickly add a mask node to the Effect/Solid/Garbage Matte inputs of a keyer node by right-clicking the header bar of that node in the Inspector, and choosing whichever mask node you want to use from the Effect Mask, SolidMatte, and GarbageMatte submenus.

You choose whether a garbage matte is applied to a keying node as opaque or transparent in the Inspector for the node to which it's connected.

Solid Matte

Solid Matte inputs (colored white) are intended to fill unwanted holes in a matte, often with a less carefully pulled key producing a dense matte with eroded edges, although you could also use a polygon 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 could also be done with a MaskPaint or Polygon node, if necessary.

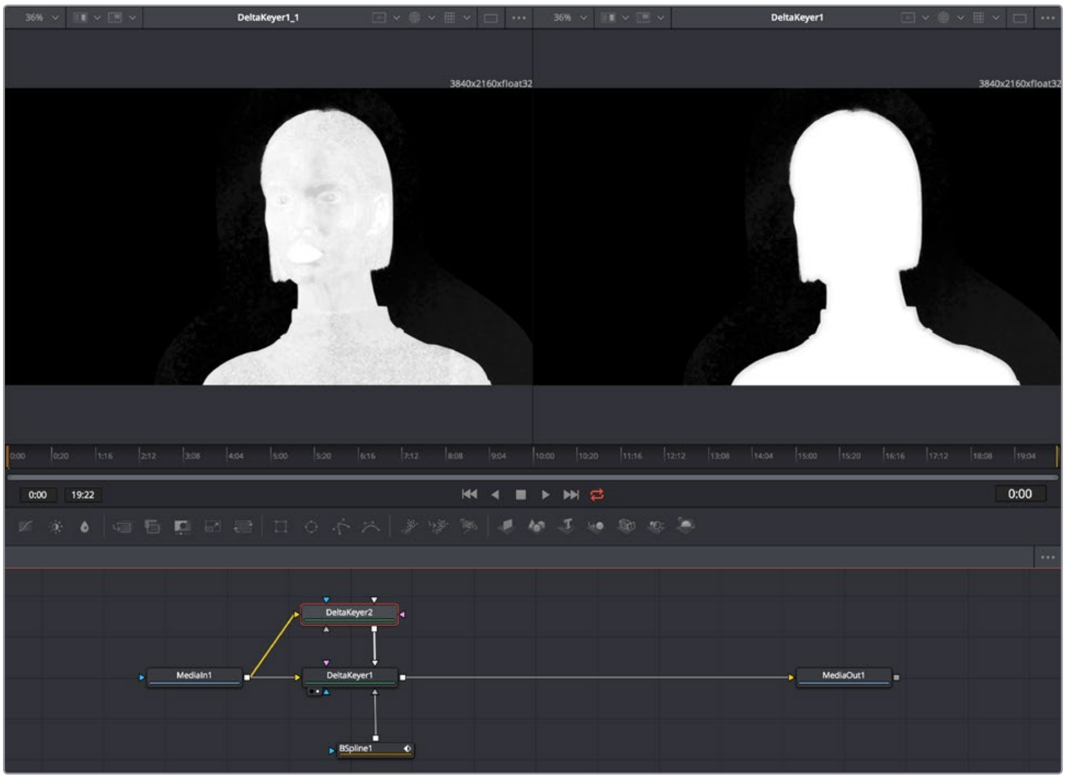

Filling in holes in the mask pulled by the DeltaKeyer1 node (at left) with another, harder but eroded key in DeltaKeyer2 that's connected to the SolidMatte input of DeltaKeyer1.

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.

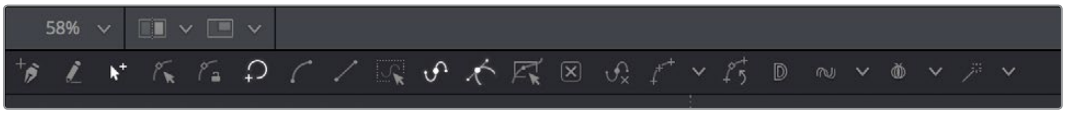

The Polyline toolbar

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 by found by right clicking on the toolbar and selecting from the options displayed in the contextual menu.

Selecting a Specific Polyline

It is possible to have several polylines in the Viewer at once if you select multiple Mask nodes in the Node Editor, so it's important to be able to switch between polylines easily.

To make a polyline active, do one of the following:

- Click one of the polyline's control points or segments.
- Press Tab and Shift-Tab to cycle between available polylines.
- Right click in the Viewer and choose the desired polyline by name from the Controls > Select menu.

Polyline Creation Modes

There are several different modes available from the toolbar for creating and modifying polylines. The specific mode used when a polyline is first added will depend on whether it is used as a path or a mask.

Each of the modes is described in more detail below.

Click Append

This mode is the default mode for mask creation. It's used to quickly define the rough shape of the mask, after which you switch to Insert and Modify mode to refine the mask further.

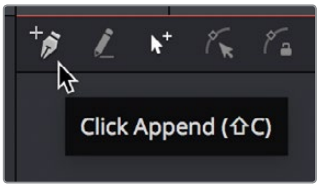

The Click Append toolbar button (Shift-C)

To create a mask using the Click Append mode, do the following:

1. Select Click Append from the toolbar or press Shift-C.
2. Click the 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 short cut. 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.

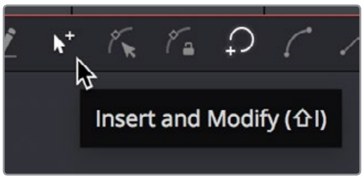

The Insert Modify toolbar button (Shift-I)

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.

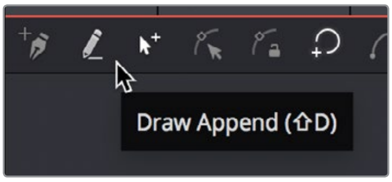

The Draw Append toolbar button (Shift-D)

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.

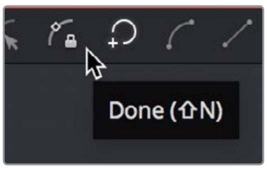

The Done toolbar button (Shift-N)

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, 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.

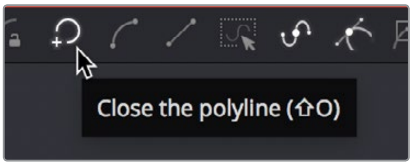

The Close toolbar button (Shift-O)

All of 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. In Fusion there are a number of simple techniques for selecting, moving and smoothing a spline, but also 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 to select a continuous range of points.
- Hold Command (macOS) or Ctrl (Windows) and click each control point you want to add or remove.
- Press Command-A or Ctrl-A to select all of 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 counter clockwise 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 (Mac OS X) or Alt (Windows) 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 (macOS) or Ctrl-Arrow Keys (Windows) 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.

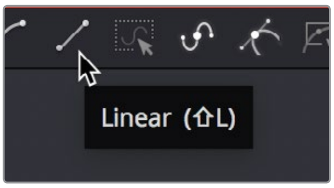

The Linear button in the toolbar (Shift-L)

Transforming Individual or Multiple Points

Select the points to be transformed, 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 the Delete (macOS) or Backspace (Windows), choosing Delete from the contextual menu, or by clicking the Del button in the toolbar. The shape of the polyline changes to reflect the removal of these points.

TIP: Deleting all of 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 only appear 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 (macOS X) or Ctrl (Windows) while dragging a handle. Subsequent changes will maintain the relationship, unless the Command or Ctrl is held during each adjustment.

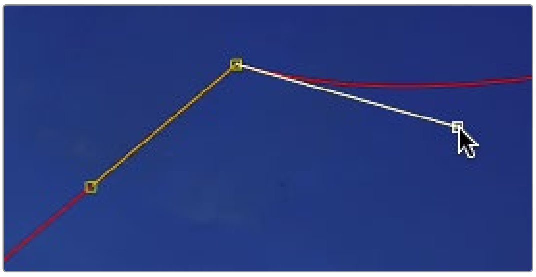

Hold Command or Ctrl to adjust one handle independently

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 dialog and allow you to reposition one or more selected control points.

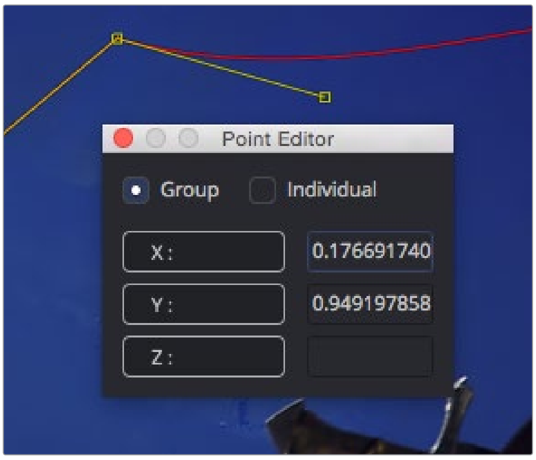

The Point Editor dialog can be used to position control points

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.

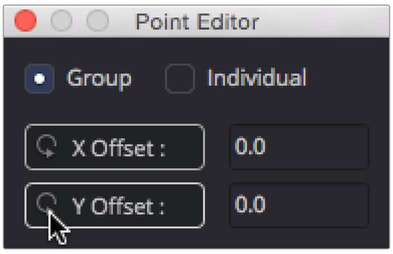

The Point Editor dialog with Offset values

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.

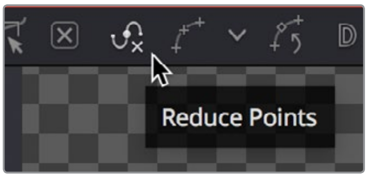

The Reduce Points button in the toolbar

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.

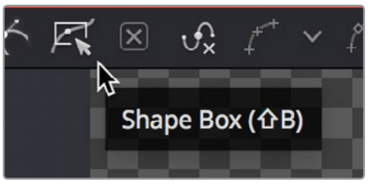

The Shape Box button in the Polyline toolbar

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.

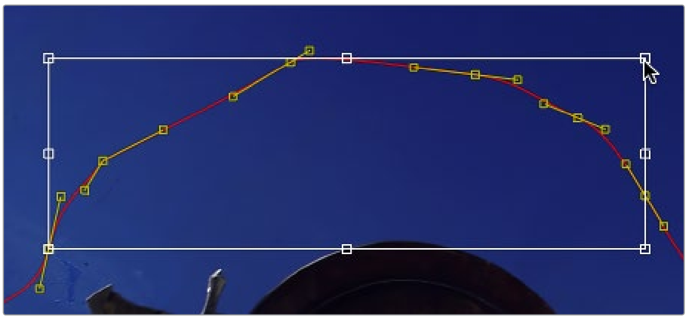

Dragging a side handle resizes along a specific axis

Holding Command or Ctrl 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.

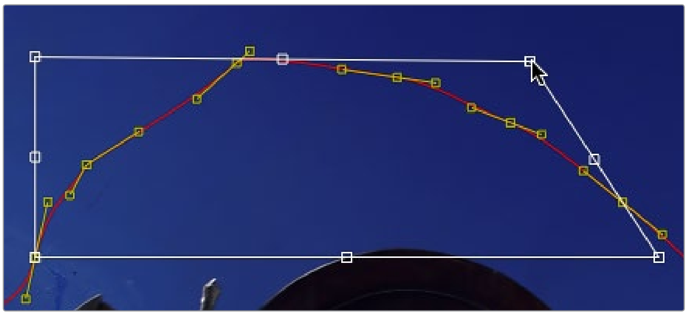

Hold shift while dragging a corner to perform non-uniform transformations

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.

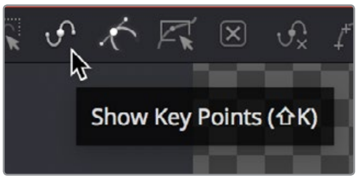

The Show Key Point and Show Handles buttons in the toolbar

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 only happen 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 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 8:** 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.

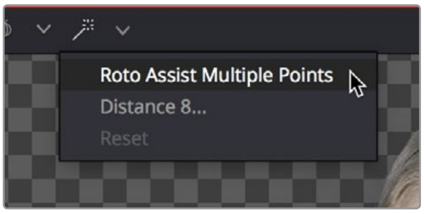

The Roto Assist options in the toolbar

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 often 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.

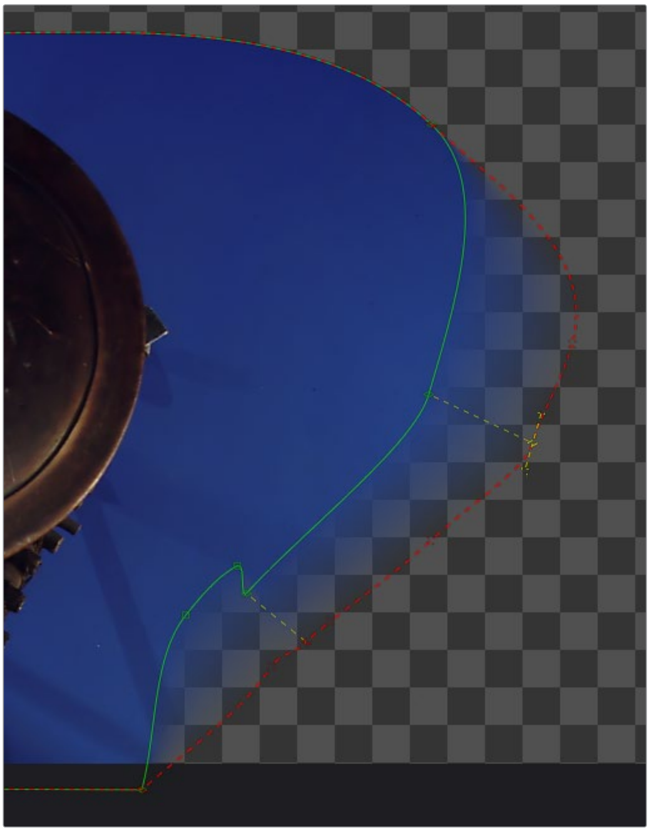

A double polyline uses an inner and outer shape for non-uniform softness

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.

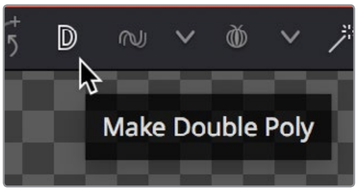

Make Double Polyline button

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 onscreen controls until the dashed outline is visible, or you can select the outer polyline using the contextual menu's Controls > Select menu.

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 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 disabling "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, they're automatically ready to be animated. All you have to do to animate a mask is move the playhead to a new frame 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.

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 only display one keyframe for the entire shape at any given frame.

This default keyframing behavior is convenient when quickly animating shapes from one form to another, but it doesn't allow for specific individual control points that need to be keyframed independently of all other control points for a particular shape. If you're working on a complex mask that would benefit from more precise timing or interpolation of individual control points, you can expose one or more specific control points on a polyline by publishing them.

Be aware that publishing a control point on a polyline removes that point from the standard animation spline. From that point forward, that control point can only be animated via its own keyframes on its own animation spline. Once removed, this point will not be connected to paths, modifiers, expressions, or trackers that are connected to the main polyline spline.

To publish a selected point or points, do one of the following:

- Click on the Publish Points button in the Polyline toolbar.
- Select Publish Points from the Polyline's contextual menu.

A new coordinate control is added to the Polyline mask controls for each published point, named Point 0, Point 1, and so on.

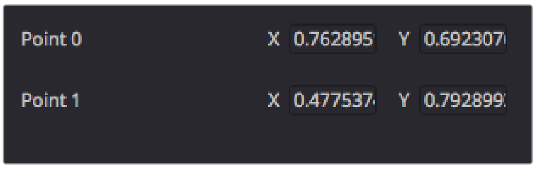

The Publish Points controls in the Inspector

The onscreen control indicates published points on the polyline by drawing that control point much larger. Once a published point is created, it can be connected to a tracker, path, expression or modifier by right-clicking on this control and selecting the desired option from the point's contextual menu.

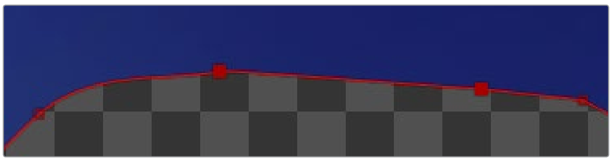

The published point in the Viewer

Using "Publish Point 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 Point 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 Point 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.

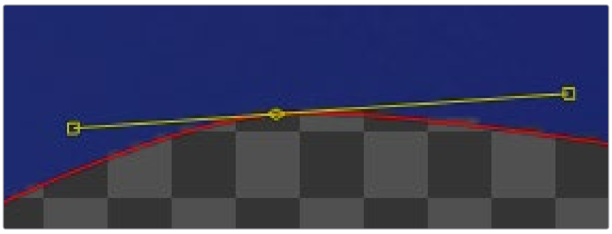

A control point set to Follow Points

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 60

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

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 of the Fusion page, 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

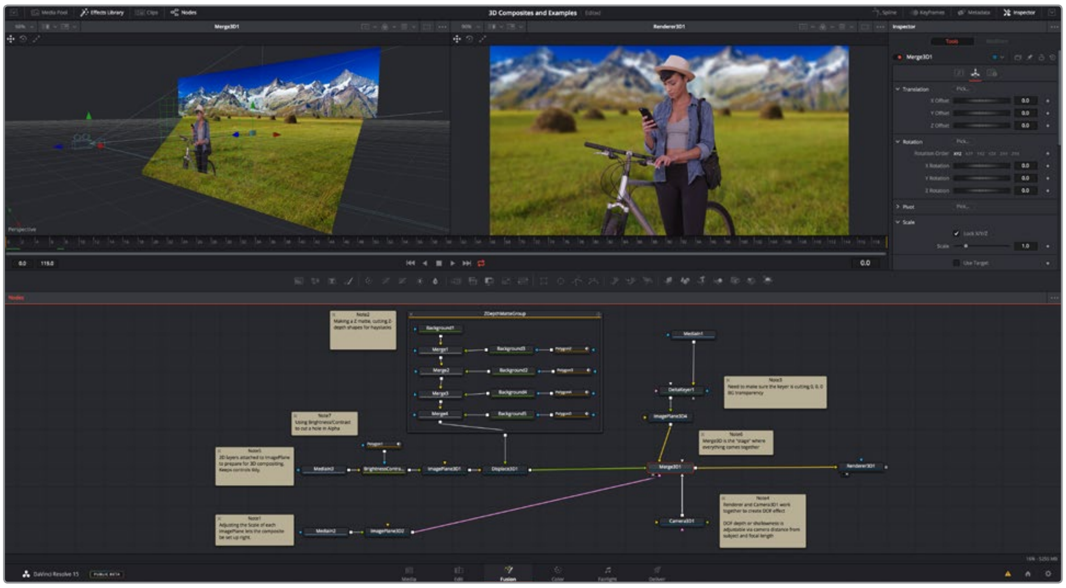

An example 3D scene in the Fusion page of DaVinci Resolve

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 of these elements into a scene. Nearly all of these nodes are collected within the 3D category of nodes found in the Effects Library.

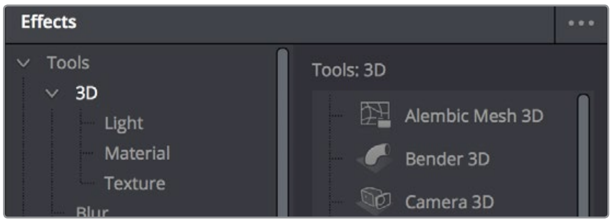

The 3D category of nodes in the Effects Library

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're only going to 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.

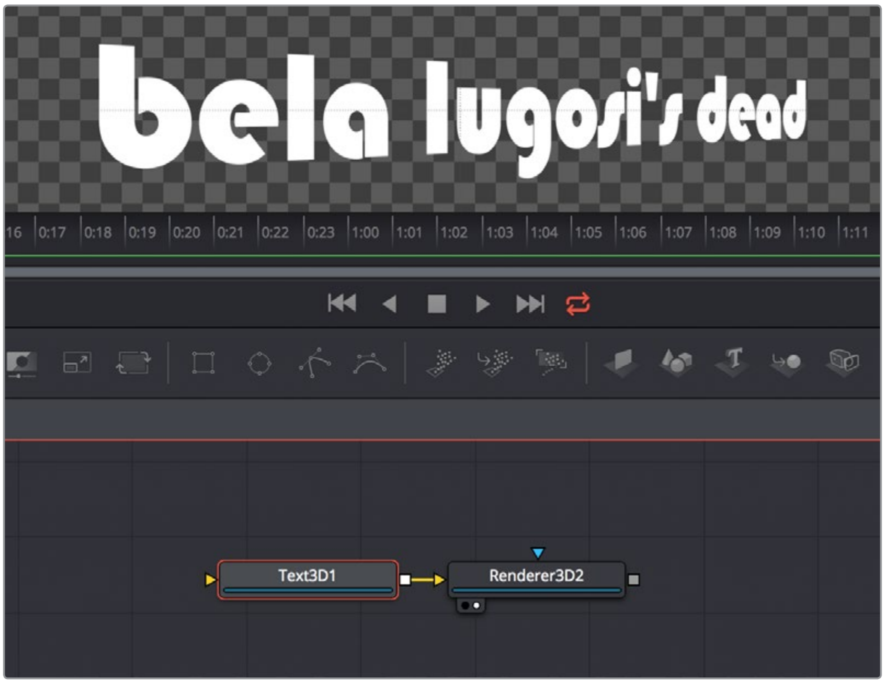

A simple 3D scene with a Text3D node connected directly to a Renderer 3D node

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 Object3D)
- A light node (such as SpotLight)
- A camera node
- A Merge3D node
- A Renderer3D node

All of 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 aMerge3D 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 Fusion page 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 right to left, correspond 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 that are 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 Merge 3D node either singly or in multiples, to create sophisticated results combining multiple elements.

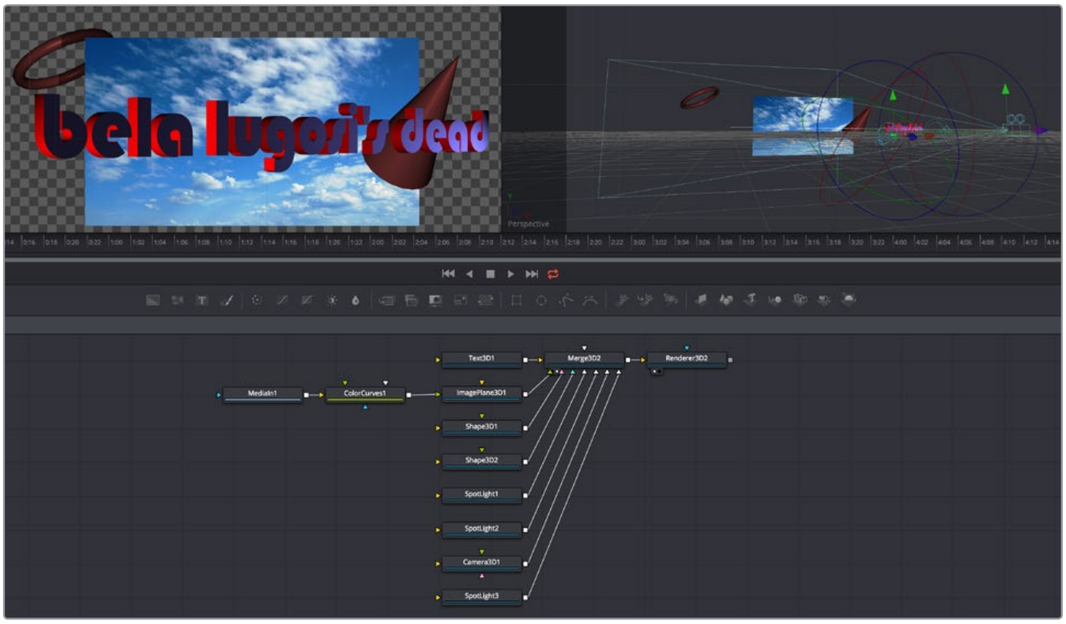

A more complex 3D scene combining several geometry nodes including the Text3D, Shape3D, and ImagePlane3D nodes

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.

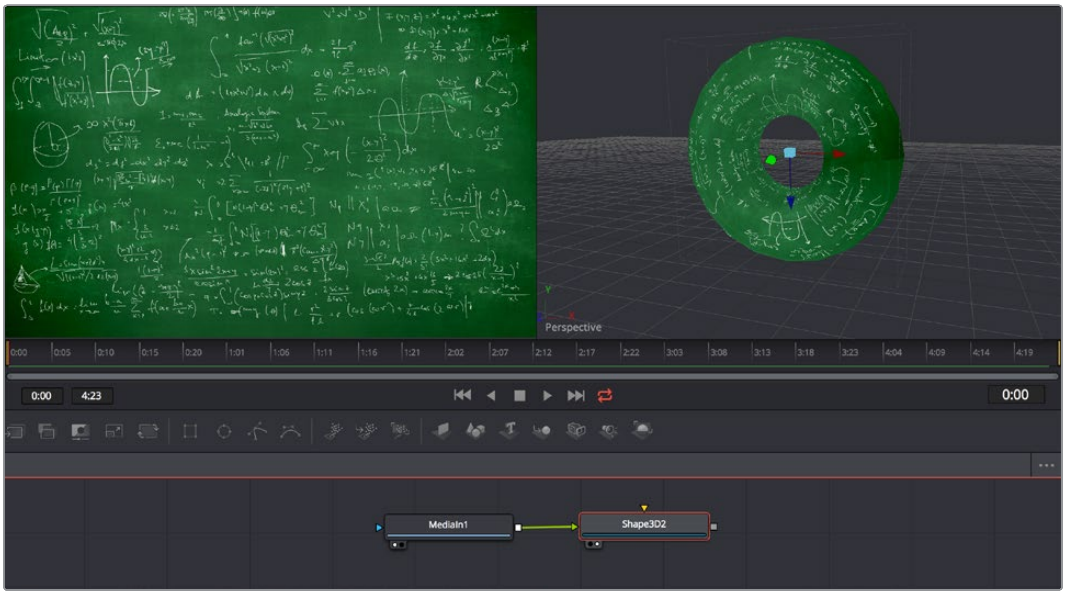

An image connected to the material input of a Shape3D node set to Taurus; (Left) the image, (Right) the shaded taurus

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 pretty powerful.

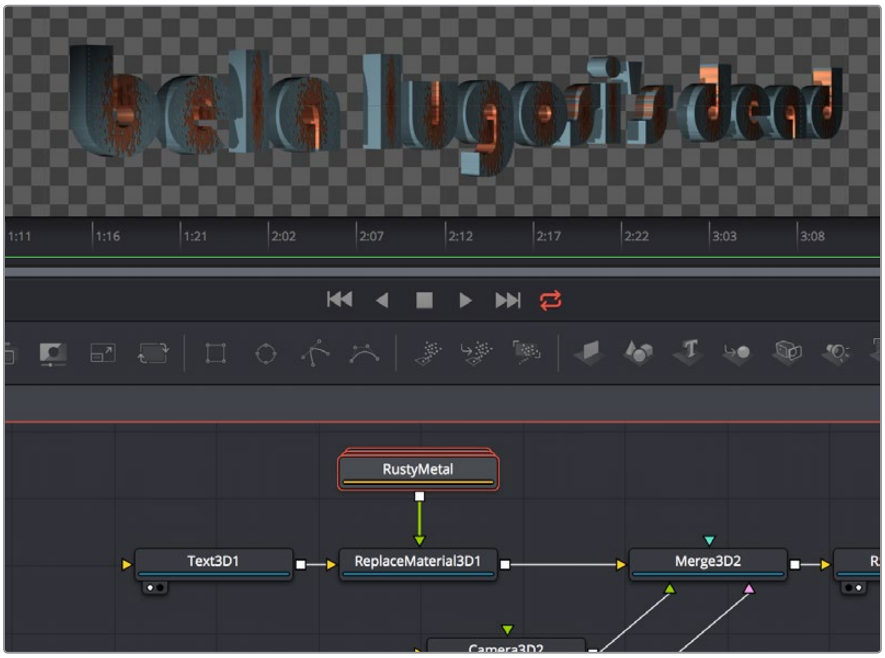

The geometry created by a Text3D node is textured using a shader connected to a ReplaceMaterial3D node that's connected downstream of the object you want to shade

The Merge 3D Node

The Merge 3D 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 Merge 3D node lets you connect an unlimited number of inputs, with the resulting output combined according to each object's absolute position in 3D space.

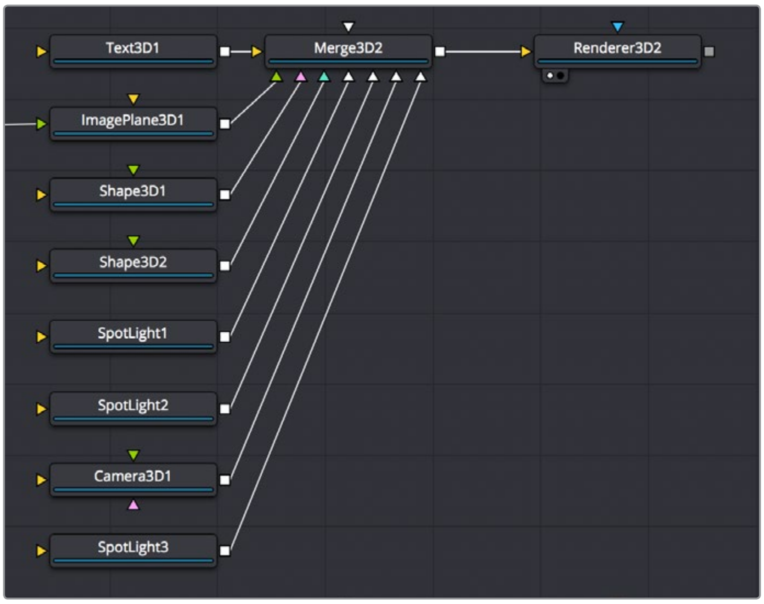

Merging many objects together in a 3D scene using the Merge3D node

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.

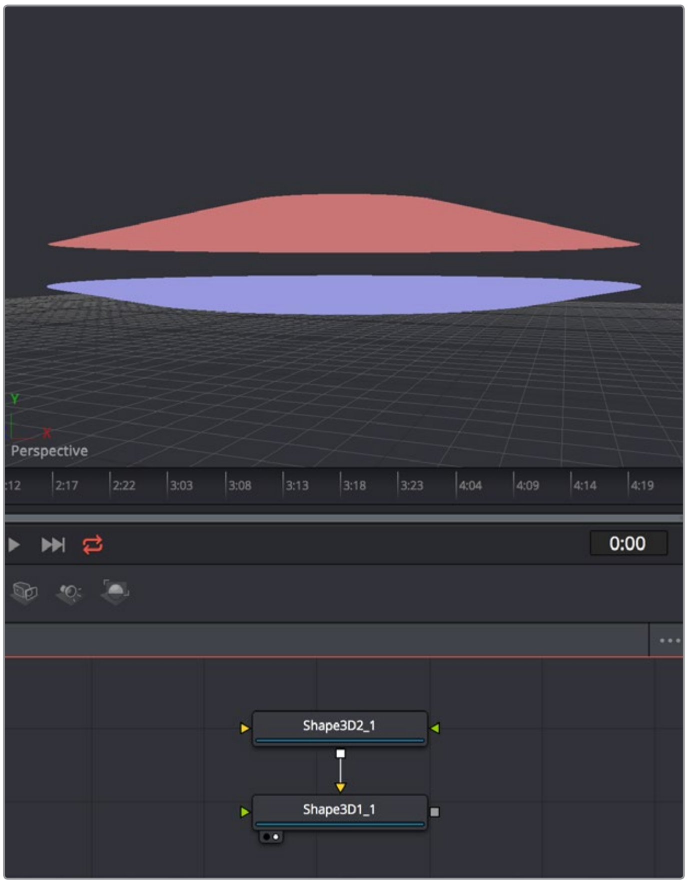

Connecting one Shape3D node to another directly to combine them, transforming the last downstream 3D object also transforms all upstream objects; the last Shape3D node is viewed, showing both

Combining Multiple Merge3D Nodes

Furthermore, Merge 3D nodes can be combined with other Merge 3D nodes, allowing you to create composite 3D scenes that are made up of multiple "sub-scenes," each put together within individual Merge 3D nodes.

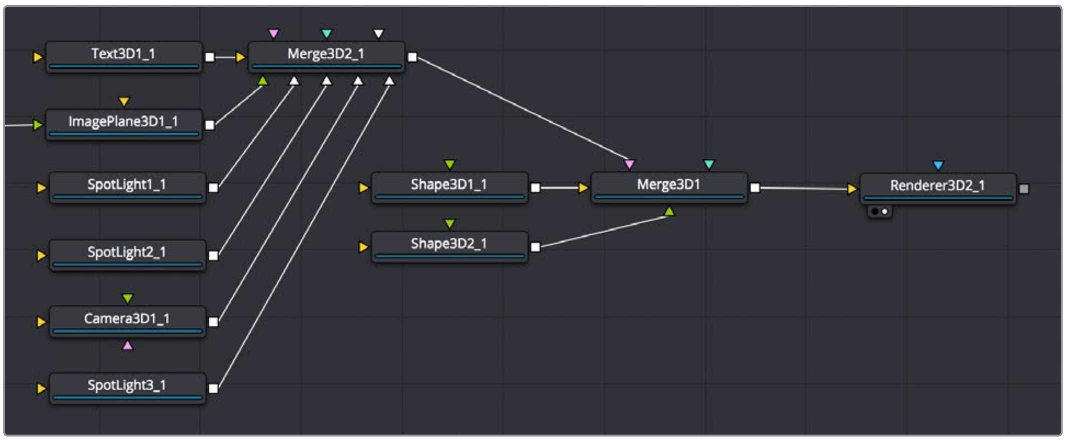

You can build elaborate scenes using multiple Merge3D nodes connected together

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.

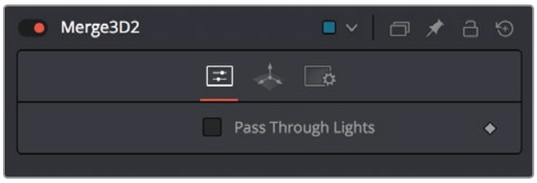

You can light downstream Merge3D scenes with lights connected to upstream Merge3D scenes by turning on "Pass Through Lights"

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.

(Left) The result of lights on the text in one Merge3D node not affecting the cone and taurus added in a downstream Merge3D node, (Right) Turning on "Pass Through Lights" in the upstream Merge3D node results in those lights also illuminating the downstream shapes

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 Merge 3D 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 that are specific to that particular node's scene.

Transforming Upstream, Lighting Downstream

It's common, when building complex scenes using multiple Merge3D nodes being combined together, 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.

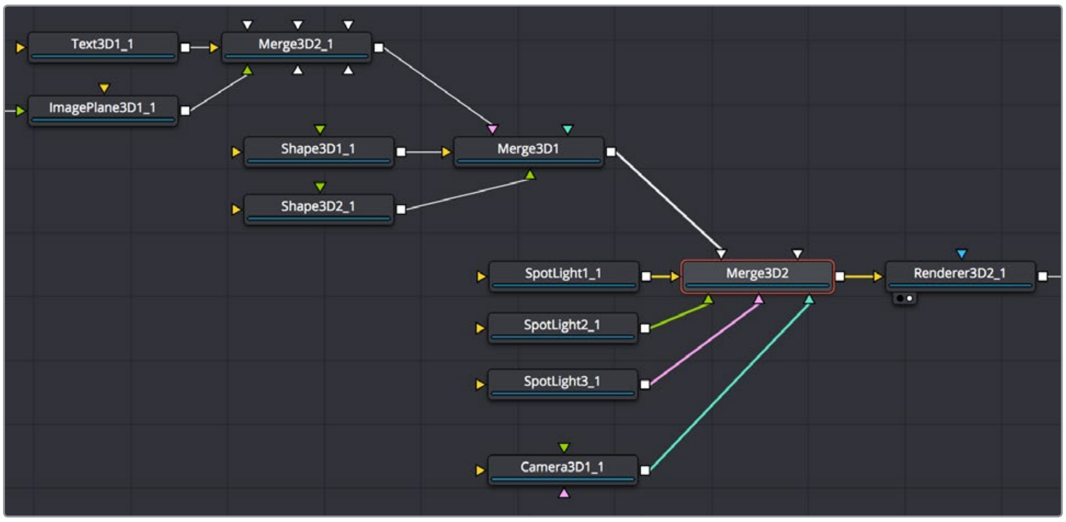

An example of a 3D scene using multiple Merge3D nodes that are working together; the upstream Merge3D nodes arrange the 3D objects placed within the scene, while the last Merge3D node (orange) lights and frames the scene

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 MediaOut node to output to the Color page 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.

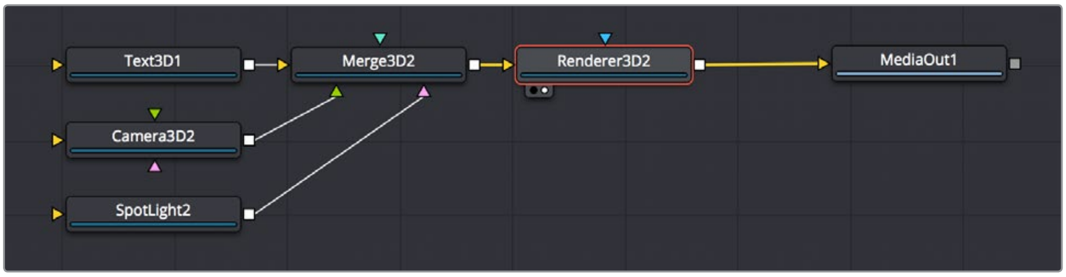

Output of a Merge 3D connected to a Renderer 3D node to output 2D image data

The Renderer 3D 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 Renderer 3D 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 naturalistic, 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.

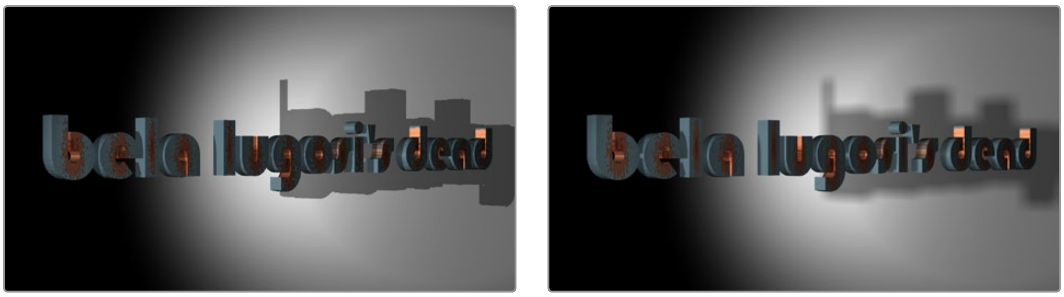

(Left) When the Renderer3D node "Renderer Type" pop-up is set to OpenGL Renderer, you cannot render soft shadows or excessively large textures, (Right) When the Renderer3D node "Renderer Type" pop-up is set to Software Renderer, you can render higher quality textures and soft shadows

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 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 that are applied to upstream objects, at the resolution that's specified within the Image tab of that Renderer3D node.

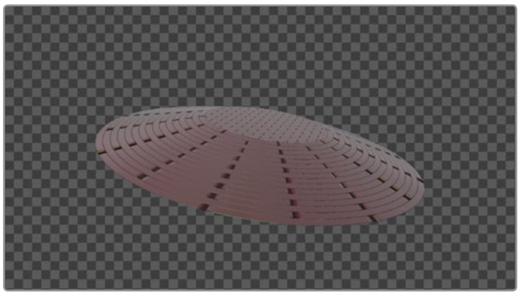

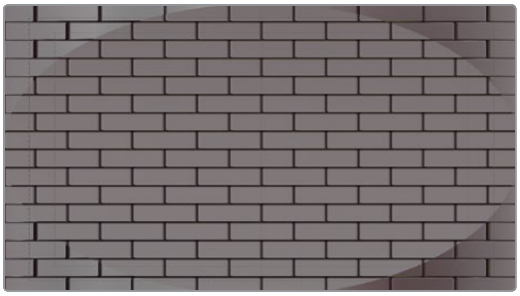

(Left) A normally rendered 3D scene, (Right) The same scene rendered using the OpenGL UV Renderer mode of the Renderer3D node

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.

For instance, 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 of those onto the mesh, UV render it out, 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, 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.

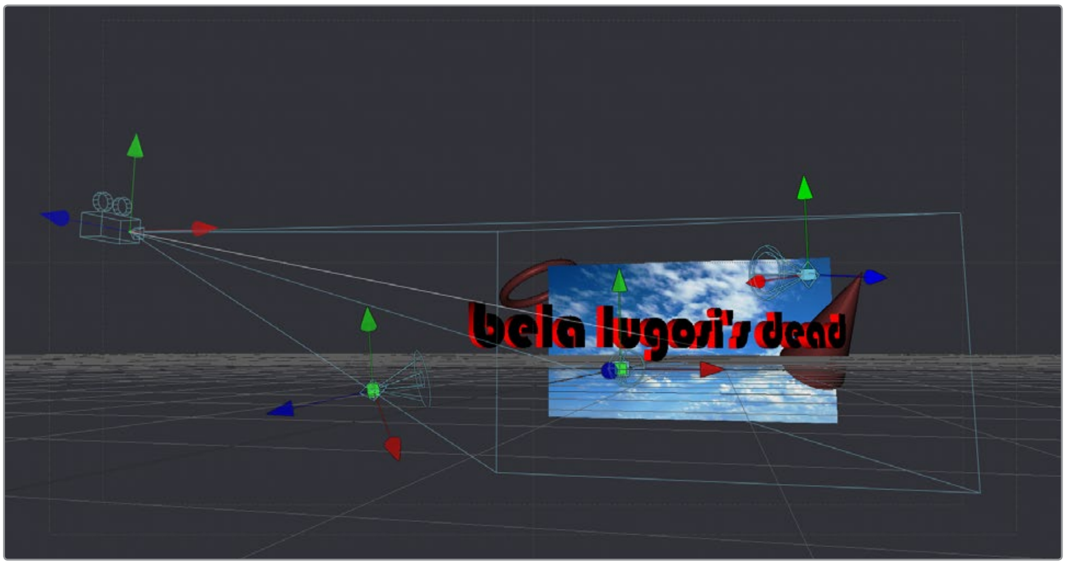

The 3D Viewer

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.

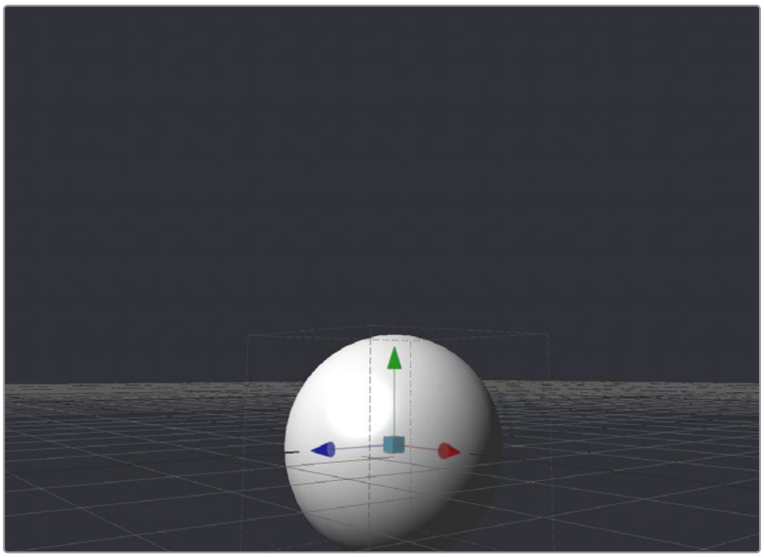

A 3D Viewer's default perspective view

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.

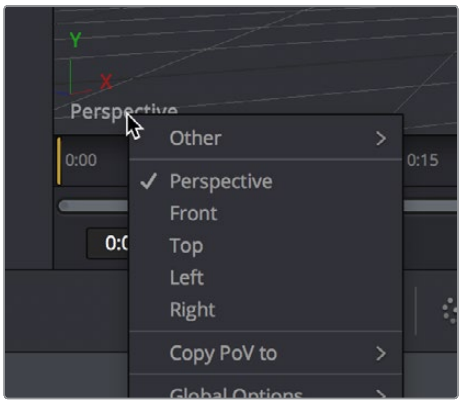

Right-click the Axis label of the Viewer to change the viewpoint

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.

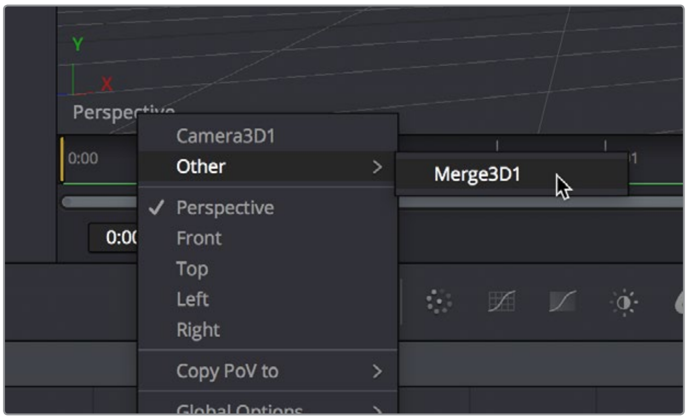

The Perspective pop-up menu also shows cameras, lights, and Merge3D and Transform3D nodes you can switch to

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 information about the options available in the 3D Viewer, see Chapter 52, "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 A to Fit all objects in the Viewer.
   - Press F to Fit to selection (or Fit All if nothing is selected).
   - Press C 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.

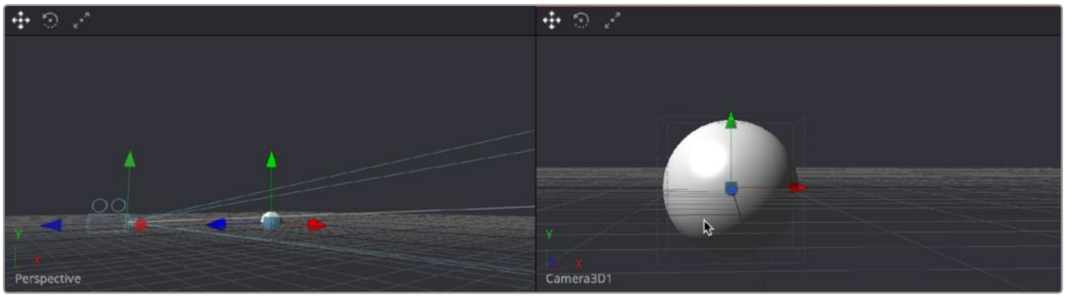

When a viewer is set to display the view of a camera or light, panning, zooming, or rotating the Viewer (seen at right) actually transforms the camera or light you're viewing through (seen at left)

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 or enabling the Fast button in the Viewer toolbar. 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 Renderer 3D 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

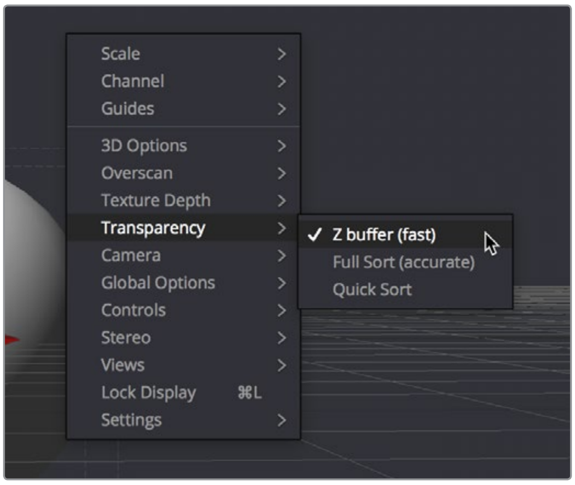

Transparency Sorting in the Viewer contextual menu

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.

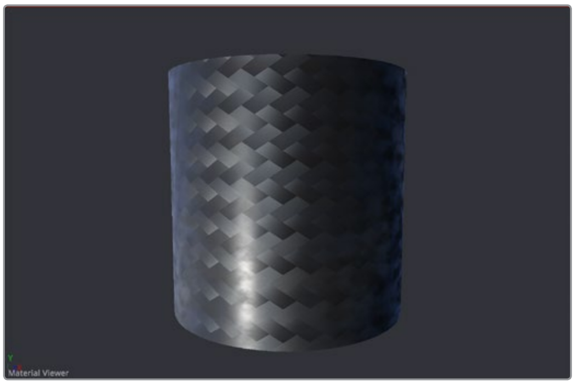

The Material Viewer mode of the Viewer

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.

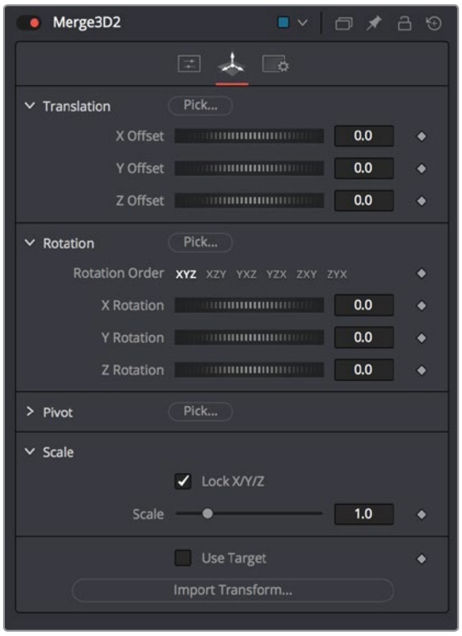

The Transform tab of a Merge3D node

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 if they’re locked together or not). The same adjustments can be made in the Viewer using on-screen controls.

Onscreen Transform Controls

When an object is selected, it displays onscreen Transform controls in the Viewers that allow you to adjust the object's position, rotation and scale. There are buttons in the Transform toolbar so you can switch modes, or you can use the keyboard short cuts.

To switch Transform modes, use the following keyboard short cuts:

- Press Q for Position
- Press W for Rotation
- Press E for Scaling

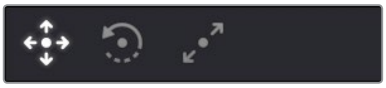

The Position, Rotation, and Scale modes in the Transform toolbar

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 onscreen control to constrain the transform to that axis, or if you drag the center of the onscreen control you can apply a transform without constraints. Holding Option and dragging in the Viewer allows you to freely translate in all three axes without clicking on a specific control.

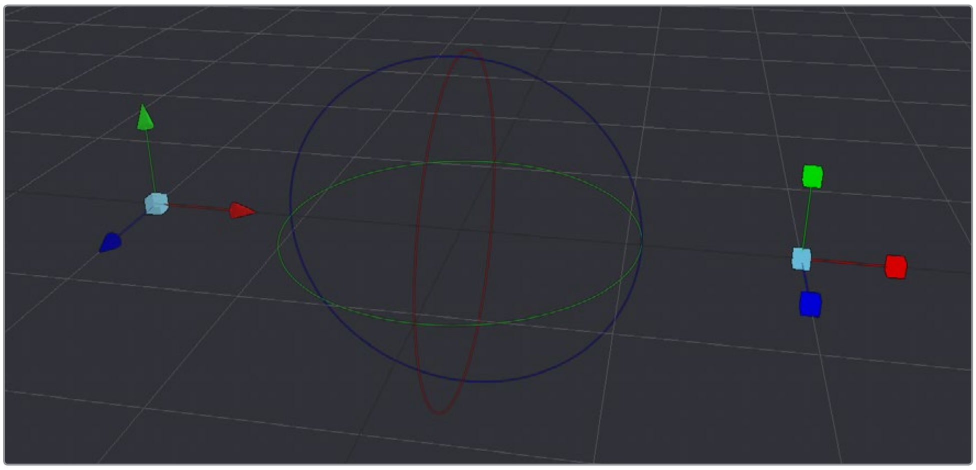

From left to right, the Position, Rotation, and Scale onscreen Transform controls

If the Scale's Lock XYZ checkbox is enabled in the Inspector, only the overall scale of the object is adjusted by dragging the red or center onscreen control while the green and blue portions of the onscreen control have no affect. If you unlock the parameters, you are able to scale an object along individual axes separately to squish or stretch the object.

Selecting Objects

With the onscreen controls visible in the Viewer, you can select any object by clicking on its center control. Alternately, you can also select any 3D object by clicking its node in the Node Editor.

Pivot

In 3D scenes, objects rotate and scale around an axis called a pivot. By default, this pivot goes through the object's center. If you want to move the pivot so it is offset from the center of the object, you can use the X, Y and Z Pivot parameters in the Inspector.

Target

Targets are used to help orient a 3D object to a specific point in the scene. No matter where the object moves, it will rotate in the local coordinate system so that it always faces its target, which you can position and animate.

To enable a target for a 3D object:

1. Select that object's node.
2. Open the object's Transform panel in the Inspector.
3. Turn on the Use Target checkbox.

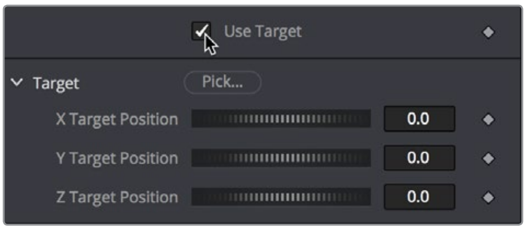

Turning on the Use Target checkbox of a Shape3D node's

4. Use the X/Y/Z Target Position controls in the Inspector or the Target onscreen control in the Viewer to position the target and in turn position the object it's attached to.

In the Viewer, a line is drawn between the target and the center of the 3D object it's attached to, to show the relationship between these two sets of controls. Whenever you move the target, the object is automatically transformed to face its new position.

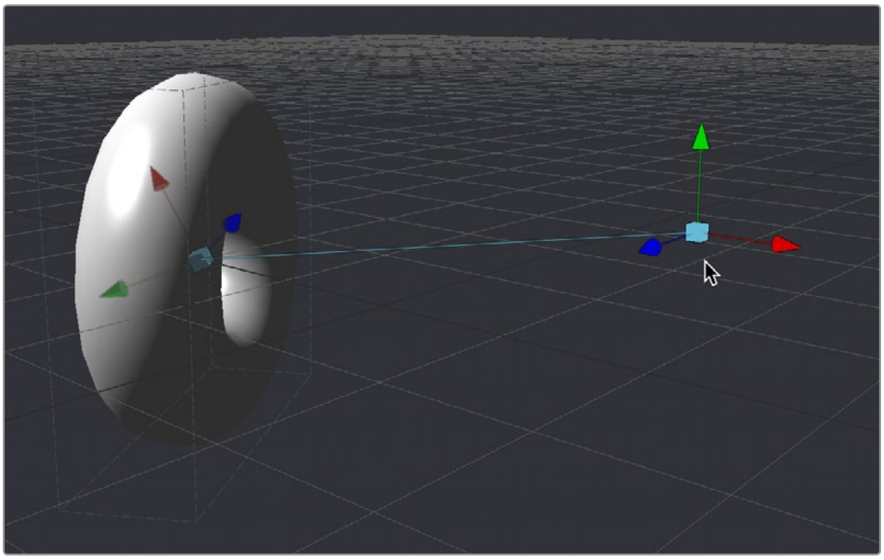

A taurus facing its onscreen Target controls

For example, if a spotlight is required in the scene to point at an image plane, enable the spotlight's target in the Transform tab and connect the target's XYZ position to the image plane's XYZ position. Now, no matter where the spotlight is moved, it will rotate to face the image plane.

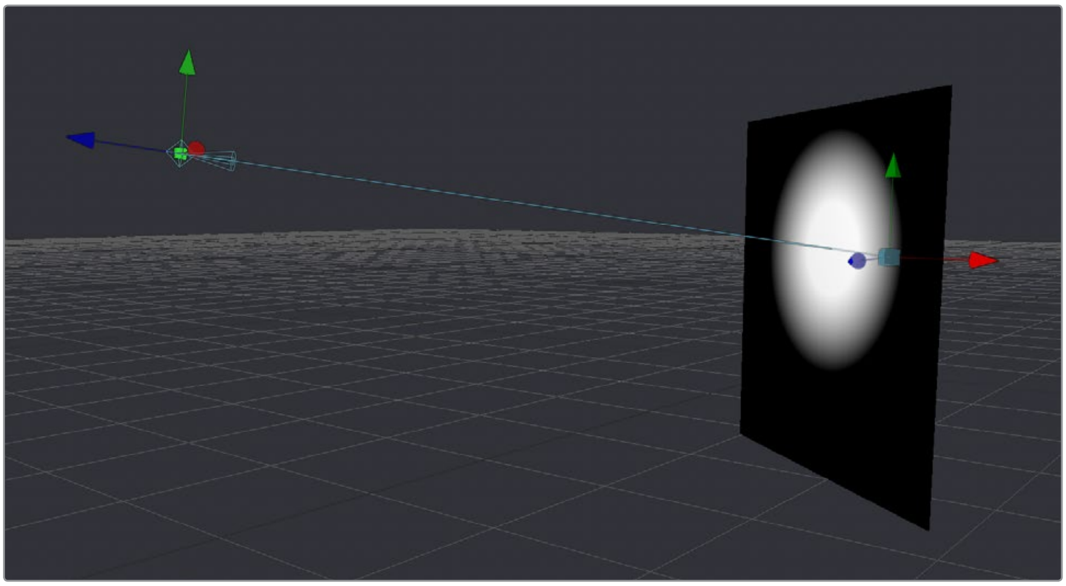

A light made to face the wall using its enabled target control

Parenting

One of the many advantages of the node based approach to 3D compositing is that parenting between objects becomes implicit in the structure of a 3D node tree. The basis for all parenting is the Merge3D node. If you're careful about how you connect the different 3D objects you create for your scene, you can use multiple Merge3D nodes to control which combinations of objects are transformed and animated together, and which are transformed and animated separately.

For example, picture a scene with two spheres that are both connected to a Merge 3D. The Merge 3D can be used to rotate one sphere around the other, like the moon around the earth. Then the Merge 3D can be connected to another Merge 3D to create the earth and the moon orbiting around the sun.

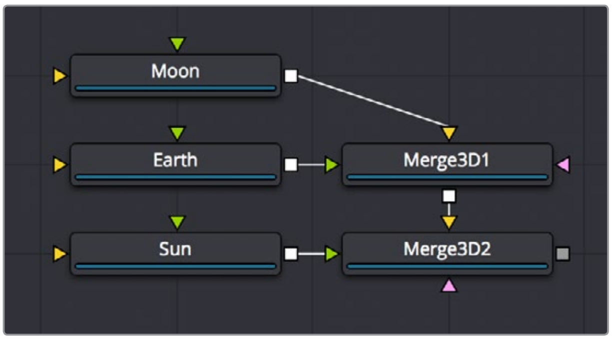

One Merge 3D with two spheres parented to another Merge 3D and parenting using three connected spheres

Here are the two simple rules of transforming parented Merge3D nodes:

- Transforms and animation applied to a Merge 3D 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 of these tasks, the Fusion page 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.

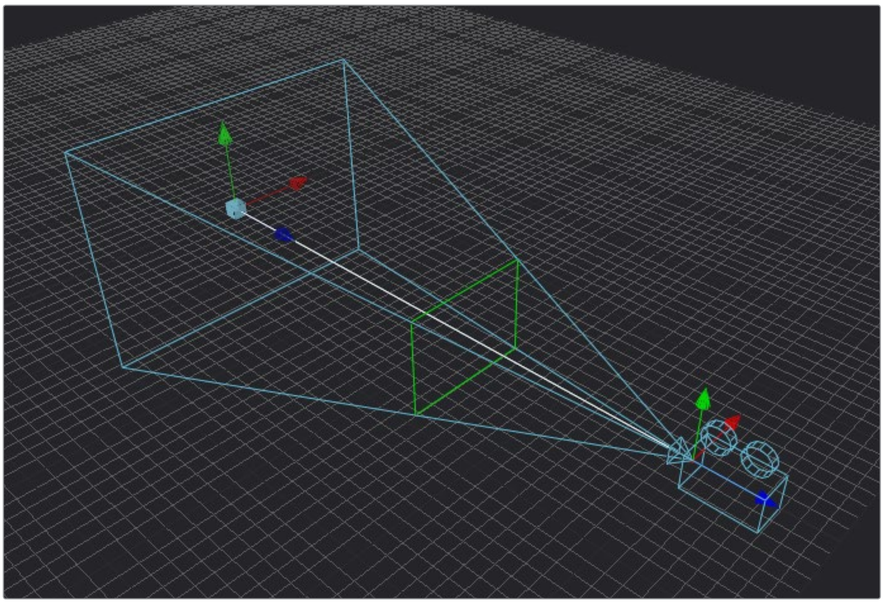

A camera displayed with on-screen Transform controls in the Viewer; the Focal Plane indicator is enabled in green

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 Merge 3D.
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. To learn more about Guides, see Chapter 52, "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.

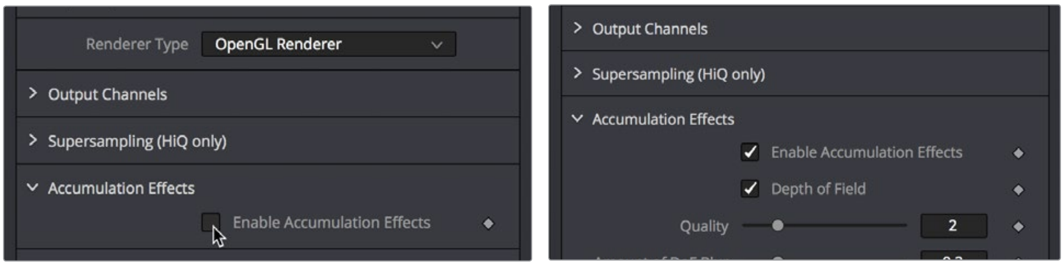

Turning on Enable Accumulation Effects enables additional depth of field controls

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 only affect 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 checkbox appears in the Control Visibility group, "Focal Plane." 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.

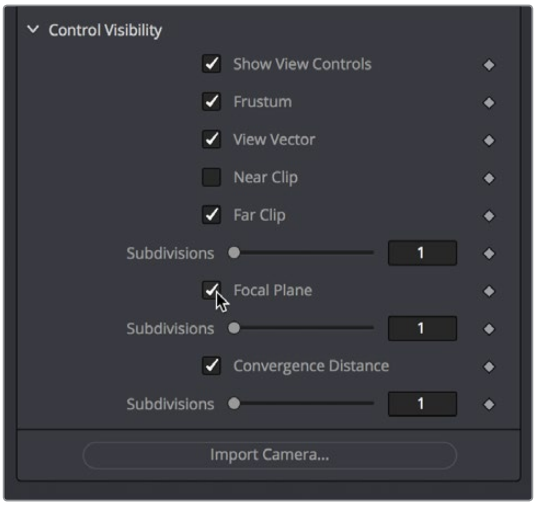

Turning on the Focal Plane checkbox in the Camera3D node

For more information about these specific camera controls, see the Camera3D section 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. The Fusion page 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, the Fusion page will sample each frame of the motion, regardless of the format.

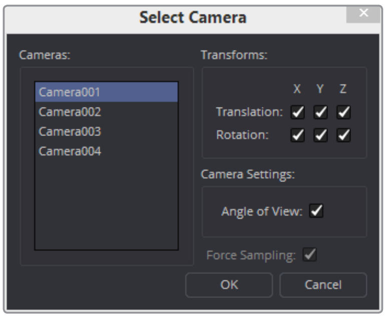

The Import Camera Control dialog

TIP: When importing parented or rigged cameras, baking the camera animation in the 3D application before importing it into the Fusion page often produces the 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 one or both on.

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 which are used to choose how each object is affected by light:

- Merge3D nodes have a “Pass Through Lights” checkbox that determines whether or not 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 the Fusion page.

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.

Spot Light

A Spot light is an advanced point light that produces a well defined cone of light with falloff. This is the only light that produces shadows.

(Left to Right) Directional Light, Point Light, and Spot Light

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.

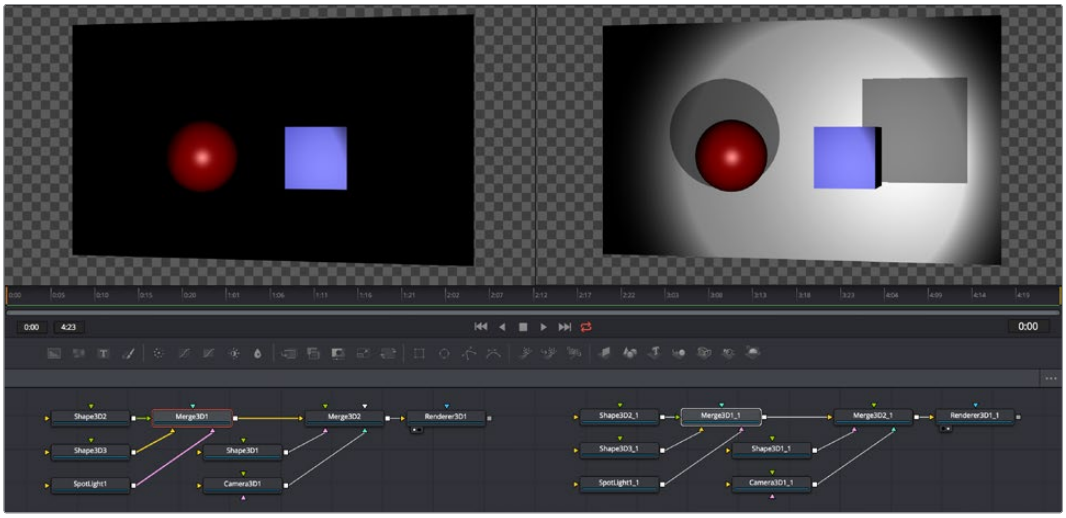

(Left) Pass Through Lights is disabled so only the front two shapes are illuminated, (Right) Pass Through Lights is enabled, so all shapes connected to both Merge3D nodes are illuminated

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 Renderer 3D 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.

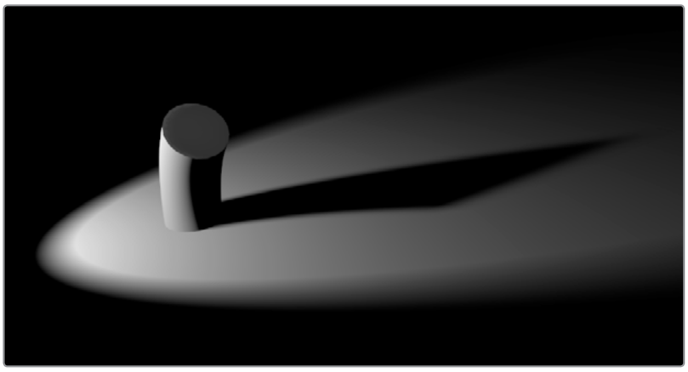

An image with spotlight casting a variable soft shadow

See the spotlight section of Chapter 67, "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 adjusting the falloff of the shadow will appear, as well as sliders for the minimum and maximum softness.

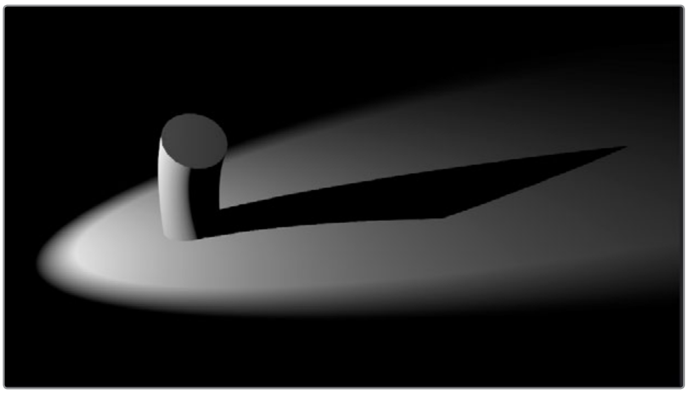

Hard shadow casted by a spotlight

Selecting the Variable option reveals the Spread, Min Softness and Filter Size sliders. A side effect of the method used to produce variable softness shadows is that the size of the blur applied to the shadow map can become effectively infinite as the shadow's distance from the geometry increases. These controls are used to limit the shadow map by clipping the softness calculation to a reasonable limit.

The filter size determines where this limit is applied. Increasing the filter size increases the maximum possible softness of the shadow. Making this smaller can reduce render times but may also limit the softness of the shadow or potentially even clip it. The value is a percentage of the shadow map size.

For more information, see Spotlight in Chapter 67, "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 renders over the top of the shadow because they effectively exist in the same exact location in 3D space.

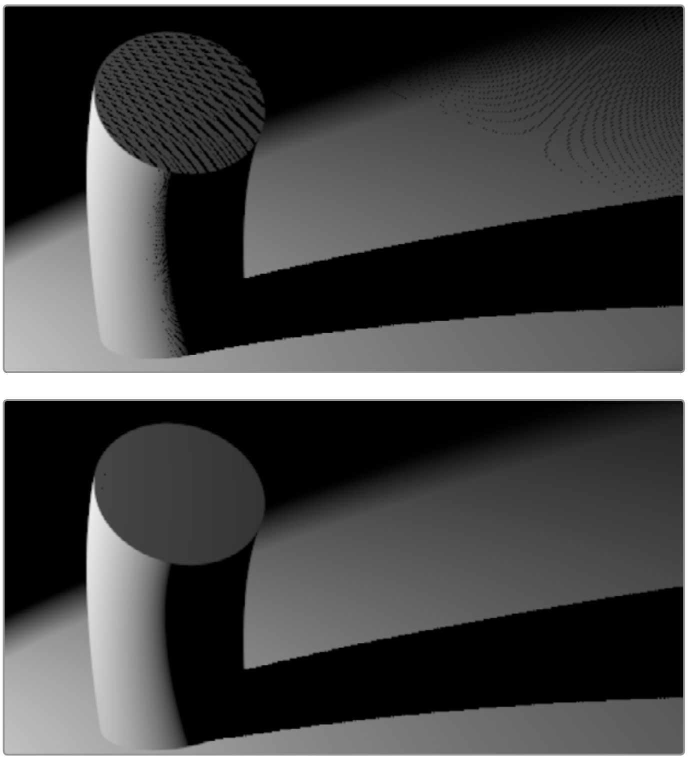

(Top) Results of shadow map Z-fighting, (Bottom) Corrected shadow shown using Biasing

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, 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

In order 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. The Fusion page has 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 effect. More complex textures like bump maps, 3D textures and reflection maps are also available and located 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.

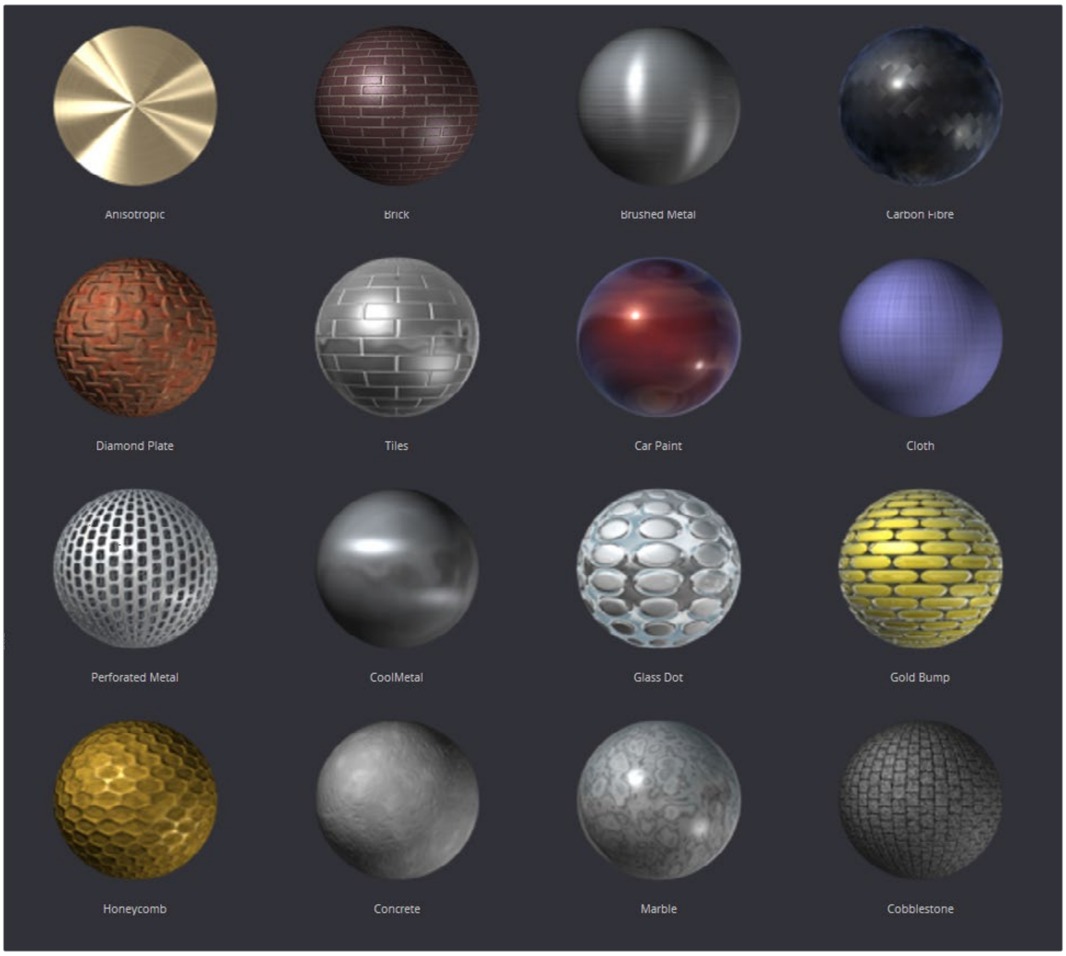

Material examples from the bin

Material Components

All of 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

This 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

This 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.

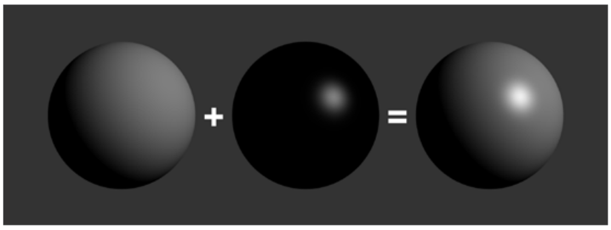

3 spheres L to R: diffuse only, specular only, and combined

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.

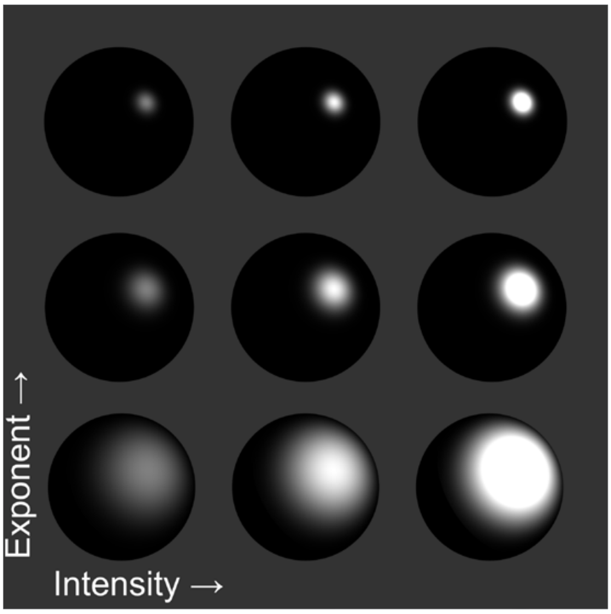

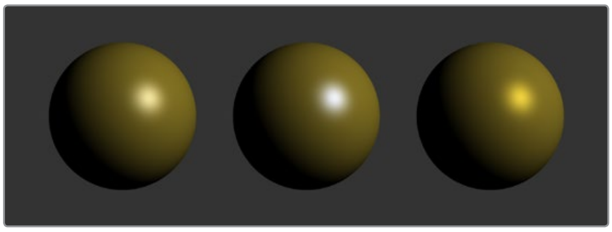

L to R: white, complimentary, and matching specular colors

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 to 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.

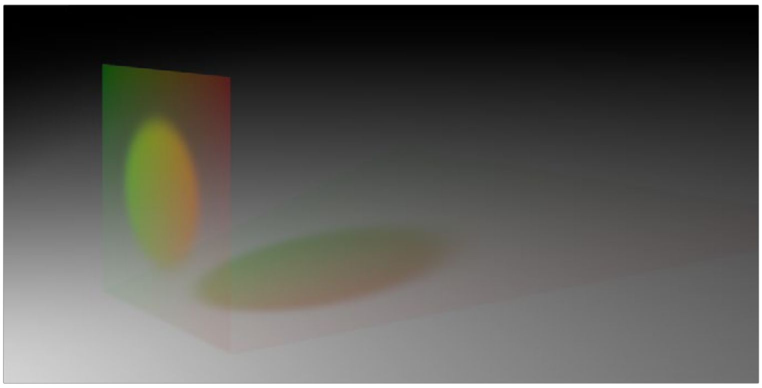

Alpha Detail set to 1, the alpha channel determines the density of the shadow

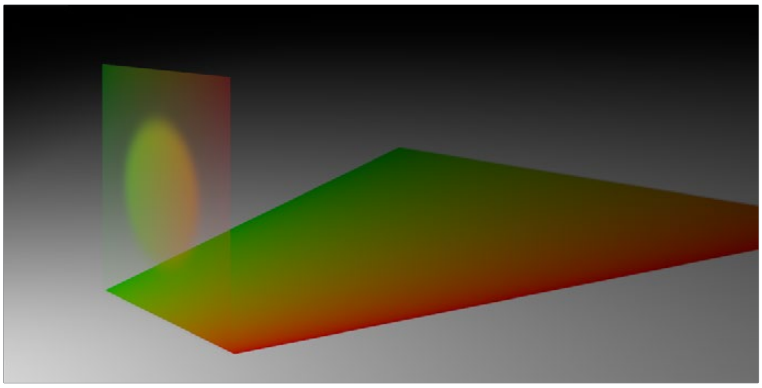

Alpha Detail set to 0, a dense colored shadow results

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 of 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.

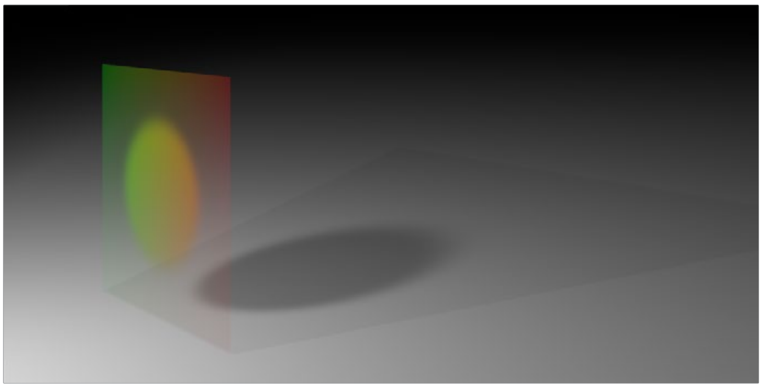

Color Detail set to 0, no color is visible in the shadow

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.

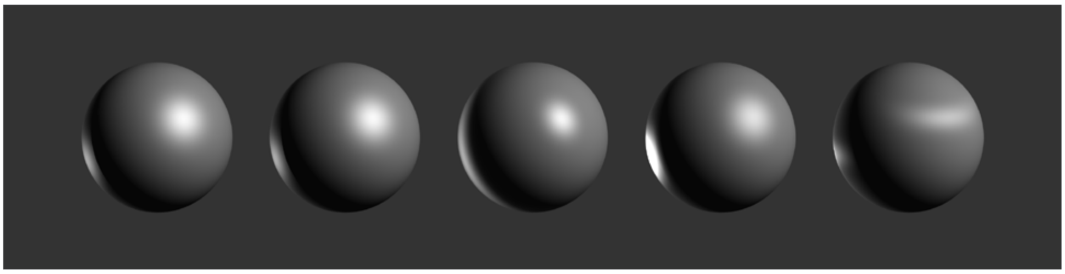

Illumination models left to right: Standard, Blinn, Phong, Cook-Torrance, and Ward

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.

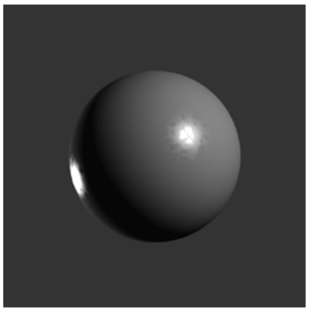

The Fast Noise texture used to control the roughness of a Cook-Torrance material

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.

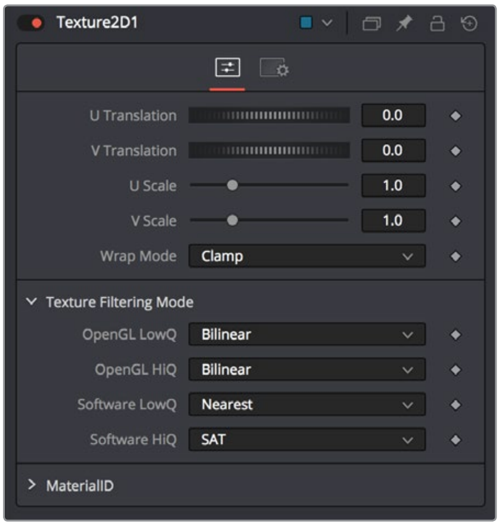

The Texture2D node is used to translate a texture in the UV space of the object, as well as set the filtering and wrap mode

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.

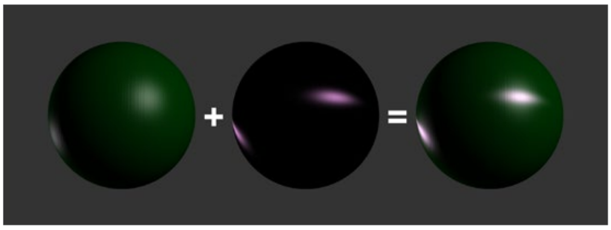

Combining an anisotropic highlight with a Blinn material using the Channel Boolean material

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.

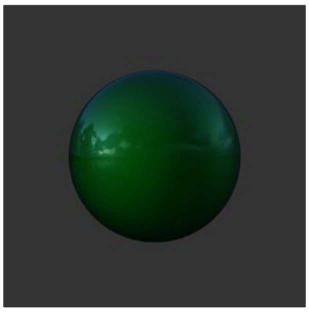

A Blinn material connected to a Background Material input of the Reflect. This causes the reflection to be added to the Blinn output.

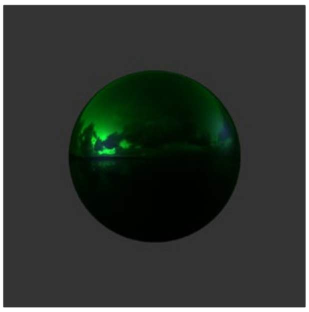

A Reflect is connected to the Diffuse Color component of the Blinn, causing the reflection to be multiplied by the diffuse color and modulated by the lighting.

Refraction

Refraction only occurs 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.

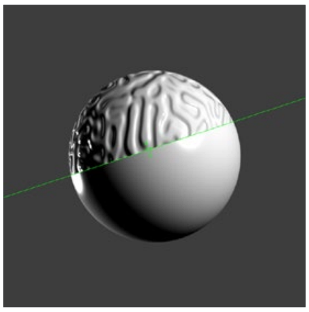

Split screen of a sphere; half with bump map, half without

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.

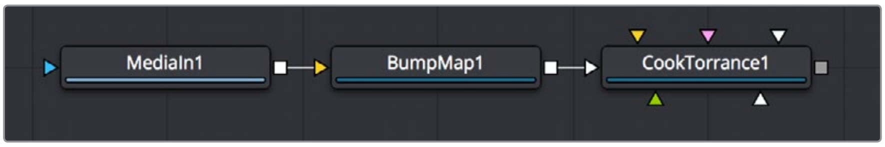

Image connected to a BumpMap connected to a CookTorrance material node

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 only uses 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 the Fusion page'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. The Fusion page prevents you from doing this, however, because the Fusion page uses a different coordinate system for doing the lighting calculation. You first must use a BumpMap which 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. The Fusion page cannot detect which type 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 the Fusion page, 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 the Fusion page.

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 if 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.

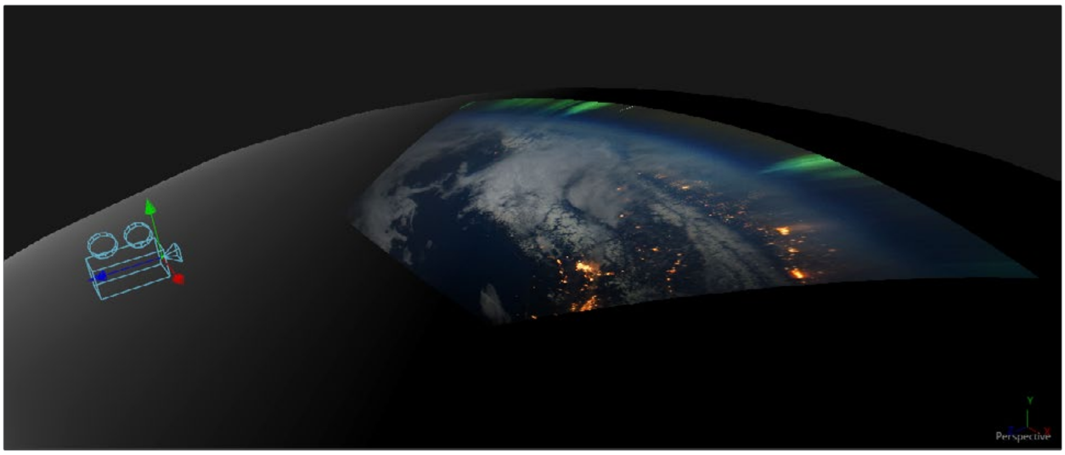

Camera node used for a projection map

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.

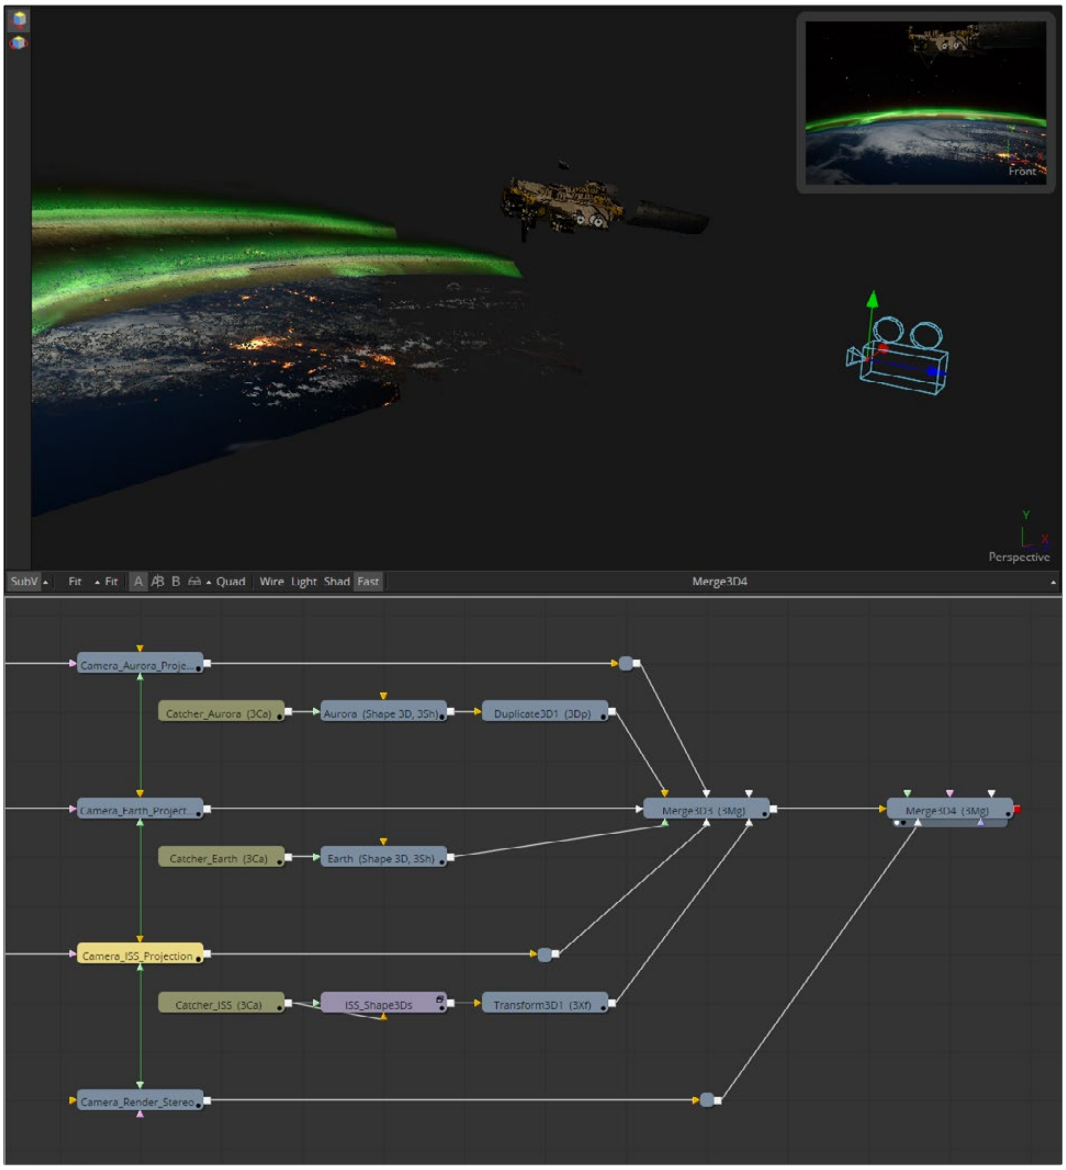

Camera projection used with a Catcher node (example from an older version of Fusion)

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; these 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, 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 Merge 3D, 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 MediaIn1 node) is projected into 3D space by mapping it onto five planes (Shape3D nodes renamed ground, LeftWall, RightWall, and so on), which are positioned as necessary within a Merge3D node to apply reflections onto a 3D car to be composited into that scene.

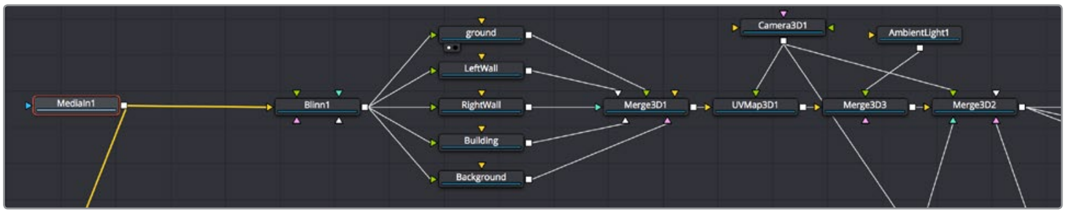

Excerpt of a composition that's projecting an image of a street scene into 3D space

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, 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.

(Left) Five planes positioning a street scene in 3D space in preparation for UV Projection, (Right) The UV Map node being used to project these planes so they appear as through a camera in the scene

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 the Fusion page. 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, “Fusion Page Effects,” 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 Renderer 3D 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.

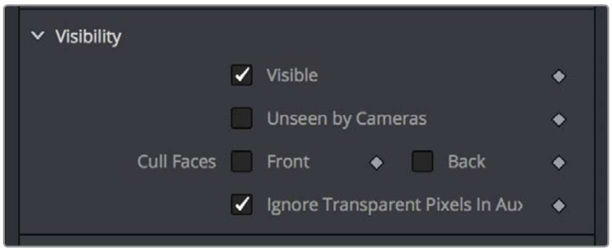

A 3D geometry node’s visibility parameters

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 Renderer 3D. 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 Renderer 3D. 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 Renderer 3D 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 of the aux channels will be filled for all of 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. The Fusion page's FBX format 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, DaVinci Resolve imports it to create a composition with the same lights, cameras, materials and geometry found in an FBX file.

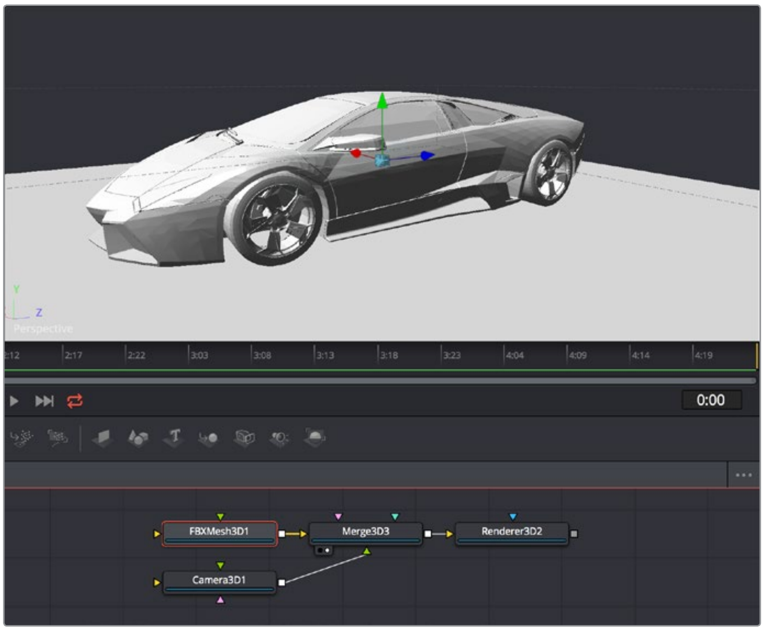

An imported model, via the FBXMesh3D node

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 the Fusion page 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 the Fusion page's 3D scene.

FBX Exporter

You can export a 3D scene from the Fusion page 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 the Fusion page. It's a powerful node filled with enough controls to create nearly any text effect you might need, all in three dimensions. This section seeks to get you started quickly with what the Text3D node is capable of. For more information, see Chapter 66, "3D Nodes" in Part 7, Fusion Page Effects.

Assembling Text Objects

Each Text3D node is a self-contained scene within which each character of text is an individual object. Because of this, the ideal way to combine numerous text objects that you might want to animate or style independently from one another is to connect as many Text3D objects as you want to be able to independently animate or style to one or more Merge3D nodes.

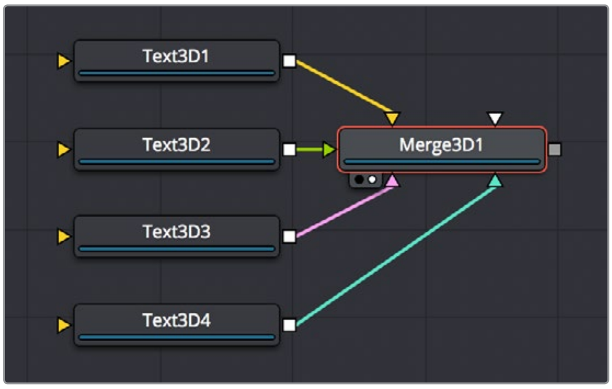

Merging multiple text objects to create an intricately styled scene

TIP: If you click the Text icon in the toolbar to create a Text3D node, and then you click it again while the Text3D node you just created is selected, a Merge3D node is automatically created and selected to connect the two. If you keep clicking the Text icon, more Text3D nodes will be added to the same selected Merge3D node.

Entering Text

When you select a Text3D node and open the Inspector, the Text tab shows a “Styled Text” text entry field at the very top into which you can type the text you want to appear onscreen. Below, a set of overall styling parameters are available to set the Font, Color, Size, Tracking, and so on. All styling you do in this tab affects the entire set of text at once, which is why you need multiple text objects if you’re going to want differently styled words in the same scene.

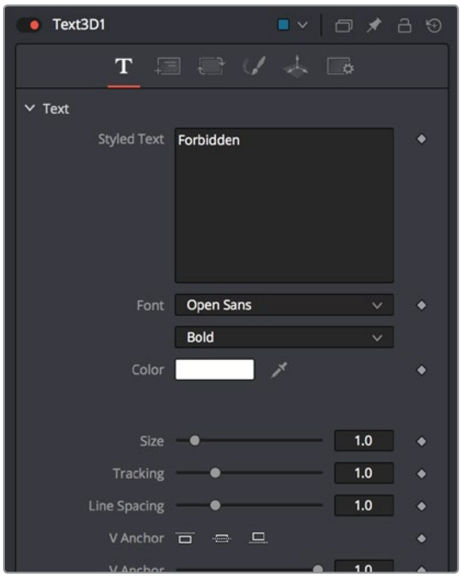

The text entry and styling parameters in the Text tab

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.

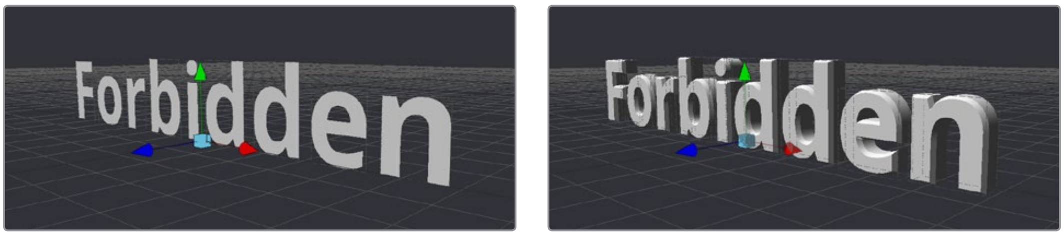

(Left) Unextruded text (Right) Extruded text

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.

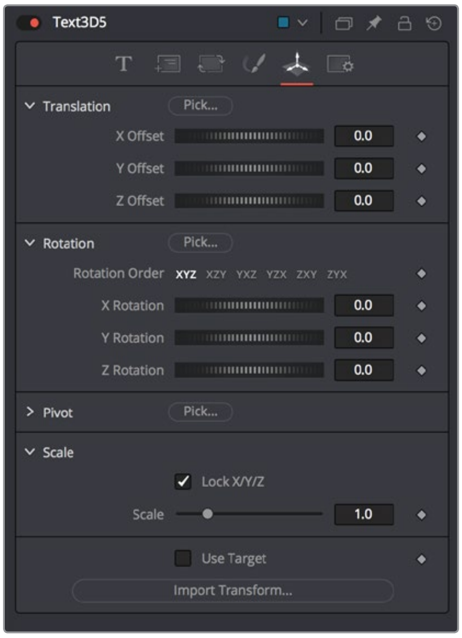

Text3D nodes also have Transform parameters built-in

Additionally, selecting a Text3D node exposes all of 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).

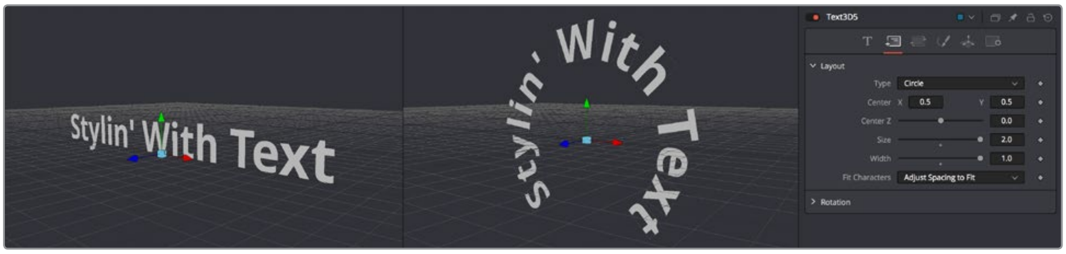

Text using two different layouts

"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 of these capabilities at once if you really need to go for it. And of course all of 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.

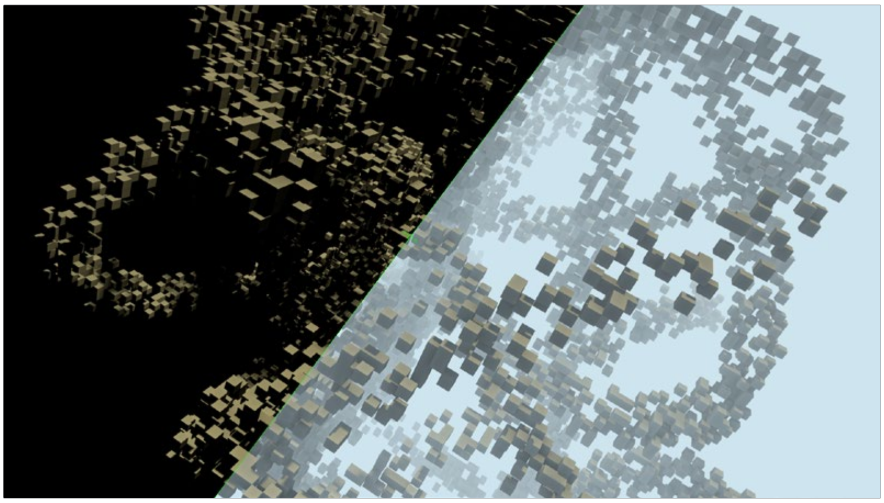

Split screen with and without fog

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 check box 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 the Fusion page 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.

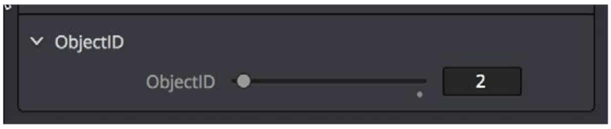

Material ID parameters in a Shape3D node's Inspector controls

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 the Fusion page 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.

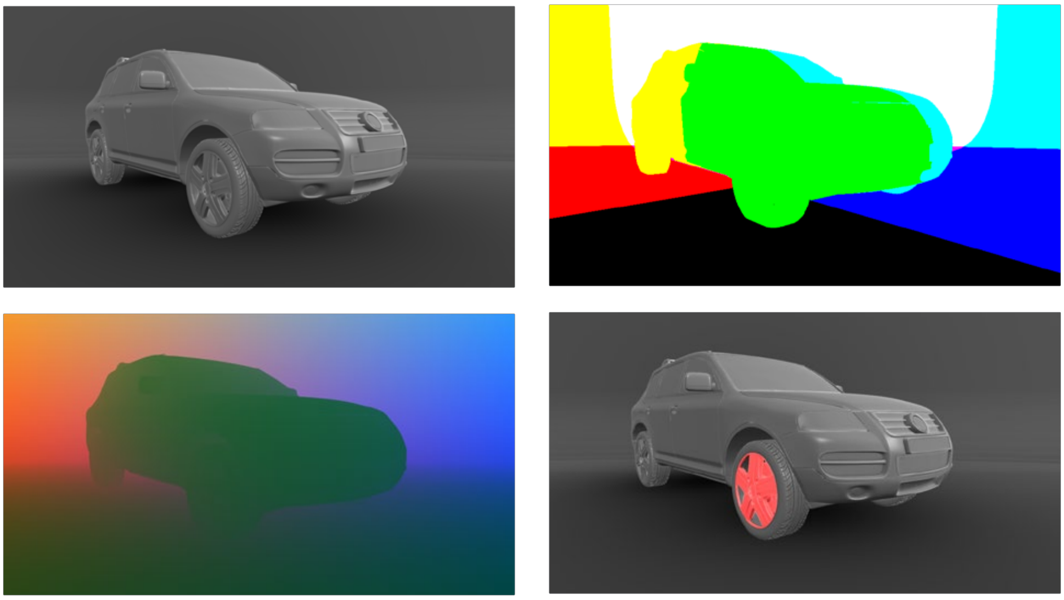

A World Position Pass Rendering of a scene with its center at (0,0,0). Actual image is on the left.

3D Scene Input

Nodes that utilize the World Position channel are located under the Position category. VolumeFog and Z to WorldPos require a camera input matching the camera that rendered the Position channels, which can either be a Camera3D or a 3D scene containing a camera. Just as in the Renderer3D, you can choose which camera to use if more than one are in the scene. The VolumeFog can render without a camera input from the Node Editor if the world space Camera Position inputs are set to the correct value. VolumeMask does not use a camera input. Nodes that support the World Position Pass, located under the Position category, offer a Scene input, which can either be 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.

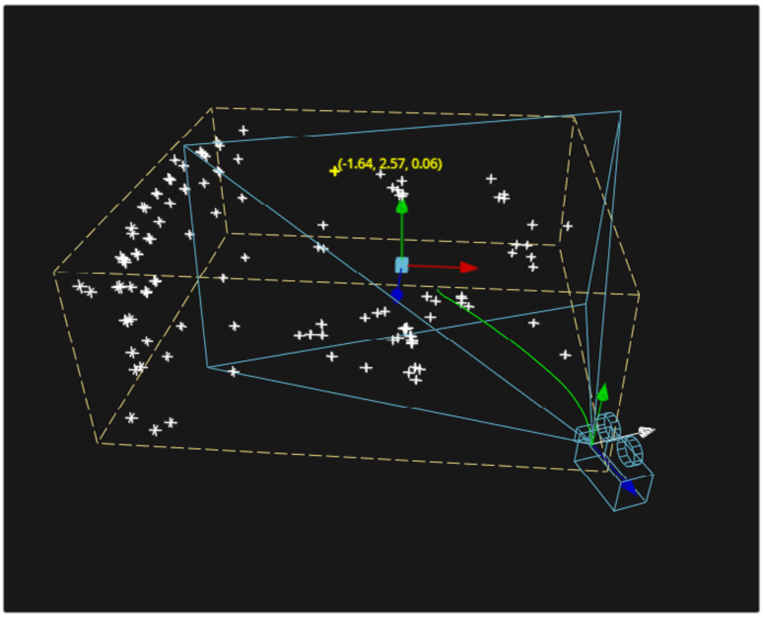

Point cloud in the Viewer

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.

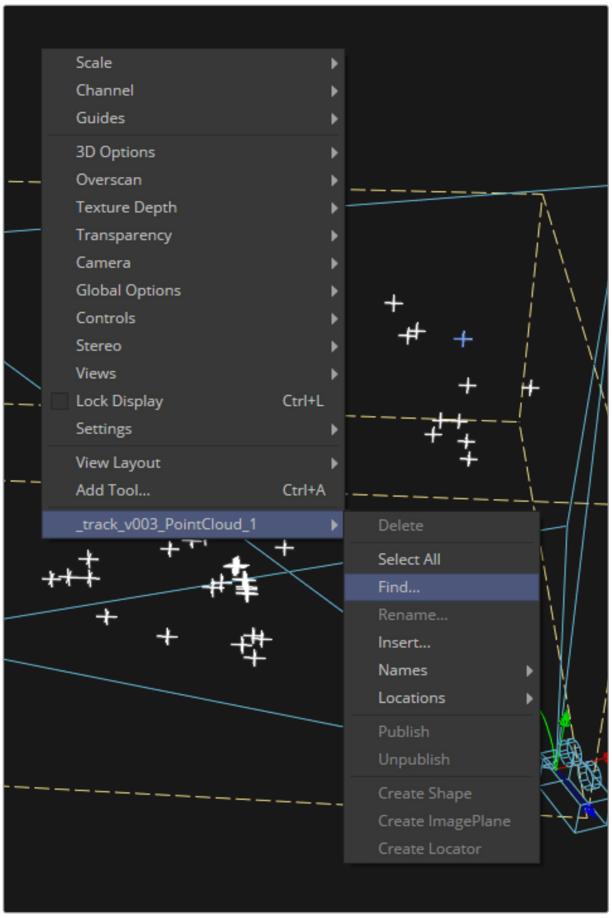

Finding a point cloud using the Viewer contextual menu.

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 only works 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 over it and choose Publish from the contextual menu

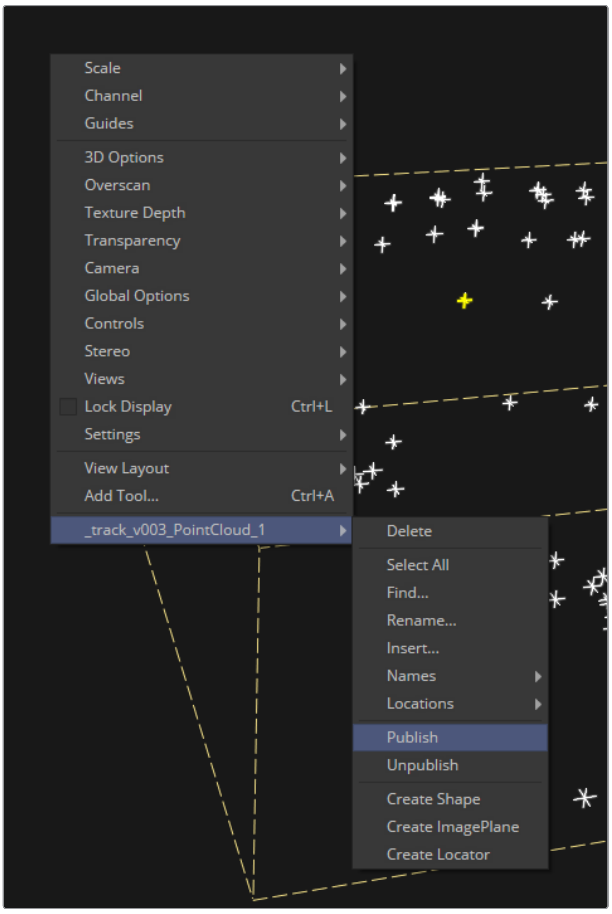

Publishing a point using the Viewer contextual menu

Chapter 61

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

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 that is 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 that are in a fixed frame of reference. Moving objects or people need to 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.

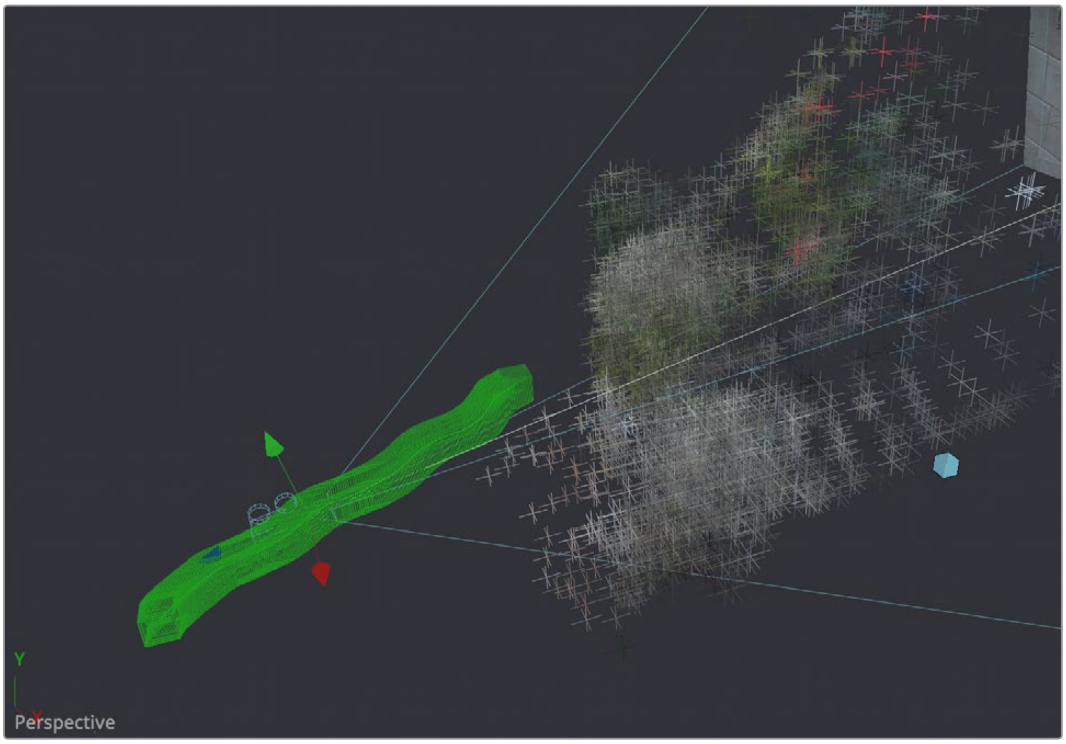

Track Data and a solved camera path

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, 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 needs to 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 the Camera Tracker can analyze. For example, suppose some footage shot with a moving camera of a subject standing in front of a green screen 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.

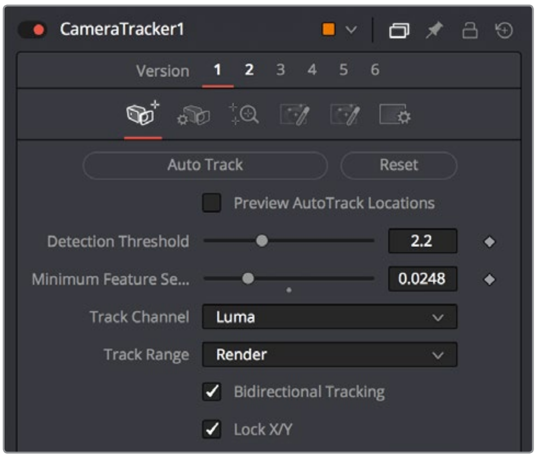

The Camera Tracker tab

Camera

This section is where the basic settings of the camera are set, like 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 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.

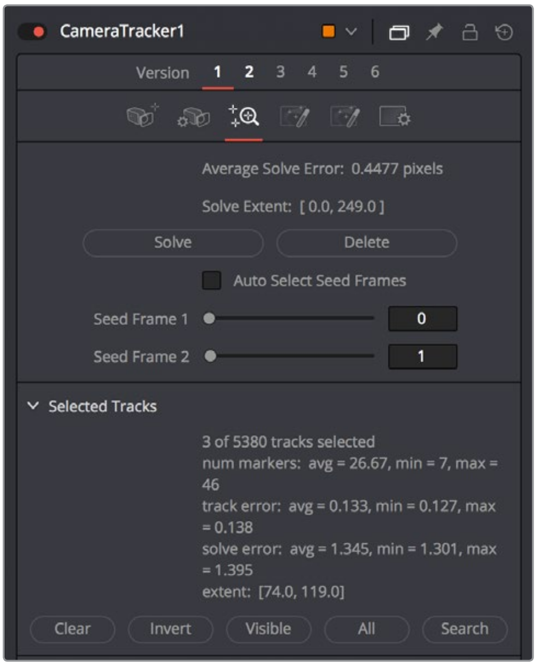

The Camera Solver tab

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 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.
- Deleter 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 which 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 useful to also:

- 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 has 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.

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.

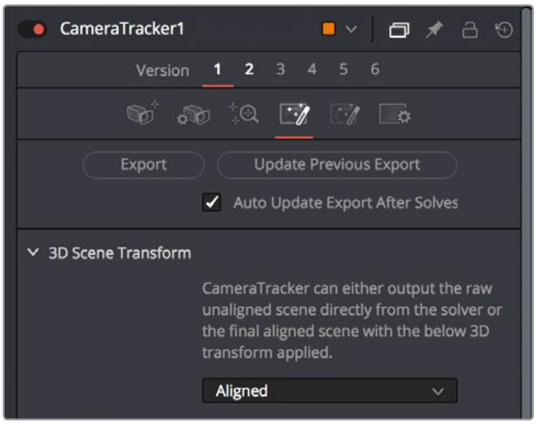

The Camera Tracker tab

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 which can be cut and paste 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 on 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 tool bar gives you access to the common features used to track and solve a shot.

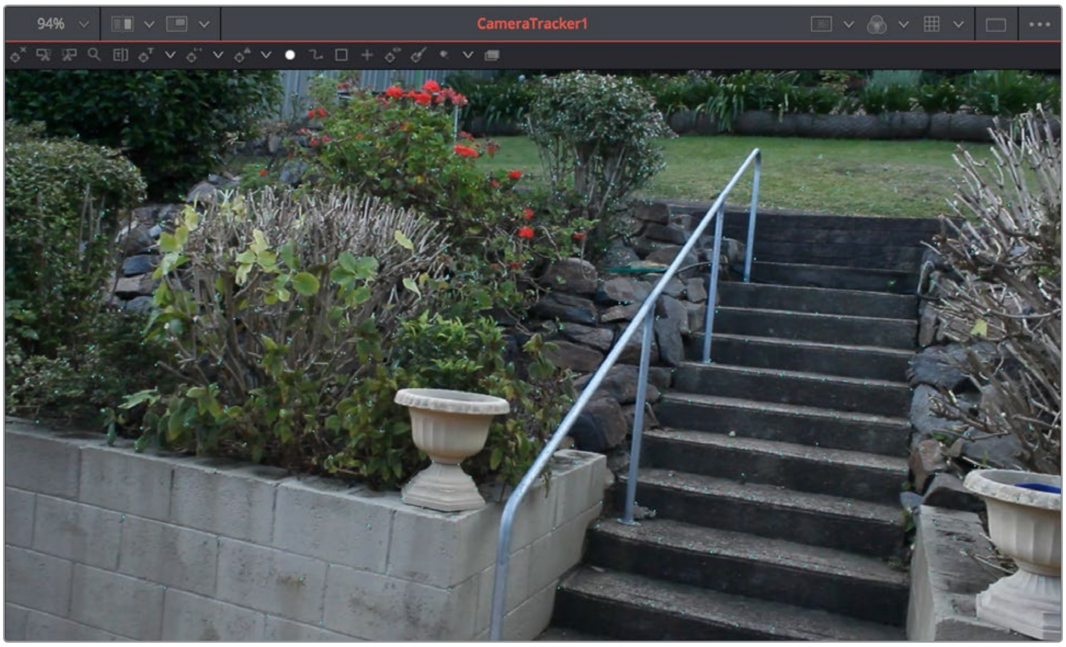

A 2D view of the camera track data

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 onscreen 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.

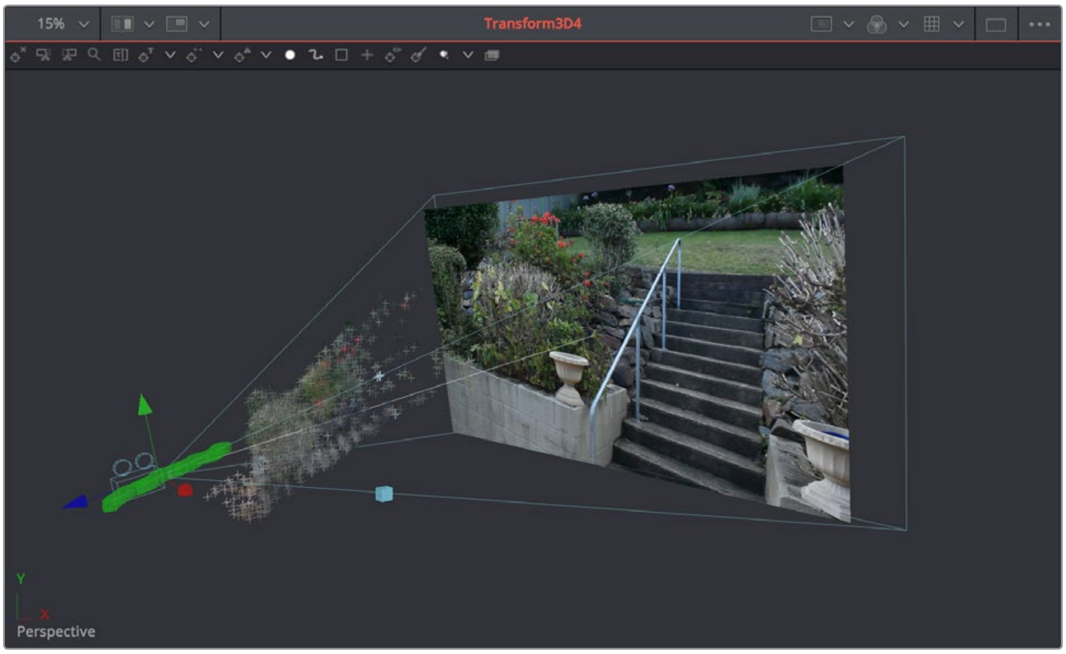

A 3D view of the camera track data

Chapter 62

Particle Systems

This chapter is designed to give you a brief introduction to the creation of fully 3D particle systems, one of the Fusion page's most powerful features. Once you understand these basics, for more information on each Particle System node that's available, see Chapter 89, "Particle Nodes."

Contents

Introduction to Particle Systems

Particle systems are computer simulations using 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 the Fusion page is a fully 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.

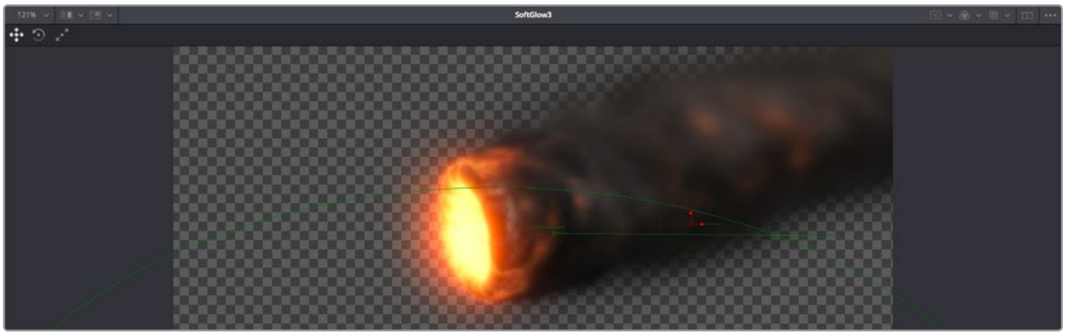

A 3D particle system, also created entirely within the Fusion page

The three most fundamental nodes required for creating particle systems are found on the Fusion page 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 to click the three particle system nodes from left to right to create a functional particle system.

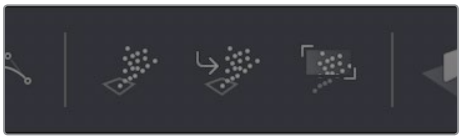

The three Particle System nodes available from the Fusion page toolbar

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.

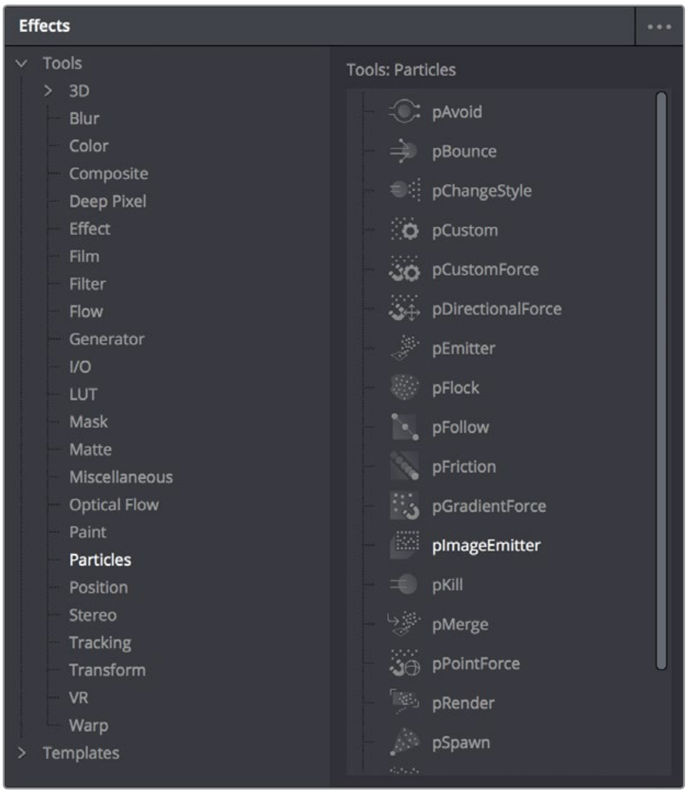

A sample of the nodes available in the Particles bin of the Effects Library

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 has 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 of 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.

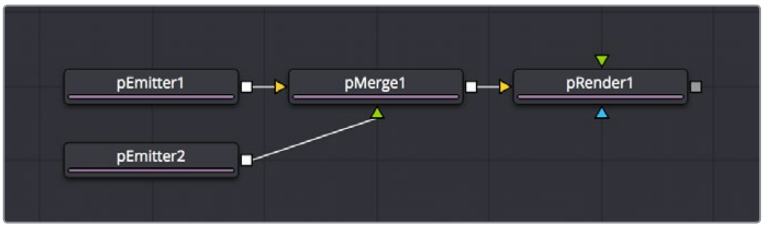

Compositing two pEmitter nodes to create a compound particle system combining two kinds of particles together

If you're trying to create particle systems with more naturalistic 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.

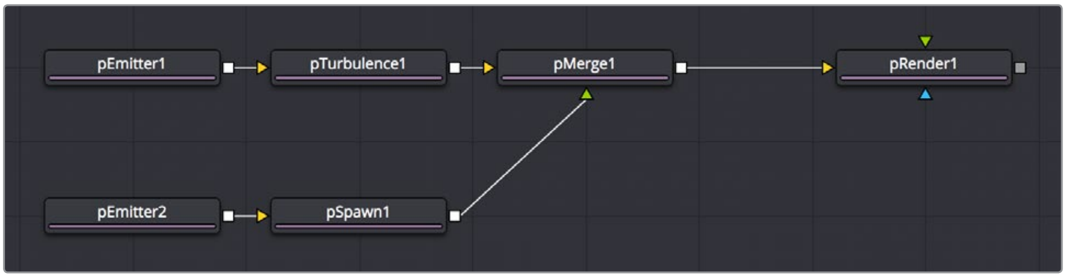

Customizing the effect of pEmitter nodes using different forces to add complexity to the particle animation

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).

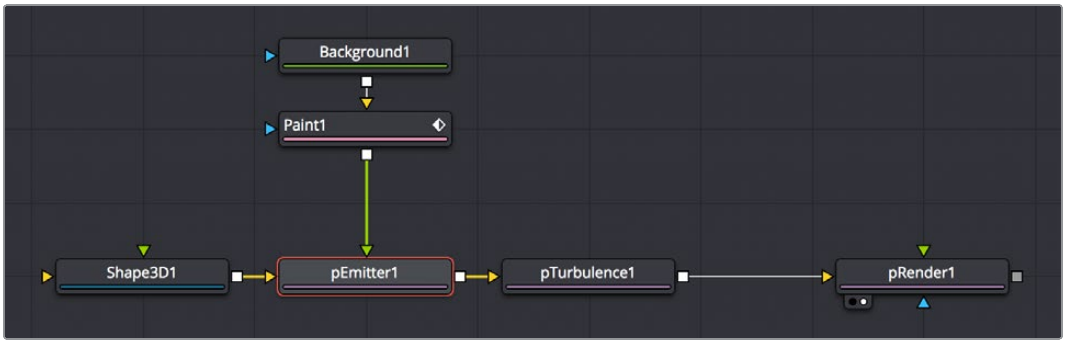

Customizing pEmitter nodes using mesh geometry to define regions and 2D images to define particle shape

The above examples assume you'll be outputting 2D renders to combine into the rest of a 2D composition. However, because particle systems in the Fusion page 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.

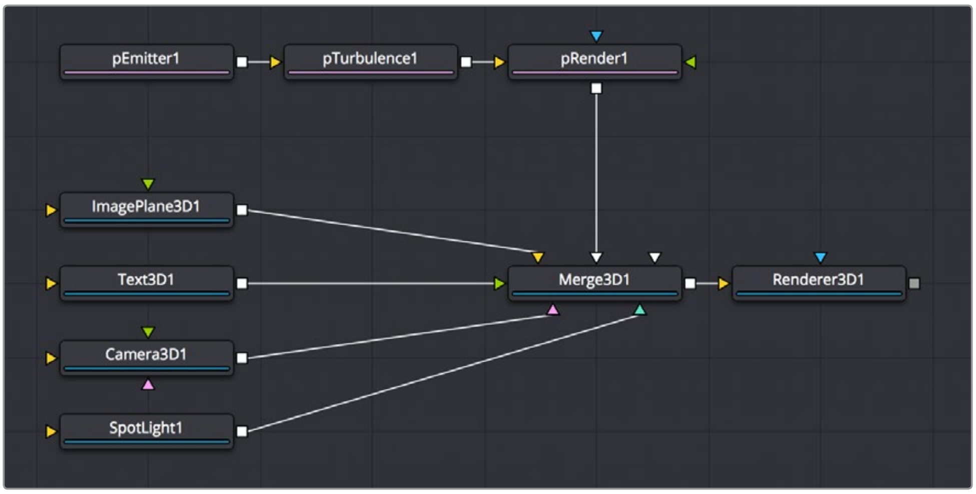

Connecting a particle system to a Merge3D node so the particles are subject to lighting and shadows within a 3D scene

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.

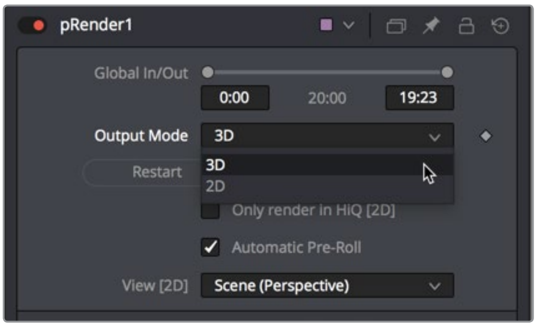

Choosing whether a particle system's output is 2D or 3D in the pRender node's Inspector controls

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.

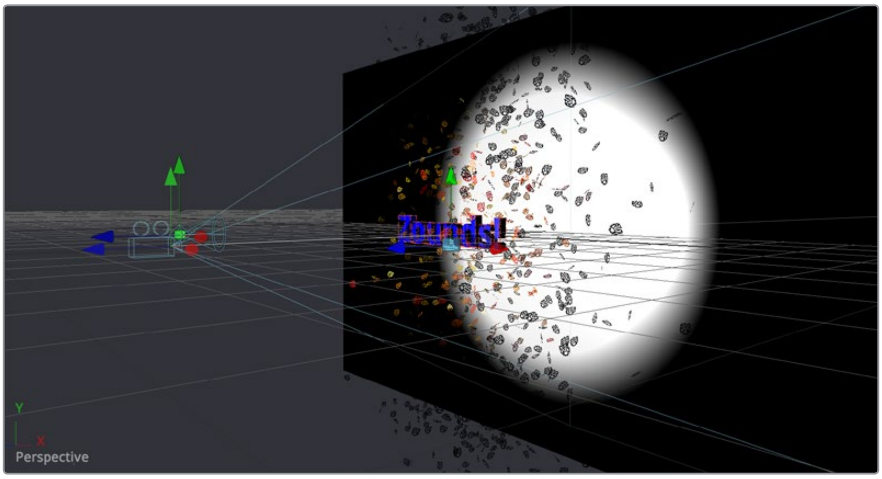

The result of using a particle system within a 3D scene

Particle System Distribution

To adjust the distribution of particles being emitted, select the pEmitter node to expose its controls into 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 of these controls can be animated.

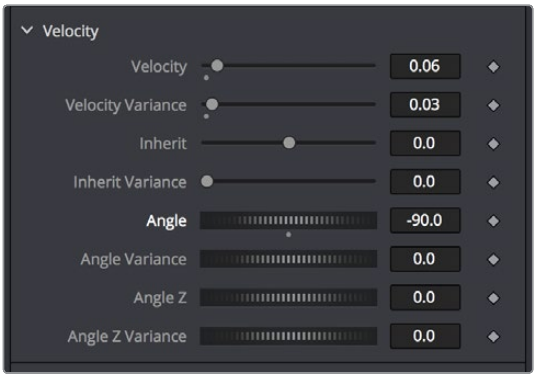

A pEmitter node's Velocity Angle and Angle Variance controls let you adjust the direction and width of particle distribution

Particle systems can be positioned and rotated by loading the pEmitter nodes that generate particles into a viewer and using the onscreen 3D position and Rotation controls provided to move the particle system around.

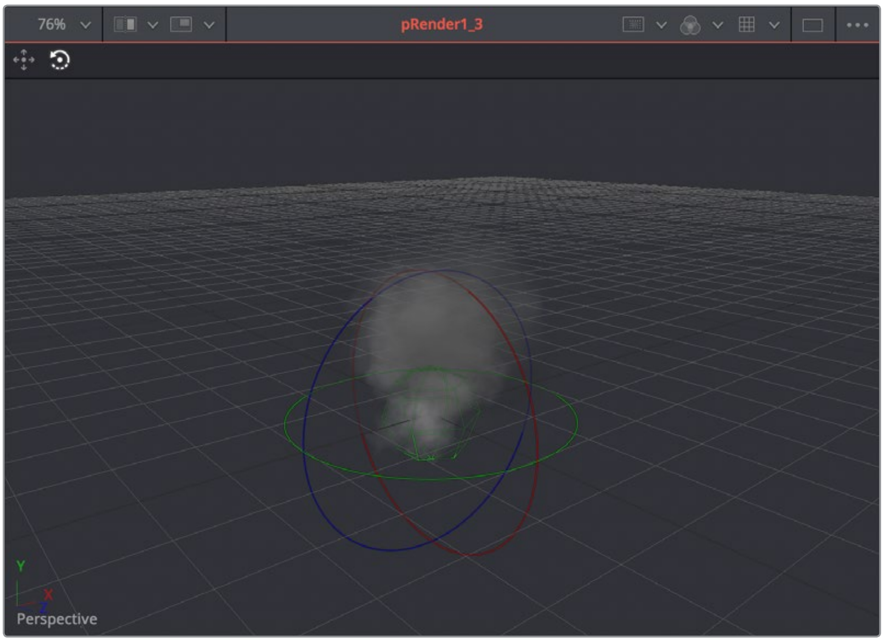

A pEmitter node loaded into the Viewer with the rotation onscreen controls enabled

Alternately, you can use the controls of the pEmitter's Region tab in the Inspector to adjust translation, rotation, and pivot. All of 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 that are 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 be creating 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 tabs worth of parameters:

- **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 naturalistic 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. They modify behavior such as velocity, spin, quantity and lifespan instead.
- **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 let 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 particle, since even geometry which 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 Particles bin of the Templates category of the Effects Library has ten different examples of particle systems creating a variety of effects. One of the best way of learning how to create and customize particle systems is to open these and investigate how they're made.

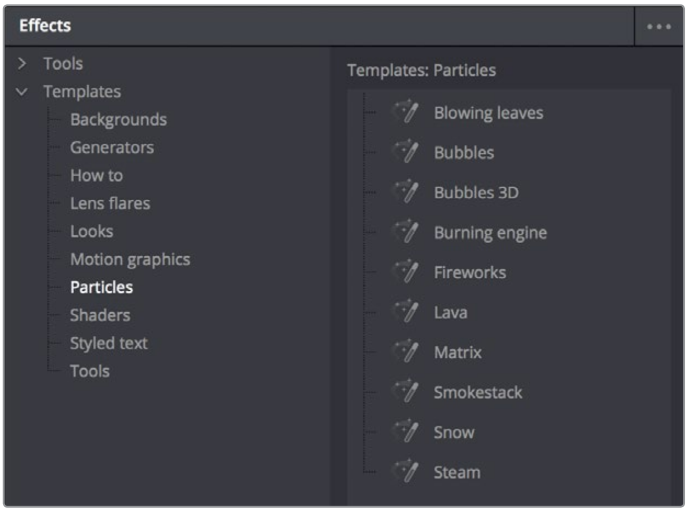

Different particle system presets in the Particles bin of the Templates category

Simply drag and drop any of the particle systems you see here into the Node Editor, load the last node into the Viewer, and you'll see how things are put together.

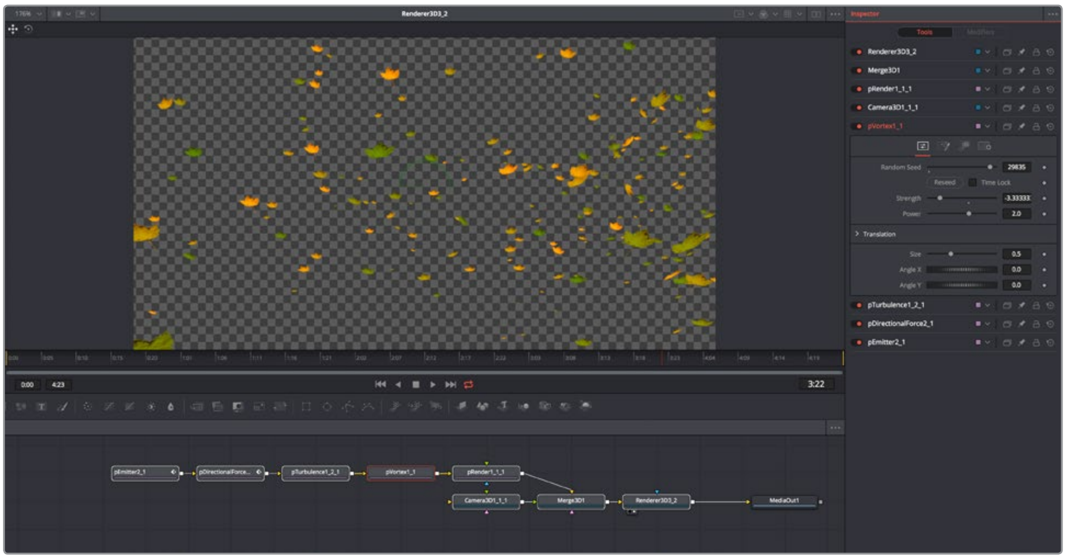

The Blowing Leaves preset from the Particles bin of the Templates category

Chapter 63

Optical Flow and Stereoscopic

This chapter covers the numerous stereoscopic and optical-flow based nodes available in Fusion and their related workflows.

Contents

Overview

Resolve Studio includes 3D stereoscopic and optical flow-based nodes, which partially overlap in functionality.

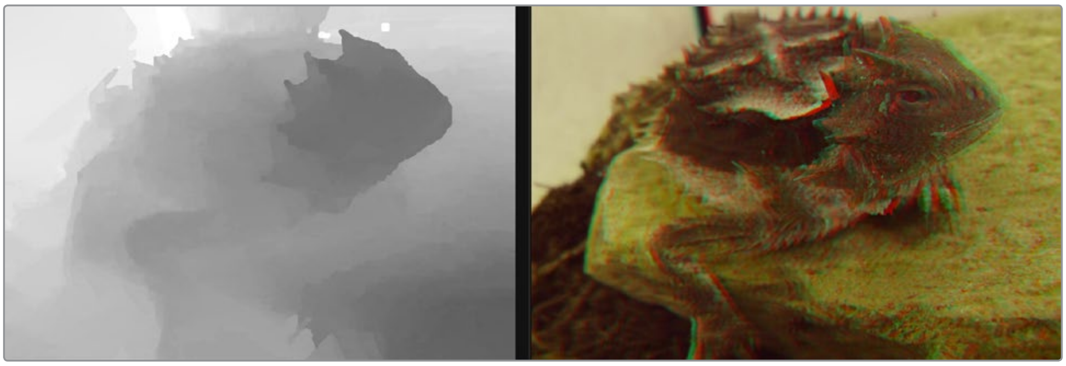

Stereoscopic comp displayed in the Viewers

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 to create smooth slow motion and variable retiming of clips, repair missing frames and even correct disparity in stereo 3D clips.

The OpticalFlow node is used to analyze a clip and generate motion vector fields. These motion vectors are then used by additional nodes to apply functions to images.

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
- 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 combing 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, eye separation
- **Stereo > ZToDisparity:** converts Z-depth to disparity

Working with Aux Deep Channels

Fusion images 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 Resolve Studio:

- **RGBA:** These are the standard colors
- **Z:** The eyespace Z coordinate is almost always negative since 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 that is 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

These extra channels are used by some 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. For example:

- The Renderer3D node is capable of generating most of these channels.
- The OpticalFlow node generates the Vector and BackVector channels, 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 node. 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 information on TimeSpeed or TimeStretcher see the Nodes Reference guide.

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 information on Tween or Repair Frame see the Nodes Reference guide.

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.

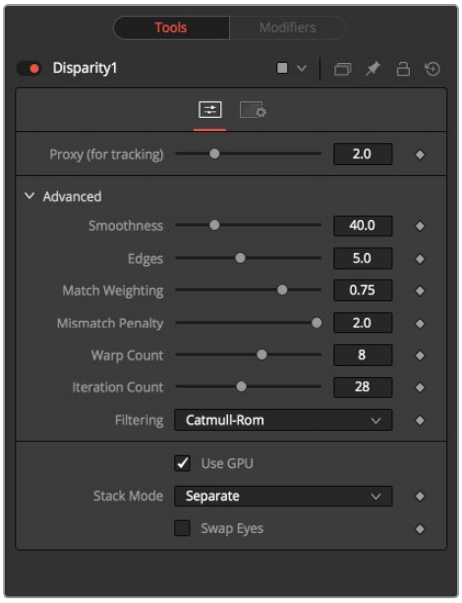

The Disparity node analyzes a stereo pair of images and generates an X&Y disparity map

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.

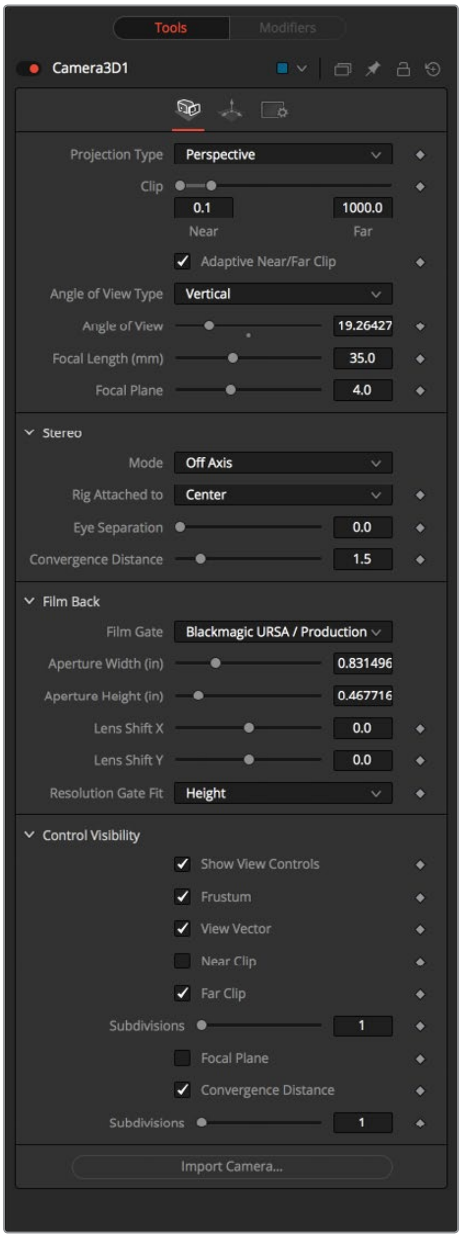

Stereoscopic cameras can be done with a single camera or two connected cameras

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.

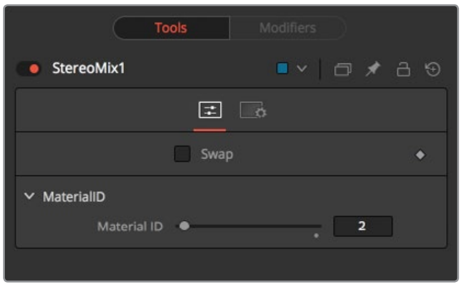

Material Viewer showing stereoscopic material

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

These 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.

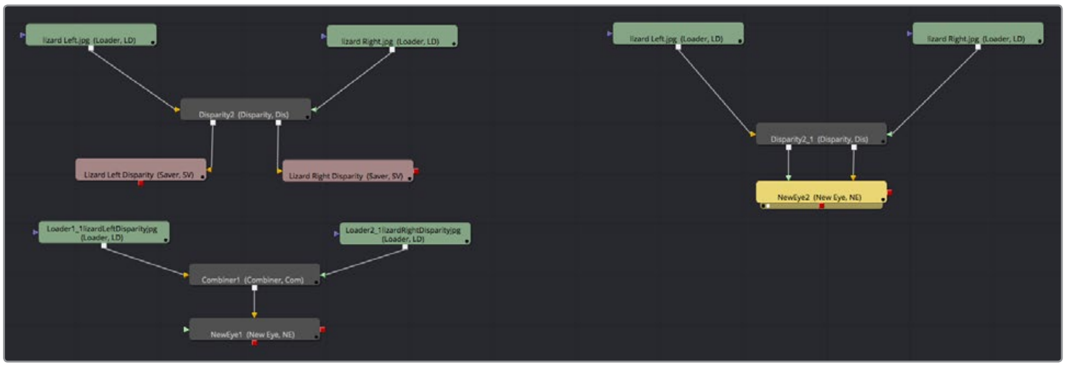

Two stereoscopic workflows (shown in a previous version of Fusion)

In the above example the workflow on the right takes the left and right eye, generates the disparity, then NewEye is used to generate a new eye for the image right away.

The example on the left renders the frames with disparity to intermediate EXR images. These images are then loaded back into Stereo nodes and used to create the NewEye images.

By using Render nodes to compute the disparity first, the later processing of the creative operations can be a much faster and interactive experience.

Although not shown in the above diagram, it is usually a good idea to color correct the right eye to be similar to the left eye before disparity generation as this helps with the disparity-tracking algorithm. The color matching does not need to be perfect, for example, it can be accomplished using the “Match” option in a Color Corrector’s histogram options.

About the Disparity Channel

The Disparity channel stores the displacement vectors that match pixels in one eye to the other eye. The left image’s Disparity channel will contain vectors that map left>right and the right image’s Disparity channel will contain vectors that map right>left. For example:

```
(xleft, yleft) + (Dleft. x, Dleft. y) -> (xright, yright)
(xright, yright) + (Dright. x, Dright. y) -> (xleft, yleft)
```

You would expect for non-occluded pixels that Dleft = -Dright, although, due to the disparity generation algorithm, this is only an approximate equality.

> **NOTE:** Disparity stores both X and Y values because rarely are left/right images perfectly registered in Y, even when taken through a carefully setup 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 (image. Width/image. OriginalWidth, image. Height/ image. OriginalHeight).

Viewing of Disparity and Vector Channels

Aux channels can be displayed directly in the Viewers through the Channel viewer button's menu.

The CopyAux node is used to copy those channels directly into the RGB channels for viewing or further processing. The advantage of using the CopyAux node is that it does static normalization, which reduces a lot of flicker that the Viewer's time-variant normalization causes. When viewing long sequences of aux channels, the CopyAux node has the option to kill off aux channels and keep only the current RGB channels, freeing up valuable memory so you can cache more frames.

> **TIP:** Although you can use the Channel Booleans to copy any aux channel into RGBA, it involves a few additional clicks when compared to CopyAux.

One thing to be aware of is that aux channels tend to consume a lot of memory. A 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 take a 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 or not 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 since they can't figure out the edges are actually two separate objects.

Where to Calculate Disparity and Optical Flow?

Where you chose 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 after 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 have to 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 of the channels except the scalar channels Z/ID are technically "vector" channels. A frames Vector aux channel will store the flow forward from the current frame to the next frame in the sequence and the BackVector aux channel will store the flow backward from the current frame to the previous frame. If either the previous or next frames do not exist (either not on disk or the global range of a loader does not allow OpticalFlow to access them), Fusion will fill the corresponding channels with zeros (transparent black).

The Disparity channel stores the displacement vectors that match pixels in one eye to the other eye. The left image's Disparity channel will contain vectors that map left > right and the right image's Disparity channel will contain vectors that map right > left.

For example:

```
(xleft, yleft) + (Dleft. x, Dleft. y) -> (xright, yright) (xright,
yright) + (Dright. x, Dright. y) -> (xleft, yleft)
```

You would expect for non-occluded pixels that Dleft = -Dright, although due to the disparity generation algorithm, this is only an approximate equality. Note that Disparity stores both X and Y values because rarely are left/right images perfectly registered in Y, even when taken through a carefully setup camera rig.

Disparity and Optical Flow values are stored as un-normalized pixel shifts. In particular, note that this breaks from Fusion's resolution independent convention. After much consideration, this convention was chosen so the user wouldn't have to worry about rescaling the Disparity/Flow values when cropping an image or working out scale factors when importing/exporting these channels to other applications. Because the Flow and Disparity channels store things in pixel shifts, this can cause problems with Proxy and AutoProxy. The convention that Fusion follows is that, for proxied images, these channels store unscaled pixel shifts valid for the full sized image. So if you wish to access the disparity values in a script or via a probe, you need to remember to always scale them by (image. Width/image. OriginalWidth, image. Height/image. OriginalHeight).

When using Vector and BackVector aux channels remember that all nodes expect these aux channels to be filled with the flow between sequential frames.

More precisely, if you have sequence of 3 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/RoI is not supported for all Fusion nodes.

Chapter 64

Using OFX and ResolveFX

The Fusion page supports third-party OFX that are compatible with DaVinci Resolve.

Contents

Using OFX

The Fusion page is able to use compatible OFX that are installed on your workstation. All OFX appear in the Tools category of the Effects Library, alongside all other effects that are available in the Fusion page.

Chapter 65

Fusion Settings

This section presents the different Fusion page options that are available in the Fusion Settings window.

Contents

Fusion Settings Overview

The Fusion Settings 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 the File menu at the top of the interface.

To open the Fusion Settings window:

Choose Fusion > Fusion Settings.

The Fusion Settings window is divided into a sidebar that appears on the left, which is used to select which of the settings panels appear on the right.

NOTE: In the standalone Fusion application, there are two levels of settings: Global and Composition. These don't exist in the Fusion page of DaVinci Resolve, but the explanation remains here. The Global preferences are used to set options that are 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 node tree that is open but not displayed.

Categories of Settings

The first entry in the Preferences sidebar is assigned to the Global preferences. Clicking the disclosure arrows for the Global preferences reveals the following sections.

3D View

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

Defaults

The Default preference is used to select default behavior for a variety of options, such as animation, global range, timecode display, and automatic node merging.

Flow

Use the Flow preference to set many of the same options found in the Flow 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 nodes like Background and Text+. It is also where the frame rate for playback is set.

General

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

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, node settings, disk caches, and more.

Script

The preference for Scripting 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 is used 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 used to create and edit Timeline/Spline filters and set default options for the Timeline Editor.

Tweaks

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 the Control Paneldisplay.

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.

Import

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

Preferences In Depth

The following sections describe all preferences found in this window.

3D View Preferences

The 3D View preferences contain settings over various defaults in the 3D Viewers, including grids, default ambient light set up, and Stereoscopic views.

Grid

The Grid section of the 3D View preference configures how the grid in 3D Viewers are drawn.

- **Grid Antialiasing:** Depending on the graphics drivers installed in your computer, the antialiased lines of the grid may sort incorrectly in the 3D Viewer. Disabling this checkbox will disable antialiasing of the grid lines, correcting the issue. To turn the grid off completely, use the 3D View > Right click > Grid option, rather than preferences.
- **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 farthest 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 define the defaults when stereo is turned on in the 3D Viewer.

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 set up in the 3D Viewer.

The default ambient light is used when lighting is turned on and you have not added light to the scene. The default light moves with the camera, so if the default light is set to "upper left," the default 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 nodes are added to the node tree and when controls are animated.

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 controls 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 positions by setting the Point with 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 nodes are added automatically for the most common operations of the Background nodes and Merge operations.

- **Background:** When set to None a standard Background node is used, however, the drop-down menu allows you to choose from a wide variety of available nodes, including 2D and 3D nodes to customize the operation to your workflow.
- **Merge:** When set to None, nothing happens. When set to Merge, connecting the outputs of two nodes or dragging multiple clips on the Node Editor area 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 Node Editor 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

Use this option to determine if 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 Flow contextual menu, like settings for Tile picture, the Navigator and pipe style, are found in this category.

Force

The Force section can set the default to display pictures in certain node tiles on the node tree, rather than showing plane tiles. The Active checkbox sets pictures for the actively selected node, the All checkbox enables pictures for all nodes, and the Source and Mask checkbox enables tile pictures for just Source and Mask nodes.

When All is enabled, the picture shown will either be a thumbnail of the image rendered by the node if the node has rendered, or if the Show Thumbnails option is disabled, the node's default icon is used. Concatenated transforms will also show a default icon.

- **Show Modes/Options:** Enabling this option will display icons in the node tile depicting various states, like Disk Caching or Locked.
- **Show Thumbnails:** When this checkbox is selected, node tiles set to show tile pictures will display the rendered output of the node. When the checkbox is cleared, the default icon for the node is used instead.

Options

The Options section has a number of settings that control or aid in the layout and alignment of nodes in the node tree.

- **Arrange to Grid:** This enables a new node tree's Snap to Grid option to force the node layout to align with the grid marks in the node tree.
- **Arrange to Connected:** Nodes snap to the vertical or horizontal positions of other nodes they are connected to.
- **Auto Arrange:** This option enables the Node Editor to shift the position of nodes as needed to make space when inserting new nodes 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 tree are considered to be 'orphaned' if the nodes 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 connecting pipes between nodes will be drawn over top of the node tiles.
- **Keep Tile Picture Aspect:** Enabling this option forces node tile thumbnail pictures to preserve the aspect of the original image in the thumbnail.
- **Full Tile Render Indicators:** Enabling this checkbox will cause the entire node to change color when it is rendering. This can make it easier to identify which nodes are rendering in a large composition. The coloring itself will form a progress bar to alert you to how close slower nodes are to finishing their renders.
- **Show Instance Links:** This option can be used to select whether Instance nodes will show links, displayed as green lines, between Instance nodes.
- **Navigator:** The Navigator is a small square overview of your entire composition, which you can use to quickly navigate to different parts of the node tree while you are zoomed in. The checkboxes in this section determine when the Navigator is displayed, if at all.
- **On:** The Node Navigator will always be visible.
- **Off:** The Node Navigator will always be hidden.
- **Auto:** The Node Navigator will only be visible when the Node Editor's contents exceed the currently visible Node area.
- **Pipe Style:** This drop-down menu selects which method is used to draw connections between nodes. The Direct method uses a straight line between nodes and Orthogonal uses horizontal and vertical lines.
- **Build Direction:** When auto-building or laying out a node tree, Build Direction controls whether nodes are organized horizontally or vertically.
- **Scale:** The Scale menu allows you to select the default zoom level of the Node Editor when a new node tree is created.

Pipe Grab Distance

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

Group Opacity

This slider controls the opacity of an expanded group's background in the node tree.

Frame Format

Frame Format preferences select the format settings for the interface. The settings in this preferences dialog are described in greater detail in the Frame Formats chapter of the manual.

Default Format

This drop-down menu is used to select the default resolution for nodes and media from a list of presets. This is only a default setting; Loaders and Creators may 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 node. It also sets the playback for the comp itself, as well as the frame to timecode conversion for nodes with temporal inputs.
- **Has Fields:** When this checkbox is enabled, any Creator or Loader node added to the node tree 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 type a name for 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 sufficient for almost all video and multimedia tasks. 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 node's Color Depth control is set to Default.

General

The sections contained in the General preferences affect the behavior of the Control Panel as well as some other user interface elements.

Usability

Usability has a number of Project, Node and User Interface settings that can make the application easier to work with, depending on your workflow.

- **Auto Clip Browse:** When this checkbox is enabled, the File Browser is automatically displayed when a new Loader or Saver is added to the node tree.
- **New Comp on Startup:** When checked, a new, empty node tree is created each time the application is launched.
- **Summarize Load Errors:** When loading comps that contain unknown nodes (e.g., comps that have been created on other computers with plug-ins not installed on the current machine), the missing nodes are summarized in the console rather than a dialogue being presented for every missing node.
- **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 node 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 Apple Mouse 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.

Controls

The Controls section includes settings for how the Control Panel is displayed.

- **Auto Control Close:** When enabled, multiple node headers will be visible, but only one node header will be opened to display parameters. When disabled, any number of nodes may be opened to display parameters at the same time. This setting has no effect unless the Auto Control Hide checkbox is disabled.
- **Auto Control Hide:** When enabled, only the parameters for the currently active node will be made visible. Otherwise, all node headers will be visible and displayed based on the Auto Control Close setting.
- **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 a node's control. 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 node's parameters are shown in the Control Panel. By default, it is enabled showing controls for the active node as well as all selected nodes.
- **Combined Color Wheel:** When the Color Corrector node is displayed in the Control Panel, 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 Control Panel.
- **Grab Distance:** This slider ranges from 1 to 10 and defaults to 5. It designates the active area around the cursor and can be modified if you have difficulties in selecting points for modification in paths and spline curves. Smaller values represent more accurate selection.

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 where as entering 10:00 causes and Auto Save every 10 minutes.

Proxy

- **Update All, Selective, No Update:** The Update mode button is located in the Time Ruler. You can use this preference to determine the default mode for all new comps. Selective is the usual default. It renders only the nodes needed to display the images in the Display view. All will render all nodes 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 beside the Render button and selecting the desired value from the contextual menu.

Path Map

Path Maps are virtual paths used to replace segments of filepaths 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.

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 in render farms 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.

  If the Reverse Pathmap option is enabled, Fusion will automatically attempt to replace portions of filenames that match known path maps. For example, a composition saved as C\Comps\example1.comp would automatically replace the path to c:\comps\example.jpg with Comp:\example.jpg.

- **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.
- **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 path map, select it from the list and click on the delete button.
- **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.

Script

The preferences for Scripting include a field for passwords used to execute scripts from the command line and programs for use when editing scripts.

Login

- **No Login Required to Execute Script:** 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 onto 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. See the Scripting documentation for details.

Options

- **Script Editor:** Use this preference to select an external editor for scripts (used when selecting Scripts > Edit from the menu).

Spline Editor

The Spline Editor preference allows you to set various spline options for Autosnap behavior, handles, markers, and more.

Spline Editor Options

Independent Handles: Enabling this option allows the In or Out direction handle on newly created keyframes to be moved independently without affecting the other.

- **Follow Active:** The Spline Editor focuses on the currently active node.
- **Show Key Markers:** Small colored triangles will be displayed at the top of the Spline Editor Time Ruler to indicate keyframes on active splines. The colors of the triangles match the colors of the splines.
- **Show Tips:** Toggles if tooltips are displayed or not.
- **Autosnap points:** When moving points in the Spline Editor, these will snap to the Fields or Frames or can be moved freely.
- **Guides:** When moving points in the Spline Editor, these will snap to Guides as well.
- **Autosnap Guides:** When moving or creating Guides, these will snap to the Fields or Frames or can be moved freely.
- **Autoscale:** Keeps the Spline View scales intact on changing the editable spline content of the view.
- **Scroll:** Scrolls horizontally and vertically to show all or most of the spline points on changing the editable spline content of the view.
- **Fit:** Zooms to fit all points within a Spline view, if necessary, on changing the editable spline content of the view.

LUT View Options

- **Independent Handles:** Enabling this option allows the In or Out direction handle on newly created keyframes 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 keyframes 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 more are found in the Splines preferences.

Autosmooth

Automatically smooths out any newly created points or keyframes on the splines selected in this section. This is valid for both animation splines as well as masks, LUTs, paths and meshes.

B-Spline Modifier Degree

This setting determines the degree used for interpolating B-Splines when used as modifiers. Cubic B-Splines determine a segment through 2 control points between the anchor points, and Quadratic B-Splines determine a segment through 1 control point between the anchor points.

B-Spline Polyline Degree

This setting is like the one above, but applies to polylines, like those used for masks.

- **Tracker Path Points Visibility:** This setting controls the control points on tracker paths. You have three options. 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 a 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 Polyline node that is selected in the toolbar once you complete the drawing of a Polyline.

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 are 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 Timeline Editor.

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 actually sets the filter in the Timeline.

Settings for Filters

To create a new filter and define its settings, you first click the New button. Creating a new filter will allow you to select any of the nodes that the filter will contain. The Settings for Filters area includes a list of all the nodes. Only nodes that are checked will appear in the Timeline 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 Timeline are enabled by default. A series of checkboxes correspond to buttons located in the Timeline, allowing you to determine the states of those button at the time a new comp is created. For more information on the Timeline functions see Chapter 19, "Creating and Working with Timelines."

- **Autosnap points:** When moving points in the Timeline Editor, the points will snap to the fields, frames, or can be moved freely.
- **Guides:** When moving points in the Timeline Editor, the point will snap to the guides that are placed in the Timeline.
- **Autosnap Guides:** When moving or creating guides, the guides will snap to the fields, frames, or can be moved freely.
- **Autoscale:** Keeps the Timeline scales intact while changing the editable spline content of the view. 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 of the view. When set to Fit, the Timeline zooms to fit all points within the Viewer, if necessary.
- **Tools Display Mode:** This menu controls the default sort order or the tools displayed in the Timeline. The default can be changed using the Sort order menu in the Timeline.

Tweaks

The Tweaks preferences handle a collection of settings for fine tuning I/O, memory, and graphics hardware behavior.

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 TIFF, VideoPump and other formats that perform a large amount of seeking.

  This option has not been tested with every hardware and OS configuration, so it is recommended to enable it only after you have thoroughly tested your hardware and OS configuration using drive loads from both local disks and network shares.
- **Enable Direct Reads:** Enabling this checkbox uses a more efficient method when loading a large chunk of contiguous data into memory by reducing I/O operations. Not every operating system employs this ability, so it may produce unknown behavior.
- **Read Ahead Buffers:** This slider determines the number of 64K buffers that are use to read ahead in a file I/O operation. The more buffers, the more efficient loading frames from disk will be, but the less responsive it will be to changes that require disk access interactively.

Area Sampling

- **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 node 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:** De-selecting 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. 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 nodes 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 node if cache memory is low. There are three settings to choose from: None, Outside and All.
  - **None:** This setting never displays the translucent overlay or controls, which can reduce the need for background renders, in some cases, resulting in a speed up of the display.
  - **Outside:** This will display only those areas of the control that are outside the bounds of the image, which can reduce visual confusion.
  - **All:** Displays all overlays of all selected nodes.
- **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, 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.

  This is only available in Resolve Studio.

User Interface

These Preferences set the appearance of the User Interface window and how the Control Panel is displayed.

View

The View preference is used to manage settings and default controls for Viewers.

View Settings

The area at the top of the View 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, see Chapter 52, "Using 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 View drop down menu is used to select the Viewer for which you want to configure the settings.

Control Colors

The Control Colors setting allows you to determine the color of the Active/Inactive onscreen control.

Color Picking Area Size

Use these width/height controls to set the number of pixels sampled when using the color picker in the Viewers.

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.

Displayed Depth Range

The Displayed Depth Range setting controls the view normalization of the Z-Channel.

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.

Headset Options

API	
Disabled	Turns off and hides all usage of headsets.
Auto	Will detect which headset is plugged in.
Occulus	Will set the VR ouput 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), i n 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 which method the stereo images are stacked.
Vstack	The stereo images are stacked vertically as left on top and right at the bottom.
Hstack	The stereo images are stacked horizontally as left and right. **Swap eyes:** if stereo is reversed this option will swap the eyes

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

Import EDL

The EDL Import options are used to determine how compositions are created from imported EDL files.

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 Dissolve node will be created with the settings imported from the EDL file.
- **Loader Per Transition:** A Loader will be created with the settings representing the settings imported from the EDL file (Loader with a Clip list created).

Use Shot Names

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