

PART 8

Fusion Fundamentals



Chapter 61

Introduction to Compositing in Fusion

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

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

Contents

What Is Fusion?	986
The Fusion Page within DaVinci Resolve	986
The Fusion Studio Stand-Alone Application	987
What Kinds of Effects Does Fusion Offer?	988
How Hard Will This Be to Learn?	989

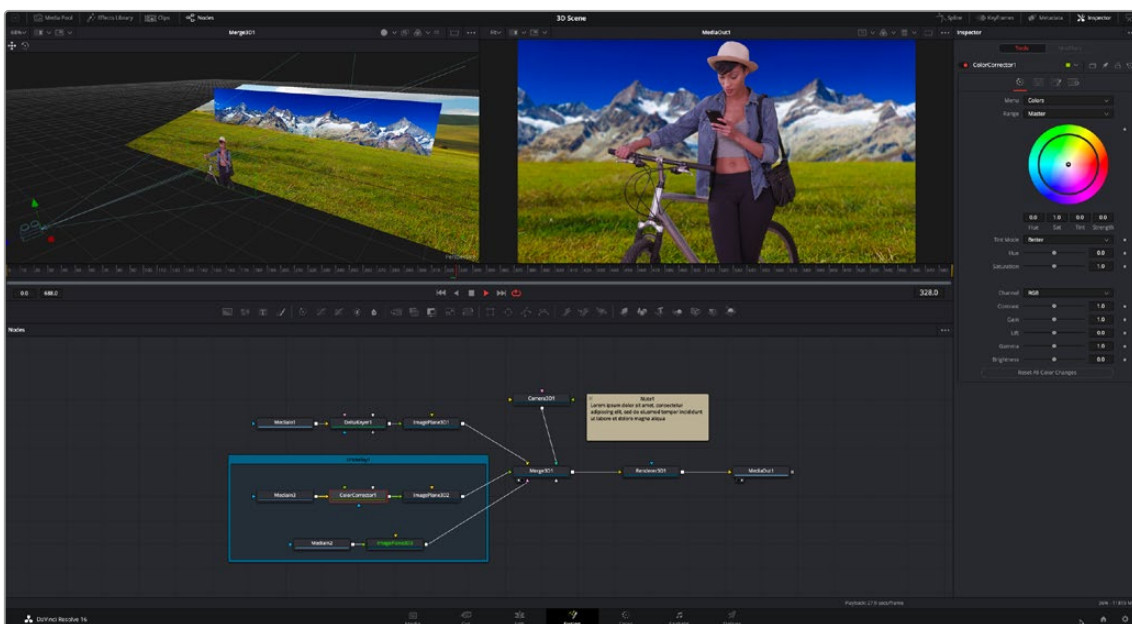
What Is Fusion?

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

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

The Fusion Page within DaVinci Resolve

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



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

How Do I Use the Fusion Page?

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

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



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

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

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

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

How Do Fusion Effects Differ from Edit Page Effects?

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

How Do Fusion Effects Differ from Color Page Effects?

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

The Fusion Studio Stand-Alone Application

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

Rendering Out Your Final Result


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

What Kinds of Effects Does Fusion Offer?

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

3D Compositing


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



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

Particles

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



A 3D particle system, also created entirely within Fusion

Text

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



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

And Lots More

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

How Hard Will This Be to Learn?

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

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

Chapter 62

Exploring the Fusion Interface

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


Contents

The Fusion User Interface	992
The Work Area	993
Interface Toolbar	993
Choosing Which Panel Has Focus	994
Viewers	994
Zooming and Panning into Viewers	996
Loading Nodes Into Viewers	996
Clearing Viewers	997
Viewer Controls	997
Time Ruler and Transport Controls	999
Time Ruler Controls in the Fusion Page	999
Time Ruler Controls in Fusion Studio	1000
The Playhead	1000
Zoom and Scroll Bar	1001
Transport Controls in the Fusion Page	1001
Audio Monitoring	1003
Transport Controls in Fusion Studio	1004
Changing the Time Display Format	1008
Keyframe Display in the Time Ruler	1008
The Fusion RAM Cache for Playback	1008

Toolbar	1010
Customizing the Toolbar	1010
Node Editor	1012
Adding Nodes to Your Composition	1012
Removing Nodes from Your Composition	1013
Identifying Node Inputs and Node Outputs	1013
Node Editing Essentials	1013
Navigating the Node Editor	1015
Vertical Node Editor Layouts	1015
Keeping Organized	1016
Status Bar	1016
Effects Library	1017
The Inspector	1018
The Tools and Modifiers Panels	1019
Parameter Header Controls	1019
Parameter Tabs	1019
Keyframes Editor	1020
Keyframes Editor Control Summary	1020
Adjusting Clip Timings	1021
Adjusting Effect Timings	1021
Adjusting Keyframe Timings	1021
Spline Editor	1022
Spline Editor Control Summary	1023
Choosing Which Parameters to Show	1023
Essential Spline Editing	1024
Essential Spline Editing Tools and Modes	1024
Thumbnail Timeline in the Fusion Page	1026
The Media Pool in the Fusion Page	1026
Importing Media Into the Media Pool on the Fusion Page	1027
Bins in Fusion Studio	1028
Bins Interface	1028
The Bin Studio Player	1029
The Console	1030
Customizing Fusion	1031
The Fusion Settings Window	1031
Showing and Hiding Panels	1032
Resizing Panels	1032
Fusion Studio Floating Frame	1032
Fusion Studio Keyboard Remapping	1033
Undo and Redo	1034

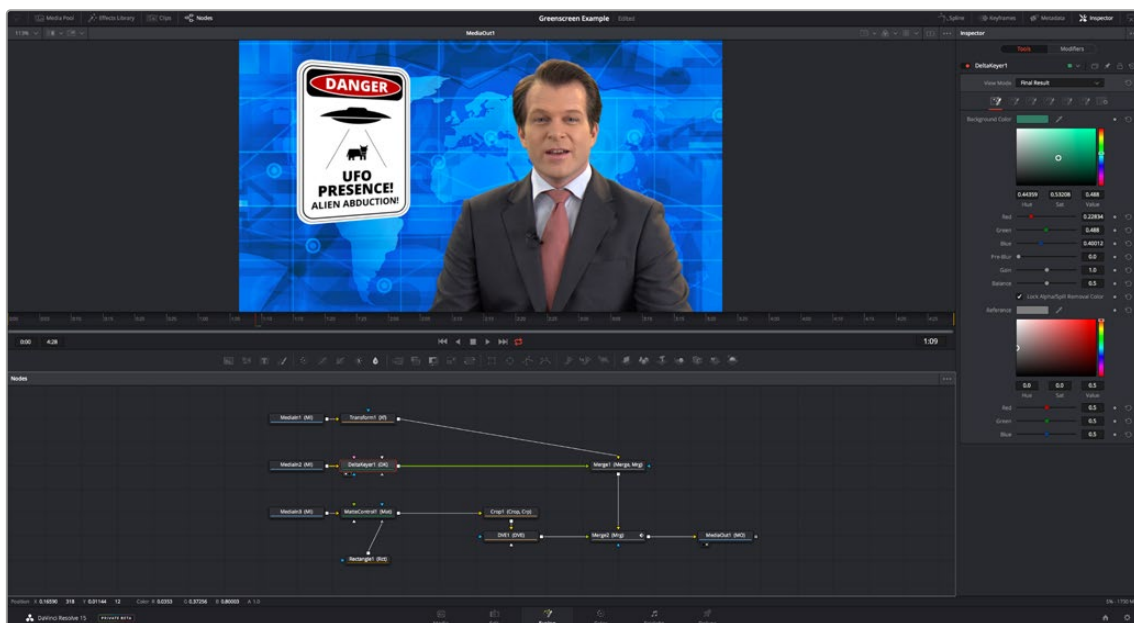
The Fusion User Interface

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



The Fusion user interface shown completely

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



A simplified set of Fusion controls for everyday working

The Work Area

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



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

Interface Toolbar

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



The UI toolbar of the Fusion page



The UI toolbar of Fusion Studio

These buttons are as follows, from left to right:


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

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

Choosing Which Panel Has Focus

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


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



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

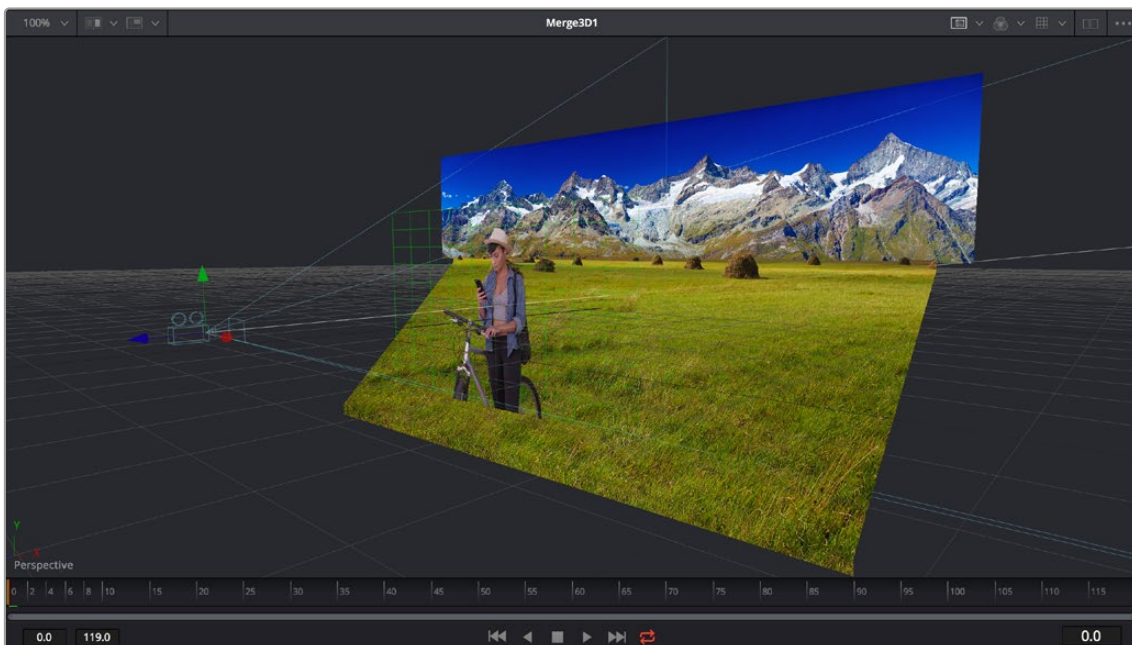
Viewers

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



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

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



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

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

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

Zooming and Panning into Viewers

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

Methods of panning viewers:

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

Methods of scaling viewers:

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

Methods of spinning 3D viewers:

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

Loading Nodes Into Viewers


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

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

To load specific nodes into specific viewers:

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

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



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

Clearing Viewers

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

Viewer Controls

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



Controls in the viewer title bar

- **Zoom menu:** Lets you zoom in on the image in the viewer to get a closer look, or zoom out to get more room around the edges of the frame for rotoscoping or positioning different layers. Choose Fit to automatically fit the overall image to the available dimensions of the viewer.
- **Split Wipe button and A/B Buffer menu:** You can actually load two nodes into a single viewer using that viewer's A/B buffers by choosing a buffer from the menu and loading a node into the viewer. Turning on the Split Wipe button (press Forward Slash) shows a split wipe between the two buffers, which can be dragged left or right via the handle of the onscreen control, or rotated by dragging anywhere on the dividing line on the onscreen control. Alternatively, you can switch between each full-screen buffer to compare them (or to dismiss a split-screen) by pressing Comma (A buffer) and Period (B buffer).
- **SubView type:** (These aren't available in 3D viewers.) Clicking the icon itself enables or disables the current "SubView" option you've selected, while using the menu lets you choose which SubView is enabled. This menu serves one of two purposes. When displaying ordinary 2D nodes, it lets you open up SubViews, which are viewer "accessories" within a little pane that can be used to evaluate images in different ways. These include an Image Navigator (for navigating when zoomed far into an image), Magnifier, 2D viewer (a mini-view of the image), 3D Histogram scope, Color Inspector, Histogram scope, Image Info tooltip, Metadata tooltip, Vectorscope, or Waveform scope. The Swap option (Shift-V) lets you switch what's displayed in the viewer with what's being displayed in the Accessory pane. When displaying 3D nodes, this button lets you 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 (Region of Interest) limiting in the viewer, while using the menu lets you choose the region of the Roi. Roi lets you define the region of the viewer in which pixels actually need to be updated. When a node renders, it intersects the current Roi with the current Domain of Definition (DoD) to determine what pixels should be affected. When enabled, you can position a rectangle to restrict rendering to a small region of the image, which can significantly speed up performance when you're working on very high resolution or complex compositions. Auto (the default) sets the region to whatever is visible at the current zoom/pan level in the viewer. Choosing Set lets you draw a custom region within the frame by dragging a rectangle that defaults to the size of the viewer, which is resizable by dragging the corners or sides of the onscreen control. Choosing Lock prevents changes from being made to the current Roi. Choosing Reset resets the Roi to the whole viewer.
- **Color controls:** Lets you choose which color and/or image channels to display in the viewer. Clicking the icon itself toggles between Color (RGB) and Alpha, the two most common things you want to see (pressing C or A also toggles between Color and Alpha). Opening the menu displays every possible channel that can be displayed for the currently viewed node, commonly including RGB, Red, Green, Blue, and Alpha (available from the keyboard by pressing R, G, B, or A). For certain media and nodes, additional auxiliary channels are available to be viewed, including Z-depth, Object ID, Material ID, XYZ Normals, and so on.
- **Viewer LUT:** Clicking the icon itself toggles LUT (LookUp Table) display on or off, while the menu lets you choose which of the many available color space conversions to apply to the viewer. The top five options let you choose Fusion controls that can be customized via the Edit item at the bottom of this menu. The rest of this menu shows all LUTs installed in the LUT directory to use for viewing.

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

- **Option menu:** This menu contains various settings that pertain to the viewers in Fusion.
 - **Snap to Pixel:** When drawing or adjusting a polyline mask or spline, the control points will snap to pixel locations.
 - **Show Controls:** Toggles whatever onscreen controls are visible for the currently selected node.
 - **Region:** Provides all the settings for the Region of Interest in the viewer.
 - **Smooth Resize:** This option uses a smoother bilinear interpolated resizing method when zooming into an image in the viewer; otherwise, scaling uses the nearest neighbor method and shows noticeable aliasing artifacts. However, this is more useful when you zoom in at a pixel level since there is no interpolation.
 - **Show Square Pixels:** Overrides the auto aspect correction when using formats with non-square pixels.
 - **Normalized Color Range:** Allows for the visualization of brightness values outside of the normal viewing range, particularly when working with floating-point images or auxiliary channels.
 - **Checker Underlay:** Toggles a checkerboard underlay that makes it easy to see areas of transparency.
 - **Gain/Gamma:** Exposes a simple pair of Gain and Gamma sliders that let you adjust the viewer's brightness.
 - **360 View:** Used to properly display spherical imagery in a variety of formats, selectable from this submenu.
 - **Stereo:** Used to properly display stereoscopic imagery in a variety of formats, selectable from this submenu.

Time Ruler and Transport Controls


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

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

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


Time Ruler Controls in the Fusion Page

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



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

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



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

Render Range

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

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

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


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

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

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

Time Ruler Controls in Fusion Studio

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



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

Global Start and End Range

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

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

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

Render Range

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

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

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

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

The Playhead

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

Zoom and Scroll Bar

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

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

Transport Controls in the Fusion Page

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

Controlling Playback

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



The Fusion page controls for playback

Navigation Shortcuts

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

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

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


Real-Time Playback Not Guaranteed

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

Frame Increment Options

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

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



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

Looping Options

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

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

Render Range Fields

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



The Render Start and Render End time fields

Audio Monitoring

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

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

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

To hear the audio from a specific Loader node:

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

Audio Toolbar Button

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

The Current Time Field

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

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

The Fusion Page Viewer Quality and Proxy Options

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

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

High Quality

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

Motion Blur

The Motion Blur button is a global setting. Turning off Motion Blur temporarily disables motion blur throughout the composition, regardless of any individual nodes for which it's enabled. This can

significantly speed up renders to the viewer. Individual nodes must first have motion blur enabled before this button has any effect.

Proxy

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

Auto Proxy

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

Selective Updates

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

The options are available in the Fusion Preferences General panel. The three options are:

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

Transport Controls in Fusion Studio

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



Fusion Studio transport controls

Controlling Playback

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



Fusion Studio transport controls

Navigation Shortcuts

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

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

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

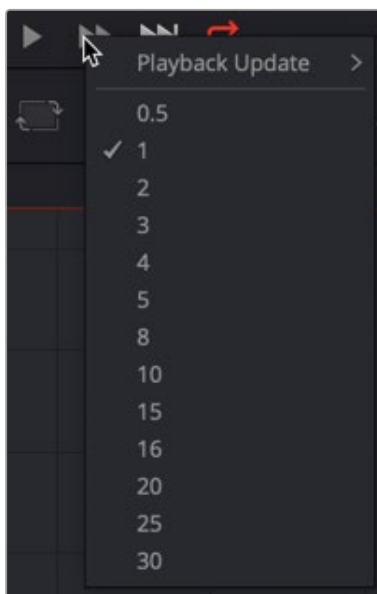
Real-Time Playback Not Guaranteed

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

Frame Increment Options

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

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



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


Looping Options

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

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

Range Fields

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



Global Time Fields to the left of the transport controls

Audio

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

Render

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

The Current Time

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


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

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

Fusion Studio Viewer Quality and Proxy Options

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

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



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

HiQ

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

MB

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

Prx

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

Aprx

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

Selective Updates

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

The three options are:

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

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

Changing the Time Display Format


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

To change the time display format:

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

Keyframe Display in the Time Ruler

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



The Time Ruler displaying keyframe marks

To move the playhead in the Time Ruler among keyframes:

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

The Fusion RAM Cache for Playback

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

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

Memory Limits of the RAM Cache

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

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


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

- **Limit Caching To:** This slider sets the maximum amount of RAM used for caching. The 60% default setting on a 32-GB system limits the cache to 19.2 GB. The maximum amount you can assign to Fusion Studio is limited to 80% of the total system memory. This leaves a minimum amount of memory for other applications and the operating system.
- **Leave at least # MBytes:** This number field further limits caching in cases where the system's available free RAM drops below the entered value. For instance, setting this to 200 MB attempts to keep 200 MB of RAM free for the OS or other applications. Setting the number field to 0 allows Fusion Studio to use the full amount of RAM specified by the Limit Caching To setting, ignoring other apps.

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

Displaying Cached Frames


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



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

Temporarily Preserving the Cache When Changing Quality or Proxy Settings

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



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

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

Toolbar

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



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

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

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

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

Customizing the Toolbar

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

To create a new toolbar, do the following:

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

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

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

To add a tool to the toolbar:

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

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

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

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

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

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

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

To prevent a custom toolbar from being modified:

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

To switch between toolbars:

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


To delete a custom toolbar, do the following:

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

Node Editor

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

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



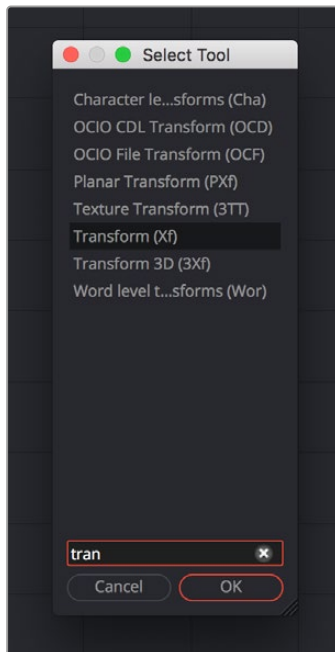
The Node Editor displaying a node tree creating a composition

Adding Nodes to Your Composition

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

Methods of adding nodes include:

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



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

Removing Nodes from Your Composition


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

Identifying Node Inputs and Node Outputs

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

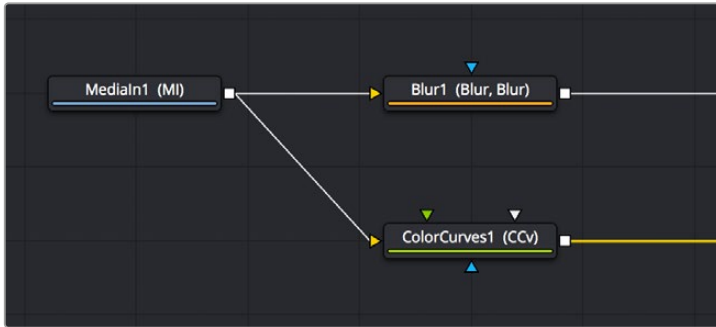
Node Editing Essentials

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




Three nodes connected together

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




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

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



Two nodes being merged together into one to create a composite

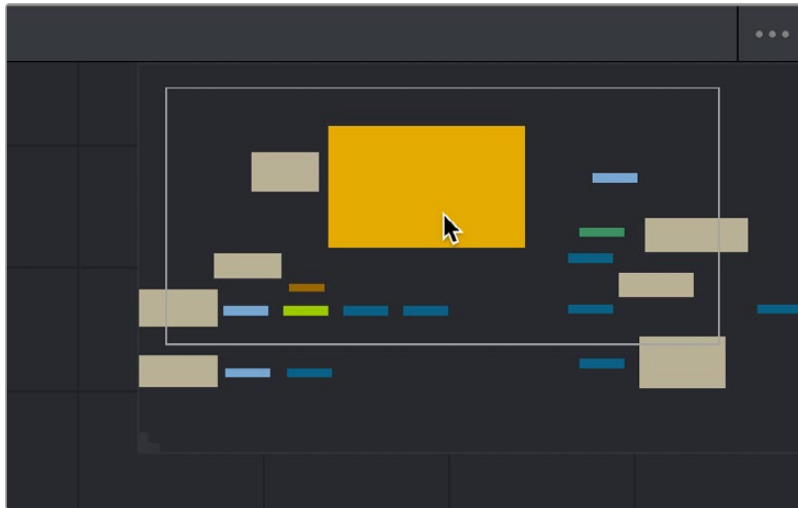
By default, new nodes are added from left to right in the Node Editor, but they can also flow from top to bottom, right to left, bottom to top, or in all directions simultaneously. Connections automatically reorient themselves along all four sides of each node to maintain the cleanest possible presentation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



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

Navigating the Node Editor

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



The Navigator pane for accessing offscreen parameters or tools

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


Methods of navigating the Node Editor:

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

Vertical Node Editor Layouts

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

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



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

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

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

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


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

Keeping Organized

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


Status Bar

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



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

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



A notification that there's a message in the Console

Effects Library

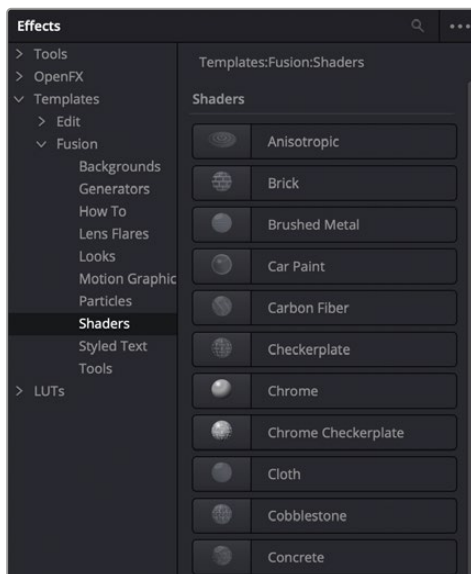
The Effects Library in Fusion shows all the nodes and effects available in Fusion, including third-party OFX plug-ins, if installed. If you are using DaVinci Resolve, Resolve FX also appear in the OFX category. While the toolbar shows many of the most common nodes you'll use in any composite, the Effects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



The Effects Library with Tools open

The hierarchical category browser of the Effects Library is divided into several sections depending on whether you are using Fusion Studio or the Fusion page within DaVinci Resolve. The Tools section is the most often used since it contains every node that represents an elemental image-processing operation in Fusion. Hovering the pointer over a specific tool will reveal a tool-tip

explaining its functionality at the bottom right of the DaVinci Resolve interface. The OpenFX section contains third-party plug-ins, and if you are using the Fusion page, it also contains ResolveFX, which are included with DaVinci Resolve. A third section, only visible when using the Fusion page in DaVinci Resolve, is the Templates section. The Template section contains a variety of additional content including templates for Lens Flares, Backgrounds, Generators, Particle Systems, Shaders (for texturing 3D objects), and other resources for use in your composites.



The Templates section of the Effects Library

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

The Inspector


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



The Inspector shows parameters from one or more selected nodes.

The Tools and Modifiers Panels

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




The Modifier panel showing a Perturb modifier

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

Parameter Header Controls

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




Common Inspector controls

- **Set Color:** A pop-up menu that lets you assign one of 16 colors to a node, overriding a node's own color.
- **Versions:** Clicking Versions reveals another toolbar with six buttons. Each button can hold an individual set of adjustments for that node that you can use to store multiple versions of an effect.
- **Pin:** The Inspector is also capable of simultaneously displaying all parameters for multiple nodes you've selected in the Node Editor. Furthermore, a Pin button in the title bar of each node's parameters lets you "pin" that node's parameters into the Inspector so that they remain there even when that node is deselected, which is valuable for key nodes that you need to adjust even while inspecting other nodes of your composition.
- **Lock:** Locks that node so that no changes can be made to it.
- **Reset:** Resets all parameters within that node.

Parameter Tabs


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



Nodes with several tabs' worth of parameters

Keyframes Editor

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



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

Keyframes Editor Control Summary

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

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

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

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

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

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

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

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

Adjusting Clip Timings

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

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


Adjusting Effect Timings

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

Adjusting Keyframe Timings

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

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



Keyframe tracks exposed

Keyframe Editing Essentials

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

Methods of adjusting keyframes:

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

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

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

Time Stretching Keyframes

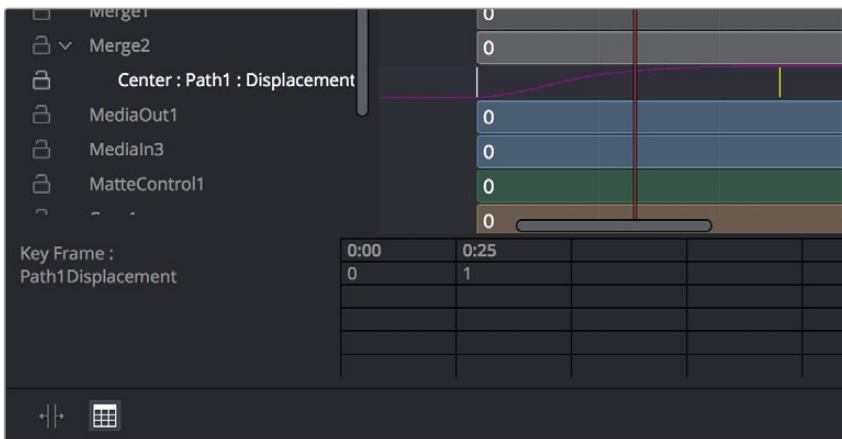
If you select a range of keyframes in a keyframe track, you can turn on the Time Stretch tool to show a box used to squeeze and stretch the entire range of keyframes relative to one another. The Time Stretcher changes the overall timing of a sequence of keyframes without losing the relative timing from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



Time stretching keyframes

The Keyframe Spreadsheet

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



Editing keyframes in the Spreadsheet

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

Spline Editor

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



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

Spline Editor Control Summary

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

- Vertical and horizontal zoom controls let you scale the size of the editor.
- A Zoom to Fit button fits the width of all curves to the current width of the Spline Editor.
- A Zoom to Rect tool lets you draw a rectangle to define an area of the Spline Editor to zoom into.
- A Sort drop-down menu lets you sort or filter the curves in various ways.
- An Option menu provides access to many other ways of filtering curves 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. The horizontal axis represents time, and the vertical axis represents the spline's value. Selected control points show their values in the edit fields at the bottom of the graph.

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

Choosing Which Parameters to Show

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

To show every parameter in every node:

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

To show splines for the currently selected node:

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

Essential Spline Editing

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

To select one or more control points:

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

To edit control points and splines:

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

To edit Bézier curves:

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

To delete control points:


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

Essential Spline Editing Tools and Modes

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

Control Point Interpolation

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



Control point interpolation controls

- **Smooth:** Creates automatically adjusted Bézier curves to create smoothly interpolating animation.
- **Flat:** Creates linear interpolation between control points.
- **Invert:** Inverts the vertical position of non-animated LUT splines. This does not operate on animation splines.
- **Step In:** For each keyframe, creates sudden changes in value at the next keyframe to the right. Similar to a hold keyframe in After Effects® or a static keyframe in the DaVinci Resolve Color page.
- **Step Out:** Creates sudden changes in value at every keyframe for which there's a change in value at the next keyframe to the right. Similar to a hold keyframe in After Effects or a static keyframe in the DaVinci Resolve Color page.
- **Reverse:** Reverses the horizontal position of selected keyframes in time, so the keyframes are backward.

Spline Loop Modes

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




Spline Loop modes

- **Set Loop:** Repeats the same pattern of keyframes over and over.
- **Set Ping Pong:** Repeats a reversed set of the selected keyframes and then a duplicate set of the selected keyframes to create a more seamless pattern of animation.
- **Set Relative:** Repeats the same pattern of selected keyframes but with the values of each repeated pattern of keyframes being incremented or decremented by the trend of all keyframes in the selection. This results in a loop of keyframes where the value either steadily increases or decreases with each subsequent loop.

Spline Editing Tools

The next five buttons provide specialized Spline editing tools.




Spline editing controls

- **Select All:** Selects every keyframe currently available in the Spline Editor.
- **Click Append:** Click once to select this tool and click again to de-select it. This tool lets you add or adjust keyframes and spline segments (sections of splines between two keyframes) depending on the keyframe mode you're in. With Smooth or Linear keyframes, clicking anywhere above or below a spline segment adds a new keyframe to the segment at the location where you clicked. With Step In or Step Out keyframes, clicking anywhere above or below a line segment moves that segment to where you've clicked.
- **Time Stretch:** If you select a range of keyframes, you can turn on the Time Stretch tool to show a box you can use to squeeze and stretch the entire range of keyframes relative to one another, to change the overall timing of a sequence of keyframes without losing the relative timing from one keyframe to the next. Alternatively, you can turn on Time Stretch and draw a bounding box around the keyframes you want to adjust to create a time-stretching boundary that way. Click Time Stretch a second time to turn it off.
- **Shape Box:** Turn on the Shape Box to draw a bounding box around a group of control points you want to adjust in order to horizontally squish and stretch (using the top/bottom/left/right handles), corner pin (using the corner handles), move (dragging on the box boundary), or corner stretch (Command-drag the corner handles).
- **Show Key Markers:** Turning on this control shows keyframes in the top ruler that correspond to the frame at which each visible control point appears. The colors of these keyframes correspond to the color of the control points they're indicating.


Thumbnail Timeline in the Fusion Page

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



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

Right-clicking on any thumbnail exposes a contextual menu.



The contextual menu for the Thumbnail timeline

To open another clip:

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

To 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.

The Media Pool in the Fusion Page

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



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, the clips are automatically connected to that line via enough Merge nodes to connect them all.

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

Importing Media Into the Media Pool on the Fusion Page

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

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

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


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

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

For more information on importing media using the myriad features of the Media page, see Chapter 18, “Adding and Organizing Media with the Media Pool” in the DaVinci Resolve Reference Manual.

Bins in Fusion Studio

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



The Bins window


To open the Bins window:

- Choose File > Bins from the menu bar.

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

Bins Interface


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



The Bin list sidebar


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

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



The Bins icon view

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




The Bins toolbar

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

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

The Bin Studio Player

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




The Bin Studio Player

For more information on Bins and the Studio Player, see Chapter 73, “Bins” in the DaVinci Resolve Reference Manual or Chapter 13 in the Fusion Reference Manual.

The Console

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

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



The Console window

A toolbar at the top of the Console contains controls governing what the Console shows. At the top left, the Clear Screen button clears the contents of the Console. The next four buttons toggle the visibility of error messages, log messages, script messages, and input echoing. Showing only a particular kind of message can help you find what you're looking for when you're under the gun at 3:00 in the morning. The next three buttons let you choose the input script language. Lua 5.1 is the default and is installed with Fusion. Python 2.7 and Python 3.6 require that you install the appropriate Python environment on your computer. Because scripts in the Console are executed immediately, you can switch between input languages at any time.


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

Customizing Fusion

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

The Fusion Settings Window

Fusion has its own settings window, accessible by choosing Fusion > Fusion Settings in DaVinci Resolve, or in Fusion Studio by choosing Fusion > Preferences on macOS or File > Preferences on Windows or Linux. This window has a variety of options for customizing the Fusion experience. For more information, see Chapter 75, "Preferences" in the DaVinci Resolve Reference Manual or Chapter 15 in the Fusion Reference Manual.



The Fusion Settings window set to the General panel

Saving Fusion Layouts

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

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

- Choose Workspace > Layout Presets > Save Layout Presets.

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

- Click the Grab Document Layout button in the Preferences > Layout panel to save the layout for all new Compositions.
- Click the Grab Program Layout button to remember the size and position of any floating views, and enable the Create Floating Views checkbox to automatically create the floating windows when Fusion restarts.

Showing and Hiding Panels

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



The UI toolbar of the Fusion page

Resizing Panels

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



Dragging the edge between two viewers to resize it

Fusion Studio Floating Frame

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


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

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

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

Fusion Studio Keyboard Remapping

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



The Fusion Studio Hotkey Manager window

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

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

- 1 Open the Keyboard Hotkey Manager.
- 2 Select Views > Effect from the Target area of the Hotkey Manager.
- 3 Below the Key/Action area, click the New button to create a new keyboard shortcut for the Node.
- 4 In the Edit Hotkey window, click the Tools disclosure arrow, and then select Blur to display all the Blur-related nodes.
- 5 Select Glow in the Action panel.
- 6 At the top of the Edit Hotkey window, type G as the shortcut for the Glow node, and then click OK.
- 7 Glow and the G hotkey will now appear in the Key/Action area on the right.
- 8 Click OK to close the Hotkey Manager.
- 9 Click in the Node Editor and press G to add a Glow node.

Undo and Redo

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

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

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

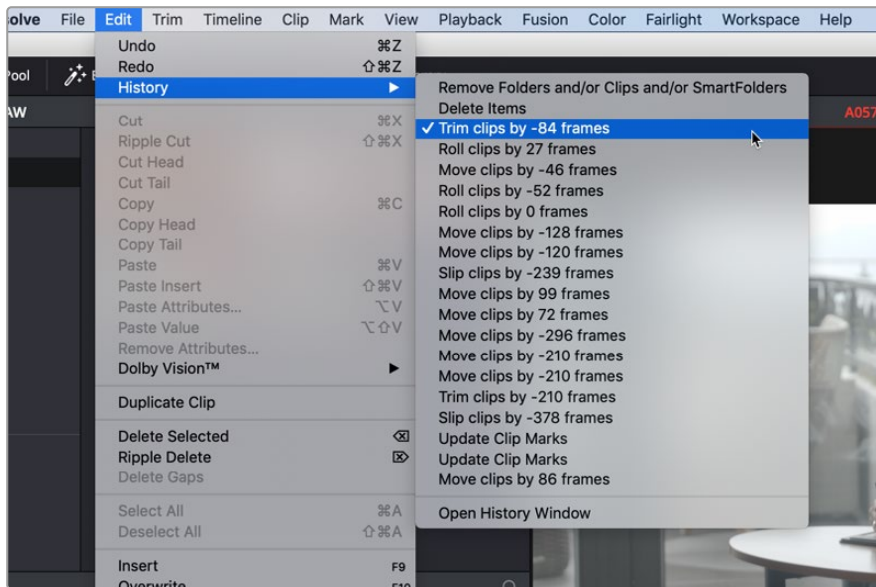
- Choose Edit > Undo (Command-Z) to undo the previous change.
- Choose Edit > Redo (Shift-Command-Z) to redo to the next change.

Undo History in DaVinci Resolve

When using the Fusion page in DaVinci Resolve, 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, Cut, Edit, and Fairlight pages.

To undo and redo using the History submenu:

- 1 Open the Edit > History submenu, which shows (up to) the last 20 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 lets you browse the entire available undo stack of the current page.

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

Sharing Memory with Other DaVinci Resolve Pages

When using the Fusion page within DaVinci Resolve, the Memory and GPU panel in the Preferences window controls the overall amount of system memory used by the application and a subset allocated to the Fusion page. The Fusion page is limited to roughly 65% of the total memory allocated to DaVinci Resolve. The memory allocated to the Fusion page is released when you are not actively working on the page.

Chapter 63

Getting Clips into Fusion

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

Contents

Preparing Compositions in the Fusion Page	1037
Working on Single Clips in the Fusion Page	1037
Turning One or More Clips into Fusion Clips	1038
Adding Fusion Composition Generators	1039
Creating a Fusion Composition Clip in a Bin	1040
Using Fusion Transitions	1040
Adding Clips from the Media Pool	1041
Adding Clips from the File System	1042
Using MediaIn Nodes	1042
MediaIn Node Inputs	1042
Inspector Properties of MediaIn Nodes	1042
Using Loader and Saver Nodes in the Fusion Page	1045
Preparing Compositions in Fusion Studio	1047
Setting Up a Composition	1049
Reading Clips into Fusion Studio	1050
Aligning Clips in a Fusion Studio Composition	1051
Loader Node Inputs	1052
Using Proxies for Better Performance	1052
Presetting Proxy Quality	1053
File Format Options	1054
Loading Audio WAV Files in Fusion Studio	1056


Preparing Compositions in the Fusion Page

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

Working on Single Clips in the Fusion Page

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

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



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

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

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



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

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

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

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

How Nodes Are Named


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


Turning One or More Clips into Fusion Clips

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

To create a Fusion clip:

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





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

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



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

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

Adding Fusion Composition Generators

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

To create a blank Fusion clip in the Edit page:

- 1 Open the Effects Library, and select the Effects category.

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

NOTE: Audio on a track beneath a Fusion Composition effect cannot be heard in the Fusion page.

Creating a Fusion Composition Clip in a Bin

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

To create a blank Fusion Composition clip in a bin:


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

Using Fusion Transitions

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

To apply and open a Fusion Transition:

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



A Noise Dissolve Transition opened in the Fusion page.

To learn about creating custom Fusion Transitions that appear in the Effects Library, go to Chapter 66, “Node Groups, Macros, and Fusion Templates” in the DaVinci Resolve Reference Manual or Chapter 6 in the Fusion Reference Manual.

Adding Clips from the Media Pool

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



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

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

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

Audio with Media Pool Clips

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

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

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

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

Adding Clips from the File System

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

Using MediaIn Nodes

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

MediaIn Node Inputs

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


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

Inspector Properties of MediaIn Nodes

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

MediaIn Node Parameters for Clips in a Timeline

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



The Inspector Image tab parameters for a clip from the Timeline

Image Tab

- **Clip Name:** Displays the name of that clip.
- **Process Mode:** Lets you choose whether the clip represented by that node will be processed as Full Frames, or via one of the specified interlaced methods.
- **MedialID:** A unique ID assigned by DaVinci Resolve for that clip.
- **Layer:** In the case of Photoshop PSD files, this selects the layer of the PSD file to use.

Source Color Space

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

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

Source Gamma Space


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

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

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

Audio Tab

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



The Medialn Audio tab

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

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

Purging the Audio Cache

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


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

- Click the Purge Audio Cache button in the Inspector.

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

Aligning Clips from the Media Pool

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



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

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

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

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

If the Global In and Out values are decreased to the point where the range between the In and Out values is smaller than the n number of available frames in the clip, Fusion automatically trims the clip by adjusting the Clip Time range control. If the Global In/Out values are increased to the point where the range between the In and Out values is larger than the number of available frames in the clip, Fusion automatically lengthens the clip by adjusting the Hold First/Last Frame controls. Extended frames are visually represented in the range control by changing the color of the held frames to purple in the control.

- **Trim:** The Trim range control is used to trim frames from the start or end of a clip. Adjust the Trim In to remove frames from the start and set Trim Out to specify the last frame of the clip. The values used here are offsets. A value of 5 in Trim In would use the 5th frame in the sequence as the start, ignoring the first four frames. A Trim Out value of 95 would stop loading frames after the 95th.
- **Hold First Frame/Hold Last Frame:** The Hold First Frame and Hold Last Frame controls will hold the first or last frame of the clip for the specified amount of frames. Held frames are included in a loop if the footage is looped.
- **Reverse:** Select this checkbox to reverse the footage so that the last frame is played first and the first frame is played last.
- **Loop:** Select this checkbox to loop the footage until the end of the project. Any lengthening of the clip using Hold First/Last Frame or shortening using Trim In/Out is included in the looped clip.

Using Loader and Saver Nodes in the Fusion Page


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

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

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

Loader Node Parameters

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



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


Outputting Images Using Saver Nodes

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

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

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

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



The Inspector parameters for a Saver node.

Manual Disk Caching Using Loader and Saver Nodes


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

Preparing Compositions in Fusion Studio

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

Opening, Closing, and Saving Compositions

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



Three compositions opened as tabs in Fusion Studio.

To create a second new composition:

- Choose File > New.

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

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

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

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

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

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

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

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

Auto Save

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

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

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

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

The Composition File Format

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

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

Importing and Exporting a Composition from DaVinci Resolve

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

To export a Fusion composition from DaVinci Resolve:

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

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

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

MediaOut nodes are converted to Saver nodes.

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

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

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

The new comp is loaded into the Node Editor, replacing the previously existing composition.


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

Setting Up a Composition

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

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

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



Frame Format preferences for images generated in Fusion.

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

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

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


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

For more information on preferences, see Chapter 75, “Preferences” in the DaVinci Resolve Reference Manual or Chapter 15 in the Fusion Reference Manual.

Reading Clips into Fusion Studio

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

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



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

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

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

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

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

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

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

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


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

For more information about the Loader node, see Chapter 104, “I/O Nodes” in the DaVinci Resolve Reference Manual or Chapter 44 in the Fusion Reference Manual.

Aligning Clips in a Fusion Studio Composition


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

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



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

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



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

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

Instead of using the Inspector to adjust timing, it is visually more obvious if you use the Keyframes Editor. For more information on the Keyframes Editor and adjusting a clip's time, see Chapter 69, "Animating in Fusion's Keyframe Editor" in the DaVinci Resolve Reference Manual or Chapter 9 in the Fusion Reference Manual.

Loader Node Inputs

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

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

Using Proxies for Better Performance


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

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

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

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


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



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

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

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



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

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

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

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

Presetting Proxy Quality

When using Fusion Studio, rather than right-clicking over the Proxy button to set the proxy quality, you can preset the standard and Auto Proxy quality in the Fusion Preferences window. The General pane in the Preferences window includes sliders for both standard Proxy files and Auto Proxy files.

These sliders designate the default ratio used to create proxies when the Proxy and Auto Proxy modes are turned on. These settings do not affect the final render quality.

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

File Format Options


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

DPX

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

OpenEXR

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




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

Photoshop PSD

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

- **Using the Media Pool in DaVinci Resolve's Fusion page:** Any PSD file added to the Media Pool in DaVinci Resolve can be accessed from the Fusion page. After dragging the PSD file from the Media Pool into the Node Editor, the image appears as a MediaIn node. From there, you can select which layer to use from the PSD file from the Layer drop-down menu in the Inspector.
- **Using a Loader node in Fusion Studio:** This lets you read in Photoshop PSD files with the ability to select the layer in the PSD file that is used in the comp. Fusion can load any one of the individual layers stored in the PSD file, or the completed image with all layers. Transformation and adjustment layers are not supported.



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

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

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

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



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

QuickTime

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

Loading Audio WAV Files in Fusion Studio

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

TIP: AIFF files can be loaded on macOS.

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

To load a WAV audio file, do the following:

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

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

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

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

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

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

Chapter 64

Rendering Using Saver Nodes

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

Contents

Rendering Overview	1058
Rendering in the Fusion Page	1058
Rendering in Fusion Studio	1059
Rendering with the Saver Node	1059
Setting Filenames for Export	1060
Using the Render Settings Dialog	1061
Render Settings Dialog Options	1062
Rendering Previews	1063
Setting Up Network Rendering in Fusion Studio	1063
Licensing for Network Rendering	1064
Configuring the Render Master and Render Nodes	1065
Setting Up the Render Manager	1067
Submitting Comps to Network Render	1068
Using the Render Settings Dialog for Network Rendering	1068
Using the Render Manager Window for Network Rendering	1069
Working with Node Groups	1070
Viewing the Render Log	1071
Using Third-Party Render Managers with Fusion Studio	1071

Preparing Compositions for Network Rendering	1073
Using Relative Paths	1073
Using Mapped Drives	1074
Installing All Fonts on Render Nodes	1075
Installing Third-Party Plug-Ins on Render Nodes	1075
Other Uses of Network Rendering	1075
Flipbook Previews	1075
Disk Cache	1075
When Renders Fail	1076
Automatic Rejoining of the Queue	1076
Relaunching Render Nodes with Fusion Server	1076
Frame Timeouts	1076
Heartbeats	1077
Managing Memory Use	1077
Override Composition Settings	1077
Render Several Frames at Once	1077
Simultaneous Branching	1078
Limitations of Render Nodes	1078
Time Stretching	1078
Linear Tools	1078
Saving to Multi-Frame Formats	1078
Troubleshooting	1078
Checking the Render Log	1079
Check the Composition	1079

Rendering Overview

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

Rendering in the Fusion Page

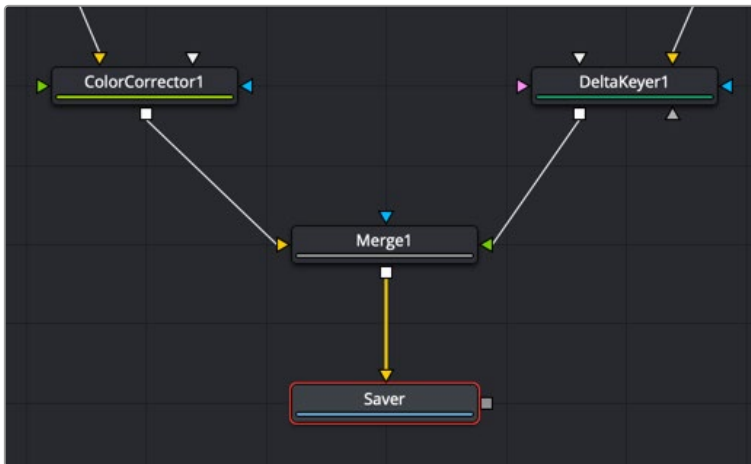
In the Fusion page, a MediaOut node is required for getting your composite from the Fusion page back into the Edit or Cut page Timeline. Whatever the MediaOut node displays when you see it in the viewer is what gets rendered back into the Edit or Cut page. This process is semi-automatic in DaVinci Resolve,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

Rendering in Fusion Studio

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

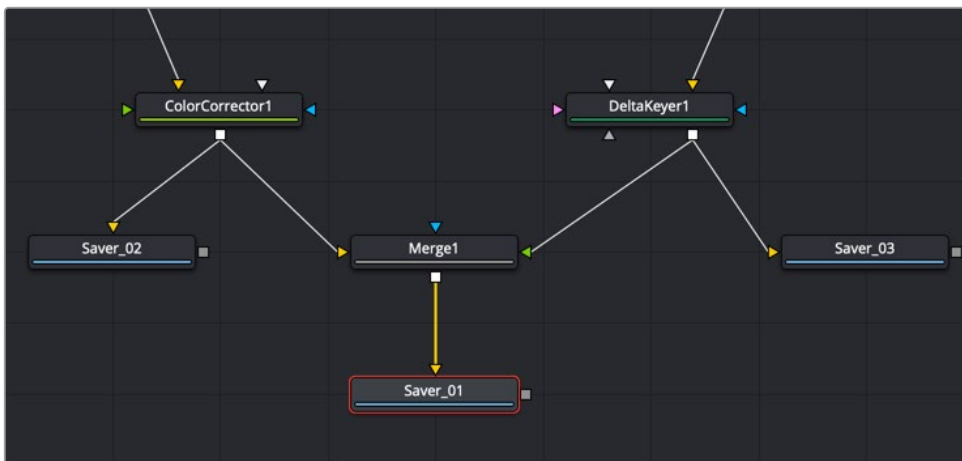
Rendering with the Saver Node

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



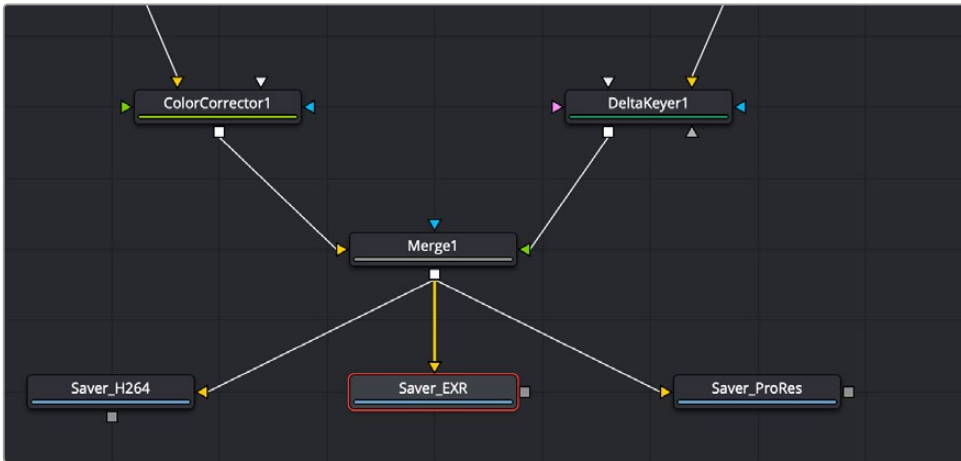
A single Saver node is added to the end of a node tree to render the final composite.

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



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

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



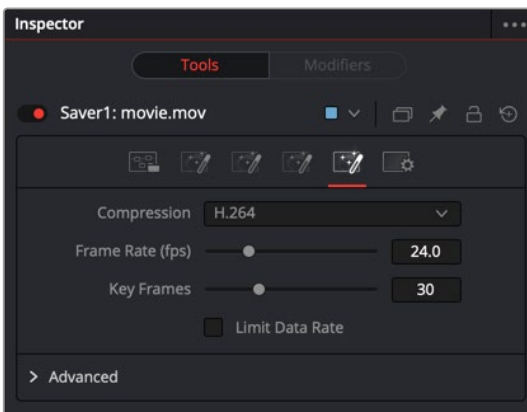
Multiple Saver nodes can be added to create different formats for output.

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

For more information on the Saver node, see Chapter 104, “I/O Nodes” in the DaVinci Resolve Reference Manual or Chapter 44 in the Fusion Reference Manual.

Setting Filenames for Export

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



The Saver’s Format tab with controls for QuickTime

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

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

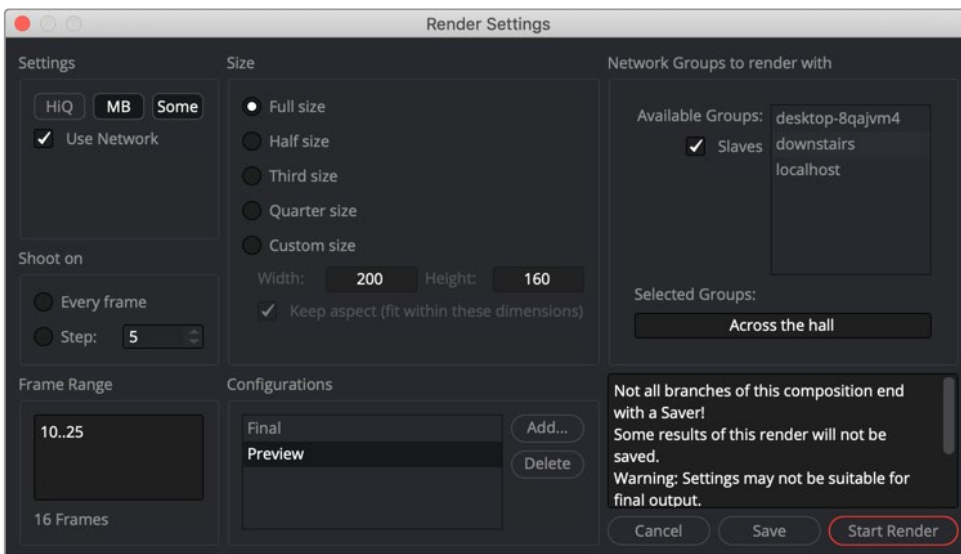
Using the Render Settings Dialog

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

To Render a comp in Fusion Studio:

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

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



The Render Settings dialog options

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

Saver nodes in the DaVinci Resolve Fusion page

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

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

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

Render Settings Dialog Options

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

Settings

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

- **HiQ:** When enabled, this setting renders in full image quality. If you need to see what the final output of a node would look like, then you would enable the HiQ setting. If you are producing a rough preview to test animation, you can save yourself time by disabling this setting.
- **MB:** The MB in this setting stands for Motion Blur. When enabled, this setting renders with motion blur applied if any node is set to produce motion blur. If you are generating a rough preview and you aren't concerned with the motion blur for animated elements, then you can save yourself time by disabling this setting.
- **Some:** When Some is enabled, only the nodes specifically needed to produce the image of the node you're previewing are rendered.

Size

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

Network

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

Shoot On

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

Frame Range

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

Configurations

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

Rendering Previews

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

For more information on rendering RAM previews, see Chapter 67, “Using Viewers,” in the DaVinci Resolve Reference Manual or Chapter 7 in the Fusion Reference Manual.

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

Setting Up Network Rendering in Fusion Studio

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

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

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

To install a Render node:

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

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

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

Licensing for Network Rendering

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

Multi-License Dongles

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

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

Setting Up a License Server

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

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

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

Setting Up a License Server with Environment Variables

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

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

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

Like most environment variables, you can put the license server in the Global Preferences via the Prefs text file. EnvironmentVars: fu:SetPrefs("Global.EnvironmentVars. FUSION_LICENSE_SERVER", "10.0.0.23;*) fu:SavePrefs()

See Chapter 75, "Fusion Preferences," in the DaVinci Resolve Reference Manual or Chapter 15 in the Fusion Reference Manual for more information on using environment variables.

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

Configuring the Render Master and Render Nodes


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

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

Preparing the Render Master

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

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



The Network panel sets up computers as Render Masters.

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

To set up the Render Master:

- 1 Install a copy of Fusion Studio on the computer you want to be the Render Master.
- 2 In Fusion Studio, choose Fusion Studio > Preferences on macOS or File > Preferences on Windows and Linux
- 3 In the Preferences dialog, select the Global > Network Preferences panel.
- 4 Enter the name of the Render Master in the Name field and enter the IP address.
- 5 Enable the Make This Machine a Render Master checkbox.
- 6 If you want to use this computer as part of the render farm, enable the Allow This Machine to Be Used as a Network Slave checkbox as well.

To have the Render Node act as the Render Manager:

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

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

Preparing Render Nodes

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

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

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

On a macOS Render Node computer:

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



The Render Node menu accessed in the macOS menu bar

On a Linux Render Node computer:


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

On a Windows Render Node computer:

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

Setting Up the Render Manager

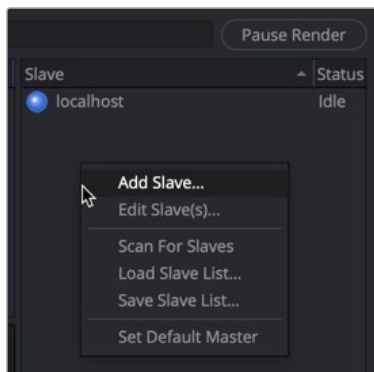
The Render Manager window is used to monitor the progress of rendering. It can be used to reorder, add, or remove compositions from a queue, and to manage the list of Render Nodes used for rendering. To open the Render Manager window in Fusion Studio, choose File > Render Manager.



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

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


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



The Render Manager is used to add Render Nodes.

Scanning for Render Nodes

With the Render Manager open, you can scan for Render Nodes by choosing Slave > Scan for Slaves in the menu bar or by right-clicking in the Render Manager's Slave list and choosing Scan for Slaves from the drop-down menu.

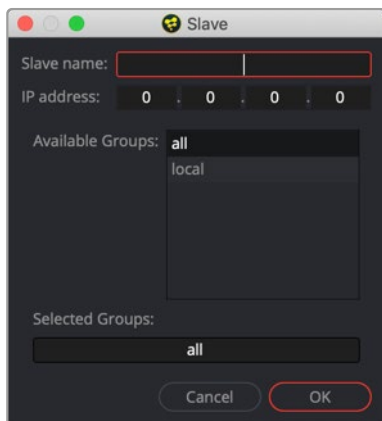


Scanning searches IP addresses on the subnet for active Render Nodes.

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

Manually Adding Render Nodes

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



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

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

Removing Render Nodes

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

Loading and Saving Render Node Lists


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

Submitting Comps to Network Render

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

Using the Render Settings Dialog for Network Rendering

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



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


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

Using the Render Manager Window for Network Rendering

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

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

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



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

Removing a Composition from the Queue

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

Saving and Reloading Queue Lists

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

Reordering the Queue Lists

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

Working with Node Groups

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

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

To assign a Render Node to a Group:

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

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

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

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

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

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

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

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

Using Multiple Node Groups

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

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

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

Viewing the Render Log

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

To view the Render Log in the console:

- Open the Render Manager and choose Misc > Show Render Log.

There are two modes for the Render Log: a Verbose mode and a Brief mode. Verbose mode logs all events from the Render Manager, while Brief mode logs only which frames are assigned to each Render Node and when they are completed.

To disable Verbose mode:

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

Using Third-Party Render Managers with Fusion Studio

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

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

Command Line Rendering

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

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

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

The following table lists additional command line features.

Command	Description
<code>"Fusion Server -i"</code>	Installs the license server as a service or daemon, launching it on startup before user login.
<code>"Fusion Server -S" (capital S)</code>	Causes the Fusion Server to run persistently in the background, until force-quit.
<code><filename.comp></code>	Full path and comp name, like <code>/storage/project/episode/shot/filename.comp</code> .
<code>-render</code>	Tells the Render Node to render.
<code>-frames <frameset></code>	Passes a series of frame ranges to be rendered—e.g., <code>101..110,120,121,130..150</code> .
<code>-start <frame></code>	Sets the start frame of the render.
<code>-end <frame></code>	Sets the last frame of the render.
<code>-step <step></code>	Normally set to 1, step skip frames for rendering. For instance, 2 would render every second frame
<code>-quit</code>	Causes the Render Node to quit after the render is complete.
<code>-join <host></code>	Prompts the node to connect to a manager at <code><hostname,IP></code> , and (re)join any ongoing renders.
<code>-listen</code>	The node remains running and waits for incoming requests from a manager.
<code>-log <filename></code>	Causes the Render Node to output information about the render to a log file. This appends to the end of an existing log file.
<code>-cleanlog</code>	Clears existing text from a log file.
<code>-verbose</code>	Outputs more detailed information into the log file.
<code>-quiet</code>	Suppresses pop-ups and interface buttons from displaying and needing interaction.
<code>-version</code>	Returns the Render Node version number.
<code>-pri high above normal below idle</code>	Sets the node's process priority to high, above, normal, below, or idle.
<code>-args <arg1> [, <arg2> ...]</code>	Allows storing custom values that can be fetched by calling the script function <code>GetArgs()</code> , which will return a table of <code>{ <arg1>, <arg2>, ... }</code>

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

Preparing Compositions for Network Rendering

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

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

Below are more details about some of these items.

Using Relative Paths

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

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

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

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

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

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

And your source media is stored here:

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


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



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

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

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



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

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

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

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

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

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

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

Using Mapped Drives

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

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

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

Installing All Fonts on Render Nodes

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

Installing Third-Party Plug-Ins on Render Nodes

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

Other Uses of Network Rendering


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

Flipbook Previews

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

Disk Cache

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



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

When Renders Fail

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

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

Automatic Rejoining of the Queue

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

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

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

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

Relaunching Render Nodes with Fusion Server

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


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

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

Frame Timeouts

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

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



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

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

Heartbeats

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

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

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

Managing Memory Use

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

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

Override Composition Settings

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

Render Several Frames at Once

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

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

Simultaneous Branching

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

Limitations of Render Nodes

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

Time Stretching

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

Linear Tools

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

Saving to Multi-Frame Formats

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

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

Troubleshooting

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

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

Checking the Render Log

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

Check the Composition

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

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

Chapter 65

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.

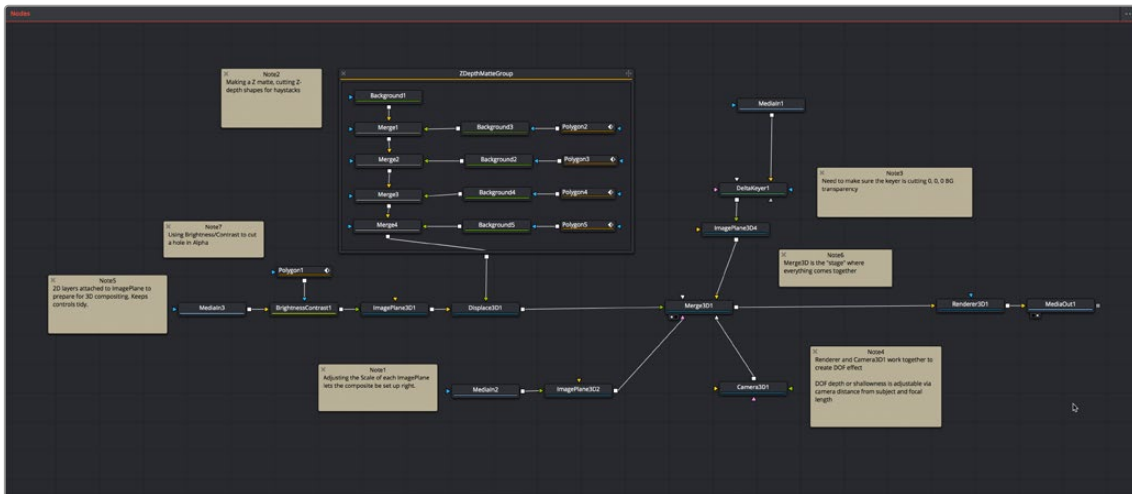
Contents

Learning to Use the Node Editor	1082
Navigating within the Node Editor	1082
Automatic Node Editor Navigation	1083
Using the Node Navigator	1083
Node View Bookmarks	1084
Adding Nodes to a Composition	1085
Adding, Inserting, and Replacing Nodes Using the Toolbar	1086
Adding Nodes Quickly Using the Select Tool Window	1087
Adding Nodes from the Effects Library	1087
Adding, Inserting, and Replacing Nodes Using the Contextual Menu	1090
Deleting Nodes	1090
Disconnected Nodes	1090
Selecting and Deselecting Nodes	1091
Selecting Nodes	1091
The Active Node	1091
Deselecting Nodes	1091
Loading Nodes into Viewers	1092
Viewed Nodes When You First Open Fusion	1092
Node View Indicators	1093
Drag and Drop Nodes into a Viewer	1093
Using the Contextual Menu	1093
Clearing Viewers	1093
Create/Play Preview	1094

Connecting and Disconnecting Nodes	1094
Node Basics	1094
How to Connect Nodes	1094
Dropping Connections on Top of Nodes	1095
Identifying Node Inputs	1097
Node Order Matters	1097
Upstream and Downstream Nodes	1099
Disconnecting and Reconnecting Nodes	1099
Tracing Connections Through the Node Tree	1100
Branching	1100
Connecting Merge Nodes	1101
Automatically Creating a Merge Node When Adding Nodes	1102
Automatically Creating a Merge Node by Connecting Two Outputs	1102
Connection Options and Routers	1103
Using Routers to Reshape and Branch Connections	1104
Swapping Node Inputs	1105
Extracting and Inserting Nodes	1105
Cut, Copy, and Paste Nodes	1106
Cut, Copy, and Paste in the Node Editor	1106
Pasting Node Settings	1107
Copying and Pasting Nodes to and from Any Text Editor	1107
Instancing Nodes	1108
Using Instanced Nodes	1109
De-Instancing and Re-Instancing Specific Parameters	1110
Keeping Node Trees Organized	1110
Moving Nodes	1110
Renaming Nodes	1111
Changing Node Colors	1112
Using Sticky Notes	1112
Using Underlay Boxes	1112
Node Thumbnails	1114
Choosing Which Nodes Show Thumbnails	1114
Switching Thumbnails between Images and Icons	1115
Finding Nodes	1116
Performing Simple Searches	1116
Using Regular Expressions	1116
Custom Node Settings	1117
Managing Saved Settings	1117
Resetting Defaults	1118
Saving and Loading Alternate Node Settings	1118
Adding Saved Settings from the File System	1119
Node Modes Including Disable and Lock	1119
Node Editor Options	1119
Node Tooltips and the Status Bar	1120

Learning to Use the Node Editor

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




The Node Editor.

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

To display the Node Editor:

- Click the Nodes button on the UI toolbar.



The Nodes button in the UI toolbar.

Navigating within the Node Editor

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

Methods of panning the Node Editor:

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

Methods of zooming the Nod4e Editor:


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

Automatic Node Editor Navigation

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

Using the Node Navigator

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



The Node Navigator.

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


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

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

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

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

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



Drag the corner to resize the Navigator.

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

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

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

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

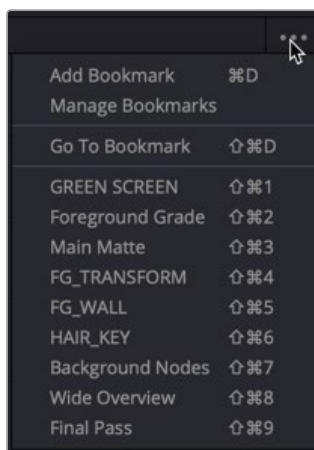
Node View Bookmarks

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

To add a bookmark, do the following:

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

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



The Node Editor Options menu with nine bookmarks added

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

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

To rename or delete a bookmark, do the following:

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

Using Bookmarks

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

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

Changing the Sort Order and Assigning Keyboard Shortcuts

Bookmarks appear in the Options menu and in the Go To Bookmarks dialog in the order they were created. The top nine bookmarks listed are assigned keyboard short cuts. If you want to change the sort order of the list you can do that in the Manage Bookmarks dialog. This can be useful if you want some bookmarks to have specific keyboard shortcuts, or you want a bookmark you made early in the process to not have a keyboard shortcut assigned to it.

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

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

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

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

Using Underlays as Bookmarks

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

Adding Nodes to a Composition

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

Make Sure You're Adding Compatible Nodes

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

Adding, Inserting, and Replacing Nodes Using the Toolbar

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



The Fusion page toolbar.

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

Methods of adding nodes by clicking toolbar buttons:


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

Methods of adding nodes by dragging toolbar buttons:

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

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

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



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

- 2 Click OK in the dialog to confirm the replacement.


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

Adding Nodes Quickly Using the Select Tool Window

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

To use the Select Tool window to add nodes:

- 1 Do one of the following to determine if you want to insert a node or create a disconnected node:
 - If you want to insert a node, select a node that's compatible with the one you'll be creating, and the new node will be inserted after it.
 - If you want to create a disconnected node, deselect all nodes.
- 2 Press Shift-Spacebar to open the Select Tool dialog.
- 3 When the window appears, type characters corresponding to the name of the node you're looking for. A list automatically appears with likely candidates, and you can use the Up and Down arrow keys to select the correct node (if it's not selected already).
- 4 When you've selected the correct node, press the Return key (or click OK), and that node will be either inserted or added.



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

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


Adding Nodes from the Effects Library

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

To open the Effects Library:

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

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



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

By default, the category list shows two primary sets of effects, Tools and Templates, with disclosure controls to the left that hierarchically show all subcategories within each category. The top two categories are:

- **Tools:** Tools consist of all the effects nodes that you use to build compositions, organized by categories such as 3D, Blur, Filter, Mask, Particles, and so on. If you have third-party OFX plug-ins on your workstation, those appear in here as well.
- **Templates:** When using the Fusion page in DaVinci Resolve, templates consist of presets, macros, and utilities that have been created to get you started quickly. For example, Backgrounds consists of a variety of customizable generators that have been created using a combination of Fusion tools. Lens flares presents a wide variety of multi-element lens flares that you can add to any composition. Particles has a selection of pre-made particle systems that you can customize for your own use. Shaders has a variety of materials that you can use as texture maps for 3D text and geometry that you create in Fusion. And there are many, many other categories' worth of useful presets and macros that you can learn from and use in your own projects.

Adding, Inserting, and Replacing Tools Using the Effects Library

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

Methods of adding nodes by clicking in the Effects Library:

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

Methods of adding nodes by dragging from the Effects Library:

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

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

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

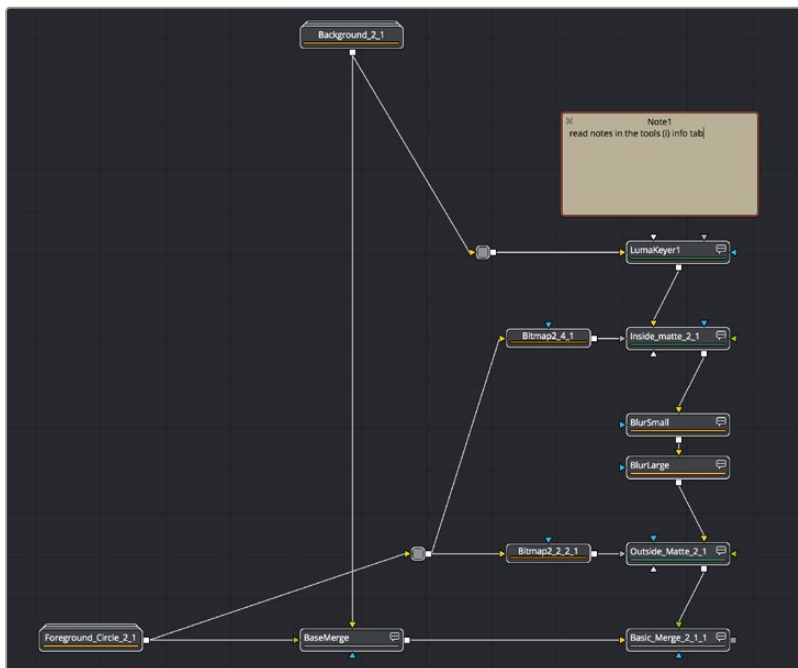
Adding, Inserting, and Replacing Templates Using the Effects Library

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



Adding a Lens Flare effect.

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



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

Adding, Inserting, and Replacing Nodes Using the Contextual Menu

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

Methods of adding nodes using the contextual menu:

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

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

Deleting Nodes

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



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

Disconnected Nodes

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

Selecting and Deselecting Nodes

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

Selecting Nodes

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


Methods of selecting nodes:

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

The Active Node

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

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



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

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

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

Deselecting Nodes

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


Methods of deselecting nodes:

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

Loading Nodes into Viewers

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

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



A node tree for doing a simple rotoscoping job.

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

Viewed Nodes When You First Open Fusion

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

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

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

Node View Indicators

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



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

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

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

For complex compositions, you may need to open additional viewers. For example, one viewer may be used to display the end result of the final comp, while another viewer displays the source, a third viewer displays a mask, and a fourth viewer might be a broadcast monitor connected via a Blackmagic DeckLink card or other display hardware. When you have more than two viewers, additional View indicators are added and each one is assigned a consecutive number between 3 and 9.

The more viewers you add, the more you may need help remembering which viewer is represented by which View indicator. Positioning the pointer over the View indicator in question will display a tooltip with the name of the viewer it represents.

Drag and Drop Nodes into a Viewer

If the View indicators are too small of a target for you to click on reliably and you are not keyboard oriented, another way to load a node into a viewer is to drag and drop it onto the viewer you want to load it into. This offers a quick explicit way to assign a node to a viewer, especially for pen and tablet users. Please note that as you drag, the node will appear to move at first, but it'll snap back into its original location once the pointer leaves the Node Editor.

Using the Contextual Menu

You can also right-click a node, and then choose View On > Left or Right to display the node on the appropriate viewer.

Clearing Viewers

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

Methods of clearing viewers:

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

Create/Play Preview

You can right-click a node, and choose an option from the Create/Preview Play On submenu of the contextual menu to render and play a preview of any node's output on one of the available viewers. The Render Settings dialog is displayed, and after accepting the settings, the tool will be rendered and the resulting frames stored in RAM for fast playback on that view.


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

Connecting and Disconnecting Nodes

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

Node Basics

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



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


Each node takes as its input the output of the node before it. By connecting a MediaIn node's output to a Blur node, you move image data from the MediaIn node to the Blur node, which does something to process the image before the Blur node's output is in turn passed to the next node in the tree.



Two nodes connected together.

How to Connect Nodes

To manually connect one node to another, click on one node's output and drag a connection line out to drop on another node's input. The order in which you drag node connections is not important; you can just as easily drag a connection from one node's input to another node's output and get the same results.



Before (top), and after (bottom) dragging a connection line and dropping it to connect two nodes.

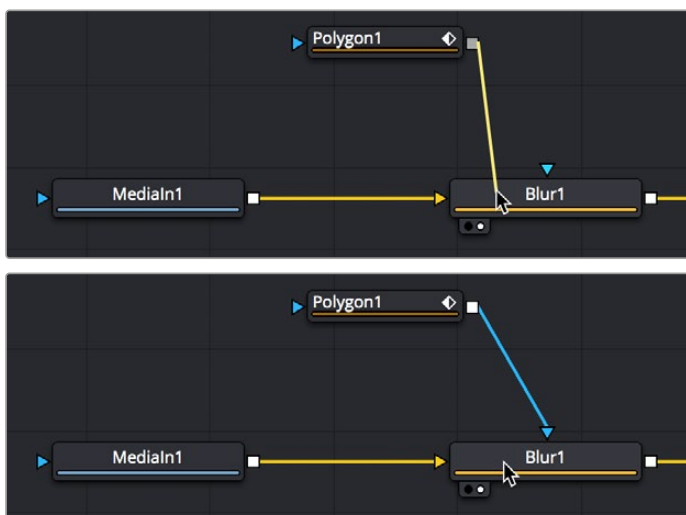
Dropping Connections on Top of Nodes

To make your life a bit easier, you can also drag a connection line and drop it directly on top of the tile of a node to automatically connect to the default input of that node, which is usually labeled “background” or “input.” In the following example, a connection is dragged from the output of a MediaIn node and dropped onto the tile of a Blur1 node, and the background input is connected first.



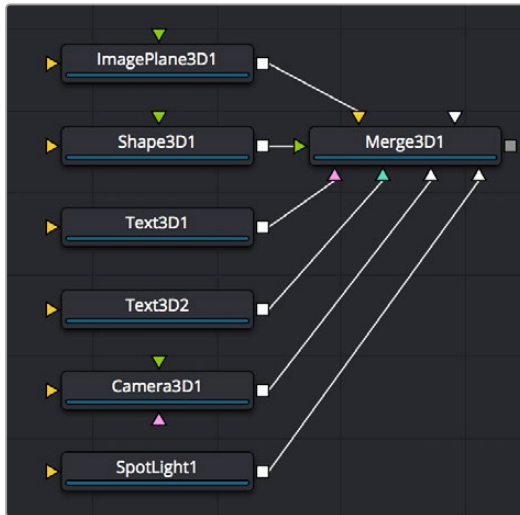
Before (top), and after (bottom) dragging a connection line and dropping it on top of a node.

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



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

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

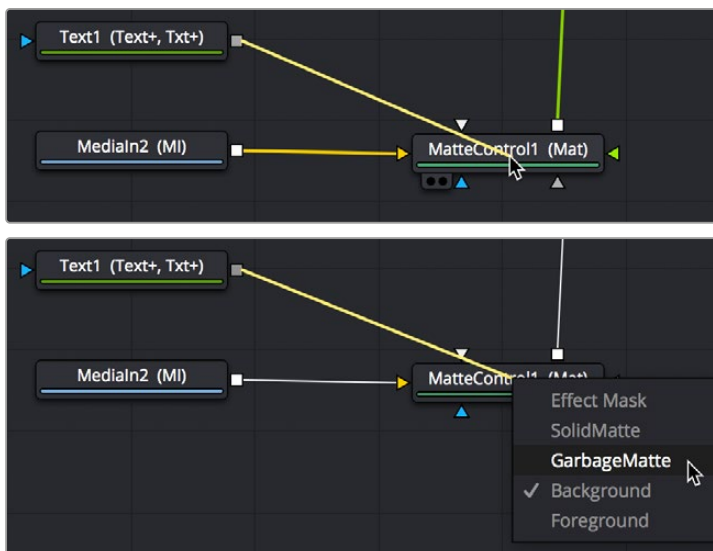


After dragging a connection line and dropping it on top of a Merge3D node.

Attaching Connections to Specific Inputs

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

However, there's an alternate method of connecting nodes together in instances where there are several inputs to choose from and you want to make sure you're choosing the correct one. Hold down the Option key while dragging a connection from one node's output and dropping it onto the body of another node. This opens a pop-up menu from which you can choose the specific input you want to connect to, by name. Please note that this menu only appears after you've dropped the connection on the node and released your pointing device's button.




Option-dragging a node connection to drop onto another node exposes a node input menu.

Automatically and Manually Attaching Mask Nodes

Mask nodes, such as the Polygon, B-Spline, Ellipse, or Rectangle, have a different automatic behavior when you connect them to other nodes. If you drag a connection from a Mask node onto the body of another node, it will automatically connect itself to the default mask input, which is usually the effect mask input. The assumption is that you're using the mask to limit the node's effect somehow. However,

this isn't always the case, so you'll need to be careful of this behavior to make sure you're attaching your mask to the input that will actually create the effect you need.



Before (left) and after (right) dragging a connection from a Mask node and dropping it on top of a MatteControl node.

Identifying Node Inputs


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

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

Node Order Matters

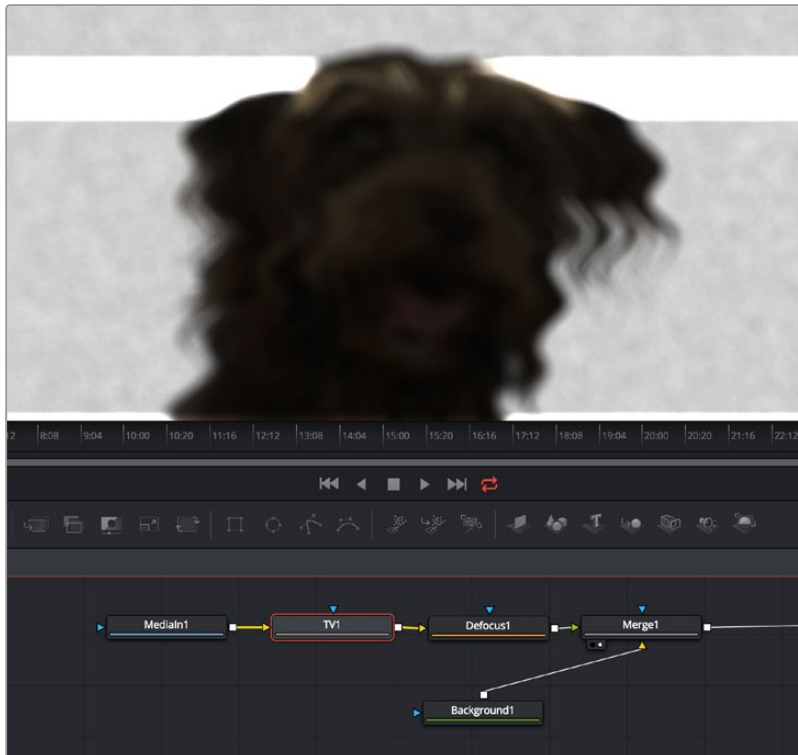
The order in which nodes are attached defines the order in which each image-processing operation is applied to the image.

In the following example, a MediaIn node adds a clip to the composition, while a Defocus node blurs the image, and then a TV node adds scanlines and vertical distortion. Those effect nodes are then connected to the MediaOut node in the Fusion page in DaVinci Resolve or a Saver node in Fusion Studio.



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


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



Adding a TV effect first, and a Defocus second.

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

In the following example, several graphics layers are individually transformed and combined with a series of Merge nodes. The result of the last Merge node is then transformed, allowing you to move the entire collection of previous layers around at once. Because each of these operations is clearly represented via the node tree, it's easy to see everything that's happening, and why.



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

Upstream and Downstream Nodes

Since nodes can be positioned anywhere in the Node Editor, and added in any direction, nodes are referred to as being upstream and downstream of one another. Once you select a node, all other nodes that directly or indirectly connect to its input are considered to be upstream. Any other nodes that are directly or indirectly connected to the output are said to be downstream.

This is an important distinction to make because, unlike layer-based systems, the visual positioning of nodes in your node tree has no bearing on the order of operations in that composition. The only thing that matters is whether nodes are upstream or downstream of each other.



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

TIP: To help you stay organized, there are Select > Upstream/Downstream commands in the Node Editor contextual menu for selecting all upstream or downstream nodes to move them, group them, or perform other organizational tasks.

Disconnecting and Reconnecting Nodes

Node trees are a continuous work in progress, requiring constant revision and rearrangement as you discover new details that need to be finessed, or things that you can do better once the overall composition has taken shape. To facilitate quick changes, each connection between two nodes is divided into two halves: the output half (connected to the upstream node's output) and the input half (connected to the downstream node's input). This can only be seen when you hover the pointer over a connection. The half your pointer is over is highlighted in blue.



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

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

To disconnect two nodes, do one of the following:

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

To overwrite a current connection:


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

To reconnect a connection from one node to another:

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


Tracing Connections Through the Node Tree

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



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


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



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


Branching

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



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


Alternatively, it lets you use one image in several different ways—for example, feeding the RGB to one branch for keying and compositing, while feeding the A channel to the Effects Mask input of another node to limit its effect, or feeding RGB to a tracker to extract motion information.



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

Connecting Merge Nodes

The Merge node is the primary tool available for compositing images together. Each Merge node is capable of combining two inputs to create a third, using standard compositing methods and composite modes. For more extensive information about the Merge node, see Chapter 104, “IO Nodes,” in the DaVinci Resolve Reference Manual or Chapter 44 in the Fusion Reference Manual. For this chapter, all you need to know is that if you attach a background image to the Background input (such as a landscape), and a foreground image with an alpha channel to the Foreground input (such as a graphic with an alpha channel), the Merge node will combine them into a single image for further compositing.



Two Medialn nodes and a DeltaKeyer node attached to a Merge node, creating a composite.

Each Merge node has three inputs:

- **Background (orange):** The default input. Whichever image is connected to this input defines the output resolution of the Merge node.
- **Foreground (green):** The secondary input, meant for whichever image you want to be “on top.”
- **Effect Mask (blue):** An optional input you can use to attach a mask or matte with which to limit the effect of the Merge node.

It’s important to make sure you’re attaching the correct nodes to the correct inputs to ensure you’re getting the result you want, and it’s important to keep these inputs in mind when you connect to a Merge node. Of course, you can always drag a connection to a specific input to make sure you’re connecting things the way you need. However, if you’re in a hurry and you simply drag connections right on top of a Merge node:


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

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

Automatically Creating a Merge Node When Adding Nodes

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

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



Dragging a node from the Media Pool onto a connection (left), and dropping it to create a Merge node composite (right).

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



By dragging three nodes from an OS window to the Node Editor (left), Merge nodes are automatically created to connect them all (right).

Automatically Creating a Merge Node by Connecting Two Outputs


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



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


Connection Options and Routers

By default, the Node Editor uses linear connections that are drawn straight between any two connected nodes. While efficient, this sometimes causes connection lines to overlap nodes, which some people feel interferes with the view of the Node Editor.



Linear connections between nodes.

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



Optional orthogonal connections between nodes.


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

To change how connections are drawn in the Node Editor:

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

Using Routers to Reshape and Branch Connections


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



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

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

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



A router branching its output to multiple nodes.

Methods of using routers:

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

Swapping Node Inputs

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

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

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



Before swapping node inputs (left), and after swapping node inputs (right), the connections don't move but the colors change.


Inputs can move freely around the node, so swapping two inputs doesn't move the connection lines; instead, the inputs change color to indicate you've reversed the background (orange) and foreground (green) connections.

Extracting and Inserting Nodes

When building a composition, you'll often find that you need to rearrange nodes that you've already added, in order to connect them in different ways to obtain a better result. Happily, this is easy to do by extracting one or more nodes from one part of a node tree and inserting them at another part of the node tree.

To extract one or more nodes from their position in the node tree:

- **To extract a single node:** Hold down the Shift key, drag a node from the node tree up or down to disconnect it, and then drop the node before releasing the Shift key. That node is now detached, and the output of the next upstream node is automatically connected to the input of the next downstream node to fill the gap in the node tree.
- **To extract multiple nodes:** Select the nodes you want to extract, hold down the Shift key, drag one of the selected nodes up or down to disconnect them, and then drop the node before releasing the Shift key. Those nodes are now detached (although they remain connected to one another), and the output of the next upstream node is automatically connected to the input of the next downstream node to fill the gap in the node tree.




Before extracting a pair of nodes (left), and after extracting a pair of nodes (right).

After you've extracted a node, you can re-insert it into another connection somewhere else. You can only insert one node at a time.

To insert a disconnected node in the Node Editor between two compatible nodes:

- 1 Hold down the Shift key and drag a disconnected node directly over a connection between two other nodes.
- 2 Once the connection highlights, drop the node, and then release the Shift key. That node is now attached to the nodes coming before and after it.



Before inserting a node (left), and after inserting a node (right).

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

Cut, Copy, and Paste Nodes

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

Cut, Copy, and Paste in the Node Editor

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

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

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

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

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

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

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

- **To paste nodes to be inserted after another node:** Select the node in the node tree you want to insert the pasted node(s) to, and choose Edit > Paste (Command-V).
- **To paste nodes to be disconnected from the rest of the node tree:** Deselect all nodes, and then choose Edit > Paste (Command-V), or right-click anywhere in the Node Editor and choose Paste from the contextual menu.
- **To paste disconnected nodes in a specific area of the Node Editor:** Deselect all nodes, and then click the place in the Node Editor where you want pasted node(s) to appear, and choose Edit > Paste (Command-V), or right-click anywhere in the Node Editor and choose Paste from the contextual menu.
- **To paste a node to replace an existing node in the Node Editor:** Right-click a node in the Node Editor that you want to replace, choose Paste from the contextual menu, and when a dialog appears asking if you want to replace that node, click OK. This only works when you use the contextual menu command.

TIP: When you paste a MediaIn, Loader, or Generator node so it will be inserted after a selected node in the node tree, a Merge tool is automatically created and used to composite the pasted node by connecting it to the foreground input. While this can save you a few steps, some artists may prefer to perform these sorts of merges manually, so this can be changed using the 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.

Note that you can paste settings between two nodes of the same type, or between two entirely different kinds of nodes that happen to have one or more of the same parameters in the Inspector. When copying settings from one type of node to another, only the settings that match between two nodes will be copied. A common example is to copy an animated Center parameter from a Transform node to the Center parameter of a Mask node.

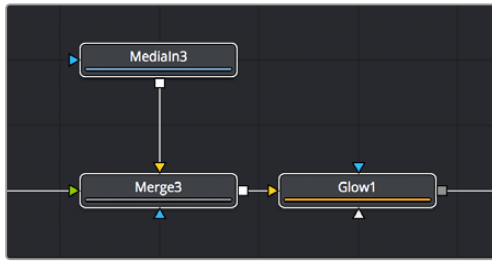
To Paste settings from one node to another:

- 1 Select a node that has settings you want to copy, and choose Copy from the Edit menu (Command-C).
- 2 Right-click a node you want to paste those settings to, and choose Paste Settings from the contextual menu.

Copying and Pasting Nodes to and from Any Text Editor

The format of nodes in the Node Editor is not binary, but is in fact a simple text format. The implications of that may not be obvious, but one example benefit is clear when you start dealing with nodes.

One or more nodes can be copied from the Node Editor and pasted directly into a text editor or email. This pastes the selection in text format, just as it's saved internally in Fusion. For example, if you copy the following set of three nodes:



A set of three nodes being copied.

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

```

{
  Tools = ordered() {
    MediaIn3 = MediaIn {
      ExtentSet = true,
      CustomData = {
        MediaProps = {
          MEDIA_HEIGHT = 1080,
          MEDIA_MARK_IN = 0,
          MEDIA_MARK_OUT = 244,
          MEDIA_NAME = "060030038-aerial-abstract-waves-crystal_PRORESHD1080.mov",
          MEDIA_NUM_FRAMES = 245,
          MEDIA_NUM_LAYERS = 1,
          MEDIA_PAN = 1,
          MEDIA_PATH = "/Volumes/Media Raid Too/DaVinci Media/Fusion Manual Example Media/Pond5 Media/Textures/060030038-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-be00-52f29d4816d5", },
        Layer = Input { Value = "", },
        ClipTimeEnd = Input { Value = "244", },
      },
      ViewInfo = OperatorInfo { Pos = { 825, 115.5 } },
    },
    Merge3 = Merge {
      Inputs = {
        Background = Input {
          SourceOp = "MediaIn3",
          Source = "Output",
        },
      },
      PerformDepthMerge = Input { Value = 0, },
      ViewInfo = OperatorInfo { Pos = { 825, 181.5 } },
    },
    Glow1 = Glow {
      CtrlWZoom = false,
      Inputs = {
        Blend = Input { Value = 0.2, },
      },
      ViewInfo = OperatorInfo { Pos = { 935, 181.5 } },
    },
  }
}

```

The same three nodes pasted into a text editor.

At this point, you have the option of editing the text (if you know what you're doing), emailing it to colleagues, or storing it in a digital notepad of some sort for future use. To use this script in Fusion again, you need only copy it and paste it back into the Node Editor.

TIP: This is a very easy way to pass specific node settings back and forth between artists who may not be in the same room, city, or country.

Instancing Nodes

Normally, when you use copy and paste to create a duplicate of a node, the new node is completely independent from the original node, so that changes made to one aren't rippled to the other. However, there are times when two nodes must have identical settings at all times. For example, when you're making identical color corrections to two or more images, you don't want to constantly have to adjust one color correction node and then manually adjust the other to match. It's a hassle, and you risk forgetting to keep them in sync if you're working in a hurry.

While there are ways to publish controls in one node and connect them to matching controls in another node, this becomes prohibitively complex and time consuming for nodes in which you're making adjustments to several controls. In these cases, creating "instanced" nodes is a real time-saver, as well as an obvious visual cue in your node tree as to what's going on.

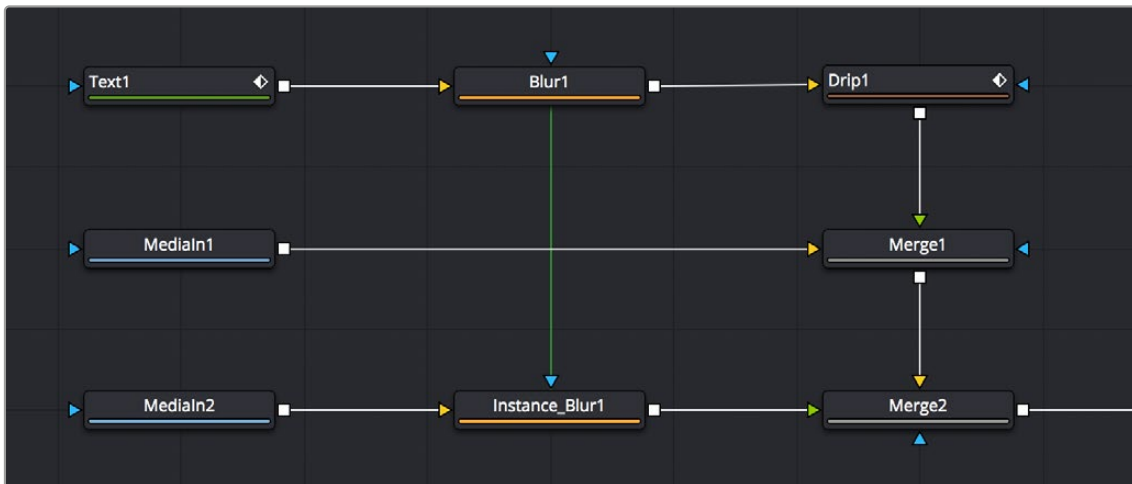
Using Instanced Nodes

Instanced nodes are nodes that have been created using the Paste Instance command, and which share settings with the original node so that a change made to one instanced node is also automatically applied to all other instances of that node (as well as the original node you copied).

To create an Instance, do the following:

- 1 Select a node you want to instance, and copy it (Command-C).
- 2 Do one of the following:
 - **To create a disconnected instance of a node:** Right-click in the background of the Node Editor, and choose Paste Instance from the contextual menu (Command-Shift-V).
 - **To insert an instanced node between two other nodes:** Select a node that's upstream of where you want to insert the instanced node, and press Command-Shift-V. Alternatively, you can right-click directly on a connection line, and choose Paste Instance from the contextual menu.

However you paste an instance, the name of that instanced node takes the form "Instance_ NameOfNode." If you paste multiple instances, each instance is numbered "Instance_ NameOfNode_01."



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

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

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

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

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

To de-instance a node, making it independent:

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

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

De-Instancing and Re-Instancing Specific Parameters

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

To de-instance a single parameter:

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

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

To re-instance a single parameter:

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

Keeping Node Trees Organized

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

Moving Nodes

Selecting one or more nodes and dragging them moves them to a new location, which is one of the simplest ways of organizing a node tree, by grouping nodes spatially according to the role they play in the overall composition.

Keep in mind that the location of nodes in the Node Editor is purely aesthetic, and does nothing to impact the output of a composition. Node tree organization is purely for your own peace of mind, as well as that of your collaborators.

TIP: Once you've arranged the nodes in a composition in some rational way, you can use the Sticky Note and Underlay tools to add information about what's going on and to visually associate collections of nodes more definitively. These tools are covered later in this section.

Snapping Nodes to the Grid

By default, you can position nodes freely wherever you want them to be. However, keeping nodes and connection lines straight and aligned can make them easier to read. To help keep them aligned, you can have nodes you're dragging automatically snap to the grid.

To have nodes snap to the grid as they're dragged:

- Right-click over an empty area of the Node Editor, and choose Arrange Tools > To Grid from the contextual menu. All nodes you drag now snap to the nearest grid coordinate.

- Right-click over an empty area of the Node Editor, and choose Arrange Tools > To Connected from the contextual menu. All nodes you drag now snap to the horizontal or vertical position of the nodes they're attached to.

TIP: You can set “Arrange to Grid” or “Arrange to Connected” as the default for new compositions by choosing Fusion > Fusion Settings in DaVinci Resolve or File > Preferences in Fusion Studio, and turning the Fusion > Node Editor > Arrange To Grid or Arrange to Connected checkboxes on.

Commands to “Clean Up” a Node Tree

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

To “clean up” an unruly node tree:

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

To “clean up” only one or more selected nodes:

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

Renaming Nodes

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

To rename a node:

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

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

Since Fusion can be scripted and use expressions, the names of nodes must adhere to a scriptable syntax. Only use alphanumeric characters (no special characters), and do not use any spaces.

Also, you cannot start a node name with a number. If you accidentally create a name that doesn't exactly follow the guidelines, spaces and invalid characters will be automatically deleted.

If you want to see the original node types instead of the node names, press and hold Command-Shift-E.

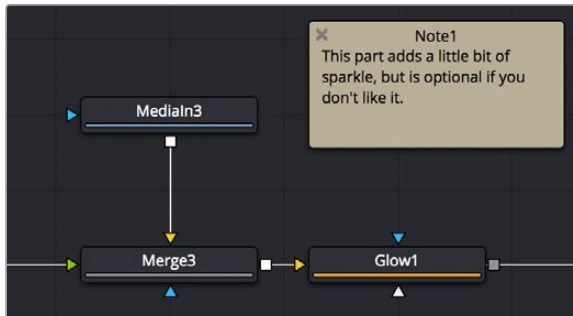
Changing Node Colors

You can change the color of any node by selecting it, opening the Inspector, and choosing a new color from the Node Color pop-up in the Inspector header for that node. Alternatively, you can right-click a node and choose a color from the Set Color submenu.

To return a node to its regular color, right-click it and choose Set Color > Clear Color from the contextual menu, or open the Node Color pop-up for a node in the Inspector, and choose Clear Color.

Using Sticky Notes

A good way to add notes about different parts of a composition, client feedback about various details, and other information you want to keep track of, is to add Sticky Notes to the Node Editor.



A Sticky Note in the Node Editor.

Sticky Notes are yellow boxes in which you can type whatever text you want. They can be resized, moved, and collapsed when they're not being edited, but once created they remain attached to the background of the Node Editor where you placed them until you either move them or delete them.

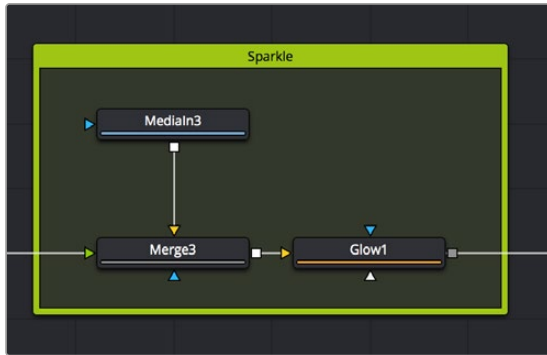
Methods of working with Sticky Notes:

- **To create a Sticky Note:** Click somewhere in the Node Editor where you want a Sticky Note to appear. Then, press Shift-Spacebar, type sticky, and press the Return key when the Sticky Note appears in the Select Tool window. Alternatively, you can open the Effects Library, open the Tools > Node Editor category, and click or drag the Sticky Notes node to create a new one.
- **To open a Sticky Note to full size:** Double-click a minimized Sticky Note and it expands to a larger, resizable yellow box.
- **To edit a Sticky Note:** If necessary, double-click a Sticky Note to open it to full size, and then click once in the body of the note to place a text cursor. You can edit text within the Sticky Note just like any other text editor.
- **To rename a Sticky Note:** Right-click a Sticky Note, choose Rename, type a new name into the Rename dialog, and click OK. Alternatively, you can select a Sticky Note, press F2 to open the Rename dialog, and press Return to close it when you're done.
- **To resize a Sticky Note:** Double-click a Sticky Note to open it to full size, and then drag any of the edges or corners to make it larger or smaller.
- **To minimize a Sticky Note:** Click the close box at the upper left-hand corner of the Sticky Note, and it collapses to a small tile.
- **To delete a Sticky Note:** Right-click any Sticky Note and choose Delete from the contextual menu or select the Sticky note in the Node Editor and press the Delete key.

Using Underlay Boxes

Underlay Boxes are a good way of associating a collection of nodes that work together to perform a specific task in your composition. They're simply colored rectangles that you can put nodes inside of.

Once you place nodes inside an Underlay, you can move the Underlay, and all the nodes within move along with it.



An Underlay in the Node Editor.

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


Methods of working with Underlay Boxes:

- **To create an Underlay Box:** Click somewhere in the Node Editor where you want the Underlay Box to appear. Then, press Shift-Spacebar, type under, and press the Return key when the Underlay Box appears in the Select Tool window. Alternatively, you can open the Effects Library, open the Tools > Node Editor category, and click or drag the Underlay Box node to create a new one.
- **To create an Underlay Box around specific nodes:** Select the nodes in the Node Editor that you want surrounded by an Underlay Box. Then, press Shift-Spacebar, type under, and press the Return key when the Underlay Box appears in the Select Tool window. Alternatively, you can open the Effects Library, open the Tools > Node Editor category, and click the Underlay Box node to have it added and sized to encompass all the selected nodes.
- **To resize an Underlay Box:** Drag any of the edges or corners to make it larger or smaller.
- **To rename an Underlay Box:** Option-click the Underlay Box to select just the box and not the contents, and then right-click it and choose Rename (or press F2). Type a new name into the Rename dialog and click OK or press Return.
- **To change the color of an Underlay Box:** Option-click the Underlay Box to select just the box and not the contents, and then right-click it and choose a color from the Set Color submenu.
- **To put nodes inside of an Underlay Box:** Select the nodes you want to place inside an Underlay Box, and then drag them to fit inside. The Underlay Box must be big enough to fit all the nodes. Alternatively, you can place an Underlay Box near a collection of nodes you want to put inside it, and then resize the Underlay Box to encompass all those nodes.
- **To move an Underlay Box and all its nodes:** Once nodes have been placed inside an Underlay Box and have been deselected, you can move the entire collection of nodes together by dragging the Underlay Box by its title bar.
- **To remove nodes from an Underlay Box:** There are two ways you can remove nodes from an Underlay Box.
 - With both the Underlay Box and nodes deselected, drag a bounding box or Command-click to select all nodes in the box you want to remove, and drag them out.
 - Resize the Underlay Box so that it's smaller than the collection of nodes it originally encompassed. Once an Underlay Box is so small that even the last node sticks out beyond its edge, those nodes are automatically removed from the Underlay Box, and you can move or delete the Underlay Box without moving those nodes.

- **To delete an Underlay Box and all nodes within:** Select an Underlay Box and press the Delete key to delete both the Underlay Box and all nodes found inside it. If you don't also want to delete the nodes, first drag the nodes out of the box.
- **To delete an Underlay Box but keep all nodes within:** Option-click the Underlay Box to select it and not the nodes, and then press the Delete key. The nodes within remain where they were.

Node Thumbnails

Once a source or an effect has been added to the Node Editor, it's represented by a node. By default, nodes are rectangular and thin, making it easier to fit reasonably complicated grades within a relatively small area. However, if you like, you can also display node thumbnails.



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

Nodes can be displayed as a small rectangle or as a larger square. The rectangular form displays the node's name in the center, while the square form shows either the tool's icon or a thumbnail of the image it is outputting.

TIP: Even if you're not displaying node thumbnails, you can quickly obtain detailed information about a node and the data it's processing by hovering your pointer over it in the Node Editor and viewing the tooltip bar below.

Choosing Which Nodes Show Thumbnails

If you want to use node thumbnails to help visually identify media and operations in your node trees, there are a variety of options for which nodes should display thumbnails in the contextual menu that appears when you right-click anywhere in the background of the Node Editor.

Force All Tile Pictures

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

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

Force Active Tile Pictures

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

Force Source Tile Pictures

This enables thumbnails for all MediaIn and Loader nodes in the Node Editor, as well as all generators, and is a great way to be able to quickly see where all the clips are in a composition.

Force Mask Tile Pictures

This enables thumbnails for all Mask nodes in a composition, which can make them easier to distinguish when you're building complex shapes made from multiple Mask nodes.

Manually Showing Tile Pictures and Node Options

You also have the option of manually choosing which nodes you'd like to show thumbnails.

For example, there may be certain key points of the node tree where you'd like to see a small visual representation of what's happening in the composition.

To toggle thumbnails for one or more specific nodes:

- 1 Select one or more nodes in the Node Editor.
- 2 Right-click one of the selected nodes, and choose one of the following from the contextual menu:
 - Show > Show Tile Pictures
 - Show > Show Modes/Options

When you've manually enabled thumbnails for different nodes, they'll remain visible whether or not those nodes are selected.

Switching Thumbnails between Images and Icons

Whenever you enable node thumbnails, you have the choice of having these thumbnails either display an image of the state of the image at that node, or you can instead choose to display the icon for that particular node. The setting for this affects all nodes at once.

To display icons instead of thumbnails:

- Right-click anywhere in the background of the Node Editor and deselect Show Thumbnails in the contextual menu.

Sometimes Nodes Only Show Icons

As you add more and more nodes to a composition, you'll notice that some nodes never display an image in their thumbnail. In these cases, the default icon for that node is displayed instead of an image.


Most nodes in the Particle and 3D categories fall into this group. The exceptions are the pRender node and the Render 3D node. These two nodes are capable of displaying a rendered thumbnail if you have the menu options set for Thumbnails to be displayed.

In other cases, whether nodes display images in their thumbnail is more situational. Some Transform nodes are able to concatenate their results with one another, passing the actual processing downstream to another node later in the node tree. In this case, upstream Transform nodes don't actually process the image, so they don't produce a thumbnail.

In other situations where the Loader is not reading in a clip or the clip is trimmed in the Keyframes Editor to be out of range, it can cause the node not to process the image, so it will not produce a rendered Thumbnail image. Also, nodes that have been set to Pass Through mode are disabled and do not display a rendered Thumbnail image.

Finding Nodes

Modern visual effects require detailed work that often results in compositions with hundreds of nodes. For such large node trees, finding things visually would have you panning around the Node Editor for a long, long time. Happily, you can quickly locate nodes in the Node Editor using the Find dialog.



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

Performing Simple Searches

To do simple searches using node names is easy.

To search for a node in the Node Editor:

- 1 Press Command-F, or right-click in an empty area of the Node Editor and choose Find from the contextual menu.
- 2 When the Find dialog appears, do the following:
 - Enter a search term in the Find field.
 - Choose search options, such as whether to match the whole phrase in the Find field, whether to match the case, whether to use a sequence number, or whether to use a regular expression in the Find field.
 - Choose what to search. Options include tool name, tool type name, or tool type ID.
- 3 To perform the find, do one of the following:
 - Click Find Next to try to select a downstream node matching the criteria.
 - Click Find Previous to try to select an upstream node matching the criteria.
 - Click Find All to try to select all nodes in the Node Editor that match the criteria.

The Find window closes. If either the Find Next, Find Previous, or Find All operations are successful, the found node or nodes are selected. If not, a dialog appears letting you know that the string could not be found.

TIP: Finding all the nodes of a particular type can be very useful if you want, for example, to disable all Resize nodes. Find All will select all the nodes based on the search term, and you can temporarily disable them by pressing the shortcut for Bypass, Command-P.

Using Regular Expressions

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

Character Sets

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

[a-z]

Finds: Every node using a lower caps letter

[a-d]

Finds: Every lower caps letter from a to d, and will find nodes with a, b, c, or d

[Tt]

Finds: Every node with an upper case T or a lower case t

[aeiou]

Finds: Every vowel

[0-9]

Finds: Every numeral

[5-7]

Finds: Every numeral from five to seven, and will find nodes numbered with 5, 6, or 7

Custom Node Settings

When a node is added to the Node Editor, its parameters are set to the default values for that type of node. If you find yourself constantly readjusting the parameters of a node to a preferred starting point as soon as it's added to the node tree, you can override the default node settings with your own custom settings.

To save new default settings for a particular type of node:

- 1 Create a new node.
- 2 Open the Inspector and customize that node's settings to the new defaults you want it to have.
- 3 Right-click that node in the Node Editor, or right-click that node's control header in the Inspector, and choose Settings > Save Default from the contextual menu.

TIP: You can also save six different settings for a node in the Node Editor using the Version buttons at the top of the Inspector. For more information, see Chapter 68, "Editing Parameters in the Inspector" in the DaVinci Resolve Reference Manual or Chapter 8 in the Fusion Reference Manual.

Managing Saved Settings

Custom node default settings are saved to a folder on your hard drive that's based on the Path Map > Defaults preference in the Fusion Settings. This path is customizable for facilities where multiple compositing artists use a common set of facility defaults, stored somewhere commonly accessible. The default paths are:

- **For macOS systems, this path defaults to:** /UserName/Library/Application Support/Blackmagic Design/DaVinci Resolve/Fusion/Defaults.
- **For Windows systems, this path defaults to C:** \Users\\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 Inspector, those values can also be saved as an alternate setting for that node, which can be reused at a later time.

To save alternate settings for a node:

- 1 Right-click on a tool, and then choose Settings > Save As from the contextual menu.
- 2 When the Save File dialog appears, enter a name for the Setting and save it to your hard drive. Unlike saved defaults, the .settings files can be saved anywhere on the file system. They do not need to be in the Default Settings folder.

To load a saved setting for one or more nodes:

- 1 Right-click a node and choose Settings > Load from the contextual menu.
- 2 Use the Open File dialog to select the settings you want to load into that node, and then click Open. Those settings are now applied to that node.

Adding Saved Settings from the File System

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

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

Node Modes Including Disable and Lock

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

- **Show Controls:** Sets whether that node reveals its parameters in the Inspector when it's selected and whether its onscreen controls appear in viewers. On by default.
- **Pass Through:** (Command-P) Identical to the toggle switch in the Inspector that turns nodes off and on. Disabled nodes are ignored as image data is passed from the next previous upstream node to the next downstream node. On by default.
- **Locked:** (Command-L) Identical to the lock button in the Inspector that prevents a node from being edited in the Inspector. Off by default.
- **Update:** (Command-U) On by default. While this option is enabled, all changes to the node will cause it to re-render. When Update is disabled, you can still change the node's parameters, but those changes will not process or update the image until Update is re-enabled. While disabled, the last processed image for that node will be displayed as a freeze frame. One example of when this is useful is when you have a large or processor-intensive composition (such as a particularly intense particle system), and disabling this option temporarily will let you quickly make several quick parameter adjustments to different nodes without forcing you to wait for the node tree to re-render after every adjustment. Another example of when this is useful is when you want to quickly see the effect of animated downstream nodes while keeping upstream nodes that are too processor-intensive to play in real time from rendering additional frames.
- **Force Cache:** When enabled, this node's output for the current frame has an extremely high cache priority, essentially forcing it to stay cached in memory. Off by default.

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

Node Editor Options

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

- **Pipes Always Visible:** Enabling this option causes a connection to cross over a node instead of beneath it, sometimes making it easier to follow the connection's path.
- **Show Hidden Connections:** When enabled, the Inspector option to Hide Incoming Connections in every node is overridden and all connections are displayed in the Node Editor.
- **Aspect Correct Tile Pictures:** Aspect Correct Tile Pictures forces the display of thumbnails to be aspect corrected, which is slower but visually more accurate. This option is enabled by default.

- **Full Tile Render Indicators:** Enabling this option causes the thumbnail to flash green when rendering, which makes it easier to identify which node is processing in a large, complex node tree.
- **Show Grid:** This option can be used to enable or disable the Node Editor’s background grid.
- **Show Instance Links:** When enabled, the Node Editor draws a green connection between an instanced node and its parent.
- **Auto Remove Routers:** If routers are disconnected from a tool, they are automatically deleted from the Node Editor. This option is enabled by default to eliminate the need to delete orphaned routers.
- **Show Navigator:** Enabling this option displays a small overview window of the entire node tree in the Node Editor’s top-right corner. For more information, see the Navigator section in this chapter.
- **Auto Navigator:** The Navigator only appears when one or more nodes is outside the visible area of the Node Editor. For more information, see the Navigator section in this chapter.
- **Build Flow Vertically/Horizontally:** Node trees can be built either horizontally from left to right or vertically from top to bottom. Enabling one of these options determines whether new nodes are added beneath the current node or to the right of the current tool.
- **Orthogonal/Direct Pipes:** Use these two options to decide whether connections between nodes are drawn as Direct (straight) lines or Orthogonal (bent) lines.


Node Tooltips and the Status Bar

Even in simple node trees, it’s easy to forget some essential detail about the nodes in your comp. To help you figure out what everything’s for, you can hover the pointer over any node in the Node Editor to display information in the Status bar at the bottom of the Node Editor consisting of that node’s name, frame size, pixel aspect, resolution, and color depth.

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

The Status bar located beneath the Node Editor.

If you wait a few moments later, a more elaborate presentation of the same information appears within a floating tooltip in the Inspector. This tooltip gives you additional information about the Domain (Image and DoD) and the data range used by that clip.



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

Chapter 66

Node Groups, Macros, and Fusion Templates

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

Contents

Groups	1122
Creating Groups	1122
Deleting Groups	1123
Expanding and Collapsing Groups	1123
Panning and Scaling within Open Group Windows	1123
Ungrouping Nodes	1123
Saving and Reusing Groups	1123
Macros	1124
Creating Macros	1124
Using Macros	1126
Re-Editing Macros	1126
Other Macro Examples	1126
Creating Fusion Templates	1126
Getting Started with a Fusion Title Template	1126
Saving a Title Macro	1127
Using Your New Title Template	1130
Getting Started with a Fusion Transition Template	1131
Creating a Fusion Transition Template	1132

Using Your New Transition Template	1134
Getting Started with a Fusion Generator Template	1134
Creating a Fusion Generator Template	1135
Using Your New Generator Template	1137
Creating a Fusion Effect Template	1137
Changing Durations of a Template	1138
Creating a Custom Template Icon	1139
Using Fusion Template Bundles	1139

Groups


When you work on complex visual effects, node trees can become sprawling and unwieldy, so grouping tools together can help you better organize all the nodes and connections. Groups are containers in your node tree that can hold multiple nodes, similar to the way a folder on your Desktop holds multiple files. There is no limit to the number of nodes that can be contained within a group, and you can even create subgroups within a group.

Creating Groups

Creating a group is as simple as selecting the nodes you want to group together and using the Group command.

To create a group:

- 1 Select the nodes you want grouped together.
- 2 Right-click one of the selected nodes and choose Group from the contextual menu (Command-G).



Several nodes selected in preparation for making a group (left), and the resulting group (right).


The selected nodes are collapsed into a group, which is displayed as a single node in the Node Editor. The Group node can have inputs and outputs, depending on the connections of the nodes within the group. The Group node only displays inputs for nodes that are already connected to nodes outside the group. Unconnected inputs inside the group will not have an Input knot displayed on the Group node.

Deleting Groups

Deleting a group is no different from deleting any other node in the Node Editor. Select a group and press Delete, Backspace, or Forward-Delete, and the group along with all nodes contained within it are removed from the node tree.

Expanding and Collapsing Groups

A collapsed group is represented by a single “stack” node in the node tree. If you want to modify any of the nodes inside the group, you can open the group by double-clicking it or by selecting the group node and pressing Command-E.



An open group window showing the minimize button.

When you open a group, a floating window shows the nodes within that group. This floating window is its own Node Editor that can be resized, zoomed, and panned independently of the main Node Editor. Within the group window, you can select and adjust any node you want to, and even add, insert, and delete nodes while it is open. When you’re ready to collapse the group again, click the minimize button at the top left corner of the floating window, or use the keyboard shortcut (Cmd-E).

Panning and Scaling within Open Group Windows

You can pan and scale an open group window using the same mouse buttons you use to pan and scale the main Node Editor. However, when you’re working in an expanded group and simultaneously making changes to the main node tree, you may want to prevent the expanded group from being individually panned or scaled. Turning off the Position button at the right of the group title bar locks the group nodes to the size of the nodes in the rest of the overall node tree. Turning on this Position button lets you size group nodes independently of the rest of the node tree.

Ungrouping Nodes

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

To ungroup nodes, do the following:

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

Saving and Reusing Groups

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

A good example of when you might want to Save and Load a group is in a studio with two or more compositing artists. A lead artist in your studio can set up the master comp and create a group specifically for keying greenscreen. That key group can then be passed to another artist who refines

the key, builds the mattes, and cleans up the clips. The setting can then be saved out and loaded back into the master comp. As versions are improved, these settings can be reloaded, updating the master comp.

Methods of saving and reusing groups:

- **To save a group:** Right-click a group and choose Settings > Save As from the contextual menu.
- **To reuse a group:** Drag it from your computer's file browser directly into the Node Editor. This creates a new group node in the node tree with all the same nodes as the group you saved.
- **To load the settings from a saved group to another group with the same nodes:** Right-click a group in the Node Editor and choose Settings > Load from the contextual menu.

In Fusion Studio, you can also save and reuse groups from the Bins window:

- **To save a group:** Drag the group from the Node Editor into the opened Bin window. A dialog will appear to name the group setting file and the location where it should be saved on disk. The .settings file will be saved in the specified location and placed in the bins for easy access in the future.

Macros

Some effects aren't built with one tool, but with an entire series of operations, sometimes in complex branches with interconnected parameter controls. Fusion provides many individual effects nodes for you to work with but gives users the ability to repackage them in different combinations as self-contained "bundles" that are either macros or groups. These "bundles" have several advantages:

- They reduce visual clutter in your node tree.
- They ensure proper user interaction by allowing you to restrict which controls from each node of the macro are available to the user.
- They improve productivity by allowing artists to quickly leverage solutions to common compositing challenges and creative adjustments that have already been built and saved.

Macros and groups are functionally similar, but they differ slightly in how they're created and presented to the user. Groups can be thought of as a quick way of organizing a composition by reducing the visual complexity of a node tree. Macros, on the other hand, take longer to create because of how customizable they are, but they're easier to reuse in other comps.

Creating Macros

While macros let you save complex functions for future use in very customized ways, they're actually pretty easy to create.


To make a macro from nodes in the Node Editor:

- 1 Select the nodes you want to include in the macro you're creating. Because the macro you're creating will be for a specific purpose, the nodes you select should be connected together to produce a particular output from a specific set of inputs.

TIP: If you want to control the order in which each node's controls will appear in the macro you're creating, Command-click each node in the order in which you want it to appear.

- 2 Right-click one of the selected nodes and choose Macro > Create Macro from the contextual menu.

A Macro Editor window appears, showing each node you selected as a list, in the order in which each node was selected.



The macro editor with a Blur node and Color Corrector node.

- 3 First, enter a name for the macro in the field at the top of the Macro Editor. This name should be short but descriptive of the macro's purpose. No spaces are allowed, and you should avoid special characters.
- 4 Next, open the disclosure control to the left of each node that has controls you want to expose to the user and click the checkbox to the right of each node output, node input, and node control that you want to expose.

The controls you check will be exposed to users in the order in which they appear in this list, so you can see how controlling the order in which you select nodes in Step 1, before you start editing your macro, is useful. Additionally, the inputs and outputs that were connected in your node tree are already checked, so if you like these becoming the inputs and outputs of the macro you're creating, that part is done for you.

For each control's checkbox that you turn on, a series of fields to the left of that control's row lets you edit the default value of that control as well as the minimum and maximum values that control will initially allow.

- 5 When you're finished choosing controls, click Close.
- 6 A dialog prompts you to save the macro. Click Yes.
- 7 A Save Macro As dialog appears in which you can re-edit the Macro Name (if necessary), and choose a location for your macro.

To have a macro appear in the Fusion page Effects Library Tools > Macros category, save it in the following locations:

- **On macOS:** Macintosh HD/Users/username/Library/Application Support/Blackmagic Design/DaVinci Resolve/Fusion/Macros/
- **On Windows:** C:\Users\username\AppData\Roaming\Blackmagic Design\DaVinci Resolve\Support\Fusion\Macros
- **On Linux:** home/username/.local/share/DaVinciResolve/Fusion/Macros

To have a macro appear in the Fusion Studio Effects Library Tools > Macros category, save it in the following locations:

- **On macOS:** Macintosh HD/Users/username/Library/Application Support/Blackmagic Design/Fusion/Macros/
- **On Windows:** C:\Users\username\AppData\Roaming\Blackmagic Design\Fusion\Macros
- **On Linux:** home/username/.fusion/BlackmagicDesign/Fusion/Macros

8 When you're done, click Save.

Using Macros

Macros can be added to a node tree using the Add Tool > Macros or Replace Tool > Macros submenus of the Node Editor contextual menu.

Re-Editing Macros

To re-edit an existing macro, just right-click anywhere within the Node Editor and choose the macro you want to edit from the Macro submenu of the same contextual menu. The Macro Editor appears, and you can make your changes and save the result.

Groups Can Be Accessed Like Macros

Groups can also be loaded from the Insert Tool > Macros submenu if you save a group's .setting file to the Macros folder in your file system. For example, on macOS, the Macintosh HD/Library/Application Support/Blackmagic Design/DaVinci Resolve/Fusion/Macros/ directory.

Other Macro Examples

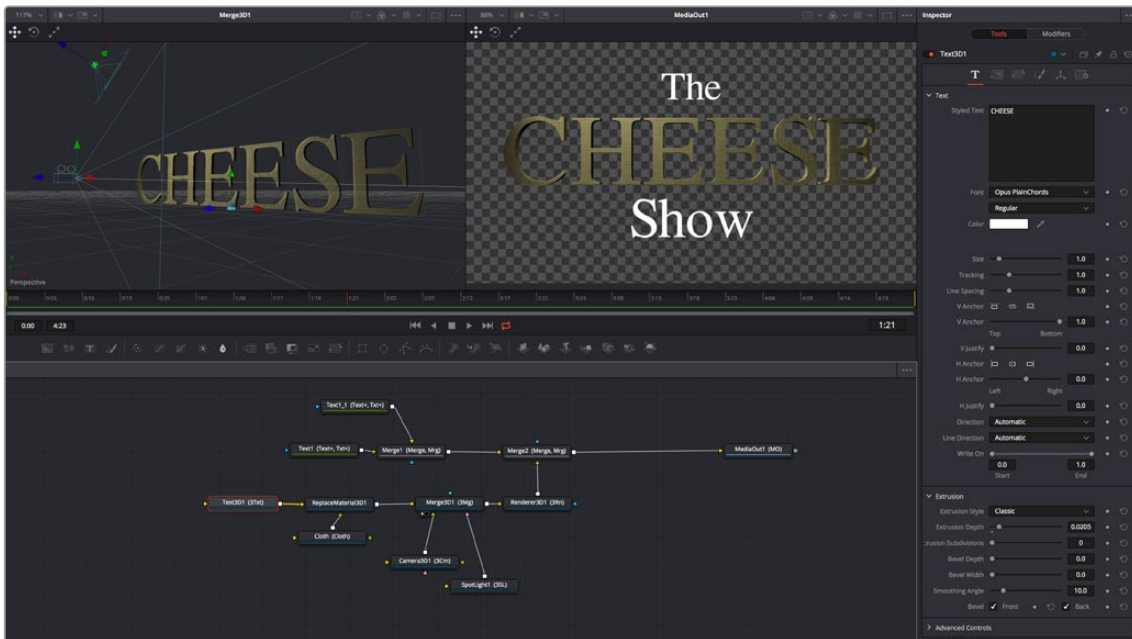
A macro can also be used as a custom LUT. Just copy the macro's .setting file to the LUTs: folder, and the macro will be selectable in the viewers as a LUT. These LUT macros can be used for more than just a color adjustment; you could make a macro that does YUV 4:2:2 resampling, a resize, a sharpening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

Creating Fusion Templates

The integration of Fusion into DaVinci Resolve has enabled the ability to create Fusion Titles, Transitions, Effects, and Generators templates for use in the Edit page. You can create these templates in the Fusion page or within Fusion Studio and then copy them into DaVinci Resolve. Fusion Titles, Generators, and Transition templates are essentially comps created in Fusion but editable in the Timeline of the Edit page with custom controls. This section shows you how it's done.

Getting Started with a Fusion Title Template

The first part of creating a Fusion title template is to create a Fusion composition consisting of Fusion-generated objects assembled to create nearly any kind of title or generator you can imagine. If you're really ambitious, it can include animation. In this example, 3D titles and 2D titles have been combined into a show opener.

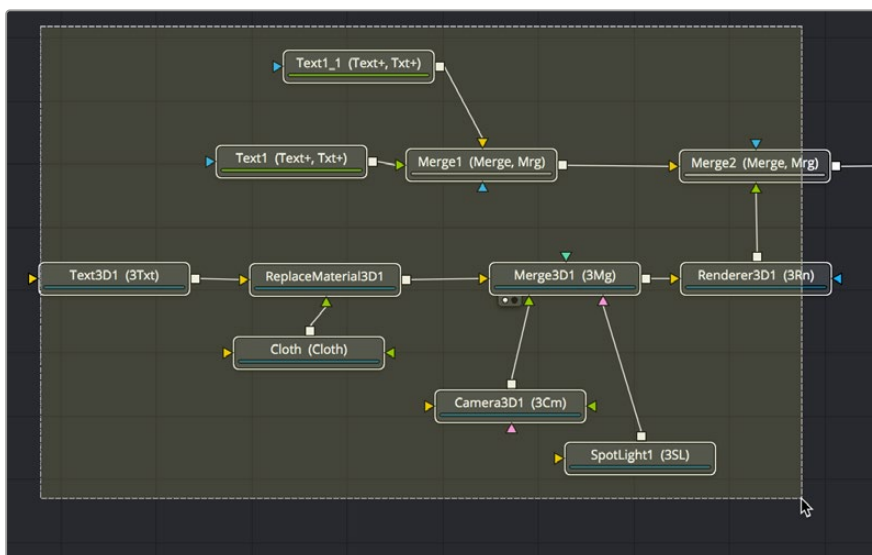


Building a composition to turn into a title template.

Saving a Title Macro

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


Having built your composition, select every single node you want to include in that template except for the MediaIn and MediaOut nodes in DaVinci Resolve or Loader and Saver nodes in Fusion Studio.



Selecting the nodes you want to turn into a title template.


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

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




Creating a macro from the selected nodes.

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



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


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



A simple list of all the nodes we've selected.

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

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




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

Once we've turned on all the parameters we'd like to use in the eventual template, we click the Close button, and a Save Macro As dialog appears.

To have the Title template appear in the Effects Library > Titles category of DaVinci Resolve, save the macro in the following locations:

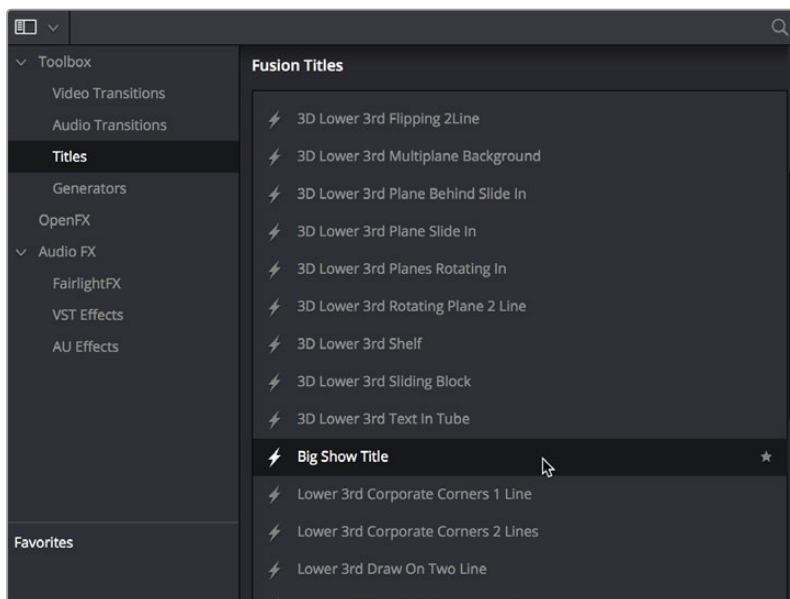
- **On macOS:** Macintosh HD/Users/username/Library/Application Support/Blackmagic Design/DaVinci Resolve/Fusion/Templates/Edit/Titles
- **On Windows:** C:\Users\username\AppData\Roaming\Blackmagic Design\DaVinci Resolve\Support\Fusion\Templates\Edit\Titles
- **On Linux:** home/username/.local/share/DaVinciResolve/Fusion/Templates/Edit/Titles



Choosing where to save a title template for the Edit page in DaVinci Resolve.

Using Your New Title Template

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



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

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




Customizing the template we made.

And that's it!

Getting Started with a Fusion Transition Template

When creating a Fusion transition template, it's easiest to start with an existing transition template and build off that. Three transitions are located in the Fusion Transitions category of the DaVinci Resolve Effects Library. The simplest transition is the Cross Dissolve, while the most complex example is the Slice Push.




The Fusion transition templates located in DaVinci Resolve's Effects Library.

Creating a Fusion Transition Template


The three Fusion Transitions located in the DaVinci Resolve Effects Library are basically Fusion compositions that have been turned into macros. Ordinarily, macros are used as building blocks inside of Fusion so that you can turn frequently-made compositing tricks that you use all the time into your own nodes. However, you can also use this macro functionality to build transition templates for the DaVinci Resolve Edit page.

Once you apply a Fusion Transition in the Edit page, you can right-click it and choose Open in Fusion Page.



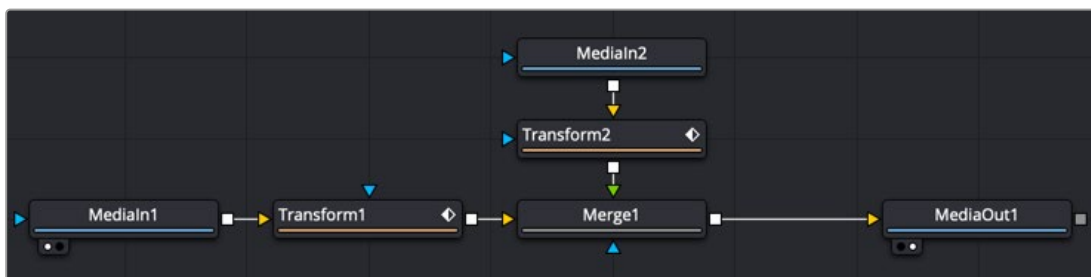
Right-clicking over a Fusion Transition in the DaVinci Resolve Edit page.

The Fusion page opens, displaying the node tree used to create the Fusion transition.



The Cross Dissolve node tree in Fusion.

The MediaIn 1 node represents the outgoing clip in the Edit page Timeline. The MediaIn 2 clip represents the incoming clip. You can modify or completely change the Cross Dissolve effect to create your own custom transition using any of Fusion's nodes.



The Fusion Cross Dissolve node tree replaced with Transforms and a Merge node.

TIP: To modify the duration of the Fusion transition from the Edit page Timeline, you must apply the Resolve Parameter Modifier to any animated parameter. In place of keyframing the transition, you create the transition using the Scale and Offset parameters of the Resolve parameter modifier.

Updating a Fusion Transition

After modifying the transition in Fusion, you can choose to update the transition in the Timeline or create a new transition, which you save into the Edit page Effects Library. To update the transition, just return to the Edit page. The transition in the Edit page Timeline reflects the changes you make in the Fusion page.

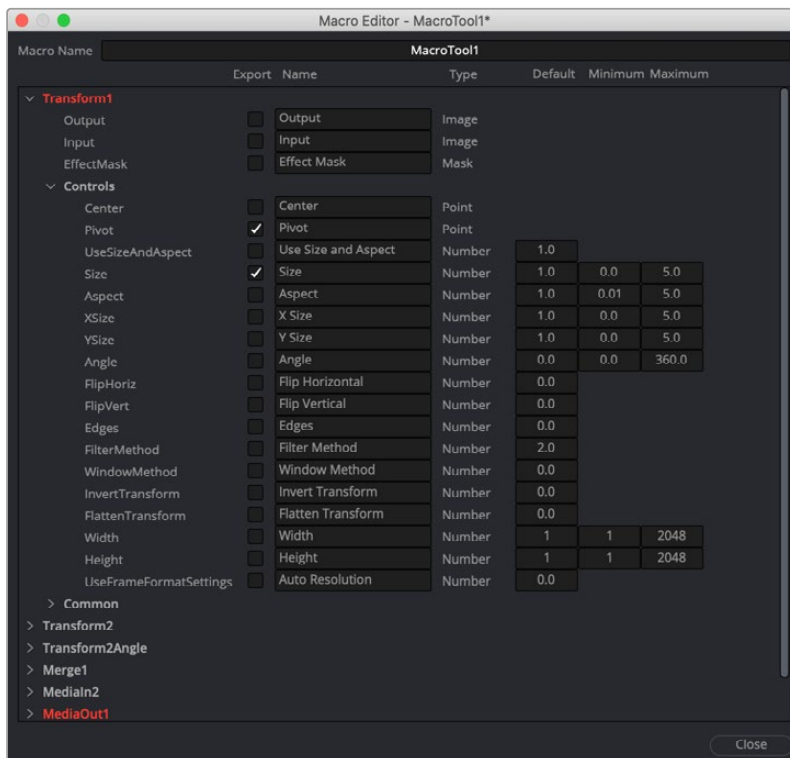
Saving a New Fusion Transition

If after modifying the transition in Fusion, you need to save it to the Effects Library to reuse it on other transitions or other projects, you must make a macro and save it to the Transitions folder.

Start by selecting every node in the Node Editor that you want to include in the transition template, including the two MediaIn nodes and the MediaOut node.

TIP: Since the transition template must include the MediaIn and MediaOut nodes, the final steps for saving a transition template must be performed in DaVinci Resolve's Fusion page and cannot be performed in Fusion Studio.

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



The Macro Editor displaying the parameters of all the nodes you selected.

The Macro Editor window appears, displaying a hierarchical list of every parameter in the composition you've just selected. The order of nodes is based on the order they were selected in the Node Editor prior to creating the macro.

The Macro Editor is designed to let you choose which parameters you want to display as custom controls in the Edit page Inspector when the transition is applied.

For transitions, you can choose not to display any controls in the Inspector, allowing only duration adjustments in the Timeline. However, you can choose a simplified set of parameters for customization by enabling the checkboxes next to any parameter name.

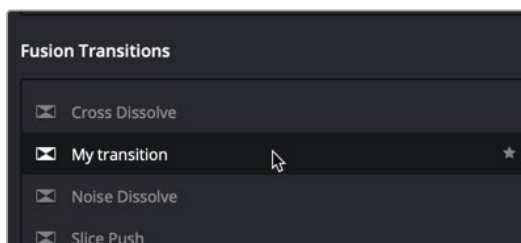
Once you enable all the parameters you want to use in the eventual template, click the Close button, and a Save Macro As dialog appears. Here, you can enter the name of the transition, as it should appear in the Edit page Effects Library.

To have the transition template appear in the Effects Library > Fusion Transitions category of DaVinci Resolve, save the macro in the following locations:

- **On macOS:** Macintosh HD/Users/username/Library/Application Support/Blackmagic Design/DaVinci Resolve/Fusion/Templates/Edit/Transitions
- **On Windows:** C:\Users\username\AppData\Roaming\Blackmagic Design\DaVinci Resolve\Support\Fusion\Templates\Edit\Transitions
- **On Linux:** home/username/.local/share/DaVinciResolve/Fusion/Templates/Edit/Transitions

Using Your New Transition Template

After you've saved your macro, you'll need to quit and reopen DaVinci Resolve. When you open the Effects Library on the Edit page, the new transition template is listed in the Video Transitions category, in the Fusion Transitions list.

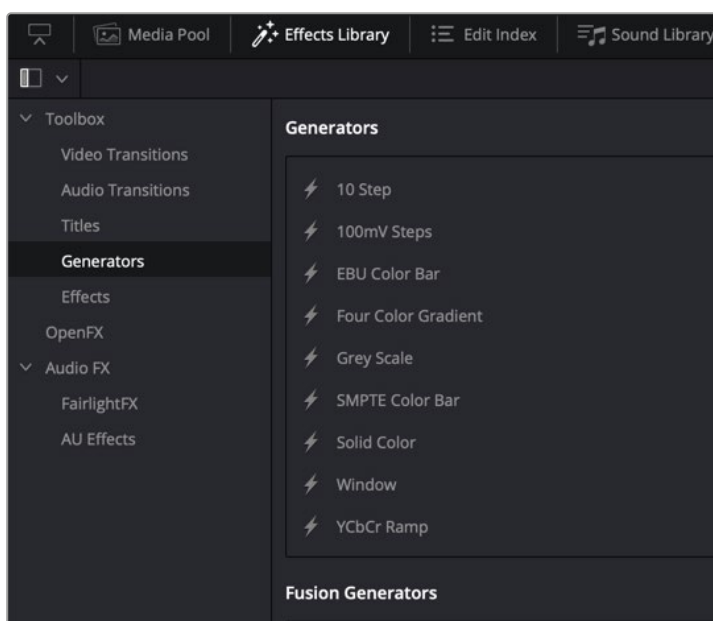


A custom Fusion Transition saved in the Edit page Effects Library.

Applying this transition to a cut in the Timeline and opening the Inspector shows the parameters you enabled for editing, if any.

Getting Started with a Fusion Generator Template

There is one simple Noise Gradient Fusion Generator located in the Effects Library that you can use as a starting point for creating your own generators.



The Fusion Generator templates located in DaVinci Resolve's Effects Library.

To open the Fusion Noise Gradient Generator in the Fusion page, do the following:

- 1 On the Edit page, drag the Fusion Noise Gradient Generator from the Effects Library to the Timeline.
- 2 Right-click over the Noise Gradient Generator and choose Open in Fusion Page from the pop-up menu.

The Fusion page opens, displaying the node tree that is used to create the Fusion Generator.

Creating a Fusion Generator Template

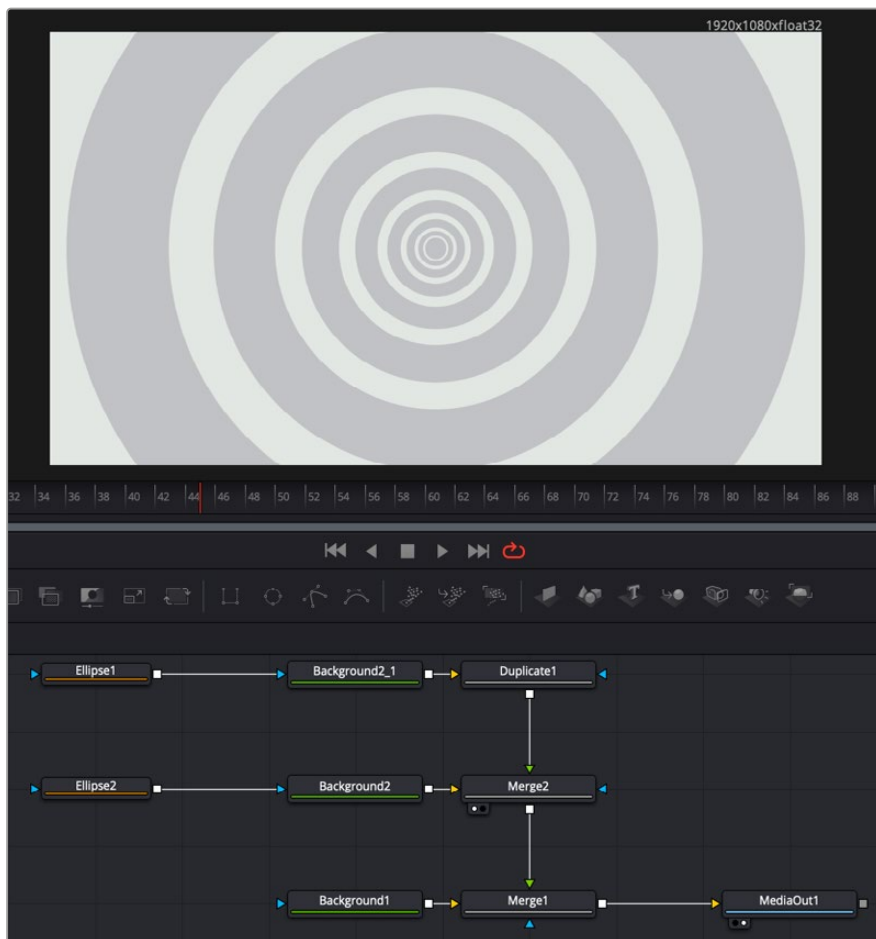
As easy as it is to begin with the Noise Gradient Generator template, you can just as easily start by adding a Fusion Composition Effect to a Timeline in the Edit page.

To begin creating a Fusion Generator Template with an empty Fusion composition, do the following:

- 1 On the Edit page, drag the Fusion Composition Effect from the Effects Library to the Timeline.
- 2 Right-click over the Composition Effect and choose Open in Fusion Page from the pop-up menu.

An empty Fusion page with a single MediaOut node opens, ready for you to create a Fusion Generator.

The Fusion Generator is a solid image generated from any number of tools combined to create a static or animated background. You can choose to combine gradient colors, masks, paint strokes, or particles in 2D or 3D to create the background generator you want.



The Fusion Generator node tree creating concentric circles.

Saving a New Fusion Generator

After creating the generator you want in Fusion, you need to save it to the Effects Library to reuse it in other projects from the Edit page. To do this, you must create a Macro and save it to the Generator folder.

Ordinarily, macros are used as building blocks inside of Fusion so that you can turn frequently-made compositing tricks that you use all the time into your own nodes. However, you also use this macro functionality to build Generator templates for the DaVinci Resolve Edit page.

Start by selecting every node in the Node Editor that you want to include in the Generator template including the MediaOut node.

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



The Macro Editor displaying the parameters of all the nodes you selected.

The Macro Editor window appears, displaying a hierarchical list of every parameter in the composition you've just selected. The order of nodes is based on the order they were selected in the Node Editor prior to creating the macro.

The Macro Editor is designed to let you choose which parameters you want to display as custom controls in the Edit page Inspector when the Generator is applied. You can choose a simplified set of parameters for customization by enabling the checkboxes next to any parameter name.


Once you enable all the parameters you want to use in the eventual template, click the Close button, and a Save Macro As dialog appears. Here, you can enter the name of the Transition, as it should appear in the Edit page Effects Library.

To have the Generator template appear in the Effects Library > Fusion Generators category of DaVinci Resolve, save the macro in the following locations:

- **On macOS:** Macintosh HD/Users/username/Library/Application Support/Blackmagic Design/DaVinci Resolve/Fusion/Templates/Edit/Generators
- **On Windows:** C:\Users\username\AppData\Roaming\Blackmagic Design\DaVinci Resolve\Support\Fusion\Templates\Edit\Generators
- **On Linux:** home/username/.local/share/DaVinciResolve/Fusion/Templates/Edit/Generators

Using Your New Generator Template

After you've saved your macro, you'll need to quit and reopen DaVinci Resolve. When you open the Effects Library on the Edit page, the new Generator template is listed in the Generators category, in the Fusion Generators list.



Custom Fusion Generator saved in the Edit page Effects Library.

Applying this Generator to the Timeline and opening the Inspector shows the parameters you enabled for editing, if any.

Creating a Fusion Effect Template

You start building a Fusion Effect template by bringing a clip from the Edit page Timeline into Fusion. This clip is only used for creating the template and will not be saved with the effect.

Once inside Fusion, you use Fusion's nodes to create the effect you want. You can use a single node or a hundred, depending on the effect you want to create. For instance, using Fusion's Color Correction nodes, you can create a simple color corrector you can use on the Edit page.

To create a simple Color Corrector effect, do the following:

- 1 Insert the Color Corrector node between the MediaIn and MediaOut nodes.
- 2 Select the Color Corrector node in the Node Editor, and then press Cmd-A to select the remaining nodes.
- 3 Right-click over any of the selected node and choose Macro > Create Macros from the contextual menu. Enabling the checkboxes in this window determines the parameters that appear in the Edit page Inspector.
- 4 The Macro Editor window opens. Here, you can enable the checkboxes for any parameters you want to be shown in the Edit page Inspector.
- 5 Enter the name of your effect at the top of the Macro Editor window.
- 6 To save the Macro, click Close at the bottom of the window, then click Yes in the dialog that appears asking you to save the changes.

The Macros must be saved into the correct folder for DaVinci Resolve to recognize the Macro as an effect.

In the save dialog, save the Macro in the following locations:

- **On macOS:** Macintosh HD/Users/username/Library/Application Support/Blackmagic Design/DaVinci Resolve/Fusion/Templates/Edit/Effects
- **On Windows:** C:\Users\username\AppData\Roaming\Blackmagic Design\DaVinci Resolve\Support\Fusion\Templates\Edit\Effects
- **On Linux:** home/username/.local/share/DaVinciResolve/Fusion/Templates/Edit/Effects

You can save and organize your Fusion Effects into separate subfolders underneath the paths above. These subfolders will show up in the Effects section in the Edit page.

To see the effect in the Edit page Effects Library, you'll need to quit DaVinci Resolve and relaunch the application.

Creating a Fusion Effect Template for Two or More Layers

If the effect you want to create requires multiple images like a video wall, you start by creating a Fusion clip on the Edit page Timeline that includes the number of layers you want the effect to have. The clips are only used to create the number of image inputs for the template and will not be saved with the effect.

Once inside Fusion, use Fusion's nodes to create the effect you want.

Save the Macro following the same steps you use for single clip effects. Enable any of the parameters you want to control in the Edit page. To be able to switch the order of video layers within the effect, make sure you have the Layer checkbox enabled for all the MediaIn nodes.

Once you've saved the Macro and relaunched DaVinci Resolve, to use the effect on multiple timeline layers, you must create a Fusion clip. The Fusion clip should contain the same number of layers the effect requires. The order of the Timeline layers, going from the bottom track to the top, matches the MediaIn numbers. For instance, video track 1 will match the position and appearance of MediaIn1, video track 2 matches MediaIn2 and so on. If you want to change how tracks map to MediaIn nodes, you can change the Layer number in the Inspector, assuming you enabled the MediaIn Layer checkbox when creating the Macro.

Changing Durations of a Template

After you make a template in Fusion, you may want to change its duration in the Edit or Cut page Timeline. Changing the duration when animation is involved can be complicated, so there are two Modifiers in Fusion that can help determine how keyframes react when the duration is updated in the Edit or Cut page Timeline.

Anim Curves Modifier

The Animation Curves Modifier (Anim Curves) is used to dynamically adjust the timing, values, and acceleration of an animation, even if you decide to change the duration of a Comp. Using this Modifier, it becomes infinitely easier to stretch or squish animations, create smooth motion, add bouncing properties, or mirror animations curves without the complexity of manually adjusting splines.

When creating Fusion templates for the Edit or Cut page in DaVinci Resolve, the Anim Curves Modifier allows the keyframed animation you've created in Fusion to stretch and squish appropriately as the transition, title, or effect's duration changes on the Edit and Cut page Timelines.

Keyframe Stretcher Modifier

The Keyframe Stretcher modifier is primarily used when creating title templates in Fusion for use in DaVinci Resolve's Edit or Cut page. The Keyframe Stretcher modifier is added to any animated parameter so that the Hold keyframes between the initial animation on screen and the final animation off-screen stretch when the template is trimmed in the Timeline. This allows the animation to retain its timing while the static portion of the title stretches to meet the new duration requirements.

Creating a Custom Template Icon

It is possible to create a custom icon that is embedded with your template thumbnail in the Effects Library, instead of the default first three letters of the template name.

To create a simple Custom Template Icon:

- 1 Create a .png file of the icon you want to use for the template. The recommended size is 104 x 58 pixels, but any image will be resized to fit.
- 2 Name the file exactly the same name as your template, just with a .png extension rather than a .setting extension.
- 3 Place the .png image in the same directory as your template.

When you restart DaVinci Resolve, the icon you created will be embedded in the template thumbnail across all the Effects Libraries in the program.



A Custom Icon added to a fisheye template, before (above) and after (below)

Using Fusion Template Bundles

For ease in distributing templates to other Fusion users, multiple templates can now be bundled together into a single .drfx file. This file can then be imported back into another Fusion workstation easily to ensure that all custom templates are the same between computers.

Creating a Fusion Template Bundle requires using a specific directory structure and using your operating systems file browser and .zip compression utility. The directory structure is listed below, and you can always find a specific folder from within Fusion by right-clicking on any bin in the Effect Library and selecting Show Folder.

Folder structure for Templates used in the Edit page:

- Edit
- Effects
- Generators
- Titles
- Transitions

Folder structure for Templates used in the Fusion page:

- Fusion

To create a Fusion Template Bundle:

- 1 In your OS, create a folder structure above that includes the specific folder for the type of templates you want to be in the bundle. For example, if you have a transition and two effects templates, you would create an Edit folder, and two subfolders inside it named Transitions and Effects.
- 2 Copy your template (.setting) files into the appropriate directories. You can also include icon files and any associated assets as well.

- 3 Use your OS zip compression utility to create a .zip file of the directory structure.
- 4 Rename the .zip file in your OS with the “.drfx” extension instead of .zip. The file icon should change to reflect the new extension.

To import a Fusion Template Bundle:

- 1 Double-click on a .drfx file in your OS. DaVinci Resolve will launch and a dialogue box will appear asking if you want to install the template bundle.
- 2 Drag to the .drfx file from your OS directly into the Fusion page in DaVinci Resolve. A dialogue box will appear asking if you want to install the template bundle.

To delete a Fusion Template Bundle:

- 1 Navigate to the appropriate template directory in your file browser.
- 2 Delete the .drfx file.

IMPORTANT

The Fusion Template Bundle contains all the templates in one file. It does not uncompress them into separate template files again. Therefore if you delete the .drfx file, all associated templates inside that bundle will be removed as well.

Chapter 67

Using Viewers

This chapter covers working with viewers in Fusion, including using onscreen controls and toolbars, creating groups and subviews, managing viewer Lookup Tables (LUTs), working with the 3D viewer, and setting up viewer preferences and options.

Contents

Viewer Overview	1142	Split Wipes between Buffers	1151
Single vs. Dual Viewers	1143	Moving the Wipe Divider	1152
Floating Viewers in Fusion Studio	1143	Subviews	1152
Video Output	1143	Showing and Hiding Subviews	1153
Clean Feed	1143	Changing the Subview Type	1153
Loading Nodes into Viewers	1143	Swapping the Subview with the Main View	1153
Clearing Viewers	1144	Viewer and Subview Types	1154
Position and Layout	1144	Viewing Selective Channels	1158
The Viewer Divider	1144	Viewing Color Channels	1158
Zooming and Panning into Viewers	1145	Viewing Auxiliary Channels	1158
Flipbook Previews	1146	The 3D Viewer	1159
Creating Flipbook Previews	1146	Panning, Scaling, and Rotating a 3D Viewer	1159
Playing Flipbook Previews	1147	Viewing Objects via Wireframe	1159
Removing Flipbook Previews	1147	Changing the POV of a 3D Viewer	1159
Flipbook Preview Render Settings	1147	Copying a Viewer's POV to a Camera	1160
Onscreen Controls	1148	Lighting and Shadows in 3D Viewers	1161
Showing and Hiding Onscreen Controls	1149	Transparency in 3D Viewers	1162
Making Fine Adjustments to Onscreen Controls	1149	Grid	1162
Toolbars	1150	Vertex Normals	1163
Viewer Toolbar	1150	Quad View	1164
Node Toolbars	1150	Quad View Layouts	1164
A/B Buffers	1150	Using Quad Views for 2D Scenes	1164
Flipping between Buffers	1151		

Guides	1165	Editing Viewer LUTs	1171
Frame Format Settings	1166	LUT Processing Order	1173
Domain of Definition and Region of Interest		Applying Multiple LUTs	1174
	1166	Saving Custom LUTs	1174
Domain of Definition (DoD)	1166	LUT Files	1175
Region of Interest (RoI)	1167	Viewer Preferences and Settings	1176
Enabling RoI Controls	1167	Viewer Settings	1176
Managing Viewer Lookup Tables (LUTs)	1169	The Viewer Options Menu	1177
How Lookup Tables Work in Fusion	1169	Locking the Viewer (Command-L)	1177
Types of Viewer LUTs	1169	Additional Viewer Options	1178
Using Viewer LUTs	1171	Status Bar Information	1178

Viewer Overview

Viewers in Fusion display the current frame of the current composition in a variety of ways to help you see what you're doing and evaluate the final result of your compositing artistry. Viewers display 2D images, but they can also display a 3D environment using a 3D View as well as a special Quad viewer to help you effectively work in three dimensions.



Side-by-side dual viewers: a 3D viewer (left), and a 2D viewer (right)

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



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

Single vs. Dual Viewers

By default, there are two viewers positioned side by side across the top of the window. However, you can use the Single/Dual Viewer button to toggle between displaying a single viewer or two viewers displayed side by side.



The Single/Dual Viewer toggle button

Floating Viewers in Fusion Studio

In addition to the dual views above the Node Editor, Fusion Studio provides the option to use an unlimited number of floating viewers. These floating viewers are excellent for taking full advantage of a dual monitor configuration. Floating viewers can also be set to full-screen mode to make the best use of screen real estate.

To create a new floating display view, select Window > New View from the menu bar at the top of the screen. The position and configuration of the new view can be saved in the layout preferences, if required.

Video Output

When using DaVinci Resolve or Fusion Studio, if Blackmagic video hardware is present in the computer, then you can select a node to preview directly on that display. While video output can't be used for manipulating onscreen controls such as center crosshairs or spline control points, they're extremely valuable for evaluating your composition via the output format, and for determining image accuracy using a properly calibrated display.

The video hardware is configured from the DaVinci Resolve and Fusion Studio preferences.

Clean Feed

When using DaVinci Resolve with dual computer monitors, a full-screen viewer can be displayed on the secondary monitor from the Fusion page. This displays a third view indicator button under each node to control what is shown on the second display. To activate this monitor, make sure you do not have Dual Screen enabled under the Workspace menu and then select Workspace > Video Clean Feed and select your second computer display from the submenu.

Loading Nodes into Viewers

The Fusion page in DaVinci Resolve and Fusion Studio show two different things when you first open each application. When you first open the Fusion page, the output of the current empty composition (the MediaOut1 node) is usually showing in viewer 2. If you're in dual-viewer mode, viewer 1 remains empty until you assign a node to one of them. In Fusion Studio, because there are no nodes when you first begin a comp, nothing is displayed in the viewers.

To load specific nodes into specific viewers, do one of the following:

- Hover the pointer over a node, and click one of two buttons that appear at the bottom-left of the node.
- Click once to select a node, and press 1 (for the left viewer) or 2 (for the right viewer).
- Right-click a node and choose View On > None/Left View/Right View in the contextual menu.

- Right-click the control header of a node in the Inspector, and choose View On > None/Left View/Right View from the contextual menu.
- Drag a node and drop it over the viewer you'd like to load it into (this is great for tablet users).

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



Viewer assignment buttons at the bottom left of nodes indicate when they're being viewed, and which dot is highlighted indicates which viewer that node is loaded into.

Clearing Viewers

To clear an image from a viewer, click in the viewer to make it active; a light purple outline is displayed around the active panel. With the viewer active, press the Tilde (~) key. This key is usually found to the left of the 1 key on U.S. keyboards. The fastest way to remove all images from all viewers is to make sure none of the viewers is the active panel, and then press the 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 Fusion Global Layout preferences.

If you want all new compositions to open with a certain viewer layout, you can configure the layout of the two primary viewers, and then use the Grab Document Layout button in the Fusion Global Layout preferences to remember the layout for any new compositions. To save the position and size of floating viewers, you use the Grab Program Layout button. Finally, if you want to have the floating viewers opened automatically when you open Fusion, enable the Create Floating Views checkbox.

The Viewer Divider

You can change the relative sizes of the left and right viewers using the horizontal viewer divider that runs between them. Drag the viewer divider to increase or decrease the amount of space used by one viewer. The adjacent viewer will adjust to accommodate the new layout.

The amount of vertical space available for both viewers can be adjusted by dragging the horizontal scrollbar between the viewers and the work area below them.



The viewer divider bar

Zooming and Panning into Viewers

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

Methods of panning viewers:

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

Methods of scaling viewers:

- Click a viewer and press the Equals key (=) to zoom in, and the Minus key (-) to zoom out.
- Press the middle and left mouse buttons simultaneously and drag left or right to resize the viewer.
- Hold down the Command key and use your pointer's scroll control to resize the viewer.
- Hold down the Command key and drag two fingers on a track pad to resize the viewer.
- Hold down the middle mouse button, and then click the left mouse button to zoom in, or click the right button to zoom out. The scaling uses a fixed amount, centered on the position of the cursor.
- Click a viewer and press Command-1 to resize the image in the viewer to 100 percent.

- Click a viewer and press Command-F or Command-1 to reset the image in the viewer to fit the viewer.
- Click the Scale Viewer menu and choose Fit or a percentage.
- Right-click on a viewer and choose an option from the Scale submenu of the contextual menu. This includes a Custom Scale command that lets you type your own scale percentage

Methods of spinning 3D viewers:

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

Flipbook Previews

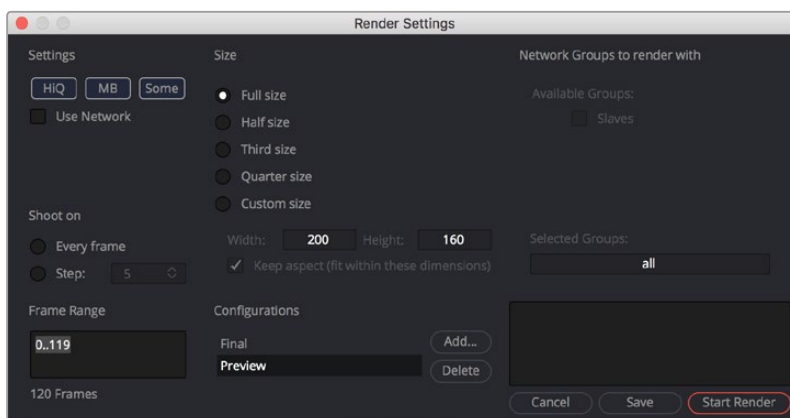
As you build increasingly complex compositions, and you find yourself needing to preview specific branches of your node tree to get a sense of how various details you're working on are looking, you may find it useful to create targeted RAM previews at various levels of quality right in the viewer by creating a RAM Flipbook. RAM Flipbook Previews are preview renders that exist entirely within RAM and allow you to render a node's output at differing levels of quality for quick processing in order to watch a real-time preview.

Creating Flipbook Previews

Creating a Flipbook Preview is relatively fast, once you know where to look.

To create a Flipbook Preview:

- 1 Choose the node in your node tree that you want to preview by doing one of the following:
 - Hold down the Option key while dragging a node into the viewer.
 - Right-click a node and choose an option from the Create/Play Preview submenu in the contextual menu.
- 2 When the Preview Render dialog opens, choose the quality, resolution, and motion blur settings you want to use for the Flipbook Preview.



The Flipbook Preview Render dialog

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

The current frame range of the Time Ruler is rendered using the settings you've selected, and the result is viewable in the viewer you selected or dragged into.

Once you've created a Flipbook Preview within a particular viewer, right-clicking that viewer presents Flipbook-specific commands and options to Play, Loop, or Ping-Pong the Flipbook, to open it Full Screen, to Show Frame Numbers, and to eliminate it.

TIP: If you want to create a Flipbook Preview and bypass the Render Settings dialog by just using either the default setting or the settings that were chosen last, hold down Shift-Option while you drag a node into the viewer. The Settings dialog will not appear, and rendering the preview will start right away.

Playing Flipbook Previews

While the Flipbook Preview is loaded into a viewer, or open in full-screen mode, you can play or scrub through it using the mouse and the keyboard.

To play back a Flipbook using the mouse, do the following:

- Double-click in the viewer to start playback.

To scrub through a Flipbook using the mouse, do the following:

- Hold down the right mouse button down and drag left or right to scrub through frames.

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

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

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

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

TIP: The mouse and keyboard shortcuts work in full-screen mode as well.

Removing Flipbook Previews

Once you create a Flipbook Preview, you need to know how to clear it from RAM.

To eliminate a Flipbook you've created:

- Right-click within a viewer containing a Flipbook Preview, and choose Remove Preview.

Flipbook Preview Render Settings

This section covers all the settings available for rendering Flipbook Previews to RAM.

Settings

The Settings section of the Preview Render dialog includes three buttons that determine the overall quality and appearance of your Flipbook Preview. These buttons also have a significant impact on render times.

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

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

Size

Since RAM Flipbook Previews use RAM, it's helpful to know how many frames you can render into RAM before you run out of memory. The Flipbook Preview dialog calculates the currently available memory and displays how many frames will fit into RAM. If you have a small amount of RAM in your computer and you cannot render the entire range of frames you want, you can choose to lower the resolution to a setting that delivers the best quality/duration ratio for your preview.

Network

Network rendering is only available in Fusion Studio. For more information on network rendering, see Chapter 64, "Rendering Using Saver Nodes" in the DaVinci Resolve Reference Manual or Chapter 4 in the Fusion Reference Manual.

Shoot On

Sometimes you may not want to render every single frame, but instead every second, third, or fourth frame to save render time and get faster feedback. You can use the Step parameter to determine the interval at which frames are rendered.

Frame Range

This field defaults to the current Render Range In/Out set in the Time Ruler to determine the start and end frames for rendering. You can modify the range to render more or fewer frames.

Configurations

Once you've created a useful preview configuration, you can save it for later use by clicking the Add button, giving it a name, and clicking OK.

Updating a Preview

This option is designed for the interactive frame-by-frame work of rotoscoping and painting. Right-click over a preview in the viewer and choose Update from its contextual menu. When active, any frames that are modified on the previewed node are automatically updated in the preview's playback. This lets you reserve the RAM for playback. You can keep it playing on a loop or ping-pong while you work in another viewer.

Onscreen Controls

When it comes to adjusting images, the Control Panel provides very precise numerical values, but sometimes visually positioning an element using onscreen controls can get you where you want to go with less tweaking. The viewers show onscreen controls for manipulating the parameters of the currently selected node. Common onscreen controls include crosshairs, angle indicators, polylines, and paint strokes. Each of these controls can be manipulated directly in the viewer using the mouse or keyboard.



The Angle preview control

The controls shown in viewers are determined by which nodes are selected, not by the node displayed in the viewer. For example, a downstream blur is easily viewed while manipulating the controls for a selected polygon mask or merge. If multiple nodes are selected, the controls for every selected node are shown simultaneously.

Showing and Hiding Onscreen Controls


The onscreen controls for a viewer can be hidden so they don't interfere with viewing the image.

To toggle the visibility of onscreen controls, do one of the following:

- Click a viewer's Option menu and choose Show Controls to toggle the controls on or off.
- Right-click in a viewer and choose Options > Show Controls from the contextual menu.
- Select a viewer and press Command-K.

Enabling/Disabling Onscreen Controls in Specific Nodes

Some nodes, like masks, allow disabling of their onscreen controls on a per-node basis, since you often use multiple Polygon nodes to organize and animate masks.



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

Making Fine Adjustments to Onscreen Controls

If you want the visual guidance of onscreen controls with the precision of the Inspector, you can use different keyboard modifiers.


- Up and Down Arrow keys can be used to adjust the vertical position of an onscreen control by small steps.
- Holding down the Command key while using the Up and Down Arrow keys reduces the scale of each step by a factor of ten. Holding Shift increases the scale of each step by a factor of ten.

Toolbars

There are two toolbars in the viewer: a viewer toolbar, which always appears at the top of each viewer and gives you control over what that viewer shows, and an optional node toolbar that appears underneath that gives you contextual controls based on the node you've selected in the Node Editor.

Viewer Toolbar

A viewer toolbar runs across the top of each viewer, providing access to many of the most commonly used viewer-related settings, as well as an indication of the status of many of the most important settings. Most of the menus and buttons found on this toolbar are described in detail throughout this chapter.



The viewer toolbar

Node Toolbars


In addition to the viewer toolbar, a node toolbar is displayed underneath, at the top of the viewer display area, whenever you select a node that exposes special nodes. Examples of nodes that expose a toolbar include the text, masks, paths, paint strokes, and the 3D environment.



The node toolbar shown for the Paint node

Customizing the Node Toolbar

If you want to change the size of the buttons that appear in the Node toolbar, or turn on text names for each node, you can right-click anywhere in the empty area of the toolbar and choose new settings from the Icon Size and Button Style submenus in the contextual menu.



The contextual menu for the node toolbar

A/B Buffers


Each viewer has two buffers, each of which can contain images from different nodes, enabling easy comparison of two different nodes within the same viewer by either toggling between buffers, or via an adjustable split-wipe. Each buffer can be considered a complete and separate viewer within the same viewer pane. The A buffer is always shown by default, so when you first load a node into a viewer, the image loads into the A buffer.

Flipping between Buffers

Switching between buffers is easy, either to view a different image while keeping another image handy, or to flip between the original image and the affected image for comparison.

To switch between buffers, do one of the following:

- Select a viewer and press comma (,) to select the A buffer or press period (.) to select the B buffer.
- Click the Buffer menu and choose either Switch to A View or Switch to B View.



The Buffer menu lets you switch between buffers


TIP: Each buffer can be set to different display settings—for example, showing different channels or different viewing LUTs, either applied to different nodes or applied to two buffered versions of the same node.

Split Wipes between Buffers

You can also wipe between both buffers, providing a more direct means of comparison.

To wipe between buffers, do one of the following:

- 1 Prepare to wipe between two images by loading different nodes into each buffer, or load the same node with different viewer options into each buffer.
- 2 To toggle the split wipe on or off, do one of the following:
 - a. Click the Switch to Split Wipe View button.
 - b. Press Forward Slash (/).
- 3 To adjust the wipe, do one of the following:
 - a. Move the center of the wipe by dragging the center handle of the wipe divider.
 - b. Press Command-Option and click anywhere in the viewer to jump the wipe divider to that location.
 - c. Change the angle or the wipe by dragging the wipe divider. Dragging the wipe divider while holding the Shift key snaps it to the nearest 45-degree angle.
 - d. Panning or zooming the viewer pans and zooms both buffers together.
- 4 (Optional) If you want to change the image that's displayed on that side of the split, you can drag new nodes onto either half of the viewer.
- 5 To turn off the wipe, click the Switch to Split Wipe View button again (or press /).



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


Even when you wipe, you can choose different display channels, view LUTs, or other display options for each buffer individually by clicking on the half of the wipe you want to alter, and then choosing the options you want that buffer to use. This allows easy comparison of different channels, LUTs, or other viewer settings while wiping the same image, or different images.

Moving the Wipe Divider

Occasionally, you will have either zoomed in or panned so far from the viewer divider that it's no longer visible in the viewer. Holding down Command-Option and clicking anywhere in the image will cause the viewer divider to jump to the current position of the pointer.

Subviews

A subview is a “mini” viewer that appears within the main viewer. A subview is usually used to show different information about the image.



The Subview menu with the Histogram subview displayed

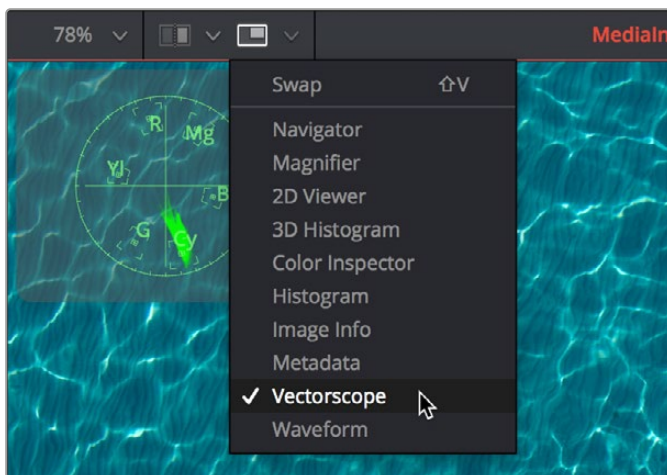
For example, the RGB channels can be viewed in the main viewer, while the alpha channel is displayed in a subview. For the most part, the subview is a fully functional miniature viewer, with its own contextual menu and options. It responds to the same keyboard shortcuts and navigation controls as any other viewer. However, there are several view types designed for use only in the subview, including the Navigator, Magnifier, Color Inspector, and Image Info.

Showing and Hiding Subviews

Subviews are easily shown and hidden.

To enable the currently selected subview in the Subview menu of a viewer, do one of the following:

- Click the Subview button in the View toolbar.
- Choose Views > Subview > Enabled from the contextual menu.
- Click a viewer, and press the V key.



The Subview button in the viewer toolbar

Changing the Subview Type

The Subview button enables and disables the subview, which usually shows the last subview you chose. You can change this at any time.

To change which subview type is displayed, do one of the following:

- Click the small arrow to the right of the Subview button to open its menu and choose which subview you want.
- Right-click within a subview to bring up the subview's contextual menu.

The Subview drop-down menu and contextual menu show all the available subview types. Once you choose an option from the list, that view will be displayed in the subview, and the Subview button will show and hide it as you wish.

Swapping the Subview with the Main View

It's possible to swap the contents of the main viewer and the subview for select subview types. However, certain view types, such as the Color Inspector and Magnifier, can only be used in the subview. In these cases, the swap will not take place.

To swap the contents of the subview with the main view, do one of the following:

- Press Shift-V.
- Right-click in a viewer and choose Views > SubView > Swap from the contextual menu.


Viewer and Subview Types

Viewers can be changed to show a variety of different information about the image, but not all view types are available at all times. For example, the 3D Viewer is not available for a 2D node, and some of the measurement viewers are available only as subviews. Below is detailed information about the different view types available.

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 Viewer is an analysis node that can be used to identify problems with the contrast and dynamic range in an image. The graph shows the frequency distribution of colors in the image, including out-of-range colors in floating-point images. The horizontal axis shows the colors from shadows to highlights. The vertical axis shows the number of pixels in the image that occur at each level.


The Histogram Viewer will also display gradient information. You can use the From Image and Perturb modifiers to output gradients. If you need to see the gradient represented in a histogram, drag the modifier's title bar into the viewer.



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

3D Histogram

The more advanced 3D Histogram Viewer shows the color distribution in an image within a 3D cube. One advantage to a 3D Histogram is that it can accurately represent the out-of-range colors commonly found in floating-point and high-dynamic-range images. It can also be used to look at vector images like position, normal, velocity, and so on.



The 3D Histogram Viewer type for evaluating out-of-range colors

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

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

Vectorscope


The Vectorscope Viewer duplicates the behavior of a specific type of video test equipment, displaying a circular graph that helps to visualize the intensity of chrominance signals.



The Vectorscope Viewer type for evaluating chrominance in an image

Waveform


The Waveform Viewer duplicates the behavior of a specific type of video test equipment, displaying a line or bar graph that helps to visualize the voltage or luminance of a broadcast signal.



The Waveform Viewer type for evaluating luminance in an image

Navigator


The Navigator can only be used in a subview. It provides a small overview of the entire image, with a rectangle that indicates the portion of the image that is actually visible in the main viewer. This is useful when zooming in on an image in the main view.



The Navigator subview for panning the image while zoomed in

Magnifier


The Magnifier can be used only in a subview. It shows a zoomed-in version of the pixels under the cursor in the main viewer.



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

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



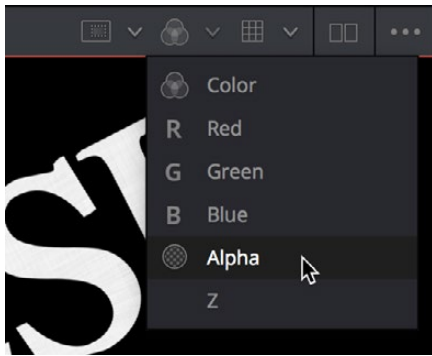
The Metadata subview for viewing embedded metadata

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 from the list of available channels in the current image.
- Click the viewer you want to toggle, and press one of the following keyboard shortcuts:
 - C - Full RGB color display
 - R - Display red channel
 - G - Display green channel
 - B - Display blue channel
 - A - Display alpha channel
 - Z - Display Z-buffer channel

Viewing Auxiliary Channels

The viewers support RGBA and Z channels using keyboard shortcuts, but they support other channels as well. File formats such as OpenEXR often include auxiliary image data that provide more control and compositing options when working with rendered 3D images. To view auxiliary image data in a viewer, click the arrow to the right of the RGB button to display the drop-down menu or right-click in the viewer and choose an option from the Channels submenu of the contextual menu.

The 3D Viewer

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

For more information on 3D controls, see Chapter 85, “3D Compositing,” in the DaVinci Resolve Reference Manual or Chapter 25 in the Fusion Reference Manual.

Panning, Scaling, and Rotating a 3D Viewer

For the most part, navigation in the 3D Viewer is similar to the navigation in the 2D Viewer. Panning and zooming work with the same controls even though you’re moving within a 3D space. However, when viewing a 3D scene, panning changes the point of view and thus the center point for scaling and rotation, too. A combination of panning and rotation will allow you to move the point of view anywhere in the scene.

Another small change is that there’s a lower limit to the scale of a 3D scene. Continuing to zoom in past this limit will instead move (“dolly”) the point of view forward. The mouse wheel will move forward slowly, and the keyboard will move more quickly.

Critically, the 3D Viewer gives you additional control to rotate the viewer within the three dimensions of the scene to better see your scene from different angles as you work.

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

- Hold down the Option key and drag left or right using the middle mouse button.
- Hold down the middle and right mouse buttons while dragging.

The rotation is centered on the middle of the view.

TIP: These rotation controls can be used with the 3D Histogram subview as well.

Viewing Objects via Wireframe

3D composites not only work with 2D images on image planes but can also integrate true geometry, such as that generated by the Particle system, Text 3D node, imported FBX meshes, and basic primitives from the 3D toolset. Using a Wireframe view helps to see through a mesh or see the density of the geometry. It is much easier to see a change in the Subdivision level of an ImagePlane3D in wireframe than viewing the rendered image.

To display 3D geometry in wireframe, do the following:

- Right-click the 3D Viewer and choose 3D Options > Wireframe from the contextual menu.

Changing the POV of a 3D Viewer

Compositing a 3D scene often requires that you view the scene from different angles to make sure the position and animation are what you want. While the 3D Viewer uses a perspective camera that allows you to look at the 3D scene from any angle, you can also switch the 3D Viewer to view from the front, top, left, or right side of the scene; these are collectively called Orthographic views.

Additionally, if you have a camera or spotlight in your scene, you can switch the viewer to face the scene from the point of view of those objects.

To change the 3D viewpoints:

- Right-click the viewer and choose an option from the Camera submenu of the contextual menu. The choices include Perspective, Front, Top, Left, and Right.

Changing Cameras in a 3D Viewer

If you have one or more camera objects in the 3D scene, they will be listed as options in the contextual menu. Spotlights and other lights or objects in the scene will be listed in the Other submenu. If you choose any one of these objects, the 3D Viewer jumps to display the scene from the point of view of the chosen object. While looking “through” an object, rotating, panning, or zooming, the viewer will instead change the rotation, position, and scale of the camera, light, or other object.

Copying a Viewer’s POV to a Camera

There are many times you will have used the default perspective view to set up your scene and then want that point of view (POV) to become your main camera’s point of view. You can easily position and orient a camera, light, or other object to match the current point of view shown in the viewer using the Copy PoV To command.

To copy the point of view in the viewer to a camera, do the following:


- 1 Set up a 3D Viewer with the point of view you want by zooming, panning, and rotating the viewer.
- 2 Add a camera to your 3D scene.
- 3 Right-click anywhere within the 3D Viewer and choose Camera > Copy PoV To > Camera3DNameOfCamera from the contextual menu.

The Camera3D’s controls will inherit the viewer’s position and angle values.

TIP: The Copy PoV To command uses the object’s own coordinate space; any transformations performed downstream by another node are not taken into account.

POV Labels

As you switch the POV of the viewer, you can keep track of which POV is currently displayed via a text label at the bottom-left corner of the viewer. Right-clicking directly on this label, or on the axis control above it, acts as a shortcut to the Camera submenu, allowing you to easily choose another viewpoint.



The Viewpoint label and Axis control in a 3D Viewer

Lighting and Shadows in 3D Viewers

Before you add lights to a 3D scene, default lighting is provided. This basic, flat lighting allows you to see the shading on objects without requiring you to add and set up lights as you work in the 3D Viewer. Additionally, shadows are hidden by default. Once you start adding lights of your own, you need to switch modes to see what they affect as you work.

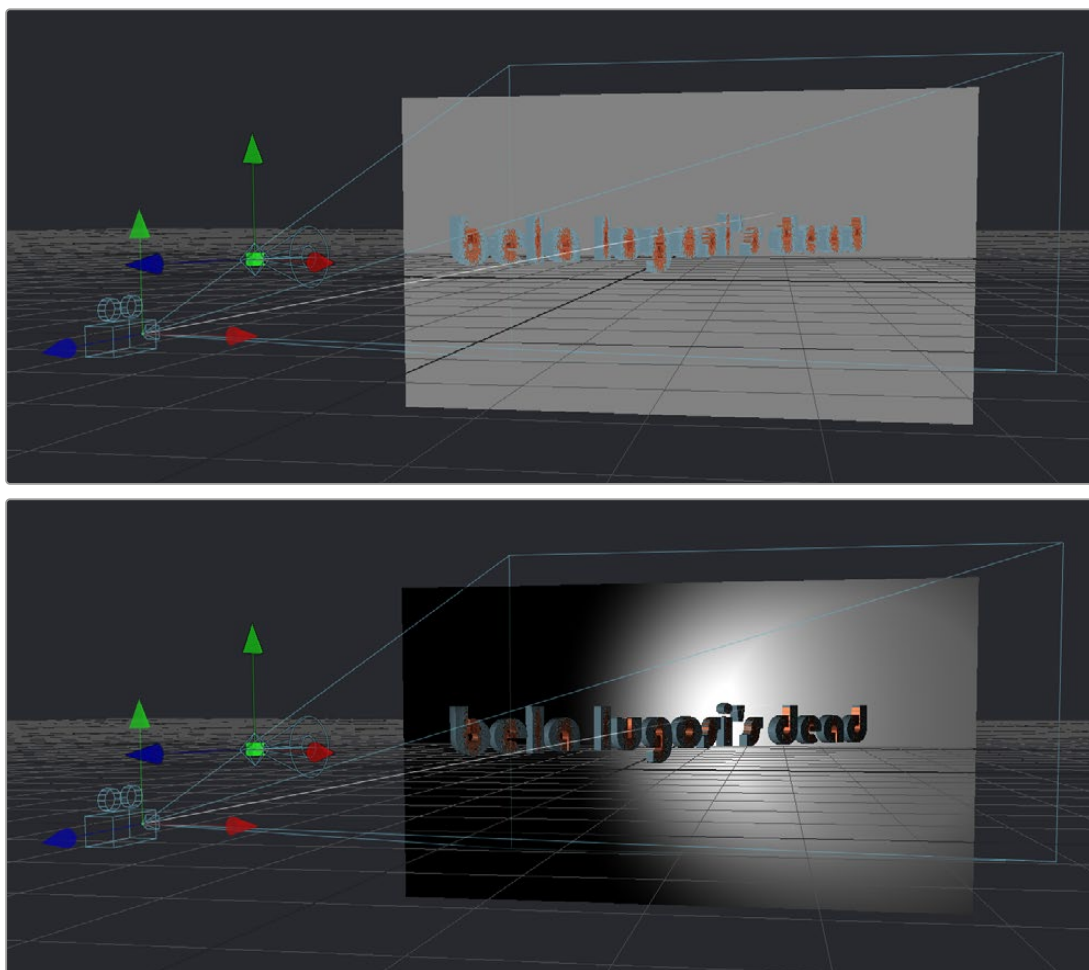
To see the effects of the default light on the scene:

- Right-click within the 3D Viewer and choose 3D Options > Default Lights from the contextual menu.

When you're ready to add your own lighting to a scene, you can connect light nodes in various ways to a Merge 3D node for the scene you're working on. Once you connect a light to a Merge 3D node, you need to switch the 3D Viewer over to showing the new, proper lighting.

To toggle lighting rendering within a 3D scene:

- Right-click within the 3D Viewer and choose 3D Options > Lighting from the contextual menu.




A 3D scene using default lights (top), and the same scene with lighting turned on (bottom)

TIP: Attempting to load a Light node into a viewer all by itself will result in an empty scene, with nothing illuminated. To see the effects of lights, you must view the Merge 3D node the light is connected to.

Similar to lights, the default 3D Viewer has shadows turned off. To see shadows cast from the lighting you've created, you must turn them on.

To toggle shadows rendering within a 3D scene:

- Right-click within the 3D Viewer and choose 3D Options > Shadows from the contextual menu. Enabling shadows will automatically turn on lighting, if it is not already turned on.



A 3D scene with shadows enabled along with the lights

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.



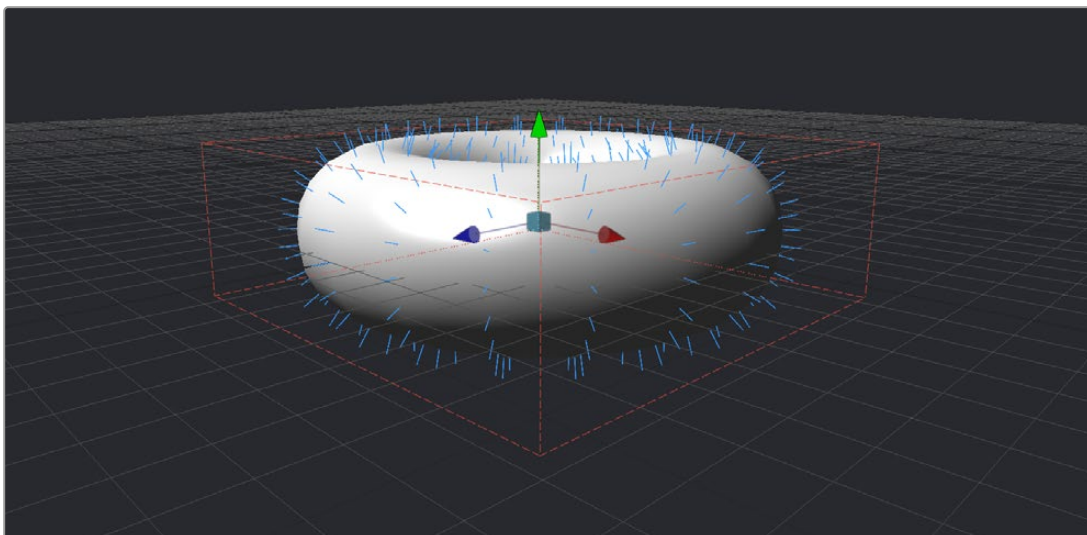
The default grid of the 3D Viewer grid with its origin at $x = 0$, $y = 0$ and $z = 0$

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.



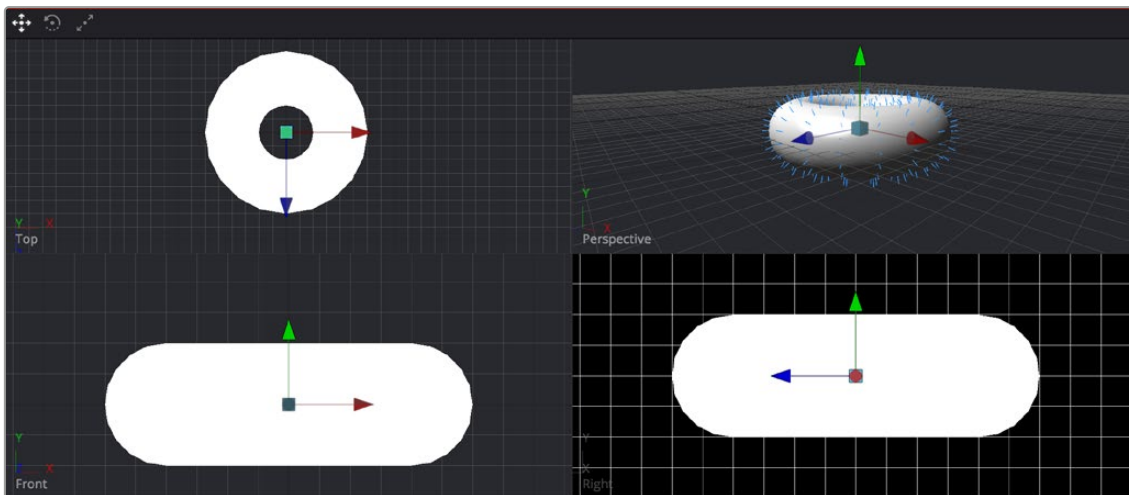
The normals viewed in a 3D scene

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, which divides the viewer into four panes. These panes can then display four different angles of the scene at one time.

To toggle the display of the Quad view, do one of the following:

- Right-click anywhere within the viewer and choose Views > Quad View from the contextual menu.
- Press Shift-Q.



A Quad view of a 3D scene

While there are four panes in the Quad view, they all show the same scene. When assigning views within a Quad view, you can choose between displaying Front, Left, Top, Bottom, and Perspective orthographic views, or you can choose the view through any camera or spotlight that's present in the scene.

To assign different views to panes of a Quad view, do one of the following:

- Right-click directly on the POV label at the bottom left of the pane you want to reassign, and choose another camera, light, or Point of View from the contextual menu.

Quad View Layouts

There are a variety of Quad view layouts, ranging from four equally sized panels to having three small panels across the bottom of a larger single panel.

To switch to a different Quad view layout, do the following:

- 1 Enable the Quad view.
- 2 Right-click anywhere within the viewer and choose an option from the Views > Quad Layouts submenu of the contextual menu.

Using Quad Views for 2D Scenes

Quad views aren't only useful for 3D scenes. They can also be used with 2D scenes, with each pane showing a different image channel or subview type. For example, one pane can show the image while the other panes show the alpha channel, a vectorscope, and a histogram.

To assign different channels or subviews to panes of a Quad view for a 2D scene:


- 1 Right-click in a viewer and choose Views > Quad View.
- 2 Click once in the pane you want to reassign.
- 3 Do one of the following:
 - a. Choose a channel from the Channel Viewer menu.
 - b. Right-click in the viewer and choose Views, and then choose a Subview from the submenu.

Guides

Guides are onscreen overlays used to help you compose elements within a boundary or along the center vertical and horizontal axes. While guides are displayed in the viewer, they're not rendered into the scene. There are four commonly used guides that can be displayed, including Monitor Safety, Safe Title, Center, and Film.

Methods of using guides:

- **To display guides in a viewer:** Right-click in the viewer and then choose Guides > Show Guides from the contextual menu, or press Command-G.
- **To change the aspect ratio of the displayed guides:** Right-click in the viewer and then choose an option from the Guides > Frame Aspect submenu. The frame aspect is usually set to Default, which forces the frame aspect to the same resolution as the image that's displayed in the view. However, when the frame aspect is set to a specific value, the guides will conform to the exact boundaries of the specified format and any image area outside of that will be dark gray.
- **To show or hide specific guides:** Right-click in the viewer and then choose an option from the Guides submenu. A variety of specific guides are provided, each of which can be individually enabled and disabled.
 - **Monitor Safety:** Monitor Safety indicates the safe action area viewable on most monitors and TV screens.
 - **Safe Title:** Safe Title indicates the safe area for titles viewable on all TV and monitor screens.
 - **Center:** Center shows a crosshair for the center point and x- and y-axis of the view.
 - **Film:** Some frame formats include film guides preset for you, whereas some will require customization. The film guides can be customized in the Preferences > Frame Format window.

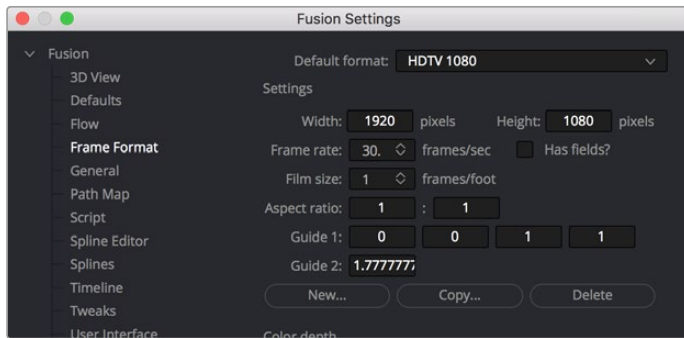


The Guides submenu in the viewer's contextual menu

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 performance.

Domain of Definition (DoD)

In compositing, the Domain of Definition, frequently abbreviated to DoD, refers to a rectangular region that defines what part of an image actually contains data. DoD makes the concept of an image's actual frame somewhat flexible, since rendering is no longer limited to the actual width and height of the image. This has two effects on the way Fusion renders images.

Firstly, nodes will no longer be required to render portions of the image that will not be affected by the node. This helps the renderer to optimize its performance. Secondly, Fusion can now keep track of and apply a node's effect to pixels that lie outside the visible portion of the image.

For example, consider the output of a Text+ node rendered against a transparent background. The text occupies only a portion of the pixels in the image. Without Domain of Definition, you would be required to process every pixel in the image needlessly. With a DoD, you are able to optimize effects applied to the image, producing faster results and consuming less memory in the process.

The following image shows an image with the DoD outlined.



The DoD is shown as two XY coordinates indicating the corners of an axis-aligned bounding box (in pixels)

For the most part, the DoD is calculated automatically and without the need for manual intervention. For example, all the nodes in the Generator category automatically generate the correct DoD. For nodes like Fast Noise, Mandelbrot, and Background, this is usually the full dimensions of the image. In the case of Text+ and virtually all of the Mask nodes, the DoD will often be much smaller or larger.

The OpenEXR format is capable of storing the data window of the image, and Fusion will apply this as the DoD when loading such an image through a Loader node and will write out the DoD through the Saver node.

When using the Fusion page in DaVinci Resolve, clips from the Edit page timeline or Media Pool will typically have the DoD default to the full image width of the source media. The exception is media stored in OpenEXR format.

The DoD is established as soon as the image is created or loaded into the composition. From there, it passes downstream, where viewers combine it with their Region of Interest in order to determine exactly what pixels should be affected by the node. As you work, different nodes will automatically shrink, expand, or move the DoD as they apply their effect to an image, causing the DoD to change from node to node.

Showing the DoD

If the current DoD for a node is different from the frame size of that image, it's shown in the tooltip that appears when the pointer hovers over a node in the Node Editor. The DoD is also visible in the viewer when you right-click in a viewer and choose Region > Show DoD from the contextual menu.

Setting the DoD Manually in the Node Editor

It is also possible to set the DoD for an image manually using the Tools > Miscellaneous > Auto Domain node in the Effects Library. This node can be useful when dealing with pre-created media that does not occupy the full image dimensions. For example, a rendering of a 3D character that walks toward the camera will frequently occupy only a portion of the image. The Auto Domain node can be used to animate a DoD that covers the character and ignores the rest of the image, making image processing more efficient.

Region of Interest (RoI)

The Region of Interest, frequently abbreviated to RoI, is a rectangular region similar to the Domain of Definition. However, unlike the DoD, which tells the node what pixels are actually present in the image, the RoI tells the node which pixels actually need to be rendered. When a node renders, it intersects the current RoI with the current DoD to determine what pixels should be affected.

Enabling RoI Controls

You can turn on the RoI controls to restrict rendering to a small region of the image to significantly improve performance when you're only working on a small part of a high-resolution or complex composition. For example, if you're using paint to clean up some holes in a matte on the floor of a composition with many, many high-resolution layers, 3D, and Lighting operations, you can use the RoI controls to isolate the part of the floor you're working on, which makes caching that part of the composition much faster.

To enable the RoI controls, do one of the following:

- Click the RoI button in the 2D Viewer toolbar.
- Right-click in a viewer and choose Region > Show 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 Rol controls, do one of the following:

- Drag any edge of the Rol rectangle to adjust one side of the Rol.
- Drag a corner to adjust the size of the Rol rectangle from that corner.
- Drag the small circle found at the top left corner of the Rol rectangle to move the Rol without adjusting its dimensions.

Sometimes, it's faster to simply draw a rectangle where you want the Rol to be.

To quickly draw the Rol at the desired size:

- 1 Choose Set from the viewer menu next to the Rol button, or right-click anywhere within the viewer and choose Region > Set Region.

- 2 When the pointer turns into an Rol drawing cursor, drag within the viewer to set a Rol rectangle.

Alternatively, an Auto command sets the Rol to fit whichever pixels are visible at the current zoom/pan level in the viewer. This lets you quickly limit the Rol to whatever part of the composition you've zoomed into.

To automatically draw the Rol:

- Choose Auto from the viewer menu next to the Rol button.
- Right-click anywhere within the viewer and choose Region > Auto Region.

When you no longer need to use the Rol, you can reset it.

To reset the Rol to the full width and height of the current image, do one of the following:

- Choose Reset from the viewer menu next to the Rol button.
- Right-click anywhere within the viewer and choose Region > Reset Region from the contextual menu or from the toolbar button menu.
- Disable the ROI control, which will also reset it.

While the Rol Is Active

The Rol is only used for previewing your composition while you work, not for output from Fusion. While the Rol is active, Fusion will only request rendering of the pixels inside the region when it displays an image in that viewer. Flipbook Previews that you create in that viewer will also respect the current Rol. MediaOut and Saver nodes will always use the full image dimensions when writing the image to disk, ignoring any Rol you've set in the viewers.

The Rol improves not only rendering speed and memory use, but it can also reduce file I/O, since Loaders and MediaIn nodes only load pixels from within the Rol, 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 Rol 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 Rol. This can be useful for comparing changes made within the Rol with a previous node state.

TIP: Right-clicking in a viewer and choosing Options > Show Controls for showing onscreen controls will override the Rol, forcing renders of pixels for the entire image.

Managing Viewer Lookup Tables (LUTs)

Lookup Tables, or LUTs, can be used to help match the appearance of a viewer to its eventual output destination. They're essentially image-processing operations that affect only the image being previewed in the viewer, not the image data itself. There are two basic ways that LUTs can calculate color transformations: The first is a simple 1D LUT, and the second is a more sophisticated 3D LUT.

- The simplest form of a LUT is a 1D LUT. It accounts for one color channel at a time, so it can make overall tonality changes but not very specific color changes.
- A 3D LUT looks at each possible color value (red, green, and blue) independently. A 3D LUT allows for large global changes as well as very specific color changes to be applied to images very quickly.

How Lookup Tables Work in Fusion

A Lookup Table (LUT) is a table of values used to transform the color and luminance of an image. A 1D LUT uses a two-column table for input color and output color, while a 3D LUT uses more of a matrix. A LUT is used primarily to correct for variances in the monitor or the source color space of the image. You can choose to apply a LUT to all the viewers or apply different LUTs to each viewer.

Image LUTs

Image LUTs can be applied to each viewer. In fact, you can even apply separate Image LUTs for the A and B buffers of a single viewer. These LUTs can only be applied to 2D images and not to 3D scenes. Image LUTs are routinely used to get from one scene-referred color space to another. For example, if you're working with log-encoded media but want to see how the image will look in the final color space, you can choose a LUT to make the image transform as a preview.

Buffer LUTs

The Buffer LUT is applied to the viewers regardless of contents, including 3D scenes, 3D materials, and subview types. Only one Buffer LUT can be applied. If a 2D image is being displayed with an Image LUT applied, then the Buffer LUT is applied to the result of the image LUT. Buffer LUTs are typically used to simulate another output color space that's unique to the display you're using—for instance, making a DCI-P3 projector show the image as it would look on an sRGB monitor.

To use a Buffer LUT:

- 1 Disable the LUT button above the viewer.
- 2 Right-click in the viewer and choose Global Options > Buffer LUT > Enable.
- 3 Right-click in the viewer and choose Global Options > Buffer LUT > Type of LUT you want to apply.

When dealing with nonlinear files from many of today's digital cinema cameras, a modern workflow would be to convert everything to linear at the beginning of the node tree, then create your composite, and then apply an Image LUT or Buffer LUT that matches the color space you want it to be in for either grading in the Color page or for final output.

However, in more elaborate production pipelines, you may need to apply multiple LUTs consecutively.

Types of Viewer LUTs

Aside from the industry standard 1D and 3D LUTs, other types of LUTs are supported, including script-based Fuse node LUTs and macros assembled from standard nodes. Generally, LUT processing is performed on the graphics card's GPU in real time, although the performance of macro-based LUTs is based on the nodes they contain.

Fusion View LUT

The Fusion View LUT is the default and is a frequently used LUT type. It provides an RGBA curve that can be used to assign IN/OUT value pairs. This control is identical to that provided by the Color Curve node.

Since the purpose of the View LUT is to provide an unchanging correction for the monitor or the file's color space, however, these splines cannot be animated.

Log-Lin View LUT

The Log-Lin LUT converts logarithmic data to linear, and vice versa. This can be particularly useful when used in conjunction with supplied LUT files that expect logarithmic data. It is similar to the Cineon Log node.

Gamut View LUT

The Gamut LUT converts a source color space to an output color space, with options to deal with gamma settings, alpha channels, and premultiplication. The Gamut LUT is a frequently used LUT type to correct the viewer when working with Linear Gamma in the Node editor.

Macro LUTs

Any macro node can also be used as a viewer LUT simply by saving the macro's .setting file to the correct Fusion directory.

In DaVinci Resolve, LUTs are saved in the following locations:

- **On macOS:** Macintosh HD/Users/username/Library/Application Support/Blackmagic Design/Fusion/LUTs/
- **On Windows:** C:\Program Files\Blackmagic Design\Fusion\LUTs
- **On Linux:** home/username/.local/share/DaVinciResolve/Fusion/LUTs

In Fusion Studio, LUTs are saved in the following locations:

- **On macOS:** Macintosh HD/Users/username/Library/Application Support/Blackmagic Design/Fusion/LUTs/
- **On Windows:** C:\Users\username\AppData\Roaming\Blackmagic Design\Fusion\LUTs
- **On Linux:** home/username/.fusion/BlackmagicDesign/Fusion/LUTs

For this to work, the macro must have one image input and one image output. Any controls exposed on the macro will be available when the Edit option is selected for the LUT. For more information about creating macros, see Chapter 66, “Node Groups, Macros, and Fusion Templates,” in the DaVinci Resolve Reference Manual or Chapter 6 in the Fusion Reference Manual.

LUT Presets


All LUTs available to DaVinci Resolve are also accessible to the Fusion page, which includes custom LUTs you've installed, as well as preset LUTs that come installed with DaVinci Resolve, such as the highly useful VFX IO category that includes a wide variety of miscellaneous to Linear and Linear to miscellaneous transforms. All of these LUTs appear by category in the viewer LUT menu.

Fuse LUTs

Fuses are scriptable 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 the Fusion Scripting Guide located on the Blackmagic Design website.

Using Viewer LUTs

Viewer LUTs can be enabled, edited, and turned off using the viewer LUT button and menu, as well as by using the viewer contextual menu. This menu shows all LUTs available to Fusion, including custom LUTs you've installed yourself.



The viewer LUT button and menu from the Fusion page in DaVinci Resolve

To turn the current viewer LUT on and off:

- Click the LUT button in the viewer toolbar to toggle the viewer LUT on and off.
- The LUT menu can also be found as a submenu in the viewer's contextual menu.

To choose another viewer LUT:

- Open the menu to the right of the viewer LUT button and choose an option from the viewer LUT menu.

To apply a Buffer LUT:

- 1 Right-click anywhere within the viewer and choose Global Options > Buffer LUT > Enable.
- 2 To choose a specific Buffer LUT, right-click again and choose a LUT from the Global Options > Buffer LUT submenu.

Buffer LUTs are often useful for applying monitor corrections, which do not usually change between projects.

To remove a Buffer LUT:

- Right-click anywhere within a viewer and choose Global Options > Buffer LUT > Enable to uncheck it.

Editing Viewer LUTs


The viewers are the primary area where composites are assessed, so it's crucial that they provide an accurate representation of what the content will look like when it's played for an audience. The LUT Editor allows you to customize your viewer's output to match the gamma and color characteristics of your eventual playback device, or to test how the current image looks in a completely different color space, or how it holds up over a range of different color spaces.

To open any editable viewer LUT option's Editor:

- 1 Click the LUT button in the viewer toolbar to enable it.
- 2 Do one of the following:
 - Choose Edit from the bottom of the viewer LUT menu.
 - Right-click in the viewer and then choose LUT > Edit from the contextual menu.

Editing the Fusion View Lookup Table

Similarly to the Color Curves node, the Fusion View LUT Editor uses spline-based color correction. In addition to the ability to modify the separate color channels, the LUT has Gain and Gamma sliders. The Gain slider is helpful for temporarily brightening or darkening the viewed image, allowing easier examination of shadow or highlight detail. The Color Gamma and Alpha Gamma sliders are used to duplicate the gamma values of the eventual output device. Video monitors, for example, commonly have a gamma of 1.7, while computer monitors can range anywhere from 1.6 to 2.2. Alpha Gamma is applied only when viewing the alpha channel of an image, or when viewing masks.



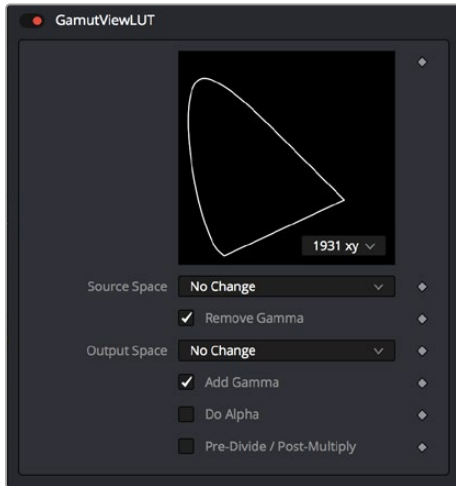
The LUT Editor for the default Fusion View LUT

Editing the Gamut View LUT

The Gamut View LUT Editor lets you choose a Source and Output color space to guide the viewer transform.

The Remove and Add Gamma checkboxes let you choose to do the gamut conversion with linear or nonlinear gamma, or they let you simply remove or add the appropriate gamma values without changing the color space.


Selecting the Pre-Divide/Post-Multiply checkbox will cause the image's pixel values to be divided by the alpha values prior to this conversion, and then re-multiplied by the alpha value after this conversion. This helps to avoid the creation of illegally additive images, particularly around the edges of a blue/green key or when working with 3D rendered objects.



The Gamut View LUT Editor

Editing the Log-Lin LUT

The Log-Lin LUT lets you apply a Log to Lin or Lin to Log operation using the Mode pop-up menu. You can choose the type of log-encoding to process from the Log Type drop-down, and choose whether to lock the R, G, and B channels together. A level adjustment lets you redefine the digital range of values used for the output, while Soft Clip (Knee), Film Stock Gamma, and Conversion Gamma sliders let you further customize the color transform. Lastly, a Conversion Table field and Browse button let you add an additional LUT as part of this operation.




The Log-Lin LUT Editor

LUT Processing Order

In elaborate workflows, facilities may apply multiple LUTs in a row before the image is seen. The order of these is important since each LUT delivers different outputs. For instance, for a Cineon file in Log color space you may often apply three LUTs. First a Log to Lin conversion, followed by a Fusion View LUT to apply a color calibration, and a third one to correct it for display on an sRGB monitor, or replace the last with a 3D DCP LUT if you are viewing on a projector.

When you select a node to be displayed, the image produced is processed before it is shown in the viewers. The processing order is slightly different for 2D images and 3D scenes.

2D images first have the image LUT applied, and the result is composited over the checker underlay. 3D scenes are instead rendered with OpenGL.



The order of processing for 2D images and 3D scenes

For either 2D or 3D, the result may be drawn to an offscreen buffer where a Buffer LUT can be applied, along with dithering, a full view checker underlay, and any stereo processing. The final result is then drawn to the viewer and any onscreen controls are drawn on top.

Applying Multiple LUTs

The viewer contextual menu can be used to apply multiple image LUTs into a processing chain.

To apply an additional LUT, do the following:

- 1 Right-click anywhere within the viewer.
- 2 From the viewer's contextual menu, choose LUT > Add New.
- 3 From the Add New submenu, choose a LUT to add.

To remove a LUT other than the first LUT, do the following:

- 1 Right-click anywhere within the viewer.
- 2 From the viewer's contextual menu, choose LUT > Delete.
- 3 From the Delete submenu, choose a LUT to remove.

A complete stacked LUT configuration can be saved to and loaded from a .viewlut file, as described below.

Saving Custom LUTs

There are a variety of ways to create and use different viewer LUTs in Fusion. You can save LUTs when you save viewer settings, you can import LUTs that have been exported from Fusion or other applications, and you can open any one of the various supported LUT file types. In addition, you can use the standard nodes in Fusion to create macros, which can then be saved and used as a LUT.

LUT Settings

The most straightforward way to save a LUT you have created using the Fusion View LUT Editor is to use the LUT > Save menu found in the viewer contextual menu. The settings are saved as an ASCII file with the extension .viewlut in the LUTs folder. Any files with this extension found in that folder will appear in the Image LUT menus for ease of loading. You can also load the settings that are not found in the menu by choosing LUT > Load from the viewer's contextual menu.

Using Viewer Settings

If you've modified a LUT, choosing Settings > Save New from the viewer's contextual menu will save all the viewer's settings, including all LUT curves and gain/gamma values. You can save these under different names, and each settings file can be reloaded at any time by choosing Settings > filename from the viewer's contextual menu. Choosing Save Default from the same menu will make these settings the standard for all new comps.

Using LUT Curves

The Viewer LUT Edit dialog can be used to import and export LUT curves. You can export the LUT curves as either ASCII or Saved format. The ASCII (.alut) file format is useful for sharing LUT curves with other software, whereas the Saved (.lut) file format is preferred for Fusion, as it is more compact, accurate, and allows further editing.

To export a LUT, do the following:

- 1 Click the viewer LUT button to enable it.
- 2 Click the viewer LUT menu, and then choose Edit.
- 3 Right-click on the LUT Curve Editor, and then choose Export LUT.
- 4 Select a LUT format at the bottom of the file browser window.
- 5 Enter a name for the LUT and click Save.

The Import LUT option will load LUT files back into the Curve Editor, or alternatively, if the file has been saved in Fusion's LUTs folder, it will appear in the LUT drop-down menu list.

TIP: This is one way to move LUTs between viewers or to and from the Color Curves node or any other LUT Editor in Fusion.

LUT Files

Any supported LUT files in the LUTs folder can be used by choosing them either from the LUT drop-down menu or the viewer's contextual menu. This includes 1D and 3D LUTs such as Fusion's .lut, .alut and .alut3 formats, as well as .cube, .shlut, .look, .3dl, and .itx formats. This is a convenient way to access standard format LUT files for different projects.

Settings and Macros

Since LUTs are a form of color correction, you can also use any node, macro, or group of nodes as a viewer LUT.

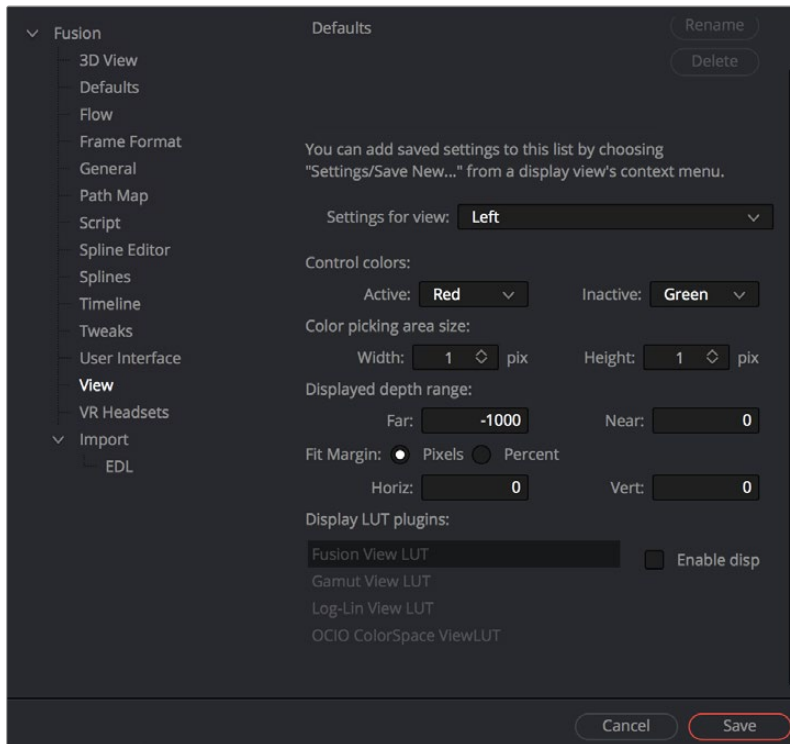
To use a node, group, or macro as a viewer LUT, do the following:

- 1 Select the node, group, or macro.
- 2 Right-click over the selected node, and then choose Settings > Save As from the menu.
- 3 In the file browser, go to the LUTs folder as set in Preferences > Global > Path Map > LUTS.
- 4 Click Save to save the .settings file.

This allows almost any combination of nodes to be used as a viewer LUT. This is the most flexible approach but is also potentially the slowest. The LUT nodes must be rendered solely on the CPU, whereas other methods are GPU-accelerated.

Setting a Default LUT

The default LUT applied when a new composition is created can be assigned in the Viewer panel of the Fusion Settings window. Clicking the Enable Display LUT checkbox allows you to select a LUT from the Display LUT plug-ins list.



The LUT default settings found in the View panel of the Fusion Settings window

Viewer Preferences and Settings

The default settings for each viewer can be changed using the Viewer panel in the Preferences. The position and size of each floating viewer can also be saved using the Layout menu in the Preferences.

Viewer Settings

It is often preferable to switch between entirely different viewer configurations while working. For example, while keying, the image may be in the main viewer, and the alpha channel may be in a subviewer. Viewer settings toward the end of a project may consist of the histogram, vectorscope, and waveform, as well as the image in a view set to Quad view.

Fusion provides the ability to quickly load and save viewer settings to help reduce the amount of effort required to change from one configuration to another.

To save a viewer setting, do the following:

- 1 Right-click over the viewer you want to save.
- 2 From the contextual menu, choose Setting > Save New.
- 3 Enter a name for the settings and click Save.

To load a viewer setting, do the following:

- 1 Right-click over the viewer you want to load a setting into.
- 2 From the contextual menu, choose Settings > filename.

Loading and Saving Defaults for a Viewer

The viewer can save new defaults and be returned to its defaults using the Load Defaults and the Save Defaults options in the Settings portion of the View contextual menu.

The Viewer Options Menu

The Options menu of the viewer contains several ways you can customize the look and behavior of the viewer. Many of these options are also in the viewer contextual menu.

Show Controls

When onscreen controls are not necessary or are getting in the way of evaluating the image, you can temporarily hide them using the Show Controls option. This option is toggled using Command-K.

Checker Underlay

The Checker Underlay shows a checkerboard beneath transparent pixels to make it easier to identify transparent areas. This is the default option for 2D viewers. Disabling this option replaces the checkerboard with black.

Pixel Grid

Enabling this option will show a light black grid that outlines the exact boundaries of pixels in the image when the image is scaled past a certain threshold. The default is Off.

Smooth Resize

The Smooth Resize option uses a smoother bilinear interpolated resizing method when zooming into an image in the viewer. When Smooth Resize is disabled, scaling uses the nearest neighbor method and shows noticeable aliasing artifacts but is more useful for seeing the actual pixels of the viewed image when you zoom all the way down to a pixel level since there is no interpolation. This option is enabled by default and can be toggled by clicking on the SmR button in the viewer toolbar.

Show Square Pixels

Depending on the frame format preferences and the type of footage loaded, many images may have pixels that are rectangular instead of square. Both the NTSC and PAL video standards, as well as some anamorphic film formats, use rectangular pixels. A computer monitor uses perfectly square pixels. To compensate for this, aspect correction is automatically performed when viewing non-square pixels. This prevents non-square pixel images from appearing squashed or stretched in the viewer.

You can enable the Show Square Pixels option to override the aspect correction. Show Square Pixels can also be toggled on and off using the 1:1 button in the viewer toolbar.

Gain/Gamma

Exposes or hides a simple pair of Gain and Gamma sliders that let you adjust the viewed image. Especially useful for “gamma slamming” a composite to see how well it holds up with a variety of gamma settings. Defaults to no change.

360° View

Sets the Fusion page viewer to properly display spherical imagery in a variety of formats, selectable from this submenu. Disable toggles 360 viewing on or off, while Auto, LatLong, Vert Cross, Horiz Cross, Vert Strip, and Horiz Strip let you properly display different formats of 360° video.

Locking the Viewer (Command-L)

You can lock a viewer to prevent it from updating. The node that's loaded into that viewer still processes and the new image is queued for display in the viewer, but until you unlock it, the viewer does not update. By default, the viewer is unlocked.

Additional Viewer Options

There are additional commands when you right-click anywhere within a viewer and choose from the generically named Options submenu.

Alpha Overlay

When you enable the alpha overlay, the viewer will show the alpha channel overlaid on top of the color channels. This can be helpful when trying to see where one image stops and another begins in a composite. This option is disabled by default.

Overlay Color

When you turn the alpha overlay on, the default color is to show white for the area the alpha covers. There are times when white does not show clearly enough, depending on the colors in the image. You can change the color by choosing a color from the list of Overlay Color options.

Follow Active

Enabling the Follow Active option will cause the viewer to always display the currently active node in the Node Editor. This option is disabled by default, so you can view a different node than what you control in the Control Panel.

Show Controls

When onscreen controls are not necessary or are getting in the way of evaluating the image, you can temporarily hide them using the Show Controls option. This option is toggled using Command-K.

Show Full Color Range


When working with floating-point images, you will occasionally need to visualize the values that fall outside the normal luminance range. Enabling the Show Full Color Range option using the toolbar button automatically normalize any image displayed in the viewer. Normalization causes the brightest pixel in a color channel to be mapped to a value of 1.0 (white) and the darkest pixel to be mapped to a value of 0.0 (black). Midrange values are scaled appropriately to fit within that range. It is also useful when viewing Z-buffer or other auxiliary channels, which often use value ranges far different from those in the color channels.

Show Labels

The Show Labels option lets you toggle the display of the text that sometimes accompanies onscreen controls in the viewer without disabling the functions that are showing those overlays, and without hiding the onscreen controls themselves.

Status Bar Information

The status bar at the bottom of the Fusion window provides the exact RGBA and Z values for the pixel beneath the pointer when it's hovering within one of the viewers. Additional information about the X and Y coordinates of the cursor and the exact pixel position are also displayed.



The status bar showing coordinates and color information

Chapter 68

Editing Parameters in the Inspector

The Inspector is where you adjust the parameters of each node to do what needs to be done. This chapter covers the various node parameters and methods for working with the available controls.

Contents

Overview of the Inspector	1180
The Tools and Modifiers Panels	1181
Customizing the Inspector	1181
Inspector Height	1181
Inspector Display Preferences	1181
Opening Nodes in the Inspector	1182
Pinning Multiple Nodes in the Inspector	1183
Hiding Inspector Controls	1184
Using the Inspector Header	1184
Selecting and Viewing Nodes in the Inspector	1185
Using Header Controls	1185
Versioning Nodes	1186
Parameter Tabs	1186
The Settings Tab	1186
Inspector Controls Explained	1190
Fusion Slider Controls	1190
Thumbwheel	1191
Range Controls	1191
Checkboxes	1191

Drop-Down Menus	1192
Button Arrays	1192
Color Chooser and Picker	1192
Gradients	1194
Modifiers	1196
Animating Parameters in the Inspector	1196
Removing Animation From a Parameter	1197
Attaching a Parameter to an Existing Animation Curve	1197
Connecting Parameters	1198
Connecting Parameters by Publishing	1198
Connecting Parameters by Pick Whipping	1198
Contextual Menus	1199
Customizing Node Parameters with User Controls	1199
An Example of Customizing Directional Blur	1200

Overview of the Inspector

While the creation and connection of nodes in the Node Editor determines the tools and order of operations that make up a composition, the Inspector (previously called the Control Panel) is where you adjust the various parameters inside each node to do what needs to be done.



Inspector displays the Brightness Contrast controls

This chapter covers methods for opening node parameters in the Inspector to edit them in different ways according to the type of available controls.

To display the Inspector:

- Click the Inspector button on the UI toolbar.

The Tools and Modifiers Panels

The Inspector is divided into two overall panels.

- The Tools panel is where the parameters of selected nodes appear so you can edit them.
- The Modifiers panel is where you edit optional extensions to the tool's standard toolset as well as automated expressions that you can attach to individual parameters to create animated effects. Additionally, certain nodes such as the Paint node generate data such as Strokes, which are saved in the Modifiers panel.




Modifiers displayed in the Modifiers panel

Customizing the Inspector

You can customize how the Inspector is presented in a variety of ways.

Inspector Height

A small arrow button at the far right of the UI toolbar lets you toggle the Inspector between full-height and half-height views, depending on how much room you need for editing parameters.



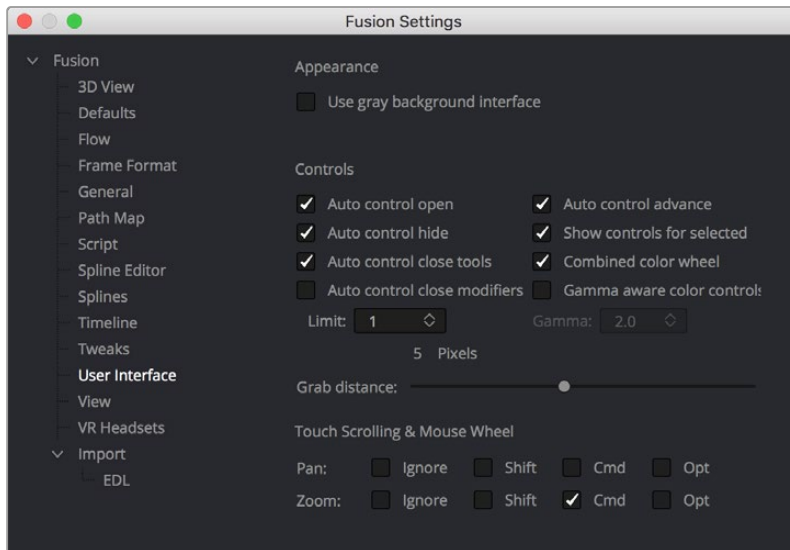
The Maximize button on the left side of the Inspector

In maximized height mode, the Inspector takes up the entire right side of the UI, letting you see every control that a node has available, or creating enough room to see the parameters of two or three pinned nodes all at once. In half-height mode, the top of the Inspector is aligned with the tops of the viewers, expanding the horizontal space that's available for the Node Editor.

Inspector Display Preferences

By default, you see only selected nodes in the Inspector, and only the Active node is expanded to show its controls. You can change this behavior by choosing Fusion > Fusion Settings in the Fusion

page or File > Preferences in Fusion Studio and opening the User Interface panel. In the User Interface, checkboxes manage the display of controls.



Control preferences in the User Interface category

- **Auto Control Open:** When enabled (the default), whichever node is active automatically opens its controls in the Inspector. When disabled, selecting an active node opens that node's Inspector header in the Inspector, but the parameters remain hidden unless you click the Inspector header.
- **Auto Control Hide:** When enabled (the default), only selected nodes are visible in the Inspector, and all deselected nodes are automatically removed from the Inspector to reduce clutter. When disabled, parameters from selected nodes remain in the Inspector, even when those nodes are deselected, so that the Inspector accumulates the parameters of every node you select over time.
- **Auto Control Close Tools:** When enabled (the default), only the parameters for the active node can be exposed. When disabled, you can open the parameters of multiple nodes in the Inspector if you want.
- **Auto Controls for Selected:** When enabled (the default), selecting multiple nodes opens multiple control headers for those nodes in the Inspector. When disabled, only the active node appears in the Inspector; multi-selected nodes highlighted in white do not appear.


Opening Nodes in the Inspector

Before you can edit a node's parameters, you need to open it in the Inspector.

To display a node's controls in the Inspector:

- Select one or more nodes from the Node Editor, Keyframes Editor, or Spline Editor.

When you select a single node so that it's highlighted orange in the Node Editor, all of its parameters appear in the Inspector. If you select multiple nodes at once, Inspector headers appear for each selected node (highlighted in white in the Node Editor), but the parameters for the active node (highlighted in orange) are exposed for editing.



Opening multiple nodes in the Inspector

Only one node's parameters can be edited at a time, so clicking another node's Inspector header opens that node's parameters and closes the parameters of the previous node you were working on. This also makes the newly opened node the active node, highlighting it orange in the Inspector.

Pinning Multiple Nodes in the Inspector

For instances where you need to work quickly by editing the parameters of multiple nodes at the same time, you can use the Pin button in the Inspector header of nodes in the Inspector to keep those parameters exposed in the Inspector, regardless of whether that node is selected and active.



The Pin button of a node's Inspector header in the Inspector

While the Pin button is on, that node's parameters remain open in the Inspector. If you select another node in the Node Editor, that node's parameters appear beneath any pinned nodes.



A pinned node on the bottom, with a selected node at the top

You can have as many pinned nodes in the Inspector as you like, but the more you have, the more likely you'll need to scroll up or down in the Inspector to get to all the parameters you want to edit. To remove a pinned node from the Inspector, just turn off its Pin button in the Inspector header.

Hiding Inspector Controls

If you like, Inspector parameters for specific nodes can be hidden so they never appear, even when that node is selected. This can be useful for preventing accidental changes by you or other compositors who may be working on a composition in situations where you don't want to lock the node.

To Toggle the Inspector controls for a node or off:

- Right-click on the node in the Node Editor, or on the Inspector header, and choose Modes > Show Controls from the contextual menu.

Using the Inspector Header

When you select a node, it populates the Inspector with a title bar, or Inspector header, that displays that node's name as well as other controls that govern that node. A node's Inspector header itself has a variety of controls, but clicking (or double-clicking) on an Inspector header also exposes that node's parameters.



A node's Inspector header

When you select multiple nodes at once, you'll see multiple headers in the Inspector. By default, only the parameters for the active node (highlighted orange in the Node Editor) can be opened at any given time, although you can change this behavior in Fusion's Preferences.

Selecting and Viewing Nodes in the Inspector

Inspector headers are click targets for selecting nodes, opening and closing node parameters, and other things.

Methods of using headers:

- **To select a node using the Inspector header:** When multiple nodes are selected, you can make a node the active node by clicking its Inspector header in the Inspector. As the actively selected node, the Inspector header and the corresponding node in the Node Editor are highlighted orange, and its parameters are exposed.
- **To load a node into the viewer using the Inspector header:** You can view a node by dragging its header into one of the viewers.
- **To view a node's splines with the control header:** If you want to view the animated curves of a node in the Spline Editor, you can add them by dragging the Inspector header into the Spline Editor. All animated splines for the parameters of that node will automatically be displayed.


Using Header Controls

The controls found in each node's Inspector header makes it fast to do simple things.

- **To turn nodes off and on:** Each Inspector header has a toggle switch to the left of its name, which can be used to enable or disable that node. Disabled nodes pass image data from the previous upstream node to the next downstream node without alteration.
- **To change the Inspector header name:** The name of the node corresponding to that Inspector header is displayed next. You can change the name by right-clicking the Inspector header to expose contextual menu commands similar to those found when you right-click a node in the Node Editor and choosing Rename. Alternatively, you can click an Inspector header and press F2 to edit its name. A Rename dialog appears, where you can enter a new name and click OK (or press Return).
- **To color-code nodes:** A color pop-up menu lets you color code with one of 16 colors. Choose Clear Color if you want to return that node to the default color.
- **To version nodes:** Turning on the Versions button displays a Version bar with six buttons. Versioning is described in the following section.
- **To pin Inspector controls:** Clicking the Pin button "pins" that node's parameters in the Inspector so they remain in place, even if you deselect that node. You can have as many pinned nodes as you like in the Inspector, but the more you have, the more likely you'll be scrolling up and down the Inspector to navigate all the available parameters.
- **To lock nodes:** Clicking the Lock button locks that node so no changes can be made to it.
- **To reset Inspector controls:** The rightmost button in the Inspector header is a Reset button that resets the entire node to the default settings for that node.

Versioning Nodes

Each button is capable of containing separate parameter settings for that node, making it easy to save and compare up to six different versions of settings for each node. All versions are saved along with the node in the Node Editor for future use.




The Version bar, underneath an Inspector header, with versions enabled

An orange underline indicates the currently selected version, which is the version that's currently being used by your composition. To clear a version you don't want to use any more, right-click that version number and choose Clear from the contextual menu.


Parameter Tabs

Underneath the Inspector header is a series of panel tabs, displayed as thematic icons. Clicking one of these icons opens a separate tab of parameters, which are usually grouped by function. Simple nodes, such as the Blur node, consist of two tabs where the first contains all of the parameters relating to blurring the image, and the second is the Settings tab.



The parameter tabs of the Blur node

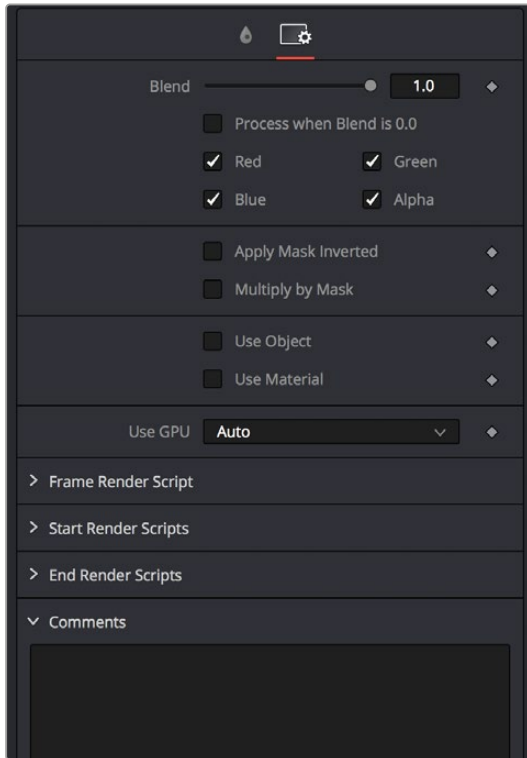
More complicated nodes have more tabs containing more groups of parameters. For example, the Delta Keyer has seven tabs: separating Key, Pre-Matte, Matte, Fringe, Tuning, and Mask parameters, along with the obligatory Settings tab. These tabs keep the Delta Keyer from being a giant scrolling list of settings and make it easy to keep track of which part of the keying process you're finessing as you work.



The parameter tabs of the Delta Keyer node

The Settings Tab

Every node that comes with Fusion has a Settings tab. This tab includes a set of standard controls that appear for nearly every node, although some nodes have special Settings tab controls that others lack.



The Settings tab in the Inspector

The following controls are common to most nodes, although some are node-specific. For example, Motion Blur settings have no purpose in a Color Space node.

Blend


The Blend control is found in all nodes, except the Loader, MediaIn, and Generator nodes. It is used to blend between the node's unaltered image input and the node's final processed output. When the blend value is 0.0, the outgoing image is identical to the incoming image. Ordinarily, this will cause the node to skip processing entirely, copying the input straight to the output. The default for this node is 1.0, meaning the node will output the modified image 100%.

Process When Blend is 0.0

This checkbox forces the node to process even when the input value is zero and the image output is identical to the image input. This can be useful on certain nodes or third-party plug-ins that store values from one frame to the next. If this checkbox is disabled on nodes that operate in this manner, the node will skip being processed when the Blend is set to 0, producing incorrect results on subsequent frames.


Red/Green/Blue/Alpha Channel Checkboxes

Most nodes have a set of RGBA checkboxes in the Settings tab. These checkboxes let you exclude any combination of these channels from being affected by that node.



The channel limiting checkboxes in the Settings panel of a Transform node set so that only the green channel is affected

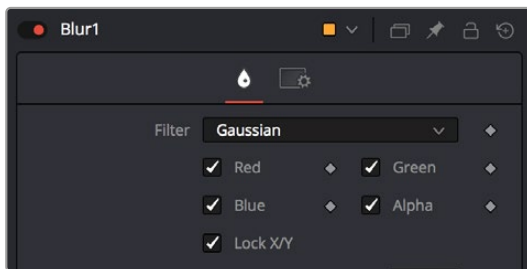
For example, if you wanted to use the Transform node to affect only the green channel of an image, you can turn off the Red, Blue, and Alpha checkboxes. As a result, the green channel is processed by this operation, and the red, blue, and alpha channels are copied straight from the node's input to the node's output, skipping that node's processing to remain unaffected.



Transforming only the green color channel of the image with a Transform effect

Skipping Channel Processing

Under the hood, most nodes actually process all channels first, but afterward copy the input image to the output for channels that have been unchecked. Modern workstations are so fast that this isn't usually noticeable, but there are some nodes where deselecting a channel actually causes that node to skip processing that channel entirely. Nodes that operate this way have a linked set of Red, Green, Blue, and Alpha checkboxes on another tab in the node.



Channel checkboxes on the Controls tab of the Blur node indicate that disabled channels won't be processed at all, to save rendering time.

In these cases, the Common Control channel checkboxes are instanced to the channel boxes found elsewhere in the node. Blur, Brightness/Contrast, Erode/Dilate, and Filter are examples of nodes that all have RGBY checkboxes in the main Controls tab of the Inspector, in addition to the Settings tab.

Apply Mask Inverted

When the Apply Mask Inverted checkbox is enabled, masks attached to the Effect Mask input of that node are inverted.

TIP: The Apply Mask Inverted checkbox option operates only on effects masks, not on garbage masks.

Multiply By Mask

Selecting this option will cause the RGB values of the masked image to be multiplied by the Mask channel's values. This will cause all pixels of the image not included in the mask (i.e., those set to 0) to become black. This creates a premultiplied image.

Use Object/Use Material (For Masking)

Some 3D animation and rendering software can output to file formats that support auxiliary channels. Notably, the OpenEXR file format supports Object ID and Material ID channels, either of which can be used as a mask for an effect. This checkbox determines whether the channels will be used if they are available. The specific Material ID or Object ID affected is chosen using the next set of controls.

Pick Controls

The Pick Controls are only displayed once the Use Object or Use Material checkbox is enabled. These controls select which ID is used to create a mask from the Object or Material channels saved in the image. You use the Pick button to grab IDs from the image in the viewer, the same way you use the Color Picker to select a color. The image or sequence must have been rendered from a 3D software package with those channels included.

Correct Edges

The Correct Edges checkbox is only displayed once the Use Object or Use Material checkbox is enabled. When the Correct Edges checkbox is enabled, the Coverage and Background Color channels are used to separate and improve the effect around the edge of the object. When disabled (or no Coverage or Background Color channels are available), aliasing may occur on the edge of the mask.

Motion Blur

For nodes that are capable of introducing motion, such as Transform nodes, Warp nodes, and so on, the Motion Blur checkbox toggles the rendering of motion blur on or off for that node. When this checkbox is enabled, the node's predicted motion is used to produce the blur caused by a virtual camera shutter. When the control is disabled, no motion blur is created.

When Motion Blur is disabled, no additional controls are displayed. However, turning on Motion Blur reveals four additional sliders with which you can customize the look of the motion blur you're adding to that node.

Quality

Quality determines the number of samples used to create the blur. The default quality setting of 2 will create two samples on either side of an object's actual motion. Larger values produce smoother results but will increase the render time.

Shutter Angle

Shutter Angle controls the angle of the virtual shutter used to produce the Motion Blur effect. Larger angles create more blur but increase the render times. A value of 360 is the equivalent of having the shutter open for one whole frame exposure. Higher values are possible and can be used to create interesting effects. The default value for this slider is 100.

Center Bias

Center Bias modifies the position of the center of the motion blur. Adjusting the value allows for the creation of trail-type effects.

Sample Spread

Adjusting Sample Spread modifies the weight given to each sample. This affects the brightness of the samples set with the Quality slider.

Scripting

Scripting fields are present on every node and contain one or more editable text fields that can be used to add scripts that process when that node is rendering. For more information on the contents of this tab, please consult the Scripting documentation.

Comments

A Comments field is found on every node and contains a single text field that is used to add comments and notes to that node. To enter text, simply click within the field to place a cursor, and begin typing.

When a note is added to a node, the comments icon appears in the Control Header and can be seen in a node's tooltip when the cursor is placed over the node in the Node Editor. The contents of the Comments tab can be animated over time, if required.

Additional controls appear under this tab if the node is a Loader. For more information, see Chapter 103, "Generator Nodes," in the DaVinci Resolve Reference Manual or Chapter 43 in the Fusion Reference Manual.


Inspector Controls Explained

Although a few nodes use fully customized interface elements that are unique to only that node, the vast majority of nodes use a mix of sliders, angle wheels, and checkboxes. This section explains how to use these controls.

Fusion Slider Controls

Slider Controls are used to select a single value from a range of values. You change the value by dragging the slider or entering a value into the edit box. This is fairly standard behavior for sliders. However, there is additional functionality that can increase your productivity when making changes with sliders.

Clicking on the gutter to the left or right of the handle will increase or decrease the value. Holding Command while clicking on the gutter will adjust the values in smaller increments. Holding Shift while clicking will adjust the value in larger increments.



Hold Command while clicking in the gutter to move in smaller increments

Once you click directly on a slider handle, you can make changes to its value using the Left and Right Arrow keys. The Command and Shift keys can again be used to modify the value in larger or smaller increments.

While slider controls use a minimum and maximum value range, entering a value in the edit box outside that range will often expand the range of the slider to accommodate the new value. For example, it is possible to enter 500 in a Blur Size control, even though the Blur Size sliders default maximum value is 100. The slider will automatically adjust its maximum displayed value to allow entry of these larger values.

If the slider has been altered from its default value, a small circular indicator will appear below the gutter. Clicking on this circle will reset the slider to its default.

Thumbwheel

A Thumbwheel control is identical to a slider except it does not have a maximum or minimum value. To make an adjustment, you drag the center portion left or right or enter a value directly into the edit box. Thumbwheel controls are typically used on angle parameters, although they do have other uses as well.



Thumbwheel controls for X,Y, and Z rotation with arrows on either end for fine-tuning adjustments

You can use the arrowheads at either end of the control to fine tune your adjustments. Once the thumbwheel has been selected either by dragging or using the arrow keys, you can use the Left and Right Arrows on your keyboard to further adjust the values. As with the slider control, the Command and Shift keys can be used to increase or decrease the change in value in smaller or larger increments.

If the thumbwheel has been altered from its default value, a small circular indicator will appear below above the thumbwheel. Clicking on this circle will reset the thumbwheel to its default.

Range Controls

The Range controls are actually two separate controls, one for setting the Low Range value and one for the High Range value. To adjust the values, drag the handles on either end of the Range bar. To slide the high and low values of the range simultaneously, drag from the center of the Range bar. You can also expand or contract the range symmetrically by holding Command and dragging either end of the Range bar. You find Range controls on parameters that require a high and low threshold, like the Matte Control, Chroma Keyer, and Ultra Keyer nodes.




A Matte Threshold Range control

TIP: You can enter floating-point values in the Range controls by typing the values in using the Low and High numeric entry boxes.

Checkboxes


Checkboxes are controls that have either an On or Off value. Clicking on the checkbox control will toggle the state between selected and not selected. Checkboxes can be animated, with a value of 0 for Off and a value of 1.0 or greater for On.



Checkboxes used to select options for tracking

Drop-Down Menus

Drop-down menus are used to select one option from a menu. Once the menu is open, choosing one of the items will select that entry. When the menu is closed, the selection is displayed in the Inspector.



Drop-down menu in the TimeSpeed node

Drop-down menu selections can be animated, with a value of 0 representing the first item in the list, 1 representing the second, and so forth.

Button Arrays


Button arrays are groups of buttons that allow you to select from a range of options. They are almost identical in function to drop-down menu controls, except that in the case of a button array it is possible to see all of the available options at a glance. Often button arrays use icons to make the options more immediately comprehensible.



The Lens Type button array in the Defocus node

Color Chooser and Picker

The Color panel is displayed wherever a parameter requires a color as its value, such as the Fill or Outline color in the Text+ node. The selected color is shown in a swatch with an Eyedropper to its right, and below the swatch is the Color Chooser.




The Color panel with transparency preview

The Color panel is extremely flexible and has four different techniques for selecting and displaying colors.


TIP: Color can be represented by 0–1, 0.255, or 0–65000 by setting the range you want in the Preferences > General panel.

macOS and Windows Color Nodes

Clicking on the color swatch will display the operating system's standard Color Selection node.



macOS Colors panel

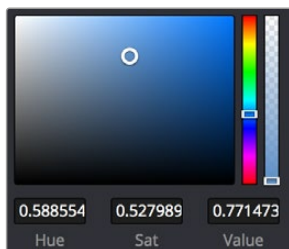


Windows Color dialog

Each operating system has a slightly different layout, but the general idea is the same. You can choose a color from the swatches provided—the color wheel on macOS, or the color palette on Windows. However you choose your color, you must click OK for the selection to be applied.

The Color Chooser


You also have access to the built-in color chooser, which includes sections for choosing grayscale values, as well as the currently chosen hue with different ranges of saturation and value. A hue bar and alpha bar (depending on the node) let you choose different values.



The color chooser in the Background node

Picking Colors from an Image

If you are trying to match the color from an image in the viewer, you can hold down the cursor over the Eyedropper, and then drag the pointer into the viewer. The pointer will change to an Eyedropper, and a pop-up swatch will appear above the cursor with the color you are hovering over and its values. When you are over the color you want, release the mouse button to set the color.



The Eyedropper with color swatch

The Color Picker normally selects from a single pixel in the image, but you can adjust the size of the selection by dragging into the viewer with the Eyedropper, and then holding Command and dragging out a rectangle for the sample size you want. The size change applies to all Color Pickers until the size is changed again.

Gradients


The Gradient Control bar is used to create a gradual blend between colors. The Gradient bar displays a preview of the colors used from start to end. By default, there are two triangular color stops: one on the left that determines the start color, and one on the right that determines the end color.



The default Gradient controls


Gradient Type

The Gradient Type button array is used to select the form used to draw the gradient. Linear draws the gradient along a straight line from the starting color stop to the ending color stop.




Linear gradient

Reflect draws the gradient by mirroring the linear gradient on either side of the starting point.




Reflect gradient

Square draws the gradient by using a square pattern when the starting point is at the center of the image.




Square gradient

Cross draws the gradient using a cross pattern when the starting point is at the center of the image.




Cross gradient

Radial draws the gradient in a circular pattern when the starting point is at the center of the image.



Radial gradient

Angle draws the gradient in a counter-clockwise sweep when the starting point is at the center of the image.



Angle gradient

Start and End Position

The Start and End Position controls have a set of X and Y edit boxes that are useful for fine-tuning the start and end position of the gradient. The position settings are also represented by two crosshair onscreen controls in the viewer, which may be more practical for initial positioning.

Gradient Colors Bar

The Gradient Colors bar is used to select the blending colors for the gradient. The default two color stops set the start and end colors. You can change the colors used in the gradient by selecting the color stop, and then using the Eyedropper or color wheel to set the new color.

You can add, move, copy, and delete colors from the gradient using the Colors bar.

To add a color stop to the Gradient Colors bar:

- 1 Click anywhere along the bottom of the Gradient Colors bar.
- 2 Use the Eyedropper or color wheel to set the color for the color stop.

To move a color stop on the Colors bar:

- Drag a color stop left or right along the Gradient Color bar.

To copy a color stop on the Colors bar:

- Hold Command while you drag a color stop.

To delete a color stop from the Colors bar, do one of the following:

- Drag the color stop up past the Gradient Colors bar.
- Select the color stop, then click the red X button to delete it.

Interpolation Space

The Gradient Interpolation Method pop-up menu lets you select what color space is used to calculate the colors between color stops.

Offset

When you adjust the Offset control, the position of the gradient is moved relative to the start and end markers. This control is most useful when used in conjunction with the repeat and ping-pong modes described below.

Once/Repeat/Ping-Pong

These three buttons are used to set the behavior of the gradient when the Offset control scrolls the gradient past its start and end positions. The Once button is the default behavior, which keeps the color continuous for offset. Repeat loops around to the start color when the offset goes beyond the end color. Ping-pong repeats the color pattern in reverse.

1x1, 2x2, 3x3, 4x4, 5x5

These buttons control the amount of sub-pixel precision used when the edges of the gradient become visible in Repeat mode, or when the gradient is animated. Higher settings will take significantly longer to render but will be more precise.

Gradient Contextual Menu

Gradients have their own contextual menu that you can bring up by right-clicking on the Gradient bar. In the Gradient contextual menu are options for animating, publishing, and connecting one gradient to another. There is also a gradient-specific modifier that builds a custom gradient by sampling colors from the output of a node in the Node Editor.

Modifiers

Modifiers are expressions, calculations, trackers, paths, and other mathematical components that you attach to a parameter to extend its functionality. When a modifier is attached to a parameter, its controls will appear separately in the Inspector Modifiers tab.

To attach a modifier:

- 1 Right-click over the parameter to which you want to attach a modifier.
- 2 Make a selection from the Modifier submenu in the contextual menu.

Animating Parameters in the Inspector

Fusion can keyframe most parameters in most nodes, in order to create animated effects such as animated transforms, rotoscoping with splines, dynamically altering warping behaviors, and so on; the list is endless. For convenience, a set of keyframing controls are available within the Inspector next to each keyframable parameter. These controls are:

- A gray Keyframe button to the right each keyframable parameter. Clicking this gray button creates a keyframe at the current position of the playhead, and turns the button orange.
- When you add a keyframe to a parameter, moving to a new frame and changing the parameter will automatically add a keyframe at the current position.

- Whenever the playhead is sitting right on top of a keyframe, this button turns orange. Clicking an orange Keyframe button deletes the keyframe at that frame and turns the button gray again.
- Small navigation arrows appear to the right and left if there are more keyframes in those directions. Clicking on navigation arrows to the right and left of keyframes jumps the playhead to those keyframes.



Orange Keyframe buttons in the Inspector show there's a keyframe at that frame

Once you've keyframed one or more parameters, the node containing the parameters you keyframed displays a Keyframe badge, to show that node has been animated.



A keyframed node displays a Keyframe badge in the Node Editor

Once you've started keyframing node parameters, you can edit their timing in the Keyframes Editor and/or Spline Editor. For more information about keyframing in Fusion, see Chapter 59, "Animating in Fusion's Keyframe Editor," in the DaVinci Resolve Reference Manual or Chapter 9 in the Fusion Reference Manual.

Removing Animation From a Parameter

To remove all keyframes from a parameter:

- 1 Right-click over the name of the keyframed parameter in the Inspector.
- 2 Choose Remove "node name:parameter name" from the contextual menu.

TIP: If you change the default spline type from Bézier, the contextual menu will display the name of the current spline type.

Attaching a Parameter to an Existing Animation Curve

Multiple parameters can be connected to the same animation curve. This can be an invaluable timesaver if you are identically animating different parameters in a node.

To connect a second parameter to the same animation curve:

- 1 Right-click on the second parameter you want to attach.
- 2 In the contextual menu, hover over the Connect To submenu.
- 3 In the Connect To submenu, choose the name of the animated parameter.

Connecting Parameters


It is often useful to connect two parameters together even without an animation curve. There are two methods you can use.

Connecting Parameters by Publishing

If you want to tie two parameters together so adjusting one adjusts the other, you must connect them together using the Publish menu command on the first parameter and the Connect menu command on the second parameter.

To Publish and Connect parameters:

- 1 Right-click the name of the parameter you want to publish, and choose Publish from the contextual menu.
- 2 Right-click on the second parameter you want to attach, and choose the name of the parameter you just published from the Connect To submenu.



The Publish contextual menu

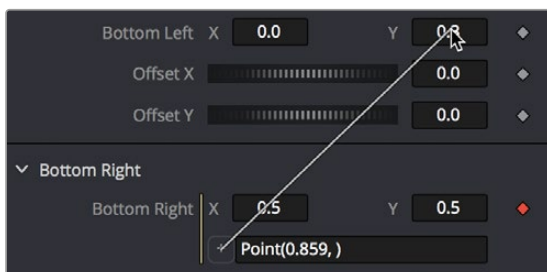
Connecting Parameters by Pick Whipping

You can also use simple expressions to link two parameters together. By using simple expressions via pick whipping, values can be connected and combined visually without the need to publish a value first. The pick whip is a temporary line drawn from one parameter to another in order to create a link between the two.

To link two parameters using a pick whip:

- 1 Double-click the field of a parameter you want to pick whip to another parameter, type `=`, and then press the Return key.
- 2 When Pick Whip controls appear underneath the parameter, drag a “whip” from the Add button to the target parameter.

Now, adjusting the target parameter automatically adjusts the original parameter.



Pick whipping one parameter to another

TIP: Disabling the Auto Control Close node’s General preference, and then selecting two nodes in the Node Editor will allow you to pick whip two parameters from different nodes.

The Expression field can further be used to add mathematical formulas to the value received from the target parameter.

For more information on Pick Whipping and Expressions, see Chapter 72, “Using Modifiers, Expressions, and Custom Controls,” in the DaVinci Resolve Reference Manual or Chapter 12 in the Fusion Reference Manual.

Contextual Menus

There are two types of contextual menus you can invoke within the Inspector.

Node Contextual Menus

To display the Node Context menu from the Inspector, right-click on the Inspector header. The node’s contextual menu includes the same menu options that are accessed by right-clicking on a node in the Node Editor. See Chapter 65, “Working in the Node Editor,” in the DaVinci Resolve Reference Manual or Chapter 5 in the Fusion Reference Manual for more information on these options.

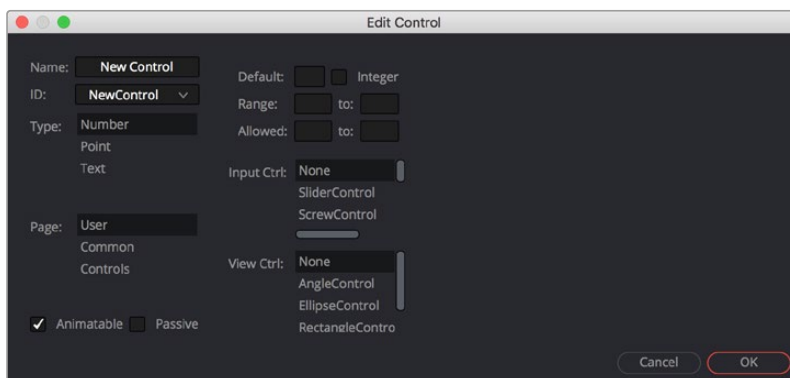
Parameter Contextual Menus

The contextual menu for individual parameters is accessed by right-clicking over the parameter’s name, slider, thumbwheel, range control, button array, or other control type. For example, right-clicking on a slider will provide the slider’s contextual menu, with options to animate the control or add additional modifiers. Many of these options were described in this chapter.

Customizing Node Parameters with User Controls

The user interface for each node in Fusion is designed to provide access to the parameters in a logical manner. Sometimes, though, you may want to add, hide, or change the controls. This is commonly done for simple expressions and macros, but it can be done for usability and aesthetic reasons for favorites and presets.

User custom controls can be added or edited via the Edit Control dialog. Right-click the name of a node in the Inspector (in the header bar) and choose Edit Control from the contextual menu. A new window will appear, titled Edit Control.



The Edit Control window


In the Input attributes, you can select an existing control or create a new one, name it, define the type, and assign it to a tab. In the Type attributes, you define the input controls, the defaults and ranges, and whether it has an onscreen preview control. The Input Ctrl attributes box contains settings specific to

the selected node control, and the View Ctrl attributes box contains settings for the preview control, if any.

All changes made using UserControls are stored in the node instance itself, so they can be copy/pasted, saved to a setting, added to the Bins, or added to your favorites.

An Example of Customizing Directional Blur

In the following example, let's suppose we wanted to create a more intuitive way of controlling a linear blur than using the Length and Angle sliders independently.



Default Directional Blur controls in the Inspector

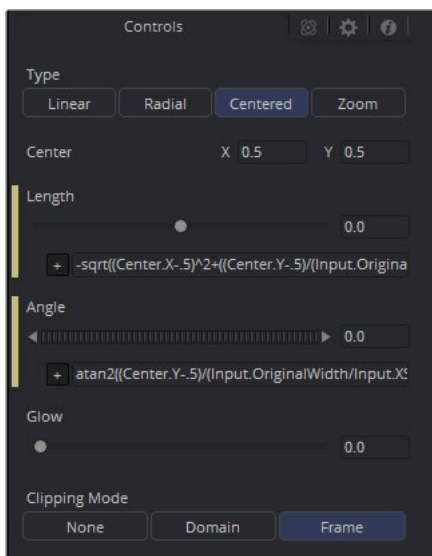
We could use the Center input control, along with its preview control, to set an angle and distance from directly within the viewer using expressions.

- 1 Right-click the label for the Length parameter, choose Expression from the contextual menu, and then paste the following expression into the Expression field that appears:

```
-sqrt(((Center.X-.5)*(Input.XScale))^2+((Center.Y-.5)*(Input.YScale)*(Input.Height/Input.Width))^2)
```

- 2 Next, right-click the label for the Angle parameter, choose Expression from the contextual menu, and then paste the following expression into the Expression field that appears:

```
atan2(Center.Y-.5)/(Input.OriginalWidth/Input.X , .5-Center.X) * 180 / pi
```



Directional Blur controlled by the Center's position

This functions fine, but the controls are confusing. The Center control doesn't work as the center anymore, and it should be named "Blur Vector" instead. The controls for the Length and Angle aren't meant to be edited, so they should be hidden away, and we're only doing a linear blur, so we don't need the buttons for Radial or Zoom. We just need to choose between Linear and Centered.

Adding Another Control

For the first task, let's rename the Center. From the Add Control window, select Center from the ID list. A dialog will appear asking if you would like to Replace, Hide, or Change ID. We'll choose Replace. Now we are editing the Center input. We'll change the Name to Blur Vector, set the Type to Point, and the Page to Controls, which is the first tab where the controls are normally. Press OK, and our new input will appear on our node in the Node Editor. The ID of the control is still Center, so our SimpleExpressions did not change.

To hide the Length and Angle, we'll run the UserControls script again. This time when we select the Length and Angle IDs, we'll choose Hide in the dialog. Press OK for each.

Finally, to change the options available in the Type, we have two options. We can hide the buttons and use a checkbox instead, or we can change the MultiButton from four entries to two. Let's try both.

To add the checkbox, run UserControls again, but this time instead of selecting an existing ID, we'll type Centered into the Name. This will set the name and the ID of our input to Centered. The Type is set to Number, and the Page is set to Controls. Now in the Type Attributes, set the Input Ctrl to be CheckboxControl. Press OK, and now we have our checkbox. To make the new control affect the Type, add a SimpleExpression to the Type:

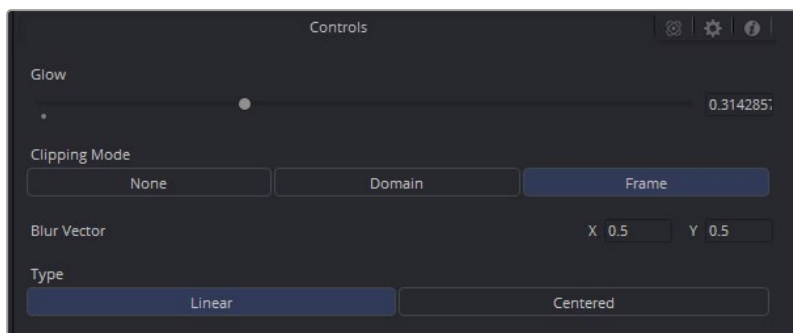
```
iif(Centered==1, 2, 0).
```

Once that's done, we can use the UserControls to hide the Type control.

To make a new MultiButton, run the UserControl script, and add a new control ID, TypeNew. You can set the Name to be Type, as the Names do not need to be unique, just the IDs. Set the Type to Number, the Page to Controls, and the Input Ctrl to MultiButtonControl. In the Input Ctrl attributes, we can enter the names of our buttons. Let's do Linear and Centered. Type them in and hit Add for each. Press OK, and we have our new buttons with the unneeded options removed. To make this new control affect the original Type, add a SimpleExpression to the Type:

```
iif(TypeNew==0, 0, 2).
```

Once that's done, we can use the UserControls to hide the original Type control.



Directional Blurs with UserControls applied

Chapter 69

Animating in Fusion's Keyframes Editor

This chapter covers how you can keyframe effects in the Inspector, and how you can edit clips, effects, and keyframes in the Keyframes Editor.

Contents

Keyframing in the Inspector	1203
Removing Animation in the Inspector	1204
Attaching a Parameter to an Existing Animation Curve	1204
Keyframes Editor Overview	1204
Keyframes Editor Tracks	1205
The Timeline Header	1205
The Playhead	1206
Spreadsheet	1206
Scaling and Panning the Timeline	1206
Working with Segments in the Timeline	1207
Moving Segments in the Timeline	1207
Trimming Segments	1207
Holding the First or Last Frame	1208
Working with Keyframes in the Timeline	1208
Drag and Drop Keyframe Editing	1208
Keyframe Editing Using the Time Editor	1209
The Keyframe Spreadsheet	1209
Duplicating Spline Keyframes	1209
Time Stretching Keyframes	1209
Showing Keyframe Values	1210


Timeline Filters	1210
Selected Filtering	1211
Sorting in the Timeline	1211
The Tree Item Order Menu	1212
The Sort Menu	1212
Markers	1213
Jumping to Markers	1213
Renaming Markers	1213
Show Marker List	1214
Deleting Markers	1214
Autosnap	1214
Autosnap Points	1214
Autosnap Markers	1215
The Spreadsheet Editor	1215
Selecting a Node to Edit	1215
Inserting Keyframes	1216
Selecting Multiple Nodes to Edit	1216
Customizing the Keyframes Editor	1216
Line Size	1216
Display Point Values	1216

Keyframing in the Inspector

Most parameters in most effects nodes can be keyframed in order to create animated effects such as animated transforms, rotoscoping with splines, dynamically altering warping behaviors, and more; the list is endless.

For convenience, a set of keyframing controls is available within the Inspector next to each keyframable parameter. These controls are:

- A gray Keyframe button to the right each keyframable parameter. Clicking this gray button creates a keyframe at the current position of the playhead, and turns the button orange.
- Whenever the playhead is sitting right on top of a keyframe, this button turns orange. Clicking an orange Keyframe button deletes the keyframe at that frame and turns the button gray again.
- Small navigation arrows appear to the right and left if there are more keyframes in those directions. Clicking on navigation arrows to the right and left of keyframes jumps the playhead to those keyframes.



Orange Keyframe buttons in the Inspector show there's a keyframe at that frame

Once you've keyframed one or more parameters, the node containing the parameters you keyframed displays a Keyframe badge to show that node has been animated.



A keyframed node displays a Keyframe badge in the Node Editor

Once you've started keyframing node parameters, you can edit their timing in the Keyframes Editor and/or Spline Editor.

Removing Animation in the Inspector

To remove a keyframed spline from a parameter:

- 1 Right-click the keyframe control of the parameter you want to remove animation from.
- 2 Choose Remove [Name of parameter] from the contextual menu.

Attaching a Parameter to an Existing Animation Curve

Multiple parameters can be connected to the same animation curve. This can be an invaluable timesaver if you are identically animating different parameters in a node.

To connect a second parameter to the same animation curve:


- 1 Right-click on the second parameter you want to attach.
- 2 In the contextual menu, hover over the Connect To submenu.
- 3 In the Connect To submenu, choose the name of the animated parameter.

Keyframes Editor Overview

The Keyframes Editor is essentially a timeline view of your composition, within which each clip and effect node in your composition is represented by a track. These tracks have the same color coding as the nodes they represent and are labeled where appropriate. A Time Ruler at the top indicates the timing of your composition, while numerous controls let you control the contents of the Keyframes Editor.

The Keyframes Editor can be used for one of two things:

- To adjust the timing of elements in a project, whether they're clips or effects. You can trim, slide, and extend clips, adjust the timing of an animation spline, or trim the duration of an effects node. You can freely rearrange the order of nodes in the Timeline without affecting the layering order of your composition. All compositing operations are handled in the Node Editor, while the Keyframes Editor manages the timing of your composition.
- To create and/or edit keyframes that you've applied to effects in a track-based manner, you can retime keyframes, add and delete keyframes, and even edit keyframe values




The Keyframes Editor

To show the Keyframes Editor, do one of the following:

- Click the Keyframes Editor button in the UI toolbar to toggle visibility of the Keyframes Editor on and off.
- Press F7 on the keyboard.

Keyframes Editor Tracks

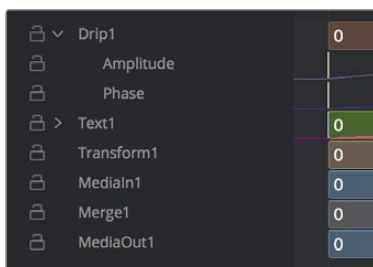
While each clip and effect node in your composition is represented by a track, keyframed parameters are exposed either as keyframes superimposed upon the track to which they're applied (as seen on the MOVEMENT track), or they can be opened up onto their own tracks for more precise editing, one keyframe track per keyframed parameter, by clicking a disclosure control to the left of that track's name in the Timeline header (as seen under the "Drip1" track).



The Timeline tracks

The Timeline Header

The Timeline header area on the left side of the Timeline is a hierarchical list of all tracks in a composition. Each track displays the name of its corresponding node, a lock button, and a disclosure control for revealing keyframe tracks for each keyframe animation, modifier, and mask that's attached to it.



The Timeline header area


Collapse/Open All

A quick way to open or close all available keyframe tracks at once is to use the Expand/Collapse Tool Controls commands in the Keyframe Timeline Option menu.

The Playhead

As elsewhere in Fusion, the playhead is a red vertical bar that runs through the Timeline view to indicate the position of the current frame or time. The Keyframes Editor playhead is locked to the viewer playhead, so the image you're viewing is in sync.

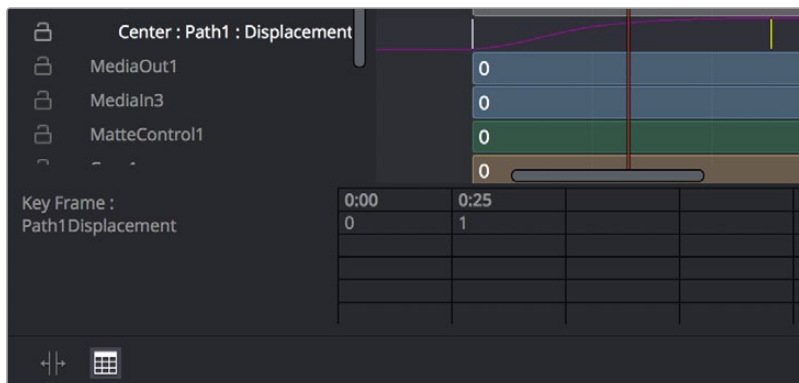
You must click on the playhead directly to drag it, even within the Timeline ruler (clicking and dragging anywhere else in the Timeline ruler scales the Timeline). Additionally, you can jump the playhead to a new location by holding down the Command-Option keys and clicking in the track area (not the Timeline ruler).



The playhead about to be dragged by the pointer

Spreadsheet

If you turn on the Spreadsheet and then click on the name of a layer in the keyframe track, the numeric time position and value (or values if it's a multi-dimensional parameter) of each keyframe appear as entries in the cells of the Spreadsheet. Each column represents one keyframe, while each row represents a single aspect of each keyframe.



Editing keyframes in the Spreadsheet

For example, if you're animating a blur, then the Key Frame row shows the frame each keyframe is positioned at, and the Blur1BlurSize row shows the blur size at each keyframe. If you change the Key Frame value of any keyframe, you'll move that keyframe to a new frame of the Timeline.

Scaling and Panning the Timeline

At the top, a series of zoom and framing controls let you adjust the work area containing the layers.

- A Horizontal zoom control lets you scale the size of the editor.
- A Zoom to Fit button fits the width of all tracks to the current width of the Keyframes Editor.
- A Zoom to Rect tool lets you draw a rectangle to define an area of the Keyframes Editor to zoom into.
- A Sort pop-up menu lets you sort or filter the tracks in various ways.
- An Option menu provides access to many other ways of filtering tracks and controlling visible options.

Working with Segments in the Timeline

Most of the work in the Timeline involves trimming and aligning clip segments.

To select a single segment in the Timeline, do one of the following:

- Click the node's name in the header.
- Click the node's segment in the Timeline.

To add another segment to the selection, do one of the following:

- Hold Command and click additional segments to select discontinuous selections.
- Select a segment, and then hold Shift and click another segment to make a contiguous selection of all segments in between.


To remove a segment from the selection, do the following:

- Hold Command and click a selected segment to deselect it.

TIP: Selecting a node's name from the Timeline header also selects the node's tile in the Node Editor, with its controls displayed in the Inspector.

Moving Segments in the Timeline

To move the position of a segment, drag on the node's segment in the Keyframes Editor. The cursor will resemble a bar with two arrows pointing in either direction. Moving a segment changes where that clip begins and ends in the composition.



The Move cursor

Trimming Segments

Trimming segments has different effects on Loaders, Medialn and Effect nodes:


- Trimming a Loader or Medialn node is similar to trimming clips in an editing application, in that you're changing the in and out points of the range of media that clip makes available to your composition.
- Trimming the segments of effect nodes instead modifies the range of that node's effect in the composition. Outside of the trimmed region, that effect node will behave as if it were disabled.

TIP: Shortening the duration of effects nodes can optimize processing. Imagine a Loader or Medialn node that represents a clip that's 100 frames long and is connected to a Defocus node that's animated from frames 80–100. There is little to no point in processing the defocus node between frames 0–79, so trimming the defocus segment to start at frame 80 in the Timeline will effectively prevent it from rendering and consuming either memory or processor time until needed.

To trim a segment in the Timeline, do the following:

- Drag on either end of the node’s segment in the Timeline.

The cursor changes to a vertical bar with a single arrow when the cursor is in the right location to trim.



The Trim cursor


Holding the First or Last Frame

If you want to hold a Loader’s first or last frame of a clip for a certain number of frames, also called a freeze frame, you can hold Command while you drag beyond the first or last of the segment in the Timeline.

Working with Keyframes in the Timeline

Keyframes can be drawn in one of two ways. When keyframe tracks are closed, they’re drawn over the node’s segment. Clicking on the disclosure icon to the left of the node’s name in the track header expands the display so each keyframed parameter has its own track in the Timeline, enabling precise editing.

Furthermore, each keyframe track, whether open or closed, exposes a miniature curve overlay that provides a visual representation of the rise and fall of keyframed values. This little overlay isn’t directly editable.



The Drip1 segment has its keyframe tracks exposed, while the Text1 segment has its keyframe tracks collapsed so they’re displayed within the segment.

Drag and Drop Keyframe Editing

Here are pointer-based keyframe editing methods that will get you started.

Methods of selecting keyframes:

- Click a single keyframe to select it.
- Drag a bounding box over a series of keyframes to select them all.
- Command-click to select discontinuous keyframes.
- Shift-click the first and last of a range of keyframes to select a contiguous range.

Methods of adjusting keyframes:


- You can drag keyframes left and right to reposition them in time.
- You can right-click one or more selected keyframes and use contextual menu commands to change keyframe interpolation, copy/paste keyframes, or even create new keyframes.

Keyframe Editing Using the Time Editor

A drop-down and editing field at the bottom right of the Keyframes Editor lets you numerically edit the timing, in frames, of any selected keyframe, making it easy to make precise adjustments.

To change the position of a keyframe using the toolbar, do one of the following:


- Select a keyframe, and then enter a new frame number in the Time Edit box.
- Choose T Offset from the Time Editor drop-down, select one or more keyframes, and enter a frame offset.
- Choose T Scale from the Time Editor drop-down, select one or more keyframes, and enter a frame offset.



The Time button can switch to Time Offset or Time Scale for moving keyframes.

The Keyframe Spreadsheet

If you turn on the Spreadsheet and then click on the name of a layer in the keyframe track, the numeric time position and value (or values if it's a multi-dimensional parameter) of each keyframe appear as entries in the cells of the Spreadsheet. Each column represents one keyframe, while each row represents a single aspect of each keyframe.



Editing keyframes in the Spreadsheet

For example, if you're animating a blur, then the Key Frame row shows the frame each keyframe is positioned at, and the Blur1BlurSize row shows the blur size at each keyframe. If you change the Key Frame value of any keyframe, you'll move that keyframe to a new frame of the Timeline.

Duplicating Spline Keyframes

Keyframes can be duplicated, either onto the same keyframe track or onto different tracks. This can save you time if you need to repeat a keyframe sequence at another time on the same segment, or even just create identically-timed keyframes on two different segments.

To duplicate keyframes, do the following:

- 1 Select one or more keyframes you want to duplicate.
- 2 Hold Command and drag one of the selected keyframes to a new position.

Time Stretching Keyframes


If you select a range of keyframes in a keyframe track, you can turn on the Time Stretch tool to show a box you can use to squeeze and stretch the entire range of keyframes relative to one another, to change the overall timing of a sequence of keyframes without losing the relative timing from one keyframe to the next. Alternatively, you can turn on Time Stretch and draw a bounding box around the keyframes you want to adjust to create a time-stretching boundary that way. Click the Time Stretch tool again to turn it off.



Time stretching keyframes

Showing Keyframe Values

When a node and its accompanying segment have animated parameters, keyframes appear as colored tick marks in keyframe tracks to indicate when animated changes occur. If the tracks and splines are open on a parameter, choosing Show Values from the Keyframes Editor Option menu shows editable fields beneath each keyframe. These fields show each keyframe's current value and allow you to edit them simply by entering a new number.



Show values from the Keyframes Editor Option menu

Timeline Filters

When a composition grows to include hundreds of nodes, locating specific node layers can quickly become difficult. Timeline filters can be created and applied to sift out nodes that are not necessary to the current operation. The Global Timeline preferences include a number of pre-made filters that you can enable, or you can create new ones as needed.

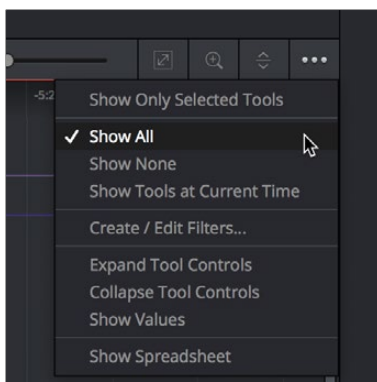
To use a Timeline filter:

Open the Keyframes Editor Option menu and choose an item from the top of the menu. Default Timeline filters include;

- Show All, which shows all node layers in the current composition.
- Show None, which hides all layers.
- Show Tools at Current Time, which only displays node layers under the playhead.
- If you've created custom filters, they appear here as well, in alphabetical order.

To go back to showing everything:


- Choose Show All from the Keyframes Editor Option menu. All layers will reappear.



Choosing a Timeline filter

To create a Timeline filter:

- 1 Choose Create/Edit Filters from the Keyframes Editor Option menu to open the Timeline panel of the Fusion Settings window. This is where you can create new Timeline filters.



The Global Timeline preferences for enabling filters

- 2 Click the New button, enter a name for your new filter setting, and click OK. The filter you created is now selected in the Filter pop-up menu at the top.
- 3 Use the “Settings for filters” list to turn on the checkboxes of nodes you want to be seen and turn off the checkboxes of nodes you want to filter out. Each category of node can be turned on and off, or you can open up a category’s disclosure control to turn individual nodes on and off. Clicking Invert All immediately turns off all node categories.
- 4 When you’re finished creating filters, click the Save button to hide the Fusion Settings window.

Filters that you’ve created in the Timeline panel of the Fusion Settings window appear in the Keyframes Editor Option menu.

To delete a filter:

- 1 Choose Create/Edit Filters from the Keyframes Editor Option menu to open the Timeline panel of the Fusion Settings window. This is where you can delete Timeline filters.
- 2 Choose the filter you want to delete from the Filter pop-up menu.
- 3 Click the Delete button, and when a dialog asks if you really want to do that, click OK.

Selected Filtering

Choosing “Show only selected tools” from the Keyframes Editor Option menu filters out all segments except for layers corresponding to selected nodes. This option can be turned on or off.

TIP: When “Show only selected tools” is enabled, you can continue to select nodes in the Node Editor to update what’s displayed in the Keyframes Editor.


Sorting in the Timeline

You can change the order in which the nodes are displayed from top to bottom in the Timeline.

- You can use the Tree Item Order Selection menu to sort the tracks by an assigned number.
- You can use the Sort pop-up menu.

The Tree Item Order Menu

Right-clicking over any track on the Keyframes Editor will display a contextual menu that contains the Tree Item Order Selection submenu. Choosing Start from the submenu allows you to start numbering each item in the track header by clicking on them. The first item you click will be #1, the second item #2, the third #3, and so on. Once you have selected all the items in the order you want them to be organized in the Keyframes Editor, right-click over a track, and from the Tree Item Order Selection submenu, choose End. The items will be ordered using the assigned numbers, with #1 appearing above #2, which appears above #3, and so on. The first items in the Keyframe track list will always be the nodes that are the root of the node tree. The numbered nodes will appear in order after the root nodes. For example, if the node tree starts with a background node, and then connects to a Fast Noise, Blur, and Color Corrector, the background node will always appear at the top of the Keyframes Editor track list because it is the root node.

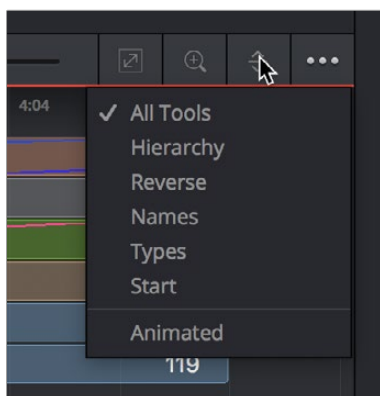


The Keyframes Editor Tree Item Order Selection menu

If you begin numbering nodes in the track header and change your mind or decide on a different order, you can choose Restart to begin numbering again or choose Cancel to keep the current order.

The Sort Menu

The Sort menu reorders how the layers of each node appear in the Keyframes Editor. Setting the menu back to All Tools will display them in a linear order, scanning the Node Editor from left to right and top to bottom. This is the default setting.




The Timeline Sort Order menu

- **All Tools:** Forces all tools currently in the Node Editor to be displayed in the Keyframes Editor.
- **Hierarchy:** Sorts with the most background layers at the top of the header, through to the most foreground layers at the bottom, following the connections of the nodes in the Node Editor.
- **Reverse:** The opposite of Hierarchy, working backward from the last node in the Node Editor toward the most background source node.
- **Names:** Sorts by the alphabetical order of the nodes, starting at the top with the beginning of the alphabet.

- **Start:** Orders layers based on their starting point in the composition. Nodes that start earlier in the Global project time are listed at the top of the header, while nodes that start later are at the bottom.
- **Animated:** Restricts the Timeline to showing animated layers only. This is an excellent mode to use when adjusting the timing of animations on several nodes at once.

Markers

Markers help identify important frames in a project that might affect how you keyframe animation. They may indicate the frame where a dragon breathes fire at a protagonist, the moment that someone passes through a portal, or any other important frame in a composition that you need to keep track of. Markers added to the Timeline in the Cut, Edit, Fairlight, or Color page will appear in the Keyframes Editor and Spline Editor of the Fusion page. They can also be added from the Keyframes Editor or the Spline Editor while working in Fusion Studio or the Fusion page. Markers in Fusion appear as a small handle with a line extending vertically through the graph view when selected.



A marker being moved in the Keyframed Editor

To create a marker, do the following:

- Right-click at a frame in the Timeline Ruler of the Keyframes Editor and choose Add Marker from the contextual menu.

The most important attribute of a marker is its position. For it to add value, a marker must be placed on the frame you intended it to be on. Hovering the cursor over a marker displays a tooltip with its current frame position. If it is on the wrong frame, you can drag it along the Time Ruler to reposition it.

Markers added to the Time Ruler are editable in the Fusion page, and the changes appear back in the other DaVinci Resolve pages. Time Ruler markers can be added, moved, deleted, renamed, and given descriptive notes from within Fusion's Keyframes or Spline Editor.

NOTE: Markers attached to clips in the Edit page Timeline are visible on MediaIn nodes in Fusion's Keyframes Editor but not editable. They are not visible in the Spline Editor.

Jumping to Markers

Double-clicking a marker jumps the playhead to that marker's position.

Renaming Markers

By default, a marker uses the frame number in its name, but you can give it a more descriptive name to go along with the frame number, making it easier to identify. To rename a marker in Fusion, right-click over the marker and choose Rename Guide from the contextual menu. Enter a name in the dialog and click OK.



The Marker contextual menu is accessed by right-clicking over a marker or Keyframe Editor Time Ruler.

Show Marker List

Markers can be used to jump to specific locations in a composition using the Marker List. If you right-click over a marker or within the Keyframe Editor Time Ruler to bring up the contextual menu, you can choose Show Marker List, or press Shift-G, to display the Marker List dialog. The Marker List is a floating dialog that will remain on top of the main window until closed.

The Marker List shows all the current markers in the composition, listed according to their position in time along with any custom name you've given them. If you double-click a marker's name from the list, the playhead jumps to the marker's location.



The Marker List dialog allows you to navigate through a composition using markers.

There is a pair of checkboxes beside the names of each marker. One is for the Spline Editor, and one is for the Keyframes Editor. By default, markers are shown in both the Spline Editor and Keyframes Editor, but you can deselect the appropriate checkbox to hide the markers in that view.

Deleting Markers

You can delete a marker by dragging it up beyond the Time Ruler and releasing the mouse. You can also use the marker's contextual menu to choose Delete Marker.

Autosnap

To help with precisely positioning keyframes and the start and end of segments as you drag in the Timeline, you can have them snap to a field, a frame, or to markers. The Autosnap option is accessed through the Keyframes Editor's contextual menu. There are two submenu options for autosnapping. One option controls the snapping behavior when you drag keyframes, control points, or the starting and ending edges of segments. The other option controls the snapping behavior of markers.

Autosnap Points

When you drag keyframes or the edges of segments, often you want them to fall on a specific frame. Autosnap restricts the placement of keyframes and segment edges to frame boundaries by default, but you have other options found in the contextual menu. To configure autosnapping on keyframes and segment edges, right-click anywhere within the Keyframes Editor and choose Options > Autosnap

Points from the contextual menu. This will display the Autosnap Points submenu with options for the snapping behavior. The options are:

- **None:** None allows free positioning of keyframes and segment edges with subframe accuracy.
- **Frame:** Frame forces keyframes and segment edges to snap to the nearest frame.
- **Field:** Field forces keyframes and segment edges to snap to the nearest field, which is 0.5 of a frame.
- **Guides:** When enabled, the keyframes and segment edges snap to markers.

Autosnap Markers

When you click to create a new marker, the default behavior is that it will snap to the closest frame. If you reposition the marker, it also snaps to the nearest frame as you drag. This behavior can be changed in the Keyframes Editor’s contextual menu by choosing from the Options > Autosnap Markers submenu. The options are:

- **None:** Markers can be placed anywhere with subframe accuracy.
- **Frame:** Frame forces all markers to snap to the nearest frame.
- **Field:** Field forces all markers to snap to the nearest field.

The Spreadsheet Editor

The Spreadsheet Editor is a separate panel that can be displayed beneath the Keyframes Editor. It is used to compactly show the numeric values of the keyframes for selected parameters in the Keyframes Editor’s header, via a table with rows and columns, showing time and value.



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 Keyframes Editor’s interface.

Selecting a Node to Edit

To display a node’s timing in the Spreadsheet, select the node’s name in the Keyframes Editor header. The Start and End points of the selected node will appear in the keyframe’s line of the Spreadsheet.

To edit an animation parameter in the Spreadsheet Editor, select the parameter in the Keyframes Editor header. The keyframe row includes a box for each frame number that contains a keyframe. The value of the keyframe is displayed in the cell below the frame number. Clicking on a cell allows you to change the frame number the keyframe is on or the parameter’s value for that keyframe.



Clicking on the parameter’s keyframe value allows you to change it

TIP: Entering a frame number using a decimal point (e.g., 10.25 or 15.75) allows you to set keyframes on a subframe level to create more natural animations.

Inserting Keyframes

You can also add new keyframes to an animation by clicking in an empty keyframe cell and entering the desired time for the new keyframe. Using the cell under the new keyframe, you can enter a value for the parameter.

Selecting Multiple Nodes to Edit

Multiple splines and nodes can be edited together in the Spreadsheet. By default, selecting a new parameter in the Timeline header will replace the parameter and keyframes currently listed in the Spreadsheet Editor. Holding down Command, you can click on additional parameters on different nodes to add them to the Spreadsheet.

TIP: You can use the Tab and Shift-Tab key shortcuts to move the selection right or left in the Spreadsheet Editor.

Customizing the Keyframes Editor

There are a few ways you can change the appearance of the Keyframes Editor to better fit your needs. All these options are found by right-clicking anywhere within the Keyframes Editor and choosing an option from the contextual menu that appears.

Line Size


The Line Size option controls the height of each Timeline segment individually. It is often useful to increase the height of a Timeline bar, especially when editing or manipulating complex splines.

Methods of increasing or decreasing the height of segments:

- **To change the height of just one segment:** Right-click anywhere within the Keyframes Editor and choose a size from the Line Size submenu. The options are Minimum, Small, Medium, Large, and Huge.
- **To change the height of all segments:** Right-click anywhere within the Keyframes Editor and choose a size from the All Line Size submenu. The options are Minimum, Small, Medium, Large, and Huge.


Display Point Values

A more traditional view of keyframes is to view them as control points instead of vertical bars, making them easier to select for some people. From the Timeline contextual menu, you can right-click anywhere within the Keyframes Editor and choose Options > Display Point Values to change how keyframes look.



The Options submenu for changing Display Point Values


Here are the two options, compared.



Keyframes displayed as bars (left), and keyframes displayed as Point Values (right).

Displaying Audio Waveforms

You can display a Medialn node's audio waveform in the Keyframes Editor and use it as a guide as you add and move keyframes.



Waveforms are displayed in the Keyframes Editor for all Medialn nodes

To display the audio waveform in the Keyframes Editor:

- 1 Open the Keyframes Editor.
- 2 Click the disclosure arrow next to the Medialn node to view the audio waveform for that clip.

To change the size of an audio waveform display:

- 1 Open the Keyframes Editor.
- 2 In the Keyframes Editor, select the audio track of the waveform you want to modify.
- 3 Right-click over the audio waveform and choose Line Size > Minimum/Small/Medium/Large/Huge.

When using Fusion Studio, you can view the audio waveforms in the Keyframes Editor by displaying the Saver node.

To view the Audio Waveform in Fusion Studio, do the following:

- 1 Open the Keyframes Editor.
- 2 Expand the Saver track to view the audio waveform.

When you want to find the precise location of an audio beat, transient, or cue, you can slowly drag over the audio waveform to hear the audio. If you need to see more resolution in the waveform display, you can increase the size.

To change the size of an audio waveform display:

- 1 Open the Keyframes Editor.
- 2 Right-click over the audio waveform and choose Line Size > Minimum/Small/Medium/Large/Huge.

TIP: Right-clicking a track in the Keyframes Editor and choosing All Line Size > Minimum/Small/Medium/Large/Huge changes all the tracks and audio waveforms in the Keyframes Editor.

Chapter 70

Animating in Fusion's Spline Editor

This chapter covers how you can keyframe effects and control animations in Fusion's Spline Editor.

Contents

Spline Editor Overview	1220	Copying and Pasting Keyframes	1232
Spline Editor Interface	1220	Time and Value Editors	1233
The Graph, Header, and Toolbar	1220	Modifying Spline Handles	1235
Renaming Splines	1222	Reducing Points	1235
Changing Spline Colors	1222	Filtering the Spline Editor	1236
Navigating Around the Spline Editor	1222	Working with Filters	1237
Markers	1224	Changing a Spline's Status	1238
Autosnap	1226	Selection States	1238
Creating Animation Splines	1226	Reshaping Splines Using the Toolbar	1239
Animating with Different Spline Types	1227	Interpolation	1239
Working with Keyframes and Splines	1229	Reversing Splines	1241
Adding Keyframes	1229	Looping Splines	1241
Locked and Unlocked Controls Points	1230	Time Stretching	1244
Selecting, Moving, and Deleting		Shape Box	1245
Keyframes	1231	Ease In/Out	1245
Showing Key Markers	1231	Importing and Exporting Splines	1246

Spline Editor Overview

The Spline Editor is the main area where animation is manipulated and refined. You primarily use the Spline Editor to show the changing values of parameters over time in the form of splines. Whereas Keyframes explicitly set the value of a parameter on a given frame, Splines are lines or curves that interpolate the values between keyframes. Once you set a keyframe on a parameter, a spline is created and can be displayed in the Spline Editor so you can make further refinements to the animation. However, the Spline Editor is more advanced than a standard curve editor since it can also display functions, which may not be splines, like changing characters within a text string or mathematical expressions that drive your animations.



The Spline Editor with three animated parameters

What Are Splines?

All animation uses splines that describe the value of a parameter at any given point in time. The Spline Editor graph shows the time of the comp along the horizontal axis and the value of the parameter along the vertical axis.

The advantage of using splines to represent animation instead of keyframes as in the Keyframes Editor is that splines allow you to manipulate the interpolation between keyframes. For example, if a keyframe is set at a value of 1.0 on a parameter for frame 1, followed by a keyframe value of 10.0 for frame 10, the values between keyframes are smoothly interpolated or calculated based on the shape of the spline. Using the functions and controls in the Spline Editor, you have a fantastic amount of control over that interpolation.

Spline Editor Interface

The Spline Editor is not visible by default, but you can show it at any time by clicking the Spline button in the user interface toolbar. You can also display the Spline Editor by right-clicking on a node in the Node Editor or a segment in the Keyframes Editor and choosing Edit Splines from the drop-down menu.

The Spline Editor can be open alongside the Node Editor or Keyframes Editor, or displayed separately in order to take up the entire work area.


The Graph, Header, and Toolbar

The Spline Editor has three main working areas: the graph, header, and toolbar. On the left side of the Spline Editor is the header, which shows a list of animated parameters. The majority of the panel is

taken up by the splines displayed in the graph area, and a toolbar runs along the bottom, providing a variety of ways to manipulate the splines.

Graph


The graph is the largest area of the interface. It is here that you see and edit the animation splines. There are two axes in the graph. The horizontal axis represents time, and the vertical axis represents the spline's value. A thin bar, called the playhead, runs vertically through the graph to represent the current time as it does in the Timeline Editor. You can drag the playhead to change the current time, updating the frame displayed in the viewers.



Removing a spline from a parameter removes the animation

Spline Editor Header

The header provides a mechanism for determining what splines are visible in the graph. It shows the name of each spline in the project beneath the tool that contains that parameter. The checkbox beside each name shows whether that spline is currently displayed in the graph and whether the spline can be edited.



The Spline header

Spline Editor Toolbar

The toolbar across the bottom of the Spline Editor represents the most common operations applied to an animation spline. The various operations represented in the toolbar are all accessible from the graph's context menu as well, but the following buttons provide a faster shortcut.



Spline toolbar buttons

Playhead

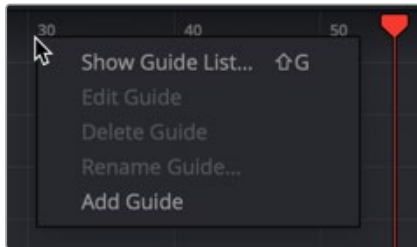
The playhead is the thin red vertical bar that runs vertically through the Spline Editor graph and represents the current time of the comp. You can drag the playhead to change the current time.

Status Bar

The status bar in the lower-right corner of the Fusion window regularly displays information about the position of the pointer, along with the time and value axes.

Contextual Menus

There are two contextual menus accessible from the Spline Editor. The Spline contextual menu is displayed by right-clicking over the graph, while the Guide contextual menu is displayed by right-clicking on the Time Ruler above the graph.



Right-click on the horizontal axis Time Ruler for the Guide menu

Renaming Splines

The name of a spline in the header is based on the parameter it animates. You can change the name of a spline by right-clicking on it in the header and choosing Rename Spline from the contextual menu.

Changing Spline Colors

Each spline in the graph is assigned a different color, making individual splines easier to identify when multiple splines are visible at once. When the spline is active, a round color swatch is displayed next to the spline's name in the header.


To change the color of a spline:

- 1 Right-click on the name of the spline in the header and choose Change Color from the contextual menu.
- 2 Select the new color from the dialog box that appears and click OK.

Navigating Around the Spline Editor

It is often necessary to magnify and pan around the graph area to ensure that the splines you want to work on are visible. In general, scaling and panning the Spline Editor works the same as in all navigable parts of the Fusion interface. However, there are several unique functions to the Spline Editor for controlling your view based on the height, width, and selection of multiple animation splines.

The most obvious navigation methods to use are the scale sliders and buttons located in the upper left of the Spline Editor panel.




The Zoom Height and Zoom Width sliders, Fit button, and Zoom to Rectangle button can be used to navigate around the graph

To scale and pan using the sliders and buttons:

- Zoom Height and Zoom Width sliders let you change the height and width of the graph area.
- The Fit button attempts to rescale the view so that all currently active splines fit within the graph.
- The Zoom to Rectangle button (Command-R) allows you to draw a bounding box around the area of the graph you want centered and scaled.

To scale using the axis labels:

- Place the mouse pointer over the rulers for the horizontal or vertical axis and drag to resize the graph on that axis only. The view is scaled centered on the original position of the pointer on the ruler.



Dragging in the Time Ruler scales the graph horizontally

To scale and pan with the mouse and/or keyboard:

- With the Spline Editor active, press the + and - keys on your keyboard to zoom in and out of the graph.
- You can also zoom to a specific control point by holding down the Command key and scrolling the middle mouse wheel. The mouse pointer location determines the area that gets magnified.
- Position the mouse pointer over the graph, and then hold down the middle mouse button. With the middle mouse button pressed down, click once on the left mouse button to zoom in and once on the right button to zoom out.

To pan the graph area:

- Drag left or right using the middle mouse button or use the scroll bar along the bottom and right side of the graph.

Using the Graph contextual menu for navigation:

There are several ways to navigate the graph area using the Spline Editor contextual menu as well.

- Choose Scale > Scale to Fit (Command-F) to fit all active splines into the graph area.
- Choose Scale > Scale to Rectangle (Command-R) to draw a bounding box around the area of the graph you want centered and scaled. This has the same effect as clicking the Zoom to Rectangle button.
- Choose Scale > Default to reset the scaling of the graph area to default values.
- Choose Scale > Zoom In/Zoom Out to scale the graph area. This performs the same functions as pressing the + and - keys on the keyboard.
- Choose Scale > Auto Fit to scale the graph to fit all splines dynamically as you make splines visible and hidden. If the scaling is changed with Auto-Fit enabled, the graph area will scroll as you play the comp to view all the keyframes.

- Choose Scale > Auto Scroll to scroll the graph area if the splines fall outside the graph horizontally as you play.
- Choose Scale > Manual to disable all automatic attempts at showing splines in the graph.
- Choose Options > Fit Times to automatically scale along the X-axis to fit the selected spline. All visible splines are taken into account, not just the newly selected spline. With this option off, activating a new spline will not change the horizontal scale.
- Choose Options > Fit Values to automatically scale along the Y-axis to fit the selected spline. All visible splines are taken into account, not just the newly selected spline. With this option off, activating a new spline will not change the vertical scale.

Markers

Markers help identify important frames in a project. They may indicate a frame where a ray gun shoots a beam in the scene, the moment that someone passes through a portal in the image, or any other important event in the composite.

Markers added to the Timeline in the Cut, Edit, Fairlight, or Color page will appear in the Keyframes Editor and Spline Editor of the Fusion page. They can also be added from the Keyframes Editor or the Spline Editor while working in Fusion Studio or the Fusion page. Markers appear along the top of the horizontal axis Spline Editor's Time Ruler. They are displayed as small blue shapes, and when selected, a line extends from each guide down vertically through the graph.


NOTE: Markers attached to clips in the Cut, Edit, Color, or Fairlight pages Timeline are not visible in Fusion's Spline Editor.



Unselected markers appear as blue shapes along the top, while selected markers display a vertical line running through the graph

Working with Markers

Markers call attention to a particular frame within a comp. They can be named, displayed in a list, and edited. After you add markers, you can easily jump the playhead between them, change their position, or delete them altogether.



Markers can be added by right-clicking in the horizontal time axis

To create a marker:

- Right-click in the horizontal axis Time Ruler and choose Add Marker.

To delete a marker, do one of the following:

- Drag the marker up outside the Spline Editor panel.
- Right-click on the marker and then choose Delete Marker from the menu.
- Select the marker and then press Delete or Backspace on the keyboard.
- From the Marker List, select a guide in the list and click the Del button.

To move a marker to a new frame, do one of the following:

- Drag the marker handle along the time axis.
- Right-click in the marker area and choose Options > Enable Guide Grabs, and then drag the marker's vertical line to move the guide.

To move the Playhead to a marker:

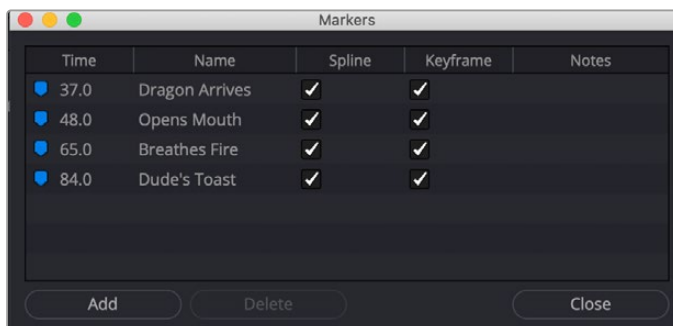
- Right-click on a marker and choose Set Current Time To [Frame number].

Using the Marker List

The Marker List is a list of all markers in the current comp. It can display the markers from either the Keyframes Editor, the Spline Editor, or both panels simultaneously. Clicking on the frame number or name of a guide causes the current time to change to that marker's frame. Since the Marker List is a floating window, it can remain open, allowing you to quickly jump to different markers while you work in the Spline Editor.

To show the Marker List:

- Right-click in the horizontal axis and choose Show Marker List, or press Shift-G.



The Marker List above shows markers in the current comp.

If markers currently exist in the comp, they are automatically displayed in the Marker List, regardless of whether they were added in the Keyframes Editor or the Spline Editor or any other page in DaVinci Resolve. You can also add markers directly from the Marker List, which can be helpful if you have multiple markers you need to add, and you know the rough timing.

To add a guide from the Marker List:

- 1 Click the Add button in the Marker List window.
- 2 Enter a frame number in the Time field.
- 3 Press Tab or click the Close button to close the Marker List.

To name a marker, do one of the following:

- In the Marker List, double-click in the Name column, to the right of the frame number, and enter a name for the marker.
- Right-click over a marker in the horizontal axis and choose Rename Marker. In the dialog that opens, enter a name for the marker.

Displaying Marker with the Timeline

The Marker List window includes checkboxes next to each marker that determines whether a marker displays in the Spline Editor, the Keyframes Editor, both, or neither. By default, when you create markers, they are active in both panels. To hide a marker from appearing in either panel, deselect the appropriate checkbox.

Autosnap

To assist in precisely positioning keyframe control points along the horizontal (time) axis, you can enable the Spline Editor's Autosnap function. Right-clicking over a spline and choosing Options > Autosnap provides a submenu with four options.

- **None:** Allows free, sub-frame positioning of the keyframes.
- **Frame:** Keyframes snap to the nearest frame.
- **Fields:** Keyframes snap to the nearest field.
- **Guides:** Keyframes snap to the nearest marker.

Autosnapping and Markers


By default, a newly-created marker snaps to the closest frame. Moving markers with the mouse also snaps them to the current frame. You can change this behavior by selecting Options > Autosnap Markers > None or by selecting Options > Autosnap Markers > Field from the contextual menu.

Creating Animation Splines

Animation splines are created automatically when you keyframe a parameter in the Inspector or the Keyframes Editor. However, you can create an animation spline without first having to add a keyframe.

To create a spline:

- Right-click on the parameter to be animated in the Inspector, and choose Animate from the contextual menu.

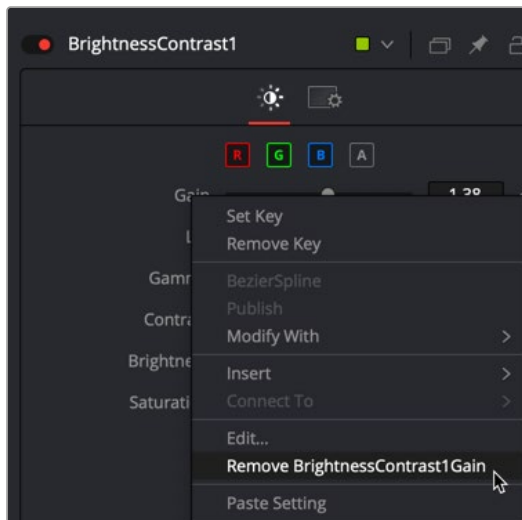


Choosing Animation from the menu displays a spline in the Spline Editor

Selecting Animate from the contextual menu connects the parameter to the default spline type. This is usually a Bézier Spline unless you change the default spline in the Defaults panel of the Fusion Preferences.

Deleting Animation Splines

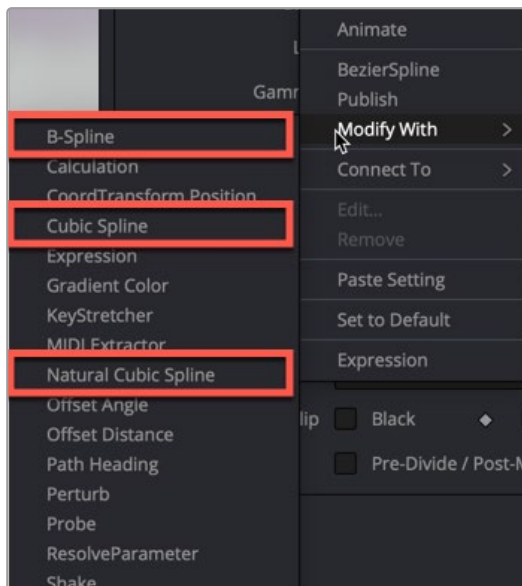
To remove an animation spline from a parameter, right-click on the control and select Remove [tool parameter's name] from the contextual menu. Removing a spline from a parameter only deletes the spline if no other tool in the composition is connected to the same spline at that time.



Removing a spline from a parameter removes the animation


Animating with Different Spline Types

A Bézier spline is the default spline unless changed in the Preferences. However, if you want to use a spline type other than Bézier for the animation curve, you can choose the spline type from the Modify With contextual submenu before creating any keyframes.



Three spline types in the Modify With menu

- **Bézier Spline:** Bézier splines are the default curve type. Three points for each keyframe on the spline determine the smoothness of the curve. The first point is the actual keyframe, representing the value at a given time. The other two points represent handles that determine how smoothly the curve for the segments leading in and out of the keyframe are drawn. Bézier is the most used spline type because Bézier splines allow you to create combinations of curves and straight lines.




Bézier spline

- **Modify with > B-Spline:** B-splines use a single point to determine the smoothness of the curve. Instead of using handles, a single control point determines the value as well as the smoothness of the curve. Holding down the W key while dragging left or right on the control point adjusts the tension of the curve.




B-spline

- **Modify with > Cubic Spline:** Cubic splines are similar to Bézier splines, in that the spline passes through the control point. However, Cubic splines do not display handles and always make the smoothest possible curve. In this way, they are similar to B-splines. This spline type is almost never used.



Cubic spline

- **Modify with > Natural Cubic Spline:** Natural Cubic splines are similar to Cubic splines, except that they change in a more localized area. Changing one control point does not affect other tangents beyond the next or previous control points.



Natural Cubic spline

Working with Keyframes and Splines

Once you animate a parameter and display the Spline Editor, you can manipulate the spline's keyframes (and thus the animation) in a variety of ways. By selecting Keyframe control points, you can move, copy, and change the interpolation of your animation.

Adding Keyframes

Once you create one keyframe, additional keyframes are automatically added to a spline whenever you move the playhead and change the value of that spline's parameter. For example, if you change the strength of an animated glow at frame 15, a keyframe with the new value occurs on frame 15.

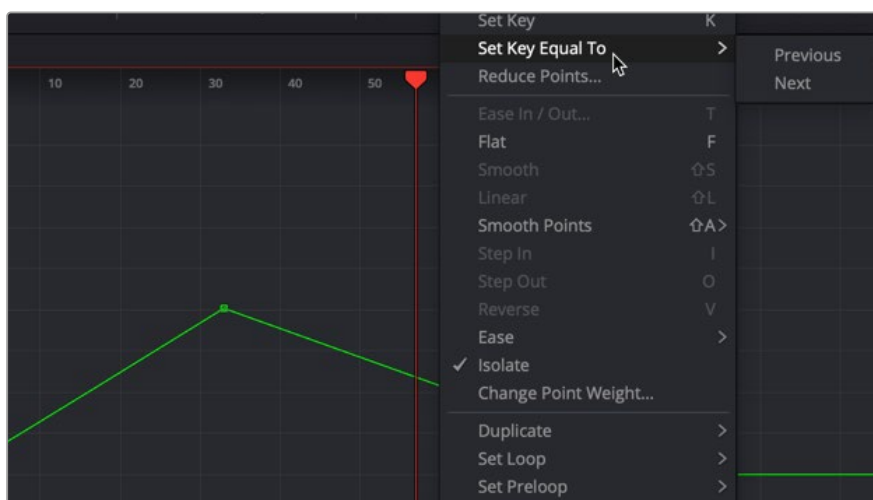
In the Spline Editor, control points can also be added directly to a spline by clicking on the spline where you want to add the new keyframe.

Adding Keyframes at the Playhead

If you want to add a new keyframe at the current playhead location, pressing Command-K on the keyboard or right-clicking in the graph and choosing Set Key adds a keyframe under the playhead.

Adding Equal Keyframes


If you want to hold a value over several frames, right-clicking in the graph area and choosing Set Key Equal To displays a submenu to add a new keyframe with a value equal to the next or the previous keyframe.



To hold a value between two keyframes choose from the Set Key Equal To submenu

Locked and Unlocked Controls Points

When animating the Center X/Y or Pivot X/Y parameters on any tool, you create a displacement spline in the Spline Editor. The displacement spline represents the relative offset position of the animated object along its path. Since the displacement spline is relative, keyframes use a value between 0.0 and 1.0. A displacement value of 0.0 in the Spline Editor indicates that the object is at the very beginning of a path. A value of 1.0 indicates that the object is positioned at the end of the path.



The displacement spline represents the relative position along a motion path

Displacement paths are composed of locked and unlocked points. Whether a point is locked is determined by how you added it to the polyline. Locked points on the spline have an associated point in the viewer's motion path; unlocked points do not have a corresponding point in the viewer's motion path. Each has a distinct behavior, as described below.

TIP: You can convert displacement splines to X and Y coordinates by right-clicking over the motion path in the viewer and choosing Path#: Polyline > Convert to X/Y Path.

Locked Points

Locked points are the motion path equivalents of keyframes. They are created by moving the playhead position and changing the parameter value. These points indicate that the animated object must be in a specified position on a specified frame. Since these keyframes are only related to position along the path, they can only be moved horizontally along the spline's time axis.

The locked points appear as larger-sized lock icons in the Spline Editor. Each locked key has an associated point on the motion path in the viewer.

Deleting a locked point changes the overall timing of the motion.

Unlocked Points

Unlocked points are created by clicking directly on the spline in the Spline Editor. These points give additional control over the acceleration along the motion path without adjusting the path itself. Conversely, you can add unlocked points in the viewer to control the shape of the motion path without changing the timing.

You can change an unlocked point into a locked point, and vice versa, by selecting the point(s), right-clicking, and choosing Lock Point from the contextual menu.

For more information on motion paths and locked keyframes, see Chapters 70 and 72 in the DaVinci Resolve manual or Chapters 9 and 11 in the Fusion Studio manual.

Selecting, Moving, and Deleting Keyframes

The placement of keyframes greatly affects the style of the animation. Using the graph, you can select keyframes and move them up or down to change their value or move them left and right to change the timing. Keyframes can be copied and pasted between splines and parameters.

Methods of selecting keyframes:

- Click directly on a keyframe on the spline, or drag a bounding box around the keyframe.
- Drag a bounding box that encompasses multiple keyframes to select more than one.
- To add or remove a keyframe from the current selection, hold down the Command key while selecting the keyframes. This will remove currently selected keyframes and add currently unselected keyframes.
- Press Command-A or right-click in the graph area and choose Select Points > Select All from the contextual menu to select all keyframes from the active splines.

Moving Keyframes

You can freely move keyframes with the mouse, keyboard, or the edit point controls. Keyframes can even pass over existing points as you move them. For instance, if a keyframe exists on frame 5 and frame 10, the keyframe at frame 5 can be repositioned to frame 15.

To move keyframes with the mouse:

- Drag the selected keyframe to its new position in the graph. If more than one keyframe is selected, all selected keyframes will be moved simultaneously.
- Hold down the Option key before dragging the keyframe to constrain its motion to a single axis.

To move keyframes with the keyboard:

- The Up and Down Arrow keys will adjust the value of the keyframes by a small amount.
- The current scale of the graph determines the degree of vertical movement applied to the value with each key press. The closer the zoom in the spline, the finer the adjustment.
- Hold down the Shift key while pressing the Up or Down Arrow keys to increase the value adjustment in larger increments.

To move keyframes using edit fields:

- The Value and Time Editors are found on the far right of the toolbar. These number fields allow explicit values to be entered for selected keyframes. These controls are explained in more depth later in this chapter.

To delete one or more keyframes:

- Select one or more keyframes and press the Delete or Backspace key on the keyboard. This only removes the keyframes; it does not remove the spline, even if there are no keyframes on the spline. To remove the spline, right-click over the parameter in the Inspector and choose Remove [parameter name].

Showing Key Markers

You can adjust the position of the keyframes in time, without worrying about manipulating splines, by using the key markers. The horizontal time axis can show markers that indicate the position of each keyframe. The display of these markers is enabled by right-clicking in the graph and choosing Show > Key Markers from the contextual menu, or by clicking on the Show Key Markers button in the toolbar.



The key markers show keyframes in the horizontal axis using the same color as the splines

Copying and Pasting Keyframes

To precisely match animation, Keyframes can be copied to a new location on the same spline or onto completely different splines and different tools.

There are two options in the graph's contextual menu for copying keyframes. Choosing Copy Points (Command-C) copies all selected points. Choosing Copy Value copies a single point identified by the pointer from multiple selected points. This does not deselect your selection set, and you can pick out numbers as needed.

To copy and paste points to a new location on the same spline:

- 1 Select the desired keyframes on the spline.
- 2 Right-click over the spline and choose Copy Points from the contextual menu or press Command-C.
- 3 Click in an empty area of the graph to deselect all the copied points.
- 4 Move the playhead to the area of the spline where you want the points pasted and press Command-V.

Or, move the pointer over the spline where you want the points pasted, and when the spline highlights, right-click and choose Paste Points/Value.

Alternatively, you can copy and paste keyframes by dragging them with the mouse. After you select the points, hold down the Command key and drag the points along the spline to where you want them pasted.

To copy and paste keyframes from one spline to another:

- 1 Make one spline the active visible spline and select the desired keyframes on the spline.
- 2 Right-click over the spline and choose Copy Points from the contextual menu or press Command-C.
- 3 Set the spline to viewed or disabled using the status checkbox next to the spline's name in the header.
- 4 Make the destination spline the active visible spline and select the keyframe on the spline where the new keyframes should be pasted.
- 5 Right-click and choose Paste Points/Value or press Command-V.

You can copy a single point's value from a group of selected points. Since this process does not deselect the selected set, you can continue picking out values as needed without having to reselect points.

To copy and paste a keyframe value:

- 1 Make one spline the active visible spline and select all the keyframes on the spline.
- 2 Right-click over a single point and choose Copy Value from the contextual menu.
- 3 Set the spline to viewed or disabled using the status checkbox next to the spline's name in the header.
- 4 Make the destination spline the active visible spline and select the keyframe on the spline where the new keyframe should be pasted.
- 5 Right-click and choose Paste Points/Value or press Command-V.

Keyframes can also be pasted with an offset, allowing you to duplicate a spline shape but increase the values or shift the timing using an offset to X or Y.


To paste keyframes points and values with an X or Y offset:

- 1 Make one spline the active visible spline and select the desired keyframes on the spline.
- 2 Right-click over the spline and choose Copy Points from the contextual menu or press Command-C.
- 3 Set the spline to viewed or disabled using the status checkbox next to the spline's name in the header.
- 4 Make the destination spline the active visible spline and select the keyframe on the spline where the new keyframes should be pasted.
- 5 Right-click and choose Paste with Offset. In the Offset dialog, enter the Y value, which will be added to the values of the pasted keyframes.

TIP: You cannot copy and paste between different spline types. For instance, you cannot copy from a Bézier spline and paste into a B-spline.

Time and Value Editors


The Time and Value Editors in the lower-right corner of the Spline Editor are used to change the position and parameter value of a keyframe by entering a number into the number field for each button. Each field can switch between three modes that help modify the time and value of a keyframe in three precise but distinct ways. The default mode for each field takes the explicit frame number or parameter value at which you want the keyframe set. The other modes offset and scale the keyframe's position or value.



Use the number fields to enter in a value or a specific time to change the selected keyframe

Time Editor

The Time Editor is used to modify the current time of the selected keyframe. You can change the Time mode to enter a specific frame number, an offset from the current frame, or spread the keyframes based on the distance (scale) from the playhead. You can select one of the three modes from the Time mode drop-down menu.



Three time editing modes are selectable from the Time mode drop-down menu

Time

The number field shows the current frame number of the selected control point. Entering a new frame number into the field moves the selected control point to the specified frame. If no keyframes are selected or if multiple keyframes are selected, the field is empty, and you cannot enter a time.

Time Offset


Selecting T Offset from the drop-down menu changes the mode of the number field to Time Offset. In this mode, the number field offsets the selected keyframes positively or negatively in time. An offset of either positive or negative values can be entered. For example, entering an offset of 2 moves a selected keyframe from frame 10 to 12. If multiple keyframes were selected in the previous example, all the keyframes would move two frames forward from their current positions.

Time Scale

Selecting T Scale from the drop-down menu changes the mode of the number field to Time Scale. In this mode, the selected keyframes' positions are scaled based on the position of the playhead. For example, if a keyframe is on frame 10 and the playhead is on frame 5, entering a scale of 2 moves the keyframe 10 frames forward from the playhead's position, to frame 15. Keyframes on the left side of the playhead would be scaled using negative values.

Value Editor

The Value Editor is used to modify the selected keyframe's parameter value using one of three Value modes. You can change the Value mode to enter a specific value for a parameter, an offset from the value, or to spread the values. The mode is chosen from the Value mode drop-down menu.



Three value editing modes are selectable from the Value mode drop-down menu

Value

The number field shows the value of the currently selected keyframes. Entering a new number into the field changes the value of the selected keyframe. If more than one keyframe is selected, the displayed value is an average of the keyframes, but entering a new value will cause all keyframes to adopt that value.

Value Offset


Choosing Offset from the drop-down menu sets the Value Editor to the Offset mode. In this mode, the value for the selected keyframes are offset positively or negatively. An offset of either positive or negative values can be entered. For example, entering a value of -2 changes a value from 10 to 8. If multiple keyframes are selected, all the keyframes have their values modified by -2.

Value Scale

Choosing Offset from the drop-down menu sets the Value Editor to the Scale mode. Entering a new value causes the selected keyframes' values to be scaled or multiplied by the specified amount. For example, entering a value of 0.5 changes a keyframe's value from 10 to 5.

Modifying Spline Handles

All Bézier spline keyframes have a pair of control handles to shape the spline as it passes through the key point. These handles are only displayed when the keyframe is selected. Initially, these handles are set to linear, creating straight line changes between keyframes. However, any control point can be made smooth by right-clicking over it and choosing Smooth or pressing Shift-S.



Bézier splines can mix linear and smooth curves

Dragging on a keyframe's handles adjusts the slope of the segments passing through the spline. By default, the two control handles on a control point are locked together so that if one moves, the one on the other side moves with it. This maintains a constant tension through the keyframe. There are situations, however, when it is desirable to modify these control handles separately for a more pronounced curve or effect.

To temporarily break the control handles on a Bézier spline, moving one independently of the other:

- 1 Select the control point to be modified.
- 2 Hold down the Command key and drag one of the control handles. They will now move independently of each other, as long as the Command key is held down.

To treat all Bézier handles as independent in the Spline Editor:

- Right-click in the graph and choose Independent Handles from the Options contextual menu.

Enabling this option causes all the Bézier handles to be independent. This is the same as using the Command key when moving a handle, except it is applied to all control points until it is disabled.

Reducing Points

When there are too many control points too close together on a spline, you can choose Reduce Points to decrease their number, making it easier to modify the remaining points. The overall shape of the spline is maintained as closely as possible while eliminating redundant points from the path.

To reduce the number of points on a spline:

- 1 Select the range of keyframes you want to thin out.
- 2 Right-click in the graph area and choose Reduce Points from the contextual menu.
- 3 When the Reduce Points dialog appears, drag the slider to a lower value.

You can set the slider value as low as possible as long as the spline still closely resembles the shape of your original spline.

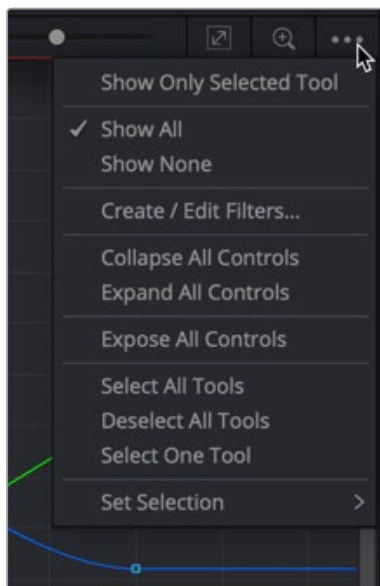
TIP: When the value is 100, no points will be removed from the spline. Use smaller values to eliminate more points.

Filtering the Spline Editor

The animation splines for multiple parameters can be displayed within the Spline Editor simultaneously, and Fusion offers several ways for you to choose which spline to view and which to edit.

A complex composition can easily contain dozens, if not hundreds, of animation curves. As a composition grows, locating a specific spline can become more difficult. There are two ways to filter the splines shown in the Spline Editor: display selected tools only or create a filter to show only certain tools.

The Spline Editor includes different ways to control which splines are displayed. The majority of these options are available in the Options menu, located in the upper-right corner of the Spline Editor panel.



The Options menu is used to control which splines are displayed in the Spline Editor

- **Show Only Selected Tools:** You can choose to limit the splines displayed in the Spline Editor by showing only the splines from selected tools. Choosing this option at the top of the Options menu displays only the splines for tools currently selected in the Node Editor.
- **Show All/None:** The default behavior of the Spline Editor displays all the splines for all the nodes with animated parameters. You can override this by enabling Show Only Selected Tools in the Options menu. You can also disable the Show All setting by choosing Show None, in which case the Spline Editor remains empty.
- **Expose All Controls:** The Expose All Controls option is a way of not filtering the parameters. Choosing this option displays all parameters in the Spline Editor header for all nodes in the Node Editor. It can be a fast way of activating one of the parameters and automatically adding an animation spline for it if one does not exist.

With a large number of nodes displayed, which themselves might have a large number of parameters, this might lead to cluttering and slowing down the interface. This option is most effective when used in conjunction with the Show Only Selected Tool option to limit the number of nodes and parameters displayed and yield optimum performance.

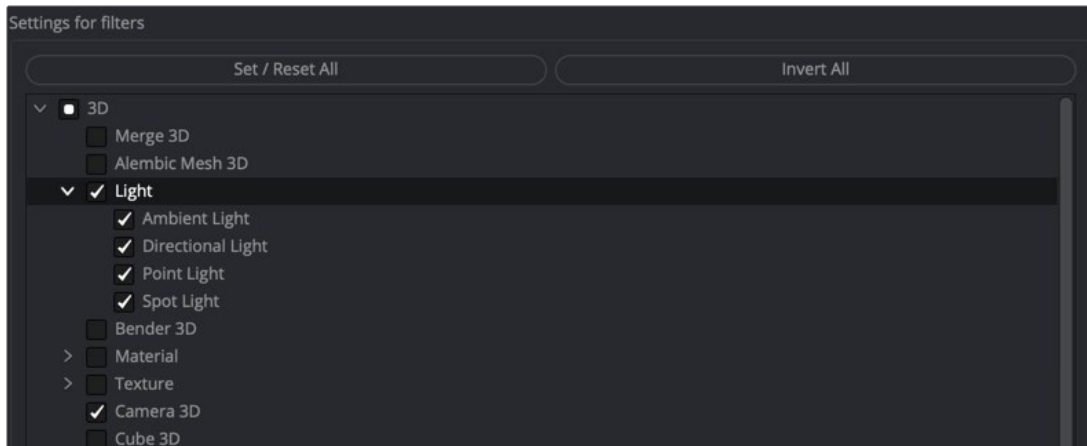
- **Follow Active:** The Follow Active option is located by right-clicking in the graph and choosing Options > Follow Active. This option provides a way to filter the splines in the graph while not filtering the header list of tools. Where the Show Only Selected Tool option hides other tools in the header, the Follow Active option leaves the header displaying all the tools but automatically enables only the splines of the Active tool.

Working with Filters

Filters allow you to select the specific types and classes of tools shown in the Spline Editor and Keyframes Editor. For example, you can make a filter that shows only particle nodes or one that only shows color correction and brightness/contrast tools.

To create a filter:

- 1 From the Options menu, choose Create/Edit Filters.
- 2 Click the New button to create a new filter and name the new filter in the dialog box.
- 3 Enable a checkbox next to the entire category or the individual tools in each category to determine the tools included in the filter.



Enable each tool you want to keep in the Spline Editor when the filter is selected

The Invert All and Set/Reset All buttons can apply global changes to all the checkboxes, toggling the selected states as described.

To switch the selection state of the categories when creating a filter list:


- 1 Click the Invert All button.
- 2 After configuring the custom filter, click the Save button to close the Settings dialog and save the filter.

To enable all checkboxes or disable all checkboxes:

- 1 Click the Set/Reset All button as many times as needed until all categories are either checked or unchecked.
- 2 After configuring the custom filter, click the Save button to close the Settings dialog and save the filter.

To apply a filter to the Spline Editor:

- Choose the desired filter by name from the Options menu. The filter applies to both the Spline Editor and the Timeline.



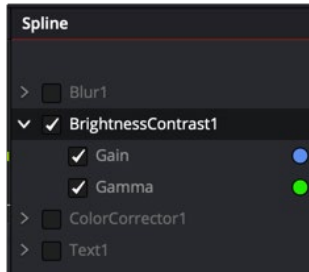
Each filter you create is listed in the Options menu

To disable a filter and show all tools in the Spline Editor again:

- Choose Show All from the Options menu.

Changing a Spline's Status

The Spline header is a hierarchical list of animated parameters and their parent nodes. Clicking the disclosure arrow next to a tool's name reveals all the names of the animated parameters on that tool. Clicking directly on the parameter name in the Spline header activates that spline for display and editing.




The Spline header with the Brightness Contrast tool and its animated Gain and Gamma parameters

Tool Status Checkbox

Next to the name of each spline is a checkbox that indicates the spline's status. When you select a parameter name, the checkbox becomes active, allowing you to see and edit the spline in the graph. There are three selection modes for each checkbox: active, viewed, and disabled. Clicking directly on the checkbox will toggle it between these three states. Changing the state of the parent node checkbox sets the state for all splines for that node.


- **Active:** When the checkbox is enabled with a check mark, the spline is displayed in the graph and can be edited.
- **Viewed:** When the checkbox is enabled with a solid gray box, the spline is visible in the graph but cannot be edited. It is read-only.
- **Disabled:** When the checkbox is clear, the spline is not visible in the graph and cannot be edited.



Checkboxes determine which splines are visible and editable

Selection States

There are three selection options, labeled Select All Tools, Deselect All Tools, and Select One Tool, that determine how the items in the Spline Editor header behave when a checkbox or label is selected to activate a spline. These states are located in the Options menu in the upper-right corner of the Spline Editor.



The Options menu selection states make it easier to select and deselect all parameters in the header

- **Select All Tools:** Choosing this option activates all splines for editing.
- **Deselect All Tools:** Choosing this option sets all spline checkboxes to disabled.
- **Select One Tool:** This option is a toggle. When Select One Tool is chosen from the menu, only one spline in the header is active and visible at a time. Clicking on any spline's checkbox will set it to active, and all other splines will be cleared. When disabled, multiple splines can be active in the header.

Selection Groups

It is possible to save the current selection state of the splines in the header, making selection groups that can easily be reapplied when needed. To create a selection group, right-click over any parameter in the header or in an empty area of the graph and choose Save Current Selection from the contextual menu. A dialog will appear to name the new selection.

To reapply the selection group, choose the selection group by name from the Set Selection menu in the same contextual menu. Other context menu options allow selection groups to be renamed or deleted.

Reshaping Splines Using the Toolbar

There are several ways to manipulate the shape of a spline, thereby altering the animation that spline generates. Other than manually adjusting Bézier handles, you can quickly squish, stretch, loop, and reverse a spline. You can also quickly change the interpolation between keyframes from the default linear motion to more natural smooth motion. All these options are provided at the bottom of the Spline Editor in the toolbar. The toolbar is divided into different groups for setting interpolation, reversing splines, looping splines, time stretching, and reshaping splines.

Interpolation

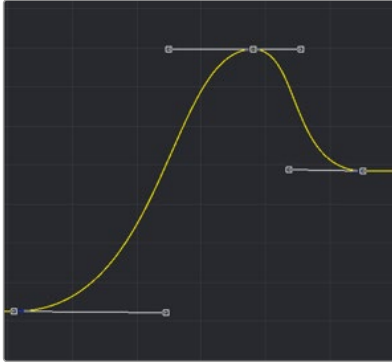
Keyframes are specific frames in an animation where control points are set to exact values on a given parameter. Interpolation is the method used to fill in the unknown values between two keyframes. Fusion automatically interpolates between two keyframes. However, you may want to modify the interpolation to achieve a specific style of animation. The Spline Editor includes several interpolation methods you can choose from using the toolbar.



Interpolation buttons in the toolbar: Smooth, Linear, Invert, Step In, and Step Out

Smooth


A smoothed segment provides a gentle keyframe transition in and out of the keyframe by slightly extending the direction handles on the curve. This slows down the animation as you pass through the keyframe. To smooth the selected keyframe(s), press Shift-S or click the toolbar's Smooth button.



Smooth interpolation between keyframes

Linear

A linear segment effectively takes the shortest route between two control points, which is a straight line. To make the selected keyframe(s) linear, press the L key or click the Linear button in the toolbar.




Linear interpolation between keyframes

TIP: Invert is used only for non-animated LUT splines, which are currently only available in the LUT Editor window.

Step In/Step Out


On occasion, it is not desirable to have any interpolation between two keyframes. Instead, the value of one keyframe may hold its value until another keyframe changes it. For these cases, use the Step In or Step Out mode.

Step In causes the value of the previous keyframe to hold, then jump straight to the value of the next keyframe.



Step In holds a value until the next keyframe is reached in the comp

Step Out causes the value of the selected keyframe to hold right up to the next keyframe.



Step Out switches immediately to the next keyframe value in a comp

Step In and Step Out modes can be set for selected keyframes by clicking on the toolbar buttons for each mode, or by right-clicking and choosing the appropriate option from the contextual menu. The keyboard shortcuts I and O can also be used to enable Step In and Step Out on selected keyframes.

Reversing Splines


Reverse inverts the horizontal direction of a segment of an animation spline. To apply reverse, choose a group of points in a spline and click the Reverse button, or right-click and choose Reverse from the contextual menu, or press the V key. The group of points is immediately mirrored horizontally in the graph. Points surrounding the reversed selection may also be affected.



The Reverse button in the toolbar

Looping Splines


It is often useful to repeat an animated section, either infinitely or for a specified number of times, such as is required to create a strobing light or a spinning wheel. Fusion offers a variety of ways to repeat a selected segment.



The various Loop buttons in the toolbar

Set Loop

To repeat or loop a selected spline segment, select the keyframes to be looped. Select Set Loop from the contextual menu or click on the Set Loop button in the toolbar. The selected section of the spline repeats forward in time until the end of the global range, or until another keyframe ends the repeating segment.



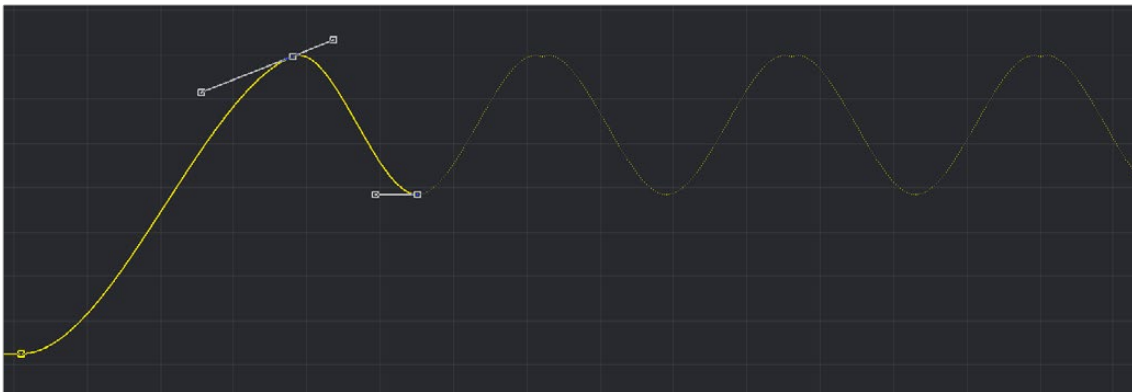
A looped section in the graph

Changing and Removing the Loop

You can change the looped segment by modifying any of the keyframes or control points originally used to create the loop. Simply select one of the originating key points, make any necessary modifications, and the looped segment updates. To remove the loop, select the keyframes you used to create the loop, and then click the Loop button in the toolbar.

Ping-Pong


The Ping-Pong Loop mode repeats the selected segment, reverses each successive loop, and then repeats. Ping-pong looping can be enabled on the selected segments from the context menu or the toolbar.



A ping-pong section in the graph

Relative Loop

The Relative Loop mode repeats the segment like the Loop, but each repetition adds upon the last point of the previous loop so that the values increase steadily over time.



A Relative Loop section in the graph

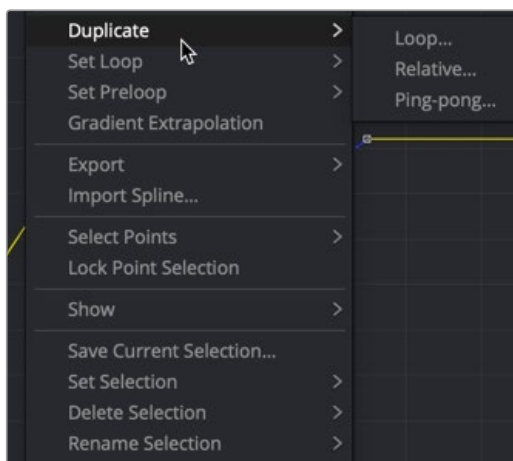
Looping Backward

You can choose Set Pre-Loop by right-clicking in the graph area and choosing it from the contextual menu. This option contains the same options for looping as the Loop option buttons in the toolbar, except that the selected segment is repeated backward in time rather than forward.

Repeating Splines X Number of Times

You can duplicate splines and repeat them a set number of times by right-clicking in the graph area and choosing Duplicate from the contextual menu. Duplicated splines are like looped splines, except that the selected segment repeats only a specified number of times, and each repetition is a copy rather than an instance. Adjustments to the original segment do not alter the shape of its repetitions.


The Duplicate modes are only accessed from the Duplicate contextual menu, which reveals a submenu with all the looping modes described above. Selecting any of these modes opens a dialog in which the number of repetitions can be entered.



The Duplicate contextual submenu

Gradient Extrapolation

You can choose Gradient Extrapolation by right-clicking in the graph area and choosing it from the contextual menu. This option continues the trajectory of the last two keyframes.



The Gradient Extrapolation applied to the spline


Time Stretching

Time Stretching allows for a selected group of keyframes to be proportionally stretched or squashed. This allows you to change the duration of the animation while keeping the relative distance between each keyframe. To enable spline stretching, select the group of keyframes that you want to time stretch, and then choose Modes > Time Stretching from the graph's contextual menu or click the Time Stretch button in the toolbar.



The Time Stretch button in the toolbar

When you have more than one keyframe selected on the spline, enabling Time Stretch surrounds the outer keyframes with two vertical white bars. Drag on the white vertical bars to stretch or shrink the timing of the spline segments within the bars. Drag these bars back and forth to stretch or squash the spline segment.



The Time Stretch bars in the graph

TIP: If no keyframes are selected when you enable Time Stretch, drag a rectangle to set the boundaries of the Time Stretch.

To disable the Time Stretching mode:

- Click on the Time Stretch button in the toolbar again or reselect Modes > Time Stretching from the contextual menu.

Shape Box

The Shape Box transform mode is similar to Time Stretching; however, it can adjust the vertical scaling of keyframe values as well as time.


To enable the Shape Box, do one of the following:

- Select a group of keyframes, and then choose Modes > Shape Box from the contextual menu.
- Select a group of keyframes, and then click the Shape Box button in the toolbar.
- Select a group of keyframes, and then press Shift-B to enable or disable the Shape Box mode.



The Shape Box button in the toolbar

A white rectangle outlines the selected points when the mode is enabled. To scale, skew, or stretch the spline, drag on any of the control points located around the box. To move all the keyframes, drag on the box edges.




The Shape Box in the graph

TIP: If no points are selected, or if you want to select a new group of keyframes, you can drag out a new rectangle at any time.

Ease In/Out

For a more precise way to adjust the length of Bézier direction handles attached to selected keyframes, you can use the Spline Ease dialog. To show the dialog, select a keyframe in the graph, and then choose Edit > Ease In/Out from the graph's contextual menu or press T on the keyboard.

The Ease In/Out controls appear above the graph area. You can drag over the number fields to adjust the length of the direction handles or enter a value in the fields.



The Ease In/Out controls above the graph area

Clicking the Lock In/Out button will collapse the two sliders into one, so any adjustments apply to both direction handles.

Importing and Exporting Splines

Spline shapes can be imported and exported from or to an ASCII text file. This makes it easier to save complex spline curves for later reuse, or to transfer tracking, path, and animation data from one application to another. Exported splines are assigned the file extension `.spl` for easy identification.

To export a spline:

- 1 Select the active spline in the Spline Editor.
- 2 Right-click on the spline in the graph area to display the contextual menu.
- 3 Choose from three format options in the submenu.
- 4 Enter a name and location in the file browser dialog, and then click Save.

Exporting a spline gives you three options. You can export the Samples, Key Points, or All Points. Samples adds a control point at every frame to create an accurate representation of the spline. Key Points replicates the control point positions and values on the spline using linear interpolation. All Points exports the spline as you see it in the Spline Editor, using the same position, value, and interpolation.

To import a spline:

- 1 Add an animation spline for the parameter.
- 2 In the Spline Editor, right-click on the animation spline and select Import Spline from the contextual menu.
- 3 In the File Browser dialog, select the spline curve `.spl` file, and then click Open.

Importing a new curve will replace any existing animation on the selected spline.

Chapter 71

Animating with Motion Paths


Layers and 3D objects can move along a designated spline shape to create motion path animations. This chapter discusses how you can create, edit, and use motion paths in Fusion.

Contents

Animating Using Motion Paths	1248
Types of Motion Paths	1249
Polyline Path	1249
Path Modifier	1252
Controlling Speed and Orientation along a Path	1253
XY Path	1255
Types of Control Points	1257
Locked Points	1257
Unlocked Points	1260
Locking and Unlocking Points	1261
Tips for Manipulating Motion Paths	1261
Compound Motion Paths Using Path Centers	1261
Copying and Pasting Motion Paths	1262
Removing Motion Paths	1262
Recording Motion Paths	1263
Importing and Exporting Polylines	1263
Native Format	1263

Animating Using Motion Paths

Motion paths are created from splines (polylines) including paint strokes and masks that have a path modifier applied to them. The path modifier defines the movement that transforms ordinary spline shapes into motion paths. You apply the path modifier either explicitly to the Center X/Y or Pivot X/Y parameters or implicitly by keyframing the Center or Pivot parameters. In either case, the spline path is displayed and then visually adjusted in the viewers. For all motion paths, the coordinate control in the viewer represents the position of an object or effect, such as a Transform or Rays node's center that moves along the path. Coordinate controls are represented onscreen with vertical and horizontal positioning arrows or an X.



A Center Offset onscreen control for an object's Transform node

The following nodes have parameters that can be animated using path modifiers to move an image around the composition. These include, but are not limited to:

- **Transform:** Center X/Y can be animated to move an image around.
- **DVE:** Center X/Y can be animated to move an image around.
- **Merge:** Center X/Y can be animated to move the Foreground connected image around.
- **Paint:** Stroke Controls > Center X/Y can be animated to move a stroke around.
- **Camera 3D node:** Translation X/Y/Z
- **Shape 3D node:** Translation X/Y/Z

The following nodes have parameters that can be animated using paths to alter the direction of a visual effect. These include, but are not limited to:

- **Directional Blur:** Center X/Y can be animated to change the direction of the blur.
- **Hot Spot:** Primary Center X/Y can be animated to move the hot spot around.
- **Rays:** Center X/Y can be animated to change the angle at which rays are emitted.
- **Polygon/BSpline/Ellipse/Rectangle/Triangle mask:** Center X/Y can be animated to move the mask.
- **Corner Positioner:** Top Left/Top Right/Bottom Left/Bottom Right X/Y can be animated to move each corner of the corner-pinned effect.
- **Vortex:** Center X/Y can be animated to move the center of the warping effect.

NOTE: It's not possible to add a motion path to a one-dimensional value, such as blur strength or merge angle. However, you can use the Spline Editor to edit these kinds of values in a visual way.

Types of Motion Paths

There are three types of motion paths: Polyline paths, XY paths, and 3D motion paths for 3D scenes.

- A Polyline path is generated by applying the path modifier. It uses two splines to define the path; one for the shape of the path displayed in the viewer, and a Displacement spline for the speed of the object along the path, displayed in the Spline Editor. The Polyline path is the default type of path modifier, and most documentation in this chapter assumes that this type is used.
- The XY path modifier employs a spline for the X position of the point and another for the Y position. The XY path modifier is explained in detail later in this chapter.
- 3D motion paths pertain only to positional controls within 3D scenes.

Polyline Path

Polyline paths are the easiest motion paths to work with. You can use the spline shape in the viewer to control the shape of the path, while a single Displacement curve in the Spline Editor is used to control the acceleration along the path. The most obvious way to create a Polyline motion path is by keyframing the Center X/Y parameter of a Transform node in the Inspector.

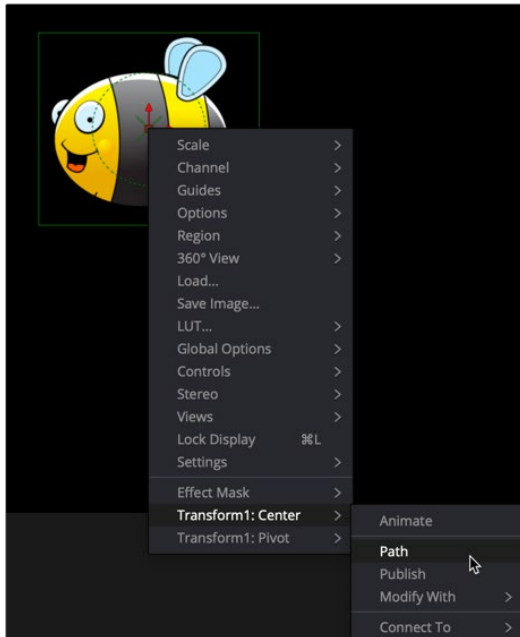
To create a Polyline motion path using the Center X/Y parameter in the Inspector:

- 1 Position the playhead on the frame where the motion will begin.
- 2 In the Inspector, click the gray Keyframe button to the right of the Center X and Y parameters.
This action applies the path modifier in the Modifiers tab in the Inspector.
- 3 Adjust the Center X and Y for the first keyframe position.
- 4 Position the playhead on the frame where the motion should change or stop.
- 5 In the Inspector, change the Center X and Y parameters to set a keyframe for the new location automatically.
- 6 In the viewer, modify and refine the motion path by selecting a control point and using any of the spline controls in the viewer toolbar.
- 7 Open the Spline Editor and adjust the Displacement spline to control the speed of the object along the path.

Keyframing Center X/Y is not the only way to apply the path modifier. An alternative method is to apply the path modifier to the Center X/Y parameter either in the Inspector or using the coordinate control in the viewer.

To apply a Polyline path to the Center X/Y parameter:


- 1 Position the playhead on the frame where the motion will begin.
- 2 Position the center coordinate control for the object, or effect, at its starting position.
- 3 Right-click the onscreen center coordinate control in the viewer (or over the Center X and Y parameter in the Inspector), and choose Path from the contextual menu for that control.



Applying a path modifier to the center of a Transform node

The object now has a path modifier applied, so without setting a keyframe you can drag the object to begin creating a motion path in the viewer.

- 4 Move the playhead to a new frame.
- 5 Drag the onscreen coordinate control or adjust the Offset or Center values in the Inspector. A keyframe is automatically created on the motion path, and a polyline is drawn from the original keyframe to the new one.



Two keyframes spaced several frames apart display a motion path showing the direction of animation.

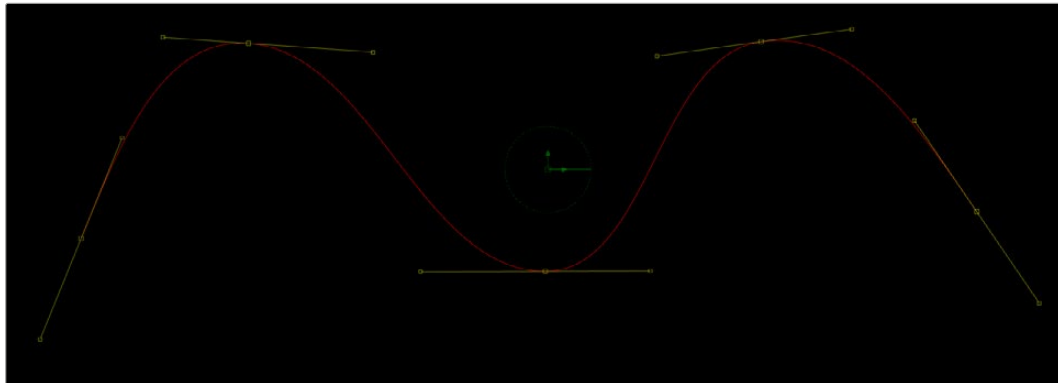
- 6 The position of the center control is interpolated between the two keyframes. You can continue adding points by moving the playhead and adjusting the object's position until the entire motion path you need is created. For motion paths, there's no need to close the spline path; you can leave it open.
- 7 Upon completion, set the polyline to Insert and Modify mode by selecting a point on the path and pressing Command-I or clicking the Insert and Modify button on the toolbar. Don't worry too much about the overall shape of the motion path at this point. The shape can be refined further by adding additional points to the polyline and modifying the existing points.

A final alternative method for creating a motion path is to draw a spline shape first and then connect a path modifier to the spline. Using any of Fusion's spline tools, you can draw the shape of the path and then connect the path modifier to the published spline. Once the path modifier and the published spline are connected, you can keyframe the Displace parameter to move an image along the path. This method is useful when you want to use a paint stroke or mask shape as a motion path.

To use a mask shape or paint stroke as a path:

- 1 Draw a polygon mask or paint stroke in the shape of the path you want to use.

You do not have to close the mask shape.




An open Polyline mask shape

- 2 When done drawing the shape, click the Insert and Modify button in the viewer toolbar to leave the mask shape as an open spline.
- 3 At this point you can select any of the control points along the spline and press Shift-S to make them smooth or Shift-L to make them linear.

All mask polylines have animation enabled by default, but that is usually not desirable for a motion path. You will need to remove this keyframe animation if you are using a mask shape.

- 4 At the bottom of the Inspector, right-click on the "Right-click here for shape animation" label and choose Remove Polygon1Polyline.




Right-clicking at the bottom of the Inspector to remove auto-keyframing from a mask shape

- 5 Right-click at the bottom of the Inspector again and select Publish to give other nodes access to this spline shape. (For a paint stroke, you will need to make the Stroke editable first by clicking the Make Editable button in the Stroke Controls.)

This enables the Modifiers tab with the Published Polyline modifier. This published spline can be used to define the shape of splines in other nodes.

- 6 Connect a Transform node to the image you want to have follow the path.
- 7 Right-click over the Center X/Y parameter in the Inspector and choose Path.
This adds a path modifier into the Modifiers tab.
- 8 In the Inspector, click the Modifiers tab and double-click the Path1 heading to open its parameters.
The Displacement parameter already has a keyframe on it automatically. You'll want to remove that so you can set your own.
- 9 Click the red keyframe button to the right of the Displacement parameter to remove it.



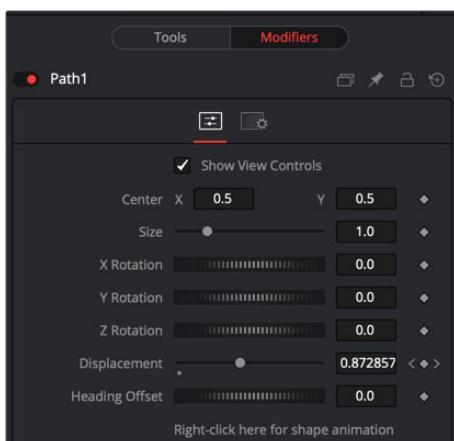
Remove the automatic displacement keyframe prior to creating your own keyframes.

- 10 At the bottom of the Modifiers tab, right-click on "Right-click here for shape animation" and choose Connect To > Polygon1Polyline.
- 11 To quickly see where your object has gone, drag the Displacement slider back and forth.
- 12 You may want to use the Size parameter to adjust the size of the overall path.

The Displacement slider is meant to be keyframed for animating the object along the path.

Path Modifier

In terms of functionality, it makes no difference which method you use to generate the path modifier. All the above methods are just different ways to get to the same point. Whichever way you decide to add the path modifier, the Modifiers tab contains controls for the path.




Creating a path adds controls to the Modifiers tab in the Inspector.

You can use the path modifier controls in the Inspector to change the position, size, and rotation of the entire path shape. The Displacement parameter is represented as a spline in the Spline Editor, which determines the object's position along the path and the Heading Offset is used for the orientation along the path.

Controlling Speed and Orientation along a Path


Every Polyline path has an associated Displacement spline in the Spline Editor. The Displacement spline represents acceleration, or the position of the animated object along its path, represented as a value between 0.0 and 1.0.



The Displacement curve of a Poly path represents the acceleration of an object on a path.

Smaller values are closer to the beginning of a path, while larger values are increasingly closer to the end of the path.

For instance, let's say you have a bumblebee that bobs up and down as it moves across the screen. To have the bee accelerate as it moves up and down but slow down as it reaches its peaks and valleys you use the Displacement curve.



A curvy path defined by a spline shape

The curved shape path does not define how fast the bee moves. The speed of the bee at any point along the path is a function of the Displacement parameter. You can control the Displacement parameter either in the Modifiers tab or in the Spline Editor.

To use the Displacement parameter to control the speed of an object on a path, do the following.

- 1 Position the playhead at the start of the animation.
- 2 In the Modifiers tab, drag the Displacement parameter to 0.0. This positions the object at the start of the path.
- 3 Click the Keyframe button to the right of the Displacement parameter.
- 4 Position the playhead somewhere further into the comp and drag the Displacement parameter until the object is where you want it to be based on the current frame.
- 5 Continue updating the playhead and the Displacement parameters at key points in the comp until you have reached the end of the path.


After the initial animation is set, you can use the Displacement curve in the Spline Editor to adjust the timing.

To adjust the timing of an object along a path:

- 1 Open the Spline Editor and enable the Displacement spline in the header.
- 2 Move control points horizontally closer together to increase the speed between the two points while maintaining the location of the object along the path.

A longer, flatter curve between two points indicates a slower rate of change.

- 3 Drag a control point up or down to change its location on the path while maintaining the timing between two points.



A Displacement curve in the Spline Editor

TIP: Holding down the Option key while clicking on the spline path in the viewer will add a new point to the spline path without adding a Displacement keyframe in the Spline Editor. This allows you to refine the shape of the path without changing the speed along the path.


Using a Path Modifier to Adjust Orientation

The Heading parameter is used to adjust the orientation of the object along the path. For instance, if you want the bee in our animation to auto-orient based on the direction of the path, you can connect the bee's angle to the Heading parameter.

To connect the Heading to an object's angle:

- 1 Right-click over an object's angle parameter in a Transform node.
- 2 Choose Connect To > Path > Header.

In our example comp, the bee now auto-oriens as the path descends and rises.




The Transform's angle parameter connected to the path modifier's Heading parameter

XY Path

Unlike a Polyline path, the XY path modifier uses separate splines in the Spline Editor to calculate position along the X-axis and along the Y-axis.

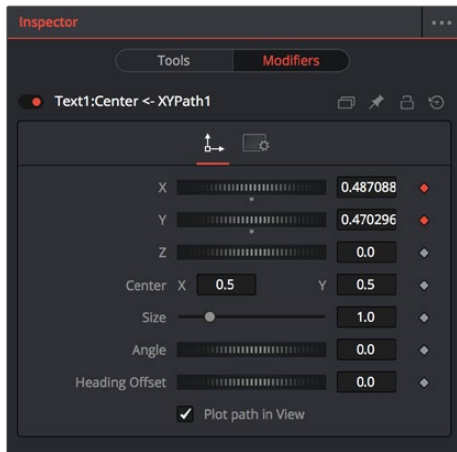
To animate a coordinate control using an XY path modifier:

- Right-click on the center coordinate control in the viewer or the Center X/ Y parameter in the Inspector, and then choose Modify With > XY Path from the contextual menu.



Adding the XY Path modifier to a Center parameter in the Inspector

At first glance, XY paths work like Polyline paths. To create the path once the modifier is applied, position the playhead and drag the onscreen control where you want it. Position the playhead again and move the onscreen control to its new position. The difference is that the control points are only there for spatial positioning. There is no Displacement parameter for controlling temporal positioning.



XY path modifier controls in the Modifiers tab of the Inspector

Instead of dragging in the viewer, you can use the controls in the Modifiers tab to create a motion path, while using the object's original Inspector controls as an offset to this motion path. You can use the XYZ parameters to position the object, the Center X/Y parameters to position the entire path, the Size and Angle to scale and rotate the path, and the Heading Offset control to adjust the orientation.



Using an XY path modifier to animate a piece of text

Using the XY Paths in the Spline Editor

The Spline Editor for the XY path displays the X and Y channel splines. Changes to the path in the viewer or the Inspector will be displayed as keyframes on these splines in the Spline Editor. Unlike a Polyline path, XY paths do not include a Displacement curve. The speed of the object along the path is tied to the path itself and cannot be separated from the timing of the keyframes that define that path.

TIP: XY path and Poly path can be converted between each other from the contextual menu. This gives you the ability to change methods to suit your current needs without having to redo animation.

The advantage of the XY path modifier is that you can explicitly set an XY coordinate at a specific time for more control.



XY Path displays X and Y curves in the Spline Editor but does not include a Displacement control.

Switching Default Paths

If you want to change the default path type to XY path, you can choose Fusion > Preferences > Globals in Fusion Studio or Fusion > Fusion Settings in DaVinci Resolve. In the Settings window, select the Defaults category and choose XY Path for the Points With drop-down menu. The next time you keyframe the Center X/Y parameter or choose Animate from CenterX/Y's contextual menu, an XY path modifier will be used instead of a Polyline path.


Types of Control Points

The control points along an XY path in the viewer are locked to the control points on the X and Y curve in the Spline Editor. The number of points are identical, and adding a control point in one place adds it to the other. That is not the case with a Polyline path. Polyline paths are composed of locked and unlocked points. Whether a point is locked is determined by how it was added to the Polyline. Locked points on the motion path in the viewer will have an associated point on the Displacement spline in the Spline Editor; unlocked points will not have corresponding points. Each has a distinct behavior, as described below.

Locked Points

Locked points are the motion path equivalents of keyframes. They are created by changing the playhead position and moving the animated control. These points indicate that the animated control must be in the specified position at the specified frame.

The locked points are displayed as larger-sized hollow squares in the viewer. Each locked key has an associated point on the path's Displacement curve in the Spline Editor.




Locked points on a path in the viewer

Deleting a locked point from the motion path will change the overall timing of the motion.

To change the duration of a path using locked points:

- 1 Connect an object to a Transform node.
- 2 Position the where you want to start the motion path.




A graphic placed on the right side of the frame

- 3 Set the playhead at frame 0.
- 4 Select the Transform node and in the Inspector, click the Keyframe button to the right of the Center X/Y parameter.

This adds the path modifier and creates the first locked point of the path.

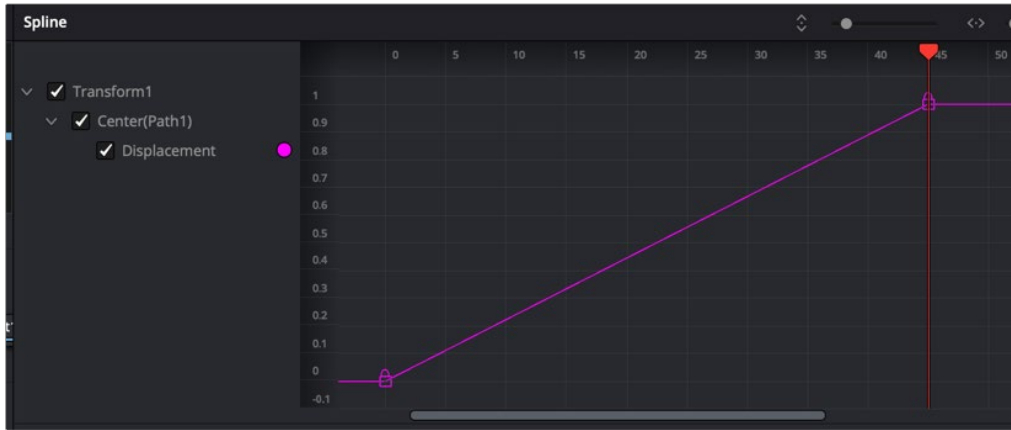
- 5 Position the playhead at frame 45.
- 6 Move the object's center to the lower center of the screen.



Moving the playhead and repositioning the bee adds a second locked point

This sets the second locked point.

- 7 View the Spline Editor and display Path1's Displacement spline.



The path's Displacement spline with locked keyframes


At a value of 0.0, the control will be located at the beginning of the path. When the value of the Displacement spline is 1.0, the control is located at the end of the path.

- 8 Select the keyframe at frame 45 in the Displacement spline and drag it to frame 50.

The motion path is now 50 frames long, without making any changes to the motion path's shape.

If you try to change this point's value from 1.0 to 0.75, it cannot be done because the point is the last in the animation, so the value must be 1.0 in the Displacement spline.

- 9 Position the playhead on frame 100 and move the bee center to the upper-left corner of the screen.



Moving locked points changes the duration of a motion path without changing its shape

This will create an additional locked point and set a new ending for the path.

Unlocked Points


Unlocked points are created when additional points are added to the motion path while in Insert and Modify modes. These points are used to adjust the overall shape of the motion path, without directly affecting the timing of the motion. This means you can add whatever unlocked points you want to reshape a motion path, without that having any effect on the timing of the animation on that path. This makes it vastly easier to fine-tune a motion path spatially after you've perfected the timing of the animation temporally.

Unlocked points *do not* have corresponding points on the path's Displacement spline. They appear in the viewer as smaller, solid square points.

To add unlocked points to a motion path, do the following:

- 1 Select a motion path spline by using the Tab key to cycle controls until the path is selected.
- 2 To insert points along a path, click the Insert and Modify button in the toolbar.
- 3 Click on the path and create two new points: one half-way between the first and the second points, and the other half-way between the second and the third points.

The two points just added are not present in the motion path's Displacement spline. These are unlocked points, used to shape the motion but unrelated to the timing of the path.




Unlocked points added to the motion path are not displayed on the Displacement spline

You can add unlocked points to the Displacement spline as well. Additional unlocked points in the Spline Editor can be used to make the object's motion pause briefly.

To pause motion along a motion path using an unlocked point:

- Select a locked point on the Displacement spline and then hold down the Command key while dragging the point horizontally to another frame. The point is copied as an unlocked point to the new frame.



Create a pause in the motion by copying locked points

Knowing the difference between locked and unlocked points gives you independent control over the spatial and temporal aspects of motion paths.

Locking and Unlocking Points

You can change an unlocked point into a locked point, and vice versa, by selecting the point(s) and choosing the Lock Point option from the contextual menu.

Tips for Manipulating Motion Paths

There are a variety of ways you can create and edit motion paths in the viewer.

Compound Motion Paths Using Path Centers


Every motion path has a defined center represented by a crosshair. Path centers allow paths to be connected to other controls and behave in a hierarchical manner, which is an exceptionally powerful way of creating complex motion by combining relatively simple paths.

A useful example of this technique would be animating the path of a bee in flight. A bee often flies in a constant figure eight pattern while simultaneously moving forward. The easy way of making this happen involves two paths working together.

In the following example, the bee would be connected to a first path in a Transform node, which would be a figure eight of the bee moving in place. This first path's center would then be connected to another path defining the forward motion of the bee through the scene via a second Transform node.

To create a compound motion path:

- 1 Create a figure-eight motion path by keyframing an object or drawing a polyline mask.
(If using a polygon mask, you'll need to remove the auto-keyframing and publish the mask prior to the next step.)
- 2 Add a polyline mask and create a smooth curve spline that travels across the screen.
- 3 At the bottom of the Inspector, right-click over the "Right-click here for shape animation" label and choose Remove Polygon Polyline to remove the auto-animation behavior.
- 4 Right-click over the label again and choose Publish.
- 5 Select the object's Transform node and click the Modifiers tab.
- 6 Right-click over the Path1 Center X/Y parameter and choose Path.
- 7 At the bottom of Path2, choose Connect To > Polygon: Polyline.
- 8 Keyframe the Path2 Displacement to move the object along the second path.



Two motion paths working together

Copying and Pasting Motion Paths

It is possible to copy an entire motion path to the clipboard and then paste it onto another node or path or composition.

Methods of copying and pasting motion paths:

- **To copy a motion path:** In the Inspector's Modifiers tab, right-click on the path's control header and choose Copy from the contextual menu.
- **To cut a motion path out of a node:** In the Inspector, right-click on the path's control header and choose Cut from the contextual menu.
- **To paste the copied path over another path:** In the Inspector, right-click on the path's control header and choose Paste from the contextual menu.

When pasting a path, the old motion path will be overwritten with the one from the clipboard.

Removing Motion Paths

There are multiple ways to remove or delete a path, and they all involve a right-click contextual menu. Removing a path modifier does not remove the object or the spline shape; it only removes the animation from the object and the modifier from the Modifiers tab in the Inspector.

To remove a path, do one of the following:

- In the Modifiers tab, right-click over the Path1 header and choose Delete from the menu.
- In the Inspector, right-click over the Center XY parameter and choose Remove Path1 from the menu.
- Right-click the center coordinate control in the viewer for the object you're animating, and choose Remove Path1 from the submenu of the NameOfObject; Center submenu.



Removing an entire motion path at once

Recording Motion Paths

You can animate both of the control's spatial and temporal information at the same time using the Record mode. This is useful when both position and speed are crucial to achieve the desired result.

Right-click on the desired path and select Record from the contextual menu. This displays a submenu of available data that may be recorded.

Use the Record Time option in conjunction with the Draw Append mode to create complex motion paths that will recreate the motion precisely as the path is drawn.

The time used to record the animation may not suit the needs of a project precisely. Adjust the path's Displacement spline in the Spline Editor to more correctly match the required motion.

Importing and Exporting Polylines

You can import and export polyline shapes into a common editable ASCII text file or its native format. These methods are used to save a particularly useful or generic mask or path for future use or for use in another application, such as Maya or LightWave. You can also import FXF, SSF, or Nuke shape files.

Native Format

To save a polyline shape in Fusion's native ASCII format, right-click on the header of the Mask node in the Inspector and select Settings > Save As from the contextual menu. Provide a name and path for the saved file and select OK to write a file with the .setting extension. This file will save the shape of a mask or path, as well as any animation splines applied to its points or controls.

To load the saved setting back into Fusion, you first create a new polyline of the same type, and then select Settings > Load from the mask's context menu or drag the .setting file directly into the Node Editor.

If you want to move a polyline from one composition to another, you can also copy the node to the clipboard, open your second composition, and paste it from the clipboard into the new composition.

Chapter 72

Using Modifiers, Expressions, and Custom Controls

Some of the most powerful aspects of Fusion are the different ways it allows you to go beyond the standard tools delivered with the application. This chapter provides an introduction to a variety of advanced features, including Modifiers, Expressions, and Scripting, which can help you extend the functionality and better integrate Fusion into your studio.

Contents

The Contextual Menu for Parameters in the Inspector	1265
Using Modifiers	1265
Adding the Right Modifier for the Job	1265
Adding Modifiers to Individual Parameters	1265
Combining Modifiers and Keyframes	1266
Publishing a Parameter	1266
Connecting Multiple Parameters to One Modifier	1267
Adding and Inserting Multiple Modifiers	1267
Performing Calculations in Parameter Fields	1268
Using SimpleExpressions	1268
Pick Whipping to Create an Expression	1271
Removing SimpleExpressions	1271
Customizing User Controls	1271
FusionScript	1274

The Contextual Menu for Parameters in the Inspector

Most of the features in this chapter are accessed via a contextual menu that appears when you right-click most parameters in the Inspector. Different contextual menus are available based on where in the Inspector you right-click. Specifically, right-clicking over parameter names or sliders displays a feature-rich contextual menu that can add animation, expression fields, modifiers to extend functionality, as well as publishing and linking capabilities, allowing you to adjust multiple controls simultaneously.

Using Modifiers


Parameters can be controlled with modifiers, which are extensions to a node's toolset. Many modifiers can automatically create animation that would be difficult to achieve manually. Modifiers can be as simple as keyframe animation or linking the parameters to other nodes, or modifiers can be complex expressions, procedural functions, external data, third-party plug-ins, or fuses.

Adding the Right Modifier for the Job

Which modifiers are available depends on the type of parameter you're right-clicking over. Numeric values, text, polylines, gradients, and points each have different sets of modifiers that will work with them, so the Modify With menu is filtered based on each parameter.

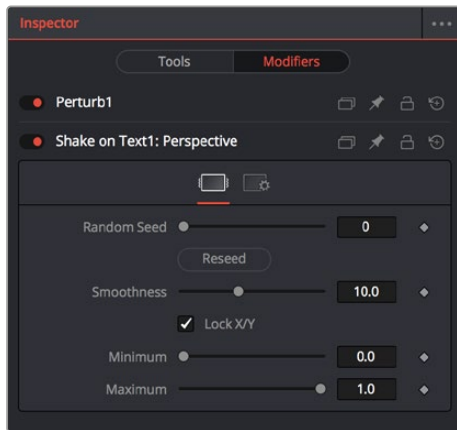
Adding Modifiers to Individual Parameters

You add modifiers to a parameter using the Inspector's contextual menu, or by right-clicking the onscreen control for a parameter in the viewer. Either way, a dynamic list of modifiers that are appropriate for the selected parameters is displayed. For instance, a Perturb modifier can be added to any slider to auto-animate the parameter by randomly wiggling the value. Once the modifier is added, controls are displayed in the Modifiers tab to adjust the speed and amplitude of the random animation.



The Inspector contextual menu is used to add a Perturb modifier to the Center X and Y parameters

A modifier's controls are displayed in the Modifiers tab of the Inspector. When a selected node has a modifier applied, the Modifiers tab will become highlighted as an indication. The tab remains grayed out if no modifier is applied.



The Modifiers tab with two modifiers applied

Modifiers appear with header bars and header controls just like the tools for nodes. A modifier's title bar can also be dragged into a viewer to see its output.

Combining Modifiers and Keyframes

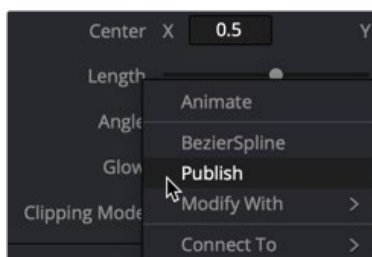
Modifiers that auto-animate parameters like Perturb and Shake can be combined with keyframes to create more natural, organic looking animations. For instance, you can create the general motion path for an element by keyframing the Center X and Y parameters, and then apply a modifier to the same parameters to create a secondary wiggling motion.

To combine a keyframed motion path with a Perturb modifier:

- 1 Add a Transform to an image like a butterfly or spaceship.
- 2 Select the Transform node in the Node Editor.
- 3 In the Inspector, click the Keyframe button to the right of the Center X and Y parameters.
- 4 Position the butterfly or spaceship image where you want the start of the animation to begin.
- 5 Continue to move the playhead in the render range and reposition the image until you create a figure-8 motion path.
- 6 Right-click over the Center X label in the Inspector and choose Modify With > Perturb.
- 7 Click the Modifiers tab at the top of the Inspector and adjust random, wiggling motion by setting the Strength, Wobble, and Speed parameters while the animation plays.

Publishing a Parameter

The Publish modifier makes the value of a parameter available, so that other parameters can connect to it. This allows you to simultaneously use one slider to adjust other parameters on the same or different nodes. For instance, publishing a motion path allows you to connect multiple objects to the same path.




Publish a parameter in order to link other parameters to it

Once a parameter is published, you can right-click another parameter and choose **Connect To > [published parameter name]** from the contextual menu. The two values are linked, and changing the parameter value of one in the Modifiers tab changes the other.

Using the pick whip between two parameters provides similar linking behavior with more flexibility. Pick whipping between parameters is covered later in this chapter.

Connecting Multiple Parameters to One Modifier

A single modifier or published parameter can be applied to multiple parameters so they all act as one. This is handled through the **Connect To** submenu. As with modifier assignment, the list is filtered to show only options that are suitable for the parameter you've right-clicked on. When you do this, the connection is bidirectional; editing either parameter will cause the other parameter to change.



The **Connect To** menu connects modifiers to multiple parameters

Adding and Inserting Multiple Modifiers

Modifiers can be connected to each other and branched, just like nodes in the Node Editor. For example, the **Calculation** modifier outputs a number, but has two **Number** parameters, both of which can have modifiers added to them. If you want to insert a modifier between the existing modifier and the modified parameter, use the **Insert** submenu of the parameter's contextual menu.

Available Modifiers in Fusion:

- **Bézier Spline:** Adds a Bézier spline to the Spline Editor for animating the selected parameter.
- **B-Spline Spline:** Adds a B-Spline spline to the Spline Editor for animating the selected parameter.
- **Calculation:** Creates an indirect link that includes a mathematical expression between two parameters.
- **CoordTransform Position:** Calculates the current 3D position of a given object even after multiple 3D transforms have repositioned the object through the node tree hierarchy.
- **Cubic Spline:** Adds a Cubic spline to the Spline Editor for animating the selected parameter.
- **Expression:** Allows you to add a variable or a mathematical calculation to a parameter, rather than a straight numeric value. The Expression modifier provides controls in the Modifiers tab, giving you more room and parameters than the SimpleExpression
- **From Image:** This modifier takes color samples of an image along a user-definable line and creates a gradient from those samples.
- **Gradient Color Modifier:** Creates a customized gradient and maps it into a specified time range to animate a value.
- **KeyStretcher:** Used to stretch keyframes in a Fusion Title template when trimming the template in the Edit page or Cut page timeline.
- **MIDI Extractor:** Modifies the value of a parameter using the values stored in a MIDI file.
- **Natural Cubic Spline:** Adds a Natural Cubic spline to the Spline Editor for animating the selected parameter.

- **Offset (Angle, Distance, Position):** The three Offset modifiers are used to create variances, or offsets, between two positional values. For instance, when this modifier is added to a size parameter, you can change the size of an object using the distance between two onscreen controls (position and offset).
- **Path:** Produces two splines to control the animation of of an object: An onscreen motion path (spatial) and a Time spline visible in the Spline Editor (temporal).
- **Perturb:** Generates smoothly varying random animation for a given parameter.
- **Probe:** Auto-animates a parameter by sampling the color or luminosity of a specific pixel or rectangular region of an image.
- **Publish:** The first step in linking two non-animated parameters is to use the Publish modifier to publish a parameter. That allows other parameters to use the Connect To submenu and link to the published parameter.
- **Shake:** Similar to Perturb, Shake generates smoothly varying random animation for a given parameter.
- **Track:** Attaches a single point tracker to the selected parameter. The tracker can then track an object onscreen to animate the parameter. This is quicker and more direct than using the normal Tracker node; however, it offers less flexibility since the resulting tracker is only a single point and can only be used for the selected parameter.
- **Vector Result:** Similar to the Offset modifier, Vector Result is used to offset position parameters using origin, distance, and angle controls to create a vector. This vector can then be used to adjust any other parameter.
- **XY Path:** Produces an X and Y spline in the Spline Editor to animate the position of an object.


For more information on all modifiers available in Fusion, see Chapter 121, “Modifiers.” In the DaVinci Resolve Reference Manual and Chapter 61 in the Fusion Reference Manual.

Performing Calculations in Parameter Fields

You can enter simple mathematical equations directly in a number field to calculate a desired value. For example, typing $2.0 + 4.0$ in most number fields will result in a value of 6.0. This can be helpful when you want a parameter to be the sum of two other parameters or a fraction of the screen resolution.


Using SimpleExpressions

Simple Expressions are a special type of script that can be placed alongside the parameter it is controlling. These are useful for setting simple calculations, building unidirectional parameter connections, or a combination of both. You add a SimpleExpression by entering an equals sign directly in the number field of the parameter and then pressing Return.



Entering an equals sign opens the SimpleExpression field with Pick Whip control

An empty field will appear below the parameter, and a yellow indicator will appear to the left. The current value of the parameter will be entered into the number field. Using Simple Expressions, you can enter a mathematical formula that drives the value of a parameter or even links two different parameters. This helps when you want to create an animation that is too difficult or impossible to set up with keyframing. For instance, to create a pulsating object, you can use the sine and time functions on a Size parameter. Dividing the time function can slow down the pulsing while multiplying it can increase the rate.



A SimpleExpression using the sine and time functions

Inside the SimpleExpression text box, you can enter one-line scripts in Lua with some Fusion-specific shorthand. Some examples of Simple Expressions and their syntax include:

Expression	Description
<code>time</code>	This returns the current frame number.
<code>Merge1.Blend</code>	This returns the value of another input, Blend, from another node, Merge1.
<code>Merge1:GetValue("Blend", time-5)</code>	This returns the value from another input, but sampled at a different frame, in this case five frames before the current one.
<code>sin(time/20)/2+.5</code>	This returns a sine wave between 0 and 1.
<code>iif(Merge1.Blend == 0, 0, 1)</code>	This returns 0 if the Blend value is 0, and returns 1 if it is not. The <code>iif()</code> function is a shorthand conditional statement, if-then-else.
<code>iif(Input.Metadata.ColorSpaceID == "sRGB", 0, 1)</code>	This returns 0 if the image connected to the current node's Input is tagged with the sRGB colorspace. When no other node name is supplied, the expression assumes the Input is coming from the current node. It is equivalent to <code>self.Input</code> . The Input in most, but not all, Fusion nodes is the main image input shown in the Node Editor as an orange triangle. Images have members that you can read, such as Depth, Width, Metadata, and so on.
<code>Point(Text1.Center.X, Text1.Center.Y-.1)</code>	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.
<code>Text1.Center - Point(0,.1)</code>	This is similar to the previous expression. This SimpleExpression returns Text instead of a Number or Point.

Expression	Description
<code>Text("Colorspace: "..(Merge1.Background.Metadata.ColorSpaceID))</code>	The string inside the quotes is concatenated with the metadata string, perhaps returning: Colorspace: sRGB
<code>Text("Rendered "..os.date("%b %d, %Y").. " at "..os.date("%H:%M").."\n on the computer "..os.getenv("COMPUTERNAME").. " running "..os.getenv("OS").."\n from the comp "..ToUNC(comp.FileName))</code>	This returns a much larger Text, perhaps something like: Rendered Nov 12, 2019 at 15:43 on the computer Rn309 running Windows_NT from the comp \\SRVR\Proj\Am109\SlateGenerator_A01.comp
<code>os.date("%H:%M")</code>	The OS library can pull various information about the computer. In the previous example, os.date gets the date and time in hours:minutes.
<code>"..os.getenv("COMPUTERNAME").. " running "..os.</code>	Any environment variable can be read by os.getenv, in this case the computer name and the operating system.
<code>"\n from the comp "..ToUNC(comp.FileName)</code>	To get a new line in the Text, \n is used. Various attributes from the comp can be accessed with the comp variable, like the filename, expressed as a UNC path.

TIP: When working with long SimpleExpressions, it may be helpful to drag the Inspector panel out to make it wider or to copy/paste from a text editor or the Console.

After setting an expression that generates animation, you can open the Spline Editor to view the values plotted out over time. This is a good way to check how your SimpleExpression evaluates over time.



A sine wave in the Spline Editor, generated by the expression used for Text1: Size

For more information about writing Simple Expressions, see the Fusion Studio Scripting Guide, and the official Lua documentation.

Pick Whipping to Create an Expression

With a SimpleExpression field open, a + button is displayed on the left side. Dragging the + button onto another control, or “pick whipping,” links the two parameters, similar to the Connect To menu. However, unlike a Connect To parameter link, the pick whip allows you to modify the connection by further editing the expression.



Pick whipping to connect one parameter to another quickly

SimpleExpressions can also be created and edited within the Spline Editor. Right-click on the parameter in the Spline Editor and select Set Expression from the contextual menu. The SimpleExpression will be plotted in the Spline Editor, allowing you to see the result over time.


Removing SimpleExpressions

To remove a SimpleExpression, right-click the name of the parameter, and choose Remove Expression from the contextual menu.

Customizing User Controls

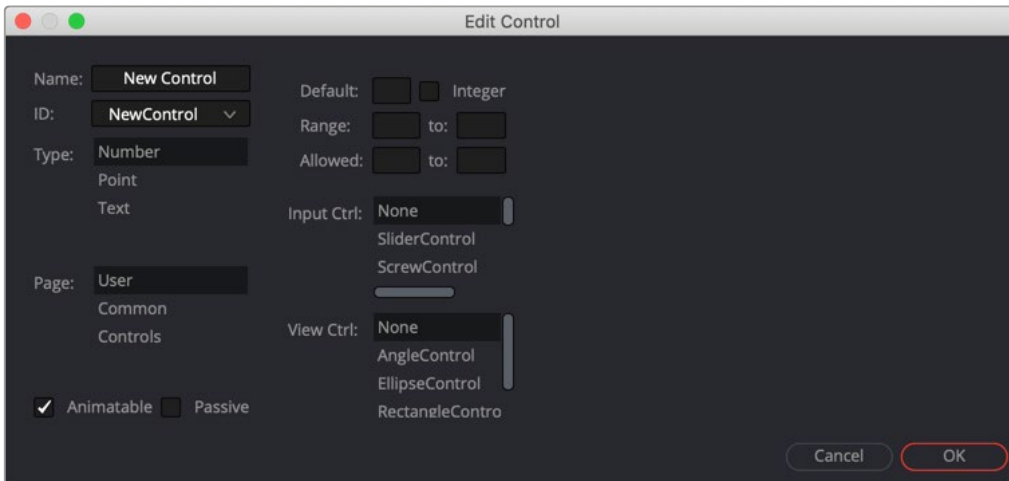
Each tool’s parameters are organized in a logical order in the Inspector. The most used controls are closer to the top, and the more subtle refinement controls are lower in the list. Sometimes, though, you may want to add, hide, or change the controls. You often need to do this for SimpleExpressions and macros, but you may also do this for usability and aesthetic reasons for favorites and presets.

Custom controls can be added or edited via the Edit Control dialog, which you access by right-clicking over the node’s name in the Inspector and choosing Edit Controls from the menu.



Choose Edit Controls to create and customize parameters

In the Edit Control dialog, you use the ID menu to select an existing parameter or create a new one. You can name the control and define whether it is a text field, number field, or a point using the Type attributes list.




Edit Control dialog

You use the page list to assign the new control to one of the tabs in the Inspector. There are also settings to determine the defaults and ranges, and whether it has an onscreen preview control. The Input Ctrl box contains settings specific to the selected Type, and the View Ctrl attributes box contains a list of onscreen preview controls to be displayed, if any.

All changes made using the Edit Controls dialog get stored in the current tool instance, so they can be copied/ pasted to other nodes in the comp. However, to keep these changes for other comps, you must save the node settings, and add them to the Bins in Fusion Studio or to your favorites.

As an example, we'll customize the controls for a DirectionalBlur:



The Inspector for a default direction blur


Let's say we wanted a more interactive way of controlling a linear blur in the viewer, rather than using the Length and Angle sliders in the Inspector. Using a SimpleExpression, we'll control the length and angle parameters with the Center parameter's onscreen control in the viewer. The SimpleExpression would look something like this:

For Length:

```
sqrt(((Center.X-.5)*(self.Input.XScale))^2+((Center.Y-.5)*(self.Input.YScale)*(self.Input.Height/self.Input.Width))^2)
```

For Angle:

```
atan2(.5-Center.Y , .5-Center.X) * 180 / pi
```




DirectionalBlur controlled by the Center's position.

This admittedly somewhat advanced function does the job fine. Dragging the onscreen control adjusts the angle and length for the directional blur. However, now the names of the parameters are confusing. The Center parameter doesn't function as the center anymore. It is the direction and length of the blur. It should be named "Blur Vector" instead. You no longer need to edit the Length and Angle controls, so they should be hidden away, and since this is only for a linear blur, we don't need the Type menu to include Radial or Zoom. We only need to choose between Linear and Centered. These changes can easily be made in the Edit Controls dialog.

To change the Inspector parameters for the example above:

- 1 In the Edit Control dialog, select the Center from the ID list.
- 2 A dialog will appear asking if you would like to Replace, Hide, or Change ID. We'll choose Replace.
- 3 Change the Name to Blur Vector.
- 4 Set the Type to Point.
- 5 Select Controls in the Page list. (Controls is the first tab in the Inspector, normally.)
- 6 Click OK.



DirectionalBlur Center parameter name changed to Blur Vector.

The new Blur Vector parameter now appears in the Inspector. The internal ID of the control is still Center, so our SimpleExpressions did not change.

To hide the Length and Angle parameters in the Inspector:

- 1 In the Edit Control dialog, select the Length from the ID list.
- 2 Select Controls from the Page list.
- 3 In the input Ctrl list, select Node.

- 4 Click OK.
- 5 In the Edit Control dialog, select the Angle from the ID list.
- 6 In the input Ctrl list, select Node.
- 7 Click OK.

Finally, to remove Radial and Zoom options from the Type menu:

- 1 In the Edit Control dialog, select the Type from the ID list.
- 2 Select Controls from the Page list.
- 3 Select Radial from the Items list and click Del to remove it.
- 4 Select Zoom from the Items list and click Del to remove it.
- 5 Click OK.

The Type menu now includes only two options.

If you want to replace the Type menu with a new checkbox control, you can do that by creating a new control and a very short expression.

To create a new control:

- 1 In the Edit Control dialog, enter Center Blur in the Name field.
- 2 Select the New Control from the ID list.
- 3 Set the Type to Number, and set the Page to Controls.
- 4 Set the Input Ctrl to CheckboxControl.
- 5 Click OK.

To make this new checkbox affect the original Type menu, you'll need to add a SimpleExpression to the Type:

```
iif(TypeNew==0, 0, 2)
```

The “iif” operator is known as a conditional expression in Lua script. It evaluates something based on a condition being true or false.

FusionScript

Scripting is an essential means of increasing productivity. Scripts can create new capabilities or automate repetitive tasks, especially those specific to your projects and workflows. Inside Fusion, scripts can rearrange nodes in a comp, manage caches, and generate multiple output files for delivery. They can connect Fusion with other apps to log artist time, send emails, or update webpages and databases.

FusionScript is the blanket term for the scripting environment in Fusion. It includes support for Lua as well as Python 2 and 3 for some contexts. FusionScript also includes libraries to make certain common tasks easier to do with Lua and Python within Fusion.

You can run interactive scripts in various situations. Common scripts include:

- Utility Scripts, using the Fusion application context, are found under the File > Scripts menu.
- Comp Scripts, using the composition context, are found under the Script menu or entered into the Console.

- Tool Scripts, using the tool context, are found in the Tool's context menu > Scripts.

Other script types are available as well, such as Startup Scripts, Scriptlibs, Bin Scripts, Event Suites, Hotkey Scripts, Intool Scripts, and SimpleExpressions. Fusion Studio allows external and command-line scripting as well and network rendering Job and Render node scripting.

FusionScript also forms the basis for Fuses and ViewShaders, which are special scripting-based plug-ins for tools and viewers that can be used in both Fusion and Fusion Studio.

For more information about scripting, see the Fusion Scripting Documentation, accessible from the Documentation submenu of the Help menu.

Chapter 73

Bins

This chapter covers the bin system in Fusion Studio. Bins allow for the storage and organization of clips, compositions, tool settings, and macros, similar to the Media Pool and Effects Library in DaVinci Resolve. It includes a built-in Studio Player for creating a playlist of multiple shots and their versions. Bins can be used in a server configuration for organizing shots and collaborating with other team members across the studio.


Contents

Bins Overview	1277
Bins Interface	1277
Viewing and Sorting Bins	1278
Organizing Bins	1279
Adding and Using Content	1281
File Type Details	1282
Using Content from Bins	1282
Jog and Shuttle	1283
Stamp Files	1283
Using the Studio Player	1284
Playing a Single Clip	1284
Creating a Reel	1285
Connecting Bins Over a Network	1292
Adding a Remote Bin Entry	1293
Accessing Remote Bins	1293
Permissions	1293
Studio Player and Bin Server	1293

Bins Overview

Bins are folders that provide an easy way of accessing commonly used tools, settings, macros, compositions, and footage. They can keep all your custom content and resources close at hand, so you can use them without searching through your hard drives. Bins can also be shared over a network to improve a collaborative workflow with other Fusion artists.

Bins are only available in Fusion Studio.




The Bins window

Bins Interface

The Bins window is separated into two panels. The sidebar on the left is a list of the bins, while the Content panel on the right displays the selected bin's contents.

To open the Bins window, do one of the following:

- Choose File > Bins from the menu bar.
- Press Command-B.




The Bins sidebar

TIP: When opening the Bins window on macOS, the window may open behind the current Fusion Studio window. Check in the dock for the Bins window icon or move the Fusion Studio window to locate the Bins window if you are working on a single monitor.

The sidebar organizes content into bins, or folders, using a hierarchical list view. These folders can be organized however they suit your workflow, but standard folders are provided for Clips, Compositions, Favorites, Projects, Reels, Settings, Templates, and Tools. The Tools category is a duplicate of all the tools found in the Effects library. The Tools bin is a parent folder, and parent folders contain subfolders that hold the content. For instance, Blurs is a subfolder of the Tools parent folder. Parent folders can be identified by the disclosure arrow to the left of their name.

When you select a folder from the sidebar, the contents of the folder are displayed in the Contents panel as thumbnail icons.



The Bins icon view

Viewing and Sorting Bins

A contextual menu is used to access most of a bin's features. You show the contextual menu by right-clicking in an empty area in the Contents panel. Right-clicking on an item will show the same menu with additional options for renaming, playing, or deleting the item.


Icon or List View

One use of the contextual menu is to switch between viewing the contents as thumbnail icons or as a list.

To view a bin's contents in List view:

- 1 Right-click in an empty area of the Contents panel.
- 2 From the contextual menu, choose View > Details.

Each bin in the sidebar can be set to List view or Icon view independently of each other. So while you may have some bins you want to see as a list, others may be easier to view as icons.




The Bins List view

Sort Order in List View

Clicking on the heading of a column in List view will sort the list in ascending order, and clicking it again will reverse the sort order.

Icon Size in Icon View

The icons can be adjusted to small, medium, large, or huge by clicking the Size button in the bottom toolbar or right-clicking in an empty area of the Contents panel to bring up the contextual menu and choosing a size from the Icon Size submenu.




Use the Size button to select the icon size in the bin

Organizing Bins

Once you begin adding your own categories and content, you can have hundreds of items that need to be organized. To keep all of your elements accessible, you'll want to use some basic organization, just like keeping files and documents organized on your computer.

To create a new folder in the sidebar:

- 1 In the sidebar, select the parent folder under which the new folder will be listed.
- 2 Right-click in an empty area of the Contents panel.
- 3 From the contextual menu, choose New > New Folder.
- 4 Enter a name for the new folder, and then click OK on the dialog.



The New Folder menu


You can also click the New Folder icon in the toolbar.

To rename a bin folder:

- 1 Right-click on the folder icon in the Contents panel.
- 2 Choose Rename from the contextual menu or press F2 on the keyboard.

To move a folder into or out of a parent folder:

- 1 Select the parent folder that contains the folder you want to move.
- 2 In the Contents panel, drag the folder into the sidebar where you want it moved.



Drag a folder from the Contents panel to move it in the sidebar.

When you drag a folder onto another folder in the sidebar, you create a hierarchical subfolder. Dragging it to the Library parent folder at the top of the sidebar will add it to the top level of the Bins window.

To remove a folder from the bins:

- 1 Select the folder in the Contents panel.
- 2 Press Command-Delete (Backspace on Windows and Linux).


TIP: You cannot undo removing a folder from the Bins window.

Adding and Using Content

You can add and use different types of content with bins. Fusion Studio compositions, tools, saved tool settings, macros, tool groups, and any file format that is supported in Fusion can be added to bins and then used in compositions at a later time.

To add an item to a bin:


- 1 Select a bin in the sidebar where you want to add the content.
- 2 Right-click in the Contents panel.
- 3 Choose New > ... from the contextual menu.
- 4 Select the media, settings, macros, or comps in the file browser, and then click Open.



Add Item in the contextual menu

To remove an item from the bins:

- 1 Select the folder in the Contents panel.
- 2 Right-click on the item and choose Delete from the contextual menu or press Command-Delete.



The Delete function in the contextual menu

TIP: Unsupported files like PDFs can also be saved in the bin, and the application that supports it will launch when you double-click the file. It's a useful location for scripts and notes.

If you have an operating system file browser window open, you can drag files directly into a bin as well. When adding an item to the bins, the files are not copied. A link is created between the content and the bins, but the files remain in their original location.

File Type Details

Some types of content have additional methods of being added to bins. Some additional methods are because of the file type and some are because of where they are located.

Projects and Media

In addition to using the New... contextual menu as explained earlier, Fusion Studio project files with the extension “.comp” and media can also be added to bins by dragging them to the Contents panel from a file browser.

Tool Settings

If you want to add a node with custom settings to a bin, first save the node as a setting by right-clicking over it in the Node Editor and choosing Settings > Save As. Once the setting is saved, you can add it to a bin by dragging it from a File Browser into the Bins window.

Image Sequences and Stills

Image sequences are automatically identified on disk and loaded as clips rather than stills, so it is not necessary to select more than one frame from an image sequence when dragging it into a bin.

To ignore the image sequence and import only a single frame, hold Shift when you drag the frame into a bin. This can be useful when trying to import a single still image from a series of still shots with a DSLR. The numbers maybe sequential, but you just need one still image from the series.

Using Content from Bins

Once you have content in your bins, you'll want to add them to a composition. In general, you can either drag the content directly into the Node Editor or double-click it to add it; however, each type of content behaves a little differently when added to the Node Editor.

Media

Dragging media into the Node Editor from a bin window creates a new Loader that points to the media on disk. Still files or photos are automatically set to loop.

Compositions

To add a composition, you must right-click and choose Open. Dragging a comp item onto an open composition will have no effect. When a composition is added, it is opened in a new window. It is not added to the existing composition.

Tools, Tool Settings, and Macros


When you add tools from the Tools categories to a composition, the methods you use and results you get are similar to adding tools using the toolbar buttons or the Effects Library. Dragging a tool allows you to place it anywhere in the Node Editor, unconnected, or, if you drag it over a connection line, inserted between two existing tools. Double-clicking a tool in the bin will insert it after the Active tool in the Node Editor. Dragging a tool from a bin into a viewer will insert that tool after the currently viewed tool.

Settings and macros work a bit differently than tools. They can only be added to the Node Editor by dragging and dropping. Dragging a setting or macro allows you to place it in the Node Editor, unconnected, or, if you drag it over a connection line, inserted between two existing tools.

Jog and Shuttle

You can scrub clips in Icon view in the bin using one of two modes. Jog mode is the default mode. It moves the clip forward and backward as long as you are dragging the mouse. Once the mouse stops, the clip pauses.

You can choose Shuttle mode by right-clicking over the clip's thumbnail in the bin and choosing Scrub Mode > Shuttle.



Jog and Shuttle contextual menu

Shuttle mode begins playing the clip forward or backward once you press the right mouse button and drag either left or right. The clip continues to play until the mouse button is released or you reach the end of the clip.

Stamp Files

Stamp files are low-resolution, local proxies of clips, which are used for playback of clips stored on a network server, or for very large clips.

To create a stamp file for a clip:

- 1 Right-click on the clip in the bin.
- 2 Choose Create Stamp from the contextual menu.

The Status bar at the top of the Bins window shows the progress of the stamp creation. Since the stamp creation is a background process, you can queue other stamps and can continue working with the composition.

Using the Studio Player

The Studio Player is a timeline-based playback interface built into the Bins window. It allows you to play and organize versions of compositions, make notes, and collaborate on shots and projects. The resolution-independent player uses any format that Fusion Studio can ingest, like EXR, ProRes, BMD RAW, QuickTime, and others. Clips use a RAM cache for playback, so even large formats can loop playback as long as there is available memory. Clips can contain audio and are output to a video monitor using Blackmagic Design Decklink or UltraStudio devices, allowing you to screen dailies and review shots. You can add annotation notes to the project, and setting up a Bin Server allows multiple artists to access the Studio Player.

Here are some Studio Player highlights:


- Clips can be played as single events, looped, or using a ping-pong playback, with a definable loop range.
- Clip metadata can be viewed, with live update during scrub/play.
- The Timeline interface can create a playlist for reviewing multiple clips.
- Per-shot color adjustment controls allow for consistent display of clips from different formats.
- Annotation notes can be typed on each clip and version, as well as the entire project.
- Audio Scratch track can be enabled for each clip during playback.
- Clip versions are stored in the same project to allow for quick access to previous work and for comparison of progress.
- Guide overlays can be customized to show monitor/title safety and show crops to various output formats.
- Blackmagic Design UltraStudio and DeckLink playback devices are supported for reviewing clips on video monitors and projectors.
- The fully collaborative workflow automatically synchronizes reel changes, annotations, and color adjustments across multiple workstations, allowing multiple artists or supervisors to access the same projects simultaneously.
- Remote sync allows multiple Studio Players to follow the master. Actions performed on the master, such as playback and scrubbing, will also be executed on the slaves, allowing the reel to be reviewed across multiple workstations or sites.
- Studios can automate tasks using the Fusion scripting engine to control the Studio Player.

Playing a Single Clip

Clips created from image sequences, MOV files, and some RAW formats can be previewed using the Studio Player without having to first add the clip to a node tree.

To play a clip in the Studio Player, do one of the following:

- Double-click a clip in a bin to open the Studio Player.
- Select the clip and click the Play button at the bottom of the bin window.



The Studio Player includes a large viewer, a Timeline, and a toolbar along the bottom.

Once you have the clip open in the Studio Player, you can click the Play button in the toolbar at the bottom of the window.

Scrubbing the Timeline

You can quickly preview the clip by scrubbing through it rather than playing it.

To scrub a clip in Studio Player:

- 1 In the Timeline, drag the playhead to the area you want to scrub over.
- 2 Use the Left and Right Arrow keys to move one frame forward or backward.

Closing the Studio Player

After you have finished previewing in the Studio Player, you can return to the Bins to switch to a new clip or continue working in Fusion.

To close the current clip in the Studio Player and return to the bins:

- Click the three-dot Options menu in the lower-left corner and choose Close.

Creating a Reel

A reel is a playlist or clip list that is viewed either as storyboard thumbnails or a timeline. In the bin, you create a new reel item to hold and play back multiple clips. The thumbnail for a reel appears with a multi-image border around it, making it easier to identify in the Bin window.

To create a reel in the current bin:

- Right-click in an empty area of the bin and choose New > Reel
- Click the Reel button along the bottom toolbar.



Use the New Reel button to create a reel in the current bin.


Once created and named, the reel appears in the current bin.

Double-clicking the reel opens the Studio Player interface along the bottom of the Bin window. An empty Studio Player appears in the top half of the window.

To add a clip to a reel:

- 1 Click the three-dot Options menu in the lower-left corner and choose Show Bins.
- 2 Drag a clip from one of the bins to the empty reel in the lower half of the window.

The toolbar across the bottom of the interface has various controls for setting a loop, showing and adjusting color, playback transport controls, collaboration sync, guide overlays, and frame number and playback speed in fps.




The toolbar along the bottom of the Studio Player includes controls to customize playback

Toolbar buttons

- **Set Loop In/Out:** Sets the start and end frame for playing a section of the Timeline in a loop.
- **Shot:** Sets the loop for the entire clip.
- **Reset:** Disables the loop mode.
- **M:** Shows Metadata of the image.
- **RGB and alpha:** Toggles between displaying the color and alpha of the clip.
- **Brightness Gamma:** Adjusts the brightness and gamma of the viewer and is applied to all clips. Individual clip color controls can also be applied using another menu.
- **Video:** Outputs the image to Blackmagic Design DeckLink and UltraStudio devices.
- **Transport controls:** Used to play forward, backward, fast forward, and fast backward, as well as go to the start and end of a clip.
- **Sync:** A three-way toggle allowing the Studio Player to be controlled or to control another player over a network. The Off setting disables this functionality.
- **Guide buttons:** These three buttons control the visibility of three customizable guide settings.

Inserting Shots

Clips or comps from the bin can be dragged to the storyboard area of the reel to add and organize a playlist.




You can insert a shot between existing clips by positioning the new clip in between existing items in the reel.

To play the clips in a reel:


- 1 Click the three-dot Options menu in the lower-left corner and choose Show Player.
- 2 Position the playhead where you want to begin playing and click the Play button.

Creating Versions

Alternatively, you can add a version to an existing clip by dragging the new item on top.




Versions of a shot will appear as stacked icons in the storyboard reel. The number of stacks in the icon indicate the number of versions included with that clip. In the example below, the first shot has two versions, the second shot has four versions, and the last clip has only one version.



Version Menu

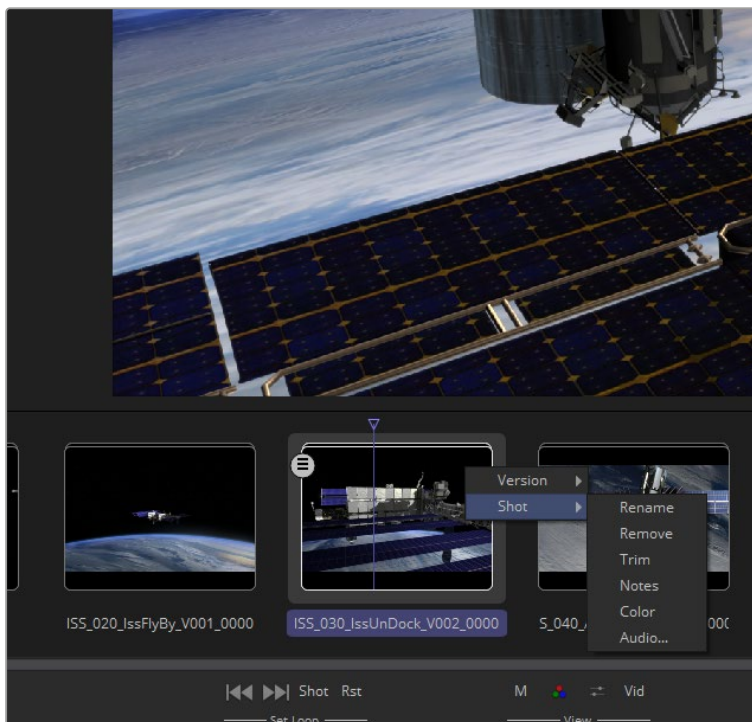
You can choose which version to view by right-clicking over the clip in the storyboard and selecting it from the Version > Select menu.




The Version menu also includes options to move or rearrange the order of the clip versions as well as remove a version, thereby deleting it from the stack.

Shot Menu

The per-clip Shot menu includes functions to Rename the shot, Remove the shot, Trim the clip's In and Out points, add notes, adjust the color, and add an audio soundtrack.




- **Rename:** Allows you to change the name of the shot.
- **Remove:** Deletes the entire shot and all the versions from the project reel.
- **Trim:** Opens the Trim dialog to adjust the clip In point and Out point on the Timeline.
- **Notes:** Opens the Notes window to the right of the interface, allowing you to add a note to the specific shot.
- **Color:** Opens the Color Adjustments panel to perform Lift, Gamma, Gain, Brightness, and Contrast adjustments to clips.
- **Audio:** Allows you to attach an audio file to a clip in the Reel.



When notes are added, they are time and date stamped as well as name stamped. The naming is from the bin login name and computer name.


Selecting Color from the Shot menu allows you to make tonal adjustments per clip using Color Decision List (CDL) style controls.



The Audio menu option can import an audio .wav file that will play back along with the selected clip.


Options Menu

The three-dot Options menu in the lower left of the interface displays the menu that can be used to switch between viewer and the bin in the top half of the window. It is also used to clear the memory used in playback by selecting Purge Cache.



The Options menu includes options to switch the top half of the window between the viewer or bin contents.

Selecting Reel > Notes opens the Notes dialog to add annotations text to the entire reel project. The Reel > Export option saves the reel to disk as an ASCII readable format so it can be used elsewhere or archived.




The Reel submenu opens an area for production notes on the entire reel.

The View menu is used when you want to switch between the reel storyboard layout and a Timeline layout.

Guides

You can assign customizable guide overlays to three Guide buttons along the bottom of the Studio Player. Fusion includes four guides to choose from, but you can add your own using the XML Guide format and style information provided at the end of this chapter. You assign a customizable guide to one of the three Guide buttons by right-clicking over a button and selecting a guide from the list. To display the guide, click the assigned button.



Three toolbar Guide buttons enable guide overlays in the viewer.

Guides can be customized and placed in the the Guides folder:

- **On macOS:** Macintosh HD/Users/username/Library/Application Support/Blackmagic Design/Fusion/Guides/
- **On Windows:** C:\Users\username\AppData\Roaming\Blackmagic Design\Fusion\Guides.
- **On Linux:** home/username/fusion/BlackmagicDesign/Fusion/Guides.

Guides are a simple XML formatted text document saved with the .guide extension, as defined below. This makes it easy to create and share guides.

Custom Guide Format

The guides are files that have drawing instructions a bit like code, like this:

```
Guide
{
    Name = "10 Pixels",
    Elements =
    {
        HLine { Y1="10T" },
        HLine { Y1="10B" },
        VLine { X1="10L" },
        VLine { X1="10R" },
    },
}
```

Or an example of safe area:

```

Guide
{
    Name = "Safe Frame",

    Elements =
    {
        HLine { Y1="10%", Pattern = 0xF0F0 },
        HLine { Y1="90%", Pattern = 0xF0F0 },
        HLine { Y1="95%" },
            HLine { Y1="5%" },
        VLine { X1="10%", Pattern = 0xF0F0 },
        VLine { X1="90%", Pattern = 0xF0F0 },
        VLine { X1="95%" },
            VLine { X1="5%" },
            HLine { Y1="50%", Pattern = 0xF0F0, Color = { R =
1.0, G = 0.75, B = 0.05, A=1.0 } },
            VLine { X1="50%", Pattern = 0xF0F0, Color = { R =
1.0, G = 0.75, B = 0.05, A=1.0 } },
    },
}

```

Guide Styles

The style of a guide is defined by a set of properties that appear in the format shown below:

```
<HLine Y1="33%" Pattern="C0C0" Color="FFFFFFF"/>
```

- **HLine**: Draws a horizontal line and requires a Y-value, which is measured from the top of the screen. The Y-value can be given either in percent (%) or in absolute pixels (px).
- **Vline**: Draws a vertical line and requires an X-value, which is measured from the left of the screen. The X-value can be given either in percent (%) or in absolute pixels (px).
- **Pattern**: The Pattern value is made up of four hex values and determines the visual appearance of the line.

Examples for such patterns include:

```

>>FFFF draws a solid line _____
>>EEEE a dashed line  - - - - -
>>ECEC dash-dot line  -. . -. . -. . -. . -. . -. .
>>ECCC dash-dot-dot  - . . . . - . . . . - . . . .
>>AAAA dotted line  .....

```

- **Color**: The Color value is composed of four groups of two hex values each. The first three groups define the RGB colors; the last group defines the transparency. For instance, the hex value for pure red would be #FF000000, and pure lime green would be #00FF0000
- **Rectangle**: Draws a rectangle, which can be empty or filled, and supports the same pattern and color settings described above.

It requires two X- and two Y-values to define the extent <Rectangle Pattern="FOFO" X1="10%" Y1="10%" X2="90%" Y2="90%">.

- **FillMode:** Applies to rectangles only and defines whether the inside or the outside of the rectangle should be filled with a color. Leave this value out to have just a bounding rectangle without any fill.

```
>>FillMode = ("None"|"Inside"|"Outside")
```

- **FillColor:** Applies to rectangles only and defines the color of the filled area specified by FillMode.


```
>>FillColor="FF000020"
```

Connecting Bins Over a Network

You can share bins among computers running Fusion on the network so that multiple visual effects artists can share assets, presets, and even entire compositions. These shared bins are called remote bins, and everyone can share one or more remote bin in a studio.

To connect to a remote system and display its bins:

- 1 Choose Fusion Studio > Preferences.
- 2 In the Preferences dialog, select Global > Bins > Servers in the list.

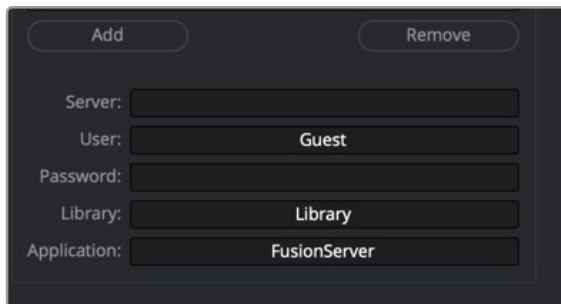


The bin servers Preferences panel

This panel shows a list of the available bin servers, with buttons below for entries to be added to or deleted from the list.

Adding a Remote Bin Entry

If you want to add another Remote bin to the list of available Remote bins, you can click the Add button in the bin servers Preferences panel. The text controls below the button will become enabled for editing. In the Server field, type the system name or IP address where the bin is hosted.



The image shows a dark-themed interface with two buttons at the top: 'Add' and 'Remove'. Below them are five text input fields with labels on the left and values on the right: 'Server:' (empty), 'User:' (Guest), 'Password:' (empty), 'Library:' (Library), and 'Application:' (FusionServer).

Add the IP address where the bin server is hosted.

Then add a User name and Password if one is needed to access the server.

The Library field lets you name the bins. So if you want to create a bin for individual projects, you would name it in the Library field, and each project would get its own bin.

The Application field allows larger studios to specify some other program to serve out the bin requests.

Once you've finished setting up the bin server information and clicked Save in the Preferences window, you can open the Bins window to test your bin server. Opening the Bins window is the first time your connection to the server will be tested. If it cannot connect, the bin server will still be listed, with access denied or unavailable marked next to the name on the bins sidebar.

There is no practical limit to the number of bins that can be accessed.

Accessing Remote Bins


Bin servers behave just like a local bins. Any bin added in the preferences show in the Bins sidebar as another top-level item. The available bins are shown by name with a status and, if required, a password. Bins unavailable to you are marked as (unavailable).

Permissions

Unlike other directories on a server, your access to bins on a network is stored in the bin document. The bins themselves contain all the users and passwords in plain text, making it easy for someone to administer the bins.


Studio Player and Bin Server

Reel projects can be shared by multiple artists across the studio via the bin server system, reviewing and adding versions and notes, all independently at the same time. With the Sync function, multiple people can collaborate together with synced playback and scrubbing.




The Sync button is a three-way toggle button: Off, Slave, and Master

When the Sync function is On, the transport controls can be set to control the playback or follow the master controller.



Master enables controlling other slave players on the network



The local transport controls can be disabled, causing playback to follow the master

Chapter 74

Fusion Connect

This chapter goes into detail on how to use the Fusion Connect AVX2 plug-in with an Avid Media Composer editing system. The Fusion Connect AVX plug-in is only available with Fusion Studio.

Contents

Fusion Connect Overview	1296	Create New Version	1299
System Requirements	1296	Version	1299
The Effects Palette	1296	About RAW Images	1299
The Layer Input Dialog	1297	About Color Depth	1300
Applying Fusion Connect to a Transition Point	1297	Manual vs. Auto-Render	1300
Export Clips	1297	Fusion/Avid Project Relationship	1302
Edit Effect	1298	Rendering with Fusion	1302
Browse for Location	1298	Directory Structure of Fusion Connect Media	1303
Auto Render in Fusion	1299	Advanced Project Paths	1304
Red on Missing Frames	1299	Configuring Paths on macOS	1304
Compress Exported Frames	1299	Configuring Paths on Windows	1305
Edit Effect Also Launches Fusion	1299	Fields and Variables	1305
Versioning	1299	Environment Variables	1305

Fusion Connect Overview

Fusion Connect is an AVX2 plug-in for Avid Media Composer. It allows editors to create a conduit between the Timeline in Avid editing products and Fusion Studio. Fusion Connect exports clips from the Avid Timeline as image sequences and assembles Fusion compositions that allow you to work your magic on the content.

Fusion can be started automatically by the plug-in if Fusion is installed on the same system, or it can be used on remote computers to modify the composition.


System Requirements

Fusion Connect has the following requirements:

- **Supported Avid products:** Media Composer 8.x
- **Supported product:** Fusion Studio 8.1 or later
- **Installation:** Two files will be installed in your Media Composer:
 - Fusion Connect.avx
 - BlackmagicFusionConnect.lua
- **Avid's default directory:** \Avid\AVX2_Plug-ins

The Effects Palette


After launching Media Composer, Fusion Connect is located in the Blackmagic Effects Palette category. As with any segment or transition effect in the Effects Palette, you can apply the Fusion Connect AVX2 plug-in to any clip. This includes filler*, edit transition point, or video track layers on the Avid Timeline.



Avid Media Composer Effects palette showing the Blackmagic category and Fusion Connect AVX plug-in.

The Layer Input Dialog

When you apply the Fusion Connect AVX plug-in to a clip or a layer in the Timeline, you are presented with an AVX Optional Inputs dialog box.




The Fusion Connect AVX plug-in can access multiple layers from the Media Composer Timeline.

Once the layer count is selected, Fusion Connect will be applied to the Timeline.

- Select the layer count equal to the number of video track layers you want to ingest into Fusion.
- Filler can be used as a layer.
- Fusion Connect will allow a maximum of eight layers.

Applying Fusion Connect to a Transition Point

If you apply Fusion Connect to a transition point, no dialog box will display, and the AVX plug-in will simply be applied to the Timeline transition point.



A Fusion Connect AVX plug-in applied as a transition to a cut in the Timeline

You can use the Avid dialog boxes or smart tools to adjust the length and offset of the transition to Start, Center, End, or Custom.


Export Clips

By pressing the Export Clips button in the Effects Editor, Fusion Connect exports all the associated clips as image sequences, to provide access to them in Fusion. Any previously exported existing images are overwritten, ensuring that all media needed by Fusion is accessible. Performing an export is desired if you want to use Fusion on a different computer than the one with Media Composer installed.

When using Fusion on the same computer as Media Composer, there is no need to export the clips explicitly by checking the Export Clips checkbox. Without this option enabled, Fusion Connect saves the source frames each time images are displayed, scrubbed, or played back from the Timeline. Depending on your Media Composer Timeline settings, these interactively exported images might be

half-resolution based on Avid proxy settings. When scrubbing around the Timeline, only a few frames—namely, those that are fully displayed during scrubbing—might be written to disk.

TIP: Set your Timeline Video Quality button to Full Quality (green) and 10-bit color bit depth. If the Timeline resolution is set to Draft Quality (green/yellow) or Best Performance (yellow), Fusion receives subsampled, lower-resolution images.



Fusion Connect AVX plug-in controls in the Effects Editor

Edit Effect

After exporting the clips, the Edit Effect button performs three subsequent functions:

- Creates a Fusion composition, with Loaders, Savers, and Merges (for layers), or a Dissolve (for transitions). This function is only performed the first time a comp is created when the Fusion Connect AVX2 plug-in is applied.
- Launches Fusion (if installed on the machine), if it is not already launched.
- Opens the Fusion comp associated with created effects.

Browse for Location

The Fusion Connect media folders are created on the drive where the associated Avid media resides, defaulting to the root level of that drive. However, you can choose a new location for the Fusion Connect media.

To choose another location to store and access the media:

- 1 In the Media Composer Effects Editor, click the Browse for Location button.
- 2 In the File Browser, change the location and create additional folders if desired.

The path settings field in the Effects Editor updates to show the current location. If you apply Fusion Connect to another clip in the Timeline, the last location is remembered.

Auto Render in Fusion

The Auto Render button is a toggle that allows you to automatically render your Fusion comp from within Avid. Please note that this method of rendering has limitations and can be significantly slower than rendering directly in Fusion. It is mostly used for batch rendering on the Avid Timeline. Auto Render also exports the necessary media without having to manually execute the Export Clips function first.

Red on Missing Frames

The Red on Missing Frames button is a toggle used to display red images within the Avid Timeline viewer (Timeline monitor) if no rendered frames from Fusion Studio can be found, or if the rendered frames are not of a high enough resolution. When disabled, the original untouched frames will be shown instead of the red frame.

Compress Exported Frames

When enabled, this button creates Fusion RAW files with compression for, both exported and rendered frames. This creates smaller file sizes, saving disk space. As with any other compression algorithm, the compression adds time to the write process of the file sequence.

Edit Effect Also Launches Fusion

When enabled, this button opens Fusion Studio when the Edit Effect button is clicked, allowing for a more direct Avid/Fusion workflow. When disabled, this button will not launch Fusion Studio when the Edit Effect button is clicked, but a Fusion .comp file is created. This is useful when working with Fusion Studio on a separate computer than the one running the Media Composer software.

Versioning

Creating visual effects is almost always an iterative process. You'll often need to create revisions after your first pass at the effect. Built into Fusion Connect is a versioning feature that lets you create multiple revisions of an effect and switch between them from within Media Composer.

Create New Version

This checkbox creates a copy of the current comp without affecting the original.

Any changes that are rendered in the copy will be written to a new folder and become another version of the rendered result played on the Avid Timeline. Previous versions of the comp and their rendered results are accessible using the Version slider.

Version

This slider selects which version of the comp is used in the Media Composer Timeline. It can be used to interactively switch from one version to the other in order to compare the results.

About RAW Images

Fusion Connect creates a Fusion RAW file image sequence for intermediate folders between Avid Media Composer and Fusion Studio, in order to preserve all the image information. This allows the images to reside on disk and not take up space in RAM. The benefits of using Fusion RAW include:

- The ability to continue the editing process while an effect is rendering
- The ability to take advantage of network rendering
- The ability to retime footage
- The ability to run Fusion Studio remotely

About Color Depth

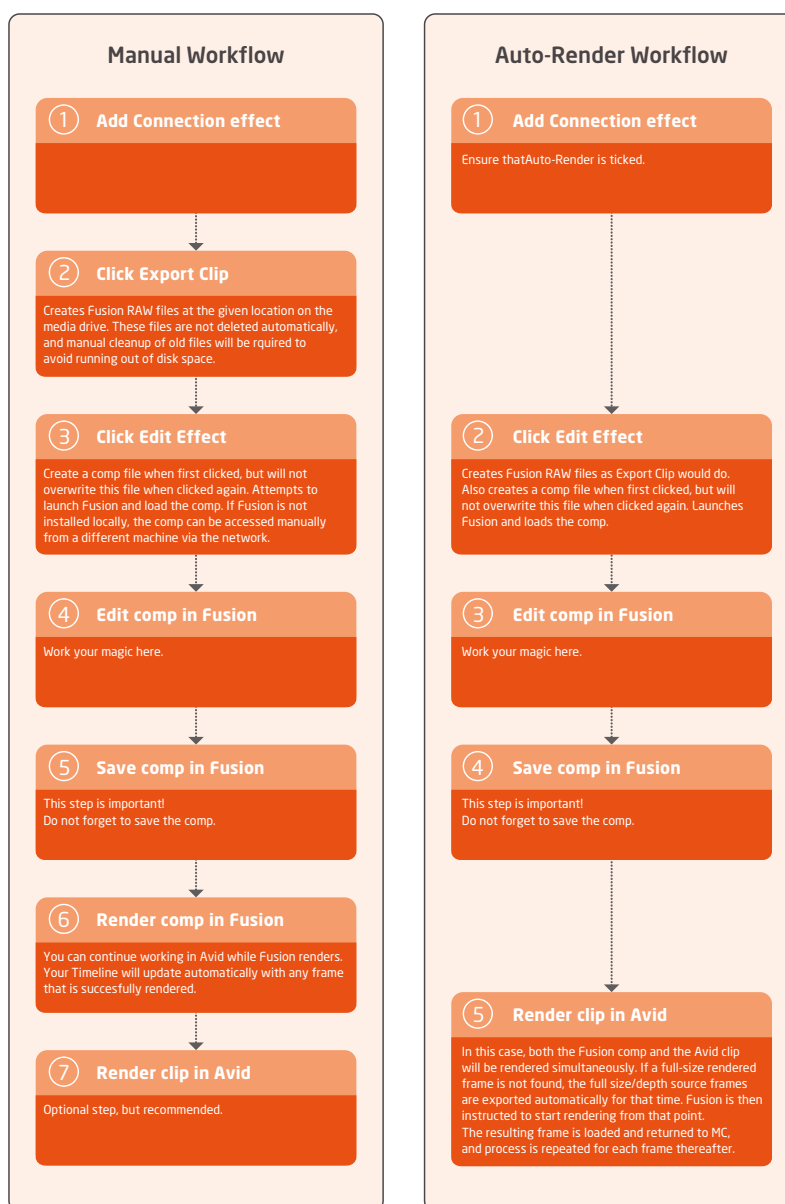
Fusion Connect derives its images directly from the RGB data within Avid Media files. This allows the images to be codec agnostic. All RAW files from Avid that begin as 8- or 10-bit images are remapped to 16-bit float in Fusion. Rendered results from Fusion are processed in 16-bit float to maintain the full color fidelity supported by Media Composer.

Manual vs. Auto-Render

While Auto-Render is the easier workflow, the manual approach offers faster renderings in Fusion and more control over the performance and memory usage on your system.

- In the manual workflow, Fusion Studio is not required to be installed on the Avid system itself but can reside on any other computer.
- For Auto-Render, Fusion Studio must be installed on the local computer.

The following diagram shows typical workflows for manual and automatic renders.




Avid/Fusion layer to comp relationship for auto and manual renders

Once the initial trip from Avid to Fusion is complete, depending on the type of clip to which you assigned the effect, in Fusion Studio you will be presented with one of the following:

- A Loader node representing a single clip.
- Two or more Loaders connected to Merge nodes representing layers.
- Two Loaders connected to a Dissolve node representing a transition.

In all three node tree layouts outlined above, there will also be a Saver node. The Saver node is automatically set to render to the directory that is connected back to the Media Composer Timeline with the correct format. If for some reason the file format or file path are changed in the Saver node, the Fusion Connect process will not render correctly.

TIP: Due to the design of the AVX2 standard, hidden embedded handles are not supported. To add handles, prior to exporting to Fusion, increase the length of your clip in the Media Composer Timeline to include the handle length.

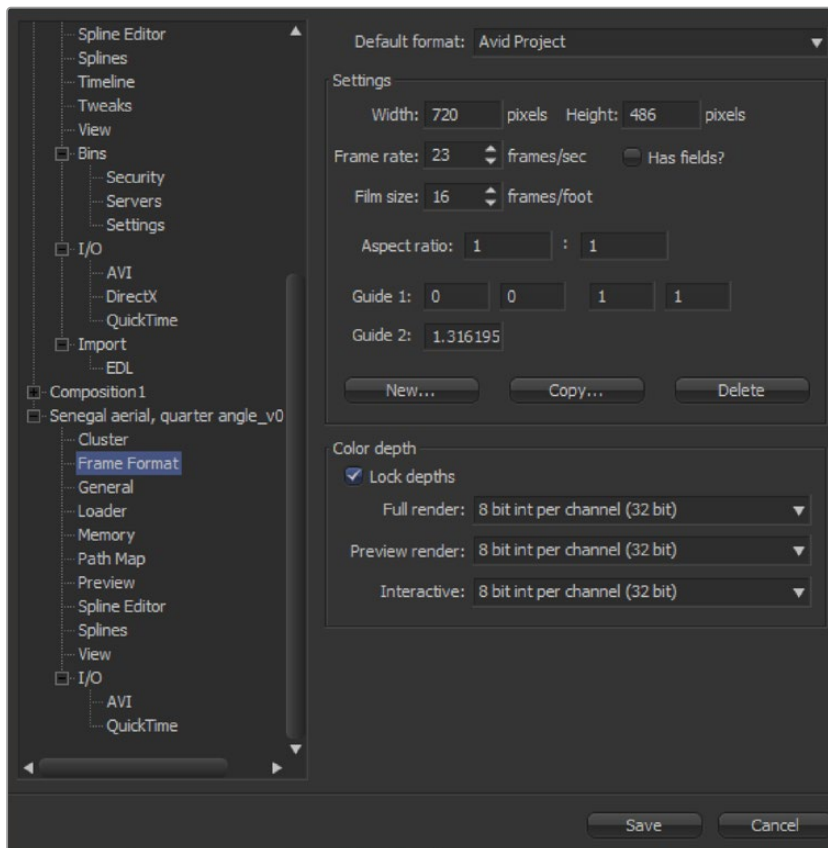


Fusion Node Editor representations of a single clip segment effect in the Media Composer Timeline (left), a multi-layered composite (center), and the transition (right)

TIP: If segments in the Avid Timeline have different durations, move the longest clip to the top layer and apply Fusion Connect to that clip. This will ensure all clips are brought into Fusion without being trimmed to the shortest clip.

Fusion/Avid Project Relationship

The frame rate and image size preferences created in Media Composer are adopted within Fusion's frame rate preferences. This allows for consistency in formats for the roundtrip process from Avid to Fusion and back to Avid. The format settings do not prevent you from using or mixing any other sized imaging within the composition as Fusion is resolution independent.



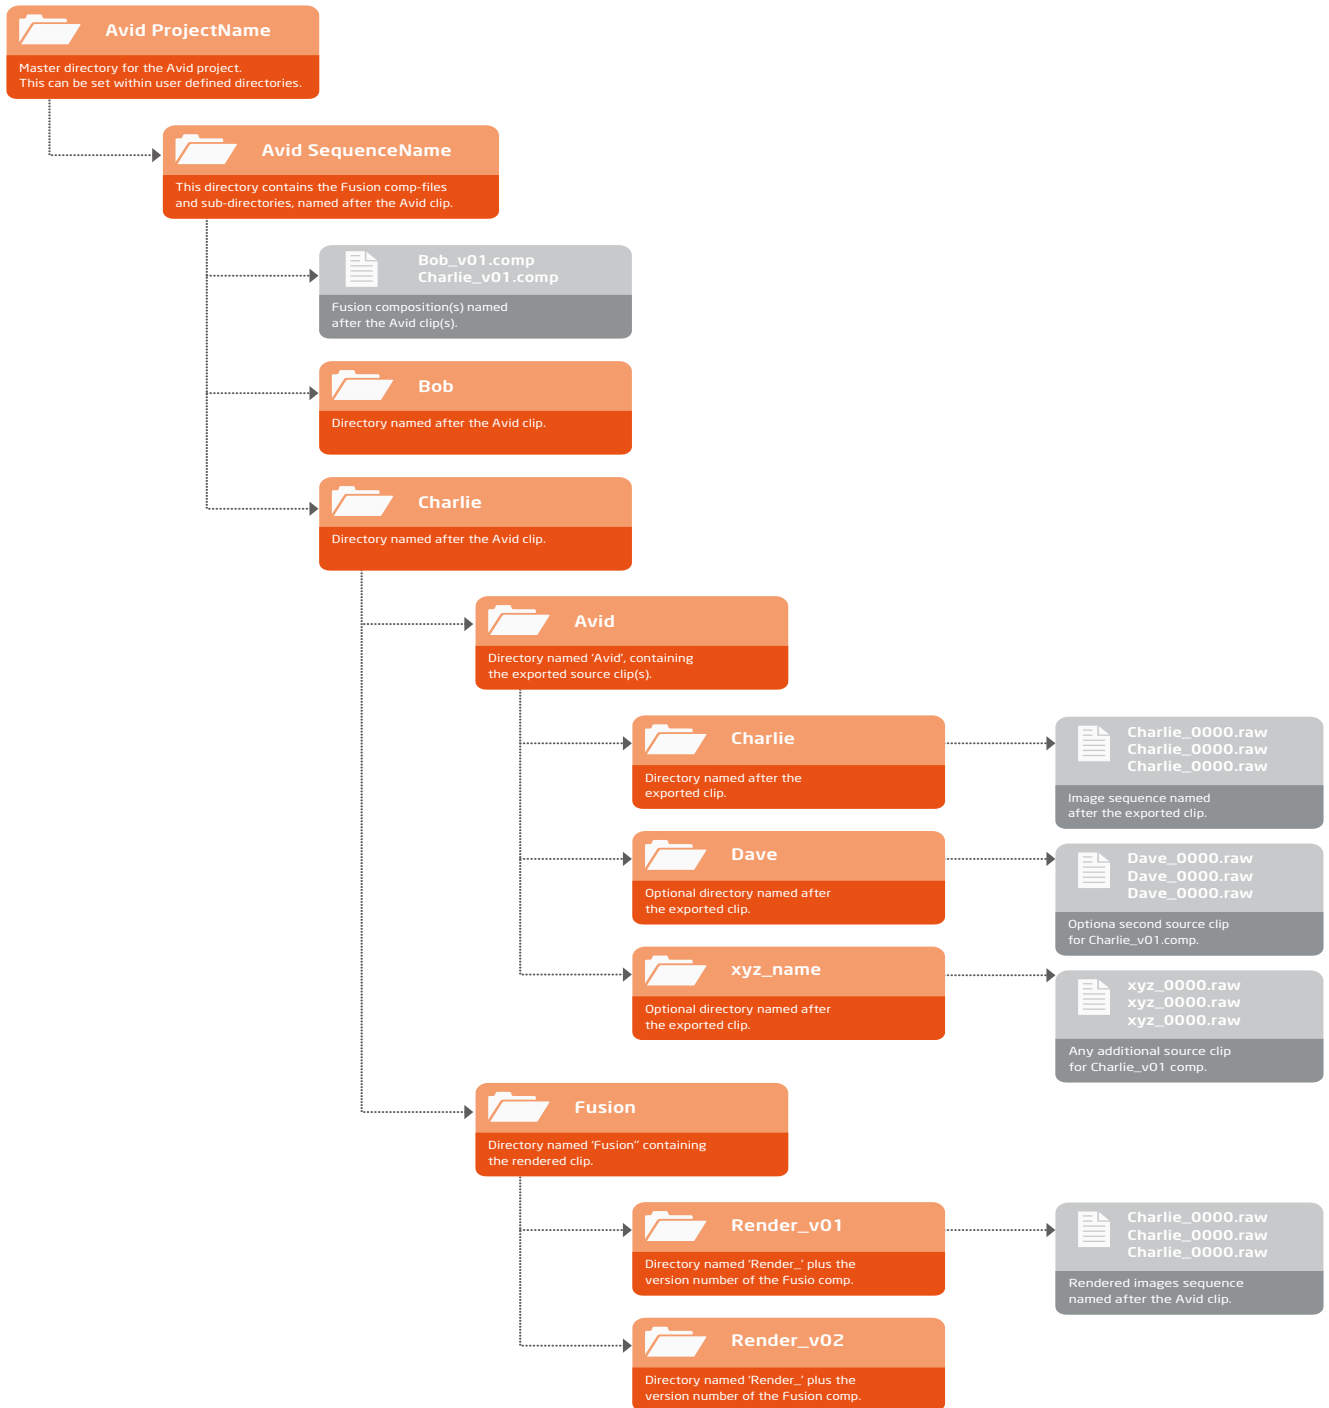
Fusion Connect AVX uses frame rate and resolution from the Media Composer Timeline.

Rendering with Fusion

When you perform a render of your comp inside Fusion Studio, the results are rendered to the output folder created by Fusion Connect during the initial application of the plug-in to the Timeline. Upon rendering, you immediately see the results of the rendered Fusion comp in the Avid Timeline. Even while Fusion is rendering, you can continue with the edit process on any clip except for the associated clip being rendered at the time.

Directory Structure of Fusion Connect Media

Fusion Connect creates a logical folder structure that is not affiliated with the Avid Media Files folder but rather the Fusion Connect AVX2 plug-in. Based on data gathered during the AVX application to the Timeline, a logical folder hierarchy is automatically created based on Avid projects, sequences, and clips. This structure allows for multiple instances of Fusion Studio to access the media and multiple instances of the AVX to relate to a single Fusion comp. In a re-edit situation, media is never duplicated but is overwritten to avoid multiple copies of identical media.



If you apply the effect to a transition, the naming behavior might be somewhat different.

By default, Media Composer refers to the two clips of a transition as “Clip_001” and “Clip_002”. Based on the naming convention, Fusion Connect will create folders with matching names. If such folders already exist, because another transition has already used up “Clip_001” and “Clip_002”, the numbers will be incremented automatically.

Likewise, “_001” will be added incrementally to the group folder name, if a folder of that name already exists. The corresponding comp file will be named accordingly.



Fusion Connect AVX creates folder structures in the OS to save media and Fusion compositions. Those names are reflected in the Timeline.

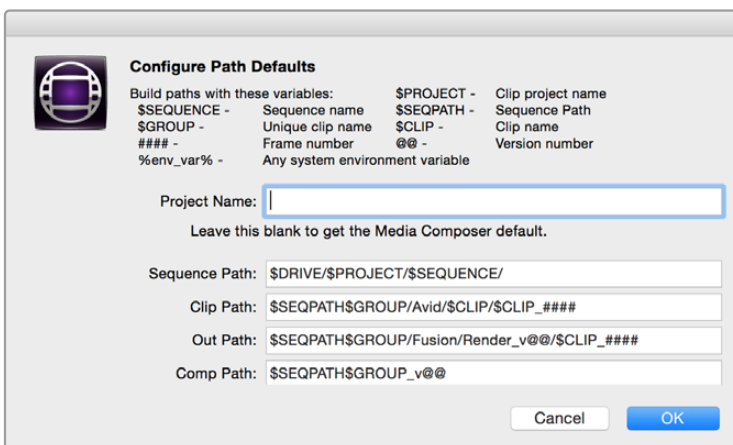
You will notice that the Fusion Connect icon is a green dot (real-time) effect. If your hardware is fast enough, the results that populate the plug-in will play in real time. However, it's recommended that you render the green dot effect, which will force an MXF precompute to be created to guarantee real-time playback.

Advanced Project Paths

The Fusion Connect AVX2 plug-in controls the pathing of Fusion Connect's .raw media and Fusion's .comp files as well as showing and hiding project and sequence level directories. This is achieved through environment variables, which are set in the operating system. This gives you the most flexible control over pathing your media, and as the name depicts, you can change variables (controls) in various application environments. This is useful for network storage and when Fusion Studio is running on other systems.

Configuring Paths on macOS

When using Fusion Connect on macOS, the Configure Path Defaults dialog looks like this:




Environment variable path maps on macOS.

Default paths can be configured using variables similarly as on Windows, but for added convenience it is possible to enter any desired path defaults directly into fields in the dialog, without the need for using environment variables.

Configuring Paths on Windows

When using Fusion Connect on Windows, the Configure Path Defaults dialog looks like this:



Environment variable path maps on Windows.

Fusion Connect can define the user variables directly in the Fusion Connect plug-in. Click the Configure Path Defaults button to launch the path defaults dialog editor. In the Options section of the Fusion Connect AVX2 plug-in, click the triangle to reveal the path details.

Fields and Variables

Variables are the control title defined specifically by the application that is being controlled, while values are the instructions that tell a variable what to do. The fields and variables that can be used on macOS and Windows are described in the following table:

Field	Variable	Environment Variable	Description
Project Name	\$PROJECT	CONNECT_PROJECT	Overrides the current Avid Project name.
	\$DRIVE	CONNECT_DRIVE	Drive or folder for all Connect projects
	\$SEQUENCE	–	Name of Avid sequence
Sequence Path	\$SEQPATH	CONNECT_SEQUENCE_PATH	Folder for all Connect files in this sequence
	\$GROUP	–	Unique name of this Connect instance
	\$CLIP	–	Name of exported clip
Clip Path	–	CONNECT_CLIP_PATH	Folder for exported clips from Avid
Out Path	–	CONNECT_OUT_PATH	Folder for rendered results from Fusion
Comp Path	–	CONNECT_COMP_PATH	Location and name of Fusion comp file

Environment Variables

The pathing can be set in the environment variables of the system, so that IT management of the project paths can be achieved.

Accessing the Environment Variables on Windows

The quickest way to access environment variables is through your Windows control panel by searching the word “env” without the quotes. You have a choice of editing at a user level or system level.

User Variables

Click on the link that says “Edit environment variables for your account.”

System Variables

Click on the link that says “Edit the system environment variables.”

Accessing the Environment Variables on macOS

To set an environment variable on macOS, you must use the Terminal window. Environment variables on macOS are added to the `.bash_profile` directory for the current user.

User Variables

For system-wide operations, place the environment variable in `~/bash_profile`

TIP: System variables control the environment throughout the operating system, no matter which user is logged in.

User variables always overrule any system variable; therefore, the user variable always wins if control for a specific function is duplicated in the user and system variable.

System Variables

For system-wide operations, place the environment variable in `/etc/profile`

Environment Variables and What They Mean


The Fusion Connect AVX plug-in can use environment variables to set different locations for certain folders or files. Each computer OS has a unique way of entering environment variables, but in every OS, the variable must be typed exactly as shown.

TIP: If you type directly in Fusion Connect’s Path Editor, you do not have to type the variable, just the value. You also can make modifications without having to restart the Media Composer! The only caveat is that in order to remove a variable, you must exit Media Composer and clear the environment variable in the Windows interface or macOS Terminal and restart the Media Composer.

Other values you can control derived from your Avid Bin include:

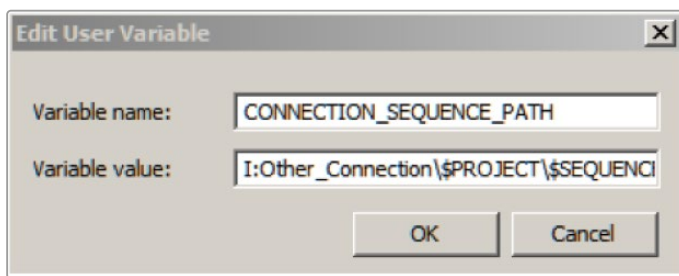
Values	Description
<code>\$DRIVE</code>	This will force the directory to the drive where the Avid media is stored.
<code>\$PROJECT</code>	This will force a directory based on the Avid project name for which the media was digitized/imported or AMA linked.
<code>\$SEQUENCE</code>	This will force a directory based on the Avid SEQUENCE name for which the media was digitized/imported or AMA linked.

Here is an example of how a variable can be set up to support project and sequence names within your directory.



Environment variable project and sequence name examples on Windows

Here is the same example in the Windows environment variable editor.



Windows environment user variable editor

Chapter 75

Preferences

This chapter covers the various options that are available from the Fusion Preferences Window.

Contents

Preferences Overview	1309	Script	1333
Categories of Preferences	1310	Spline Editor	1334
Preferences In Depth	1312	Splines	1335
AVI	1314	Timeline	1336
Defaults	1314	Tweaks	1337
Flow	1315	User Interface	1340
Frame Format	1317	Video Monitoring	1341
General	1318	View	1342
GPU	1320	VR Headsets	1343
Layout	1322	Bins/Security	1345
Loader	1323	Bins/Server	1346
Memory	1324	Bins/Settings	1347
Network	1326	EDL Import	1348
Path Maps	1327	Customization	1348
Preview	1331	Shortcuts Customization	1349
QuickTime	1332	Customizing Preferences	1350

Preferences Overview


The Preferences window provides a wide variety of optional settings available for you to configure Fusion's behavior to better suit your working environment. These settings are accessed via the Preferences window. The Preferences window can be opened from a menu at the top of the interface.

In DaVinci Resolve, to open the Fusion Preferences window, do one of the following:

- On macOS, switch to the Fusion page and choose Fusion > Fusion Settings.
- On Windows, switch to the Fusion page and choose Fusion > Fusion Settings.
- On Linux, switch to the Fusion page and choose Fusion > Fusion Settings.

In Fusion Studio, to open the Fusion Preferences window, do one of the following:

- On macOS, choose Fusion Studio > Preferences.
- On Windows, choose File > Preferences.
- On Linux, choose File > Preferences.



The Preferences window includes composition settings


Global and Composition Preferences

The Preferences window is divided into a category sidebar on the left and the settings panel on the right. In Fusion Studio, there are two levels of preferences: Global and Composition. The Fusion page in DaVinci Resolve uses just a single Global preference that affects every project, new and existing.

The Global preferences are used to set options specific to Fusion's overall behavior as well as defaults for each new composition. The Composition preferences in Fusion Studio can further modify the currently open composition without affecting the Global preferences or any other composition that is open but not displayed.

Categories of Preferences

The first entry in the Preferences sidebar is assigned to the Global preferences. Clicking the Global and Default Settings disclosure arrow reveals the following sections.



Fusion Studio Global and Default settings

3D View

The 3D View preferences offer control over various parameters of the 3D Viewers, including grids, default ambient light setup, and Stereoscopic views.

AVI (Windows Fusion Studio Only)

The AVI preferences configure the default compression options when creating AVI files from a Saver node. These settings can be overridden on a case-by-case basis using the Format tab in the Saver's Inspector.

Defaults

The Defaults preferences are used to select default behavior for a variety of options, such as animation, global range, timecode display, and automatic tool merging.

Flow

You use the Flow preferences to set many of the same options found in the Node Editor's contextual menu, like settings for Tile picture, the Navigator, and pipe style.

Frame Format

The Frame Format preferences are used to create new frame formats as well as select the default image height and width when adding new creator tools like Background and Text+. You also set the frame rate for playback.

General

The General preferences contain options for the general operation, such as auto save, and gamma settings for color controls.

GPU (Fusion Studio Only)

The GPU preferences include options for selecting specific GPU acceleration methods based on your computer platform and hardware capabilities. It is also used for enabling caching, and debugging GPU devices and tools.

Layout (Fusion Studio Only)

You can use the Layout preferences to save the exact layout of Fusion's windows.

Loader (Fusion Studio Only)

Using the Loader preferences, you can set options for the default Loader's depth and aspect ratio as well as define the local and network LoaderCache settings.

Memory (Fusion Studio Only)

Memory management for multi-frame and simultaneous branch rendering is configured in the Memory preferences.

Network (Fusion Studio Only)

The Network rendering preferences are used to configure options such as selecting a render master, email notification, and whether the machine can be used as a render slave.

Path Map

Path Map preferences are used to configure virtual file path names used by Loaders and Savers as well as the folders used by Fusion to locate comps, macros, scripts, tool settings, disk caches, and more.

Preview (Fusion Studio Only)

The Preview preferences is where you configure the Preview creation and playback options.

QuickTime (macOS Fusion Studio Only)

This section lets you preconfigure the QuickTime codec used for rendering.

Script

The Script preferences include a field for passwords used to execute scripts externally, programs to use for editing scripts, and the default Python version to use.

Spline Editor

The Spline Editor preferences allow you to set various spline options for Autosnap behavior, handles, markers, and more.

Splines

Options for the handling and smoothing of animation splines, Tracker path defaults, onion-skinning, roto assist, and more are found in the Splines preference.

Timeline

The Timeline preferences is where you create and edit Timeline/Spline filters and set default options for the Keyframes Editor.

Tweaks

The Tweaks preferences handle miscellaneous settings for modifying the behavior when loading frames over the network and queue/network rendering.

User Interface

These preferences set the appearance of the user interface window and how the Inspector is displayed.

Video Monitoring (Fusion Studio Only)

The Video Monitoring preferences is where you can configure your Blackmagic video display hardware for monitoring on an HD, Ultra HD, or DCI 4K display.

View

The View preferences are used to manage settings for viewers, including default control colors, Z-depth channel viewing ranges, default LUTs, padding for fit, zoom, and more.

VR Headsets

The VR Headsets preferences allow configuration of any connected Virtual Reality headsets, including how stereo and 3D scenes are viewed.

Bins (Fusion Studio Only)

There are three panels as part of the Bins preferences: a Security panel where you set users and passwords for serving the local bins; a Servers panel used to select which remote Bin servers are connected; and a Settings panel for stamp rendering.

Import

The Import settings contain options for EDL Import that affect how flows are built using the data from an EDL.

Preferences In Depth

Within each category is a deep set of controls for configuring Fusion so that it better fits your working environment. The preferences contain both software and hardware options that affect all newly created comps. The following section explains every option located in the preferences categories.

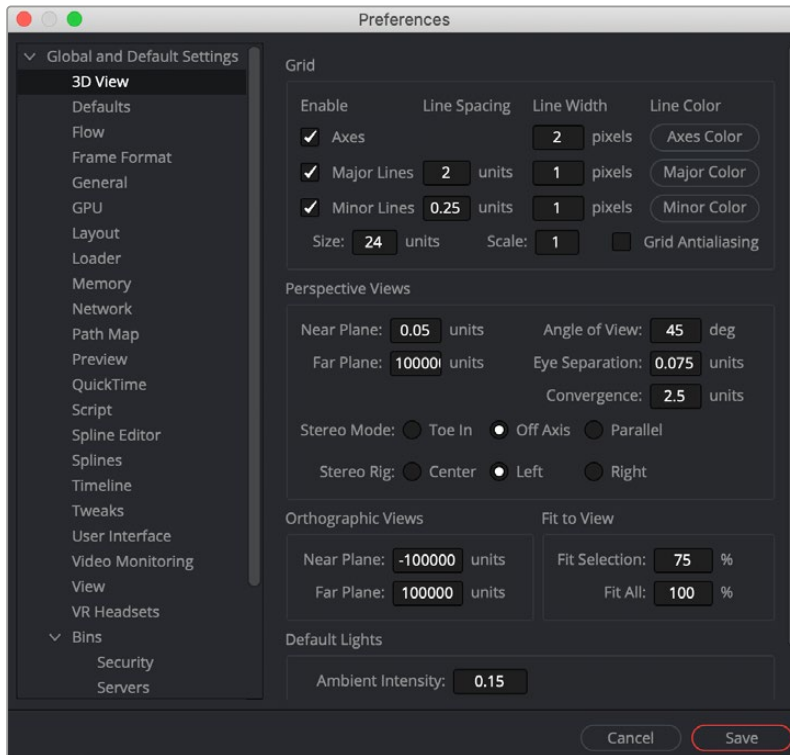
3D View

The 3D View preferences contain settings for various defaults in the 3D Viewers, including grids, default ambient light setup, and Stereoscopic views.

Grid

The Grid section of the 3D View preferences configures how the grid in 3D Viewers are drawn.

- **Grid Antialiasing:** Some graphics hardware and drivers do not support antialiased grid lines, causing them to sort incorrectly in the 3D Viewer. Disabling this checkbox will disable antialiasing of the grid lines. To turn off the grid completely, right-click in a 3D Viewer and choose 3D Options > Grid.
- **Size:** Increasing the Size value will increase the number of grid lines drawn. The units used for the spacing between grid lines are not defined in Fusion. A “unit” is whatever you want it to be.
- **Scale:** Adjusting the overall scaling factor for the grid is useful, for example, if the area of the grid appears too small compared to the size of your geometry.



The 3D View preferences

Perspective Views

The Perspective Views section handles the appearance of the perspective view in both a normal and stereoscopic project.

- **Near Plane/Far Plane:** These values set the nearest and furthest point any object can get to or from the camera before it is clipped. The minimum setting is 0.05. Setting Near Plane too low and Far Plane too far results in loss of depth precision in the viewer.
- **Eye Separation/Convergence/Stereo Mode:** This group of settings defines the defaults when stereo is turned on in the 3D Viewer.

Orthographic Views

Similar to the Perspective Views, the Orthographic Views (front, top, right, and left views) section sets the nearest and furthest point any object can get to or from the viewer before clipping occurs.

Fit to View

The Fit to View section has two value fields that manage how much empty space is left around objects in the viewer when the F key is pressed.

- **Fit Selection:** Fit Selection determines the empty space when one or more objects are selected and the F key is pressed.
- **Fit All:** Fit All determines the empty space when you press F with no objects selected.

Default Lights

These three settings control the default light setup in the 3D Viewer.

The default ambient light is used when lighting is turned on and you have not added a light to the scene. The directional light moves with the camera, so if the directional light is set to “upper left,” the light appears to come from the upper-left side of the image/camera.


AVI

The AVI preference is only available in Fusion Studio on Windows. It configures the default AVI codec settings when you select AVI as the rendering file format in the Saver node.

- **Compressor:** This drop-down menu displays the AVI codecs available from your computer. Fusion tests each codec when the application opens; therefore, some codecs may not be available if the tests indicate that they are unsuitable for use within Fusion.
- **Quality:** This slider determines the amount of compression to be used by the codec. Higher values produce clearer images but larger files. Not all codecs support the Quality setting.
- **Key Frame Every X Frames:** When checked, the codec creates keyframes at specified intervals. Keyframes are not compressed in conjunction with previous frames and are, therefore, quicker to seek within the resulting movie. Not all codecs support the keyframe setting.
- **Limit Data Rate To X KB/Second:** When checked, the data rates of the rendered file are limited to the amount specified. Not all codecs support this option. Enter the data rate used to limit the AVI in kilobytes (kB) per second, if applicable. This control does not affect the file unless the Limit Data Rate To option is selected.

Defaults

The choices made here are used to determine Fusion's behavior when new tools are added to the Node Editor and when parameters are animated.



The Defaults preferences

Default Animate

The Default Animate section is used to change the type of modifier attached to a parameter when the Animate option is selected from its contextual menu. The default option is Nothing, which uses a Bézier spline to animate numeric parameters and a path modifier for positional controls.

- **Number With and Point With:** Drop-down lists are used to select a different modifier for the new default. For example, change the default type used to animate position by setting the Point with the drop-down menu to XY Path.

Choices shown in this menu come from installed modifiers that are valid for that type of parameter. These include third-party plug-in modifiers, as well as native modifiers installed with Fusion.

Auto Tools

The Auto Tools section determines which tools are added automatically for the most common operations of the Background tools and Merge operations.

- **Background:** When set to None, a standard Background tool is used; however, the drop-down menu allows you to choose from a variety of tools including 2D and 3D tools to customize the operation to your workflow.
- **Merge:** When set to None, nothing happens. When set to Merge, connecting the outputs of two tools or dragging multiple clips on the Node Editor uses a standard Merge. Other valid options for this are Anaglyph, Channel Booleans, and Dissolve.
- **Use Merge Only When Connecting Outputs Directly:** When this option is active, Merges are not automatically added when you drag multiple clips from the Finder or Windows Explorer onto the Flow area.

Global Range


Using the Start and End fields, you can define the Global Start and End frames used when creating new compositions.

Time Code

You use this option to determine whether new compositions will default to showing SMPTE Time Code or frames (Feet + Frames) to represent time.

Flow

Many of the same options found in the Node Editor's contextual menu, like settings for Tile Picture, the Navigator, and Pipe Style, are found in this category.



The Flow preferences

Force

The Force section can set the default to display pictures in certain tool tiles in the Node Editor rather than showing plane tiles. The Active checkbox sets pictures for the actively selected tool, the All checkbox enables pictures for all tiles, and the Source and Mask checkbox enables tile pictures for just Source and Mask tools.

When All is enabled, the picture shown will either be a thumbnail of the image rendered by the tool if the tool has rendered, or if the Show Thumbnails option is disabled, the tool's default icon is used. Concatenated transforms will also show a default icon.

- **Show Modes/Options:** Enabling this option will display icons in the tool tile depicting various states, like Disk Caching or Locked.
- **Show Thumbnails:** When this checkbox is selected, tool tiles set to show tile pictures will display the rendered output of the tool. When the checkbox is cleared, the default icon for the tool is used instead.

Options

The Options section includes several settings that control or aid in the layout and alignment of tools in the Node Editor.

- **Arrange to Grid:** This enables a new node tree's Snap to Grid option to force the tool layout to align with the grid marks in the flow.
- **Arrange to Connected:** Tools snap to the vertical or horizontal positions of other tools they are connected to.
- **Auto Arrange:** This option enables the Node Editor to shift the position of tools as needed to make space when inserting new tools or auto-merging layers.
- **Show Grid:** This enables or disables the display of the Node Editor's background grid.
- **Auto Remove Routers:** Pipe Routers or "elbow nodes" in the Node Editor are considered "orphaned" if the tools connected to either the input or output are deleted. When this option is enabled, Orphaned Routers are automatically deleted.
- **Pipes Always Visible:** When enabled, the connection lines between tools are drawn over the top of the tool tiles.
- **Keep Tile Picture Aspect:** Enabling this option forces tool tile thumbnail pictures to preserve the aspect of the original image in the thumbnail.
- **Full Tile Render Indicators:** Enabling this checkbox causes the entire tile to change color when it is processing. This can make it easier to identify which tools are processing in a large composition. The coloring itself will form a progress bar to alert you to how close slower tools are to finishing their process.
- **Show Instance Links:** This option is used to select whether Instance tools will show links, displayed as green lines, between Instance tools.
- **Navigator:** The Navigator is a small square overview of the entire composition. It is used to quickly navigate to different parts of a node tree while you are zoomed in. The checkboxes in this section determine when the Navigator is displayed, if at all.
 - **On:** The Navigator will always be visible.
 - **Off:** The Navigator will always be hidden.
 - **Auto:** The Navigator will only be visible when the Node Editor's contents exceed the currently visible Work area.
- **Pipe Style:** This drop-down menu selects which method is used to draw connections between tools. The Direct method uses a straight line between tools, and Orthogonal uses horizontal and vertical lines.

- **Build Direction:** When auto-building or laying out a node tree, Build Direction controls whether tools are organized horizontally or vertically.
- **Scale:** The Scale menu allows you to select the default zoom level of the Node Editor when a new composition is created.

Pipe Grab Distance

The Pipe Grab Distance slider allows you to choose how close the pointer must be (in pixels) to a connection line in the node tree when selecting them.

Link Grab Distance

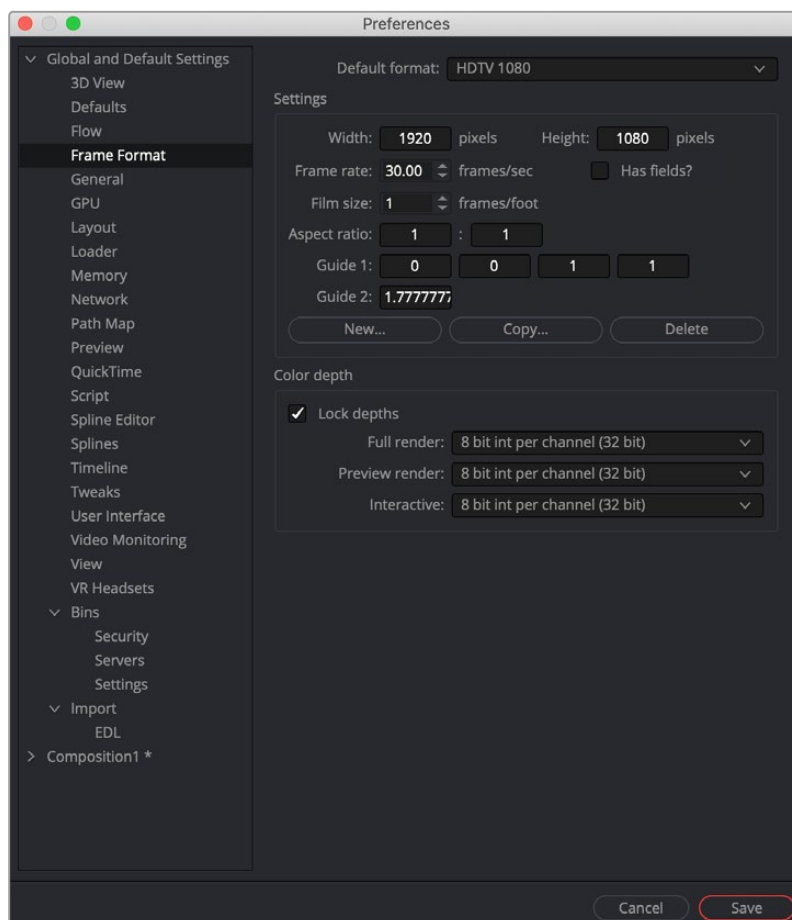
The Link Grab Distance slider allows you to choose how close the pointer must be (in pixels) to a knot on a node before a connection is made or removed. 0 pixels means you must be directly on the knot, while a maximum setting of 20 means you can be 20 pixels from a knot and still connect to it.

Group Opacity

This slider controls the opacity of an expanded group's background in the Node Editor.

Frame Format

Frame Format preferences allow you to select the resolution and frame rate for the nodes that generate images like Background, fast noise, and Text+. It also sets the color bit depth for final renders, previews, and interactive updates in the viewer. The color bit depth settings only apply to Fusion Studio. Rendering in DaVinci Resolve always use 32-bit float.



The Frame Format preferences

Default Format

This drop-down menu is used to select the default resolution for Generator tools from a list of presets. This is only a default setting; these settings can be overridden using the Resolution settings in a node's Inspector.

Use the Edit boxes to change any of the default settings. When creating a new setting, press the New button and enter a name for the setting in the dialog box that appears and enter the parameters.

Settings

The Settings section defines the format that is selected in the Default Format menu. You can modify an existing format or create a new one.

- **Width/Height:** When creating a new format for the menu or modifying an existing menu item, you specify the Width or Height in pixels of the format using these fields.
- **Frame Rate:** Enter or view the frames per second played by the format. This sets the default Frame Rate for previews and final renders from the Saver tool. It also sets the playback for the comp itself, as well as the frame to time code conversion for tools with temporal inputs.
- **Has Fields:** When this checkbox is enabled, any Creator or Loader tool added to the Node Editor will be in Fields process mode.
- **Film Size:** This field is used to define how many frames are found in one foot of film. The value is used to calculate the display of time code in Feet + Frames mode.
- **Aspect Ratio:** These two fields set the pixel aspect ratio of the chosen frame format.
- **Guide 1:** The four fields for Guide 1 define the left, top, right, and bottom guide positions for the custom guides in the viewer. To change the position of a guide, enter a value from 0 to 1. The bottom-left corner is always 0/0, the top-right corner is always 1/1. If the entered value's aspect does not conform to the frame format as defined by the Width and Height parameters, an additional guide is displayed onscreen. The dotted line represents the image aspect centered about Guide 1's Center values.
- **Guide 2:** This setting determines the image aspect ratio in respect to the entire frame format width and height. Values higher than 1 cause the height to decrease relative to the width. Values smaller than 1 cause height to increase relative to width.
- **New:** You use the New button to create a new default setting in the drop-down menu. Once you click the button, you can name the setting in the dialog box that appears.
- **Copy:** The Copy button copies the current setting to create a new one for customization.
- **Delete:** The Delete button will remove the current setting from the default drop-down list.


Color Depth

The three menus in the Color Depth section are used to select the color mode for processing preview renders, interactive renders, and full (final) renders. Processing images at 8-bit is the lowest color depth and is rarely sufficient for final work these days but is acceptable for fast previews. 16-bit color has much higher color fidelity but uses more system resources. 16-bit and 32-bit float per channel uses even more system resources and is best for digital film and HDR rendered images.

Generally, these options are ignored by the composition unless a Loader or Creator tool's Color Depth control is set to Default.

General

The sections contained in the General preferences affect the behavior of the Inspector as well as some other user interface elements.



The General preferences

Usability

Usability has a number of project, Node Editor, and user interface settings that can make the application easier to work with, depending on your workflow.

- **Auto Clip Browse:** When this checkbox is enabled, the File Browser is automatically displayed when a new Loader or Saver is added to the Node Editor.
- **New Comp on Startup:** When checked, a new, empty project is created each time Fusion Studio is launched. This has no effect in DaVinci Resolve’s Fusion page.
- **Summarize Load Errors:** When loading node trees or “comps” that contain unknown tools (e.g., comps that have been created on other computers with plug-ins not installed on the current machine), the missing tools are summarized in the console rather than a dialog being presented for every missing tool.
- **Save Compressed Comps:** This option enables the saving of compressed node trees, rather than ASCII based text files. Compressed node trees take up less space on disk, although they may take a moment longer to load. Node trees containing complex spline animation and many paint strokes can grow into tens of megabytes when this option is disabled. However, compressed comps cannot be edited with a text editor unless saved again as uncompressed.
- **Show Video I/O Splash:** This toggles whether the Splash image will be displayed over the video display hardware. This is only applies to Fusion Studio.
- **Use Simplified Copy Names:** This option reduces the occurrence of underscores in tool names when copying.

- **Show Render Settings:** When this checkbox is selected, the Fusion Render Settings dialog will be displayed every time a render is started in Fusion Studio. Holding Shift while starting a render will prevent the display of the dialog for that session, using whatever settings were applied during the last render. Disabling this option reverses this behavior.
- **Mouse Wheel Affects the Window Under the Pointer:** Normally the mouse wheel or trackpad swiping works in the currently active window. With this option enabled, it will work in the window underneath the cursor, so you don't have to click into a window first to make it active.
- **Frames Start From:** This designates the starting frame number for clip times in the Loader and its Clip list.
- **Show Color As:** This setting determines the numeric scale used to represent colors. The available options are Normalized (0 to 1), 8-bit (0 to 255), and 16-bit (0 to 65,535). This does not affect the actual processing or quality of the image, but it can make the mental math sometimes used to figure out adjustments a bit easier.

Auto Save

The Auto Save settings only apply to Fusion Studio. To set auto backups for the Fusion page in DaVinci Resolve, use the DaVinci Resolve Project Load and Save Preferences.

When Auto Save is enabled in Fusion Studio, comps are automatically saved to a backup file at regular intervals defined by the Delay setting. If a backup file is found when attempting to open the comp, you are presented with the choice of loading either the backup or the original.

If the backup comp is opened from the location set in the Path Map preference, saving the backup will overwrite the original file. If the backup file is closed without saving, it is deleted without affecting the original file.

- **Save Before Render:** When enabled, the comp is automatically saved before a preview or final render is started.
- **Delay:** This preference is used to set the interval between Auto Saves. The interval is set using mm:ss notation, so entering 10 causes an Auto Save to occur every 10 seconds, whereas entering 10:00 causes an Auto Save every 10 minutes.

Proxy

- **Update All, Selective, No Update:** The Update mode button is located above the toolbar. You can use this preference to determine the default mode for all new comps. Selective is the usual default. It renders only the tools needed to display the images in the Display view. All will render all tools in the composition, whereas None prevents all rendering.
- **Standard and Auto:** These sliders designate the default ratio used to create proxies when the Proxy and Auto Proxy modes are turned on. These settings do not affect the final render quality.


Even though the images are being processed smaller than their original size, the image viewing scales in the viewers still refer to original resolutions. Additionally, image processing performed in Proxy Scale mode may differ slightly from full-resolution rendering.

The Proxy and Auto Proxy size ratios may be changed from within the interface itself by right-clicking on the Prx and APrx buttons above the toolbar and selecting the desired value from the contextual menu.

GPU

The GPU preference is only available in Fusion Studio. In DaVinci Resolve, you can configure the GPU processing in Resolve's Memory and GPU preferences.

In Fusion Studio, the GPU preference is used to specify the GPU acceleration method used for processing, based on your computer platform and hardware capabilities. It is also used for enabling caching and debugging GPU devices and tools.



The GPU preferences

Options

The GPU options include radio buttons to select whether the GPU is used when processing and, if so, which computer framework is used for communicating with the GPU.

- **GPU Tools:** This preference has three settings: Auto, Disable, and Enable. When set to Disable, no GPU acceleration is used for tools or third-party plug-ins. Fuses may still require GPU acceleration. If Enable is selected, GPU acceleration is available for tools and plug-ins, if appropriate drivers are installed.
- **API:** The API setting selects the GPU processing method to use.
- **Device:** The Device setting determines which GPU hardware to use in the case of multiple GPUs. The Auto setting gives priority to GPU processing; however, if it is unavailable, Fusion uses the platform default. Currently, both the AMD and CPU options require either the AMD Catalyst 10.10 Accelerated Parallel Processing (APP) technology Edition driver or the ATI Stream SDK 2.1 or later to be installed. The Select setting allows you to choose the device explicitly.

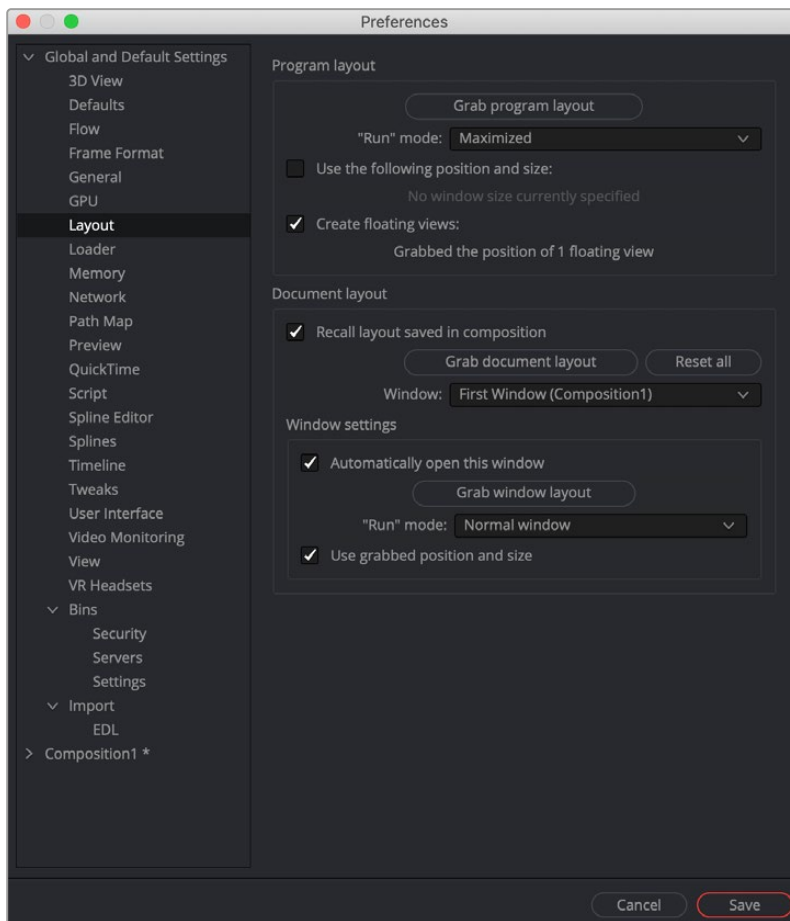
Debugging

The more advanced preferences located in this section are designed for diagnostics and analyzing GPU operations.

- **Verbose Console Messages:** Enabling this option causes information to be shown in the Console. For example, Startup Logs, Compiler Warnings, and Messages.
- **OpenGL Sharing:** Enabling this option shares system RAM with onboard GPU RAM to create a larger, but slower, OpenGL memory pool.
- **Clear Cache Files:** This option will clear already compiled GPU code and then recompile the kernels.

Layout

The Layout preferences are only available in Fusion Studio. To save a Layout in DaVinci Resolve's Fusion page, use the Workspace > Layout Presets menu. The Layout options are used to control the layout, size, and position of various windows in Fusion's interface at startup or when a comp is created.



The Layout preferences

There are a lot of options, but in practice, you simply organize the interface the way you prefer it on startup and when a new composition is created, then open this Preferences panel and click on the three buttons to grab the Program Layout, the Document Layout and the Window Settings.

Program Layout

The Program Layout is used to save the overall Fusion interface window and any open floating windows. Each new composition you open within the larger overall Fusion interface window will adhere to these preferences.

- **Grab Program Layout:** Pressing this button stores the application's overall current position and size.
- **Run Mode:** This menu is used to select the application's default mode at startup. You choose between a Maximized application window, a Minimized application, or a Normal application display.
- **Use the Following Position and Size:** When checked, the values stored when Grab Program Layout was selected will be used when starting Fusion Studio.
- **Create Floating Views:** When checked, the position and size of the floating viewers will be saved when the Grab Program Layout button is used.

Document Layout

The Document Layout is used to save the layout of panels and windows for the current Fusion comp.

- **Recall Layout Saved In Composition:** When checked, all Document Layout settings in the controls below will be recalled when a saved composition is loaded.
- **Grab Document Layout:** Pressing this button stores the entire interface setup, including all the internal positions and sizes of panels and work areas.
- **Window:** When multiple windows on the same composition are used, this menu is used to select the window to which the Window Settings will apply.


Window Settings

Rather than saving entire comp layouts, you can save position and size for individual floating windows and panels within a comp using the Window Settings.

- **Automatically Open This Window:** When checked, the selected window will automatically be opened for new flows.
- **Grab Window Layout:** Pressing this button stores the size and position of the selected window.
- **Run Mode:** Select the default run mode for the selected window. You can choose between a Maximized window, a Minimized window, or a Normal window display.
- **Use Grabbed Position and Size:** When checked, the selected window will be created using the stored position and size.

Loader

The Loader preferences are only available in Fusion Studio. Using the Loader preferences, you can set options for the default Loader's color depth and aspect ratio as well as define the local and network cache settings.



The Loader preferences

Defaults

The Defaults section includes two settings to determine how color depth and aspect ratio are handled for Loaders.

- **Loader Depth:** The Loader Depth defines how color bit depth is handled when adding a Loader. Choosing Format means that the correct bit depth is automatically selected, depending on the file format and the information in the file's header. Choosing Default sets the bit depth to the value specified in the Frame Format preferences.

Cache

The Cache preferences allow you to control how disk caching operates in Fusion. You can set how and where the cache is generated, when the cache is removed, how the cache reacts when source files are not available, as well as many other cache related options. This is not to be confused with RAM cache, which is controlled in the Memory preferences.

- **Disable All Local Caching:** This setting disables local caching.
- **Cache Files from Network DiskCaches:** If a tool has disk caching enabled, and the disk cache files are stored remotely on the network, then enabling this option will use a local copy of those cache files, similarly to the local cache on a networked Loader.
- **Enable Local Caching of Loaders:** Files will be copied into the LoaderCache path set below or in the Path Maps preferences.
- **Cache Multi-Frame Files:** Files like AVI or QuickTime will be copied into the LoaderCache path. This may take some time if the file is large.
- **Don't Cache Files from Local Disks:** Files that do not sit on a network drive will not be copied into the LoaderCache path. You can disable this option if you have, for example, a fast SSD cache drive and want to use it for local files as well to speed up file access while working interactively.
- **Only Files Smaller Than xxx MB.:** Files larger than the value set here will not be copied into the LoaderCache path.
- **Cache Path Separator Character:** When Enable Local Caching of Loaders is enabled, you can use this setting to rebuild the path of the original files in LoaderCache.


For instance, given the default “!” character, the original path X\Project\MyShots\Shot0815\ will be translated into X!Project!MyShots!Shot0815! in the LoaderCache path. Other separator characters may be used, including the “\” character, which will use subdirectories in LoaderCache: X\Project\MyShots\Shot0815\.
- **If Original File Is Missing:** This setting provides three options to determine the caching behavior when the original files can't be found. The Fail option behaves exactly as the Default Loader in Fusion. The Loader will not process, which may cause the render to halt. The Load Cache option loads the cache even though no original file is present. The Delete Cache option clears missing files from the cache.
- **Cache Location:** For convenience, this is a copy of the LoaderCache path set in the Path Maps preferences.
- **Explore:** This button opens the LoaderCache path in the macOS X Finder window or a Windows Explorer window.
- **Clear All Cache Files:** This button deletes all cached files present in the LoaderCache path.

Memory

The Memory preferences are only available in Fusion Studio. To control Fusion's memory when using the Fusion page in DaVinci Resolve, open DaVinci Resolve's Memory and GPU preferences.

Occasionally, it will be necessary to adjust the Memory preferences in order to make the best use of available memory on the computer. For example, some people prefer a higher cache memory for

faster interactive work, but for final renders the cache memory is often reduced, so there's more memory available for simultaneous processing of tools or multiple frames being rendered at once.



The Memory preferences

Caching Limits

The Caching Limits include options for Fusion's RAM cache operation. Here, you can determine how much RAM is allocated to the RAM cache for playing back comps in the viewer.

- **Limit Caching To:** This slider is used to set the percentage of available memory used for the interactive tool cache. Available memory refers to the amount of memory installed in the computer.
- When the interactive cache reaches the limit defined in this setting, it starts to remove lower priority frames in the cache to clear space for new frames.
- **Automatically Adjust In Low Memory Situations:** This checkbox will set the caching to adjust when memory is low. The console will display any cache purges.
- **Leave At Least X MBytes:** This setting is used to set the hard limit for memory usage. No matter what the setting of the Cache Limit, this setting determines the amount of physical memory available for use by other applications. Normally, this value should not be smaller than 25 MBytes.

Interactive Render

The Interactive Render option allows you to optimize Fusion's processing based on the amount of RAM you have installed in your system.

- **Simultaneous Branching:** When checked, more than one tool will be processed at the same time. Disable this checkbox if you are running out of memory frequently.

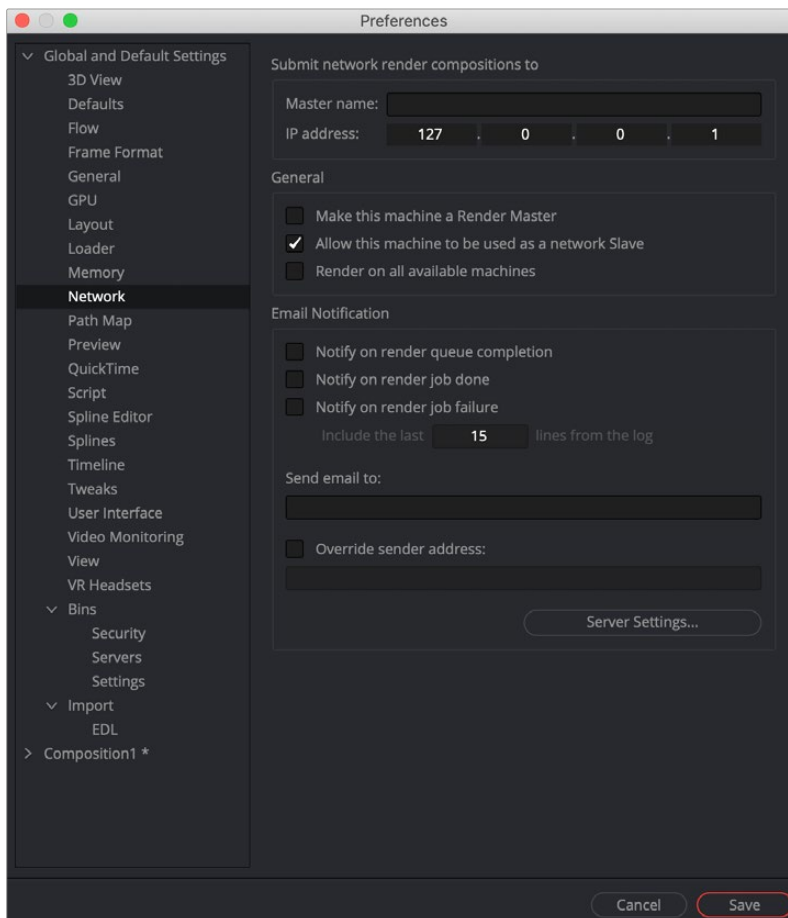
Final Render

These settings apply to memory usage during a rendering session, either preview or final, with no effect during an interactive session.

- **Render Slider:** This slider adjusts the number of frames that are rendered at the same time.
- **Simultaneous Branching:** When checked, more than one branch of a node tree will be rendered at the same time. If you are running low on memory, turn this off to increase rendering performance.

Network

The Network preferences are only available in Fusion Studio. These preferences are used to set up and control network rendering in Fusion Studio. The majority of settings are found in the Render Manager dialog.



The Network preferences

Submit Network Render Compositions

In these fields, you enter the Master Name and IP address of the computer that will manage all network renders sent from this machine. If a standalone render master is in use on the network, these fields may be pre-filled and may not be editable. This is done to prevent multiple unauthorized render masters from being created by each person in a facility.

To re-enable editing of the master name and IP, create the environment variable `FUSION_NO_MANAGER` and set the value to `True`. Check your operating system user guide for how to create environment variables.

General

The General preferences are designed with the most used options at the top in the General section. These options determine in what capacity the system is used during network rendering.

- **Make This Machine a Render Master:** When enabled, Fusion will accept network render compositions from other computers and manage the render. It does not necessarily mean that this computer will be directly involved in the render, but it will submit the job to the render nodes listed in the Render Manager dialog.
- **Allows This Machine to Be Used as a Network Slave:** When enabled, this computer can be used as a Render node and will accept compositions for network rendering. Deselect it to prevent other people from submitting compositions to render on this computer.
- **Render on All Available Machines:** Enable this checkbox to ignore groups and priorities configured in the Render Manager. Compositions submitted from this computer for network rendering will always be assigned to every available slave.

Email Notification

You can use the Email Notification section to set up who gets notified with status updates regarding the render jobs and the network.

- **Notify Options:** These checkboxes cause emails to be sent when certain render events take place. The available events are Queue Completion, Job Done, and Job Failure.
- **Send Email to:** Enter the address or addresses to which notifications should be sent. You separate multiple addresses with a semicolon.
- **Override Sender Address:** Enter an email address that will be used as the sender address. If this option is not selected, no sender address is used, which may cause some spam filters to prevent the message from being delivered to the recipient.

Server Settings

This section covers Clustering and Network Rendering. For more information on these settings and clustering, see Chapter 64, “Rendering Using Saver Nodes,” in the DaVinci Resolve Reference Manual or Chapter 4 in the Fusion Reference Manual.


Path Maps

Path Maps are virtual paths used to replace segments of file paths with variables. For example, define the path ‘movie_x’ as actually being in X\Shows\Movie_X. Using this example, Fusion would understand the path ‘movie_x\scene_5\scan.000.cin’ as actually being X:\Shows\Movie_X\scene_5\scan.000.cin.

For Fusion Studio, there are two main advantages to virtual path maps instead of actual file paths. One is that you can easily change the path to media connected to Loaders (for example, when moving a comp from one drive to another), without needing to make any changes in the composition. The other advantage is when network rendering, you can bypass the different OS filename conventions.

- **Enable Reverse Mapping of Paths Preferences:** This checkbox is at the top of the Path Map settings. When enabled, Fusion uses the built-in path maps for entries in the path’s settings when applying mapping to existing filenames. The main benefit is for Fusion Studio. Enabling this checkbox causes Loaders to automatically use paths relative to the location of the saved composition when they are added to the Node Editor. For more information on using relative paths for Loaders, see Chapter 104, “IO Nodes,” in the DaVinci Resolve Reference Manual or Chapter 44 in the Fusion Reference Manual.

As with other preferences in Fusion Studio, paths maps are available in both Global and Composition preferences. Global preferences are applied to all new compositions, while Composition path maps are only saved with the current composition. Composition path maps will override Global path maps with the same name.



The Path Map preferences

The Global paths maps are divided into three sections:

- **System Path Maps:** The operating system determines system path maps, and they define Fusion’s global locations. You can override specific System path maps using the Defaults or current Composition Path Map settings. If you change your mind at a later time, you are always able to return to Fusion’s “factory” defaults using the System path maps. There are several top-level path maps established in the System Path Map settings.
 - **AllData:** The folder where Fusion saves all shared application data.
 - **AllDocs:** The folder where Fusion saves the public/shared document folder.
 - **AllLUTs:** The nested LUTs path in the Defaults section, where Fusion saves LUTs.
 - **Fusion:** The folder where Fusion Studio app is installed. For example, if you open Fusion from C:\Program Files\Fusion, then the path Fusion:\Help refers to C:\Program Files\Fusion\Help. If you instead used a copy of Fusion found in \\post-server\fusion\16, then Fusion:\Help would expand to \\post-server\fusion\16\Help.
 - **FusionLibs:** The Fusion libraries used for the application.
 - **Profile:** The folder where default Fusion preferences file is saved.
 - **Profiles:** The folder where Fusion individual user preferences are saved.
 - **Programs:** The location of Fusion Studio or DaVinci Resolve.
 - **SystemFonts:** The folder where the OS saves fonts that appear for Text+ and Text 3D nodes.
 - **Temp:** The system’s temporary folder.

- **UserData:** The folder where Fusion saves all user-specific miscellaneous roaming data. The individual elements included in the roaming data are listed in the Default Path Maps section. For Fusion Studio on Windows, this is “C:\Users\username\AppData\Roaming”. On Linux, this will be “\$HOME/Blackmagic/Fusion”. On macOS, this is “Users/UserName/Library/Application Support/Blackmagic Design/Fusion”.
- **UserDocs:** The folder where Fusion saves the user’s document folders.
- **Default Path Maps:** The Defaults are user-editable path maps. They can reference the System paths, as part of their paths. For instance, the Temp folder is defined in the System path and used by the Default DiskCache path map to refine the nested location (Temp:DiskCache). Default path maps can also redirect paths without using the Global System path maps. After you change a Default, the updated setting can be selected in the Preferences window, and a Reset button at the bottom of the Preferences window will return the modified setting to the System default.
 - **AutoSaves:** This setting determines the Fusion Comp AutoSave document’s location, set in the Fusion General preferences.
 - **Bins:** Sets the location of Fusion Studio bins. Since the bins use pointers to the content, the content is not saved with the bin. Only the metadata and pointers are saved in the bins.
 - **Brushes:** Points Fusion to the folder that contains custom paintbrushes.
 - **Comps:** The folder where Fusion Studio compositions are saved. On macOS or Windows, the default location is in Users/YourUserName/Documents/Blackmagic Design/Fusion.
 - **Config:** Stores Configuration files used by Fusion Studio during its operation.
 - **Defaults:** Identifies the location of node default settings so they can be restored if overwritten.
 - **DiskCache:** Sets the location for files written to disk when using the Cache to Disk feature. This location can be overridden in the Cache to Disk window.
 - **Edit templates:** The location where Fusion macros are saved in order to appear as templates in the DaVinci Resolve Effects Library.
 - **Filters:** Points to a folder containing Convolution filters like sharpen, which can be used for the Custom Filter node.
 - **Fonts:** The default path map for Fonts points to the operating system fonts folders. Changing this will change the fonts that are available in the Text+ or Text 3D nodes as well as any Fusion Title Template. In DaVinci Resolve. This path map does not affect the five additional Edit page titles (L Lower 3rd, R Lower 3rd, M Lower 3rd, Scroll, and Text.)
 - **Fuses:** Points to a folder containing Fusion Fuses plug-ins.
 - **FusionTemplates:** Location where Fusion macros are saved in order to appear as templates in Fusion’s Effects Library.
 - **Guides:** Location where custom viewer guide overlays are stored.
 - **Help:** Identifies where Fusion Studio PDF files are located.
 - **Layouts:** Location where Fusion Studio custom window layouts are saved.
 - **Libraries:** Points to a support folder where custom Effects Library items can be stored.
 - **LoaderCache:** The Fusion Studio Loader preferences allow the Loader to cache when reading from a slow network. This path map point to the local drive location for that cache.
 - **LuaModules:** Location for Lua Scripting modules.
 - **LUTs:**
 - **Macros:** Points to the location for user created macros. The macros saved to this location appear in the macros category of the Effects Library and in the right-click Edit Macro contextual menu.

- **Plugins:** This refers to user specific OpenFX plug-ins that you do not want loaded for all users.
- **Previews:** Path map used for the older style, file sequence flipbook previews.
- **Queues:** Location of the Render manager list.
- **Scripts:** Location of Lua and Python scripts. This path can be further refined into specific scripts for tools (nodes), comps, and other specific script types.
- **Settings:** Location where custom Node settings are saved.
- **Stamps:** Location for preview movies generated in a Fusion Studio bin. This is an outdated path map since bins now include the Studio Player.
- **Templates:** Location of the Templates folder. Saving Macros to the Template folder will cause them to appear in the Effects Library in a Templates category. In Fusion Studio, the Templates category does not appear until a Macros is saved into the folder.
- **Thumbs:** Location for clip thumbnails generated in a Fusion Studio bin. This is an outdated path map since bins include now include the Studio Player.
- **UserPaths:** Used for locations of studio- or facility-specific tools, like custom plug-ins and scripts located on a central server.
- **User Path Maps:** User paths are new paths that you have defined that do not currently exist in the Defaults settings.
 - Comp refers to the folder where the current composition is saved. For instance, saving media folders within the same folder as your Fusion Studio comp file is a way to use relative file paths for Loaders instead of actual file paths.

Modifying a System Path Map

To modify an existing System path map, select the path map in the System section. Click the folder icon at the bottom of the Preferences window, and enter the name of the path map in the From field below. Enter the value of the path map in the To: field.

Modifying a Default Path Map

To modify an existing Default path map, select the path map in the Default section. Click the folder icon at the bottom of the Preferences window, and enter the name of the path map in the From field below. Enter the value of the path map in the To: field.

Creating a User Path Map

To create a path map, click on the New button and enter the name of the path map in the From field below. Enter the value of the path map in the To: field.

Deleting a Path Map

To delete a user-created path map, select it from the list and click the Delete button. System and Default path maps cannot be deleted; only user created path maps can be removed from the Path Maps list.


Nesting Path Maps

When defining your own path map, you can use an existing path map in the new definition. For example, define a path map called 'Episode' that maps to MyDrive\ Projects\Episode1. Then create new path maps called Renders and Stills that map to Episode\ Renders_v1 and Episode\Stills_v1.

Preview

Preview is only available in Fusion Studio. Previews in DaVinci Resolve use the Scratch Disk setting in the Media Storage preferences.

In the Preview preferences, you configure the creation and playback options for preview renders.



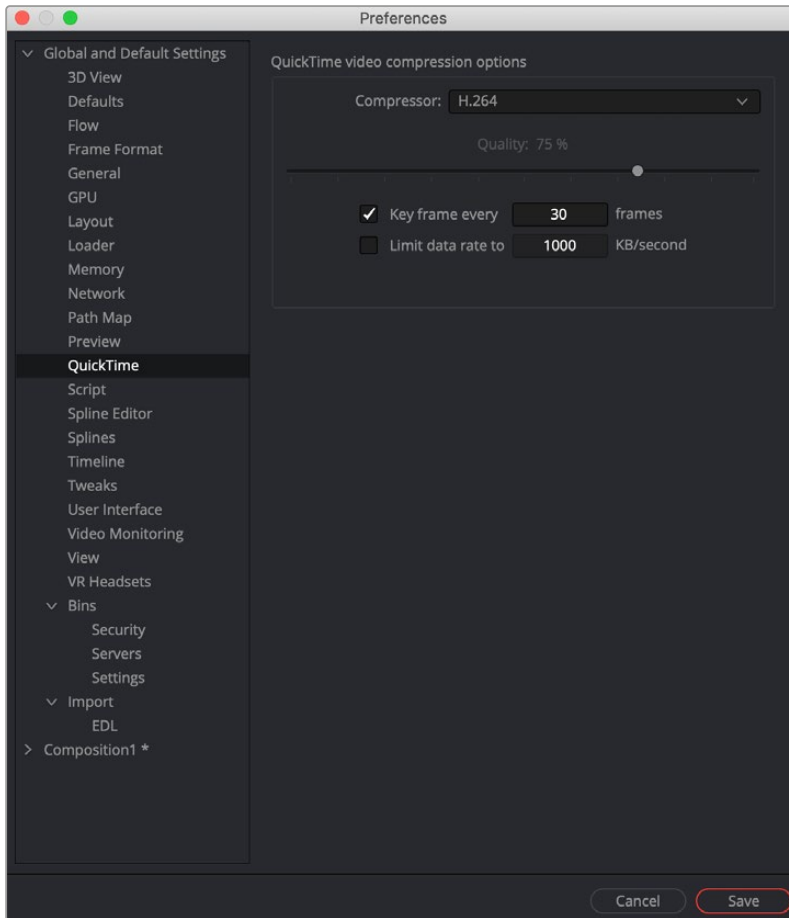
The Preview preferences

Options

- **Render Previews Using Proxy Scaling:** When checked, this option scales down the images to the preview size for the Loader and Creator tools. This causes much faster rendering. If this option is disabled, frames will be rendered at full size and are then scaled down.
- **Skip Frames to Maintain Apparent Framerate:** When checked, frames are skipped during playback of Flipbooks and file sequences to maintain the frame rate setting.
- **Show Previews for Active Loaders:** This setting determines whether the preview playback controls are shown below the Inspector when a Loader with a valid file is activated.
- **Show Previews for Active Savers:** This setting determines whether the preview playback controls below the Inspector are shown when a Saver with a valid file is activated.
- **Display File Sequences On:** This setting determines which viewer or external monitor is used for the interactive and file sequence playbacks as well as for the scrubbing function in the bins.

QuickTime

The QuickTime preferences are only available in Fusion Studio on macOS. These settings configure the default QuickTime codec settings when you select QuickTime as the rendering file format in the Saver node.

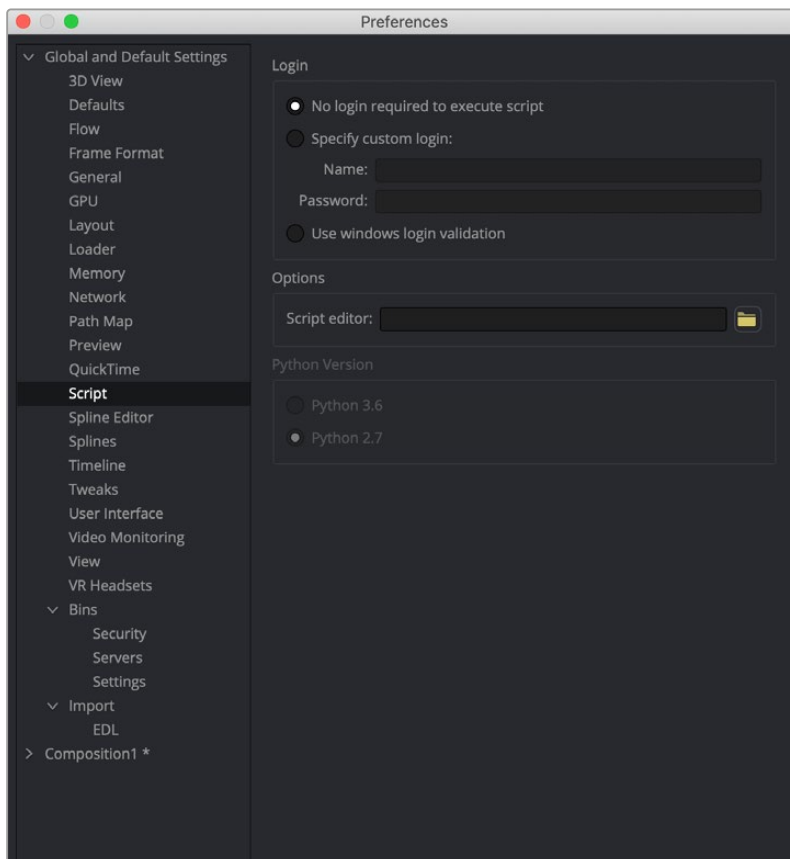


The QuickTime preferences

- **Compressor:** This drop-down menu displays the QuickTime codecs available from your computer. Fusion tests each codec when the program is started; therefore, some codecs may not be available if the tests indicate that they are unsuitable for use within Fusion.
- **Quality:** This slider is used to determine the amount of compression to be used by the codec. Higher values produce clearer images but larger files. Not all codecs support the Quality setting.
- **Key Frame Every X Frames:** When checked, the codec will create key frames at specified intervals. Key frames are not compressed in conjunction with previous frames and are, therefore, quicker to seek within the resulting movie. Not all codecs support the key frame setting.
- **Limit Data Rate To X KB/Second:** When checked, the data rates of the rendered file will be limited to the amount specified. Not all codecs support this option. Enter the data rate used to limit the QuickTime in kilobytes (kB) per second, if applicable. This control will have no effect if the Limit Data Rate To option is not selected.

Script

The preferences for Scripting include a field for passwords used to execute scripts from the command line and applications for use when editing scripts.



The Script preferences

Login

There are three login options for running scripts outside of the Fusion application.

- **No Login Required to Execute Script:** When enabled, scripts executed from the command line, or scripts that attempt to control remote copies of Fusion, do not need to log in to the workstation in order to run.
- **Specify Custom Login:** If a username and password are assigned, Fusion will refuse to process incoming external script commands (from FusionScript, for example), unless the Script first logs in to the workstation. This only affects scripts that are executed from the command line, or scripts that attempt to control remote copies of Fusion. Scripts executed from within the interface do not need to log in regardless of this setting. For more information, see the Scripting documentation.
- **Use Windows Login Validation:** When using Fusion on Windows, enabling this option verifies the user name and password (also known as credentials) with the operating system before running the script.

Options

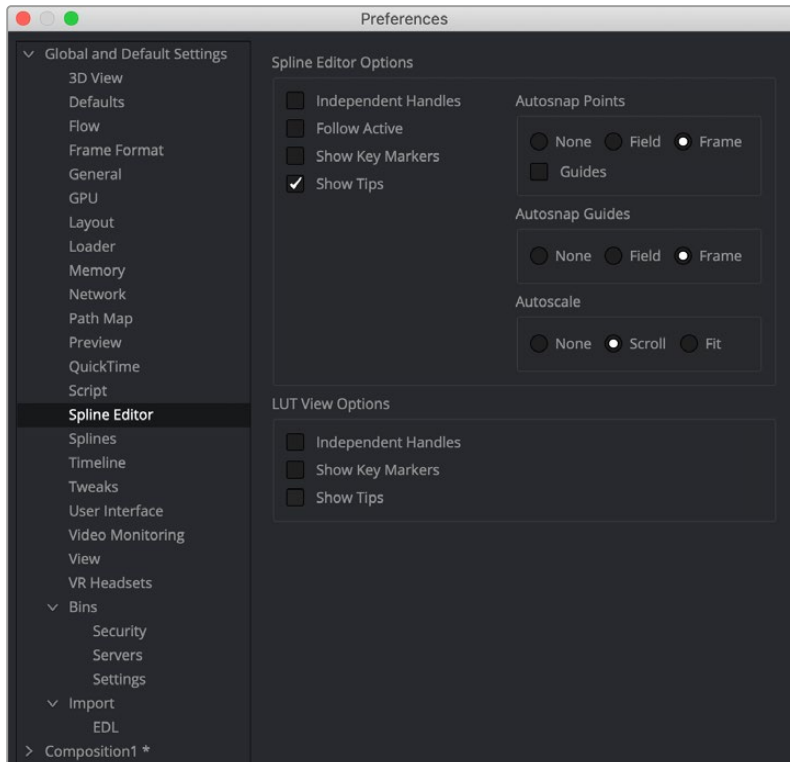
- **Script Editor:** Use this preference to select an external editor for scripts. This preference is used when selecting Scripts > Edit.

Python Version

- Two options are presented here for selecting the version of Python that you plan on using for your scripts.

Spline Editor

The Spline Editor preferences allow you to set various spline options for Autosnap behavior, handles, markers, and more. This only affects splines displayed in the Spline Editor, not splines created in the viewer using the polygon tool or paths.



The Spine Editor preferences

Spline Editor Options

These settings control the spline behavior in the Spline Editor, as well as the appearance of the graph area.

- **Independent Handles:** Enabling this option allows the In or Out direction handle on newly created key frames to be moved independently without affecting the other. This option is also available via the Options submenu when right-clicking in the Spline Editor graph.
- **Follow Active:** The Spline Editor focuses on the currently active tool. This option is also available via the Options submenu when right-clicking in the Spline Editor graph.
- **Show Key Markers:** Small colored triangles will be displayed at the top of the Spline Editor Time Ruler to indicate key frames on active splines. The colors of the triangles match the colors of the splines. This option is also available via the Show submenu when right-clicking in the Spline Editor graph.
- **Show Tips:** Toggles if tooltips are displayed or not. This option is also available via the Show submenu when right-clicking in the Spline Editor graph.
- **Autosnap Points:** When moving points in the Spline Editor, these will snap to the fields or frames or can be moved freely. This option is also available via the Options submenu when right-clicking in the Spline Editor graph.
- **Guides:** When moving points in the Spline Editor, these will snap to guides as well. This option is also available via the Options submenu when right-clicking in the Spline Editor graph.
- **Autosnap Guides:** When moving or creating guides, these will snap to the fields or frames or can be moved freely. This option is also available via the Options submenu when right-clicking in the Spline Editor graph.

- **Autoscale:** Keeps the Spline Editor scales intact on changing the editable spline content of the graph. This scale is also available via the Options submenu when right-clicking in the Spline Editor graph.
- **Scroll:** Scrolls horizontally and vertically to show all or most of the spline points. This option is also available via the Scale submenu when right-clicking in the Spline Editor graph.
- **Fit:** Zooms to fit all points within the spline graph, if necessary. This option is also available via the Scale submenu when right-clicking in the Spline Editor graph.

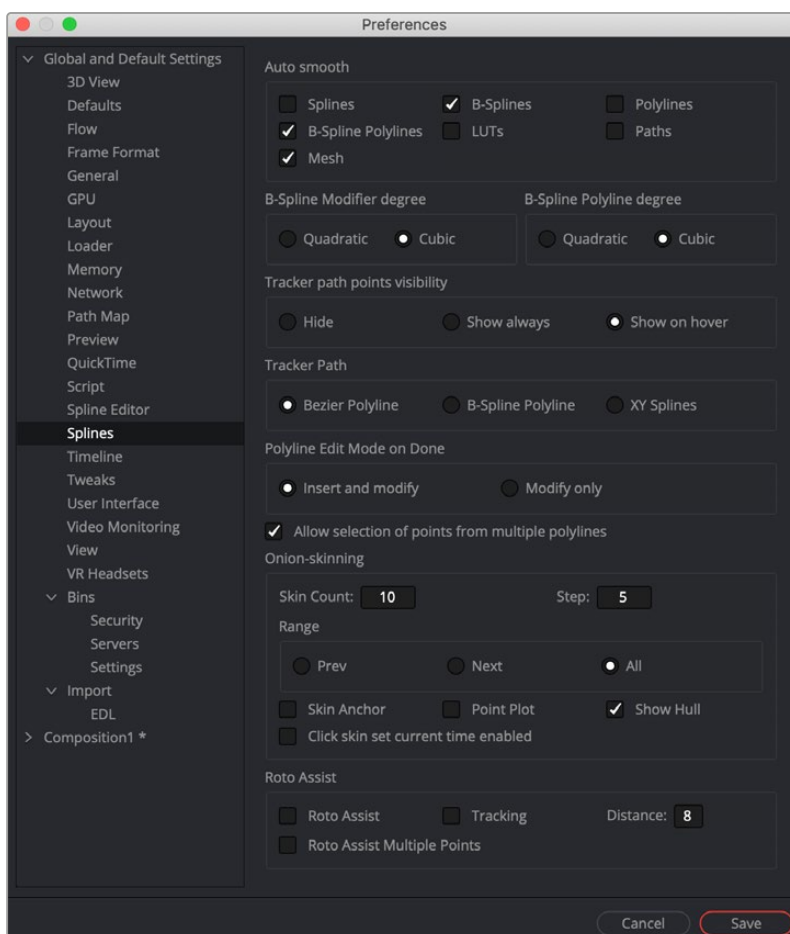
LUT View Options

LUT splines are not currently available in the Spline Editor, and therefore these options are not operational in the current version.

- **Independent Handles:** Enabling this option allows the In or Out direction handle on newly created key frames to be moved independently without affecting the other.
- **Show Key Markers:** Small colored triangles will be displayed at the top of the Spline Editor Time Ruler to indicate key frames on active splines. The colors of the triangles match the colors of the splines.
- **Show Tips:** Toggles whether tooltips are displayed.

Splines

Options for the handling and smoothing of animation splines, tracker path defaults, and rotoscoping are found in the Splines preferences.

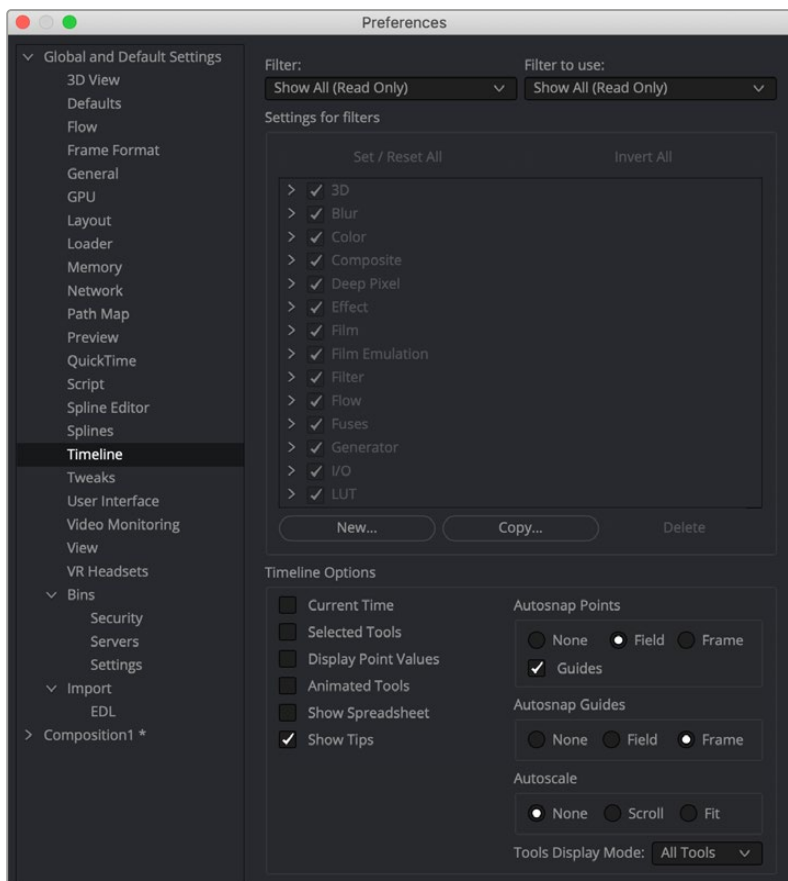


The Splines preferences

- **Autosmooth:** Automatically smooths out any newly created points or key frames on the splines selected in this section. You can choose to automatically smooth animation splines, B-Splines, polyline matte shapes, LUTs, paths, and meshes.
- **B-Spline Modifier Degree:** This setting determines the degree to which the line segments influence the resulting curvature when B-Splines are used in animation. Cubic B-Splines determine a segment through two control points between the anchor points, and Quadratic B-Splines determine a segment through one control point between the anchor points.
- **B-Spline Polyline Degree:** This setting is like the one above but applies to B-Splines used for masks.
- **Tracker Path Points Visibility:** This setting determines the visibility of the control points on tracker paths. You can show them, hide them, or show them when your cursor hovers over the path, which is the default behavior.
- **Tracker Path:** The default tracker creates Bézier-style spline paths. Two other options in this setting allow you to choose B-Spline or XY Spline paths.
- **Polyline Edit Mode on Done:** This setting determines the state of the Polyline tool after you complete the drawing of a polyline. It can either be set to modify the existing control points on the spline or modify and add new control points to the spline.
- **Onion Skinning:** The Onion Skinning settings determine the number of frames displayed while rotoscoping, allowing you to preview and compare a range of frames. You can also adjust if the preview frames only from the frame prior to the current frame, after the current frames, or split between the two.

Timeline

The Timeline preferences is where you create and edit Keyframes Editor/Spline Editor filters and set default options for the Keyframes Editor.



The Timeline preferences

Filter/Filter to Use

The Filter menu populates the hierarchy area below the menu with that setting. It lets you edit the filters. The Filter to Use menu selects the default filter setting located in the Keyframes Editor Options menu.

Settings for Filters

This area is used to create a new filter and define its settings. You start by first clicking the New button and entering the name of the new Filter. You then select any of the tools that you want the filter to contain. Only tools that are checked will appear in the Keyframes Editor or Spline Editor when the filter is selected. You can also create a copy of the filter using the Copy button or remove a filter from the list by clicking the Delete button.

Timeline Options

The Timeline Options configure which options in the Keyframe Editor are enabled by default. A series of checkboxes correspond to buttons located in the Timeline, allowing you to determine the states of those buttons at the time a new comp is created. For more information on the Keyframes Editor functions, see Chapter 69, “Animating in Fusion’s Keyframes Editor,” in the DaVinci Resolve Reference Manual or Chapter 9 in the Fusion Reference Manual.

- **Autosnap Points:** When moving points in the Keyframes Editor, the points will snap to the fields or to the frames, or they can be moved freely.
- **Guides:** When moving points in the Keyframes Editor, the point will snap to the guides that are placed in the Timeline graph.
- **Autosnap Guides:** When moving or creating guides, the guides will snap to the fields or to the frames, or they can be moved freely.
- **Autoscale:** Keeps the Timeline scales intact while changing the editable spline content in the graph. When set to scroll, the Timeline scrolls horizontally and vertically to show all or most of the spline points when changing the editable spline content in the graph. When set to Fit, the Timeline zooms to fit all points within the graph, if necessary.
- **Tools Display Mode:** This menu controls the default sort order of the tools displayed in the Keyframes Editor. The default can be changed using the Sort order menu in the upper right of the Keyframes Editor.


Tweaks

The Tweaks preferences handle a collection of settings for fine-tuning Network rendering in Fusion Studio and graphics hardware behavior.

Network

The Network section is used to control and monitor the health of communication packets over TCP/IP when rendering over a network in Fusion Studio.

- **Maximum Missed Heartbeats:** This setting determines the maximum number of times the network is checked before terminating the communication with a Render node.
- **Heartbeat Interval:** This sets the time between network checks.
- **Load Composition Timeout:** This timeout option determines how long the Render Manager will wait for a composition to load before moving on to another task.
- **Last Slave Restart Timeout:** This timeout option determines how long the Render Manager will wait for a render slave to respond before using another render slave.



The Tweaks preferences

File I/O

The File I/O options are used to control the performance when reading frames or large media files from both direct and networked attached storage.

- **I/O Canceling:** This option enables a feature of the operating system that allows queued operations to be canceled when the function that requested them is stopped. This can improve the responsiveness, particularly when loading large images over a network.

Enabling this option will specifically affect performance while loading and accessing formats that perform a large amount of seeking, such as the TIFF format.

This option has not been tested with every hardware and OS configuration, so it is recommended to enable it only after you have thoroughly tested your hardware and OS configuration using drive loads from both local disks and network shares.

- **Enable Direct Reads:** Enabling this checkbox uses a more efficient method when loading a large chunk of contiguous data into memory by reducing I/O operations. Not every operating system employs this ability, so it may produce unknown behavior.
- **Read Ahead Buffers:** This slider determines the number of 64K buffers that are use to read ahead in a file I/O operation. The more buffers, the more efficient loading frames from disk will be, but the less responsive it will be to changes that require disk access interactively.

Area Sampling

The Area Sampling options allow you to fine-tune the RAM usage on Render nodes by trading off speed for lower RAM requirements.

- **Automatic Memory Usage:** This checkbox determines how area sampling uses available memory. Area sampling is used for Merges and Transforms. When the checkbox is enabled (default), Fusion will detect available RAM when processing the tool and determine the appropriate trade-off between speed and memory.

If less RAM is available, Fusion will use a higher proxy level internally and take longer to render. The quality of the image is not compromised in any way, just the amount of time it takes to render. In node trees that deal with images larger than 4K, it may be desirable to override the automatic scaling and fix the proxy scale manually. This can preserve RAM for future operations.

- **Pre-Calc Proxy Level:** Deselecting the Automatic Memory will enable the Pre-Calc Proxy Scale slider. Higher values will use less RAM but take much longer to render.


Open GL

This section controls how Fusion makes use of your graphics card when compositing in 3D with the Renderer 3D node. Most settings may be left as they are, but since OpenGL hardware varies widely in capabilities and different driver revisions can sometimes introduce bugs, these tweaks can be useful if you are experiencing unwanted behavior.

- **Disable View LUT Shaders:** OpenGL shaders can often dramatically accelerate View LUTs, but this can occasionally involve small trade-offs in accuracy. This setting will force Fusion to process LUTs at full accuracy using the CPU instead. Try activating this if View LUTs do not seem to be giving the desired result.
- **Use Float16 Textures:** If your graphics hardware supports 16-bit floating-point textures, activating this option will force int16 and float32 images to be uploaded to the viewer as float16 instead, which may improve playback performance.
- **Texture Depth:** Defines in what depth images are uploaded to the viewer.
 - **Auto:** The Auto option (recommended) lets Fusion choose the best balance of performance and capability.
 - **int8:** Similar to the Use Float16 Textures switch, this option can be used to force images to be uploaded to the Display View as int8, which can be faster but gives less range for View LUT correction.
 - **Native:** The Native option uploads images at their native depth, so no conversion is done.
- **Image Overlay:** The Image Overlay is a viewer control used with Merge and Transform tools to display a translucent overlay of the transformed image. This can be helpful in visualizing the transformation when it is outside the image bounds but may reduce performance when selecting the tool if cache memory is low. There are three settings to choose from: None, Outside, and All.
 - **None:** This setting never displays the translucent overlay or controls, which can reduce the need for background renders, in some cases resulting in a speed up of the display.
 - **Outside:** This will display only those areas of the control that are outside the bounds of the image, which can reduce visual confusion.
 - **All:** Displays all overlays of all selected tools.
- **Smooth Resize:** This setting can disable the viewer's Smooth Resize behavior when displaying floating-point images. Some older graphics cards are not capable of filtering floating-point textures or may be very slow. If Smooth Resize does not work well with float images, try setting this to flt16 or int.
- **Auto Detect Graphics Memory (MB):** Having Fusion open alongside other OpenGL programs like 3D animation software can lead to a shortage of graphics memory. In those cases, you can manually reduce the amount of memory Fusion is allowed to use on the card. Setting this too low or too high may cause performance or data loss.
- **Use 10-10-10-2 Framebuffer:** If your graphics hardware and monitor support 30-bit color (Nvidia Quadro/AMD Radeon Pro, and some Nvidia GeForce/AMD Radeon), this setting will render viewers with 10 bits per primary accuracy, instead of 8 bits. Banding is greatly reduced when displaying 3D renders or images deeper than 8-bit.

User Interface

The User Interface preferences set the appearance of the user interface window and how the Inspector is displayed.



The User Interface preferences

Appearance

When enabled, the Use Gray Background Interface checkbox will change the color of the background in Fusion's panels to a lighter, more neutral shade of gray.

Controls

This group of checkboxes manages how the controls in the Inspector are displayed.

- **Auto Control Open:** When disabled, only the header of the selected node is displayed in the Inspector. You must double-click the header to display the parameters. When enabled, the parameters are automatically displayed when the node is selected.
- **Auto Control Hide:** When enabled, only the parameters for the currently active tool (red outline) will be made visible. Otherwise, all tool headers will be visible and displayed based on the Auto Control Open setting.
- **Auto Control Close Tools:** When enabled, only the active (red outlined) tool in the Node Editor will have controls displayed. Any previous active node's tools will be closed in the Inspector. When disabled, any number of tools may be opened to display parameters at the same time. This setting has no effect if the Auto Control Hide checkbox is enabled.
- **Auto Control Close Modifiers:** When enabled, only one modifier's parameters will be displayed for the active node. Any additional modifiers for the active node will show only their header.
- **Auto Control Advance:** If the Auto Control Advanced checkbox is enabled, the Tab key and Return/Enter key will cause the keyboard focus to advance to the next edit box within the Inspector. When disabled, Return/Enter will cause the value entered to be accepted, but the keyboard focus will remain in the same edit box of the control. The Tab key can still be used to advance the keyboard focus.


- **Show Controls for Selected:** When this option is disabled, only the active tool’s parameters are shown in the Inspector. By default, it is enabled, showing controls for the active tool as well as all selected tools.
- **Combined Color Wheel:** When the Color Corrector tool is displayed in the Inspector, enabling this checkbox will show one color wheel with buttons to switch between the master, shadow, midtones, and highlight channels. Otherwise, four color wheels are displayed in the Inspector.
- **Gamma Aware Color Controls:** This setting adjusts color correction nodes when working with Rec. 709 images in a non-color managed project. Rec. 709 images appear correct on the computer monitor because monitors have a gamma adjustment built in. When working in the Rec. 709 color space without color management, enabling Gamma Aware color removes the gamma, applies the color correction as if it were linear, and then reapplies the gamma. For Rec. 709 images, enable the Gamma Aware setting and enter a Gamma value of 2.4. In a color managed linear project, this should be set to Off or a value of 1.0. When dealing with mixed color spaces, Fusion reads the metadata from the image and sets the Aware gamma value based on the metadata available.
- **Grab Distance:** This slider ranges from 1 to 10 and defaults to 5. It designates the active area around the mouse pointer and can be modified if you have difficulties in selecting points for modification in paths and spline curves. Smaller values will require a more accurate selection with the mouse pointer.

Touch Scrolling and Mouse Wheel

This group of settings allows you to configure which, if any, keyboard modifiers are needed to pan or zoom a panel when using a trackpad or middle mouse wheel.

Video Monitoring

This setting is only available in Fusion Studio. Control over video hardware for the Fusion Page is done in the DaVinci Resolve preferences. The Video Monitoring preferences are used to configure the settings of Blackmagic Design capture and playback products such as DeckLink PCIe cards and UltraStudio i/O units.



The Video Monitoring preferences

Video Output

This group of drop-down menus allows you to select the type of video I/O device you have installed, the output resolution, and the pixel format. These settings have nothing to do with your rendered output; it is only for your display hardware.

The Output HDR over HDMI settings are used to output the necessary metadata when sending high dynamic range signals over HDMI 2.0a and have it correctly decided by an HDR capable video display.

The Auto setting detects the image's values and outputs HDR. This will not affect non HDR images.

The Always setting forces HDR on all the time. This can be useful when checking non HDR and HDR grades.

When Auto or Always is selected, you can then set the “nit” level (slang for cd/m^2) to whatever peak luminance level your HDMI connected HDR display is capable of.


Stereo Mode

This group of settings configures the output hardware for displaying stereo 3D content.

- Mono will output a single non stereo eye.
 - Auto will detect which method with which the stereo images are stacked.
 - Use the Vstack option if the stereo images are stacked vertically as left on top and right at the bottom.
 - Use the Hstack option if the stereo images are stacked horizontally as left and right.
- The Swap eyes checkbox will swap the eyes if stereo is reversed.

View

The View preferences are used to manage settings and default controls for viewers.



The View preferences

Saved View Settings

The area at the top of the view preferences lists the currently saved settings that you create from the viewer's contextual menu. You can use the Rename and Delete buttons to manage the selected entries in the list. For more information on the viewer and its contextual menu, see Chapter 67, "Using Viewers," in the DaVinci Resolve Reference Manual or Chapter 7 in the Fusion Reference Manual.

Settings for View

Each viewer has its own preferences. The Settings for View drop-down menu is used to select the viewer you want to configure.

Control Colors

The Control Colors setting allows you to determine the color of the active/inactive onscreen controls.

Color Picking Area Size

You can use these width/height controls to set the number of pixels sampled when using the Color Picker in the viewers.

Displayed Depth Range

The Displayed Depth Range setting controls the view normalization of the Z-Channel.

Fit Margin

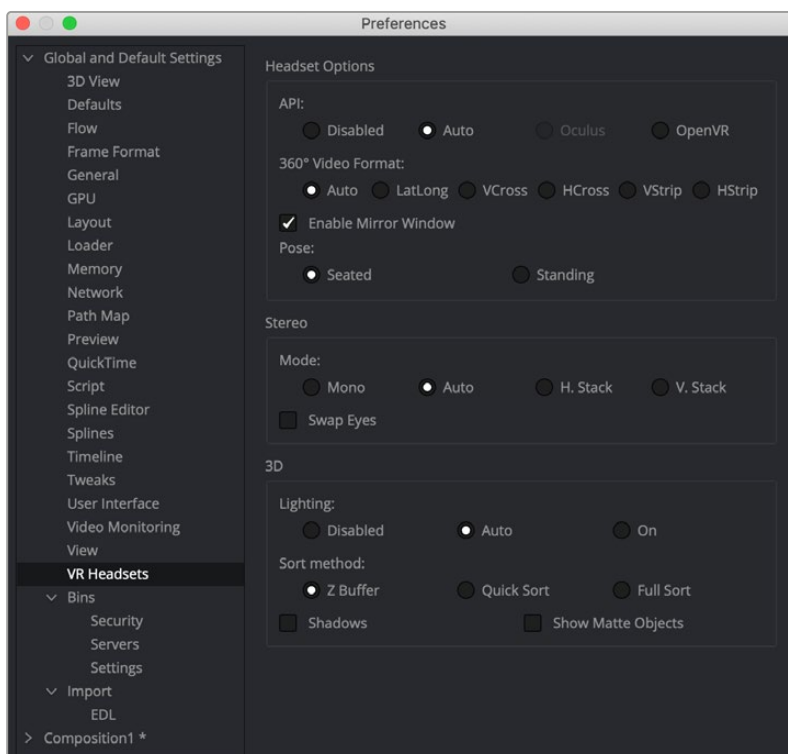
The Fit Margin setting determines how much padding is left around the frame when the Fit button is pressed or Fit is selected from the viewer's contextual menu.

Display LUT Plug-Ins

This list shows the available display LUTs and activates the selected one as default.

VR Headsets

The VR Headsets preferences allow configuration of any connected Virtual Reality headsets, including how stereo and 3D scenes are viewed.



The VR Headsets preferences

Headset Options

The Headset options are used to select the type of VR headset you are using to view the composite as well as the video layout of the 360° view.

API

- **Disabled:** Disabled turns off and hides all usage of headsets.
- **Auto:** Auto will detect which headset is plugged in.
- **Oculus:** Oculus will set the VR output to the Oculus headset.
- **OpenVR:** OpenVR will support a number of VR headsets like the HTC Vive.

360° Video Format

- **Auto:** Auto will detect the incoming image layout from the metadata and image frame aspect.
- **VCross and HCross:** VCross and HCross are the six square faces of a cube laid out in a cross, vertical or horizontal, with the forward view in the center of the cross, in a 3:4 or 4:3 image.
- **VStrip and HStrip:** VStrip and HStrip are the six square faces of a cube laid vertically or horizontally in a line, ordered as Left, Right, Up, Down, Back, Front (+X, -X, +Y, -Y, +Z, -Z), in a 1:6 or 6:1 image.
- **LatLong:** LatLong is a single 2:1 image in equirectangular mapping.
- **Enable Mirror Window:** Enable Mirror Window will show a window displaying the headset user's live view.

Stereo

Similar to normal viewer options for stereo 3D comps, these preferences control how a stereo 3D comp is displayed in a VR headset.

Mode

- **Mono:** Mono will output a single non stereo eye.
- **Auto:** Auto will detect the method with which the stereo images are stacked.
- **Vstack:** Vstack stereo images are stacked vertically as left on top and right at the bottom.
- **Hstack:** Hstack stereo images are stacked horizontally as left and right.
- **Swap Eyes:** Swap eyes will swap the eyes if stereo is reversed.

3D

Similar to normal viewer options for 3D comps, these preferences control how a 3D comp is displayed in a VR headset.

Lighting

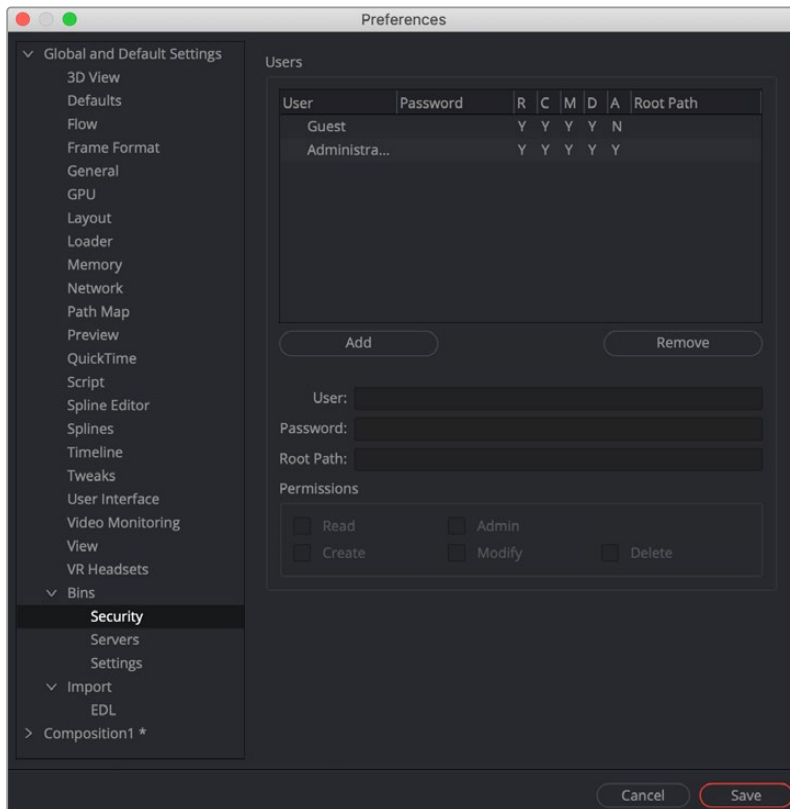
- Disabled lighting is off.
- Auto will detect if lighting is on in the view.
- On will force lighting on in the VR view.

Sort Method

- Z buffer sorting is the fast OpenGL method of sorting polygons.
- Quick Sort will sort the depth of polygons to get better transparency rendering.
- Full Sort will use a robust sort and render method to render transparency .
- Shadows can be on or off.
- Show Matte Objects will make matte objects visible in view or invisible.

Bins/Security

Bins preferences are only available in Fusion Studio. These preferences are used to manage the Bin users and their permissions.



The Bins Security preferences

Users List

The Users List is a list of the users and their permissions. You can select one of the entries to edit their settings using the User and Password edit boxes.

- **Add:** The Add button is used to add a new user to the list by entering a username and password.
- **Remove:** Click this button to remove the selected entry.

User

This editable field shows the username for the selected Bin Server item. If the username is unknown, try “Guest” with no password.

Password

Use this field to enter the password for the Bin user entered in the Users list.

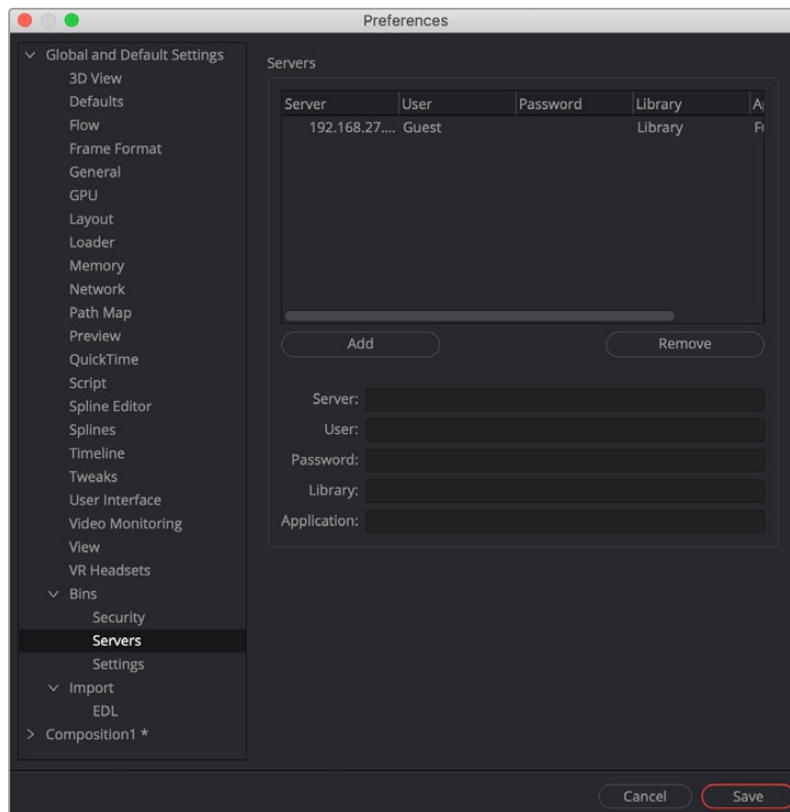
Permissions

The administrator can set up different permission types for users.

- **Read:** This will allow the user to have read-only permission for the bins.
- **Create:** This will allow the user to create new bins.
- **Admin:** This gives the user full control over the bins system.
- **Modify:** This allows the user to modify existing bins.
- **Delete:** This allows the user to remove bins.

Bins/Server

These preferences are used to add Bin Servers to the list of bins Fusion will display in the Bins dialog.



The Bin Servers preferences

Servers

This dialog lists the servers that are currently in the connection list. You can select one of the entries to edit its settings.

- **Add:** Use this button to add a new server to the list.
- **Remove:** Click this button to remove the selected entry.

Server

This editable field shows the name or IP address of the server for the selected entry in the list.

User

This editable dialog shows the username for the selected Bin Server item.

Password

Use this field to enter the password for the server entered in the Server list.

Library

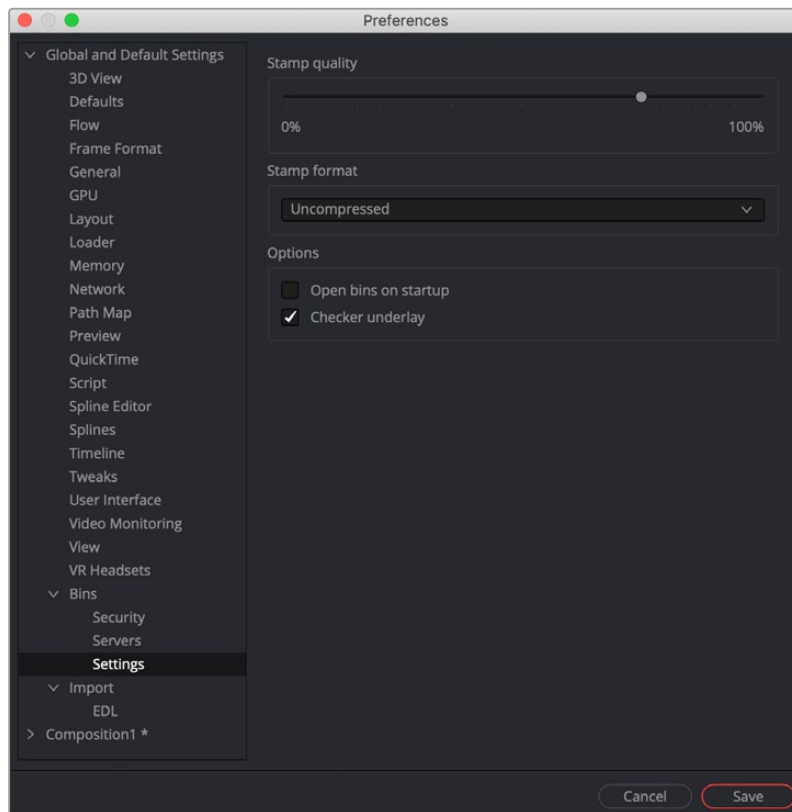
The Library field lets you name the bins. If you wanted to create a bin for individual projects, you would name it in the Library field and each project would get its own bin.

Application

The Application field allows larger studios to specify some other program to serve out the Bin requests.

Bins/Settings

These preferences are used to control the default behavior of bins.



The Bins Settings preferences

Stamp Quality

The Stamp Quality is a percentage slider that determines the compression ratio used for Stamp thumbnail creation. Higher values offer better quality but take up more space.

Stamp Format


This drop-down list determines whether the Stamp thumbnails will be saved as compressed or uncompressed.

Options

- **Open Bins on Startup:** When Open Bins on Startup is checked, the bins will open automatically when Fusion is launched.
- **Checker Underlay:** When the Checker Underlay is enabled, a checkerboard background is used for clips with alpha channels. When disabled, a gray background matching the Bin window is used as the clip's background.

EDL Import

The EDL Import options are used to determine how compositions are created from imported CMX-formatted EDL files.



The EDL Import preferences

Flow Format

This drop-down menu provides three options that determine how the node tree is constructed for the imported EDL file.

- **Loader Per Clip:** A Loader will be created for each clip in the EDL file.
- **A-B Roll:** A node tree with a Dissolve tool will be created automatically.
- **Loader Per Transition:** A Loader with a Clip list will be created, representing the imported EDL list.

Use Shot Names


When checked, shot names stored in the EDL file are used to locate the footage.

Customization

The following section covers the customization of preferences that are not technically part of the Preferences window. Using Fusion Studio's Hotkey Manager window, you can customize the keyboard shortcuts, making the entire process of working in Fusion not only faster but potentially more familiar if you are migrating from another software application. You can also customize Fusion with environment variables to switch between different preferences files, allowing different working setups based on different users or job types. Both of these customization options are only available in Fusion Studio.


Shortcuts Customization

Keyboard shortcuts can be customized in Fusion Studio. You can access the Hotkey Manager by choosing Customize HotKeys from the View menu.



Fusion has active windows to focus attention on those areas of the interface, like the Node Editor, the viewers, and the Inspector. When selected, a gray border line will outline that section. The shortcuts for those sections will work only if the region is active. For example, Command-F in the View will scale the image to fit the view area; in the Flow view, Command-F will open the Find tool dialog; and in the Spline editor, it will fit the splines to the window.

On the right is a hierarchy tree of each section of Fusion and a list of currently set hotkeys. By choosing New or Edit, another dialog will appear, which will give specific control over that hotkey.



Creating a new keyframe will give you the key combo to press, and this Edit Hotkey dialog will appear where the Action can be defined at top right: pressed, repeated, or released. The Name and abbreviated Short Name can be set, as can the Arguments of the action.

Customizing Preferences

Fusion Studio's preferences configure Fusion's overall application default settings and settings for each new composition. Although you access and set these preferences through the Preferences window, Fusion saves them in a simple text format called Fusion.prefs.

These default preferences are located in a \Profiles\Default folder and shared by all Fusion users on the computer. However, you may want to allow each user to have separate preferences and settings, and this requires saving the preferences to different locations based on a user login.

To change the saved location of the preferences file requires the use of environment variables.

Setting the Preferences Location

When you first open Fusion, the environment variable FUSION_PROFILE_DIR defines the folder that contains the Profiles folder. If this variable defines a valid path, then the preferences are saved to this folder. If the FUSION_PROFILE_DIR does not exist, then Fusion attempts to create it. If it cannot create the path, then the preferences are stored in the default Path Map location: AllData:\Profiles.

Typically, all users share the same preferences. If you want each user to save separate preferences within their home folder, you must create another environment variable with the name FUSION_PROFILE (e.g., FUSION_PROFILE=jane). Using this second environment variable, Fusion will look for the preferences in the PROFILE_DIR of the user profile. Using a login script, you can make sure the FUSION_PROFILE is set to the name of the logged in user.

Creating a Master Preferences File

When working with multiple Fusion users in a studio, you may want to standardize on a few settings. Using the FUSION_MasterPrefs environment variable, you can create one or more site-wide preferences in addition to your local personal preferences.

FUSION_MasterPrefs must contain the full path to at least one preferences file. If you have multiple preferences paths, separate them using semicolons. Fusion does not write to these prefs files, and they may contain a subset of all available settings. You may change settings in these files and use them only where local prefs do not already exist unless you set the Locked flag.

Locking Preferences

If the line "Locked = true," appears in the main table of a master file, all settings in that file are locked and override any other preferences. Locked preferences cannot be altered by the user.

Chapter 76

Controlling Image Processing and Resolution

This chapter covers the overall image-processing pipeline. It discusses color bit-depth and how to control the output resolution in a resolution-independent environment.

Contents

Fusion's Place in the DaVinci Resolve Image-Processing Pipeline	1352
Source Media into the Fusion Page	1352
Forcing Effects into the Fusion Page	1352
Output from the Fusion Page to the Color Page	1353
What Viewers Show in Different DaVinci Resolve Pages	1353
Managing Resolution In Fusion	1353
Changing the Resolution of a Clip	1354
Compositing with Different-Resolution Clips	1354
Sizing Between DaVinci Resolve Pages	1355
Color Bit Depths	1355
Understanding Integer vs. Float	1355
Setting Color Depth in Fusion Studio	1356
Combining Images with Different Color Depths	1357
Advantages of Floating-Point Processing	1358

Fusion's Place in the DaVinci Resolve Image-Processing Pipeline

When working in a single unified environment like DaVinci Resolve, it is important to understand the order of operations among the pages. DaVinci Resolve exposes some of this via the order of the page buttons at the bottom of the screen, with the Media, Cut, and Edit page at the beginning of the chain and the Color, Fairlight, and Deliver page at the end. However, this isn't the whole story, especially when it comes to the Fusion page. The following sections describe where the Fusion page fits in the image-processing chain of DaVinci Resolve.

Source Media into the Fusion Page

For ordinary, single clips coming in from the Edit or Cut page, the MediaIn node in the Fusion page represents the source media, as modified by the Clip Attributes window. Although you select the clip from the Edit or Cut page Timeline, in the Fusion page, the clip is accessed from the Media Pool.

TIP: The decoding or debayering of RAW files occurs prior to all other operations, and as such, any RAW adjustments will be displayed correctly in the Fusion page.

This means you have access to the entire source clip in the Fusion page, but the render range is set to match the duration of the clip in the Timeline. You also use the full resolution of the source clip, even if the Timeline is set to a lower resolution. However, none of the Edit or Cut page Inspector adjustments carry over into the Fusion page, with the exception of the Lens Correction adjustment.

When you make Zoom, Position, Crop, or Stabilization changes in the Edit or Cut page, they are not visible in the Fusion page. The same applies to any Resolve FX or OpenFX third-party plug-ins. If you add these items to a clip in the Edit or Cut page, and then you open the Fusion page, you won't see them taking effect. All Edit and Cut page timeline effects and Inspector adjustments, with the exception of the Lens Correction adjustment, are computed after the Fusion page but before the Color page. If you open the Color page, you'll see the Edit and Cut page transforms and plug-ins applied to that clip, effectively as an operation before the grading adjustments and effects you apply in the Color page Node Editor.

With this in mind, the order of effects processing in the different pages of DaVinci Resolve can be described as follows:




TIP: Retiming applied to the clip in the Edit page Timeline is also not carried over into the Fusion page.

Forcing Effects into the Fusion Page

There is a way you can force clips with Edit page Inspector adjustments, plug-ins, retiming, 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. However, bringing a compound clip into the Fusion page does change the resolution of the source clip to match the Timeline resolution. For more information, see the section “Sizing Between DaVinci Resolve Pages” in this chapter.

Output from the Fusion Page to the Color Page

The composition output from the Fusion page's MediaOut node are passed on via the Color page's source input, with the sole exception that if you've added plug-ins to that clip in the Edit or Cut page, then the handoff from the Fusion page to the Color page is as follows:



What Viewers Show in Different DaVinci Resolve Pages

Owing to the different needs of compositing artists, editors, and colorists, the viewers show different states of the clip.

- **The Edit page source viewer:** Always shows the source media, unless you're opening a compound clip that's been saved in the Media Pool. If Resolve Color Management is enabled, then the Edit page source viewer shows the source media at the Timeline color space and gamma.
- **The Edit page Timeline viewer:** Shows clips with all Edit page effects, Color page grades, and Fusion page effects applied, so editors see the program within the context of all effects and grading.
- **The Fusion page viewer:** Shows Media Pool source clips at the Timeline color space and gamma, but no Edit page Inspector adjustments or Resolve FX effects and no Color page grades.
- **The Color page viewer:** Shows clips with all Edit page effects, Color page grades, and Fusion page effects applied.

Managing Resolution In Fusion

There is no formal resolution to a comp in Fusion. Even though opening Fusion > Fusion Settings in the Fusion page or Preferences in Fusion Studio allows you to set the Width and Height in the Frame Format panel, those settings only affect the size of Fusion-generated images, like the Background tool, Fast Noise, and Text+ tool. The actual resolution of your composition is initially determined by the source resolution of the input image. However, it can be modified at any time using a variety of operations and nodes. For example, if you read in a full HD 1920 x 1080 resolution image, your comp starts at full HD 1920 x 1080 resolution. This is regardless of the Timeline resolution when you are using the Fusion page in DaVinci Resolve. The initial resolution of the Fusion comp is the size of the source media. Depending on how you combine images and the nodes you use, the output comp resolution can be maintained or modified.

TIP: The output of the Fusion page is placed back into the Edit page Timeline based on DaVinci Resolve's Image Sizing setting. By default, DaVinci Resolve uses an image sizing setting called Scale to Fit. This means that even if the Fusion page outputs a 4K composition, it conforms to 1920 x 1080 if that is what the project or a particular Timeline is set to. Changing the image sizing setting in DaVinci Resolve's Project Settings affects how Fusion compositions are integrated into the Edit page Timeline.

Changing the Resolution of a Clip

If your comp uses a single image, you can change the pixel output resolution in several ways. Three common tools that change the pixel resolution of a clip are the Resize, Scale, and Crop nodes. A fourth node, Letterbox, is less commonly used but also changes the pixel resolution of a clip.

These four nodes are located in the Transform category of the Effects library. Resize is also located in the toolbar.

- **Crop:** Sets the output resolution of the node using a combination of X and Y size along with X and Y offset to cut the frame down to the size you want. Crop removes pixels from the image, so if you later use a Transform node and try to move the image, those pixels are not available.
- **Letterbox:** Sets the output resolution of the node by adding horizontal or vertical black edges where necessary to format the frame size and aspect ratio.
- **Resize:** Sets the output resolution of the node using absolute pixels.
- **Scale:** Sets the output resolution of the node using a relative percentage of the current input image size.

TIP: To change resolution and reposition a frame without changing the pixel resolution of a clip, use the Transform node.

Compositing with Different-Resolution Clips

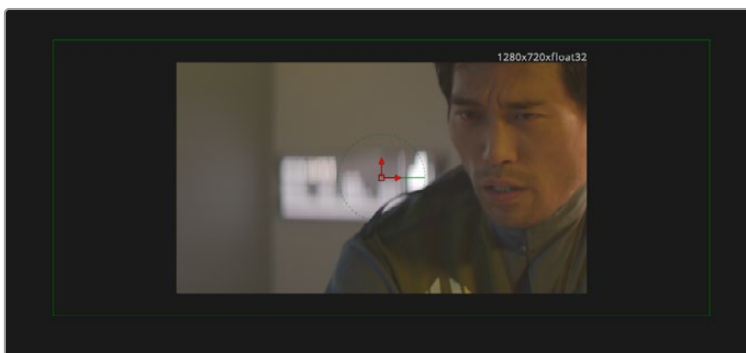
When you composite images with different resolutions using the Merge node, the image that's connected to the orange background input determines the output resolution of the Merge node.

Often, it's easiest to control the comp resolution right at the start by connecting a node with the desired output resolution you want to the orange background input on the Merge node. A Background node is often used in this situation because it consumes meager system resources.



A Background node determines the output resolution of the merge

The Background node sets the output size, and the foreground image is cropped if it is larger.



A Background node created at 1280 x 720 crops the larger foreground. However, all the pixels of the larger foreground are available for repositioning.

Sizing Between DaVinci Resolve Pages

The order of sizing operations between DaVinci Resolve pages is a bit more nuanced. It's important to understand which sizing operations happen in the Fusion page, and which happen after, so you know which effects alter the image that's input to the Fusion page, and which effects alter the page's output. For example, lens correction, while not strictly sizing, is nonetheless an effect that changes how the image begins in your Fusion composition. However, the Edit or Cut page stabilization function is an effect that comes after the Fusion page, so it does not appear in the composition you're creating.

The order of sizing effects in the different pages of DaVinci Resolve can be described as follows:



Sizing with Compound and Fusion Clips

Another way to modify the resolution before clips get handed off from the Edit page to the Fusion page is to create a compound clip or a Fusion clip. Both compound clips and Fusion clips change the working resolution of the individual clips to match the Timeline resolution. For instance, if two 4K clips are stacked one on top of the other in an HD timeline, creating a compound or Fusion clip resizes the clips to HD. The full resolution of the individual 4K clips is not available in Fusion and is therefore handed off to the Color page at the rescaled size. To maintain the full resolution of source clips, bring only one clip into the Fusion page from the Edit or Cut page Timeline, and then bring other clips into the Fusion composition using the Media Pool. Of course, if your clips are full HD and your timeline is full HD, then creating a Fusion clip or compound clip does not affect the resolution.

Color Bit Depths

The term bit depth describes how many colors are available in the color palette used to make up an image. The higher the bit depth, the greater the precision of color in the image, and therefore the greater the color reproduction. The higher precision is most apparent in gradients with subtle changes. Lower bit-depth gradients have noticeable banding artifacts, whereas higher bit-depth images can reproduce more colors, so fewer, if any, banding artifacts occur. The Fusion page within DaVinci Resolve always uses 32-bit float bits per channel precision to process images. However, in Fusion Studio you can choose to process images with 8-bit integer, 16-bit integer, 16-bit float, and 32-bit float bits per channel. Although always working at 16-bit float or 32-bit float will produce the best quality, it may be more efficient to use a lower bit depth if your images are 8-bit or 16-bit integer formats to begin with.


Understanding Integer vs. Float

Generally, 8-bit integer color processing is the lowest bit depth you'll come across for video formats. 8-bit images come from older or consumer-grade video equipment like mobile phones and camcorders. If you try to perform any significant gamma or color correction on 8-bit images, you can often see more visible banding.

16-bit integer color depth doubles the amount of precision, eliminating problems with banding. Although you can select 16-bit integer processing for an 8-bit clip, it does not reduce banding that already exists in the original file. Still, it can help when adding additional effects to the clip. This sounds like the best solution until you realize that many digital cameras like Blackmagic Design URSA Mini Pro and others record in formats that can capture over-range values with shadow areas below 0.0 and super highlights above 1.0, which are truncated in 16-bit integer.

The 16-bit float color depth sacrifices a small amount of the precision from standard 16-bit integer color depth to allow storage of color values less than 0 and greater than 1.0. 16-bit float, sometimes called half-float, is most often found in the OpenEXR format and contains more than enough dynamic range

for most film and HDR television purposes yet requires significantly less memory and processing time than is required for full float, 32-bit images.




Preserving over-range values allows you to change exposure while maintaining highlights

Processing at 32-bit float can work with shadow areas below 0.0 and highlights above 1.0, similar to 16-bit float, except with a much greater range of precision but also much greater memory and processing requirements.

Setting Color Depth in Fusion Studio

As we said earlier, DaVinci Resolve always processes at 32-bit float bits per channel; however, you can use less memory and still achieve more-than-acceptable results using the Performance Mode setting located in the User > Playback Preferences panel.

Fusion Studio automatically uses the color depth that makes the most sense for each file format. For example, if you read in a JPEG file from disk, then the color depth for the Loader is set to 8 bits per channel. Since the JPEG format is an 8-bit format, loading the image at a greater color depth would generally be wasteful. If a 16-bit TIFF is loaded, the color depth is set to 16 bits. Loading a DPX file defaults to 32-bit float, whereas OpenEXR generally defaults to 16-bit float. However, you can override the automatic format color depth using the settings found in the Import tab of the Loader node's Inspector. The Loader's Inspector, as well as the Inspector for images generated in Fusion (i.e., text, gradients, fast noise, and others), has a Depth menu for 8-bit, 16-bit integer, 16-bit float, and 32-bit float.

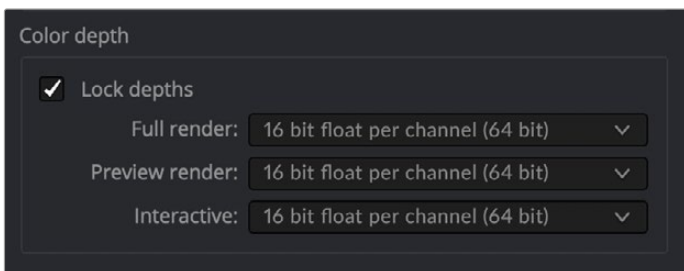


The Loader's Inspector Color Bit Depth settings

Configuring Default Color Depth Preferences

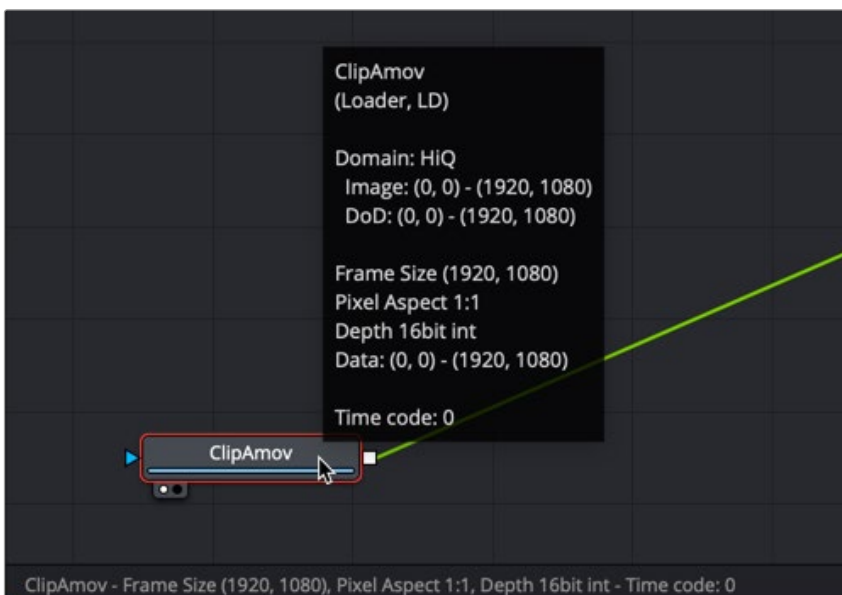
The default color depth setting forces the tool to process based on the settings configured in the Node Editor's Frame Format preferences. These are used to set a default value for color depth, applied when a Generator tool is added to the Node Editor. There are three drop-down menus to configure color depth in the preferences. They specify the different color depths for the interactive session, final renders, and preview renders.

To improve performance as you work on your comp, you can set the Interactive and Preview depth to 8-bits per channel, while final renders can be set to 16-bit integer. However, if your final render output is 16-bit float or 32-bit float, you should not use the integer options for the interactive setting. The final results may look significantly different from interactive previews set to integer options.



The Frame Format Color Depth settings

If you aren't sure what the color depth process is for a tool, you can position the pointer over the node's tile in the Node Editor, and a tooltip listing the color depth for that node will appear on the Status bar.



Hover over a node to view its Color Bit Depth setting.

TIP: When working with images that use 10-bit or 12-bit dynamic range or greater, like Blackmagic RAW or Cinema DNG files, set the Depth menu in the Inspector to 16-bit float or 32-bit float. This preserves highlight detail as you composite.

Combining Images with Different Color Depths

You can combine images with different color depths in a single composition. When images of different color depths are combined, the image from the background input of the node determines the bit depth output, and the foreground image is adjusted to match.

Advantages of Floating-Point Processing

There are two major advantages to floating-point processing that make the additional RAM requirements and longer render times worth your while. The first benefit is that floating-point values are more accurate than integer values. The second benefit is the preservation of shadow and highlight values that go beyond the normal tonal range.

Greater Accuracy

Using 16- or 32-bit floating-point processing prevents the loss of accuracy that can occur when using 8- or 16-bit integer processing. The main difference is that integer values cannot store fractional or decimal values, so rounding occurs in all image processing. Floating-point processing allows decimal or fractional values for each pixel, so it is not required to round off the values of the pixel to the closest integer. As a result, color precision remains virtually perfect, regardless of how many operations are applied to an image.

If you have an 8-bit pixel with a red value of 75 (dark red) and that pixel is halved using a Color Correction tool, the pixel's red value is now 37.5. Since you cannot store decimal or fractional values in integers, that value is rounded off to 37. Doubling the brightness of the pixel with another Color Correction tool should bring back the original pixel value of 75 but because of rounding 37×2 is 74. The red value lost a full point of precision due to integer rounding on a very simple example. This is a problem that can result in visible banding over several color corrections. Similar problems arise when merging images or transforming them. The more operations that are applied to an image, the more color precision is lost to rounding when using 8- or 16-bit integer processing.

Accessing Extended Highlights and Shadows

Increasingly more productions are capturing out-of-range images thanks to digital cinema cameras like the Blackmagic URSA Mini Pro and even the Pocket Cinema 6K camera. These cameras capture very high dynamic range RAW images and maintain color detail even in heavily over or underexposed frames. The extended white color detail can also give very nice, natural results when blurred, glowed, color corrected, or even just when faded or dissolved. While it is possible to work with these RAW images using integer data, doing so results in the loss of the extended range values, losing all detail in the highlights and shadows. Float processing makes working with logarithmic RAW images considerably easier by preserving highlight and shadow detail.

If you have an 8-bit pixel that has a red value of 200 (bright red) and a Color Gain tool is used to double the brightness of the red channel, the result is 200×2 , or 400. However, 8-bit color values are limited to a range of 0 through 255. So the pixel's value is clipped to 255, or pure red. If now the brightness is halved, the result is half of 255, or 127 (rounded), instead of the original value of 200.

When processing floating-point colors, pixel values brighter than white or darker than black are maintained. There is no value clipping. The pixel is still shown in the viewer as pure red, but if float processing is used instead of 8-bit, the second operation where the gain was halved would have restored the pixel to its original value of 200.

Using Float with 8-Bit HD Video


There is also some value to using float color depths with an 8-bit HD video when the images require a lot of color correction. Using float helps maintain precision by avoiding the rounding errors common to 8-bit processing, as described above.

Detecting Extended Highlight and Shadow Values

Although floating-point processing preserves extended values below 0.0 and greater than 1.0, also called "out-of-range values," the viewer still displays them as black or white. This can make it difficult for you to determine the overall dynamic range of an image.

To discover whether there are out-of-range values in a viewed image:

- Right-click in the viewer and choose Options > Show Full Color Range.



Use the Show Full Color Range pop-up menu to detect out-of-range images.


Enabling this display mode rescales the color values in the image so that the brightest color in the image is remapped to a value of 1.0 (white), and the darkest is remapped to 0.0 (black).

The 3D Histogram subview can also help visualize out-of-range colors in an image. For more information, see Chapter 67, “Using Viewers,” in the DaVinci Resolve Reference Manual, or Chapter 7 in the Fusion Reference Manual.

Clipping Out-of-Range Values

When processing in floating point, there may be situations where the out-of-range values in an image need to be clipped. The Brightness/Contrast tool provides checkboxes that can be used to clip out-of-range values to 0 or 1.

For example, there may be files that contain out-of-range alpha values. Since the alpha channel represents the opacity of a pixel, it makes little sense to be more than completely transparent or more than fully opaque, and compositing such an image may lead to unexpected results. To easily clip alpha values below 0 and above 1, add a Brightness/Contrast toolset to Clip Black and Clip White, with only the Alpha checkbox selected.



Clip White and Clip Black settings in Brightness/Contrast can be used to clip mattes.

Alternatively, you can clip the range by adding a Change Depth node and switching to 8-bit or 16-bit integer color depths.

Chapter 77

Managing Color for Visual Effects

This chapter discusses LUTs, color space conversions, and the value of compositing with linear gamma while previewing the image in the viewer using the gamma of your choice.

Contents


Color Management	1361
All Compositing Is Math	1361
Introducing Color Management in Fusion	1362
Converting to Linear Gamma	1362
Applying LUTs to a Viewer	1365
Using Resolve Color Management	1367
Using ACES Color Management in Resolve	1368
Using OCIO for ACES Color Management in Fusion	1369
Applying OCIO LUTs in the Viewer	1370

Color Management

The simplified goal of color management is to make sure that the image you see on your computer screen is what your audience sees when they view it on a television, cinema screen, or mobile device. For such a simple goal, problems arise that you, as the compositor, must manage. These problems begin with the fact that our eyes see luminance one way, and a computer display represents luminance differently.

Each capture device records images using a nonlinear tonal curve or gamma curve to compensate for this difference. Specifically, Rec. 709 HD gamma curves are designed so that when shown on HD displays, the images have built-in compensation for the display. The result is that HD images on HD displays appear normal to us.

Digital cinema cameras have taken the concept of gamma curves further. They use gamma curves as a way to maximize the bit depth of an image and store a wider dynamic range. Digital cinema cameras' gamma curve (often collectively referred to as *log gamma*), give more attention to the darker mid-tones where the human eye is most sensitive. This allows them to save images with brighter highlights and more detail in shadows.



A Rec. 709 HD gamma curve (left) and a nonlinear, or log gamma, curve (right)

The problem is that these images do not look normal on any monitor. Clips recorded with a log gamma curve typically have a low contrast, low saturated appearance when viewed on an sRGB computer display or Rec. 709 HD video monitor. This problem is easy to fix using a LookUp Table, or LUT. A LUT is a form of gamma and color correction applied to the viewer to normalize how the image is displayed on your screen.



A clip displayed with a nonlinear, log gamma curve (left) and corrected in the viewer using a LUT (right)

All Compositing Is Math

The LUT applied to the viewer only solves the problem in the viewer. Now we come to the larger problem. The image data is still using a log gamma curve. Fusion, and every other image-processing application, operates with the assumption that the image data has linear gamma. The image-processing filters you apply to images use standard math functions, like $1 + 1 = 2$. Common operations such as those that add pixels like Brightness, or divide pixels (a.k.a. “unpremultiply”), or composite


modes that include multiplication such as “screen,” and many other compositing tasks assume that $1 + 1$ always equals 2. In other words, if you perform an operation that doubles the amount of brightness, then every pixel should be twice as bright. However, if you are starting with a nonlinear gamma curve, pixels are not being adjusted linearly, so some pixels might end up 1.2 x as bright, 1.7 x as bright, or 2.4 x as bright. Now the math is $1 + 1 = 3$. The further your images are from linear gamma, the more pronounced the math error. A Rec. 709 HD clip shows less error than a log gamma clip from a digital cinema camera. However, an error is still an error, and the more compositing operations you perform on the image, the more the error is compounded.

You can see a more practical example when you apply filtering effects, such as a blur, to an image with any gamma setting. The image probably looks fine. However, if you convert the image to a linear gamma first and then apply the blur, the images (especially those with extremely bright areas) are processed with greater accuracy, and you should notice a different and superior result.

The answer to these problems is to manage your color before compositing.

Introducing Color Management in Fusion

Images loaded into Fusion by default are not color managed. The image is displayed directly from the file to the viewer without any interpretation or conversion. However, Fusion includes nodes that convert the output of each image to linear gamma at the beginning of your composite. The same nodes can convert from linear back to your desired output gamma at the end of your composite, just prior to the Saver or MediaOut node.



A log clip converted to linear and then converted for output

To manually set up a linear gamma workflow in Fusion:

- 1 Use a Gamut or CineonLog node after all MediaIn or Loader nodes to convert them to linear.
- 2 Apply a GAMUT View LUT to the viewers to correct the display of a linear image to sRGB or Rec. 709.
- 3 Before a Saver or MediaOut node, insert a Gamut or CineonLog node to convert from linear to your target output format.


Converting to Linear Gamma

Whether an image comes from the Edit page in DaVinci Resolve, or from a Loader in Fusion Studio, the color and gamma are read directly into Fusion, with no modification. For some simple operations on sRGB or Rec 709 clips, this may be fine, but it's not always the ideal way to work, especially for log-encoded media. The ideal way to work with log-encoded media is to convert images to linear gamma, since the majority of image-processing operations in Fusion expect gamma to be linear and will produce superior results.

TIP: 3D rendered CGI images are often generated as EXR files with linear gamma, and converting them is not necessary. However, you should check your specific files to make sure they are using linear gamma.


Fusion includes several kinds of nodes you can use to convert the image out of each MediaIn or Loader node to linear gamma at the beginning of your composite, and then convert from linear back to your desired output gamma at the end of your composite. These include:

- **CineonLog node:** The CineonLog node, found in the Film category of the Effects Library, performs a conversion from any of the formats in the Log Type menu to linear, and also reverses the process, adding log gamma back to a clip. This is most often used for images coming from common digital cinema cameras like BlackMagic Design, Arri, or Red. The CineonLog node is added directly after a MediaIn or Loader node. The Mode menu chooses the direction of the conversion to or from linear.



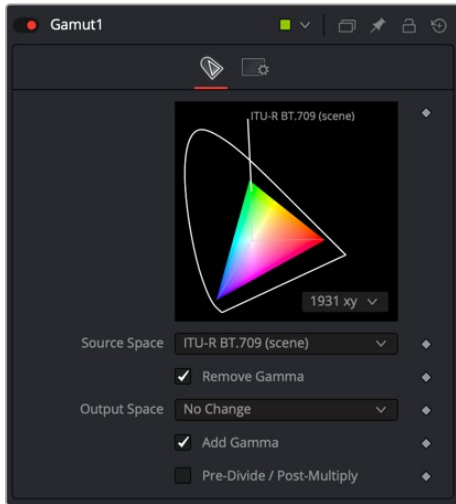
Add a CineonLog tool to convert log gamma curves to linear and vice versa.

- **Gamut node:** The Gamut node, found in the Color category of the Effects Library, lets you perform linear conversions based on color space. This node converts to linear or from linear and is often inserted after a MediaIn or Loader node or just before a MediaOut or Saver node. Depending on where you insert the node, you either choose from the Source Space controls or the Output Space controls.



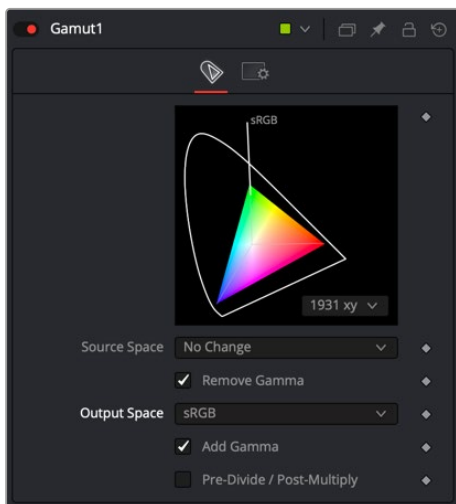
Add a Gamut tool to convert gamma curves to linear based on color space.

When converting media to linear gamma, set the Source Space menu to the color space of your source material. For instance, if your media is full 1080 HD ProRes, then choose ITU-R BT.709 (scene) for gamma of 2.4. Then, enable the Remove Gamma checkbox if it isn't already enabled, to use linear gamma.




Source Space is used to convert to linear gamma.

When converting from linear gamma for output, you insert the Gamut node before your output node, which is a Saver in Fusion Studio or a MediaOut node in DaVinci Resolve's Fusion page. Make sure the Source Space menu is set to No Change, and set the Output Space to your output color space. For instance, if your desired output is full 1080 HD, then choose either sRGB or ITU-R BT.709 (scene) for gamma of 2.4. Then, enable the Add Gamma checkbox if it isn't already enabled, to format the output of the Gamut node for your final output.



Output Space is used to convert from linear gamma.

- **Medialn and Loader nodes:** Medialn and Loader nodes have Source Gamma Space controls in the Inspector that let you identify and remove the gamma curve without the need to add another node. If your files include gamma curve metadata like RAW files, the Auto setting for the Curve Type drop-down menu reads the metadata and uses it when removing the gamma curve. When using intermediate files or files that do not include gamma curve metadata, you can choose either a log gamma curve by choosing Log from the Curve Type menu or a specific color space using the Space option from the menu. Clicking the Remove Curve checkbox then removes the gamma curve, converting the image to linear gamma.



MediaIn and Loader nodes include a Remove Curve checkbox in the Inspector.

- **FileLUT node:** The FileLUT node, found in the LUT category of the Effects Library, lets you do a conversion using any LUT you want, giving you the option to manually load LUTs in the ALUT3, ITX, 3DL, or CUBE format to perform a gamma and gamut conversion. Although LUTs are very commonly placed at the end of a node tree for final rendering, you'll get more accurate gamma and color space conversions using the Gamut and CineonLog nodes to transform your MediaIn and Loader nodes into linear.

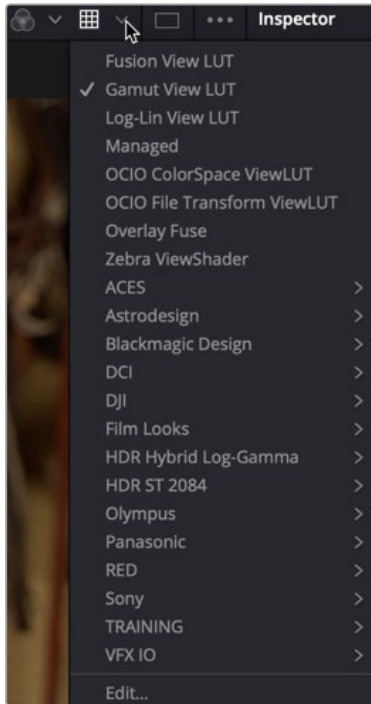
Applying LUTs to a Viewer

Images converted to a linear gamma don't look correct. They usually look very dark, with extremely bright highlights and oversaturated colors. Happily, even though the image may appear to be incorrect, the fact that Fusion can work entirely with floating-point color data means that you're not actually clipping or losing any image data. It just looks completely wrong when viewing the linear state of your image data directly.



A clip displayed with a nonlinear, log gamma curve (left) and the clip transformed to linear gamma (right)


It would be impossible to work if you couldn't view the image as it's supposed to appear within the final gamut and gamma you'll be outputting. For this reason, each viewer has a LUT menu that lets you enable a "preview" color space and/or gamma conversion, while the node tree is processing correctly in linear gamma.



Apply a Gamut View LUT to preview the image in your intended output color space.

To preview the images in the viewer using sRGB or Rec. 709 color space:

- 1 Enable the LUT button above the viewer.
- 2 From the Viewer LUT drop-down menu, choose either a Gamut View LUT, or a LUT from the VFX IO category that transforms linear to Rec. 709 or sRGB.
- 3 If you choose the Gamut View LUT, then choose Edit from the bottom of the LUT menu to configure the LUT.
- 4 In the LUT Editor, set the Output Space to the target color space you want.
- 5 Enable the Add Gamma checkbox to apply the gamma curve based on the selected color space.



Gamut View LUT Output Space set to Rec. 709 for HD with the Add Gamma checkbox enabled

TIP: If your monitor is calibrated differently, you need to select a LUT that matches your calibration.

Whether you use the Gamut View LUT or a LUT for your specific monitor calibration, you can save the viewer setup as the default.

To Save the Gamut LUT setup as the default viewer setup:

- Right-click in the viewer, and then choose Settings > Save Defaults.

For every comp, the viewer will now be preconfigured based on the saved defaults.

For more information on Viewer LUTs, see Chapter 67, “Using Viewers,” in the DaVinci Resolve Reference Manual or Chapter 7 in the Fusion Reference Manual.


Using Resolve Color Management

If you’re using the Fusion page within Davinci Resolve, you have the option of enabling DaVinci Resolve’s scene-referred color management, instead of inserting Gamut and CineonLog nodes. When DaVinci YRGB Color Managed (RCM) is enabled, the color of MediaIn nodes in the Fusion page is handled differently. The RCM Timeline color space setting is used for all MediaIn nodes in the Fusion page, which then automatically get converted to linear gamma. The MediaOut node then gets converted back into the Timeline color space to get graded in the Color page or further edited in the Edit page Timeline.

Despite the seeming complexity of color management, using RCM is actually simple. In essence, all you have to do is (A) turn on RCM, (B) assign the appropriate Input Color Space to each clip in the Media Pool based on its recording device, and (C) choose the Timeline and Output Color Space combination you want to use.

To enable Resolve Color Management:

- 1 Open the Color Management panel of the Project Settings.
- 2 Choose DaVinci YRGB Color Managed from the Color Science drop-down menu.
- 3 Choose the Input Color Space for the clips in the Media Pool.
- 4 Set the desired options for the Timeline Color Space and Output Color Space RCM settings.



When DaVinci YRGB Color Managed is enabled, Timeline color space is used for all MediaIn nodes in the Fusion page.

To override the input color space for differently recorded clips in the Media Pool:

- 1 Enable DaVinci YRGB Color Management as explained above.
- 2 Save and close the Settings dialog.
- 3 In the Media Pool, select the clip or clips you want to assign a new Input Color space.
- 4 Right-click one of the selected clips.
- 5 Choose the Input Color Space that corresponds to those clips from the contextual menu.

Using RCM eliminates a few steps, since the input color space math used to transform the source preserves all wide-latitude image data, making highlights easily retrievable without any extra steps. With RCM enabled, there is no need to insert CineonLog or Gamut nodes while in the Fusion page. The transforms from and to linear are done automatically based on the RCM settings. Switching to the Fusion page converts the images to linear and enables the LUT button in the viewers with the Managed LUT selected. The Managed LUT uses the RCM settings to take a linear image and display it based on the RCM output color space.

For more information on Resolve Color Management, see Chapter 9, “Data Levels, Color Management, and ACES,” in the DaVinci Resolve Reference Manual.

Using ACES Color Management in Resolve


The ACES (Academy Color Encoding Specification) color space is another standard for managing color throughout an entire production. It’s designed to make start-to-finish, scene-referred color management a reality for digital cinema workflows. Just like DaVinci’s RCM, ACES makes it easier to extract high-precision, wide-latitude image data from raw camera formats, to preserve high-quality image data from acquisition through the color grading process and to output high-quality data for broadcast viewing, film printing, or digital cinema encoding.

ACES works by assigning an IDT (Input Device Transform) to every camera and acquisition device. The IDT specifies how media from that device is converted into the ACES color space. At the end of the pipeline, an ODT (Output Device Transform) is applied to convert the image data from ACES color space into the gamut of your final output.

Similar to setting up RCM, DaVinci Resolve’s color management project settings can be configured for ACES, which carries through the Edit, Fusion, and Color pages.

NOTE: When using Fusion Studio, the OpenColorIO (OCIO) framework is used for ACES color management.

The Color Science drop-down menu in the Color Management panel of the Project Settings is used to set up the ACES color management in DaVinci Resolve.



When ACES is enabled, IDT and ODT are used to identify input and output devices.

- **Color Science:** Using this drop-down menu, you can choose either ACEScct or ACEScc color science. This is primarily a personal preference since they are mostly identical, but the shadows respond differently to grading operations. In the Fusion page, images are automatically converted to linear, so whoever does the grading has more of a reason to choose one or the other.
 - **ACEScc:** Choose ACEScc color science to apply a standard Cineon style log encoding to the ACES data before it is processed by DaVinci Resolve.
 - **ACEScct:** This variation of ACEScc adds a roll-off at the toe of the image to make color correction lift operations “feel” more like they do with film scans and LogC encoded images.
- **ACES Version:** When you’ve chosen one of the ACES color science options, this menu becomes available to let you choose which version of ACES you want to use. As of DaVinci Resolve 16, you can choose either ACES 1.0.3 or ACES 1.1 (the latest version).
- **ACES Input Device Transform:** This menu lets you choose which IDT (Input Device Transform) to use for the dominant media format in use.
- **ACES Output Device Transform:** This menu lets you choose an ODT (Output Device Transform) with which to transform the image data to match your required deliverable.
- **Process Node LUTs In:** This menu lets you choose how you want to process LUTs in the Color page and does not affect the Fusion page.

For more information on ACES within Davinci Resolve, see Chapter 9, “Data Levels, Color Management and ACES,” in the DaVinci Resolve Reference Manual.

Using OCIO for ACES Color Management in Fusion


When using Fusion Studio or not using color management in DaVinci Resolve, you have the option to use OpenColorIO nodes in Fusion to composite in an ACES color space.

OpenColorIO (OCIO) is an open-source color management framework for visual effects and computer animation. OCIO is compatible with the Academy Color Encoding Specification (ACES). Three OCIO

nodes located in the Color category of the Effects Library allow you to use OCIO color space transforms in Fusion.

- OCIO CDL Transform node allows you to create, save, load, and apply a Color Decision List (CDL) grade.
- OCIO Color Space allows sophisticated color space conversions based on an OCIO config file.
- OCIO File Transform allows you to load and apply a variety of LookUp Tables (LUTs).

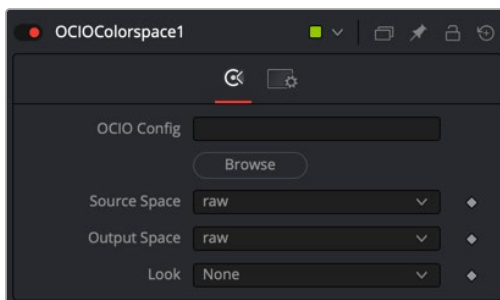
Using OCIO for converting MediaIn or Loader nodes to linear gamma is based on the OCIO Color Space node. Placing the OCIO Color Space node directly after a Loader (or MediaIn in DaVinci Resolve) displays the OCIO Source and Output controls in the Inspector.



OCIO Color Space nodes can be used to work in an ACES color managed environment.

Within the Inspector for OCIO Color Space node, Fusion includes default Source and Output transforms for standard color spaces. However, to use the full OCIO standard, you'll need to download and install the OCIO config file. You can download the config file from the OCIO website. <https://opencolorio.org>

Clicking the Browse button in the Inspector will allow you to navigate to the downloaded config file. From the download, locate the ACES 1.0.3 or later folder and select the file config.ocio.




Source and Output menus are populated based on the config.ocio file that you download.

The Source menu is used to choose the color profile for your Loader or MediaIn node. The default raw setting shows an unaltered image, essentially applying no color management to the clip. The selection you make from the menu is based on the recording profile of your media.

The Output menu is set based on your deliverables. When working in Fusion Studio, typically the Output selected is ACEScg, to work in a scene linear space.

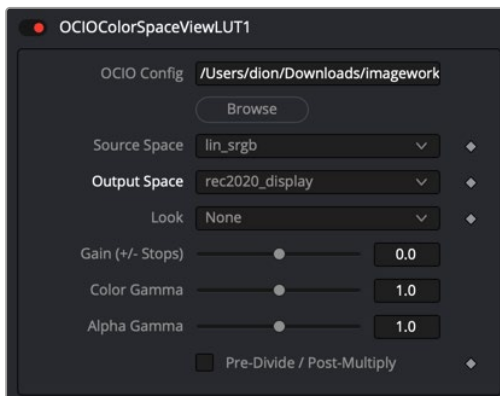
Applying OCIO LUTs in the Viewer

The viewer also includes OCIO View LUTs to calibrate the viewers. Once the OCIO Color Space View LUT is selected from the LUT menu above the viewer, choosing Edit for the same menu opens a dialog where you can load the OCIO config file.



The OCIO Color Space View LUT is located in the Viewer LUT menu in Fusion Studio.

By default, the same standard options are available in the View LUT. However, clicking the Browse button allows you to load the same config file you loaded into the OCIO Color Space node. Once loaded, all the expanded OCIO options are available. If you selected the OCIO Color Space node to output ACEScg, you use the OCIO View LUT to go from a source setting of linear sRGB to an output setting of sRGB or Rec. 709 in most cases.



The OCIO Color Space View LUT dialog is used to configure the viewer when using the OCIO Color Space node in the Node Editor.

TIP: If your monitor is calibrated differently, you will need to select a LUT that matches your calibration.

To set the LUT using OCIO:

- 1 Click the LUT menu and choose the OCIO Color Space View LUT.
- 2 From the same menu, select Edit.
- 3 In the View LUT editor that opens, set the source's color space to lin sRGB
- 4 Set the output space to sRGB or REC 709, assuming you are viewing on a standard computer monitor. You now see a normalized image in the viewer, but all color operations will be on linear images.

Whether you use the OCIO Color Space LUT or a LUT for your specific monitor calibration, you can save the viewer setup as the default.

To save the OCIO ColorSpace LUT setup as the default viewer setup:

- Right-click in the viewer, and then choose Settings > Save Defaults. Now, for every comp, the viewer is preconfigured based on the saved defaults.

Chapter 78

Understanding Image Channels

This chapter seeks to demystify how Fusion handles image channels and, in the process, show you how different nodes need to be connected to get the results you expect. It also explains the mysteries of premultiplication, and presents a full explanation of how Fusion is capable of using and even generating auxiliary data.

Contents

Channels in Fusion	1373
Types of Channels Supported by Fusion	1373
Fusion Node Connections Carry Multiple Channels	1374
Node Inputs and Outputs	1374
Node Colors Tell You Which Nodes Go Together	1377
Using Channels in a Composition	1379
Channel Limiting	1379
Adding Alpha Channels	1380
How Channels Propagate During Compositing	1381
Rearranging or Combining Channels	1382
Understanding Premultiplication	1382
The Rules of Premultiplication	1385
Alpha Channel Status in MediaIn and Loader Nodes	1387
Controlling Premultiplication in Color Correction Nodes	1387
Controlling Premultiplication With Alpha Divide and Alpha Multiply	1387
Multi Channel Compositing	1388
Compositing with Beauty Passes	1388
Working with Auxiliary Channels	1393
Auxiliary Channels Explained	1395
Propagating Auxiliary Channels	1402
Nodes That Use Auxiliary Channels	1403
Image Formats That Support Aux Channels	1404
Creating Auxiliary Channels in Fusion	1404

Channels in Fusion

Fusion introduces some innovative ways of working with the many different channels of image data that modern compositing workflows encompass. This chapter's introduction to color and data channels and how they're affected by different nodes and operations is a valuable way to begin the process of learning to do paint, compositing, and effects in Fusion.

If you're new to compositing, or you're new to the Fusion workflow, you ignore this chapter at your peril, as it provides a solid foundation to understanding how to predictably control image data as you work in this powerful environment.

Types of Channels Supported by Fusion

Digital images can be divided into separate streams of image data called Channels, each of which carries a specific kind of image data. Nodes that perform different image processing operations typically expect specific channels to provide predictable results. You are probably familiar with the three standard color channels of red, green, and blue, but there are many others. This section describes the different kinds of channels that Fusion supports.

RGB Color channels

The red, green, and blue channels of any still image or movie clip combine additively to represent everything we can see via visible light. Each of these three channels is a grayscale image when seen by itself. When combined additively, these channels represent a full-color image.

Alpha Channels

An alpha channel is an embedded fourth channel that defines different levels of transparency in an RGB image. Alpha channels are typically embedded in RGB images that are generated from computer graphics applications. In Fusion, white denotes solid areas, while black denotes transparent areas. Grayscale values range from more opaque (lighter) to more transparent (darker).

If you're working with an imported alpha channel from another application for which these conventions are reversed, never fear. Every node capable of using an alpha channel is also capable of inverting it.

Single-Channel Masks

While similar to alpha channels, mask channels are single channel images, external to any RGB image and typically created by Fusion within one of the available Mask nodes. Mask nodes are unique in that they propagate single-channel image data that defines which areas of an image should be solid and which should be transparent. However, masks can also define which parts of an image should be affected by a particular operation, and which should not. Mask channels are designed to be connected to specific mask inputs of nodes including Effect Mask, Garbage Mask, and Solid Mask inputs.

Auxiliary Channels

Auxiliary channels (covered in more detail later in this chapter), describe a family of special-purpose image data that typically expose 3D data in a way that can be used in 2D composites. For example, Z-Depth channels describe the depth of each pixel in an image along a Z axis, while an XYZ Normals channel describes the orientation (facing up, facing down, or facing to the left or right) of each pixel in an image. Auxiliary channel data is generated by rendering 3D images, so it usually accompanies or is embedded with RGB images generated by 3D modeling and animation applications. These channels can also be generated from within Fusion via the Renderer 3D node, which outputs a 3D scene that you've assembled and lit as 2D RGBA channels, with optionally accompanying auxiliary channels.


The reason to use auxiliary data is that 3D rendering is computationally expensive and time-consuming, so outputting descriptive information about a 3D image that's been rendered empowers compositing artists to make sophisticated alterations in 2D. You can add motion blur, perform relighting, and composite with depth information faster than re-rendering the 3D source material over and over.

TIP: You can view any of a node's channels in isolation using the Color drop-down menu in the viewer. Clicking the Color drop-down menu reveals a list of all channels within the currently selected node, including red, green, blue, or auxiliary channels.

Fusion Node Connections Carry Multiple Channels

The connections that pass image data from one node to the next in Fusion's Node Editor are capable of carrying multiple channels of image data. That means that a single connection may route RGB, or RGBA, or RGBA-Z-Depth, or even just Z-Depth, depending on the Input you connect to and the node's function.

In the following example, the two MediaIn nodes output RGB data. However, the Delta Keyer creates an alpha channel and combined it with MediaIn2's RGB image. The RGB-A of the Delta Keyer becomes the foreground image that the Merge node can use to create a two-layer composite.




The alpha channel generated by the Delta key is used for compositing by the foreground input of the Merge node.

NOTE: Node trees shown in this chapter may display MediaIn nodes found in DaVinci Resolve's Fusion page; however, Fusion Studio Loader nodes are interchangeable unless otherwise noted.

Running multiple channels through single connection lines makes Fusion node trees simple to read, but it also means you need to keep track of which nodes process which channels to make sure that you're directing the intended image data to the correct operations.

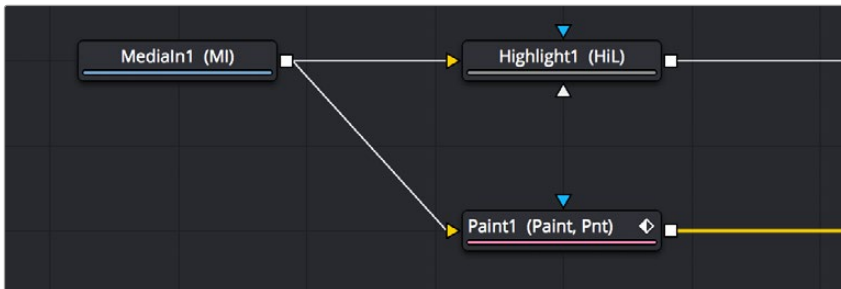
Node Inputs and Outputs

By default, MediaIn nodes in the Fusion page and Loader nodes in Fusion Studio output RGBA channels. When you connect one node's output to another node's input, the active channels are passed from the upstream node to the downstream node, which then processes the image according to that node's function. Only one node output can be connected to a node input at a time. In this simple example, a MediaIn node's output is connected to the input of a Highlight node to create a sparkly highlight effect.



MediaIn node connected to a Highlight node connected to MediaOut node in the Fusion page.

When connecting nodes, a node's output carries the same channels no matter how many times the output is "branched." You cannot send one channel out on one branch and a different channel out on another branch.




The Medialn node's output is branched but carries the same RGB channels to both inputs.

Using Multiple Inputs

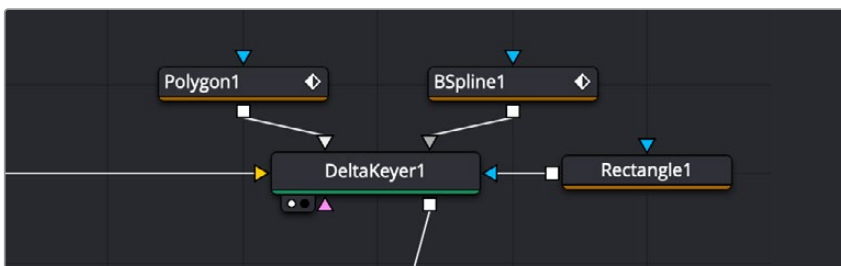
Most nodes have two inputs, one for RGBA and another for an effect mask that can be optionally used to limit the effect of that node to a particular part of the image. However, some nodes have three or even more inputs, and it's important to make sure you connect the correct image data to the appropriate input to obtain the desired result. If you connect one node's output to another node's input and nothing happens, chances are you've connected to the wrong input.

For example, the MatteControl node has a background input and a foreground input, both of which accept RGBA channels. However, it also has SolidMatte, GarbageMatte, and EffectsMask inputs that accept alpha or mask channels to modify the transparency of the Node's output. If you want to perform the extremely common operation of using a MatteControl node to create an alpha channel using a Polygon node for rotoscoping an image, you need to make sure that you connect the Polygon node to the GarbageMatte input to obtain the correct result. The GarbageMatte input is automatically set to alter the alpha channel of the foreground image. If you connect to any other input, your Polygon mask may not produce expected results.




Polygon node connected to the GargabeMatte input of a MatteControl node for rotoscoping

In another example, the DeltaKeyer node has a primary input (labeled "Input") that accepts RGBA channels, but it also has three Matte inputs. These SolidMatte, GarbageMatte, and EffectsMask inputs on the Delta Keyer accept alpha or mask channels to modify the matte being extracted from the image in different ways.



The DeltaKeyer combines the multiple mask nodes in different ways..

If you position your pointer over any node's input or output, a tooltip appears in the Tooltip bar at the bottom of the Fusion window, letting you know what that input or output is for, to help guide you to using the right input for the job. If you pause for a moment longer, another tooltip appears in the Node Editor itself.



(Left) The node input's tooltip in the Tooltip bar, (Right) The node tooltip in the Node Editor

Connecting to the Correct Input

When you're connecting nodes, pulling a connection line from the output of one node and dropping it right on top of the body of another node makes a connection to the default input for that node, which is typically the "Input" or "Background" input.



Side by side, dropping a connection on a node's body to connect to that node's primary input

However, if you drop a connection line right on top of a specific input, then you'll connect to that input, so it's important to be mindful of where you drop connection lines as you wire up different node trees together.



Side by side, dropping a connection on a specific node input, note how the inputs rearrange themselves afterwards to keep the node tree tidy-looking

TIP: If you hold the Option key down while you drag a connection line from one node onto another, and you keep the Option key held down while you release the pointer's button to drop the connection, a menu appears that lets you choose which specific input you want to connect to, by name.

Inputs Require Specific Channels

Usually, you're prevented from connecting a node's output to another node or node input that's not compatible with it. For example, if you try to connect a Text3D node's output directly to the input of a regular Merge node, it won't work; 3D nodes do not produce RGB images, they produce 3D geometry data, so you must first connect to a Renderer3D node that creates the RGB output appropriate for 2D compositing operations.

In other cases, connecting the wrong image data to the wrong node input won't give you any error, it simply fails to produce the result you were expecting, necessitating you to troubleshoot the composition. If this happens to you, check the Fusion Effects section of this manual to see if the node you're trying to connect to has any limitations as to how it must be attached.

TIP: This chapter tries to cover many of the easy-to-miss exceptions to node connection that are important for you to know, so don't skim too fast.

Always Connect the Background Input First

Many nodes combine images in different ways, using “background” and “foreground” inputs. This includes the Merge node, the Matte Control node, and the Channel Booleans node, to cite common examples. The color of inputs on a node can help you make the right corrections. For instance, background inputs are always orange, and foreground inputs are always green.

When you first connect any node's output to a multi-input node, you usually want to connect the background input first. This is handled for you automatically when you first drop a connection line onto the body of a new multi-input node. The orange-colored background input is almost always connected first (the exception is Mask nodes, which always connect to the first available Mask input). This is good because you want to get into the habit of always connecting the background input first.

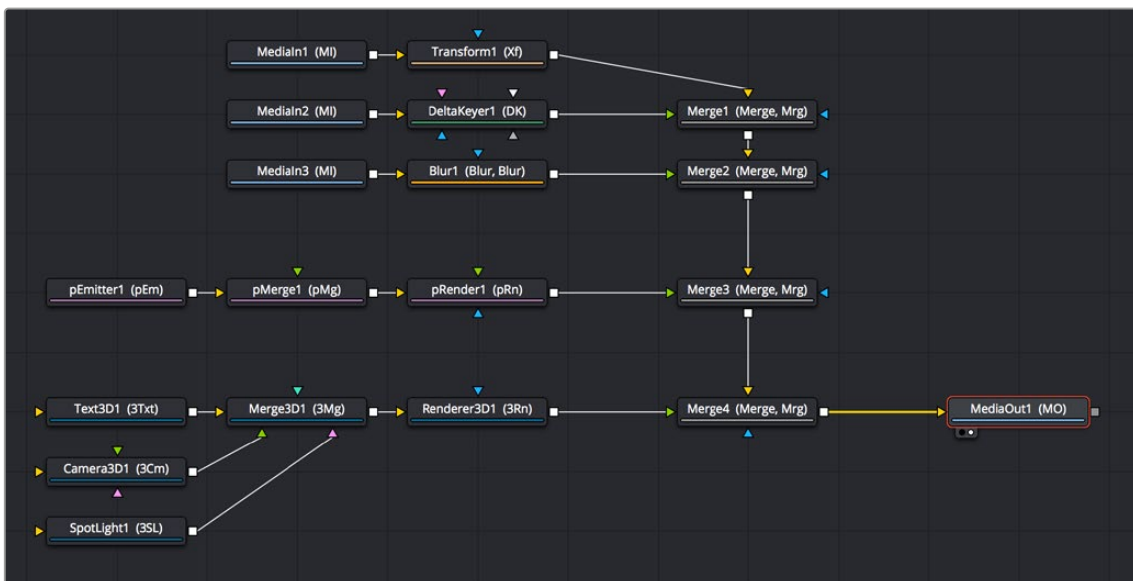
If you connect to only one input of a multi-input node and you don't connect to the background input, you may find that you don't get the results you wanted. This is because each multi-input node expects that the background is connected before anything else so that the internal connections and math used by that node can be predictable.

TIP: The only node to which you can safely connect the foreground input prior to the background input is the Dissolve node, which is a special node that can be used to either dissolve between two inputs, or automatically switch between two inputs of unequal duration.

Node Colors Tell You Which Nodes Go Together

Each node in Fusion accomplishes a single type of effect or operation. These single-purpose nodes make it easier to decipher a complex composition when examining its node tree. Single-purpose nodes also make it easier to focus on fine-tuning specific adjustments, one at a time, when assembling an ever-growing tree.

Because each Fusion node has a specific function, they're categorized by type to make it easier to keep track of which nodes require what types of image channels as input, and what image data you can expect each node to output. These general types are described here.



A node tree showing the main categories of node colors

Blue MediaIn and Loader Nodes, and Green Generator Nodes

Blue MediaIn nodes and blue Loader nodes add clips to a composite, and green Generator nodes create images. Both types of nodes output RGBA channels (depending on the source and generator), and may optionally output auxiliary channels for doing advanced compositing operations.

Because these are sources of images, both kinds of nodes can be attached to a wide variety of other nodes for effects creation besides just 2D nodes. For example, you can also connect MediaIn nodes to Image Plane 3D nodes for 3D compositing, or to pEmitter nodes set to “Bitmap” for creating different particle systems. Green Generator nodes can be similarly attached to many different kinds of nodes, for example, attaching a FastNoise node to a Displace 3D node to impose undulating effects to 3D shapes.

Shape nodes are also green, although they must be attached to a specialized set of gray modifier and render nodes (all of which begin with the letter “s” and appear in the Shape category of the Effects Library).

2D Processing Nodes, Color Coded by Type

These encompass most 2D processing and compositing operations in Fusion, all of which process RGBA channels and pass along auxiliary channels. These include:

- Orange Blur nodes
- Olive Color Adjustment nodes (color adjustment nodes additionally concatenate with one another)
- Pink Paint nodes
- Dark orange Tracking nodes
- Tan Transform node (transform nodes additionally concatenate with one another)
- Teal VR nodes
- Dark brown Warp nodes
- Gray, which includes Compositing nodes as well as many other types.

Additionally, some 2D nodes such as Fog and Depth Blur (in the Deep Pixel category) accept and use auxiliary channels such as Z-Depth to create different perspective effects in 2D.

TIP: Two 2D nodes that specifically don't process alpha channel data are the Color Corrector node and the Gamut node. The Color Correction node lets you color correct a foreground layer to match a background layer without affecting an alpha channel. The Gamut node lets you perform color space conversions to RGB data from one gamut to another without affecting the alpha channel.

Purple Particle System Nodes

These are nodes that connect to create different particle systems, and they're incompatible with other kinds of nodes until you add a pRender node that outputs 2D RGBA and auxiliary data that can be composited with other 2D nodes and operations.

Dark Blue 3D Nodes

These are 3D operations, which generate and manipulate 3D data (including auxiliary channels) that are incompatible with other kinds of nodes until processed via a Renderer 3D node, which then outputs RGBA and auxiliary data.


Brown Mask Nodes

Masks output single-channel images that can only be connected to one another (to combine masks) or to specified Mask inputs. Masks are useful for defining transparency (Alpha masks), defining which parts of an image should be cropped out (Garbage masks), or defining which parts of an image should be affected by a particular node operation (Effects masks).

Using Channels in a Composition

When you connect one node's Output to another node's Input, you feed all of the channels that are output from the upstream node to the downstream node. 2D nodes, which constitute most simple image processing operations in Fusion, propagate all channel data from node to node, including RGB, alpha, and auxiliary channels, regardless of whether or not that node actually uses or affects a particular channel.

2D nodes also typically operate upon all channel data routed through that node. For example, if you connect a node's output with RGBA and XYZ Normals channels to the input of a Vortex node, all channels are equally transformed by the Size, Center, and Angle parameters of this operation, including the alpha and XYZ normals channels, as seen in the following screenshot.



(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, do not affect auxiliary channels. This makes them convenient for color-matching foreground and background layers you're compositing, without worrying that you're altering the depth information accompanying that layer.

TIP: If you're doing something exotic and you actually want to operate on a channel that's usually unaffected by a particular node, you can always use the Channel Booleans node to reassign the channel. When doing this to a single image, it's important to connect that image to the background input of the Channel Booleans node, so the alpha and auxiliary channels are properly handled.

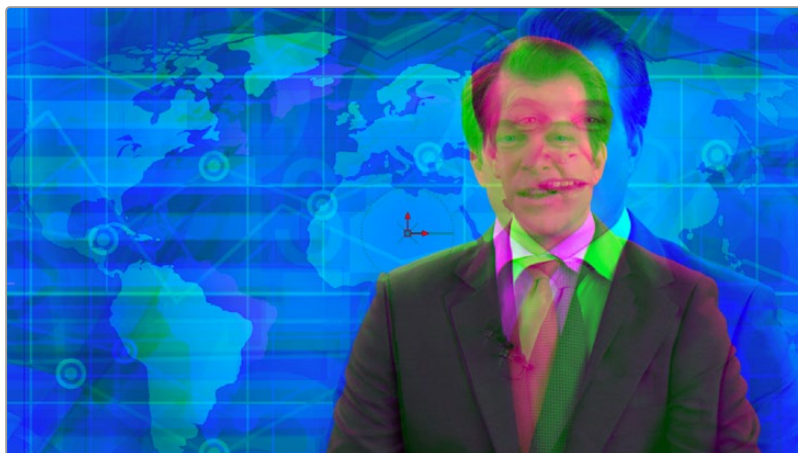
Channel Limiting

Most nodes have a set of Red, Green, Blue, and Alpha buttons in the Settings tab of that node's Inspector. These buttons let you exclude any combination of these channels from being affected by that node.



The channel limiting buttons in the Settings panel of a Transform node, so only the Green channel is affected

For example, if you wanted to use the Transform node to affect only the green channel of an image, you can turn off the Green, Blue, and Alpha buttons. As a result, the green channel is processed by this operation, and the red, blue, and alpha channels are copied straight from the node's input to the node's output, skipping that node's processing to remain unaffected.



Transforming only the green color channel of the image with a Transform effect

Skipping Channel Processing


Under the hood, most nodes process all channels first, but afterward copy the input image to the output for channels that have been enabled. Modern workstations are so fast that this isn't usually noticeable, but there are some nodes where deselecting a channel actually causes that node to skip processing that channel entirely. Nodes that operate this way have a linked set of Red, Green, Blue, and Alpha buttons on another tab in the node. In these cases, the Common Control channel buttons are instantiated to the channel buttons found elsewhere in the node.

Blur, Brightness/Contrast, Erode/Dilate, and Filter are examples of nodes that all have RGBA buttons in the main Controls tab of the Inspector, in addition to the Settings tab.

Adding Alpha Channels

Much of visual effects compositing has to do with placing a foreground subject over a background. Possibly the most fundamental method is through the use of an alpha or matte channel. If an alpha channel is not contained within the clip, you add one via keying or roto-scoping. While more specific methods are covered in detail in later chapters, here's an example of how this is handled within Fusion.


In the case of extracting an alpha matte from a green screen image, you typically connect the image's RGB output to the "Input" input of a Keyer node such as the Delta Keyer, and you then use the keyer's controls to extract the matte. The Keyer node automatically inserts the alpha channel that's generated alongside the RGB channels, so the output is automatically RGBA. Then, when you connect the keyer's output to a Merge node to composite it over another image, the Merge node automatically knows to use the embedded alpha channel coming into the foreground input to create the desired composite, as seen in the following screenshot.



A simple node tree for keying; note that only one connection links the DeltaKeyer to the Merge node

Rotoscoping, or manually drawing a mask shape using a Polygon or other Mask node is another technique used to create the matte channel. There are many ways to configure the node tree for this task, but the simplest setup is just to connect a Polygon or B-Spline mask node to the Effect Mask input of a MediaIn or Loader node.

TIP: When rotoscoping, it is best to leave the Mask node disconnected from the image while you draw the shape. This allows you to view the MediaIn node while drawing. Connect the Matte node once you have finished drawing the shape.

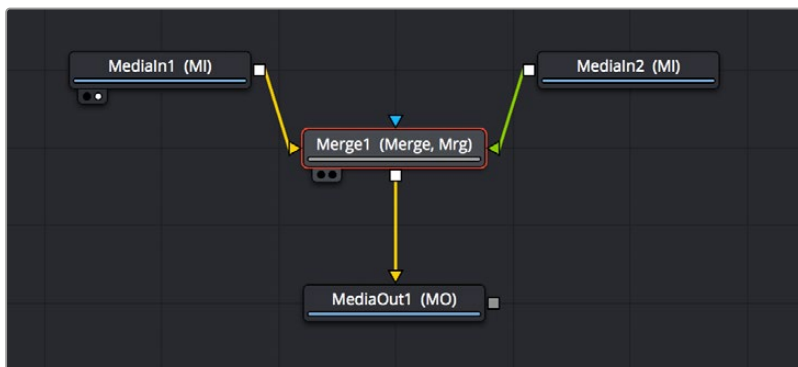


A simple rotoscoping setup using a Polygon node directly connected to the Effect Mask input of a MediaIn.

In both cases, you can see how the node tree’s ability to carry a single channel or multiple channels of image data over a single connection line simplifies the compositing process.

How Channels Propagate During Compositing

Images are combined or composited together using the Merge node. The Merge node takes two RGBA inputs labeled “Foreground” (green) and “Background” (orange) and combines them into a single RGB output (or RGBA if both the foreground and background input images have alpha), where the foreground image is in front (or on top, depending on what you’re working on), and the background image is, you guessed it, in back.



A simple Merge node composite

Auxiliary channels, on the other hand, are handled in a much more specific way. When you composite two image layers using the Merge node, auxiliary channels only propagate through the image that's connected to the background input. The rationale for this is that in most CGI composites, the background is most often the CG layer that contains auxiliary channels, and the foreground is a live-action green screen plate.

Since most compositions use multiple Merge nodes, it pays to be careful about how you connect the background and foreground inputs of each Merge node to make sure that the correct channels flow properly.

TIP: Merge nodes are also capable of combining the foreground and background inputs using Z-Depth channels using the “Perform Depth Merge” checkbox, in which case every pair of pixels are compared. Which one is in front depends on its Z-Depth and not the connected input.

Rearranging or Combining Channels

Last, but certainly not least, it's also possible to rearrange and re-combine channels in any way you need, using one of four different node operations. For example, you might want to combine the red channel from one image with the blue and green channels of a second image to create a completely different channel mix. Alternatively, you might want to take the alpha channel from one image and merge it with the alpha channel of a second image in different ways, adding, subtracting, or using other intersection operations to create a very specific blend of the two.

The following nodes are used to re-combine channels in different ways:

- **Channel Boolean:** This is a 3D node used to remap and modify channels of 3D materials using a variety of simple pre-defined math operations.
- **Channel Booleans:** Used to shuffle or rearrange RGBA and auxiliary channels within a single input image, or among two input images, to create a single output image. If you only connect a single image to this node, it must be connected to the background input to make sure everything works.
- **Copy Aux:** The Copy Aux node is used to remap channels between RGBA channels and auxiliary data channels in a single 2D image. The Copy Aux node is mostly a convenience node, as the copying can also be accomplished with more effort (and flexibility) using a Channel Booleans node.
- **Matte Control:** Designed to do any combination of the following: (a) re-combining mattes, masks, and alpha channels in various ways, (b) modifying alpha channels using dedicated matte controls, and (c) copying alpha channels into the RGB stream of the image connected to the background input in preparation for compositing. You can copy specific channels from the foreground input to the background input to use as an alpha channel, or you can attach masks to the garbage matte input to use as alpha channels as well.


Understanding Premultiplication

Now that you understand how to direct and recombine RGB images and alpha channels in Fusion, it's time to go more deeply into alpha channels to make sure you always combine RGB and alpha channels correctly for each operation you perform in your composite. This might seem simple, but small mistakes are easy to make and can result in unsightly artifacts. This is arguably one of the most confusing areas of visual effects compositing, so don't skip this section.

When alpha channel and RGB pixels are both contained within a media file, such as a 3D rendered animation that contains RGB and transparency, or a motion graphics movie file with transparency baked in, there are two different ways they might be combined, and it's important to know which is in use.

- **Unmultiplied (Straight):** An RGB image unaltered by the semi-transparency information in a fourth channel (alpha channel)
- **Premultiplied:** An RGB image that has each channel multiplied by its alpha channel before compositing.

The term Premultiplied alpha is a term that has historically been used by editors, visual effects artists, and motion graphics designers, but it's imprecise. The alpha channel itself is not multiplied. The R, G, and B channels are multiplied by the alpha. In the end, the alpha channel stays the same, but the values contained in the R, G, and B channels are modified.




A RGB image (left) and its alpha channel (right)

Non-premultiplied images, sometimes called “straight” alpha channels, have RGB channels that are unaltered (not multiplied) by the alpha channel. The result is that the RGB image has no anti-aliased edges and no semi-transparency. It's usually obvious where the RGB image ends and the alpha matte begins. The image below is an example of the ragged edges seen in the RGB channels when using a non-premultiplied alpha channel. But notice the smooth semi-transparent edges found in the alpha.



A detailed view of a non-premultiplied RGB image (left) and its alpha channel (right)

A premultiplied alpha channel means the RGB pixels are multiplied by the alpha channel. This method guarantees that the RGB image pixels include semi-transparency where needed, like anti-aliased edges. Most computer-generated images are premultiplied for convenience, because they're easier to review smoothly without actually being placed inside of a composite.



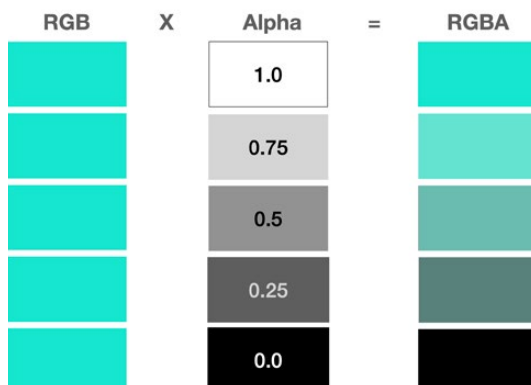
A detailed view of a premultiplied image (left) and its alpha channel (right)

What does this mean for compositing? The edges of a premultiplied image look much smoother, making them the preferred choice when compositing foreground over the background in a merge. Consequently, all alpha channels are made to be premultiplied in compositing operations if they weren't already.

On the other hand, it is always preferred to color correct a non-premultiplied RGBA image, because you don't want to alter the pixel values of an image after the RGB channels have been multiplied by the alpha channel.

If you think of this from a mathematical perspective:

- **RGB pixel value x 0 = 0:** The black transparent areas of an alpha channel have a pixel value of 0. When you take the value of an RGB pixel and multiply it by 0 ($n \times 0 = 0$) then by the laws of multiplication, the RGB value becomes 0, or fully transparent.
- **RGB Pixel value x 1 = RGB Pixel:** The solid or opaque white areas have a value of 1.0. When you take the value of an RGB pixel and multiply it by 1 ($n \times 1 = n$), then the RGB value stays the same, fully opaque.
- **RGB Pixel value x 0.3 = A different color:** Along the edges of an alpha channel are gray pixels, indicating semi-transparency. These semi-transparent pixels have a value falling somewhere between 1.0 and 0.0. To apply the alpha channel's anti-aliased edges to the RGB channels, you multiply the pixel values. The multiplication process mixes some percentage of the transparent pixels (black) with the RGB pixels. Although this is desired to get good anti-aliased edges, you can not color correct the image because it alters the smooth semi-transparency you created once it is done.



RGB pixels are multiplied by varying degrees of transparency and the result is a different RGB value.

The Rules of Premultiplication

Following from the information presented above, when you're compositing multiple images together and one or more has a built-in alpha channel, you want to make sure you follow these general rules to avoid problems:

- Always use premultiplied images with a Merge node.
- Only color-correct images that are not premultiplied.
- Always filter and transform images that are premultiplied.
- Never double premultiply an image.

Premultiplication and the Merge Node

The foreground input of a Merge node expects a premultiplied RGBA image. It is an additive merge, meaning the semi-transparent areas of the foreground are added over the background. However, if the image is not premultiplied, the pixels that should be transparent are still added, which typically results in an unwanted bright fringe around the edges of your foreground subject.

If you are compositing with a non-premultiplied alpha, you can fix these bright edges by changing the Merge to perform a Subtractive merge in the Inspector.




Nonpremultiplied edges in an additive Merge (left) and premultiplied edges in an additive Merge (right)

TIP: When an RGB image and a Mask node are combined using, for instance, a Matte Control node, if the RGB image is not multiplied by the mask in the Matte control, the checkerboard background pattern in the viewer will appear only semi-transparent when it should be fully transparent.

Color Correcting Premultiplied RGBA Images

When you premultiply an image, you tie any changes made to the brightness of the RGB pixels to the alpha channel pixels as well. If, for example, you raise the brightness of the foreground image in any way, you'll also raise the brightness of the alpha channel, which will likely not be desirable as this will change the transparency the alpha channel creates. The visible result of this when viewing the Merge node is that the entire background will become brighter (or darker if you'd lowered the RGB brightness) based on your color adjustment.

For this reason, the rule is always to divide the semi-transparent pixels before performing any color correction on an image with an alpha channel. You can do this turning on the Pre-Divide/Post Multiply checkbox in any node that performs color correction. Alternatively, you can use the Alpha Divide and Alpha Multiply nodes to do the same thing. These methods are covered in more detail later in this chapter.



Color correcting a pre-multiplied foreground incorrectly alters the background (left). Color correcting a nonpre-multiplied foreground works correctly (right).

Double Premultiplied RGBA Means Double Trouble

A common mistake made by many artists is to over-compensate for pre-multiplication. As important as it is to pre-multiply the alpha before compositing in a Merge node, it's just as important not to double pre-multiply the alpha. Performing a pre-multiply operation two times in a row can create a darkening halo effect around your images. You are effectively multiplying by the gray semi-transparent pixels twice; this is not optimal.



Double pre-multiplied image displays dark edges (left); pre-multiplied image with correct edges (right)

Premultiplied Alpha Channels and Filtering

When dealing with filtering, the state of the RGBA channels shouldn't matter for most composites. However, an exception might be if the filter algorithm you choose includes color modification. For instance, if a filter attempts to simulate a defocus by blooming highlights like a real light source, that filter might over-brighten pixels near a transparent edge, which will result in some manner of artifact when that image is composited.


Alpha Channel Status in MediaIn and Loader Nodes

When using a Loader node to add a clip to the composite, the Import tab in the Inspector includes a group of checkboxes that let you determine how an embedded alpha channel is handled. There are checkboxes to make the alpha channel solid (ignore transparency), to Invert the alpha channel, and to Post-Multiply the RGB channels with the alpha channel, should that be necessary.

When using a MediaIn node, you can modify how an embedded alpha channel is interpreted using the Clip attributes window. The Clips Attributes window includes an Alpha Mode menu setting to choose if the alpha channel is ignored, treated as premultiplied, inverted, or treated as nonpremultiplied (straight).

Controlling Premultiplication in Color Correction Nodes

Most nodes that require you to explicitly deal with the state of premultiplication of an RGBA image have a “Pre-Divide, Post-Multiply” checkbox. This includes simple color correction nodes such as Brightness Contrast and Color Curves, as well as the Color Correct node, which has the “Pre-Divide/Post-Multiply” checkbox in the Options panel of its Inspector settings.



The Pre-Divide/Post-Multiply checkbox of the Color Curves node, seen in the Inspector

This checkbox allows you to connect an RGBA premultiplied image to the node and perform a color correction operation. It takes the RGBA image input, performs a divide operation to remove the semi-transparency and then performs a multiplication operation before outputting the color corrected image. This way, the color correction is done using a nonpremultiplied image but the resulting output is a Merge-friendly premultiplied image.

Controlling Premultiplication With Alpha Divide and Alpha Multiply

The Alpha Divide and Alpha Multiply nodes, found in the Matte category of the Effects Library, are available when multiple operations in a row expect a “straight” alpha channel. Instead of performing repetitive Pre-Divide/Post Multiply operations on each node, you can use these two nodes to book-end the other nodes. Simply add the Alpha Divide node when you want the RGBA image data not to be premultiplied, and add the Alpha Multiply node when you want the image data to be premultiplied again. For example, if you’re using third-party OFX nodes that make color adjustments, you may need to control premultiplication before and after such an adjustment manually.



A node tree with explicit Alpha Divide and Alpha Multiply nodes

Multi Channel Compositing

If you've gone to the cinema and seen any recent superhero movie, you will have seen the results of sophisticated 3D rendering, combined with a whole lot of compositing. 3D applications can render very realistic images but the time it takes to render each frame of those realistic images can be measured in hours, not minutes. Any change to the 3D images, even relatively simple operations like color adjustments, changes to focus, filtering, or additional masking, means these images will need to be re-rendered entirely, which means you waiting, often many times over. For efficiencies sake, it's much faster to make iterative changes that can be accomplished with 2D image processing operations in Fusion instead.

To have the flexibility you need to make common changes to 3D images after-the-fact, the various attributes that make up the 3D scene are separated and rendered as different image sequences, often referred to as render passes. For example, render passes are often created for attributes like raw color, shadows, and reflections, that can be recombined as a 2D composite to produce the final result. Having different attributes rendered into different image sequences gives you a significant amount of flexibility, since now each image attribute can be color corrected, blurred, or further processed independently of the other attributes of the image, with fast-processing operations in Fusion.

The most common render passes that are typically generated come from the RGBA channels of the 3D scene. These are collectively called beauty passes and can consist of attributes like color, shadows, lighting, reflections, environment, and others.


Render passes can also contain non-RGB data. Different effects applications have different names for these passes, such as Data Channels, or AOVs (Arbitrary Output Variables). In Fusion, these channels are called Auxiliary Channels, and they contain 3D data such as Depth, Normals, Motion Vectors, and UV Coordinates (to name just a few).

When compositing a 3D render consisting of multiple render passes, the beauty passes are handled using one technique, and the Auxiliary Channels are handled with another. Since Fusion nodes carry RGBA channels by default, we'll cover beauty passes first, and then explain how to work with Auxiliary Channels later in this chapter.

Compositing with Beauty Passes

Each attribute of a beauty pass can be rendered into individual image sequences, so you end up with a series of numbered images, one for the diffuse pass, one for the reflection pass, another of the shadow pass, and so on. Alternatively, all the passes can be contained in a multi-part EXR image sequence. Multi-part EXRs benefit from requiring a bit less file management, but the passes are handled similarly in Fusion using either method.

A single MediaIn or Loader node only handles a single beauty pass since only one set of RGBA channels gets output per node. Setting up your composite in Fusion requires you to use a separate MediaIn or Loader node for each pass.

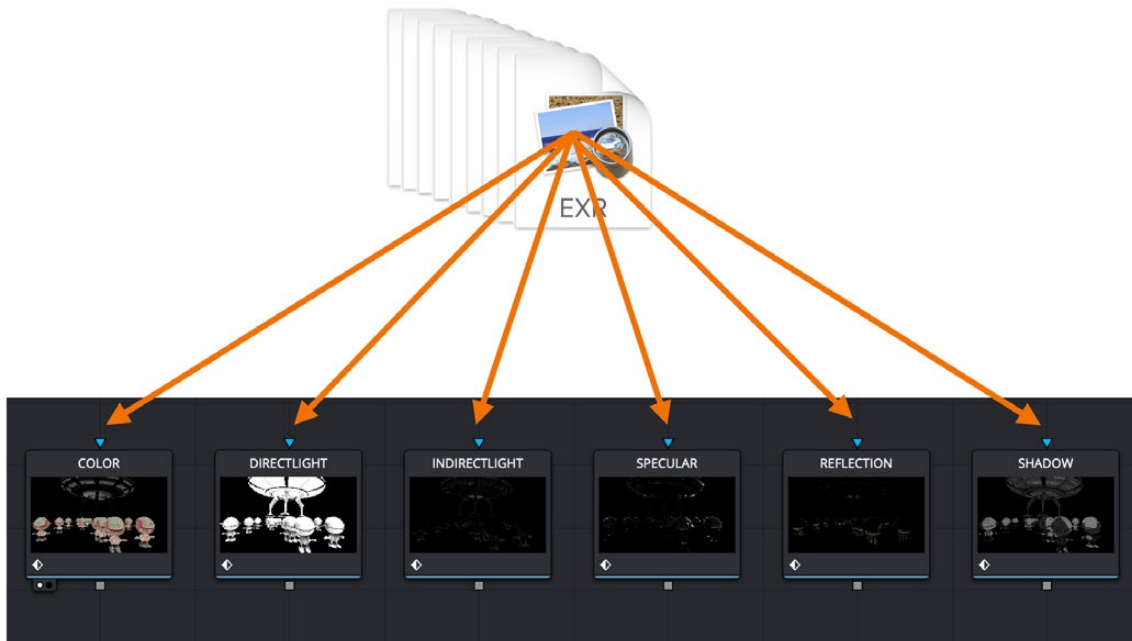


A MediaIn or Loader node is needed for each beauty pass.

TIP: The EXR format in Fusion is optimized when multiple Loaders are used to read the same EXR file. The file is only loaded once to access all the channels.

Beauty Pass Setup

Separated components of 3D renders, such as diffusion, shadows, or reflections, can be rendered individually using RGB channels. If you're provided with individual image sequences for each image component, they'll import and open in Fusion and can be displayed in the viewer, just like any other clip. If you're using multi-part EXR files, then you essentially have multiple RGB images contained within a single file. However, the MediaIn and Loader nodes can only make use of one set of RGB channels at a time. If you want to composite multiple RGB beauty passes together with one another, you must use multiple MediaIn or Loader nodes that point to the same image sequence but are assigned to different passes contained in the EXR file.




Multiple Loader or MediaIn nodes connect to a multi-part EXR image sequence

TIP: It is wise to rename each Loader or MediaIn to represent the beauty pass it contains.


Mapping Each Set of Beauty Passes to a Particular Node

Depending on whether you're using a MediaIn node or a Loader node, beauty passes can be mapped to RGB channels using either the Image tab, the Channels tab, or the Format tab in the Inspector. When using a Loader node, the Loader's Format tab is used.



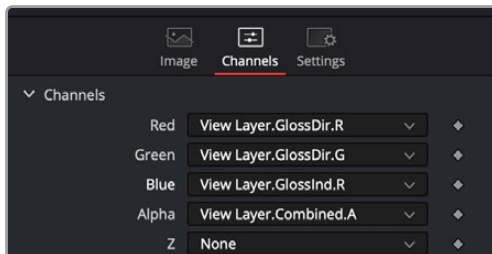
Beauty passes mapped to red, green, and blue channels in a Loader node

The MediaIn's Image tab includes a Layer menu. Any pass included in a multi-part EXR image sequence can be selected from this menu and automatically assigned to the RGBA channels.




The Layer menu in a MediaIn node showing headings for combined channel passes

In most cases, the menu shows the combined channel passes, meaning the individual red, green, blue, and alpha channels cannot be selected. Because the alpha channel is not included in many beauty passes, you sometimes need to borrow the alpha channel from a different beauty pass. For this reason, it's often better to use the Channels tab for mapping the individual channels of a beauty pass to the channels of the MediaIn node.



The MediaIn node's Channels tab or the Loader's Format tab provides access to individual channels.

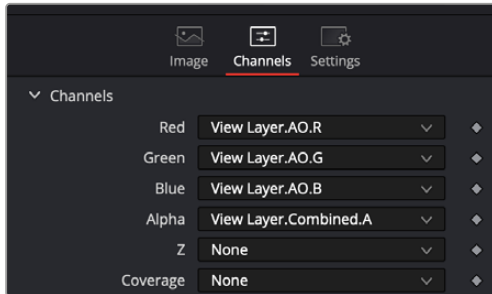
The MediaIn node in DaVinci Resolve and the Format tab in a Loader node include the same channel mapping functionality. The Channels tab and the Format tab include individual RGBA menus at the top of the tab. You can use these menus to map the RGBA channels from any pass contained in the multi-part EXR. For instance, if you want to map the Ambient Occlusion pass to the RGB channels, choose AO. R (Red) from the Red channel menu, AO. G (Green) from the Green channel menu, and AO. B (Blue) from the Blue channel menu.



A pass' individual channels mapped to the red, green, blue, and alpha channels in the MediaIn node

TIP: Different 3D applications will label beauty passes in different ways. For instance, the name for an Ambient Occlusion beauty pass may be AO, AM_OCC, or some other abbreviation.

The Ambient Occlusion beauty pass does not include an alpha channel. To composite it, you can reuse the alpha channel pass from another beauty pass. In the image below, the alpha channel is mapped using the combined render pass' alpha channel.




The alpha channel from a different beauty pass combined with the Ambient Occlusion pass

TIP: When using the Format tab in the Loader node, the checkbox next to each channel needs to be turned on for the corresponding channel to become available in the node's output.

Compositing Multiple Beauty Passes in the Node Editor

Once all of your passes are brought in and mapped to RGBA channels, you'll end up with a series of MediaIn or Loader nodes. How many MediaIn and Loader nodes is really up to your workflow. There is no predefined number of passes you will use. Every studio decides for themselves the standard set. However, there are common render passes that are involved in most composites. The following is a list of commonly used render passes and their generic names.



Commonly used beauty passes, compared

Compositing multiple beauty passes into a single output image is relatively straight forward. 3D rendering applications typically output linear gamma, so no Gamut or other color space conversion nodes are required if you're keeping the image in a linear color space for ease of compositing.

The basic compositing is accomplished with either a Merge node or a Channels Booleans node. Both allow for additive combining of render passes. There's no strict requirement for compositing each pass in any particular way, although in most situations a simple additive composite should work just fine.

To begin compositing render passes using a Merge node:

- 1 Connect a Color pass to the background input of a Merge node.
- 2 Connect a Direct Lighting pass to the foreground input.
- 3 Adjust the Alpha gain and Blend parameters to get the look you desire.




Compositing beauty passes starts by connecting a background and foreground into a Merge node

If you prefer, you can use a Channels Booleans node to make the same composite. In this case, there is no technical difference between the two nodes.

To composite render passes with a Channels Booleans node:

- 1 Connect a Color pass to the background input of a Channels Booleans node.
- 2 Connect a Direct Light pass to the foreground input.
- 3 Choose Add from the Operations menu
- 4 Choose Do Nothing from the Alpha To menu.



Channel Booleans is set to Add to combine the foreground input and the background input.

One of the exceptions to the steps above are Shadow passes, such as Ambient Occlusion. In that case, a *multiply* Apply mode is usually employed.

To composite an Ambient Occlusion render pass using a Merge node:

- 1 Connect last of a sequence of Merge node render passes to the background input of a Merge node.
- 2 Connect an Ambient Occlusion pass to the foreground input.
- 3 Choose Multiply from the Apply Mode menu.
- 4 Adjust the Gain and Blend parameters to get the look you desire.


As straightforward as this sounds, compositing using a recipe doesn't always work for every shot. When using different images, you may need to experiment with varying techniques of compositing for the best results.

Embedding Alpha into Beauty Passes

Alpha channels are not included in all beauty passes. If your shot requires that you composite your assembled beauty passes over live-action or some other background, it may be up to you to add the alpha channel from a pass that does include it.

To add an alpha channel into your assembled beauty pass composite, do the following:

- 1 Connect the last Merge or Channel Booleans output into the background input of a Matte Control node.
- 2 Connect the render pass that contains the alpha into the green Foreground input of the Matte Control node.
- 3 In the Matte Control's Inspector, choose Combine Alpha from the Combine menu.
- 4 Choose Copy from the Combine Op menu.



An alpha channel from the color pass added back into the completed beauty pass node tree

TIP: Alpha channels from 3D renderings are typically premultiplied. That being the case, be sure to turn on the Pre Divide/Post Multiply checkbox on any node that performs color correction. If using more than one node in a row to perform color correction, use the Alpha Divide and Alpha Mult nodes instead.

Working with Auxiliary Channels


Auxiliary channels are not RGB images; they are a family of special-purpose 3D image data that typically describes position, orientation, and object information for use in 2D composites. For example, Z-Depth channels describe the depth of each pixel of an image along a Z-axis, while an XYZ Normals channel describes the orientation (facing up, facing down, facing to the left or right) of every pixel in an image.

Similar to the use of multiple beauty passes, one of the most common reasons to use auxiliary data is to eliminate the need to re-render computationally expensive 3D imagery, by enabling even more aspects of rendered images to be manipulated after-the-fact. 3D rendering is computationally expensive and time-consuming, so outputting descriptive information about a 3D image allows sophisticated alterations to occur in 2D compositing, which is faster to perform and adjust.

There are two ways of obtaining auxiliary channel data:

- First, auxiliary data may be embedded within a clip rendered from a 3D application, most often using the EXR file format. In this case, it's best to consult your 3D application's documentation to determine which auxiliary channels can be generated and output.

- You may also obtain auxiliary channel data by generating it within Fusion, via 3D operations output by the Renderer 3D node, by the Optical Flow node, or by the Disparity node.

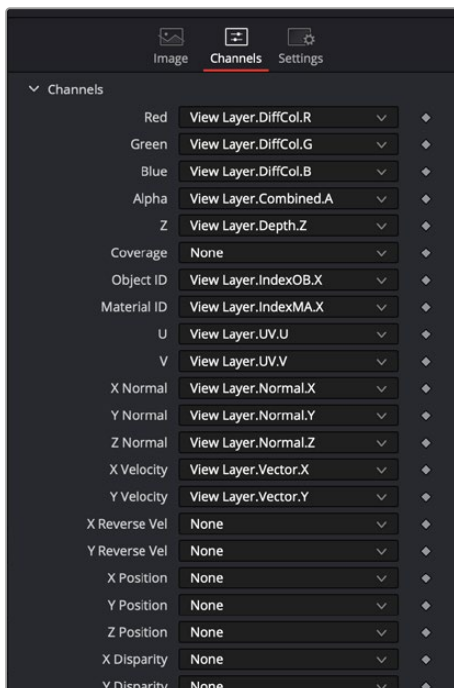


An RGBA 3D rendered scene that also contains auxiliary channels

TIP: When trying to locate information about auxiliary channels in other software, some 3D applications refer to auxiliary channels as Arbitrary Output Variables (AOVs), render elements, or secondaries.

Aux Channel Setup


When using a MediaIn or Loader node that links to a multi-part EXR file with auxiliary channel data, the Inspector's Channels or Format tab includes a pre-defined set of auxiliary channels for mapping purposes. Each pre-defined channel includes a menu that displays every attribute included in a multi-part EXR. From the menu, you select the render pass that should be assigned to the corresponding channel. As described earlier, RGB beauty passes like diffuse, shadow, and reflection are mapped to the red, green, and blue channels. Auxiliary passes include preset mappings.



A multi-part EXR file with embedded render passes mapped to their aux channels in a MediaIn node

Displaying Channels in the Viewer

Once you map an auxiliary channel in the Inspector, you can display the data as an RGB image in the viewer. Clicking the drop down Color menu at the top of the viewer reveals a list of every active auxiliary channel for the currently viewed node.



Select an aux channel from the color drop-down menu to display the aux channel in the viewer

TIP: The Color Inspector SubView can be used to read numerical values from all of the channels.


Auxiliary Channels Explained

Fusion is capable of using auxiliary channels to perform depth-based compositing, to create masks and mattes based on Object or Material IDs, and for texture replacements. Nodes that work with auxiliary channel information have been specifically developed to work with this data. The auxiliary channels that are supported in Fusion are described below.

Z-Depth

Each pixel in a Z-Depth channel contains a value that represents the relative depth of that pixel in the scene. In the case of overlapping objects in a model, most 3D applications take the depth value from the object closest to the camera when two objects are present within the same pixel since the closest object typically obscures the farther object.

When present, Z-Depth can be used to perform depth merging using the Merge node or to control simulated depth-of-field blurring using the Depth Blur node.




The rendered Z-Depth channel for the previous RGBA image

For this example, we'll examine the case where the Z-Depth channel is provided as a separate file. The Z- channel can often be rendered as an RGB image. You'll need to combine the beauty and Z pass using a Channel Booleans node. When the Z pass is rendered as an image in the RGB channels, the Channels Booleans node is used to re-shuffle the Lightness of the foreground RGB channel into the Z-channel.


To combine the Z-pass and beauty pass:

- 1 Connect the MediaIn node containing the the beauty pass to the background input of the Channel Booleans node.
- 2 Connect the MediaIn node containing the Z-Depth pass to the green foreground input of the Channel Booleans node.
- 3 Select the Channel Booleans node, and use the Inspector to set the To Red, To Green, To Blue, and To Alpha menus to Do Nothing.
- 4 Select the Aux tab, and set the To Z Buffer menu to Lightness FG.
- 5 Connect the output of the Channels Booleans node into the Depth Blur node.



The Aux tab configured to shuffle the Foreground Lightness to the Z-Depth channel.


The Depth Blur node is one of the nodes that take advantage of a Z-channel in order to create blurry depth-of-field simulations. To set this up, the output of the MediaIn node connects to the background input on the Depth Blur.



The Depth Blur uses the Z-channel that is enabled in the Channel Booleans node.


The Depth Blur's controls in the Inspector are very dependent on the type of image you're using. It can be easier to begin by adjusting the controls in the Inspector to some better defaults. Start by increasing the Blur Size to 10. This will make it easier to see even the smallest of changes. Next, instead of using the Focal Point, you should pick a focal point in the image by dragging the Sample button into the viewer and selecting a pixel that determines the part of the picture to keep in focus.

The final setup steps are to lower Z Scale to somewhere around 0.2 (if you're using a floating-point image), and leave the Depth of Field alone for now. This should show you some blurring in the image.



Start by improving the defaults if your image is 16- or 32-bit floating point.

Once you see these experimental results, you can return to each parameter and refine it as needed to achieve the actual look you want.



An image using a Z-Depth channel for blurring

TIP: Z-Depth channels often contain negative values. If this causes problems, you can choose Normalize Color Range from the viewer's Options menu to apply a normalization to the viewer, keeping the image within a range from 0 to 1.

Z-Coverage

The Z-Coverage channel is a somewhat extinct render pass in most 3D applications. It was a way of restoring antialiasing to rendered color masks and Z-Depth passes. It indicated pixels in the Z-Depth that contained two objects. The value was used to indicate, as a percentage, how transparent that pixel was in the final depth composite. It can still be used today if you are rendering files from one of the few applications that can produce them.

TIP: The wide adoption of an open-source matte creation technology called Cryptomatte, has somewhat superseded mattes created from Coverage, Background, Object ID, and Material ID passes.

Background RGBA

This channel is a somewhat extinct render pass in most 3D applications. It contained the color values from the objects behind the pixels described in the Z coverage.


Object ID

Most 3D applications are capable of assigning ID values to objects in a scene. Each pixel in the Object ID channel will be identified by that ID number, allowing for the creation of masks.

If you want to use an Object ID in a comp, like all aux channels you must map the Object ID pass to the Object ID channel in the MediaIn or Loader Node.

To use an ObjectID pass, do the following:

- 1 In the MediaIn or Loader node, use the Channels or Format tab to map the Object ID pass to the Object ID aux channel.
- 2 In whatever node you want to have affected by the ObjectID matte, select the Settings tab, turn on the Object ID checkbox, and select the ID number assigned to the object.



The common Settings tab on most nodes contains ObjectID controls.


Material ID

Most 3D applications are capable of assigning ID values to materials in a scene. Each pixel in the Material ID channel will be identified by that ID number, allowing for masks based on materials.

You can set up Material IDs using the Settings tab, similarly to how ObjectIDs are set.

UV Texture

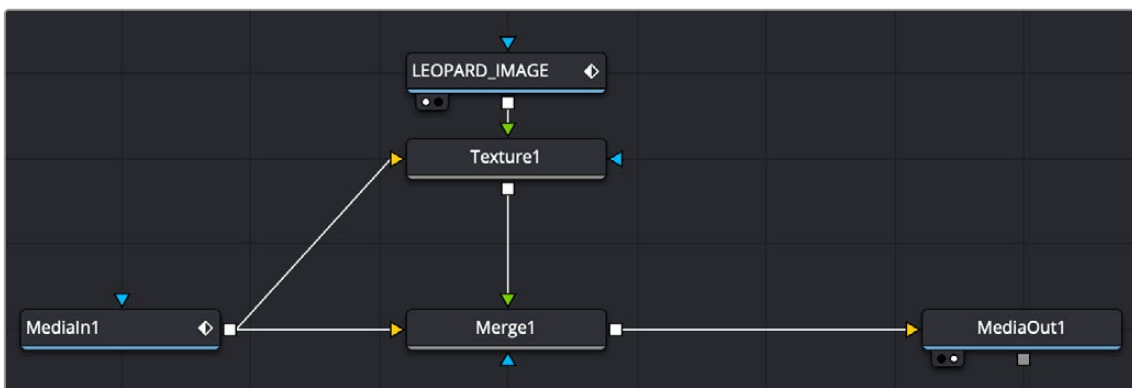
The UV Texture channels contain information about how pixels in an image map to texture coordinates. This is used to retexture an object in a 2D image. For instance, If you want to apply a logo onto a rendered object, you can use the UV aux channel with the Texture node.



Texture (left) applied to 2D image (right) using UV channels and texture node.

To use a UV pass, do the following:

- 1 In the MediaIn or Loader node, use the Channels or Format tab to map the U and V pass to the U and V aux channels.
- 2 Connect the output of the MediaIn or Loader node to the background input of the Texture node.
- 3 Connect the texture image you want to use to the foreground input of the Texture node.
- 4 If you want to combine the original texture with the new texture, use a merge with the background input from the original image and the foreground input from the Texture node.
- 5 Adjust Merge's Apply mode, Alpha Gain, and blend to get the desired mix of the two textures.




UV channels from a MediaIn node used in a Texture node and merged over the original image

TIP: If you are using a separate UV render pass with the UV data in the RGB channels, map red to U and green to V in a Channel Booleans node.

X, Y, and Z Normals

The X, Y, and Z Normal channels contain information about each pixel's orientation (the direction it faces) in 3D space. The normals are often displayed as lines coming out from your object perpendicular to the surface, letting you visualize the relationship between the surface and camera.



The Normals display the direction of the surface.

The Normals X, Y, and Z channels are often used with a Shader node to perform relighting adjustments on a 2D rendered image.

To setup a Shader node to use XYZ Normals, do the following:


- 1 In the MediaIn or Loader node, use the Channels or Format tab to map the individual X, Y, and Z Normals pass to the X Normal, Y Normal, and Z Normal channels.
- 2 Connect the output of the MediaIn or Loader node to the background input of the Shader node.
- 3 Optionally, connect a floating-point EXR image to be used as a reflection image to the Shader node's reflection input.
- 4 Adjust the Shader controls to perform relighting.



Original 2D image (left) and Normals used for relighting (right)

XY Vector and XY BackVector

The Vector channels indicates the pixel's motion from frame to frame. It can be used to apply motion blur to an image or to generate optical flow analysis for retiming. The XY Vector points to the next frame, while the XY BackVector points to the previous frame.

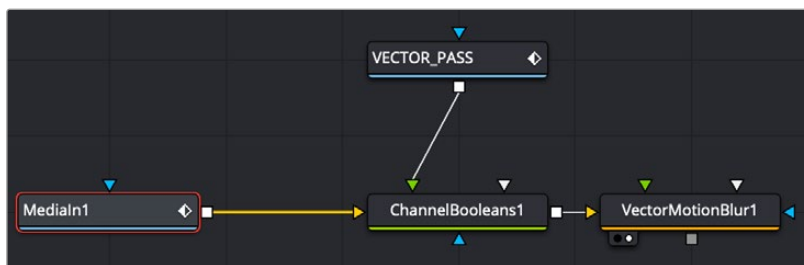


XY Vector pass (left) used with Vector Motion Blur to generate motion blur on spaceship (right)

Often the vector pass will be rendered in a separate pass as an RGB image. The X and Y vector data is located in the R and G channels. In order to place them in the vector channels, you can use a Channel Booleans node.

To use a Motion Vector pass to create motion blur, do the following:

- 1 Add a MediaIn or Loader node for the image and the Vector render pass.
- 2 Connect the output of the image into the background of a Channel Booleans node.
- 3 Connect the output of the Vector render pass to the Channel Boolean's foreground
- 4 In the Channel Booleans inspector, set the To Red, To Green, To Blue, and To Alpha all to Do Nothing.
- 5 Select to the Aux tab.
- 6 Turn on Enable Extra Channels.
- 7 Set the To X Vector drop-down menu to Red FG, and then set the To Y Vector drop-down menu to Green FG.
- 8 Connect the Channel Booleans node's output to the yellow background input on a Vector Motion Blur node.



The Vector render pass is combined with the beauty image using the Channels Booleans node, which then feeds the Vector Motion Blur node.

World Position

World Position Pass (WPP) is an auxiliary channel, sometimes referred to as Point Position, XYZ pass, or WPP. It's used to represent each pixel's 3D (XYZ) position as an RGB color value. The result is data that can be viewed as a very colorful RGB image. Like Z-Depth, this can be used for compositing via depth. However, it can also be used for masking based on 3D position, regardless of camera transforms.


The colors correspond to a pixel's position in 3D, so if a pixel sits at 0/0/0 in a 3D scene, the resulting pixel's will have an RGB value of 0/0/0 or black. If a pixel sits at 1/0/0 in the 3D scene, the resulting pixel is fully red. Due to the huge extent, 3D scenes can have the WPP channel should always be rendered in 32-bit floating-point to provide the accuracy needed.



XYZ Position

XY Disparity

XY Disparity is the only channel listed here that is not generated in a 3D application. These channels indicate where each pixel's corresponding matte can be found in a stereo image. Each eye, left and right, will use this vector to point to where that pixel would be in the other eye. This can be used for adjusting stereo effects, or to mask pixels in stereo space.



Combined X and Y Disparity channels

Using Cryptomatte in Fusion

Cryptomatte is an open-source technology that has been widely adopted in 3D applications. Unlike Z-Depth mattes or Object IDs, Cryptomatte automatically generates anti-aliased ID mattes from 3D renders with support for motion blur, transparency, and depth of field.

Fusion does not natively support the Cryptomatte format. However, using a free plug-in from 3rd party developers, you can use Cryptomatte render passes in Fusion.

Cryptomatte for Fusion can be downloaded and installed for free: <https://github.com/Psyop/Cryptomatte>

Or to use an easier installer, you can download Reactor, which comes bundled with Cryptomatte and offers many other free, useful Fusion plug-ins. Reactor can be found at: <https://www.steakunderwater.com>

Propagating Auxiliary Channels

Ordinarily, auxiliary channels are propagated along with RGBA image data, from node to node, among gray-colored nodes, including those in the Blur, Filter, Effect, Transform, and Warp categories. Basically, most nodes that simply manipulate channel data propagate (and potentially manipulate) auxiliary channels with no problems.

However, when you composite two image layers using the Merge node, auxiliary channels only propagate through the image that's connected to the background input. The rationale for this is that in most composites that include computer-generated imagery, the background is most often the CG layer that contains auxiliary channels, while the foreground is a live-action green screen plate with subjects or elements that are combined against the background, which lack auxiliary channels.

Nodes That Use Auxiliary Channels

The availability of auxiliary channels opens up a world of advanced compositing functionality. This section describes every Fusion node that has been designed to work with images that contain auxiliary channels.

- **Copy Aux:** The Copy Aux tool can copy auxiliary channels to RGB and then copy them back. It includes some useful options for remapping values and color depths, as well as removing auxiliary channels.
- **Channel Booleans:** The Channel Boolean tool can be used to combine or copy the values from one channel to another in a variety of ways.
- **Custom Tool, Custom Vertex 3D, pCustom:** The “Custom” tools can sample data from the auxiliary channels per pixel, vertex, or particle and use that for whatever processing you would like.
- **Depth Blur:** The Depth Blur tool is used to blur an image based on the information present in the Z-Depth. A focal point is selected from the Z-Depth values of the image and the extent of the focused region is selected using the Depth of Field control. The Scale value default is based on an 8-bit image so it is important to lower the scale value when using the Depth Blur with 16- or 32-bit float files.
- **Disparity to Z, Z to Disparity, Z to WorldPos:** These tools use the inherent relationships between depth, position, and disparity to convert from one channel to another.
- **Fog:** The Fog tool makes use of the Z-Depth to create a fog effect that is thin closer to the camera and thickens in regions farther away from the camera. You use the Pick tool to select the Depth values from the image and to define the Near and Far planes of the fog’s effect.
- **Lumakeyer:** The Lumakeyer tool can be used to perform a key on the Z-Depth channel by selecting the Z-Depth in the channel drop-down list.
- **Merge:** In addition to regular compositing operations, Merge is capable of merging two or more images together using the Z-Depth, Z-Coverage, and BG RGBA buffer data. This is accomplished by enabling the Perform Depth Merge checkbox from the Channels tab.
- **New Eye:** For stereoscopic footage, New Eye uses the Disparity channels to create new viewpoints or to transfer RGBA data from one eye to the other.
- **Shader:** The Shader tool applies data from the RGBA, UV and the Normal channels to modify the lighting applied to objects in the image. Control is provided over specular highlights, ambient and diffuse lighting, and position of the light source. A second image can be applied as a reflection or refraction map.
- **Shadow:** The Shadow tool can use the Z-Depth channel for a Z-Map. This allows the shadow to fall onto the shape of the objects in the image.
- **Smooth Motion:** Smooth Motion uses Vector and Back Vector channels to blend other channels temporally. This can remove high frequency jitter from problematic channels such as Disparity.
- **SSAO:** SSAO is short for Screen Space Ambient Occlusion. Ambient Occlusion is the lighting caused when a scene is surrounded by a uniform diffuse spherical light source. In the real world, light lands on surfaces from all directions, not from just a few directional lights. Ambient Occlusion captures this low frequency lighting, but it does not capture sharp shadows or specular lighting. For this reason, Ambient Occlusion is usually combined with Specular lighting to create a full lighting solution. The SSAO tool uses the Z-Depth channel but requires a Camera3D input.
- **Stereo Align:** For stereoscopic footage, the Disparity channels can be used by Stereo Align to warp one or both of the eyes to correct misalignment or to change the convergence plane.
- **Texture:** The Texture tool uses the UV channels to apply an image from the second input as a texture. This can replace textures on a specific object when used in conjunction with the Object ID or Material ID masks.

- **Time Speed and Time Stretcher:** These tools can use the Vector and BackVector channels to retime footage.
- **Vector Distortion:** The forward XY Vector channels can be used to warp an image with this tool.
- **Vector Motion Blur:** Using the forward XY Vector channels, the Vector Motion Blur tool can apply blur in the direction of the velocity, creating a motion blur effect.
- **Volume Fog:** Volume Fog is a raymarcher that uses the World Position channels to determine ray termination and volume dataset placement. It can also use cameras and lights from a 3D scene to set the correct ray start point and Illumination parameters.
- **Volume Mask:** Volume Mask uses the World Position channels to set a mask in 3D space as opposed to screen space. This allows a mask to maintain perfect tracking through a camera move.

TIP: The Object ID and Material ID auxiliary channels can be used by some tools in Fusion to generate a mask. The “Use Object” and “Use Material” settings used to accomplish this are found in the Settings tab of that node’s controls in the Inspector.

Image Formats That Support Aux Channels

Fusion supports auxiliary channel information contained in a variety of image formats. The number of channels and methods used are different for each format.

- **OpenEXR (*.exr):** The OpenEXR file format is the primary format used to contain an arbitrary number of additional image channels. Many renderers that will write to the OpenEXR format will allow the creation of channels that contain entirely arbitrary data. For example, a channel with specular highlights might exist in an OpenEXR. In most cases, the channel will have a custom name that can be used to map the extra channel to one of the channels recognized by Fusion.
- **SoftImage PIC (*.PIC, *.ZPIC and *.Z):** The PIC image format (used by SoftImage) is an older image format that can contain Z-Depth data in a separate file marked by the ZPIC file extension. These files must be located in the same directory as the RGBA PIC files and must use the same names. Fusion will automatically detect the presence of the additional information and load the ZPIC images along with the PIC images.
- **Wavefront RLA (*.RLA), 3ds Max RLA (*.RLA) and RPF (*.RPF):** This is an older image format capable of containing any of the image channels mentioned above. All channels are contained within one file, including RGBA, as well as the auxiliary channels. These files are identified by the RLA or RPF file extension. Not all RLA or RPF files contain auxiliary channel information, but most do. RPF files have the additional capability of storing multiple samples per pixel, allowing different layers of the image to be loaded for very complex depth composites.
- **Fusion RAW (*.RAW):** Fusion’s native RAW format is able to contain all of the auxiliary channels as well as other metadata used within Fusion.

Creating Auxiliary Channels in Fusion

The following nodes create auxiliary channels:

- **Renderer 3D:** Creates these channels in the same way as any other 3D application would, and you have the option of outputting every one of the auxiliary data channels that the Fusion page supports.
- **Optical Flow:** Generates Vector and Back Vector channels by analyzing pixels over consecutive frames to determine likely movements of features in the image.
- **Disparity:** Generates Disparity channels by comparing stereoscopic image pairs.

Chapter 79

Compositing Layers in Fusion

This chapter is intended to give you a solid base for making the transition from a layer-based compositing application to Fusion's node-based interface. It provides practical information about how to start structuring a node tree for simple layered composites.


Contents

Applying Effects	1406
Adding a Node to the Tree	1406
Editing Parameters in the Inspector	1406
Replacing Nodes	1407
Adjusting Fusion Sliders	1408
Compositing Two Clips Together	1408
Adding Additional Media to Compositions	1408
Automatically Creating Merge Nodes	1409
Fixing Problem Edges in a Composite	1410
Using Composite Modes in the Merge Node	1411
Creating and Using Text	1412
Creating Text Using the Text+ Node	1412
Styling and Adjusting Text	1413
Using Text as a Mask	1414
Using Transform Controls in the Merge Node	1417
Building a Simple Green-Screen Composite	1418
Mapping Timeline Layers to Nodes in Fusion	1418
Pulling a Green-Screen Key Using the Delta Keyer	1420
Dealing with Spill	1423
Masking a Graphic	1424

Applying Effects

Before we dive into multi-layered composites, let's start by looking at some very simple effects and build up from there. Opening the Effects Library, then clicking the disclosure control to the left of Tools, reveals a list of categories containing all the nodes available in Fusion. As mentioned before, each node does one thing, and by using these nodes in concert you can create extremely complex results from humble beginnings.


Clicking the Effect category reveals its contents. In this example we'll use the TV effect.



Browsing the Effect category to find the TV node

Adding a Node to the Tree

Assuming that the MediaIn node in the Fusion page or the Loader node in Fusion Studio is the currently selected node in the Node Editor, clicking once on the TV node, for example, in the Effects Library automatically adds that node to the node tree to the right of the selected node. In the Fusion page, it immediately takes effect in the viewer thanks to the fact that the MediaOut1 node is what's loaded in the viewer, since that means that all nodes upstream of the MediaOut1 node will be processed and shown.



A new node added from the Effects Library

In Fusion Studio, you must press the 1 or 2 key on the keyboard to load the selected node in the viewer.

There are many other ways of adding nodes to your node tree, but it's good to know how to browse the Effects Library as you get started.


Editing Parameters in the Inspector

To improve an effect, you can make an adjustment to a node's parameters in the Inspector at the right. The selected node shows its controls in the Inspector, in which most nodes have several tabs of controls seen as little icons just underneath that node's title bar.




The Inspector showing the parameters of the TV effect

Clicking the last panel on any node opens the Settings panel. Every node has a Settings panel, and this is where the parameters that every node shares, such as the Blend slider and RGBA buttons, are found. These let you choose which image channels are affected, and let you blend between the effect and the original image.



The Settings panel, which includes channel limiting and mask handling controls that every node shares

In the case of the TV effect, for example, the resulting image has a lot of transparency because the scan lines being added are also being added to the alpha channel, creating alternating lines of transparency. Turning off the Alpha checkbox results in a more solid image, while opening the Controls tab (the first tab) and dragging the Scan Lines slider to the right to raise its value to 4 creates a more visible television effect.



The original TV effect (left), and modifications to the TV effect to make the clip more solid (right)

Replacing Nodes

In the Effect category of the Effects Library, you'll also find a Highlight node that adds glints to the highlights of an image.

Instead of clicking the Highlight node, which would add it after the currently selected node, dragging and dropping a node from the Effects Library on top of a node in the Node Editor replaces the node in the Node Editor.



Dragging a node from the Effects Library onto a node in the Node Editor to replace it

In our example, the Highlight1 node takes the TV node's place in the node tree, and the new effect can be seen in the viewer, which in this example consists of star highlights over the lights in the image.

It's time to use the Inspector controls to customize this effect.

Adjusting Fusion Sliders

When you drag a slider in the Fusion Inspector—in this case, the Number of Points slider—a little dot appears underneath it. This dot indicates the position of the default value for that slider and also serves as a reset button if you click it.



Adjusting a slider reveals a reset button underneath it

Each slider is limited to a different range of minimum and maximum values that is particular to the parameter you're adjusting. In this case, the Number of Points slider maxes out at 24. However, you can remap the range of many (but not all) sliders by entering a larger value in the number field to the right of that slider. Doing so immediately repositions the slider's controls to the left as the slider's range increases to accommodate the value you just entered.



Entering a larger value to expand the range over which a slider will operate


Compositing Two Clips Together

As entertaining as it is adding individual nodes to create simple effects, eventually you'll need to start adding additional layers of media in order to merge them together as composites. Let's turn our attention to another example in which we need to combine a background clip with a foreground clip that includes a built-in alpha channel, to see simple layering in action.

Adding Additional Media to Compositions

You'll often find that even though you start out wanting to do something relatively simple, you end up adding additional media to create the effect that you need.

- In Fusion Studio, you do this by adding additional Loader nodes. If you add a new Loader node to an empty area of the Node Editor, you'll add an unconnected Loader2 node (incremented to keep it unique) that you can then connect how you want.
- In the Fusion page, you can open the Media Pool and drag clips directly to the Node Editor to add them to your node tree. If you drag a clip from the Media Pool to an empty area of the Node Editor, you'll add an unconnected MediaIn2 node (incremented to keep it unique) that you can then connect in any way you want.




The Media Pool as seen in the Fusion page

Automatically Creating Merge Nodes

If you want to connect the incoming clip immediately to your node tree as the top layer, or foreground, of a composite, in Fusion Studio select the Loader1 node and then add a second Loader node. In the Fusion page, drag the new clip from the Media Pool right on top of any connection line.

In both cases, the new MediaIn or Loader node automatically becomes the “foreground input”.



Dragging a node from the Media Pool onto a connection (left), and dropping it to create a Merge node composite (right)

The Node Editor is filled with shortcuts like this to help you build your compositions more quickly. Here's one for when you have a disconnected node that you want to composite against another node with a Merge node. Drag a connection from the output of the node you want to be the foreground layer, and drop it on top of the output of the node you want to be the background layer. A Merge node will be automatically created to build that composite. Remember: background inputs are orange, and foreground inputs are green.



Dragging a connection from a disconnected node to another node's output (left), and dropping it to create a Merge node composite (right)

Adding Clips to a Fusion Composite From the File System

If you drag clips from the file system directly into the Node Editor, they'll be added to the DaVinci Resolve Media Pool automatically. So, if you have a library of stock animated background textures and you've just found one you want to use using your file system's search tools, you can simply drag it straight into the Node Editor and it'll be added to the currently selected bin of the Media Pool.


Fixing Problem Edges in a Composite

Most often, the Merge node does a perfectly good job when handed a foreground image with premultiplied alpha transparency to composite against a solid background image. However, from time to time, you may notice a small bit of fringing at the edge of the border of a foreground element and transparent area, such as seen in the following close-up. This slight lightening at the edge is a tell-tale sign that the clip probably wasn't premultiplied. The Merge node expects all foreground images with alpha channels to be premultiplied. But this is something that's easily fixed.



A bit of fringing at the edge of a foreground element surrounded by transparency

Click to select the Merge node for that particular composite, and look for the Subtractive/Additive slider.



The Subtractive/Additive slider, which can be used to fix or improve fringing in composites

Drag the slider all the way to the left, to the Subtractive position, and the fringing disappears.



A clip with alpha exhibits fringing (left), and after fixing fringing by dragging the Subtractive/Additive slider to the left (right)

The Subtractive/Additive slider, which is only available when the Apply mode is set to Normal, controls whether the Normal mode performs an Additive merge, a Subtractive merge, or a blend of both. This slider defaults to Additive merging, which assumes that all input images with alpha transparency are premultiplied (which is usually the case). If you don't understand the difference between Additive and Subtractive merging, here's a quick explanation:

- An Additive merge, with the slider all the way to the right, is necessary when the foreground image is premultiplied, meaning that the pixels in the color channels have been multiplied by the pixels in the alpha channel. The result is that transparent pixels are always black, since any number multiplied by 0 is always 0. This obscures the background (by multiplying with the inverse of the foreground alpha), and then simply adds the pixels from the foreground.
- A Subtractive merge, with the slider all the way to the left, is necessary if the foreground image is not premultiplied. The compositing method is similar to an Additive merge, but the foreground image is first multiplied by its own alpha to eliminate any background pixels outside the alpha area.

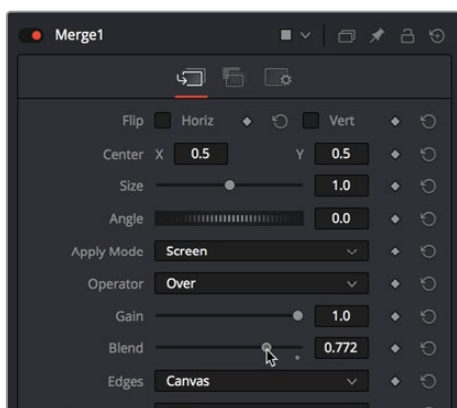
The Additive/Subtractive slider lets you blend between two versions of the merge operation, one Additive and the other Subtractive, to find the best combination for the needs of your particular composite. Blending between the two is occasionally useful for dealing with problem composites that have edges that are calling attention to themselves as either too bright or too dark.

For example, using Subtractive merging on a premultiplied image may result in darker edges, whereas using Additive merging with a non-premultiplied image will cause any non-black area outside the foreground's alpha to be added to the result, thereby lightening the edges. By blending between Additive and Subtractive, you can tweak the edge brightness to be just right for your situation.

Using Composite Modes in the Merge Node

To create a more convincing composite in layer-based systems, you often use Blend modes. Blend modes are located in the Merge node since that is where one layer gets composited over another. Let's take an example where you want to use the Screen mode to make a foreground image look more like a reflection.

The Merge node has a variety of controls built into it for creating just about every compositing effect you need. Items you may be familiar with as Blend modes are located in the Apply Mode pop-up menu. You can use these mathematical compositing modes to combine the foreground and background layers together. A Blend slider allows you to fade the foreground input with the background.



Adjusting the Apply Mode and Blend slider of the Merge node in the Inspector

NOTE: The Subtractive/Additive slider disappears when you choose any other Apply Mode option besides Normal, because the math would be invalid. This isn't unusual; there are a variety of controls in the Inspector that hide themselves when not needed or when a particular input isn't connected.

The Screen node is perfect for simulating reflections, and lowering Blend a bit lets you balance the foreground and background images. It's subtle, but helps sell the shot.

TIP: You may have noticed that the Merge node also has a set of Flip, Center, Size, and Angle controls that you can use to transform the foreground image without needing to add a dedicated Transform node. It's a nice shortcut for simplifying node trees large and small.

Creating and Using Text


In this next example, we'll look at how to create a simple text object using the Text+ node. Then, we'll see how to use the text generator's alpha channel in another image to create a more complex composite.

Creating Text Using the Text+ Node

The Text+ node is the primary tool for creating 2D text in the Fusion page. If you are using DaVinci Resolve, this is also the same Text+ generator available in the Edit page. It is easily accessible right in the toolbar.


The Text+ node is an incredibly deep tool for creating text effects, with six tabs of controls for adjusting everything from text styling, to different methods of layout, to a variety of shading controls including fills, outlines, shadows, and borders. As sophisticated a tool as this is, we'll only be scratching the surface in this example.

We'll start with a MediaIn node that will serve as our background selected in the Node Editor. Clicking the Text+ button automatically creates a new Text+ node connected as the foreground input of a Merge node. The same behavior occurs if you are using Fusion Studio, with a Loader node.




Selecting a node you want to append another node to (top) Clicking the Text+ button on the toolbar automatically creates a Merge composite with the text as the foreground input connection (bottom)

Selecting a Text node opens the default Text panel parameters in the Inspector, and it also adds a toolbar at the top of the viewer with tools specific to that node. Clicking on the first tool at the left lets you type directly into the viewer, or you can type into the Styled Text field in the Inspector.



The viewer toolbar for the Text node with tools for text entry, kerning, and outline controls


If you're viewing the Merge, the text appears in the viewer superimposed against the background clip. Onscreen controls appear that let you rotate (the circle) and reposition (the red center handle and two arrows) the text, and we can see a faint cursor that lets us edit and kern the text using other tools in the viewer toolbar.



Text that's been typed into the viewer, with onscreen text transform controls

Styling and Adjusting Text

To style the text, you use the controls in the Inspector, modifying text style controls such as Font, Size, and Tracking to change the spacing between the letters.



The restyled text

TIP: Holding down the Command key while dragging any control in the Inspector “gears down” the adjustment so that you can make smaller and more gradual adjustments.

Selecting the Manual Kerning tool in the viewer toolbar (second tool from the left) reveals small red dots underneath each letter of text.




The Manual Kerning tool in the viewer toolbar

Clicking a red dot under a particular letter puts a kerning highlight over that letter.

To make manual kerning adjustments:

- 1 Option-drag the red dot under any letter of text to adjust that character's kerning while constraining letter movement to the left and right. You can also drag letters up and down for other effects. Depending on your system, the kerning of the letter you're adjusting might not update until you drop the red dot in place.
- 2 If you don't like what you've done, you can open the Advanced Controls in the Inspector and clear either the kerning of selected letters or all manual kerning before starting over again.



Option-dragging the little red dot revealed by the Manual Kerning tool to manually adjust kerning

Using Text as a Mask

You can fill the text with a color or gradient, or you can use the text as a matte to cut the letters out of another image. First, we'll drag another clip of a chalkboard covered with math from the Media Pool to the Node Editor as a disconnected MediaIn2 node.


Disconnecting and Reconnecting Nodes

If we start from the previous example that had a Medialn as the background and the Text+ as the foreground to a Merge, we'll need to do a little rearranging. Clicking the last half of the connection from the Text1 node to the Merge foreground input disconnects it.



Clicking the second half of a connection to disconnect it (top), and the result with the text node disconnected (bottom)

Connecting a Medialn2 or Loader2 node onto the Merge1 node's foreground input causes the entire viewer to be filled with the Medialn2 (assuming we're still viewing the Merge node). At this point, we need to insert the Text1 node's image as an alpha channel into the Medialn2 node's connection, and we can do that using a MatteControl node.

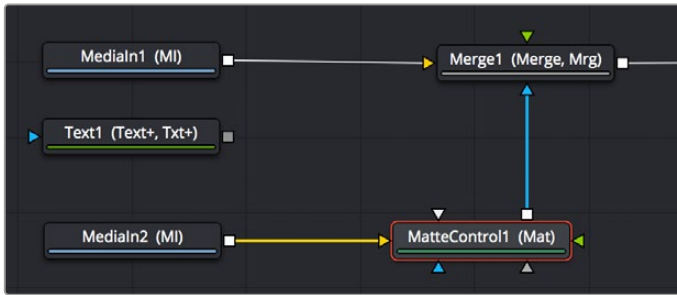


The updated composite, with two video images connected and the text node disconnected

Using Matte Control Nodes

Select the Medialn2 node, and click the Matte Control button in the toolbar to add it between the Medialn2 and Merge1 nodes. (To tidy up, I've moved the nodes around a bit in the screenshot.)

The MatteControl node has numerous uses. Among them is taking one or more alpha channels, mattes, or images that are connected to the Garbage Matte, Solid Matte, and/or foreground inputs, combining them, and using the result as an alpha channel for the image that's connected to the background input. It's critical to make sure that the image you want to add an alpha channel to is connected to the background input of the MatteControl node, or the MatteControl node won't work.

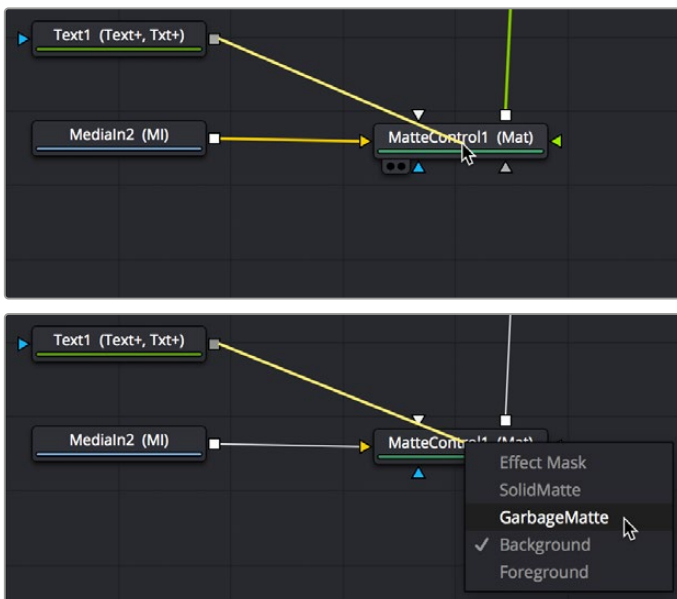


The second image properly connected to the MatteControl node's background input

With this done, connecting the Text+ node's output, which has the alpha channel, to the MatteControl node's Garbage Matte input, is a shortcut we can use to make a mask, matte, or alpha punch out a region of transparency in an image.



Keep in mind that it's easy to accidentally connect to the wrong input. Because inputs rearrange themselves depending on what's connected and where the node is positioned (and, frankly, the colors can be hard to keep track of when you're first learning), it's key to make sure that you always check the tooltips associated with the input you're dragging a connection over to make sure that you're really connecting to the correct one. If you don't, the effect won't work, and if your effect isn't working, the first thing you should always check is whether you've connected the proper inputs.

One alternate method of connecting nodes together is to hold down the Option key while dragging a connection from one node's output and dropping it onto the body of another node. This opens a pop-up menu from which you can choose the specific input you want to connect to, by name. Note that the menu only appears after you've dropped the connection on the node and released your pointing device's button.



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, a text-shaped area of transparency is displayed for the graphic if you load the MatteControl node into the viewer.



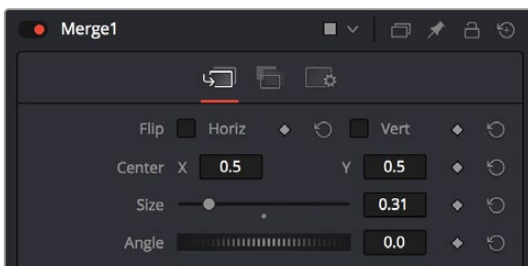
Connecting the Text node to the MatteControl node's Garbage Matte input (top), and the resulting hole punched in the image (bottom)

Customizing Matte Control Nodes

You can use the Inspector to change some parameters to get the result you want. In the Inspector controls for the MatteControl node, revealing the Garbage Matte controls exposes parameters for modifying how the Garbage Matte input is applied to the image. For example, you can choose to have the text mask filled with the image instead of cutting a hole in the image.

Using Transform Controls in the Merge Node

The Merge node includes a set of transform parameters in the Inspector that specifically affect the foreground input's image. This makes it quick and easy to adjust a foreground image to match the background without requiring another node.



The Merge node transform controls that affect the foreground input's image

NOTE: When connecting two images of different sizes to a Merge node, the resolution of the background image defines the output resolution of that node. Keep that in mind when you run into resolution issues.

You can use the Size slider to resize the foreground.



The final composite

Building a Simple Green-Screen Composite

Taking another step forward in compositing, the next example shows how you can equate a multilayered Timeline like the one in DaVinci Resolve's Edit page to nodes in Fusion's node tree. We'll use DaVinci Resolve, but understanding how layers map to nodes can be helpful for anyone new to dealing with a node-based interface. In this example, we'll create a simple composite using a green-screen key and two other layers to create a news story.

Mapping Timeline Layers to Nodes in Fusion

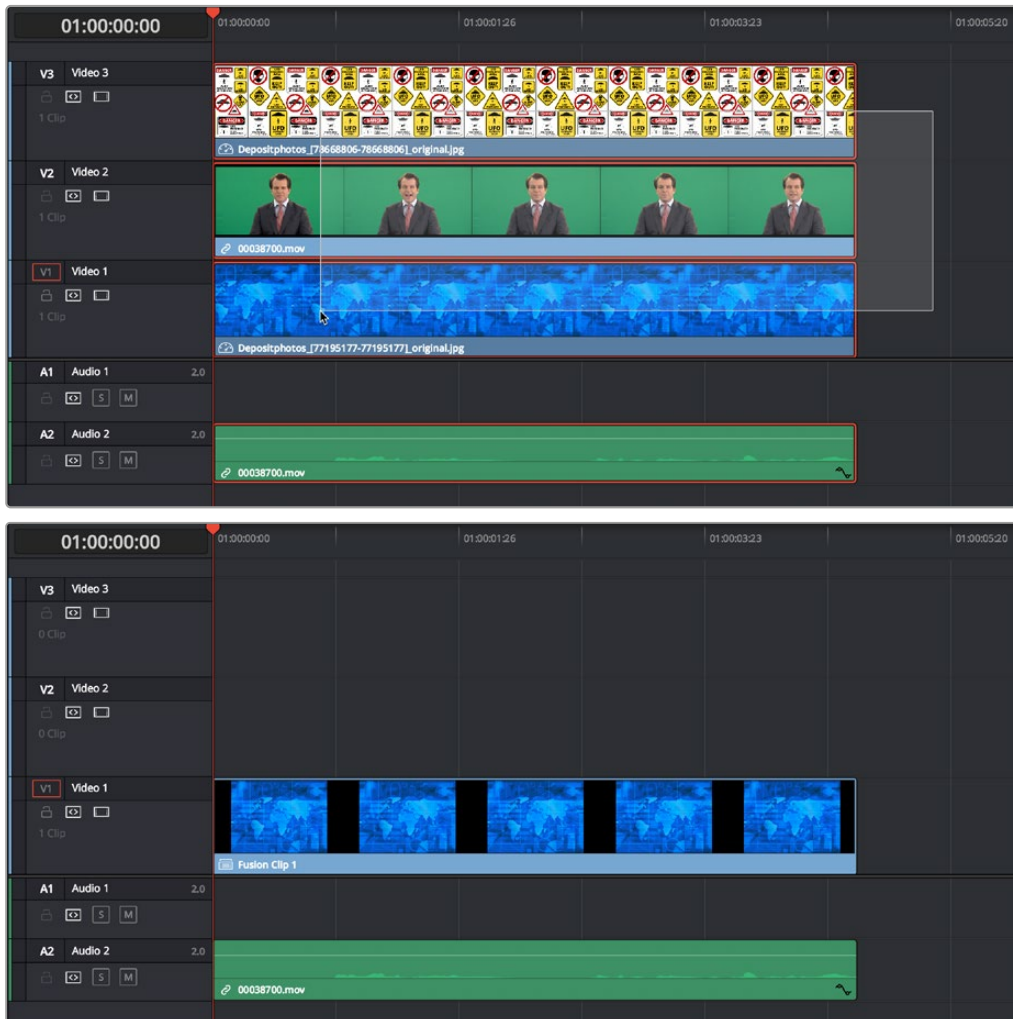
This composite involves three layers in a Timeline. The Timeline consists of a background graphic on video track 1, a green-screen clip on video track 2, and a foreground graphic on video track 3.



Background on video track 1 (top left), green-screen clip on video track 2 (bottom), and graphic file on video track 3 (top right)

Implied in a timeline-based system is that higher numbered video tracks appear as the more forward, or frontmost, element in the viewer. Video track 1 is the background to all other video tracks. Video track 3 is in the foreground to both video track 1 and video track 2.

TIP: If using DaVinci Resolve, you can bring all three layers from the Edit page into Fusion by creating a Fusion clip. For more information on creating Fusion Clips, see Chapter 63, “Getting Clips into Fusion,” in the DaVinci Resolve Reference Manual or Chapter 3 in the Fusion Reference Manual.



A stack of clips to use in a composite (top), and turning that stack into a Fusion clip in DaVinci Resolve's Edit page (bottom)

In Fusion, each video clip is represented by a MediaIn in the Fusion page or a Loader in Fusion Studio.

In our example below, the MediaIn2 is video track 2, and MediaIn 1 is video track 1. These two elements are composited using a Merge node (foreground over background, respectively). The composite of those two elements becomes the output of the first Merge node, which becomes the background to a second Merge. There is no loss of quality or precomposing when you chain Merges together. MediaIn3 represents video track 3 and is the final foreground in the node tree since it is the topmost layer.



The initial node tree of the three clips we turned into a Fusion clip

With this node tree assembled to mimic the video layers, we can focus the rest of this example on adding the nodes we'll need to each branch of this tree to create the green-screen composite.

Pulling a Green-Screen Key Using the Delta Keyer

To make this composite, you need to create transparency behind the newscaster. When working in a node tree, you must become accustomed to rearranging existing nodes to make room for new ones. You'll often move nodes off to the side or up above to make room for the additional nodes.



Creating space after the MediaIn nodes and selecting the second one in preparation for adding a node


The DeltaKeyer node is the main tool used for green-screen keying. It attaches to the output of the node that represents the green screen—in our example, that is the MediaIn2 node. With the MediaIn2 selected, pressing Shift-Space opens the Select Tool dialog where you can search for and insert any node. Below we have added the DeltaKeyer after the MediaIn2 node but prior to being merged with the background.



Adding a DeltaKeyer node inline after the MediaIn2 node

The DeltaKeyer node is a sophisticated keyer that is capable of impressive results by combining different kinds of mattes and a clean-plate layer, but it can also be used very simply if the background that needs to be keyed is well lit. And once the DeltaKeyer creates a key, it embeds the resulting alpha channel in its output, so in this simple case, it's the only node we need to add. It's also worth noting that, although we're using the DeltaKeyer to key a green screen, it's not limited to keying green or blue only; the DeltaKeyer can create impressive keys on any color in your image.

With the DeltaKeyer selected, we'll use the Inspector controls to pull our key by quickly sampling the shade of green from the background of the image. To sample the green-screen color, drag the Eyedropper from the Inspector over the screen color in the viewer.




Dragging the Eyedropper to the viewer samples the screen color

As you drag in the viewer, an analysis of the color picked up by the location of the Eyedropper appears within a floating tooltip, giving some guidance as to which color you're really picking. Meanwhile, if viewing the Merge in a second viewer, we get an immediate preview of the transparency and the image we've connected to the background.




The original image (left), and after sampling the green screen using the Eyedropper from the Inspector (right)

When we're happy with the preview, releasing the pointer button samples the color, and the Inspector controls update to display the value we've chosen.



The DeltaKeyer Inspector updates with the sampled color

No matter how good the composite may look, once you've selected the screen color to pull a key, you need to load the DeltaKeyer node into the viewer itself. This allows you to evaluate the quality or density of the alpha channel created by the key. Above the viewer, click the Color button in the viewer toolbar, or click in the viewer and press C to switch the viewer between the RGB color channels of the image and the alpha channel.



Loading the DeltaKeyer into the viewer and clicking the Color button to view the alpha channel being produced

Black in a matte represents the transparent areas, while white represents the opaque areas. Gray areas represent semi-transparency. Unless you are dealing with glass, smoke, or fog, most mattes should be pure white and pure black with no gray areas. If a close examination of the alpha channel reveals some fringing in the white foreground of the mask, the DeltaKeyer has integrated controls for post-processing of the key and refining the matte. Following is a quick checklist of the primary adjustments to make.


After making the screen selection with the Eyedropper, try the following adjustments to improve the key.

- Adjust the Gain slider to boost the screen color, making it more transparent. This can adversely affect the foreground transparency, so adjust with care.
- Adjust the Balance slider to tint the foreground between the two non-screen colors. For a green screen, this pushes the foreground more toward red or blue, shifting the transparency in the foreground.

Clicking the third of the seven tabs of controls in the DeltaKeyer Inspector opens up a variety of controls for manipulating the matte.


Initial adjustments in the matte tab may include the following parameters:

- Adjust the lower and upper thresholds to increase the density in black and white areas.
- Very subtly adjust the Clean Foreground and Clean Background sliders to fill small holes in the black and white matte. The more you increase these parameters, the more harsh the edges of your matte become.



Adjusting the Clean Foreground slider in the Matte panel of the DeltaKeyer controls

In this case, raising the Clean Foreground slider a bit eliminates the inner fringing we don't want, without noticeably compromising the edges of the key.



The original key (left), and the key after using the Clean Foreground slider (right)

With this accomplished, we're happy with the key, so we load the Merge1 node back into the viewer, and press C to set the Color control of the viewer back to RGB. We can see the graphic in the background, but right now it's too small to cover the whole frame, so we need to make another adjustment.



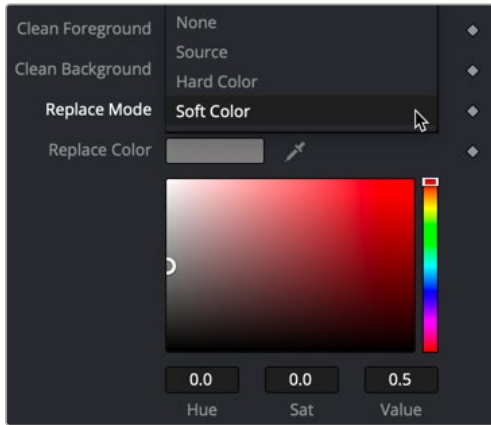
The final key is good, but now we need to work on the background

Dealing with Spill

The DeltaKeyer can handle any screen color that may bounce off the background and land on your subject. This bounce screen light is called spill. However, instead of using one node to handle the matte creation and the spill correction, consider disabling this in the DeltaKeyer and remove the spill using a separate color correction node.

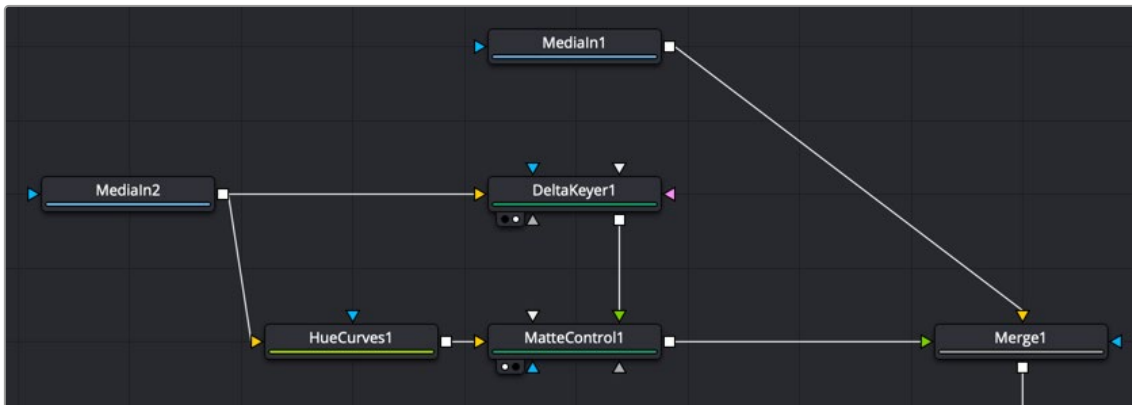
To disable spill suppression in the DeltaKeyer, do the following.

- Select the matte tab in the Inspector.
- From the Replace Mode drop-down menu, choose Source. This uses the color from the original source image, effectively disabling any spill correction in the DeltaKeyer.



The Replace Mode menu determines how spill suppression is handled in the DeltaKeyer

Spill can now be handled using a color correction node placed directly after the DeltaKeyer or branched from the original MediaIn or Loader node and combined with a MatteControl.



Branching the original image with one branch for the DeltaKeyer and a second for color correction


Masking a Graphic

Next, it's time to work on the top video track: the news graphic that will appear to the left of the newscaster. The graphic we will use is actually a sheet of different logos, so we need to cut one out using a mask and position it into place.



A graphic of multiple logos that must be cropped down to isolate just one

The easiest way to crop a MediaIn or Loader node is to add one of the mask shapes from the toolbar directly to it. Selecting the MediaIn or Loader node and clicking the Rectangle mask from the toolbar will crop, or mask off, the graphic.



Masking the logo using a Rectangle mask connected directly to a Merge node


Now, all we need to do is to use the onscreen controls of the Rectangle mask to crop the area we want to use, dragging the position of the mask using the center handle, and resizing it by dragging the top/bottom and left/right handles of the outer border.

As an extra bonus, you can take care of rounded corners when masking a graphic by using the Corner Radius slider in the Inspector controls for the Rectangle mask to add the same kind of rounding.




Moving and resizing the mask to fit our logo, and rounding the edges using the Corner Radius Inspector control

For a simple over-the-shoulder graphic, masking the image may be all you need to do, but masking an image does not change the actual dimensions of the graphic. It only changes the area you see. So, accurately positioning the graphic based on the center of the composite becomes more difficult, and any type of match moving would give incorrect results because the graphic has a different resolution than the background. To fix this resolution mismatch, you can place a Crop node after the MediaIn to change the actual dimensions of the graphic layer.



Adding a Crop node after the masked MediaIn to center the cropped logo on the background

With the Crop node selected, the viewer toolbar includes a Crop tool.



Selecting the crop tool in the viewer toolbar

You can crop the image by dragging a bounding box around it. Unlike a mask which creates a small window you view the image through, a crop effectively changes the resolution of the graphic to the crop bounding box size.



Dragging a bounding box using the Crop tool (left), and the cropped logo now centered on the frame (right)

NOTE: The Resize, Letterbox, and Scale nodes also change the resolution of an image.

At this point, we're all set to move the logo into place. Because the logo is the foreground input to a Merge, you can select the Merge2 node, load it into the viewer, and use the built-in Center X and Y controls or the onscreen controls to place the logo where you want it and make it a suitable size.



Placing the logo using the foreground input transform controls of the Merge2 node

Chapter 80

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	1428
Mask Nodes	1428
Polyline Types	1429
Converting Polylines from One Type to Another	1430
How to Use Masks with Other Nodes	1431
Attaching Masks to an Image for Rotoscoping	1431
Combining Multiple Masks	1433
Mask Inputs on Other Nodes	1433
Creating and Editing Polylines In-Depth	1436
The Polyline Toolbar	1436
Selecting a Specific Polyline	1436
Polyline Creation Modes	1437
Protection Modes	1438
Closing Polylines	1438
Selecting and Adjusting Polylines	1439
Polyline Points Selection	1439
Moving Polyline Points	1439
Smoothing a Polyline Segment	1440
Linearizing a Polyline Segment	1440
Transforming Individual or Multiple Points	1440
Deleting Selected Points	1440

Editing Bézier Handles	1441
Point Editor	1441
Reduce Points	1442
Shape Box	1442
Showing and Hiding Onscreen Polyline Controls	1443
Stop Rendering	1444
Roto Assist	1444
Creating Softness Using Double Polylines	1444
Converting a Single Polyline to a Double Polyline	1445
Adding Softness to a Segment	1446
Adding Additional Points to the Shape	1446
Locking/Unlocking Point Pairs	1446
Animating Polyline Masks	1446
Removing Animation from a Polyline Mask	1447
Adding and Removing Points from an Animated Mask	1447
Publishing Specific Control Points	1447

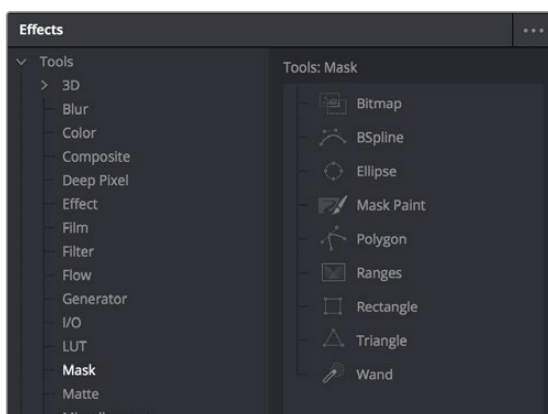
Introduction to Masks and Polylines

Polylines are splines that are used whenever a control is animated with a motion path or when a node's effect is masked with a drawn shape. They are also used in the Paint and Grid Warp nodes. In a more basic form, polylines are used to control the animation in the Spline Editor. Since these splines are used for just about everything, they are extremely flexible, with a considerable amount of controls, modes, and options. This chapter offers an overview of polylines and their operation, with specific information on how to use them for masks.

Mask Nodes

Mask nodes create an image that is used to define transparency in another image. Unlike other image creation nodes in Fusion, mask nodes create a single channel image rather than a full RGBA image. The most used mask tool, the Polygon mask tool, is located in the toolbar.

For more information on these mask tools, see Chapter 106, “Mask Nodes,” in the DaVinci Resolve Reference Manual or Chapter 46 in the Fusion Reference Manual.



The available nodes in the Mask bin of the Effects Library

Polygon Mask

Polygon masks are user-created Bézier shapes. This is the most common type of polyline and the basic workhorse of rotoscoping. Polygon mask tools are automatically set to animate as soon as you add them to the Node Editor.

B-Spline Masks

B-Spline masks are user-created shapes made with polylines that are drawn using the B-Splines. They behave identically to polyline shapes when linear, but when smoothed the control points influence the shape through tension and weight. This generally produces smoother shapes while requiring fewer control points. B-Spline mask tools are automatically set to animate as soon as you add them to the Node Editor.

Bitmap Masks

The Bitmap mask allows images from the Node Editor to act as masks for nodes and effects. Bitmap masks can be based on values from any of the color, alpha, hue, saturation, luminance, and the auxiliary coverage channels of the image. The mask can also be created from the Object or Material ID channels contained in certain 3D-rendered image formats.

Mask Paint

Mask Paint allows a mask to be painted using Fusion's built-in vector paint nodes.

Wand Mask

A Wand mask provides a crosshair that can be positioned in the image. The color of the pixel under the crosshair is used to create a mask, where every contiguous pixel of a similar color is also included in the mask. This type of mask is ideal for isolating color adjustments.

Ellipse, Rectangle, and Triangle Masks

These are primitive shape masks. For more information on these mask tools, see Chapter 106, "Mask Nodes," in the DaVinci Resolve Reference Manual or Chapter 46 in the Fusion Reference Manual.

Ranges Mask


Similar to the Bitmap mask, the Ranges mask allows images from the Node Editor to act as masks for nodes and effects. Instead of creating a simple luminance-based mask from a given channel, Ranges allows spline-based selection of low, mid, and high ranges, similar to to the Color Corrector node.

Polyline Types

You can draw polylines using either B-Spline or Bézier spline types. Which you choose depends on the shape you want to make and your comfort with each spline style.


Bézier Polylines

Bézier polylines are shapes composed of control points and handles. Several points together are used to form the overall shape of a polyline.



Bézier control point with direction handles extended to create a smooth curve

Each control point has a pair of handles used to define the exact shape of the polyline segments passing through each control point. Adjusting the angle or length of the direction handles will affect whether that segment of the polyline is smooth or linear.




Bézier control point with direction handles aligned to create a linear segment

If you're familiar with applications such as Adobe Photoshop or Illustrator, you'll already be familiar with many of the basic concepts of editing Bézier polylines.

B-Spline Polylines

A B-Spline polyline is similar to a Bézier spline; however, these polylines excel at creating smooth shapes. Instead of using a control point and direction handles for smoothness, the B-Spline polyline uses points without direction handles to define a bounding box for the shape. The smoothness of the polyline is determined by the tension of the point, which can be adjusted as needed.



B-Splines excel at creating smooth curves

Converting Polylines from One Type to Another

Just because you created a shape using a B-Spline or polyline, that doesn't mean you're stuck with the controls you started with. You can convert any shape from B-Spline to Bézier, or Bézier to B-Spline, as needed.


To switch a shape between Polyline and B-Spline controls:

- Right-click a shape in the viewer and choose Convert Bézier Spline to B-Spline or Convert B-Spline to Bézier from the spline's contextual menu (only the appropriate option will be displayed).

When converting from one type to another, the original shape is preserved. The new polyline generally has twice as many control points as the original shape to ensure the minimum change to the shape. While animation is also preserved, this conversion process will not always yield perfect results. It's a good idea to review the animation after you convert spline types.

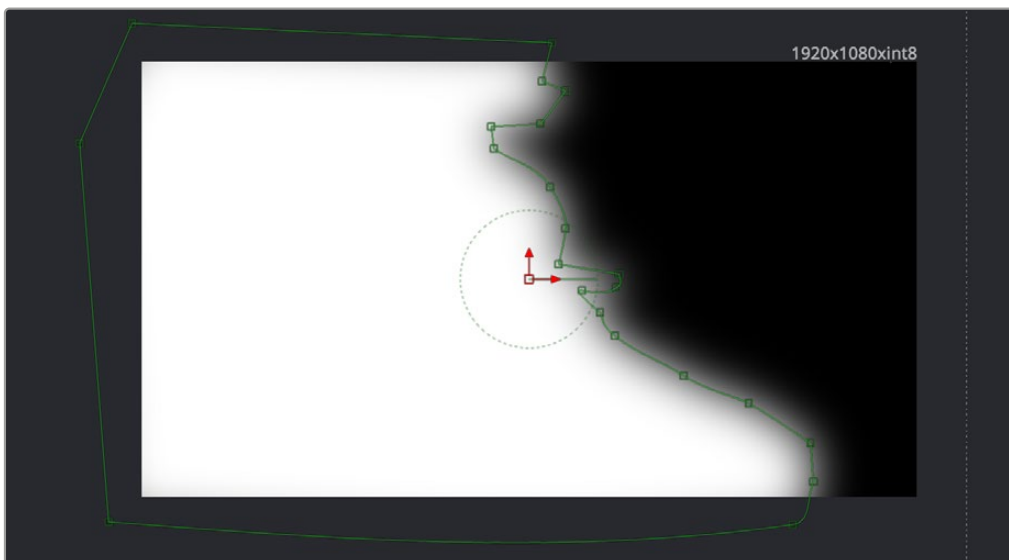
How to Use Masks with Other Nodes

Typically, a node applies its effect to every pixel of an image. However, many nodes have mask inputs that can be used to limit the effect that node has on the image.



A Blur node with a Polygon node masking its effect

Masks are single-channel images that can be used to define which regions of an image you want to affect. Masks can be created using primitive shapes (such as circles and rectangles), complex polyline shapes that are useful for rotoscoping, or by extracting channels from another image.



A Polygon node's mask seen in the viewer

Each mask node is capable of creating a single shape. However, Mask nodes are designed to be added one after the other, so you can combine multiple masks of different kinds to create complex shapes. For example, two masks can be subtracted from a third mask to cut holes into the resulting mask channel.


Fusion offers several different ways you can use masks to accomplish different tasks. You can attach Mask nodes after other nodes in which you want to create transparency, or you can attach Mask nodes directly to the specialized inputs of other nodes to limit or create different kinds of effects.

Attaching Masks to an Image for Rotoscoping

There are two ways you'll typically attach a Mask node, such as a Polygon node, so that it adds an alpha channel to an image for compositing later in the node tree.

Using a MatteControl Node

The MatteControl node is the main node used for combining masks in different ways and inserting the result into an image stream. The MatteControl node is attached downstream of the node outputting the image you want to rotoscope. You'll typically attach a Polygon or B-Spline node to the Garbage Matte input of the MatteControl node to use the spline as an alpha channel.



Feeding a Polygon node to a MatteControl node to perform rotoscoping

To use this setup, you'll load the MatteControl node into the viewer and select the Polygon node to expose its controls so you can draw and modify a spline while viewing the image you're rotoscoping. The MatteControl node's Garbage Matte > Invert checkbox lets you choose which part of the image becomes transparent.


Connecting a Mask to a MediaIn or Loader Node's Input

This method is a bit simpler but requires you to know that you can view one node while adjusting another node, even if that other node is disconnected. If you add an unattached Mask node such as a Polygon or B-Spline node, and then place a MediaIn or Loader node directly into the viewer while selecting the Mask node, you can draw a spline to rotoscope the image.



Rotoscoping a MediaIn node using a disconnected Polygon node

When you're finished rotoscoping, you simply connect the Polygon node's output to the Loader node's input, and an alpha channel is automatically added to that node.

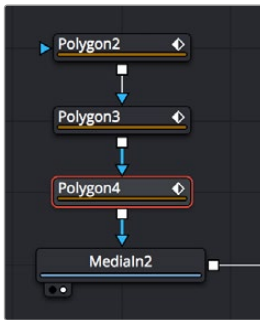


Connecting a Polygon node to a MediaIn node to use a spline as an alpha channel

TIP: If you connect a Mask node to a MediaIn or Loader node's effect input without any shapes drawn, that mask outputs full transparency, so the immediate result is that the image output by the MediaIn or Loader node becomes completely blank. This is why when you want to rotoscope by connecting a mask to the input of a MediaIn or Loader node, you need to work within a disconnected Mask node first. Once the shape you're drawing has been closed, connect the Mask node to the MediaIn or Loader's input, and you're good to go.


Combining Multiple Masks

Masks are designed to be added one after the other, with each Mask node acting as an additional layer of masking.




Combining multiple Polygon nodes one after the other in the node tree

When a Mask node's input is attached to another mask, a Paint Mode drop-down menu appears, which allows you to choose how you want to combine the two masks.



The Paint Mode 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 include to accommodate these different uses. Incidentally, in most cases you can connect either masks or mattes to a mask input to take advantage of that input's functionality.

TIP: If you select a node with an empty effect mask input, adding a Mask node automatically connects to the open effect mask input.

Effects Mask Inputs

Almost every node in Fusion has an Effect mask input (colored blue), which lets you choose which parts of the image will or will not be affected by that node.



A Blur node with a Polygon node masking its effect

While masks (or mattes) are connected via an input, they are actually applied “post effect,” which means the node first applies its effect to the entire image, and then the mask is used to limit the result by copying over unaffected image data from the input.



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 light stands, rigging, and boom microphones that intrude upon masks being pulled via blue-screen and green-screen keys. In the following example, a lighting stand to the left is removed from the image via a B-Spline node's mask connected to the Garbage Matte input of the DeltaKeyer node.




A B-Spline node is connected to the Garbage Matte input of a DeltaKeyer node to eliminate a light stand at the left of frame

TIP: You can quickly add a mask node to the Effect/Solid/Garbage Matte inputs of a keyer node by right-clicking the header bar of that node in the Inspector and choosing whichever mask node you want to use from the Effect Mask, SolidMatte, and GarbageMatte submenus.

You choose whether a garbage matte is applied to a keying node as opaque or transparent in the Inspector for the node to which it's connected.

Solid Matte

Solid Matte inputs (colored white) are intended to fill unwanted holes in a matte, often with a less carefully pulled key producing a dense matte with eroded edges, although you could also use a polygon or mask paint to serve this purpose. In the example below, a gentle key designed to preserve the soft edges of the talent's hair leaves holes in the mask of the woman's face, but using another DeltaKeyer to create a solid matte for the interior of the key that can be eroded to be smaller than the original matte lets you fill the holes while leaving the soft edges alone. This is also sometimes known as a hold-out matte.



Filling in holes in the mask pulled by the DeltaKeyer1 node (left) with another, harder but eroded key in DeltaKeyer2 that's connected to the SolidMatte input of DeltaKeyer1 (right)

Creating and Editing Polylines In-Depth

This section covers the Polygon node's capabilities in depth.

The Polyline Toolbar

Whenever a node that contains one or more polylines is selected, the polyline is shown on all viewers and the Polyline toolbar is displayed along the side of each viewer. The toolbar contains several buttons that make switching polyline modes and options easy to access.



The Polyline toolbar

If you hover the pointer over any of the Polyline toolbar buttons, a tooltip that describes the button's function appears. Clicking on a button will affect the currently active polyline or the selected polyline points, depending on the button.

You can change the size of the toolbar icons, add labels to the buttons, or make other adjustments to the toolbar's appearance in order to make polylines easier to use. All the options can be found by right-clicking on the toolbar and selecting from the options displayed in the contextual menu.

Selecting a Specific Polyline

It is possible to have several polylines in the viewer at once if you select multiple Mask nodes in the Node Editor, so it's important to be able to switch between polylines easily.

To make a polyline active, do one of the following:

- Click one of the polyline’s control points or segments.
- Press Tab and Shift-Tab to cycle between available polylines.
- Right-click in the viewer and choose the desired polyline by name from the Controls > Select menu.


Polyline Creation Modes

There are several different modes available from the toolbar for creating and modifying polylines. The specific mode used when a polyline is first added will depend on whether it is used as a path or a mask.

Each of the modes is described in more detail below.

Click Append

This mode is the default mode for mask creation. It’s used to quickly define the rough shape of the mask, after which you switch to Insert and Modify mode to refine the mask further.



The Click Append toolbar button (Shift-C)

To create a mask using the Click Append mode, do the following:

- 1 Select Click Append from the toolbar or press Shift-C.
- 2 Click the pointer where you want to start the shape.
- 3 Move and click the pointer to append a point to the last one.
- 4 To close the shape, place the mouse pointer over the first point created and click when the pointer changes shape.


When a shape is closed, the polyline is automatically switched to Insert and Modify mode.

Although the Click Append mode is rarely used with paths, it can be helpful when you know the overall shape of a motion path, but you don’t yet know the timing.

TIP: Holding Shift while you draw a mask constrains subsequent points to 45-degree angles relative to the previous point. This can be very helpful when drawing regular geometry.

Insert and Modify

Masks, which are created in Click Append mode, automatically switch to Insert and Modify mode when the mask shape is closed. You can also manually switch to this mode by clicking the Insert and Modify button in the toolbar or using the Shift-I keyboard shortcut. This mode makes it easier to add additional points and refine the shape of the mask. Dragging the control points or direction handles modifies existing points on the polyline.




The Insert Modify toolbar button (Shift-I)

Insert and Modify mode is also the default mode for creating motion paths. A new control point is automatically added to the end of the polyline, extending or refining the path, any time a parameter that is animated with a motion path is moved.

Draw Append

The Draw Append mode creates a freehand polyline shape directly on the viewer, like drawing with a pencil or a paintbrush. This mode is ideal to use in conjunction with a tablet and for the creation of garbage mattes and effect masks around complex shapes.




The Draw Append toolbar button (Shift-D)

Protection Modes

In addition to the modes used to create a polyline, two other modes are used to protect the points from further changes after they have been created.

Modify Only

Modify Only mode allows existing points on the polyline to be modified, but new points may not be added to the shape.




The Modify Only toolbar button (Shift-M)

TIP: Even with Modify Only selected, it is still possible to delete points from a polyline.

Done

The Done mode prohibits the creation of any new points, as well as further modification of any existing points on the polyline.



The Done toolbar button (Shift-N)


Closing Polylines

There are several ways to close a polyline, which will connect the last point to the first.

To close a polyline, do one of the following:

- Hover the pointer over the first point created, and then click on the point.
- Press Shift-O on the keyboard.

- Click the Close button on the polyline toolbar.
- Draw a polyline until you are ready to close the shape, and then right-click and choose Polygon:Polyline > Closed.



The Close toolbar button (Shift-O)

All these options are toggles that can also be used to open a closed polygon.

Selecting and Adjusting Polylines

To create the shape you need for a mask or a motion path, you need to know how to manipulate the splines. Fusion provides a number of simple techniques for selecting, moving, and smoothing a spline, but also includes more complex adjustment techniques for scale, skewing, and twisting a spline.

Polyline Points Selection

To select one or more control points on a polyline, do one of the following:

- Click directly on the control points.
- Lasso around the points.

To add or remove points from the current selection, do one of the following:

- Hold the Shift key to select a continuous range of points.
- Hold Command and click each control point you want to add or remove.
- Press Command-A to select all the points on the active polyline.

TIP: Once a control point is selected, you can press Page Down or Page Up on the keyboard to select the next control point in a clockwise or counterclockwise rotation. This can be very helpful when control points are very close to each other.

Moving Polyline Points

The selected polyline points can be moved using either the keyboard or the mouse.

To move selected control points using the pointer, do one of the following:

- Drag on the selected points anywhere in the viewer.
- Hold Shift while dragging to restrict movement to a single axis.
- Hold Option and drag anywhere in the viewer to move the selected control point.

To move selected control points using the keyboard, do one of the following:


- Press the Up or Down Arrow keys on the keyboard to nudge a point up or down in the viewer.
- Hold Command-Up or Down Arrow keys to move in smaller increments.
- Hold Shift-Up or Down Arrow keys to move in larger increments.

Smoothing a Polyline Segment

If you want to shape the polyline and control its slope, you can choose to smooth a spline segment by adjusting the Bézier direction handles.

To smooth the selected points on an active polyline, do one of the following:

- Press Shift-S.
- Click the Smooth button on the Polyline toolbar.
- Choose Smooth from the polyline’s contextual menu.



The Smooth button in the toolbar (Shift-S)

Linearizing a Polyline Segment

To make certain that a polyline segment is perfectly straight, that segment must be linearized. A linear segment aligns the Bézier direction handles with the segment and therefore has no curvatures. The segment is always drawn in a straight line between two points on the polyline.

To linearize the selected points on an active polyline, do one of the following:

- Press Shift-L.
- Click the Linear button on the polyline’s toolbar.
- Choose Linear from the polyline’s contextual menu.



The Linear button in the toolbar (Shift-L)

Transforming Individual or Multiple Points

Select the points to be transformed, and then do one of the following:

- Hold T and drag to twist.
- Hold S and drag to scale.
- Hold X and drag to scale horizontally only.
- Hold Y and drag to scale vertically only.
- Hold O and drag to offset the points perpendicular to the tangent.

The position of the pointer when the transformation begins becomes the center used for the transformation.

Deleting Selected Points

You can delete a selected point or group of points by pressing Delete or Backspace, choosing Delete from the contextual menu, or by clicking the Delete Point button in the toolbar. The shape of the polyline changes to reflect the removal of these points.


TIP: Deleting all the points in a polyline does not delete the polyline itself. To delete a polyline, you must delete the node or modifier that created the polyline.

Editing Bézier Handles

For Bézier polylines, each control point has two direction handles that adjust the slope of a curve through the control point. These direction handles appear only when the point is selected.

Dragging a direction handle makes adjustments to the curve of the segment that emerges from the control point. The direction handle on the opposing side of the control point will also move to maintain the relationship between these two handles.

To break the relationship between direction handles and adjust one independently, hold Command while dragging a handle. Subsequent changes will maintain the relationship, unless Command is held during each adjustment.




Hold Command to adjust one handle independently

If you want to adjust the length of a handle without changing the angle, hold Shift while moving a direction handle.

Point Editor

The Point Editor dialog can be used to reposition control points using precise X and Y coordinates. Pressing the E key on the keyboard will bring up the Point Editor dialog and allow you to reposition one or more selected control points.




The Point Editor dialog can be used to position control points

The dialog box contains the X- and Y-axis values for that point. Entering new values in those boxes repositions the control point. When multiple control points are selected, all the points move to the same position. This is useful for aligning control points along the X- or Y-axis.

If more than one point is selected, a pair of radio buttons at the top of the dialog box determines whether adjustments are made to all selected points or to just one. If the Individual option is selected, the affected point is displayed in the viewer with a larger box. If the selected point is incorrect, you can use the Next and Previous buttons that appear at the bottom of the dialog to change the selection.

In addition to absolute values for the X- and Y-axis, you can adjust points using relative values from their current position. Clicking once on the label for the axis will change the value to an offset value. The label will change from X to X-offset or from Y to Y-offset.




The Point Editor dialog with Offset values

If you are not sure of the exact value, you can also perform mathematical equations in the dialog box. For example, typing 1.0-5 will move the point to 0.5 along the given axis.

Reduce Points

When freehand drawing a polyline or an editable paint stroke, the spline is often created using more control points than you need to efficiently make the shape. If you choose Reduce Points from the polyline's contextual menu or toolbar, a dialog box will open allowing you to decrease the number of points used to create the polyline.



The Reduce Points button in the toolbar


The overall shape will be maintained while eliminating redundant control points from the path. When the value is 100, no points are removed from the spline. As you drag the slider to the left, you reduce the number of points in the path.

Shape Box

If you have a polyline shape or a group of control points you want to scale, stretch, squish, skew, or move, you can use the shape box to easily perform these operations.

To enable the shape box, do one of the following:


- Click the Shape Box toolbar button.
- Choose Shape Box from the contextual menu.
- Press Shift-B.



The Shape Box button in the Polyline toolbar

If there are selected points on the polyline when the Shape Box mode is enabled, the shape box is drawn around those points. Otherwise, you can drag the shape box around the area of control points you want to include.

If you want to freely resize the shape box horizontally and vertically, you can drag a corner handle. Dragging a handle on the side of the shape box resizes the polyline along a specific axis.



Dragging a side handle resizes along a specific axis


Holding Command while dragging a shape box handle will apply adjustments from the center of the shape box, constraining the transformation to the existing proportions of the shape box. Holding Shift while dragging a corner handle affects only that handle, allowing skewed and non-uniform transformations.



Hold Shift while dragging a corner to perform non-uniform transformations

Showing and Hiding Onscreen Polyline Controls

It is often difficult to identify individual points when they are placed closely together. You can choose to display both points and their direction handles, just points, or just handles. These display mode options are selected using the Show Key Points and Show Handles toolbar buttons, or from the polyline's context menu.



The Show Key Points and Show Handles buttons in the toolbar

You use these options to simplify the screen display when adjusting control points placed closely together and to avoid accidentally modifying controls and handles that are adjacent to the intended target.

Stop Rendering


While points along the polyline are being moved, the results are rendered to the viewer to provide constant interactive feedback. Although extremely useful, there are situations where this can be distracting and can slow down performance on a complex effect. To disable this behavior so renders happen only when the points stop moving, you can toggle the Stop Rendering button in the toolbar or select this option from the polyline contextual menu.

Roto Assist

You can enable the Roto Assist button in the toolbar when you begin drawing your shape to have points snap to the closest high-contrast edge as you draw the shape. The points that have snapped to an edge are indicated by a cyan outline.

There are three main Roto Assist options.

- **Multiple Points:** Allows adding multiple points along an entire edge with a single click instead of having to add each point individually.
- **Distance:** Defines the pixel range within which searching for an edge will take place.
- **Reset:** Used for resetting the snap attribute of the snapped points. After resetting, the points will become unavailable for tracking.




The Roto Assist options in the toolbar

Creating Softness Using Double Polylines

The standard soft edge control available in all Mask nodes softens the entire mask equally. However, there are times, particularly with a lot of motion blur, when softening part of the curve while keeping other portions of the curve sharp is required.

This form of softness is called non-uniform softness, which is accomplished by converting the shape from a single polyline to a double polyline. The double polyline is composed of two shapes: an inner and an outer shape. The inner shape is the original shape from the single polyline, whereas the outer shape is used to determine the spread of the softness. The further the outer shape gets from the inner shape, the softer that segment of the shape becomes.




A double polyline uses an inner and outer shape for non-uniform softness

Converting a Single Polyline to a Double Polyline

To convert a mask into a double polyline, click the Double Polyline button in the Polyline toolbar or right-click in the viewer and select Make Outer Polyline from the mask's contextual menu.

The shape will be converted into an inner and an outer polyline spline. Both polylines start with exactly the same shape as the original single polyline. This keeps the mask sharp to start with and allows any animation that may have already been applied to the shape to remain.



Make Double Polyline button

The control points on the outer shape are automatically parented to their matching points on the inner shape. This means that any changes made to the inner shape will also be made to the outer shape. The relationship is one-way; adjustments to the outer shape can be made without affecting the inner shape.

A dashed line drawn between the points indicates the relationship between the points on the inner and outer shapes.

Adding Softness to a Segment

The outer shape is drawn using a green dashed line instead of a solid line to help distinguish it from the inner shape.

To select the outer soft edge shape, do one of the following:

- Use the Tab key to cycle between the onscreen controls until the dashed outline is visible
- Right-click over a spline in the view and choose Controls > Select > Polygon: Outer Polygon.

Once the outer polyline is selected, you can drag any of the points away from the inner polyline to add some softness to the mask.

TIP: Press Shift-A to select all the points on a shape, and then hold O and drag to offset the points from the inner shape. This gives you a starting point to edit the falloff.

The farther the outer shape segment is from the inner shape, the larger the falloff will be in that area.

Adding Additional Points to the Shape

It is not necessary for every point on the inner shape to have a match on the outer shape, or vice versa. You can add additional control points to refine the shape of either shape.

Each polyline stores its animation separately; however, if a point is adjusted on the inner shape that is parented to a point on the outer shape, a keyframe will be set for both splines. Adjusting a parented point on the outer shape only sets a keyframe for the outer shape's spline. If a point that is not parented is adjusted, it will only set a keyframe on the relevant spline. You can disable this behavior entirely for this polyline by selecting Polygon: Outer Polygon > Follow Inner Polyline from the contextual menu.

Locking/Unlocking Point Pairs

If you want to parent additional control points, you can select the points, right-click in the viewer, and choose Lock Point Pairs from the contextual menu for either spline. This will cause the selected point on the outer shape to become parented to the selected point on the inner shape.

Any animation already applied to either point is preserved when the points become parented.

To unlock a point so it is no longer parented, select the point, right-click in the viewer, and deselect Lock Point Pairs from the contextual menu.

Animating Polyline Masks

Animating masks is surprisingly easy. When Polygon or B-Spline masks are added to the Node Editor, the spline's control points are automatically ready to be animated. All you have to do to animate a mask is move the playhead to a new frame and then change the shape of the mask. A new keyframe is added in the Spline Editor and Timeline Editor. This one keyframe controls the position of all control points for that mask at that frame. Once two or more keyframes have been created, the shape of the polygon or B-Spline is automatically interpolated from one keyframe to the next.

TIP: The center point and rotation of a shape are not auto-animated. Only the control points are automatically animated. To animate the center position or rotation, enable keyframes for that parameter in the Inspector.

To adjust the overall timing of the mask animation, you edit the Keyframe horizontal position spline using the Spline Editor or Timeline Editor. Additional points can be added to the mask at any point to refine the shape as areas of the image become more detailed.

Removing Animation from a Polyline Mask

If you want a Polyline mask to remain static, you can remove the automatic animation setting. In the Inspector for the mask, right-click in the bottom of the panel where it says Right Click Here For Shape Animation. From the contextual menu, choose Remove Bézier Spline. If you decide you need to animate the mask at a later time, right-click in the same area again and choose Animate.

Adding and Removing Points from an Animated Mask

When adding points to an animated mask, the new point is fit into the shape at all keyframes. Deleting a point removes that point from all keyframes in the animated mask.

Publishing Specific Control Points

Although you can rapidly animate the entire shape of a polyline using a single keyframe, by default the Spline Editor and Timeline display only one keyframe for the entire shape at any given frame.


This default keyframing behavior is convenient when quickly animating shapes from one form to another, but it doesn't allow for specific individual control points that need to be keyframed independently of all other control points for a particular shape. If you're working on a complex mask that would benefit from more precise timing or interpolation of individual control points, you can expose one or more specific control points on a polyline by publishing them.

Be aware that publishing a control point on a polyline removes that point from the standard animation spline. From that point forward, that control point can only be animated via its own keyframes on its own animation spline. Once removed, this point will not be connected to paths, modifiers, expressions, or trackers that are connected to the main polyline spline.

To publish a selected point or points, do one of the following:

- Click on the Publish Points button in the Polyline toolbar.
- Select Publish Points from the Polyline's contextual menu.

A new coordinate control is added to the Polyline mask controls for each published point, named Point 0, Point 1, and so on.



The Publish Points controls in the Inspector

The onscreen control indicates published points on the polyline by drawing that control point much larger. Once a published point is created, it can be connected to a tracker, path, expression, or modifier by right-clicking on this control and selecting the desired option from the point's contextual menu.



The published point in the viewer


Using “Publish to Path” to Preserve Animation

When a point is published, any animation already applied to that point is removed. However, if you need to keep the animation, you can use the “Publish to Path” option. This Polyline contextual menu option publishes the selected points and converts their existing animation to a path. You can also use the Publish to Path button in the Polyline toolbar.

Using “Follow Published Points” to Add Points

There are times when you will need to have control points that lie between two other published points follow the motion of the published points, while still maintaining their relative offset and shape. For this reason, points in a Polyline mask can be set to “Follow Published Points” using the Polyline's contextual menu.

When a point of an effect mask is set to follow points, the point will be drawn as a diamond shape rather than a small box.



A control point set to Follow Published Points

When this mode is enabled, the new “following” control points will maintain their position relative to the motion of any published points in the mask, while attempting to maintain the shape of that segment of the mask. Unlike published points, the position of the following points can still be animated to allow for morphing of that segment's shape over time.

Chapter 81

Paint

This chapter describes how to use Fusion's non-destructive Paint tool to repair images, remove objects, and add creative elements.

Contents

Paint Overview	1450
Types of Paint Nodes	1450
Setting Up the Paint Node	1450
Setting the Paint Node's Resolution	1451
Paint Node Workflow	1451
Select the Correct Paint Stroke Type	1452
Setting the Brush Size	1454
Choosing an Apply Mode	1454
Editing Paint Strokes	1458
Editing Paint Strokes in the Modifiers Tab	1459
Deleting Strokes	1459
Animating and Tracking Paint Strokes	1460
Animating with Write-On Controls	1460
Tracking a Paint Stroke	1460
Using the Planar Tracker with the Paint Tool	1462
Inverting the Steady Effect to Put the Motion Back In	1468
Painting a Clean Plate	1469

Paint Overview


The Paint node is a procedural paint tool, which means that each paint stroke is a live, editable object that's drawn with properties that you can mix and match to address a wide variety of painting tasks. You can use it to paint masks, retouch images, perform beauty work, clone out objects, or even create motion graphics. Each element of a paint stroke can be altered long after you apply it. Since strokes are editable, you can apply, change, ignore, delete, and even reorder them in a node tree.

Types of Paint Nodes

There are two types of paint nodes in Fusion. The Paint node is a full-featured creative and retouch vector-based paint tool that requires an input on which to paint. The Mask Paint node lets you specifically paint an alpha channel to limit the area of an effect. It allows you to create paint strokes on an alpha channel without needing to have an input.

- The Paint node is located in the Paint category of the Effects Library.
- The Mask Paint node is located in the Mask category of the Effects Library.

The main difference between these two Paint tools is that the Mask Paint tool only paints on the Alpha channel, so there are no channel selector buttons. The Paint tool can paint on any or all channels. The majority of this chapter covers the Paint node, since it shares identical parameters and settings with the Mask Paint node.



The Paint node available in the Paint category of the Effects Library

Setting Up the Paint Node

The Paint node has two inputs. Typical of most Fusion nodes, the orange input background is the primary input for connecting the “canvas” or image to paint on, while the second blue input is an Effect Mask. Unlike the Mask Paint node, the Paint node requires a background input to begin painting.

Setting the Paint Node's Resolution

No matter how you use the Paint node in your node tree, the Paint node assumes the resolution of the background input image as your working resolution for that operation. Although the Paint tool is in fact resolution independent and you can change this assigned resolution at any time, it's essential to make sure you properly set the resolution of the media you're working with, because it will affect operations such as motion tracking that you'll want to use as part of your paint workflow.

Because of this, the Paint node requires a background input to set the resolution of the "canvas" you'll be painting upon. To do this, you can set up a Paint node in the node tree in one of two ways: painting directly on an image or using Paint as the foreground.

Painting Directly on an Image


You can connect the image you want to paint on directly to the background input of the Paint node. This is the easiest and cleanest node tree construction, but it doesn't provide much in terms of compositing flexibility.



The Paint node is inserted directly after the node it is painting on

Using Paint as the Foreground in a Merge Composite

An alternative setup is to use a Background node to set the resolution for the Paint node, compositing the result over the actual background using a Merge node. Working this way lets you use the Merge node's Apply Mode setting (also referred to as composite modes) to control how your paint strokes are composited against the image, but it does require a bit more setup.



The Paint node is composited over the image you want to paint on using the merge

Setting this up requires some configuration of the nodes. The Background node must be fully transparent and, unless you are doing something simple like using the Stroke tool set to Color to paint over an image, you must drag the image you want to clone or smudge into the Source Tool field in the Paint node's inspector. These steps are described in more detail later in this chapter.

Paint Node Workflow

You begin painting by first selecting the paint stroke type from the Paint toolbar above the viewer. There are ten stroke types to choose from as well as two additional tools for selecting and grouping paint strokes. The stroke types and tools are described in detail in Chapter 111, "Paint Node" in the DaVinci Resolve Reference Manual or Chapter 51 in the Fusion Reference Manual. The primary tool for

painting and cloning is the Stroke tool. The Stroke tool is a fully animatable and editable vector-based paint stroke and initially uses a duration of the entire global range.



The Stroke tool is most often used for cloning, beauty work, and creative paint

Select the Correct Paint Stroke Type

Instead of having multiple dedicated Paint nodes for different operations, the procedural nature of Fusion's Paint tool means there is just one tool with a variety of different stroke types. Each stroke type serves a different purpose. The Paint toolbar above the viewer allows you to choose between different stroke types and drawing tools. These are grouped into a few categories.

Single-Frame Stroke Types

The Multistroke and Clone Multistroke are explicitly designed for single-frame retouching paint jobs like removing raindrops in a shot that's supposed to be a sunny day or removing scratch marks and dust when restoring vintage content. When handling these types of jobs, these two Multistroke options are faster than the other stroke types, but they're not editable later on. This means you must set up the size and function of the brush as well as the duration for the paint stroke before you paint.

The duration of each of these stroke types is one frame by default, but this can be changed using the Stroke Duration slider in the Inspector. Multistroke and Clone Multistroke are basically the same tools, except Clone Multistroke automatically configures the tool for cloning. In contrast, the Multistroke requires you to set up the tool for cloning manually.



The Multistroke and Clone Multistroke are non-editable single frame stroke types

Editable Stroke Types

The Stroke and Polyline are similar in that they can be modified and animated at any time. Also, they both begin with a duration lasting the entire comp, but that also can be changed using the Keyframes Editor.




The Stroke and Polyline strokes are editable and last for the entire comp

The Stroke Tool

One of the most flexible editable stroke types you'll use for many tasks is the Stroke, because it is fully animatable and editable. You can animate all elements of the Stroke, and you can use the Write-on/Write-off parameters to control how the stroke appears onscreen. You can also connect to a tracker from the Center point of the Stroke if you want to make the stroke follow specific onscreen motion.

By default, the Stroke type does not expose control points for the shape of the path. You can move and track the center and rotation of the Stroke, but the individual control points that create the spline are hidden. To reveal the control points, you can open the Stroke Controls at the bottom of the Inspector and click the Make Editable button.




The Stroke's control points can be revealed using the Make Editable button

Although the Stroke type is the most flexible, that flexibility can come at a performance penalty if you're painting hundreds of strokes on a frame. For larger numbers of strokes that do not need to be animated, it's better to use Multistroke or Clone Multistroke, as these are more processor efficient.

The Polyline Stroke Tool

The Polyline Stroke acts more like a drawing tool than a paintbrush. It includes the same functionality as the Stroke tool, except that it is created not by dragging or "painting" like a paintbrush, but by clicking to create a spline path, as you do with masks and motion paths. Without even creating a stroke in the viewer, the Polyline Stroke can connect to existing polylines like a mask or a motion path.

If a motion path is published, right-clicking on the Shape Animation label at the bottom of the Polyline Stroke's Stroke Controls allows you to use the Connect To menu to assume the shape of a motion path or mask. You can also use this method if you import SVG graphics and want to "paint-on" the outlines.



The Polyline Stroke Shape Animation label

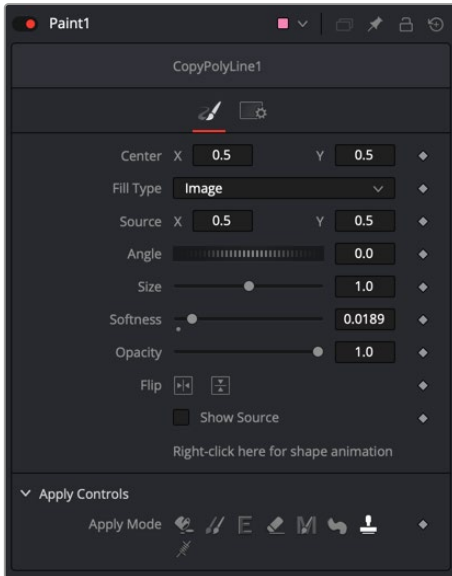
Shape Drawing Tools

Five shape-based drawing tools allow you to draw shapes and either fill them with a color or clone an area from a source image. All these tools act similarly to the Stroke and Polyline stroke type in that they are editable at any time and have a default duration spanning the entire global range of the comp. However, you can edit the duration at any time in the Keyframes Editor.



The Shape Strokes are used to create shapes or clone areas based on shapes

All of the Copy [Shape Name] stroke types require that you connect the source node you are cloning from into the Paint node, and set the Fill Type menu to Image.



The Copy Shapes require a Source to be directly connected to the Paint tool and set to Image

Setting the Brush Size

After selecting the Stroke type, the Brush size can be set in the Inspector or more intuitively in the viewer. With the Paint node selected in the Node Editor and the pointer positioned over the viewer, you can see an outline of the current brush size. To change the brush size, hold down the Command key and drag. The circle changes size, so you can set it relative to other objects you may be painting over.



Brush size can be changed interactively in the viewer

Choosing an Apply Mode

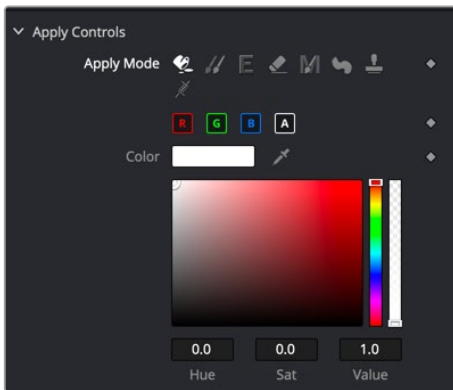
The Apply Mode buttons determine the functionality of the paintbrush. There are eight Apply modes that set the brush to do things like paint a color, clone from a source, smudge an area, or remove thin wires.



Apply Mode buttons determine the paint brush functionality

Picking a Paint Color

There are several ways to pick the paint color and opacity for a colored brush stroke. You use the Fill button in the row of Apply modes when you want to paint with a solid color.



The Color swatch shows the current color
Clicking it opens the OS Color Picker window

To select a color for the paint brush, do one of the following:

- Click the color swatch to open a standard OS Color Picker window.
- Drag the Eyedropper into the viewer.
- Drag inside the color chooser to select a saturation and luminance. Drag on the sidebars to change the hue and transparency.

When you paint, each stroke is unpremultiplied, so adjusting the Alpha slider in the Inspector does not affect what you apply to the RGB channels. However, changing opacity affects all four channels.

Cloning from the Frame

Choosing the Clone Apply Mode allows you to paint from one area of an image over another area. This is the most common use of the Paint tool. It allows you to remove objects or artifacts from a clip by covering them up with another area of the frame. Depending on the Stroke type chosen, you may clone on either a single frame or for the entire duration of the clip.



The Clone Apply Mode allows you to sample from one area and use it as a source to paint over another area

You can use the Clone Apply Mode to clone from the same image connected to the Paint node's background input or a different source from the node tree.

To clone from a different area of the same frame:

- 1 Select a Stroke type from the Paint toolbar above the viewer.
- 2 Using the size slider in the Brush controls section of the Inspector, set the size of the brush.
- 3 From the Apply Mode buttons, select the Clone mode.
- 4 Option-click over the area in the viewer you want to use as the source. A dot appears showing you the center of what you're sampling from.



The Clone source starting area identified by the X and the paint brush size represented by the circle

- 5 Paint over the area you want to cover up using the source pixels.




The Clone completed after selecting the source area and painting over the flag pole

When trying to erase objects or artifacts from a clip using the Clone Apply Mode, it can sometimes be easier if you sample from a different frame on the same clip. This works well when the object you are trying to clone out moves during the clip, revealing the area behind the object. Sampling from a different frame can use the revealed background by offsetting the source frame.

To clone from a different frame of the same clip:

- 1 Select a Stroke type from the Paint toolbar above the viewer.
- 2 Set the size of the brush.
- 3 From the Apply Mode buttons, select the Clone mode.
- 4 Drag the clip (MediaIn or Loader) from the Node Editor into the Source Tool field in the Inspector.



The Paint Inspector with the MediaIn1 dragged from the Node Editor into the Source Tool field

- 5 Click the Overlay checkbox to see the current frame and the offset frame superimposed.
- 6 Drag the Time Offset slider to select the source frame you want to use.
- 7 Option-click over the area in the viewer you want to use as the source or to offset the source frame's position.




Overlay shows two frames overlapped with Time Offset, allowing you to clone from one frame onto another

- 8 Paint over the area you want to cover up using the source pixels.



The plane is half painted out using on the Overlay with Time Offset

- 9 Disable the Overlay checkbox.




The Clone Apply Mode can use a different frame from the same clip

TIP: When using a Clone Apply Mode, you can hold down the O key instead of clicking the Overlay checkbox in Inspector to see the Overlay. Releasing the O key will return to normal viewing without the Overlay.

Editing Paint Strokes

Once you've painted using the Stroke or Polyline stroke type, you can change the look of the stroke by selecting it and updating the parameters in the Inspector. Selecting the stroke requires you to switch to the selection tool in the Paint toolbar above the viewer. Using the Paint node's selection tool, you can either click once on a stroke or drag a bounding box around a stroke to select it for editing.

TIP: To select multiple strokes, you can Shift-click or Command-click to select and deselect multiple specific strokes, or you can drag a selection box around all strokes you want to select.



The Stroke or Polyline Stroke type can be edited by selecting the stroke in the viewer


Although you can make changes in the Tools tab in the Inspector, the Paint node uses both the Tools tab and the Modifiers tab. In the Tools tab, you can create new brush strokes and select a stroke in the viewer to edit. The Modifiers tab presents a list of all the strokes for the selected Paint node, which makes it easy to modify any previously created paint stroke.

NOTE: Multistroke and Clone Multistroke each only appear as one item in the Modifiers tab no matter how many strokes you create using those tools. Those two tools are not editable after creating them.

The same controls you used in the Tools tab to create the strokes are located in the Modifier's tab to modify them. You can also animate each individual stroke.

Editing Paint Strokes in the Modifiers Tab

When you paint a stroke, the settings for that stroke appear in the Inspector's Modifiers tab. You can then change the settings in the Tools tab for the next stroke you are about to paint. Each time you click, drag, and release the pointer button, you create a new stroke. Each stroke is numbered in the Modifiers tab, where it can be selected and edited.



Each stroke is listed in the Modifiers tab for editing

Once you stop painting a stroke, it's added to the Modifiers tab along with an additional Stroke modifier that represents the next stroke. For instance, if you paint your first stroke, the Modifiers tab shows your stroke as Stroke1 and then a stroke 2 as well, which represents the next stroke you create. You always have one more stroke in the Modifiers tab than strokes in the viewer.

Deleting Strokes

There are two ways you can delete paint strokes.

To delete any individual stroke, do the following:

- 1 Select the Paint node.
- 2 Click the Modifiers tab.
- 3 Right-click over the Stroke header you want to delete and choose Delete from the menu.

To delete all paint strokes you've made on every frame, do one of the following:

- Click the reset button in the upper-right corner of the Inspector.
- Delete the Paint node in the Node Editor.

Animating and Tracking Paint Strokes

In some ways, animating paint strokes is no different than animating any other effect in the Inspector. Each parameter that can be animated includes a gray diamond Keyframe button along the right side. Clicking the Keyframe button sets a keyframe on the current frame and enables auto-keyframe mode for the parameter. However, more commonly, the paint stroke is tracked using one of Fusion's trackers, or for motion graphics, animated using the Write-On Start and End sliders.

Animating with Write-On Controls

The Stroke and Polyline stroke types include Write-On controls located in the Stroke controls section of the Inspector. These Write-On controls animate the appearance of a stroke along the path. You can animate the Write-On controls using either the Stroke Animation drop-down menu or using the Start and End sliders.

Stroke Animation Drop-Down Menu

The Stroke Animation drop-down menu includes six options for auto-animating a paint stroke. The first two options do not truly animate the stroke as much as set a duration. The Limited Duration option uses the Duration slider to set the number of frames the stroke is onscreen.

To auto-animate the stroke, you can choose one of the three Write options or the Trail option.

Choosing Write On automatically creates a write-on animation. The duration is set by two keyframes that get added when you choose Write On from the menu. The Start keyframe is set on the frame where you first created the stroke. The End keyframe is added on the current frame when you choose Write On from the menu. The remaining options in the menu set their Start and End keyframes similarly but change the direction of the animation based on the menu selection.

Write-On Start and End Parameters

The Write-On Start and End parameters allow you to manually control the start point and end point along any stroke's path, and use keyframes to animate each parameter individually. The Start parameter determines the point at which the stroke begins, measured as a percentage offset from the beginning of the stroke's path. For example, a Start value of 50 moves the starting point of the stroke to the middle of the stroke's path. The End parameter works the same way but from the other end of the stroke. You can animate a stroke onscreen, creating a handwriting effect by setting keyframes for the End parameter from 0 to 100 over several frames.

Tracking a Paint Stroke


You can animate the position of a paint stroke using any of Fusion's trackers. For instance, if you cloned out a flag pole from a clip, but the camera moves, you can track the flag pole and attach the resulting path to the paint stroke.



Trackers can be attached to the Center parameters of a paint stroke


To attach a tracker to a paint stroke:

- 1 With the Paint node, select the Stroke brush type and clone out an object on a frame.
- 2 In the Paint toolbar above the viewer, click the Select tool.
- 3 Drag a selection box around the stroke to select it.
- 4 Right-click the center control on the stroke, and then choose Stroke1:Center > Modify With > Tracker Position.



Right-clicking over the Stroke's Center control in the view allows you to apply a tracker modifier

- 5 Click the Modifiers tab to view the Tracker controls.
- 6 From the Node Editor, drag the Medialn for the image you painted on, and drag it into the Tracker Source field in the Inspector.



Drag the Medialn you want to track into the Tracker Source field in the Inspector


- 7 Click the Track Forward button.
- 8 After tracking, at the bottom of the Inspector, use the Tracker 1 X Offset/Y Offset controls to reposition the paint stroke, if necessary.

Tracking a Group of Paint Strokes

You can assign a tracker to multiple strokes by adding the stroke to a group and connecting the tracker to the group. Instead of connecting each individual stroke, the group's center is used for all of the strokes. Assuming the motion of each object is consistent in the same direction, as it would be objects "nailed to the set", then applying the tracker to the group makes cloning multiple objects out with a single paint node very easy.


To group paint strokes, do the following:

- 1 Drag a bounding box, Shift-click, or Command-click to select every stroke that you want to group together.
- 2 Click the Paint Group button in the Paint toolbar.



Selecting all the strokes and then clicking the Paint Group button collects all the strokes into a single group

The group's onscreen controls replace the controls for each paint stroke, and the Modifiers tab in the Inspector shows the group's parameters. The individual strokes are still editable by selecting Show Subgroup Controls in the Modifiers tab of the Inspector. The group then comes with a Center, Angle, and Size control for connecting to a tracker.



Use the onscreen controls, or in the Modifiers tab, right-click over the Paint Group's Center X label to connect the tracker

Using the Planar Tracker with the Paint Tool

Here's an example that dives deeper into a workflow where we use the Paint tool with the Planar Tracker for retouching a clip. We'll eliminate some facial scars on an actor's forehead in a commercial by combining the Paint node with the PlanarTracker node, illustrating a common way of using these two powerful tools together.



The actor has some scars on his forehead that the director would like painted out

Because this is a clip in motion, we can't just paint out the scars on the man's forehead; we need to deal with the motion so that the paint work we do stays put on his face. In this case, a common workflow is to analyze the motion in the image and use it to apply a "steady" operation, pinning down the area we want to paint in place so we can paint on an unmoving surface.

Setting Up the Planar Tracker for Stabilization

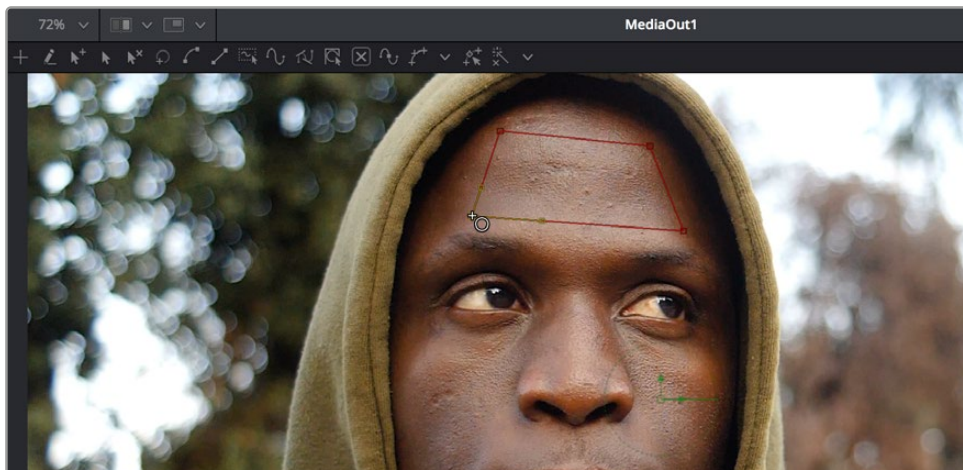
The best way to do this is to use the Planar Tracker, so we'll add the PlanarTracker node after the MediaIn1 node, such that the image connects to the background input of the PlanarTracker node. As always, it's important to be careful about which input you connect for the effect to work properly.



Adding a PlanarTracker node to analyze and steady the part of the image we want to paint on


With the PlanarTracker node selected and loaded in the viewer, a viewer toolbar appears with a variety of tools for drawing shapes and manipulating tracking data. The Planar Tracker works by tracking flat surfaces that you define by drawing a shape around the feature you want to track. When you first create a PlanarTracker node, you can immediately begin drawing a shape, so in this case, we draw a simple polygon over the man's forehead since that's the feature we want to steady in preparation for painting.

We draw a simple box by clicking once each on each corner of the man's forehead to create control points, and then clicking the first one we created to close the shape.



Drawing a shape over the man's forehead to prepare for Planar Tracking

In the Inspector, the PlanarTracker node has tracking transport controls that are similar to those of the Tracker. However, there are two buttons, Set and Go, underneath the Operation Mode menu, which defaults to Track, since that's the first thing we need to do. The Set button lets you choose which frame to use as the "reference frame" for tracking, so you click the Set button first before clicking the Track Forward button below.



Setting a reference frame at the beginning of the range of frames we want to track

TIP: The Set button lets you supervise a Planar Track in progress and stop it if you see it slipping, making adjustments as necessary before clicking Set at the new frame to set a new reference before continuing to track forward towards the end of the clip.

The Pattern controls let you set up how you want to handle the analysis. Of these controls, the Motion Type menu is perhaps the most important. In this particular case, Perspective tracking is the analysis we want. Still, in other situations, you may find you get better results with the Translation, Translation/Rotation, and Translation/Rotation/Scale options.


Once you initiate the track, a series of dots appears within the track region shape you created to indicate trackable pixels found. A green progress bar at the bottom of the Timeline ruler lets you see how much of the shot is remaining to track.



Clicking the Track from First Frame button to set the Planar Track in progress; green dots on the image and a green progress bar let you know the track is happening

NOTE: If nothing happens when you track, or it starts to track and then stops, that's your cue that there isn't enough trackable detail within the shape you've drawn for the Planar Tracker to work, and your best bet is to choose a different location of the image to track.

Once the track is complete, you can set the Operation Mode of the PlanarTracker node's controls in the Inspector to Steady.



Setting the PlanarTracker node to Steady

You'll immediately see the image warped as much as is necessary to pin the tracked region in place for whatever operation you want to perform. If you scrub through the clip, you should see that the image dynamically cornerpin-warps as much as is necessary to keep the forehead region within the shape you drew pinned in place. In this case, this sets up the man's head as a canvas for paint.

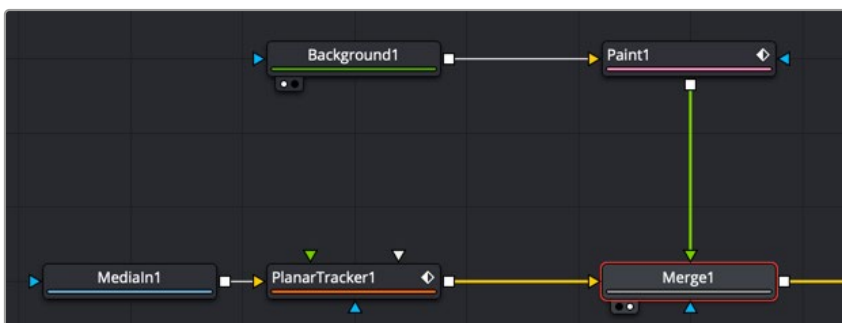


Steadying the image results in warping as the forehead is pinned in place for painting

At this point, you're ready to paint out those scars.

Connecting the Paint Node

Although you could paint directly on the image by connecting the Paint node after the Planar Tracker, it gives you more control over the process if you merge the Paint node over the top of the steadied image. To do that, you'll add a Merge after the Planar Tracker and then connect a Background node into the foreground of the Merge. Insert a Paint node after the background, and you are ready for this process.



Paint node merged over the top of a steadied Planar Tracker

Making the Background Node Transparent


The Background node must be fully transparent; otherwise, you are painting on the solid color instead of the Merge background image. To make the background transparent, drag the Alpha channel sidebar all the way up in the Color Picker.



The Color Picker in the Background node set to be transparent

Selecting the Stroke and Clone Mode

With the Paint node selected and set up, the next thing we want to do is to select the Stroke tool. The Stroke tool is the tool of choice when you want to paint out features or paint in fixes to subjects within the frame that need to remain in place for the entire shot.



Choosing the Stroke tool from the Paint node's tools in the viewer toolbar

Next, choose the Clone mode from the Apply Controls. In this example, we'll clone part of the man's face over the scars to get rid of them. Choosing the Clone mode switches the controls of the Paint node to those used for cloning.



Choosing the Clone mode in the Inspector

Setting a Clone Source

The Paint node clones from the connected Input image unless you instruct it otherwise. If you are just painting a color over the background image, then you are good to go. However, if you plan on using the merge's steadied background image for cloning with the Paint tool, you must set that image as the Paint node's source. To set the clone source for the Paint node, you drag the PlanarTracker node into the Source Tool field in the Inspector.



Any node can be dragged into the Source Tool field when cloning with the Paint tool

With the Stroke tool selected in the Paint toolbar, the Clone mode selected in the Inspector controls, and the Source for cloning added to the Source Tool field, we're ready to start painting. If we move the pointer over the viewer, a circle shows us the paint tool, ready to go.

To use the clone brush, first hold down the Option key and click somewhere on the image to identify the source area of the clone. In this example, we'll sample from just below the first scar we want to paint. After Option-clicking to sample the image, you can click to begin painting anywhere in the frame.



Setting an offset to sample for cloning (left), and dragging to draw a clone stroke (right)

If you don't like the stroke you've created, you can undo with Command-Z and try again. We repeat the process with the other scar on the man's forehead, possibly adding a few other small strokes to make sure there are no noticeable edges, and in a few seconds, we've taken care of the issue.



Original image (left), and after painting out two scars on the man's forehead with the Stroke tool set to Clone

TIP: You can adjust the size of the brush right in the viewer, if necessary, by holding down the Command key and dragging the pointer left and right. You'll see the brush outline change size as you do this.

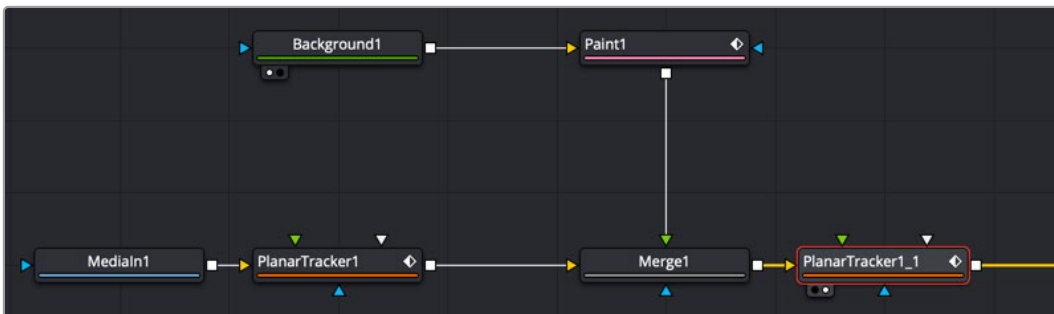
Inverting the Steady Effect to Put the Motion Back In

At this point, scrubbing through the clip shows that the paint strokes we've made are indeed sticking to the man's forehead as we need them to do. Now we just have to invert the transform that the Planar Tracker applied to restore the clip back to the way it was, except now with the painted fix attached in the process. This ends up being a two-part process, but the first part is the simplest.




Scrubbing through the steadied clip shows the paint fix is “sticking” to the man's forehead

We select and copy the PlanarTracker node coming before the Merge node, and paste a copy of it after. This copy has all the analysis and tracking data of the original PlanarTracker node.



Pasting a second copy of the PlanarTracker node after the Paint node

With the second PlanarTracker node selected, we go into the Inspector and turn on the Invert Steady Transform checkbox, which inverts the steady warp transform to restore the image back to the way it was.



Turning on Invert Steady Transform to try to restore the image back to the way it was

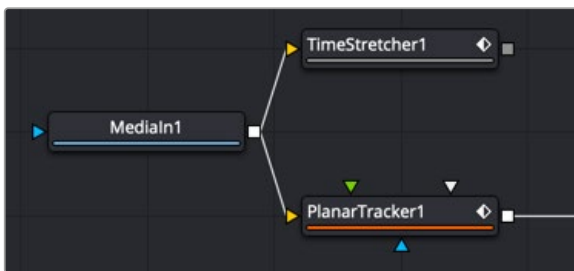
This is just one example of how to set up a Planar Tracker and Paint node. In some instances, you made need to do more work with masks and layering, but the above example gives you a good starting point.

Painting a Clean Plate

On simple clips, planar tracking the clone paint strokes may work fine. In other cases, you may not be able to steady the clip, or the strokes might appear like they are “bubbling.” Paint is just a single stroke repeated over multiple frames. Each painted frame has the potential to appear differently than the one before or after. Often, the more reliable way to use paint is by creating what is referred to as a “clean plate.” Create one good painted frame and freeze it. Then use the Planar Tracker to track the cleaned area over the top of the original. The benefit is that there is only a single paint stroke instead of repeated paint strokes that potentially show unwanted artifacts. For this next example, we continue with the previous clip, but now the technique is to freeze, mask, and composite just a single frame of the fixed forehead over the original clip.

Creating a Freeze Frame and Clean Plate

As in the previous example, you still track the clip using the Planar Tracker. After tracking, you can branch out from the MediaIn (or Loader node if you are using Fusion Studio) into a Time Stretcher to create the freeze frame.



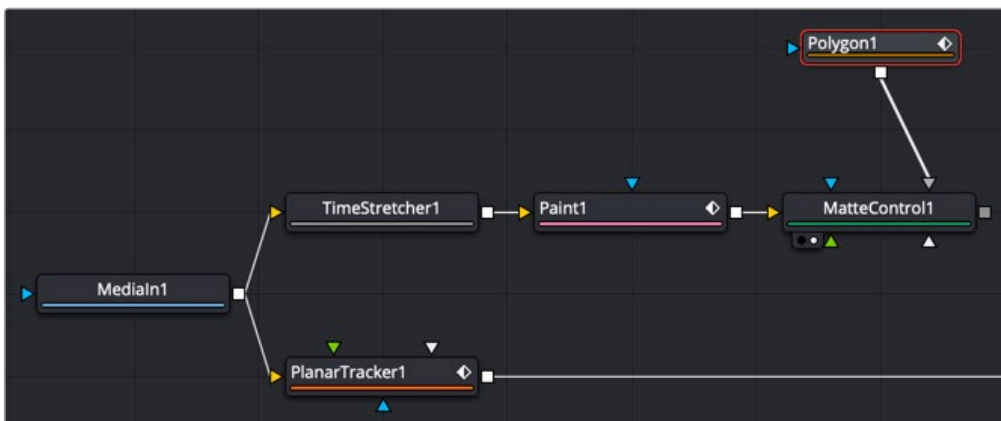
Time Remap is used to retime or freeze a clip

Disable the default Keyframe in the Time Stretcher and enter the frame you want to freeze. If you have already performed a Planar Track, then entering the frame you set as the Reference Frame is usually a good frame to freeze.

To create the clean plate, you connect the paint node to the output of the Time Stretcher. Clone over the areas you want to hide, and you now have a single clean frame. Now you need to composite the clean area over the original.

Isolating the Painted Forehead

First, we need to mask out just the man’s painted forehead. We can do this by connecting a MatteControl node to the output of the Paint node and then connect a Polygon node to the garbage matte input. This lets us draw a shape with the Polygon node and use it as a mask to crop out the man’s painted forehead.



Add a MatteControl node with a garbage mask to cut out the painted forehead

TIP: When it comes to using masks to create transparency, there are a variety of ways to connect one—for example, (a) attach the image to the background input of a Brightness/Contrast node and attach a Polygon mask node to the effect mask input. On the Brightness/Contrast node, enable the Alpha channel and lower the Gain slider to darken a hole, or (b) using Channel Booleans to copy channel data to the alpha from a Polygon node attached to the foreground input and the image attached to the background input.


Drawing a Polygon Mask

After moving the playhead to the first frame of the clip, we're ready to draw a mask to isolate the "clean plate" forehead. Loading the MatteControl1 into the viewer and selecting the Polygon1 node so that we see its tools in the viewer toolbar sets us up for drawing a polygon.

Drawing shapes using the Polygon node is similar to shape drawing in other spline-based environments, including the Color page:

- Clicking once draws a corner control point.
- Clicking and dragging creates a Bézier curve.
- Clicking the first control point you created closes a shape.

We click and drag to create a shape that outlines the man's forehead, and when we close the shape, we see exactly the opposite of what we want, a hole in the middle of the image.

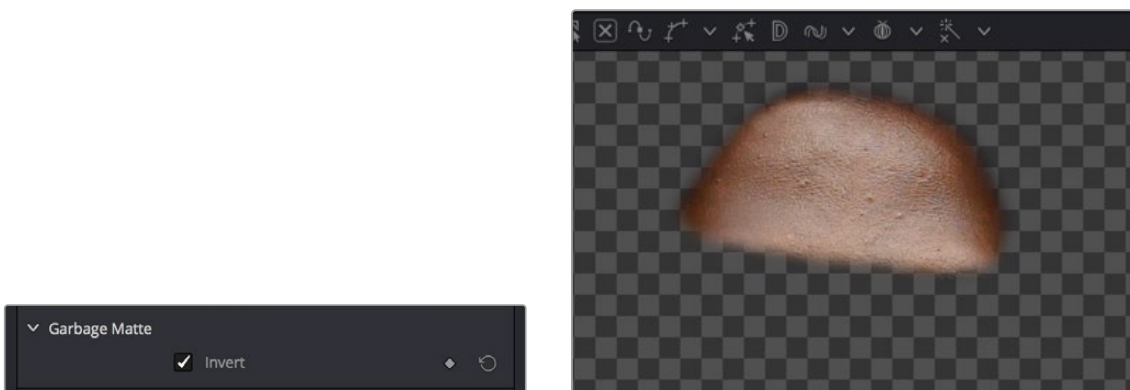


Drawing a shape to isolate the forehead gives an inverted result at first when using the Garbage Matte input of the MatteControl node to attach the Polygon to the MatteControl node

Before fixing this, we drag the Soft Edge slider in the Inspector to the right to blur the edges just a bit.

Inverting the Garbage Input

Selecting the MatteControl1 node, we open the Garbage Matte controls and click the Invert checkbox, which immediately gives us the result we want, of the forehead in isolation, ready for compositing.




Inverting the Garbage Matte input (Left), and the resulting inverted mask inverting the forehead (right)

Compositing the Painted Forehead Against the Original Image

Almost finished. Next, we'll add a Merge node to layer the fixed forehead over the original image being output by the MediaIn node.

We create a Merge node connected to the output of the PlanarTracker node, and then we connect the MatteControl's output to the green foreground input of the Merge node. This puts the cropped and fixed forehead on top of the original image.



The painted forehead composited against the original image

Match Moving the Mask to the Shot

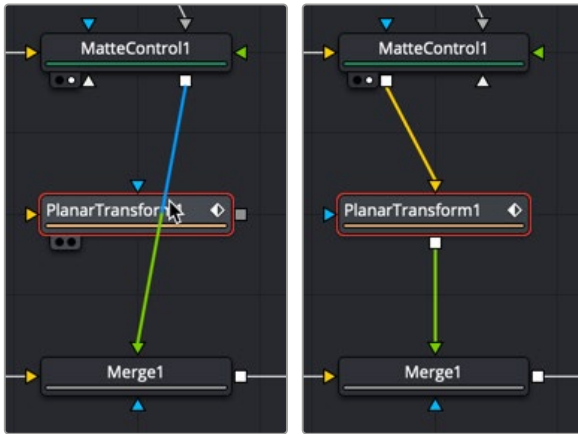
Now we have the best of both worlds: a fixed forehead and the background of the shot looking good. However, if we select the Polygon node and then scrub forward in the clip, the fixed forehead mask drifts out of sync with the motion of the shot, so we have one remaining issue. Fortunately, match moving the mask to move with the shot is really simple.



Because the Polygon isn't animated to match the motion of the shot, it goes out of sync

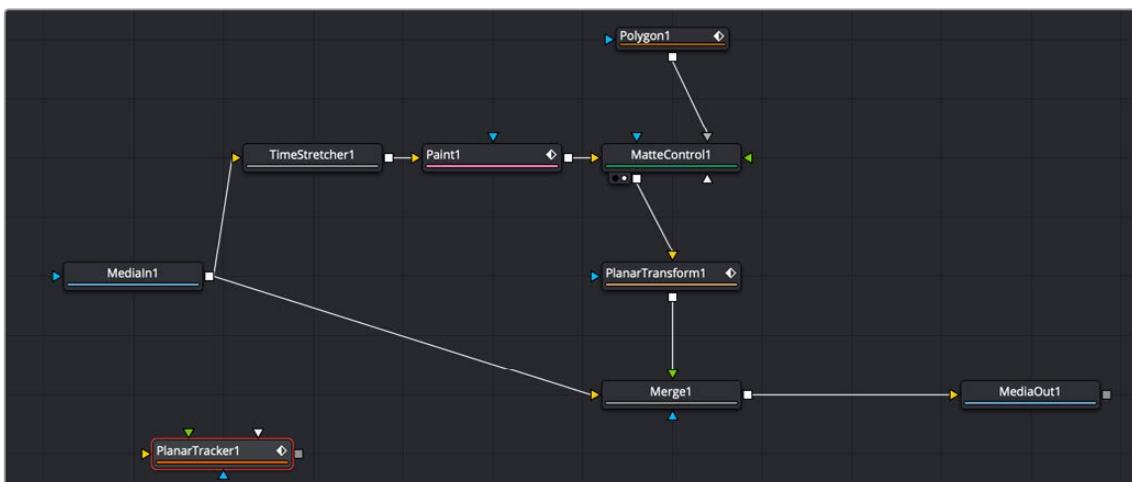
Selecting the first PlanarTracker node that comes right after the MediaIn node, and choosing Track from the Operation Mode menu, reveals a Create Planar Transform button at the bottom of the listed controls. Clicking this button creates a new, disconnected Planar Transform node in the Node Editor, which has the transforms from the Planar Tracker baked in. Unlike the Planar Tracker, the Planar Transform is intended for use with masks. With the Planar Transform created, there is no longer a use for the Planar Tracker, and it can be disconnected or deleted from the Node Editor.

We can insert this new node into the node tree to use it by holding down the Shift key and dragging the node over the connection between the Polygon node and the MatteControl node, dropping it when the connection highlights.



Inserting a PlanarTransform node by holding down the Shift key while dropping over a connection (left), and after inserting the PlanarTransform node (right)

With the new Planar Transform node inserted, the Polygon automatically moves to match the motion of the forehead that was tracked by the original PlanarTracker node, and it animates to follow along with the movement of the shot. At this point, we're finished!



The final painted image, along with the final node tree

Chapter 82

Using the Tracker Node

This chapter shows the many capabilities of the Tracker node in Fusion, starting with how trackers can be connected in your node trees, and finishing with the different tasks that can be performed.

Contents

Introduction to Tracking	1474
Tracker Node Overview	1475
Modes of the Tracker Node	1475
Basic Tracker Node Operation	1475
Connect to a Tracker's Background Input	1475
Analyze the Image to be Tracked	1476
Apply the Tracking Data	1477
Viewing Tracking Data in the Spline Editor	1479
Tracker Inspector Controls	1479
Motion Tracking Workflow In Depth	1481
Connect the Image to Track	1481
Add Trackers	1481
Refine the Search Area	1483
Perform the Track Analysis	1484
Tips for Choosing a Good Pattern	1485
Using the Pattern Flipbooks	1486
Using Adaptive Pattern Tracking	1486
Dealing with Obscured Patterns	1487
Dealing with Patterns That Leave the Frame	1487
Setting Up Tracker Offsets	1488

Stabilizing with the Tracker Node	1489
Stabilization Using the Tracker Match Move Mode	1489
Smoothing Motion	1490
Using the Tracker Node for Match Moving	1491
Simple Match Moving	1491
Corner Positioning Operations	1491
Perspective Positioning Operations	1492
Connecting to Trackers' Operations	1492
Steady Position	1492
Steady Angle	1492
Offset Position	1493
Unsteady Position	1493
Steady Size	1493
Using the Outputs of a Tracker	1493
Using the Tracker as a Modifier	1495
Match Moving Text Example	1497
Adding a Layer to Match Move	1497
Setting Up Motion Tracking	1498
A Simple Tracking Workflow	1499
Connecting Motion Track Data to Match Move	1502
Offsetting the Position of a Match Moved Image	1503

Introduction to Tracking

Tracking is one of the most useful and essential techniques available to a compositor. It can be roughly defined as the creation of a motion path from analyzing a specific area in a clip over time. Fusion includes a variety of different tracking nodes that let you analyze different kinds of motion. Once you have tracked motion on a clip, you can then use the resulting data for stabilization, motion smoothing, matching the motion of one object to that of another, and a host of other essential tasks.

Types of tracking nodes in Fusion:

- **Tracker:** Follows a relatively small, identifiable feature or pattern in a clip to derive a 2D motion path. This is sometimes referred to as point tracking.
- **Planar Tracker:** Follows a flat, unvarying surface area in a clip to derive a 2 ½D motion path including perspective. A planar tracker is also more tolerant than a point tracker when some tracked pixels move offscreen or become obscured.
- **Camera Tracker:** Tracks multiple points or patterns in a clip and performs a more sophisticated analysis by comparing those moving patterns. The result is a precise recreation of the live-action camera in virtual 3D space.

Each tracker type has its own chapter in this manual. This chapter covers the tracking techniques with the Tracker node.

Tracker Node Overview

The Tracker node is a single node that actually performs tracking, stabilizing, matching moving, and corner-pinning operations. Since the Tracker node can transform the foreground input, it can be used to generate tracks and then operate as a Merge in a match move or corner-pin setup. Or you can use it to produce tracking data only and then publish that data to other nodes in the Node Editor.

Modes of the Tracker Node

The Tracker node is an incredibly flexible tool often used multiple times in a composite to help with dozens of tasks. However, most of those tasks can be boiled down into just a few operations. The Tracker node has four operation modes that cover the majority of tracking situations.

Stabilizing

You can use one or more tracked patterns to remove all the motion from the sequence or to smooth out vibration and shakiness. When you use a single tracker pattern to stabilize, you stabilize only the X and Y position. Using multiple patterns together, you are able to stabilize position, rotation, and scaling.

Match Moving

The reverse of stabilizing is match moving, which detects position, rotation, and scaling in a clip using one or more patterns. Instead of removing that motion, it is applied to another image so that the two images can be composited together.

Corner Positioning

Corner positioning tracks four patterns that are then used to map the four corners of a new foreground into the background. This technique is generally used to replace signs or mobile phone screens. The Planar Tracker node is often a better first choice for these types of tracking tasks.

Perspective Positioning

Perspective positioning again tracks four patterns to identify the four corners of a rectangle. Each corner is then mapped to a corner of the image, rescaling and warping the image to remove all apparent perspective. The Planar Tracker node is often a better first choice for removing perspective from a clip.

Basic Tracker Node Operation

All tracking workflows consist of three fundamental steps.

- 1 Attach an image you want to track to the yellow background input of the Tracker node.
- 2 Set the tracking pattern and analyze the clip to create a path.
- 3 Apply the tracking data to stabilize, match move, corner pin, or remove perspective.

Connect to a Tracker's Background Input


You start by connecting the output of the image you want to track to a Tracker node's background input. The Tracker node analyzes the image that's attached to its background input.

You can insert the Tracker node serially with other nodes if you intend to use the Tracker node itself to do a simple stabilization operation or if you want to use it to perform the function of a Merge node in a match move or corner-pin operation.



Tracker node connected serially so it can both track and transform the input images

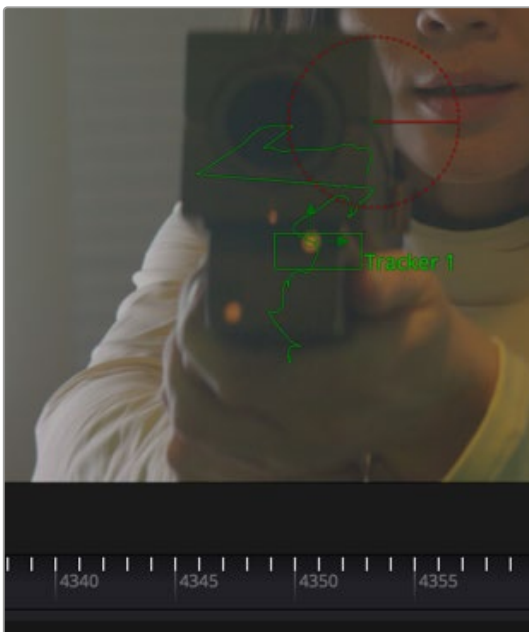
However, if you're just using a Tracker node to analyze data for use with multiple nodes elsewhere in the comp, you could choose to branch it and leave its output disconnected to indicate that Tracker node is a data repository. Please note that this is not necessary; serially connected Tracker nodes can be linked to multiple other nodes as well.



Tracker connected as a branch to indicate it is linked to other nodes and not used directly

Analyze the Image to be Tracked

After constructing the node tree and inserting the Tracker where you want, you can set up the tracker in the viewer. You identify one or more features in the image that you wish to track (referred to as patterns) by adding trackers (there's one by default) and positioning them using the onscreen controls in the viewer. After the Tracker node analyzes the clip, the resulting tracking data is stored within that Tracker node. Keyframes, one per frame, indicate the Tracked Center X and Y data that has been saved, while a motion path shows the path of tracked data in the viewer.



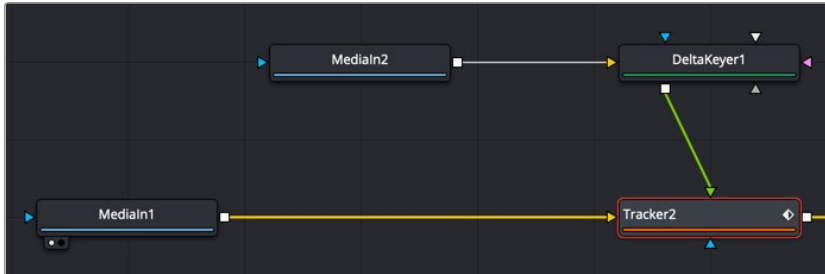
A motion path that indicates the tracked motion, and tic marks that indicate tracking data keyframes

Apply the Tracking Data

The resulting tracking data stored within the Tracker node is used to stabilize, match move, corner pin, or remove perspective in one of two ways.

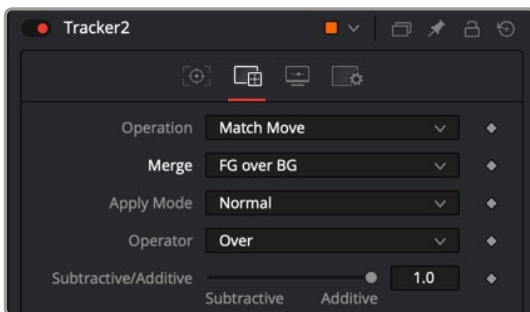
Method 1: Use the Tracker node to match move and merge

You can connect a foreground image to the Tracker node and apply the motion from the analyzed background image.



Using a Tracker node in line for a match move


Setting the Operation parameter in the Operation tab in the Inspector to Match Move, Corner Position, or Perspective Position always applies the motion to the foreground input (if one is connected). This is an easy workflow for simple situations. In this scenario, you can use the Tracker node to replace a Merge node since Tracker nodes include all the same functionality as a Merge.



Using a Tracker node to do a match move and merge, all in one


Method 2: Connect specific parameters to the Tracker Node

Alternatively, you can connect the tracking data from the Tracker node to the specific parameters of other nodes that will actually do the work, for instances where setting up a match move isn't just a matter of transforming a foreground image. Each Tracker node and each pattern within the Tracker node publishes its data for other nodes to use without directly linking to them in the node tree. For example, in the following node tree, an Ellipse node is being used to isolate a glow effect for the ray gun prop.




The Tracker set up as a branch and connected using the Connect To menu

The ellipse needs to follow the motion of the ray gun, so a Tracker node is used to analyze the movement of the gun tip so that tracking data can be used to animate the ellipse. The ellipse is not connected to the tracker directly via the foreground input but indirectly through the Connect To contextual menu.



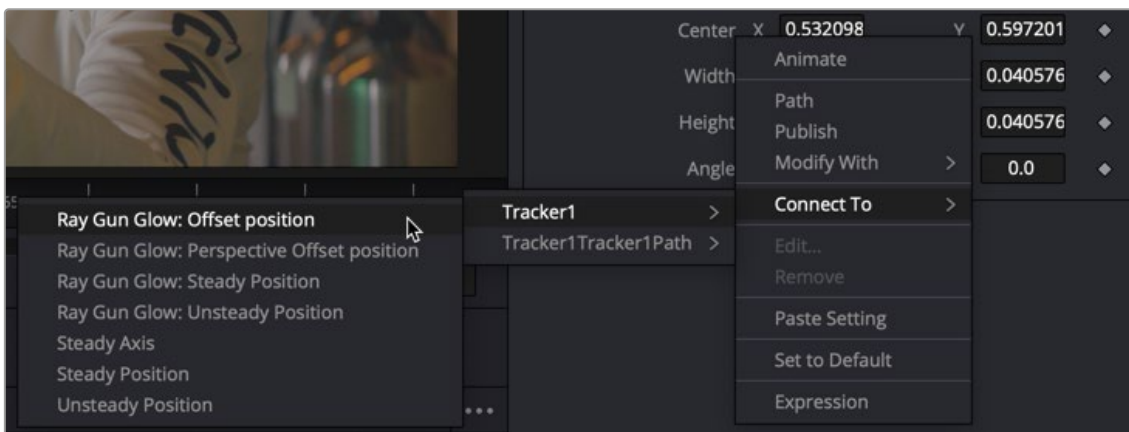
Applying the light of a ray gun by connecting tracking data to the center position of an Ellipse node

This is made easier by renaming the Tracker you created to something descriptive of what's being tracked.



You can rename trackers in the Tracker List by double-clicking them and typing descriptive text

Once the tip of the ray gun has been tracked, this tracking data is then connected to the Center parameter of an Ellipse node that's limiting a Glow effect by right-clicking the label of the Center parameter in the Inspector, and choosing Tracker1 > Ray Gun Glow: Offset position from the Connect to submenu of the contextual menu. All the data from every Tracker node in your node tree and every tracking pattern appears within this submenu, and since we named the Tracker, it's easy to find. Choosing Offset position because it will place the center of the ellipse directly over the path. However, it also gives us the flexibility to offset the ellipse if need be, using the offset controls in the Inspector.



Right-clicking a parameter's label lets you connect tracking data to it to animate it

You can connect the data from a Tracker node to any other node's parameter; however, you'll most typically connect track data to center, pivot, or corner X/Y style parameters. When you use tracking data this way, it's not necessary to connect the output of the Tracker node itself to anything else in your node tree; the data is passed from the Tracker to the Center parameter by linking it with the Connect To submenu.

Viewing Tracking Data in the Spline Editor

Tracking data can be seen in the viewer as a path or as a displacement spline in the Spline Editor. You can manipulate the tracking data in either place.

The Tracker uses a displacement spline by default that indicates how far the tracking point is based on the original location. It is great for modifying velocity, but it doesn't tell you anything about direction. If you need to nudge a few points in a certain direction, you can convert the displacement spline to an X and Y coordinate spline.



Right-clicking a parameter's label lets you connect tracking data to it to animate it


Right-click in the viewer to bring up a contextual menu. At the very bottom is a reference to the path the Tracker created, called Tracker1Tracker1Path:Polyline. Choosing it calls up a longer submenu where you can choose Convert to XY Path.

For more information on Displacement Splines, see Chapter 70, "Animating in Fusion's Spline Editor," in the DaVinci Resolve Reference Manual or Chapter 10 in the Fusion Reference Manual.

Tracker Inspector Controls


The layout of the Tracker node's tabs in the Inspector reflects this workflow. It's divided into three main Tracker tabs, as well as the common Settings tab.

- **The Tracker Control tab:** This is where you create onscreen trackers with which to target patterns, and where the controls appear that let you perform the required track analysis.




The Tracker Control tab

- **The Operations tab:** This is where you decide how the tracking data is used.



The Tracker Operations tab

- **The Display Options tab:** This is where you can customize how the onscreen controls look in the viewer.



The Tracker Display Options tab

Motion Tracking Workflow In Depth

Tracker nodes serve two purposes. They provide a method to analyze an object you want to follow, and they serve as a container for the resulting track data. This allows you to use one node for analysis and to pass on that analysis to any other node that requires it. Following is a more detailed breakdown of the tracking process.

Connect the Image to Track


Regardless of whether you actually use the Tracker node itself to do anything with the tracking data, the image you want to track must be connected to the background input (yellow) of a Tracker node for there to be a successful analysis. While the Tracker node has a foreground input, it is initially ignored for purposes of tracking analysis, so even if you connect a foreground, the background is the only input used during the analysis process.



Connecting a MediaIn node's output to a Tracker node's background input

Add Trackers

Although each Tracker node starts with a single tracker pattern, a single node is capable of analyzing multiple tracking patterns that have been added to the Tracker List, enabling you to track multiple features of an image all at once for later use and to enable different kinds of transforms. Additional trackers can be added by clicking the Add button immediately above the Tracker List control.



To add an additional tracker, click the Add button in the Inspector

Multiple patterns are useful when stabilizing, match moving, or removing perspective from a clip. They also help to keep the Node Editor from becoming cluttered by collecting into a single node what would otherwise require several nodes.

Working in the Tracker list:


- **To select a tracker:** Click the name of the Tracker you want to select.

- **To rename a tracker:** You can rename trackers to make it easier to reference them later. For example, if you're tracking a car door handle, you can name the Tracker "Car Handle" so it's easy to find later. To do so, just double-click the default name of the Tracker in the Tracker list, type a new one, and press Return.
- **To delete a tracker:** Select a tracker and click the Delete button.
- **To disable, suspend, or re-enable trackers:** Click the checkbox to the left of each tracker's name in the Tracker List. It's a three-way toggle that switches between Enabled, Suspended, and Disabled.
- **Enabled:** An Enabled tracker will re-track its pattern every time the track is performed. Its path data is available for use by other nodes, and the data is available for stabilization and corner positioning.
- **Suspended:** When the checkbox appears gray, it is Suspended. In this state, it does not re-track its pattern when the track is performed. The data is locked to prevent additional changes, but the data from the path is still available for other nodes. The data is also available for advanced tracking modes like stabilization and corner positioning.
- **Disabled:** A Disabled tracker does not create a path when tracking is performed. Its data is not available to other nodes or for advanced tracking operations like stabilization and corner positioning.

Position the Pattern Boxes

A pattern is the region of pixels that are targeted for tracking within an image using a pattern box. The pattern box is defined in the viewer by a rectangle when a tracker is active. A single Tracker node can have many pattern boxes, each targeting different patterns. Each tracked pattern will produce its own motion path.

Clicking any part of a tracker's onscreen controls will select it. Selected pattern boxes are red, while deselected pattern boxes are green.



A pattern box positioned over an eye you want to track


When you add a Tracker node to the Node Editor, you start with one pattern box displayed in the viewer as a small rectangle. When the cursor is placed over the pattern rectangle, the control expands and two rectangles appear. The outer rectangle has a dashed line, and the inner rectangle has a solid line. The outer rectangle is the search area, and the inner rectangle is the pattern.

If you need to select a new pattern, you can move the pattern box by dragging the small (and easily missed) handle at the top left of the inner pattern box.




To move a pattern in the viewer, drag it from the upper-left corner

While moving the pattern box, an overlay pop-up appears, showing a zoomed version of the pixels contained within the rectangle to help you precisely position the pattern via the crosshairs within.



A zoomed view appears while you drag a pattern box to help you position it

The pattern rectangle can also be resized by dragging on the edges of the rectangle. You want to size the pattern box so that it fits the detail you want to track, and excludes area that doesn't matter. Ideally, you want to make sure that every pixel of the pattern you're tracking is on the same plane, and that no part of the pattern is actually an occluding edge that's in front of what you're really tracking. When you resize the pattern box, it resizes from the center, so one drag lets you create any rectangle you need.




Resizing a pattern box to fit the tracking point on the ray gun

TIP: The magnified pattern box does not take viewer LUTs into account. When using Log content, it may make it easier to position the tracker if you temporarily insert a Brightness Contrast node between the source content and the yellow input of the tracker. You can use the Brightness Contrast node to temporarily increase the visibility of the region you are tracking.

Refine the Search Area

A second rectangle with a dotted border surrounds the pattern box. This is the search area. When progressing from one frame to another while tracking, the Tracker analyzes the region defined by the search area, which surrounds the last known tracker position in an attempt to relocate the pattern. The larger the search area, the better chance you have of successfully tracking fast moving objects, but the longer the track will take. However, there are some ways to optimize tracking for specific content.

For example, tracking a pattern that is moving quickly across the screen from left to right requires a wide search area but does not require a very tall one, since all movement is horizontal. If the search area is smaller than the movement of the pattern from one frame to the next, the Tracker will likely fail and start tracking the wrong pixels, so it's important to take the speed and direction of the motion into consideration when setting the search area.




You can resize the search area by dragging the edges of the dotted outline

Perform the Track Analysis

Before you begin analyzing, you'll need to make sure you've set a render range in the Time Ruler that corresponds to the range of frames during which the pattern is visible. This may be an entire clip or only a small portion of that clip. Depending on the type of motion you're tracking, you may want to use the Adaptive Mode option to aid the analysis (see below for more details).

Once your options are set, you can use any of the tracking transport buttons at the top of the Inspector to start tracking. Once tracking has started, you cannot work in the Node Editor until it has completed.



The tracking transport buttons and analysis parameters

To begin tracking, do one of the following:

- Click the Track Reverse button to track from the very end of the render range.
- Click the Track Backward from Current Frame button to track backward from the current playhead position.
- Click the Track Forward button to track from the very start of the render range.
- Click the Track Forward from Current Frame button to track forward from the current playhead position.

Pattern tracking will stop automatically when it reaches the end of the render range (or the start when tracking backward), but you can also interrupt it and stop tracking at any time.

To stop tracking, do one of the following:

- Click the Stop Tracking button in the tracker transports.
- Click Stop Render at the bottom of the Fusion window.
- Press the Escape key.

When tracking is complete, the path will be connected to the pattern. The path from that pattern can now be connected to another node or used for more advanced operations like stabilization and corner positioning.

Once the track is complete, assuming it's good, you can use the various techniques in this chapter to use the track in your composition.

Tips for Choosing a Good Pattern

The Tracker works by searching each frame for the pixels contained in the pattern. In order for a track to be successful, a fairly high contrast and unique region of the image must be located in the footage. This process is known as pattern selection.

The first step in pattern selection is to review the footage to be tracked several times. Watch for candidate patterns that are visible through the entire range of frames, where the contrast is high and the shape of the pattern does not change over time. The more unique the pattern, the more likely the track is to be successful.

In addition to locating high contrast, defined patterns, watch for the frames where the pattern moves the most. Identifying the maximum range of a pattern's motion will help to determine the correct size for the pattern search area.

It is not uncommon to have a scene that requires the use of several different patterns to generate a single path. This most often occurs because the pattern moves out of frame or is temporarily obscured by another scene element. Combining patterns into a single pattern is described later in the chapter.

Selecting the Pattern's Image Channels

When a pattern of pixels is selected, the Tracker automatically selects the color channel used for tracking the pattern based on an analysis of each channel for contrast, clarity, and reliability. The channels selected are highlighted in the bars to the right of the Pattern display window in the node controls.



Highlighted channel bars indicate which channel is selected for tracking

You can override the automatic channel selection by clicking the buttons beneath the bars for each channel to determine the channel used for tracking.

You can choose any one of the color channels, the luminance channels, or the alpha channel to track a pattern.

When choosing a channel, the goal is to choose the cleanest, highest contrast channel for use in the track. Channels that contain large amounts of grain or noise should be avoided. Bright objects against dark backgrounds often track best using the luminance channel.

Selecting Patterns for Stabilization

Selecting patterns for stabilization can be a tricky business. The location of the pattern, when it is selected, is used to determine precisely how the image will be stabilized. At least two patterns are required to correct for rotation; using three patterns will correct for scaling, and more will usually improve the quality of the solution.

Try not to select just any potentially valid pattern in the sequence, as some patterns will make the solution worse rather than better. To help with your selection, use the following guidelines when selecting patterns for stabilization.

- Locate patterns at the same relative depth in the image. Objects further in the background will move in greater amounts compared to objects in the foreground due to perspective distortion. This can confuse the stabilization calculations, which do not compensate for depth.
- Locate patterns that are fixed in position relative to each other. Patterns should not be capable of moving with reference to each other. The four corners of a sign would be excellent candidates, while the faces of two different people in the scene would be extremely poor choices for patterns.

Using the Pattern Flipbooks

Each pattern has a pair of thumbnail windows shown in the Inspector. The left window shows the selected pattern, while the right window is updated during the track to show the actual pattern that has been acquired for each frame.




The Tracker Pattern Selection and Flipbook thumbnails

Each pattern that's stored is added to a Flipbook. Once the render is complete, you can play this Pattern Flipbook to help you evaluate the accuracy of the tracked path. If you notice any jumps in the frames, then you know something probably went wrong.

Using Adaptive Pattern Tracking

Even the most ideal pattern will usually undergo shifts in profile, lighting conditions, and other variables. These can adversely affect pattern recognition to the point that a pattern becomes unusable. The Tracker offers three modes of pattern acquisition during tracking that can help to correct these conditions. The modes can be set using the Adaptive Mode options in the Inspector.



The Adaptive Mode options

None

When the Adaptive mode is set to None, the pattern within the rectangle is acquired when the pattern is selected, and that becomes the only pattern used during the track.

Every Frame

When Every Frame is chosen, the pattern within the rectangle is acquired when the pattern is selected, and then reacquired at each frame. The pattern found at frame 1 is used in the search on frame 2, the pattern found on frame 2 is used to search frame 3, and so on. This method helps the Tracker adapt to changing conditions in the pattern.

Every Frame tracking is slower and can be prone to drifting from sub-pixel shifts in the pattern from frame to frame. Its use is therefore not recommended unless other methods fail.

Best Match Tracking

Best Match tracking works in much the same way as Every Frame tracking; however, it will not reacquire the pattern if the difference between the original pattern and the new one is too great. This helps to prevent cases where transient changes in the image cause the Tracker to become confused.

As a comparison between the two Adaptive modes, if a shadow passes over the tracker point, the Every Frame tracking mode may start tracking the shadow instead of the desired pattern. The Best Match mode would detect that the change from the previous frame's pattern was too extreme and would not grab a new pattern from that frame.


The Adaptive mode is applied to all active patterns while tracking. If you only want some patterns to use the Adaptive mode, disable all other patterns in the list before tracking.

Dealing with Obscured Patterns

Often, an otherwise ideal pattern can be temporarily obscured (occluded) or blocked from tracking—for example, when tracking a car that passes behind a telephone pole.

In these situations, you divide the render range up into two ranges, the range before the pattern is obscured and the range after the pattern becomes visible again. After tracking the two ranges individually, the Tracker will automatically interpolate between the end of the first range and the start of the second.


If you need to edit the resulting motion path to account for any non-linear motion that takes place between the two tracked ranges, you can select the track path to expose a Node toolbar with controls for adjusting the control points on this path. For example, you can choose Insert and Modify mode to insert points in the non-tracked range to compensate for any nonlinear motion in the tracked pattern.



Tools for modifying tracker paths in the Node toolbar of the viewer

Dealing with Patterns That Leave the Frame

There are two options when a tracker leaves the frame. If the pattern re-enters the frame, you can treat it like an obscured pattern. If the pattern does not re-enter the frame, or it is undesirable to hand track portions of the movement, you can use the Track Center (Append) mode to select a new pattern.



The Track Center (Append) mode pop-up menu

The Track Center (Append) mode selects a new pattern that will continue to add keyframes to the existing path. The offset between the old pattern and the new pattern is automatically calculated to create one continuous path.


To use the Track Center (Append) mode, do the following:

- 1 When the pattern has become untrackable for some reason, stop analysis and move the playhead to the last frame that tracked successfully.
- 2 Choose Track Center (Append) from the Path Center pop-up menu in the Inspector.
- 3 Now, drag the Pattern selector to a new pattern that can be tracked from that point onward.
- 4 Restart tracking from the current frame.

When selecting a pattern to use in appending to an existing path, a pattern that is close to the old pattern and at the same apparent depth in the frame generates the best results. The further away the new pattern is, the more likely it is that the difference in perspective and axial rotation will reduce accuracy of the tracked result.


Setting Up Tracker Offsets

Often, it's impossible to track the thing you want to apply an effect to. For example, the only pattern available for an accurate track is a button on an actor's sleeve. However, the effect requires the person's hand to be glowing. To cause the glow's effect mask to be centered on the actor's hand, it's necessary to use the Tracker Offset control.



The Tracker Offset controls in the Inspector

The X and Y Offset controls allow for constant or animated positional offsets to be created relative to the actual Tracker's pattern center. The position of the offset in the viewer will be shown by a dashed line running from the pattern center to the offset position. You can also adjust the offset in the viewer using the Tracker Offset button. Clicking the button enables you to reposition the path while keeping the Tracker pattern in place.



The Tracker Offset tool in the Node toolbar of the viewer; a track of the orange dot is being offset to the center of the ray gun


Once an offset for a pattern is set, you can connect other positional controls to the Tracker's Offset menu using the Connect To > Tracker: Offset Position option in the control's contextual menu. The path created during the track remains fixed to the center of the pattern.

Stabilizing with the Tracker Node

When a Tracker node is set to Match Move in the Operations tab, it is capable of a variety of functions. Applying the motion from the background clip to the foreground clip is the obvious functionality. However, the Match Move operation is also used for stabilizing footage to either completely remove motion from the scene or smooth existing motion.

Here are some common scenarios for stabilization that are handled when the Tracker is set to Match Move.

- A sequence that should be steady has vibrations or undesirable movement.
- A sequence that requires a smooth camera move suffers from jarring.



The Tracker Operation tab Match Move set to BG Only for stabilization


Stabilization Using the Tracker Match Move Mode

Stabilizing motion completely removes the appearance of motion from the image. The motion from frame to frame is calculated, and the contents of the frame are transformed to return the image to a reference position. This position can be either the start or end of the sequence or a manually selected frame from the sequence.

Stabilization can correct for position with as little as one pattern. Two or more patterns are required to correct for rotation or scaling within the image.

When the Operation menu is set to Match Move, choosing BG only from the Merge operation menu stabilizes the background (yellow input) clip. Only the controls that are applicable for stabilization operations will appear in the Operation tab.

Several of the stabilization controls are always available, collected under the Match Move Settings disclosure button. These controls are available at all times because the Steady and Unsteady positions of a tracker are always published. This makes them available for connection by other controls, even when the Tracker's operation is not set to match moving.



The Match Move settings

Merge

The Merge menu determines to which input connection the Tracking data is applied. When stabilizing an image to remove all motion, or smooth the motion, the Merge button must be set to BG Only.

Edges

The Edges menu determines whether the edges of an image that leave the visible frame are cropped, duplicated, or wrapped when the stabilization is applied. Wrapping edges is often desirable for some methods of match moving, although rarely when stabilizing the image for any other purpose. For more information on the controls, see Chapter 117, “Tracker Nodes,” in the DaVinci Resolve Reference Manual or Chapter 57 in the Fusion Reference Manual.

Position/Rotation/Scaling

Use the Position, Rotation, and Scaling checkboxes to select what aspects of the motion are corrected.

Match Move Settings

Options for the Match Move settings include Pivot and Reference.

Pivot Type

The Pivot Type for the stabilization is used to calculate the axis of rotation and scaling calculations. This is usually the average of the combined pattern centers but may be changed to the position of a single tracker or a manually selected position.

Reference


The Reference controls establish whether the image is stabilized to the first frame in the sequence, the last frame, or to a manually selected frame. Any deviation from this reference by the tracked patterns is transformed back to this ideal frame.

As a general rule, when tracking to remove all motion from a clip, set the Merge mode to BG Only, the Pivot Type to Tracker Average or Selected Tracker, and the Reference control to Start, End, or Select Time.

Smoothing Motion

When confronted with an image sequence with erratic or jerky camera motion, instead of trying to remove all movement from the shot, you often need to preserve the original camera movement while losing the erratic motion.

The Start & End reference option is designed for this technique. Instead of stabilizing to a reference frame, the tracked path is simplified. The position of each pattern is evaluated from the start of the path and the end of the path along with intervening points. The result is smooth motion that replaces the existing unsteady move.



The Reference Intermediate Points slider is displayed when Start & End is selected to enable the smoothing of motion

To preserve some of the curvature of the original camera motion, you can increase the value of the Reference Intermediate Points slider that appears when the Start & End reference mode is selected.

When tracking to create smooth camera motion, ensure that the Start & End reference mode is enabled and set the Merge mode to BG Only. It is recommended to leave the Pivot Type control set to Tracker Average.

Using the Tracker Node for Match Moving


A simple match moving example is shown at the beginning of this chapter, but this section presents additional details that you may not have been aware of. Examples of match moving include:

- A static CG element must be believably added to a moving sequence.
- Two sequences with different motions must be composited together.

Some clips may need to be stabilized so that an element from another source can be added to the shot. After the element or effect has been composited, the stabilization should be removed to make the shot look natural again.

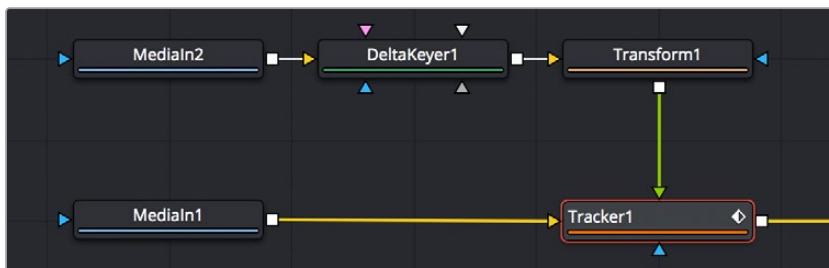
Simple Match Moving

Match moving essentially applies the movement from the tracked clip to another clip. There are two ways to perform match moving. One method involves connecting other nodes, such as Transform or Merge, to a Tracker's outputs. The other method is to stabilize an image by trying to remove all motion, but instead of setting the Merge menu to BG Only, set it to FG Over BG, FG Only, or in rare occasions, BG Over FG.



Set the Merge menu to BG Only, FG Over BG, or BG Over FG

When using this Merge menu, you connect a foreground image to the Tracker node's input connection in the Node Editor.



Connect a foreground image to the Tracker's foreground input

Enabling the FG Only mode will apply the motion from the background to the foreground, and the Tracker will only output the modified FG image. This result can later be merged over the original, allowing further modifications of the foreground to be applied using other nodes before merging the result over the background clip.

Corner Positioning Operations

The Corner Positioning operation maps the four corners of a foreground image to four patterns within the Tracker. This operation, or technique, is most commonly used for sign replacements.

The Corner Positioning operation of the Tracker requires the presence of a minimum of four patterns. If this operation mode is selected and there are not four patterns set up in the Tracker already, additional patterns will automatically be added to bring the total up to four.

When this mode is enabled, a set of drop-down boxes will appear to select which tracker relates to each corner of the rectangle. It has no effect when the Merge control option is set to BG Only.

Perspective Positioning Operations

The Perspective Positioning operation is used to remove perspective from a foreground image or apply the perspective from one sequence to another. This can be useful when you need to paint out an area that is distorted by perspective. Removing the perspective flattens the images for painting, and then another tracker adds the perspective back.

The Perspective Positioning operation of the Tracker requires the presence of a minimum of four patterns. If this operation mode is selected and there are not four patterns set up in the Tracker already, additional patterns will automatically be added to bring the total up to four.

When this mode is enabled, a set of drop-down boxes will appear to select which tracker relates to each corner of the rectangle. It has no effect when the Merge control option is set to BG Only.

Connecting to Trackers' Operations

One of the most common applications for a tracked pattern is using the tracked position or path to drive the position of another node's parameters. For example, tracking an eye in order to color correct the eye to blue using an effect mask. You start off by tracking the eye, and then create a color corrector with the desired settings. You create a mask in the shape of the eye and connect the Tracker's position to the Center of the mask.

In addition to the path (called Offset Position), each pattern in a tracker publishes four other values for use as connections that are available to other nodes in the Node Editor.

You connect a node's position parameters to a tracker by selecting the connection type from the controls contextual menu (for example, Transform 1: Center > Connect To > Tracker 1 > Offset Position).

There are five connection types automatically published by the tracker to connect to a position parameter in another node.

Steady Position

Steady Position can be used to stabilize footage in both X and/or Y to remove camera shake and other unwanted movement. The connection inverts the output of the tracked pattern's motion. When you connect a Center parameter to the Steady Position of the Tracker, it will be placed at 0.5/0.5 (the center of the screen) by default at frame 1. You can change this using the Reference mode in the Tracker's Operation tab.

Steady Angle

The Steady Angle mode can be used to stabilize footage in both X and/or Y to remove camera shake and other unwanted movement. When you connect a control, for example the Angle of a Transform, to the Steady Angle of the Tracker, it will be placed at 0 degrees by default at frame 1. This can be changed by means of the Reference mode in the Tracker's Operation tab. From there on, the resulting motion of the Steady Angle mode will rotate into the opposite direction of the original motion.

So if the angle at frame 10 is 15 degrees, the result of the Steady Angle will be -15 degrees.

To use Steady Angle, you need at least two tracked patterns in your tracker. With just one point, you can only apply (Un)Steady Position.

Offset Position

An Offset Position is essentially the path generated by the tracker. It is the one you select when you want an object to follow the path. It is available for each single tracker in the Tracker node and refers to that single tracker only. When you connect the Center X and Y parameters to the offset position of the Tracker, the node's center will follow exactly the path of that tracker. Connecting to single trackers is always useful when you want to match elements with object motion in your footage. For example, you could track a hand of your actor and attach a ball to the Tracker's offset position, so that the ball follows the exact motion of the hand. Or you could track an element that needs rotoscoping and attach the mask's center to the Tracker's offset position.

Unsteady Position

After using the Steady Position, the Unsteady Position is used to reintroduce the original movement on an image after an effect or new layer has been added. The resulting motion from Unsteady Position is basically an offset in the same direction as the original motion.


Steady Size

The Steady Size connection outputs the inverse of the tracked pattern's scale. When you connect a parameter, for example the Size of a Transform, to the Steady Size of the Tracker, it will be placed with a Size of 1 (i.e., the original size) by default at frame 1. This can be changed by means of the Reference mode in the Tracker's Operation tab. The resulting size of the Steady Size mode will then counteract the size changes of the original motion. So if the actual size at frame 10 is 1.15, the result of the Steady Size will be $1 - (1.15 - 1) = 0.85$.

To use Steady Size, you need at least two tracked patterns in your tracker. With just one point, you can only apply (Un)Steady Position.

Using the Outputs of a Tracker

The tracker outputs described above are published by each tracker pattern created in the Tracker node. Each Tracker node itself also publishes a Steady Position, Angle, Size and an Unsteady Position. The values of these Tracker node outputs are calculated using all the patterns in that tracker, as configured by the Match Move Settings controls in the Tracker's Operation tab.



The Match Move settings determine the reference frame for stabilization and the point around which stabilization pivots

Rather than using the Tracker node to perform the Merge operation, an alternative and common way to use these published outputs is to create a match move by connecting the outputs to multiple nodes. A tracker is used to track a pattern, and then that data can be connected to multiple other nodes using the Connect To submenu.


As an example, to use the Connect To menu to perform a match move, do the following:

- 1 Track the background clip using at least two tracking patterns in the tracker.
- 2 In a different branch, add a Transform node to the background clip.




The background branched to go to a Tracker and a Transform

- 3 Right-click over the Transform's Center and choose Connect to > Tracker1 > Steady Position.



The tracker publishes its output for other nodes to connect to, as done here to stabilize the clip

- 4 Connect the foreground to a corner-positioned node, so you can position the corners of the foreground appropriately over the background.
- 5 Add another Transform node to the Node Editor after the Merge.



A second Transform after the Merge is used to add back in the original motion with Unsteady Position


- 6 Connect the new Transform's Center to the Tracker's Unsteady Position. The image will be restored to its original state with the additional effect included.

To better understand how this works, imagine a pattern that is selected at frame 1, at position 0.5, 0.5. The pattern does not move on frame 2, so its position is still 0.5, 0.5. On the third frame, it moves 10 percent of the image's width to the right. Now its position is 0.6, 0.5.

If a transform center is connected to the Steady Position output of the Tracker, the Transform node's center is 0.5, 0.5 on the first and second frames because there has been no change. On frame 3, the center moves to 0.4, 0.5. This is the inverse of the horizontal motion that was tracked in the pattern, moving the image slightly to the right by 10 percent of the image width to counteract the movement and return the pattern of pixels back to where they were found.

Using the Tracker as a Modifier

Another technique for adding a tracker directly to a control is to add it as a modifier. Choosing the Tracker from the Modify With contextual menu does not use a Tracker node; it adds a modifier in the Inspector with a set of parameters almost identical to those found in the Tracker node itself. The benefit here is that the object that you want to follow the tracked path is automatically connected to the tracker modifier when you apply it.



Applying the tracker as a modifier

The differences between a Tracker modifier and a Tracker node are as follows:

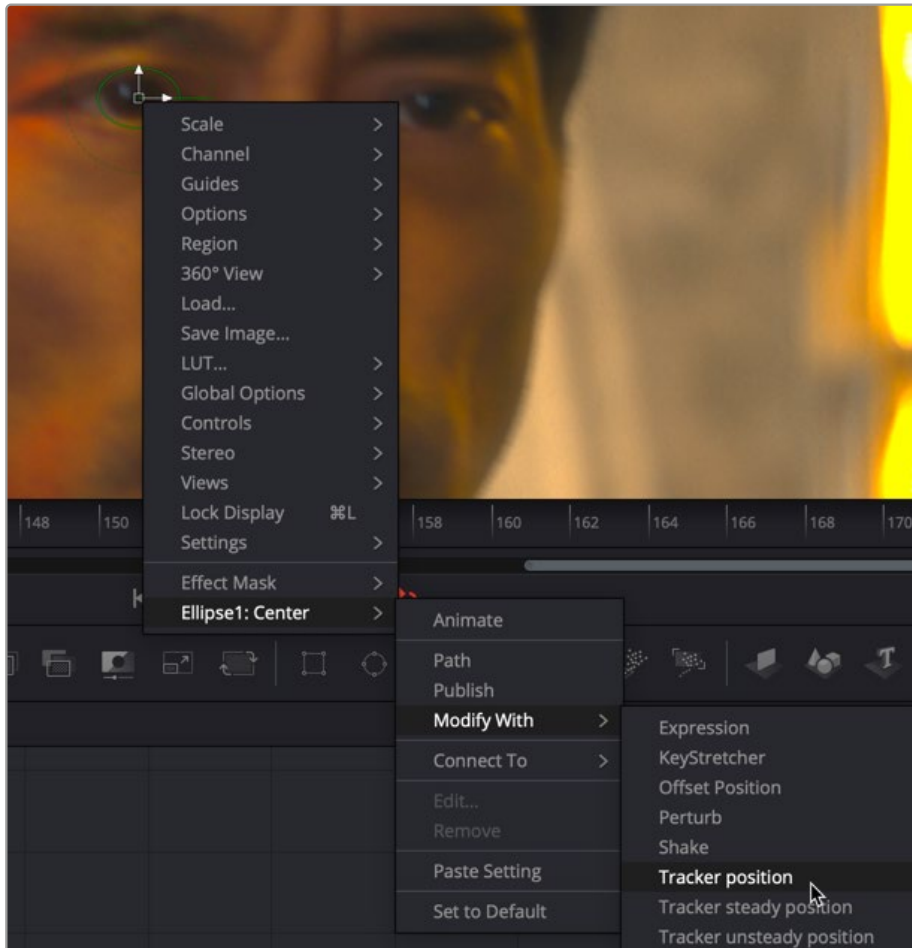
- The Tracker modifier can only track a single pattern.
- A source image must be set for the Tracker modifier.

The Tracker modifier can only output a single value and cannot be used for complex stabilization procedures, but it is a nice quick way to apply a tracker to a point that you need to follow.

As an example, to apply the Tracker as a modifier, do the following:

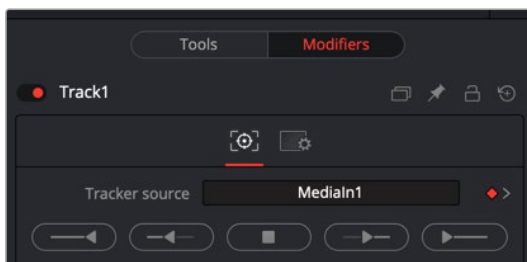
Imagine that you needed to track an actor's eyes so that an unearthly, alien glow could be applied to the eyes.

- 1 Add an ellipse mask node to cover an actor's eye.
- 2 In the Inspector, right-click on the mask's Center parameter and from the contextual menu choose Ellipse1 Center > Modify With > Tracker Position.



Right-click over any Center Coordinate Control and choose Modify With > Tracker Position to add a Tracker modifier

- 3 Click the Modifiers tab in the Inspector and drag the MediaIn1 node that you want to track into the Tracker Source field.



The Modifiers tab includes the tracking controls and a Tracker Source field to indicate which node to use for tracking

- 4 Click the Track Forward button to begin tracking the person's eye.
- 5 Insert a Soft Glow node directly after the MediaIn and connect the Ellipse Mask to the white Glow Mask input.



A Tracker modifier applied to the Ellipse to create a green glow on an actor's pupil

You can set a different source image for the Tracker modifier by typing in the name of the node or dragging and dropping the node from the Node Editor into the Tracker Source field control. If you have a node (let's call it node#1) connected to the node that contains the modifier (let's call it node#2), the source image for the Tracker modifier will automatically be the node #1

For more information on the Tracking parameters, see Chapter 117, "Tracker Nodes," in the DaVinci Resolve Reference Manual or Chapter 57 in the Fusion Reference Manual.

Match Moving Text Example

This example takes you through a complete motion tracking task, and shows how you can create a very simple match-moving effect using the Tracker node, which is the Swiss army knife of trackers in Fusion.

Adding a Layer to Match Move

In the example composition, we have a Text1 node that's creating a "Switzerland" title that's composited over a drone shot flying over and around a mountain bridge. With the Text1 node selected, the onscreen controls that let you position the text it's generating are visible in the viewer, and the text is positioned where we'd like it to start. Note that, with the Text node selected, even the part of the text that's offscreen can still be seen as an outline showing us where it is.

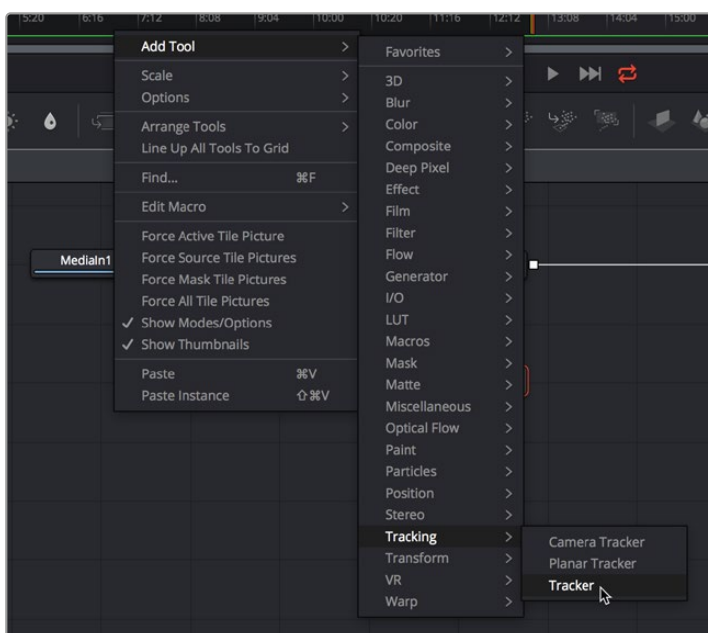


Some text superimposed against a background, ready to track

Our goal for this composition is to motion track the background image so that the text moves along with the scene as the camera flies along.


Setting Up Motion Tracking

To set up for the motion track, we'll begin by creating a disconnected Tracker node, using another method other than those seen previously. Right-click anywhere in the background of the Node Editor (preferably where you want the new node to appear), and choose Add Tool > Tracking > Tracker from the contextual menu to create a new Tracker1 node underneath the MediaIn node (or Loader node if you are using Fusion Studio).



Creating a new node using the Node Editor contextual menu

Next, we'll drag a connection from the Medialn1 node to the Tracker1 node to automatically connect the source clip to the Tracker1 background input. This branches the output from the Medialn1 node to the Tracker node so that the Tracker1 node processes the image separately from the rest of the node tree. This is not required, but it's a nice organizational way to see that the Tracker node is doing an analysis that must be referred to in a way other than a "physical" connection.



Branching a Tracker node to use to analyze an image

A Simple Tracking Workflow


The Tracker node is the simplest tracking operation the Fusion page has, and while there are several ways of using it, an extremely common workflow is to use the Tracker node controls to analyze the motion of a subject in the frame with motion you want to follow, and then use the resulting motion path data by "connecting" it to the Center parameter of another node that's capable of transforming the image you want to match move.

Positioning the Tracker Onscreen Control

When the Tracker node is selected, a single green box appears in the viewer, which is the default onscreen control for the first default tracker that node contains (seen in the Tracker List of the Inspector controls). Keep in mind that you only see onscreen controls for nodes that are selected, so if you don't see the onscreen tracker controls, you know you need to select the tracker you want to work with. Loading the tracker you want to work with into the viewer is also the safest way to make sure you're positioning the controls correctly relative to the actual image that you're tracking.

If you position your pointer over this box, the entire onscreen control for that tracker appears, and if you click the onscreen control to select that tracker, it turns red. As with so many other tracker interfaces you've likely used, this consists of two boxes with various handles for moving and resizing them:

- The inner box is the "pattern box," which identifies the "pattern" in the image you're tracking and want to follow the motion of. The pattern box has a tiny handle at its upper-left corner that you use to drag the box to overlap whatever you want to track. You can also resize this box by dragging any corner, or you can squish or stretch the box by dragging any edge to make the box better fit the size of the pattern you're trying to track. The center position of the tracker is indicated via X and Y coordinates.
- The outer box is the "search box," which identifies how much of the image the tracker needs to analyze to follow the motion of the pattern. If you have a slow-moving image, then the default search box size is probably fine. However, if you have a fast-moving image, you may need to resize the search box (using the same kind of corner and side handles) to search a larger area, at the expense of a longer analysis. The name of that tracker is shown at the bottom right of the search box.



The onscreen controls of a selected tracker seen in isolation

It's worth saying a second time that the handle for moving a tracker's onscreen control is a tiny dot at the upper-left corner of the inner pattern box. It's really easy to miss if you're new to Fusion. You must click on this dot to drag the tracker around.



The handle for dragging the tracker boxes to move them around


In this example, we'll drag the onscreen control so the pattern box overlaps a section of the bridge right over the leftmost support. As we drag the onscreen control, we see a zoomed-in representation of the part of the image we're dragging over to help us position the tracker with greater precision. For this example, the default sizes of the pattern and search box are fine as is.



The zoomed-in preview that helps you position the pattern box as you drag it

Using the Tracker's Inspector Controls to Perform the Analysis


At this point, let's look at the Tracker node's controls in the Inspector. There are a lot of controls, but for this simple example we only care about the main Tracker panel, with the tracking analysis buttons at the top, the tracking options below those, and the Tracker List underneath those. The Tracker List also has buttons for adding and deleting trackers; you have the option of adding multiple trackers that can be analyzed all at once for different workflows, but we don't need that for now.



Tracker Inspector controls, with the tracking analysis buttons at top, the tracker options in the middle, and the Tracker List below

Additional controls over each tracker and the image channels being analyzed appear at the bottom, along with offset controls for each tracker, but we don't need those now (at least not yet).

Again, this track is so simple that we don't need to change the default behaviors that much, but because the drone is flying in a circular pattern, the shape of the pattern area is changing as the clip plays. Fortunately, we can choose Every Frame from the Adaptive Mode menu to instruct the tracker to update the pattern being matched at every frame of the analysis, to account for this.




Changing the Adaptive Mode of the Tracker node to Every Frame to account for the camera's shift of perspective

Now, we just need to use the tracker analysis buttons at the top to begin the analysis. These buttons work like transport controls, letting you start and stop analysis as necessary to deal with problem tracks in various ways. Keep in mind that the first and last buttons, Track from Last Frame and Track from First Frame, always begin a track at the last or first frame of the composition, regardless of the playhead's current position, so make sure you've placed your tracker onscreen controls appropriately at the last or first frame.



The analysis buttons, left to right: Track from Last Frame, Track Backward, Stop Tracking, Track Forward, Track from First Frame

For now, clicking the Track from Beginning button will analyze the entire range of this clip, from the first frame to the last. A dialog lets you know when the analysis is completed, and clicking the OK button dismisses it so you can see the nice clean motion path that results.



The analyzed motion path resulting from tracking a section of the bridge as the camera flies past

Viewing Motion Track Data in the Spline Editor


This is not a necessary part of the tracking workflow, but if you have an otherwise nice track with a few bumps in it, you can view the motion tracking data in the Spline Editor by viewing that tracker's Displacement parameter curve. This curve is editable, so you can massage your tracking data in a variety of ways, if necessary.



Viewing motion tracking analysis data in the Spline Editor

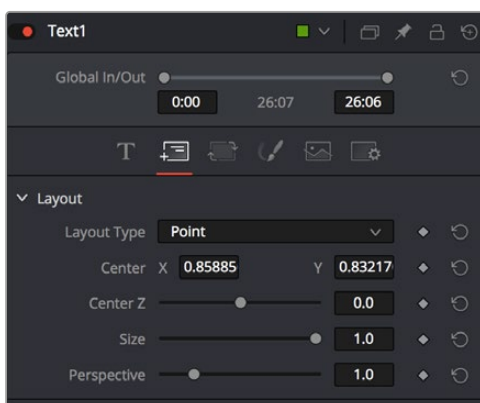
Connecting Motion Track Data to Match Move

Now that we have a successful analysis, it's time to use it to create the Match Move effect. To make this process easier, we'll double-click the tracker's name in the Tracker List of the Inspector, and enter a new name that's easier to keep track of (heh). Adding your own names make that tracker easier to find in subsequent contextual menus and lets you keep track of which trackers are following which subjects as you work on increasingly complex compositions.



Renaming a tracker to make it easier to find


Now it's time to connect the track we've just made to the text in order to start it in motion. After loading the Merge1 node into the viewer to see the text in context with the overall composite we're creating, we'll select the Text1 node to open its parameters in the Inspector, and click the Layout panel icon (second button from the left) to expose the Layout controls, which are the text-specific transform controls used to position the text object in the frame. These are the controls that are manipulated when you use the Text node onscreen controls for repositioning or rotating text.



The Layout controls for a Text node, in the Layout panel

The Center X and Y parameters, while individually adjustable, also function as a single target for purposes of connecting to tracking to quickly set up match moving animation. You set this up via the contextual menu that appears when you right-click any parameter in the Inspector, which contains a variety of commands for adding keyframing, modifiers, expressions, and other automated methods of animation including connecting to motion tracking.

If we right-click anywhere on the line of controls for Center X and Y, we can choose Connect To > Tracker1 > Bridge Track: Offset position from the contextual menu, which connects this parameter to the tracking data we analyzed earlier.



Connecting the Center X and Y parameter to the Bridge Track: Offset position motion path we analyzed

Immediately, the text moves so that the center position coincides with the center of the tracked motion path at that frame. This lets us know the center of the text is being match moved to the motion track path.



The text now aligns with the motion track coordinate

Offsetting the Position of a Match Moved Image

In fact, we want to offset the match-moved text, so it's higher up in the frame. To do this, we select the Tracker1 node again and use the Y Offset 1 dial control to move the text up, since now any changes we make to the Bridge Track dataset now apply to the center of the text that's connected to it.



Using the X and Y Offset controls in the Tracker1 node to offset the text layer's position from the tracked motion path

The offset we create is shown as a dotted red line that lets us see the actual offset being created by the X and Y Offset controls. In fact, this is why we connected to the Bridge Track: Offset position option earlier.



The text offset from the tracked motion path; the offset can be seen as a dotted red line in the viewer

Now, if we play through this clip, we can see the text moving along with the bridge.



Two frames of the text being match moved to follow the bridge in the shot

Chapter 83

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 117, “Tracker Nodes,” in the DaVinci Resolve Reference Manual or Chapter 57 in the Fusion Reference Manual.

Contents

Introduction to Tracking	1506
Using the Planar Tracker	1506
Different Ways of Using the Planar Tracker Node	1506
Setting Up to Use the Planar Tracker	1507
Check for Lens Distortion	1507
A Basic Planar Tracker Match Move Workflow	1507
Tips for Choosing Good Planes to Track	1510

Introduction to Tracking

Fusion includes three different Tracking nodes that let you analyze different kinds of motion. Once you have tracked motion on a clip, you can then use the resulting data for stabilization, motion smoothing, matching the motion of one object to that of another, and a host of other essential tasks. Each tracker type has its own chapter in this manual. This chapter covers the tracking techniques with the Planar Tracker node.

Using the Planar Tracker

The Planar Tracker node is designed to deal with match moving issues on flat, unvarying surfaces that commonly arise during post-production. Examples of flat, unvarying surfaces include clips containing a license plate, a road sign, or a brick wall that often need images merged on top of them, such as replacing the numbers in the license plate, changing the city's name in the road sign, or placing a billboard poster on the empty brick wall.

The Planar Tracker automates this process by analyzing the perspective distortions of a planar surface on a background plate over time, and then re-applying those same perspective distortions to a different foreground.

TIP: Part of using the Planar Tracker is also knowing when to give up and fall back to using Fusion's Tracker node or to manual keyframing. Some shots are simply not trackable, or the resulting track suffers from too much jitter or drift. The Planar Tracker is a time-saving node in the artist's toolbox and, while it can track most shots, no tracker is a 100% solution.

Different Ways of Using the Planar Tracker Node


Like the other tracking nodes found in Fusion, the Planar Tracker can both analyze and contain the resulting image tracking data interior to the node, and it can also use that tracking data to transform either another image, paint strokes, or a polygon mask shape.

The Planar Tracker provides four modes of operation:

- **Track:** Used to isolate a planar surface and track its movement over time. Then, you can create a Planar Transform node that uses this data to match move another clip in various ways.
- **Steady:** After analyzing a planar surface, this mode removes all motion and distortions from the planar surface, usually in preparation for some kind of paint or roto task, prior to “unsteadying” the clip to add the motion back.
- **Corner Pin:** After analyzing a planar surface, this mode computes and applies a matching perspective distortion to a foreground image you connect to the foreground input of the Planar Tracker node, and merges it on top of the tracked footage.
- **Stabilize:** After analyzing a planar surface, allows smoothing of a clip's translation, rotation, and scale over time. Good for getting unwanted vibrations out of a clip while retaining the overall camera motion that was intended.

Setting Up to Use the Planar Tracker

Similar to the Tracker node, to do a planar track, you need to connect the output of the image you want to track to the background input of a Planar Tracker node.




Connecting an image to the background input of a PlanarTracker node

Check for Lens Distortion

If the image has barrel distortion, or any other kinds of lens distortion, it can adversely affect your track. The more lens distortion in the footage, the more the resulting track will slide and wobble. If you can see distortion in the image or you're having problems with the track, you'll want to try inserting the Lens Distort node between the image and the Planar Tracker to eliminate this problem.

Fusion's Lens Distort node can be used to remove or add lens distortion in an image. Connecting the Medialn or Loader node to the Lens Distort node displays controls for manually correcting lens distortion. If you use Synth Eyes, PFTTrack or 3D Equalizer software, you can also import lens data from those applications to make the adjustments more automatic.



A Lens Distort node inserted between a Medialn1 and Planar Tracker to remove lens distortion

For more information about using the Lens Distort node, see Chapter 120, "Warp Nodes," in the DaVinci Resolve Reference Manual or Chapter 60 in the Fusion Reference Manual.


If you are using DaVinci Resolve, you can use the Lens Corrections control in the Cut page or Edit page. This adjustment carries over into the Fusion page. Lens correction in DaVinci Resolve automatically analyzes the frame in the Timeline viewer for edges that are being distorted by a wide angle lens. Clicking the Analyze button moves the Distortion slider to provide an automatic correction. From there, the Medialn node in the Fusion page will have the correction applied, and you can begin planar tracking.

A Basic Planar Tracker Match Move Workflow

Using the Planar Tracker is a process, but it's straightforward once you've learned how to use it. The following procedure tries to make this process as clear as possible.

To track a surface using the Planar Tracker:

- 1 Make sure the Operation Mode is set to Track, as you need to analyze an image to track a surface before you do anything else.
- 2 With the background input of the Planar Tracker connected to an image, and the Planar Tracker open in a viewer, move the playhead to a frame of video where the planar surface you want to track is at its largest, is unoccluded, and is clearly a plane, and then click the Set button in the Track panel of the Inspector to make this the reference frame that will be used to guide the track.



Clicking the Set button to set the reference frame to use for analysis

- 3 Next, you'll need to identify the specific pattern within the image that you want to track. In most cases, this will probably be a rectangle, but any arbitrary closed polygon can be used. The pixels enclosed by this region will serve as the pattern that will be searched for on other frames. Please note that it is important that the pattern is drawn on the reference frame. In this example, we want to track the wall behind the man, so we draw a polygon around part of the wall that the man won't pass over as he moves during the shot.



Drawing a polygon to identify the part of the image you want to track, which should be a flat, trackable plane

TIP: Do not confuse the pattern you're identifying with the region you're planning to corner pin (which always has four corners and is separately specified in Corner Pin mode).

- 4 (Optional) If moving objects partially cover up or occlude the planar surface, you may wish to connect a mask that surrounds and identifies these occlusions to the white "occlusion mask" input of the Planar Tracker. This lets the Planar Tracker ignore details that will cause problems.

When using the Hybrid Tracker, providing a mask to deal with occluding objects is nearly mandatory, while with the Point Tracker it is recommended to try tracking without a mask.

- 5 If necessary, move the playhead back to the reference frame, which in this case was the first frame. Then, click the Track To End button and wait for the track to complete.




The Analyze buttons of the Planar Tracker

As the clip tracks, you can see track markers and trails (if they're enabled in the Options tab of the Inspector) that let you see how much detail is contributing to the track, and the direction of motion that's being analyzed.




During tracking, you can see track markers and trails to follow how well the track is going

- 6 Once the track is complete, play through the clip to visually inspect the track so you can evaluate how accurate it is. Does it stick to the surface? Switching to Steady mode can help here, as scrubbing through the clip in Steady mode will help you immediately see unwanted motion in the track.
- 7 Since we're doing a match move, click the Create Planar Transform button to export a Planar Transform node that will automatically transform either images or masks to follow the analyzed motion of the plane you tracked.



Clicking Create Planar Transform to create a node to use to transform other images or masks

In this case, the Planar Transform node will be inserted after a pair of Background and Paint nodes that are being used to put some irritatingly trendy tech jargon graffiti on the wall. The Planar Transform will automatically transform the Paint node's output connected to its background input to match the movement of the wall.



Adding the PlanarTransform node after a Paint node to match move it to the background image, combining it via a Merge node

The result is a seamless match move of the fake graffiti married to the wall in the original clip.



The final result; the paint layer is match moved to the background successfully

TIP: If you want to composite semi-transparent paint strokes on the wall, or use Apply modes with paint stroke, you can attach a Paint node to a Background node set to 100 transparency. The resulting image will be whatever paint strokes you make against transparency and is easy to composite.

Tips for Choosing Good Planes to Track

The region to track is specified by drawing a polygon on the reference frame. Make sure the region selected belongs to a physically planar surface in the shot. Sometimes a region that is only approximately planar can be used. In general, the less planar the surface, the poorer the quality of the resulting track.

As a rule of thumb, the more pixels in the pattern, the better the quality of the track. In particular, this means on the reference frame, the pattern to be tracked should:

- Be as large as possible.
- Be as much in frame as possible.
- Be as unoccluded as possible by any moving foreground objects.
- Be at its maximum size (e.g., when tracking an approaching road sign, it is good to pick a later frame where it is 400 x 200 pixels rather than 80 x 40 pixels).
- Be relatively undistorted (e.g., when the camera orbits around a flat stop sign, it is better to pick a frame where the sign is face on parallel to the camera rather than a frame where it is at a highly oblique angle).

If the pattern contains too few pixels or not enough trackable features, this can cause problems with the resulting track, such as jitter, wobble, and slippage. Sometimes dropping down to a simpler motion type can help in this situation.

Chapter 84

Using Open FX, Resolve FX, and Fuse Plug-Ins

Fusion's capabilities can be extended using different kinds of plug-ins. All compositions in Fusion Studio and in the Fusion page of DaVinci Resolve support third-party Open FX plug-ins. Additionally, the Fusion page of DaVinci Resolve provides access to all of the Resolve FX that come with DaVinci Resolve.

Lastly, you can develop your own plug-ins without using a computer development environment by scripting Fusion's native Fuse plug-ins.

Contents

What Are Open FX?	1512
What Are Resolve FX?	1512
Applying Open FX and Resolve FX Plug-Ins	1512
Introduction to Fuse Plug-Ins	1513

What Are Open FX?

Fusion is able to use compatible Open FX (OFX) plug-ins that are installed on your computer. Open FX is an open standard for visual effects plug-ins. It allows plug-ins written to the standard to work on both DaVinci Resolve and Fusion Studio as well as other applications that support the standard.

OFX plug-ins can be purchased and downloaded from third-party suppliers such as BorisFX, Red Giant, and RE:Vision Effects. All OFX appear in the Open FX category of the Effects Library, alongside all other effects that are available in Fusion.

What Are Resolve FX?

The Fusion page in DaVinci Resolve is also able to access Resolve FX. Resolve FX are filter effects within DaVinci Resolve. Although most Resolve FX have the same capabilities on all DaVinci Resolve pages, some are specific to the Color page and require the use of the Color page tracker for full functionality. All Resolve FX appear in the Fusion page Effects Library in the Open FX category alongside third-party OFX plug-ins. Resolve FX are not available in Fusion Studio. For more information about Resolve FX, see Chapter 152, “Resolve FX,” in the DaVinci Resolve Reference Manual.

Applying Open FX and Resolve FX Plug-Ins

Resolve FX and OFX plug-ins are applied to the Node Editor just as you would apply native Fusion nodes. They can help you create fast, interesting looks and effects using the Color page, or imaginative transitions and effects on your clips on the Edit page. Resolve FX are installed with DaVinci Resolve,

After installing a set of OFX plug-ins, you can access them or Resolve FX plug-ins in Fusion by opening the Open FX category in the Effects Library.

To add a plug-in to the Node Editor, either click the Open FX or Resolve FX plug-in name in the Effects Library or drag and drop the plug-in onto a connection line to insert it into the node tree. If the plug-in has editable settings, you can adjust these in the Inspector.

Introduction to Fuse Plug-Ins

Fuses are plug-ins developed for Fusion using the Lua built-in scripting language. Being script-based, Fuses are compiled on-the-fly in Fusion without the need of a computer programming environment. While a Fuse may be slower than an identical Open FX plug-in created using Fusion's C++ SDK, a Fuse will still take advantage of Fusion's existing nodes and GPU acceleration.

To install a Fuse:

- 1 Use the .fuse extension at the end of the document name.
- 2 For DaVinci Resolve, save it in one of the following locations:
 - **On macOS:** Macintosh HD/Users/username/Library/Application Support/Blackmagic Design/DaVinci Resolve/Fusion/Fuses
 - **On Windows:** C:\Users\username\AppData\Roaming\Blackmagic Design\DaVinci Resolve\Support\Fusion\Fuses
 - **On Linux:** home/username/.local/share/DaVinciResolve/Fusion/Fuses

For Fusion Studio, save it in one of the following locations:

- **On macOS:** Macintosh HD/Users/username/Library/Application Support/Blackmagic Design/Fusion/Fuses/
- **On Windows:** C:\Users\username\AppData\Roaming\Blackmagic Design\Fusion\Fuses
- **On Linux:** home/username/.fusion/BlackmagicDesign/Fusion/Fuses

You can open and edit Fuses by selecting the Fuse node in the Node Editor and clicking the Edit button at the top of the Inspector. The Fuse opens in the text editor specified in the Global Preferences/Scripting panel.

TIP: Changes made to a Fuse in a text editor do not immediately propagate to other instances of that Fuse in the composition. Reopening a composition updates all Fuses in the composition based on the current saved version. Alternatively, you can click the Reload button in the Inspector to update the selected node without closing and reopening the composition.

Chapter 85

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	1515
3D Compositing Fundamentals	1516
Creating a Minimal 3D Scene	1517
The Elements of a 3D Scene	1518
Geometry Nodes	1518
The Merge3D Node	1520
The Renderer3D Node	1523
Software vs. GPU Rendering	1523
Software Renderer	1523
OpenGL Renderer	1524
OpenGL UV Renderer	1524
Loading 3D Nodes into the Viewer	1525
Navigating the 3D View	1527
Transforming Cameras and Lights Using the Viewers	1527
Transparency Sorting	1528
Material Viewer	1529
Transformations	1529
Onscreen Transform Controls	1530
Pivot	1531
Target	1531
Parenting	1533

Cameras	1533
Quickly Viewing a Scene Through a Camera	1534
Plane of Focus and Depth of Field	1534
Importing Cameras	1535
Lighting and Shadows	1536
Enabling Lighting in the Viewer	1536
Enabling Lighting to Be Rendered	1536
Controlling Lighting within Each 3D Object	1537
Lighting Types Explained	1537
Materials and Textures	1541
Material Components	1542
Alpha Detail	1544
Illumination Models	1545
Textures	1546
Reflections and Refractions	1547
Bump Maps	1549
Projection Mapping	1550
Project a Texture onto a Catcher Material	1551
Project Using the UVMap Node	1552
Geometry	1553
Common Visibility Parameters	1553
Adding FBX Models	1554
Using Text3D	1555
Fog 3D and Soft Clipping	1558
Material and Object IDs	1560
World Position Pass	1561
Point Clouds	1562


An Overview of 3D Compositing

Traditional image-based compositing is a two-dimensional process. Image layers have only the amount of depth needed to define one as foreground and another as background. This is at odds with the realities of production, since all images are either captured using a live-action camera with freedom in all three dimensions, in a shot that has real depth, or have been created in a true 3D modeling and rendering application.

Within the Fusion Node Editor, you have a GPU-accelerated 3D compositing environment that includes support for imported geometry, point clouds, and particle systems for taking care of such things as:

- Converting 2D images into image planes in 3D space
- Creating rough primitive geometry
- Importing mesh geometry from FBX or Alembic scenes
- Creating realistic surfaces using illumination models and shader compositing


- Rendering with realistic depth of field, motion blur, and supersampling
- Creating and using 3D particle systems
- Creating, extruding, and beveling 3D text
- Lighting and casting shadows across geometry
- 3D camera tracking
- Importing cameras, lights, and materials from 3D applications such as Maya, 3ds Max, or LightWave
- Importing matched cameras and point clouds from applications such as SynthEyes or PF Track



An example 3D scene in Fusion

3D Compositing Fundamentals

The 3D category of nodes (which includes the Light, Material, and Texture subcategories) work together to create 3D scenes. Examples are nodes that generate geometry, import geometry, modify geometry, create lights and cameras, and combine all these elements into a scene. Nearly all these nodes are collected within the 3D category of nodes found in the Effects Library.




The 3D category of nodes in the Effects Library

Conveniently, at no point are you required to specify whether your overall composition is 2D or 3D, because you can seamlessly combine any number of 2D and 3D “scenes” together to create a single output. However, the nodes that create these scenes must be combined in specific ways for this to work properly.

Creating a Minimal 3D Scene

Creating a 3D scene couldn't be easier, but you need to connect the required nodes in the right way. At minimum, you need only connect a geometry node (such as a Text3D node) to a Renderer3D node to output a 2D image that can be combined with other 2D images in your composition, as seen below. However, you'll only get a simply shaded piece of geometry for your trouble, although you can color and transform it in the Inspector using controls internal to whichever geometry node you're using.




A simple 3D scene with a Text3D node connected directly to a Renderer3D node

More realistically, each 3D scene that you want to create will probably have three to five nodes to give you a better lit and framed result. These include:

- One of the available geometry nodes (such as Text3D or Image Plane 3D)
- A light node (such as DirectionalLight or SpotLight)
- A camera node
- A Merge3D node
- A Renderer3D node

All these should be connected together as seen below, with the resultantly more complex 3D scene shown below.



The same text, this time lit and framed using Text3D, Camera, and SpotLight nodes to a Merge3D node

To briefly explain how this node tree works, the geometry node (in this case Text3D) creates an object for the scene, and then the Merge3D node provides a virtual stage that combines the attached geometry with the light and camera nodes to produce a lit and framed result with highlights and shadows, while the aptly named Renderer3D node renders the resulting 3D scene to produce 2D image output that can then be merged with other 2D images in your composition.

In fact, these nodes are so important that they appear at the right of the toolbar, enabling you to quickly produce 3D scenes whenever you require. You might notice that the order of the 3D buttons on the toolbar, from left to right, corresponds to the order in which these nodes are ordinarily used. So, if you simply click on each one of these buttons from left to right, you cannot fail to create a properly assembled 3D scene, ready to work on, as seen in the previous screenshot.



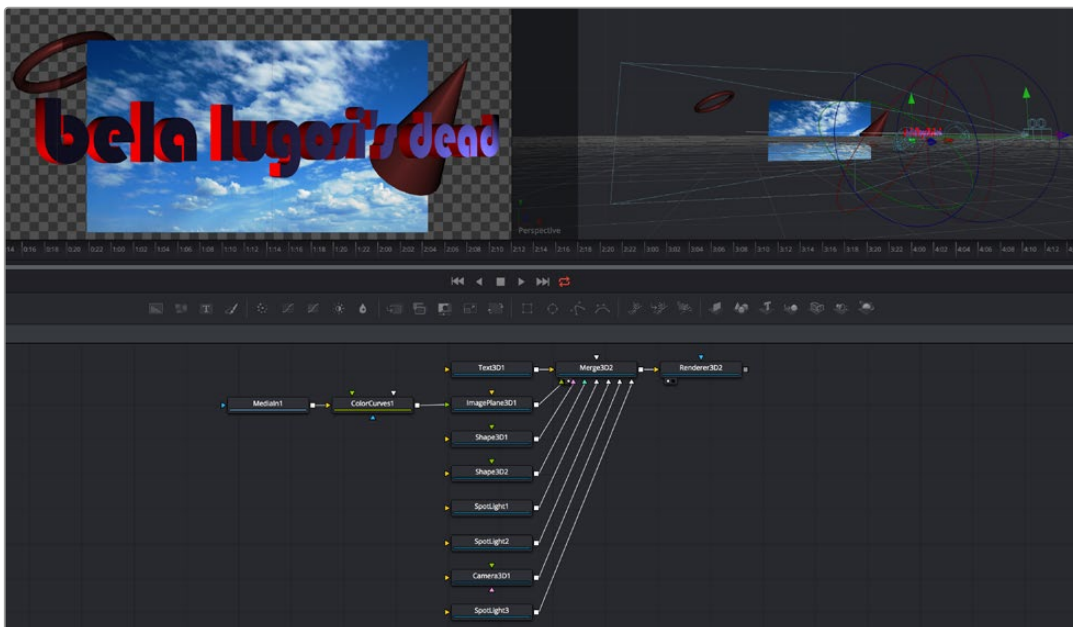
The 3D nodes available from the toolbar include the ImagePlane3D, Shape3D, Text3D, Merge3D, Camera3D, SpotLight3D, and Renderer3D nodes

The Elements of a 3D Scene

All 3D nodes can be divided into a number of categories.

Geometry Nodes

You can add 3D geometry to a composition using the ImagePlane3D node, the Shape3D node, the Cube3D node, the Text3D node, or optionally by importing a model via the FBX Mesh 3D node. Furthermore, you can add particle geometry to scenes from pEmitter nodes. You can connect these to a Merge3D node either singularly or in multiples to create sophisticated results combining multiple elements.



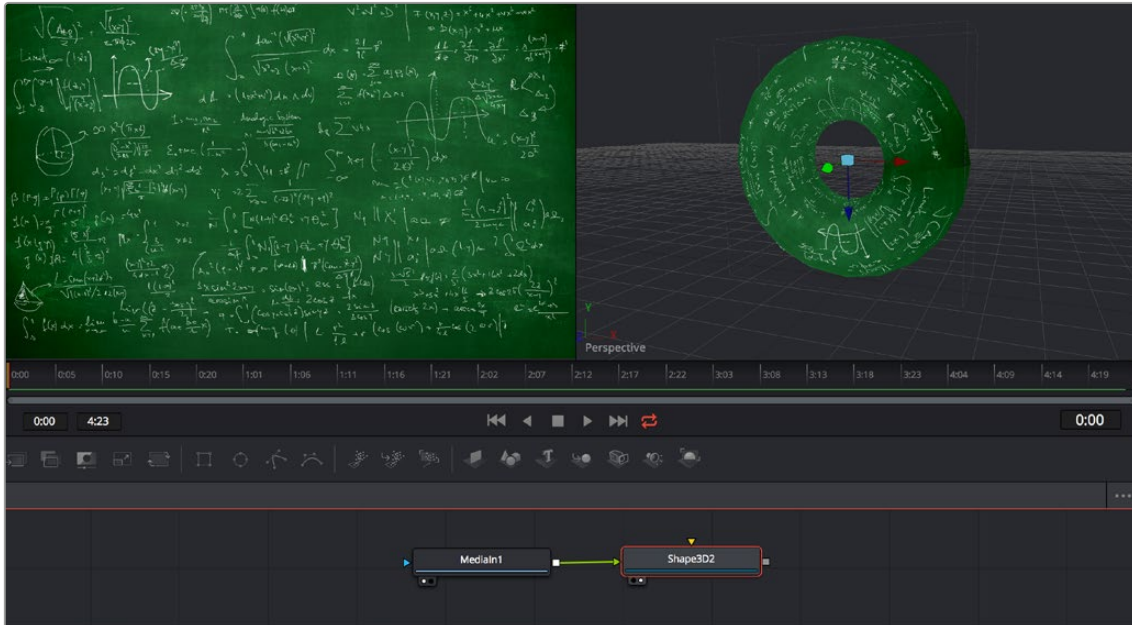
A more complex 3D scene combining several geometry nodes including the Text3D, Shape3D, and ImagePlane3D nodes

Texturing Geometry

By itself, geometry nodes can only consist of a simple flat color. However, you can alter the look of 3D geometry by texturing it using clips (either still images or movies), using material nodes such as the

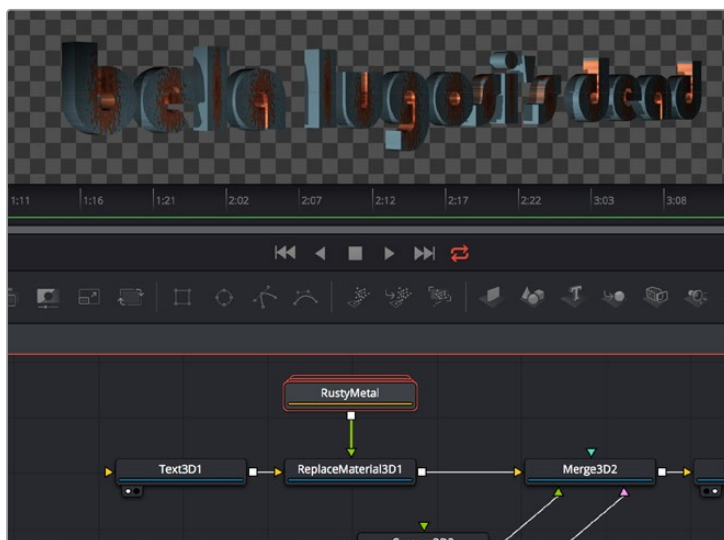
Blinn and Phong nodes to create more sophisticated textures with combinations of 2D images and environment maps, or you can use a preset shader from the Templates > Shader bin of the Effects Library, which contains materials and texture presets that are ready to use.

If you're working with simple geometric primitives, you can texture them by connecting either an image (a still image or movie) or a shader from the Templates bin of the Effects Library directly to the material input of a Shape3D, Cube3D, or other compatible node, as shown below.



An image connected to the material input of a Shape3D node set to Taurus, with the image (left), and the shaded taurus (right)

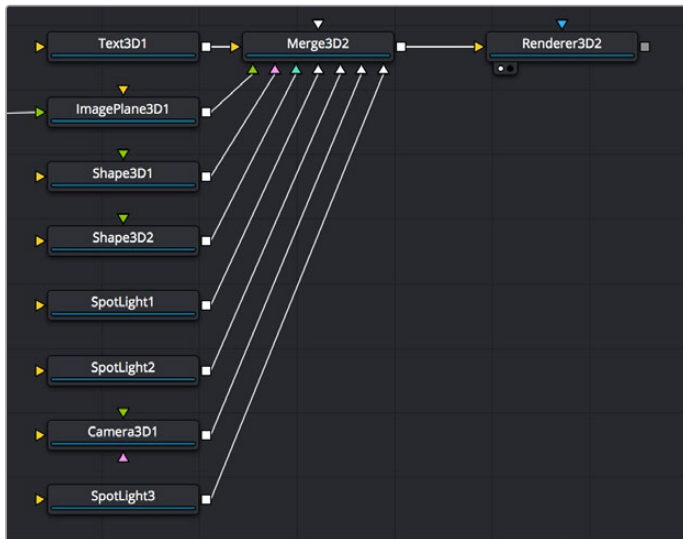
If you're shading or texturing Text3D nodes, you need to add a texture in a specific way since each node is actually a scene with individual 3D objects (the characters) working together. In the following example, the RustyMetal shader preset is applied to a Text3D node using the ReplaceMaterial3D node. The interesting thing about the ReplaceMaterial3D node is that it textures every geometric object within a scene at once, meaning that if you put a ReplaceMaterial3D node after a Text3D node, you texture every character within that node. However, if you place a ReplaceMaterial3D node after a Merge3D node, then you'll end up changing the texture of every single geometric object being combined within that Merge3D node, which is quite powerful.



The geometry created by a Text3D node is textured using a shader connected to a ReplaceMaterial3D node that's connected downstream of the object you want to shade

The Merge3D Node


The Merge3D node combines the output of one or more 3D nodes into a single scene. Unlike the Merge2D node, the ordering of elements in the scene is not restricted to only background and foreground inputs. Instead, the Merge3D node lets you connect an unlimited number of inputs, with the resulting output combined according to each object's absolute position in 3D space.



Merging many objects together in a 3D scene using the Merge3D node

Combining Objects Directly


While the Merge3D node provides a structured way of combining objects, you can also combine 3D objects such as Text3D and Shape3D nodes by connecting the output of one 3D object node to the input of another, as seen in the following screenshot. When you do this, you must use each node's internal transform parameters to transform their position, size, and rotation directly, but the transform control of downstream 3D object nodes also transforms all upstream 3D object nodes. This even works for lights and the Camera3D node, giving you a fast way of combining a set of objects that always go together, which you can later connect to a Merge3D node for additional lighting and eventual connection to a Renderer3D node.



Connecting one Shape3D node to another directly to combine them. Transforming the last downstream 3D object also transforms all upstream objects; the last Shape3D node is viewed, showing both

Combining Multiple Merge3D Nodes


Furthermore, Merge3D nodes can be combined with other Merge3D nodes, allowing you to create composite 3D scenes made up of multiple “sub-scenes,” each put together within individual Merge3D nodes.



You can build elaborate scenes using multiple Merge3D nodes connected together


Lighting Multiple Merge3D Nodes

Once you’ve combined multiple Merge3D nodes, there’s an easy way to control how lights that are connected to upstream Merge3D nodes affect the results of other Merge3D nodes connected downstream. Each Merge3D node’s Controls tab contains a single checkbox, Pass Through Lights, which enables lighting to pass through the output of an upstream Merge3D node in order to shine onto objects connected to downstream Merge3D nodes.



You can light downstream Merge3D scenes with lights connected to upstream Merge3D scenes by turning on Pass Through Lights

This checkbox is disabled by default, which lets you light elements in one Merge3D scene without worrying about how the lighting will affect geometry attached to other Merge3D nodes further downstream. For example, you may want to apply a spotlight to brighten the wall of a building in one Merge3D node without having that spotlight spill over onto the grass or pavement at the foot of the wall modeled in another Merge3D node. In the example shown below, the left image shows how the cone and taurus connected to a downstream node remain unlit by the light in an upstream node with Pass Through Lights disabled, while the right image shows how everything becomes lit when turning Pass Through Lights on.



The result of lights on the text in one Merge3D node not affecting the cone and taurus added in a downstream Merge3D node (left) Turning on Pass Through Lights in the upstream Merge3D node results in those lights also illuminating the downstream shapes (right)

Transforming Merge3D Scenes

Each Merge3D node includes a Transform tab. These transform parameters adjust the position, scale, and rotation of all objects being combined within that Merge3D node together, including lighting and particles. All transformations take place around a common pivot point. This forms the basis of parenting in the 3D environment.




The Transform tab of a Merge3D node

If you transform a Merge3D node that's connected to other Merge3D nodes, what happens depends on which node you're transforming, an upstream node or the downstream node:

- If you transform a downstream Merge3D node, you also transform all upstream nodes connected to it as if they were all a single scene.
- If you transform an upstream Merge3D node, this has no effect on downstream Merge3D nodes, allowing you to make transforms specific to that particular node's scene.

Transforming Upstream, Lighting Downstream

When building complex scenes using multiple Merge3D nodes being combined together, it's common to use one last downstream node to combine light and camera nodes to illuminate the final scene, while leaving the upstream Merge3D nodes free for controlling object transforms and animation. This way, you can transform and animate subsets of your overall scene without worrying about accidentally altering the overall lighting scheme or cameras for that scene, unless you've specifically connected lights or cameras upstream that are meant to be attached to the geometry you're transforming.




An example of a 3D scene using multiple Merge3D nodes working together; the upstream Merge3D nodes arrange the 3D objects placed within the scene, while the last Merge3D node (orange) lights and frames the scene.

The Renderer3D Node

Every 3D node you add outputs a complete 3D scene. This is unlike most traditional 3D modeling and animation programs, where all objects reside within a global scene environment. This means that the scenes created by a Camera 3D node and an image plane are separate until they're combined into the same scene via a Merge3D node, which itself outputs a complete 3D scene. However, this 3D scene data can neither be composited with other 2D images in your composition nor rendered out without first being rendered within the node tree using a Renderer3D node.

To be more specific, 3D nodes that output 3D scenes cannot be connected directly to inputs that require 2D images. For example, the output of an ImagePlane3D node cannot be connected directly to the input of a Blur node, nor can the output of a Merge3D node be directly connected to a regular Merge node. First, a Renderer3D node must be placed at the end of your 3D scene to render it into 2D images, which may then be composited and adjusted like any other 2D image in your composition.




Output of a Merge3D connected to a Renderer3D node to output 2D image data

The Renderer3D uses one of the cameras in the scene (typically connected to a Merge3D node) to produce an image. If no camera is found, a default perspective view is used. Since this default view rarely provides a useful angle, most people build 3D scenes that include at least one camera.

The image produced by the Renderer3D can be any resolution with options for fields processing, color depth, and pixel aspect.

Software vs. GPU Rendering

The Renderer3D node lets you choose between using a software renderer or an OpenGL renderer, trading off certain aspects of rendered image quality for speed, and trading off depth of field rendering for soft shadow rendering, depending on the needs of a particular element of your composition. To choose which method of rendering to use, there's a Renderer Type pop-up menu in the Controls tab of each Renderer3D node's parameters in the Inspector. The default is Software Renderer.




The Renderer Type option in the Controls tab of a Renderer3D node

Software Renderer

The software renderer is generally used to produce the final output. While the software renderer is not the fastest method of rendering, it has twin advantages. First, the software renderer can easily handle textures much larger than one half of your GPU's maximum texture size, so if you're working with texture images larger than 8K you should choose the software renderer to obtain maximum quality.

Second, the software renderer is required to enable the rendering of "constant" and "variable" soft shadows with adjustable Spread, which is not supported by the OpenGL renderer. Soft shadows are

more natural, and they're enabled in the Shadows parameters of the Controls tab of light nodes; you can choose Sampling Quality and Softness type, and adjust Spread, Min Softness, and Filter Size sliders. Additionally, the software renderer supports alpha channels in shadow maps, allowing transparency to alter shadow density.



When the Renderer3D node “Renderer Type” drop-down is set to OpenGL Renderer, you cannot render soft shadows or excessively large textures (left). When the Renderer3D node “Renderer Type” drop-down is set to Software Renderer, you can render higher-quality textures and soft shadows (right).

OpenGL Renderer

The OpenGL renderer takes advantage of the GPU in your computer to render the image; the textures and geometry are uploaded to the graphics hardware, and OpenGL shaders are used to produce the result. This can produce high-quality images that can be perfect for final rendering, and can also be potentially orders of magnitude faster than the software renderer, but it does pose some limitations on some rendering effects, as soft shadows cannot be rendered, and the OpenGL renderer also ignores alpha channels during shadow rendering, resulting in a shadow always being cast from the entire object.


On the other hand, because of its speed, the OpenGL renderer exposes additional controls for Accumulation Effects that let you enable depth of field rendering for creating shallow-focus effects. Unfortunately, you can't have both soft shadow rendering and depth of field rendering, so you'll need to choose which is more important for any given 3D scene you render.

Don't Forget That You Can Combine Rendered Scenes in 2D

While it may seem like an insurmountable limitation that you can't output both soft shadows and depth of field using the same renderer, don't forget that you can create multiple 3D scenes each using different renderers and composite them in 2D later on. Furthermore, you can also render out auxiliary channels that can be used by 2D image processing nodes such as AmbientOcclusion, DepthBlur, and Fog to create pseudo-3D effects using the rendered images.

OpenGL UV Renderer

When you choose the OpenGL UV Renderer option, a Renderer3D node outputs an “unwrapped” version of the textures applied to upstream objects, at the resolution specified within the Image tab of that Renderer3D node.



A normally rendered 3D scene (left), and the same scene rendered using the OpenGL UV Renderer mode of the Renderer3D node (right)

This specially output image is used for baking out texture projections or materials to a texture map for one of two reasons:

- Baking out projections can speed up a render.
- Baking out projections lets you modify the texture using other 2D nodes within your composition, or even using third-party paint applications (if you output this image in isolation as a graphics file) prior to applying it back onto the geometry.


Suppose, for instance, that you have a scene on a street corner, and there's a shop sign with a phone number on it, but you want to change the numbers. If you track the scene and have standing geometry for the sign, you can project the footage onto it, do a UV render, switch the numbers around with a Paint node, and then apply that back to the mesh with a Texture2D.

The UV renderer can also be used for retouching textures. You can combine multiple DSLR still shots of a location, project all those onto the mesh, UV render it out, and then retouch the seams and apply it back to the mesh.

You could project tracked footage of a road with cars on it, UV render out the projection from the geometry, do a temporal median filter on the frames, and then map a “clean” roadway back down.

Loading 3D Nodes into the Viewer


When you load a 3D node into the viewer, it switches to a 3D Viewer, which lets you pan, zoom, and rotate the scene in 3D, making it easy to make adjustments in three dimensions.



The 3D Viewer

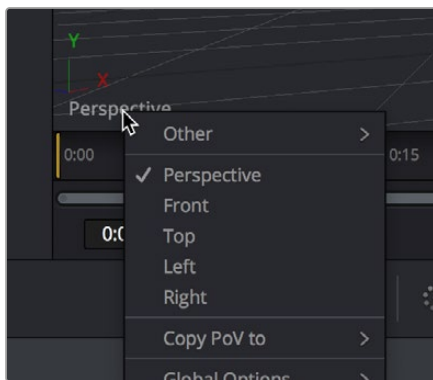
The interactive 3D Viewer is highly dependent on the computer's graphics hardware, relying on support from OpenGL. The amount of onboard memory, as well as the speed and features of your workstation's GPU, make a huge difference in the speed and capabilities of the 3D Viewer.

Displaying a node with a 3D output in any viewer will switch the display type to a 3D Viewer. Initially, the contents of the scene will be displayed through a default perspective view.



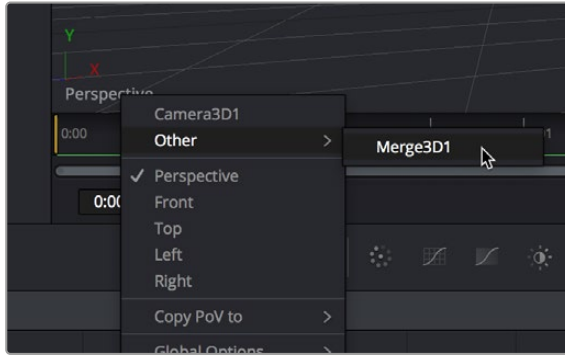
A 3D Viewer's default perspective view

To change the viewpoint, right-click in the viewer and choose the desired viewpoint from the ones listed in the Camera submenu. A shortcut to the Camera submenu is to right-click on the axis label displayed in the bottom corner of the viewer.



Right-click the Axis label of the viewer to change the viewpoint

In addition to the usual Perspective, Front, Top, Left, and Right viewpoints, if there are cameras and lights present in the scene as potential viewpoints, those are shown as well. It's even possible to display the scene from the viewpoint of a Merge3D or Transform3D by selecting it from the contextual menu's Camera > Other submenu. Being able to move around the scene and see it from different viewpoints can help with the positioning, alignment, and lighting, as well as other aspects of your composite.



The Perspective drop-down menu also shows cameras, lights, and Merge3D and Transform3D nodes you can switch to.

Navigating the 3D View

For the most part, panning and scaling of the 3D Viewer uses the same controls as the 2D Viewer. For more information about the options available in the 3D Viewer, see Chapter 67, “Using Viewers,” in the DaVinci Resolve Reference Manual or Chapter 7 in the Fusion Reference Manual.

To pan in a 3D Viewer, do the following:

- Hold the middle mouse button and drag in the viewer.

To dolly (zoom) in the 3D Viewer, do one of the following:

- Hold down the middle and left mouse buttons and drag left or right in the viewer.
- Hold down the Command key and use your pointing device’s scroll control.

To rotate around the 3D Viewer, do the following:

- Hold down the Option key and middle-button-drag left and right in the viewer.

If you want to frame certain objects in the viewer:

- 1 Select the viewer you want to work in.
- 2 Do one of the following:
 - Press Shift-F to Fit all objects in the viewer.
 - Press F to Fit to selection (or Fit All if nothing is selected).
 - Press D to Rotate the viewer to look at the center of the currently selected object without moving the viewer’s position.

Furthermore, selecting a 3D node in the Node Editor also selects the associated object in the 3D Viewer.

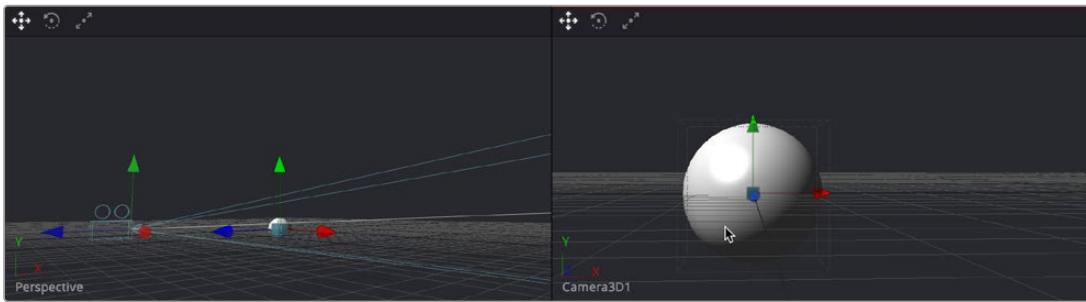
Transforming Cameras and Lights Using the Viewers

When the viewer is set to look through a 3D object in the scene, such as a camera or spotlight, the usual controls for panning and rotating the viewer will now directly affect the position of the camera or spotlight you’re viewing through. Here’s an example.

To adjust a camera’s position when looking through it in a viewer:

- 1 Right-click the viewpoint label, and choose a camera from the contextual menu. (Optional) If you’re in dual-viewer mode, you can load the camera you’ve selected in one viewer into the other viewer to see its position as you work.
- 2 Move the pointer into the viewer that’s displaying the camera’s viewpoint.

- 3 Hold the middle and left mouse buttons down and drag to zoom the viewer, or middle-click-drag to pan the viewer, or option-middle-click-drag to rotate the viewer, all while also moving the camera.



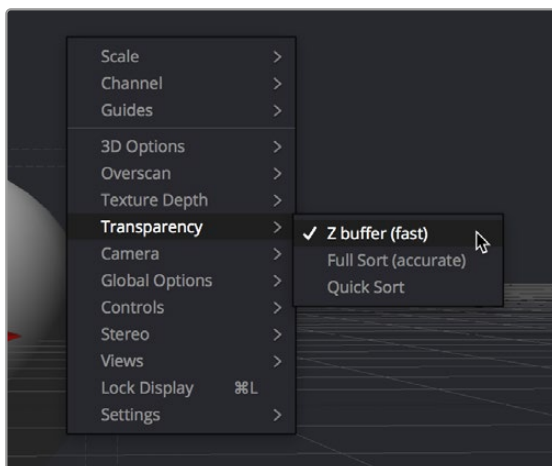
When a viewer is set to display the view of a camera or light, panning, zooming, or rotating the viewer (seen at right) actually transforms the camera or light you're viewing through (seen at left)

It is even possible to view the scene from the perspective of a Merge3D or Transform3D node by selecting the object from the Camera > Others menu. The same transform techniques will then move the position of the object. This can be helpful when you are trying to orient an object in a certain direction.

Transparency Sorting

While generally the order of geometry in a 3D scene is determined by the Z-position of each object, sorting every face of every object in a large scene can take an enormous amount of time. To provide the best possible performance, a Fast Sorting mode is used in the OpenGL renderer and viewers. This is set by right-clicking in the viewer and choosing Transparency > Z-buffer. While this approach is much faster than a full sort, when objects in the scene are partially transparent it can also produce incorrect results.

The Sorted (Accurate) mode can be used to perform a more accurate sort at the expense of performance. This mode is selected from the Transparency menu of the viewer's contextual menu. The Renderer3D also presents a Transparency menu when the Renderer Type is set to OpenGL. Sorted mode does not support shadows in OpenGL. The software renderer always uses the Sorted (Accurate) method.




Transparency Sorting in the viewer contextual menu

The basic rule is when a scene contains overlapping transparency, use the Full/Quick Sort modes, and otherwise use the Z-buffer (Fast). If the Full Sort method is too slow, try switching back to Z-buffer (Fast).

Material Viewer

When you view a node that comes from the 3D > Material category of nodes in the Effects Library, the viewer automatically switches to display a Material Viewer. This Material Viewer allows you to preview the material applied to a lit 3D sphere rendered with OpenGL by default.



The Material Viewer mode of the viewer

The type of geometry, the renderer, and the state of the lighting can all be set by right-clicking the viewer and choosing options from the contextual menu. Each viewer supports A and B buffers to assist with comparing multiple materials.

Methods of working with the Material Viewer:

- You can change the shape of the previewed geometry by right-clicking the viewer and choosing an option from the Shape submenu of the contextual menu. The geometry that the material is applied to is locked to the center of the viewer and scaled to fit. It is not possible to pan or scale the Material Viewer.
- The Material Viewer can be rotated to provide a different angle on the material by holding Option while pressing the middle mouse button and dragging to the left and right.
- You can adjust the position of the light used to preview the material by dragging with the middle mouse button. Or, you can right-click the viewer and choose an option from the Lighting > Light Position submenu of the contextual menu.
- You can also toggle lighting off and on by right-clicking the viewer and choosing Lighting > Enable Lighting from the contextual menu.
- You can choose the renderer used to preview the material by right-clicking the viewer and choosing an option from the Renderer submenu of the contextual menu.

Transformations

Merge3D, 3D Objects, and Transform3D all have Transform parameters that are collected together into a Transform tab in the Inspector. The parameters found in this tab affect how the object is positioned, rotated, and scaled within the scene.



The Transform tab of a Merge3D node


The Translation parameters are used to position the object in local space, the Rotation parameters affect the object's rotation around its own center, and the Scale slider(s) affect its size (depending on whether or not they're locked together). The same adjustments can be made in the viewer using onscreen controls.

Onscreen Transform Controls

When an object is selected, it displays onscreen Transform controls in the viewers that allow you to adjust the object's position, rotation, and scale. Buttons in the Transform toolbar allow you to switch modes, or you can use the keyboard shortcuts.

To switch Transform modes, use the following keyboard shortcuts:


- Press Q for Position
- Press W for Rotation
- Press E for Scaling



The Position, Rotation, and Scale modes in the Transform toolbar

Using Onscreen Transform Controls

In all three modes, red indicates the object's local X-axis, green the Y-axis, and blue the Z-axis, respectively (just remember RGB = XYZ). You can drag directly on the red, green, or blue portion of any onscreen control to constrain the transform to that axis, or if you drag the center of the onscreen control, you can apply a transform without constraints. Holding Option and dragging in the viewer allows you to freely translate in all three axes without clicking on a specific control.



From left to right, the Position, Rotation, and Scale onscreen Transform controls

If the Scale's Lock XYZ checkbox is enabled in the Inspector, only the overall scale of the object is adjusted by dragging the red or center onscreen control, while the green and blue portions of the onscreen control have no effect. If you unlock the parameters, you are able to scale an object along individual axes separately to squish or stretch the object.

Selecting Objects

With the onscreen controls visible in the viewer, you can select any object by clicking on its center control. Alternatively, you can also select any 3D object by clicking its node in the Node Editor.

Pivot


In 3D scenes, objects rotate and scale around an axis called a pivot. By default, this pivot goes through the object's center. If you want to move the pivot so it is offset from the center of the object, you can use the X, Y, and Z Pivot parameters in the Inspector.

Target

Targets are used to help orient a 3D object to a specific point in the scene. No matter where the object moves, it will rotate in the local coordinate system so that it always faces its target, which you can position and animate.

To enable a target for a 3D object:


- 1 Select that object's node.
- 2 Open the object's Transform panel in the Inspector.
- 3 Turn on the Use Target checkbox.



Turning on the Use Target checkbox of a 3D object


- 4 Use the X/Y/Z Target Position controls in the Inspector or the Target onscreen control in the viewer to position the target and in turn position the object it's attached to.

In the viewer, a line is drawn between the target and the center of the 3D object it's attached to, to show the relationship between these two sets of controls. Whenever you move the target, the object is automatically transformed to face its new position.



A 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.

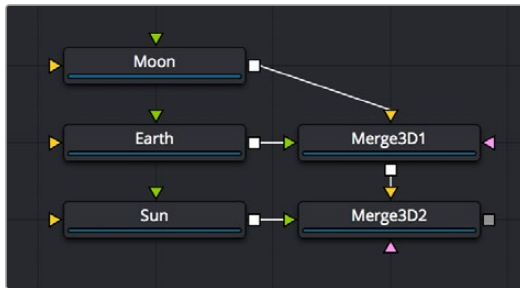


A light made to face the wall using its enabled target control

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 Merge3D. The Merge3D can be used to rotate one sphere around the other, like the moon around the earth. Then the Merge3D can be connected to another Merge3D to create the earth and the moon orbiting around the sun.



One Merge3D with two spheres parented to another Merge3D and parenting using three connected spheres


Here are the two simple rules of transforming parented Merge3D nodes:

- Transforms and animation applied to a Merge3D are also applied to all 3D objects connected to that Merge3D node, including cameras, lights, geometry, and other merge nodes connected upstream.
- Transforms and animation applied to upstream merge nodes don't affect downstream merge nodes.

Cameras

When setting up and animating a 3D scene, the metaphor of a camera is one of the most comprehensible ways of framing how you want that scene to be rendered out, as well as animating your way through the scene. Additionally, compositing artists are frequently tasked with matching cameras from live-action clips, or matching cameras from 3D applications.

To accommodate all these tasks, Fusion provides a flexible Camera3D node with common camera controls such as Angle of View, Focal Length, Aperture, and Clipping planes, to either set up your own camera or to import camera data from other applications. The Camera3D node is a virtual camera through which the 3D environment can be viewed.



A camera displayed with onscreen Transform controls in the viewer; the Focal Plane indicator is enabled in green

Cameras are typically connected and viewed via a Merge3D node; however, you can also connect cameras upstream of other 3D objects if you want that camera to transform along with that object when it moves.

Quickly Viewing a Scene Through a Camera

When you've added a camera to a scene, you can quickly view the scene "through the camera" by setting up the following.

To view the scene through the camera:

- 1 Select the Merge3D node that the camera is connected to, or any node downstream of that Merge3D.
- 2 Load the selected Merge3D or downstream node into a viewer.
- 3 Right-click on the axis label in the bottom corner of the viewer and choose the camera name.

The viewer's frame may be different from the camera frame, so it may not match the true boundaries of the image that will be rendered by the Renderer3D node. If there is no Renderer3D node added to your scene yet, you can use Guides that represent the camera's framing. For more information about Guides, see Chapter 67, "Using Viewers," in the DaVinci Resolve Reference Manual or Chapter 7 in the Fusion Reference Manual.


Plane of Focus and Depth of Field

Cameras have a plane of focus, for when depth of field rendering is available. Here's the procedure for enabling depth of field rendering in your scenes.

To render depth of field in a 3D scene:

- 1 You must add a Renderer3D node at the end of your 3D scene.
- 2 Select the Renderer3D node, and set the Renderer Type to OpenGL Renderer.

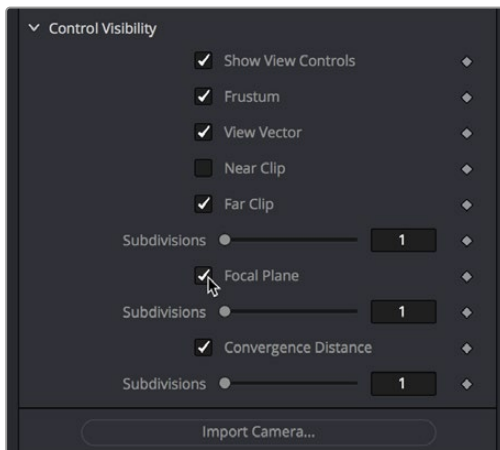
- 3 Open the Accumulation Effects disclosure control that appears, and turn on the Enable Accumulation Effects checkbox in the OpenGL render.



Turning on Enable Accumulation Effects enables additional depth of field controls

Turning on “Enable Accumulation Effects” exposes a Depth of Field checkbox along with Quality and Amount of DoF Blur sliders that let you adjust the depth of field effect. These controls affect only the perceived quality of the depth of field that is rendered. The actual depth of field that’s generated depends solely on the setup of the camera and its position relative to the other 3D objects in your scene.

When you select your scene’s Camera3D node to view its controls in the Inspector, a new Focal Plane checkbox appears in the Control Visibility group. Turning this on lets you see the green focal plane indicator in the 3D Viewer that lets you visualize the effect of the Focal Plane slider, which is located in the top group of parameters in the Camera3D node’s Controls tab.



Turning on the Focal Plane checkbox in the Camera3D node

For more information about these specific camera controls, see Chapter 89, “3D Nodes,” in the DaVinci Resolve Reference Manual or Chapter 29 in the Fusion Reference Manual.

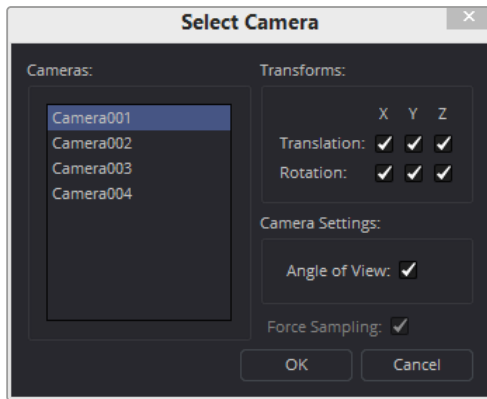
Importing Cameras

If you want to match cameras between applications, you can import camera paths and positions from a variety of popular 3D applications. Fusion is able to import animation splines from Maya and XSI directly with their own native spline formats. Animation applied to cameras from 3ds Max and LightWave are sampled and keyframed on each frame.

To import a camera from another application, do the following:

- 1 Select the camera in the Node Editor.
- 2 At the bottom of the Inspector, click the Import Camera button.
- 3 In the file browser, navigate to and select the scene that contains the camera you want to import.

A dialog box with several options will appear. When the Force Sampling checkbox is enabled, Fusion will sample each frame of the motion, regardless of the format.



The Select Camera dialog

TIP: When importing parented or rigged cameras, baking the camera animation in the 3D application before importing it into Fusion often produces more reliable results.

Lighting and Shadows


You can add light sources to a scene to create very detailed lighting environments and atmosphere. There are four different types of lights you can use in 3D scenes: ambient, directional, point, and spotlights.

Enabling Lighting in the Viewer

A scene without lights uses a default directional light, but this automatically disappears once you add a 3D light object. However, even when you add light objects to your scene, lighting and shadows won't be visible in the viewer unless you first enable lighting in the viewer contextual menu by right-clicking anywhere within a viewer and choosing 3D Options > Lighting or Shadows to turn on one or both.

Enabling Lighting to Be Rendered

Lighting effects won't be rendered in the Renderer3D node until the Enable Lighting and/or Shadows checkboxes are checked in the Inspector.




The Lighting button under the viewer

NOTE: When lighting is disabled in either the viewer or final renders, the image will appear to be lit by a 100% ambient light.

Cotrolling Lighting within Each 3D Object

All nodes that create or merge geometry also include lighting options that are used to choose how each object is affected by light:

- Merge3D nodes have a Pass Through Lights checkbox that determines whether lights attached to an upstream Merge3D node also illuminate objects attached to downstream Merge3D nodes.
- ImagePlane3D, Cube3D, Shape3D, Text3D, and FBXMesh3D nodes have a set of Lighting controls that let you turn three controls on and off: Affected by Lights, Shadow Caster, and Shadow Receiver.



3D objects have individual lighting controls that let you control how each object interacts with light and shadows

Lighting Types Explained

Here's a more detailed explanation of each type of light in Fusion.

Ambient Light

You use ambient light to set a base light level for the scene, since it produces a general uniform illumination of the scene. Ambient light exists everywhere without appearing to come from any particular source; it cannot cast shadows and will tend to fill in shadowed areas of a scene.

Directional Light


A directional light is composed of parallel rays that light up the entire scene from one direction, creating a wall of light. The sun is an excellent example of a directional light source.

Point Light

A point light is a well defined light that has a small clear source, like a light bulb, and shines from that point in all directions.

Spotlight

A spotlight is an advanced point light that produces a well defined cone of light with falloff. This is the only light that produces shadows.




From left to right: Directional light, point light, and spotlight

All of the Light nodes display onscreen controls in the viewer, although not all controls affect every light type. In the case of the ambient light, the position has no effect on the results. The directional light can be rotated, but position and scale will be ignored. The point light ignores rotation. Both position and rotation apply to the spotlight.

Lighting Hierarchies


Lights normally do not pass through a Merge, since the Pass Through Lights checkbox is off by default. This provides a mechanism for controlling which objects are lit by which lights. For example, in the following two node trees, two shapes and an ambient light are combined with a Merge3D node, which is then connected to another Merge3D node that's also connected to a plane and a spotlight. At the left, the first Merge3D node of this tree has Pass Through Lights disabled, so you can only see the two shapes lit. At the right, Pass Through Lights has been enabled, so both the foreground shapes and the background image plane receive lighting.



Pass Through Lights is disabled, so only the front two shapes are illuminated (left) Pass Through Lights is enabled, so all shapes connected to both Merge3D nodes are illuminated (right)

Lighting Options

Most nodes that generate geometry have additional options for lighting. These options are used to determine how each individual object reacts to lights and shadows in the scene.




3D objects have individual lighting controls that let you control how each object interacts with light and shadows

- **Affected By Lights:** If the Affected By Lights checkbox is enabled, lights in the scene will affect the geometry.
- **Shadow Caster:** When enabled, the object will cast shadows on other objects in the scene.
- **Shadow Receiver:** If this checkbox is enabled, the object will receive shadows.

Shadows

The only light that can cast shadows is the spotlight. Spotlight nodes cast shadows by default, although these shadows will not be visible in the viewer until shadows are enabled using the viewer toolbar button. Shadows will not appear in the output of the Renderer3D unless the Shadows option is enabled for that renderer. If you want to prevent a spotlight from casting shadows, you can disable the Enable Shadows checkbox in the node's Inspector.



An image with spotlight casting a variable soft shadow

For more information on shadow controls, see the “Spotlight” section of Chapter 90, “3D Light Nodes,” in the DaVinci Resolve Reference Manual or Chapter 30 in the Fusion Reference Manual.

Shadow Maps

A shadow map is an internal depth map that specifies each pixel's depth in the scene. This information is used to assemble the shadow layer created from a spotlight. All the controls for the shadow map are found in the Spotlight Inspector.


The quality of the shadow produced depends greatly on the size of the shadow map. Larger maps generate better-looking shadows but will take longer to render. The wider the cone of the spotlight, or the more falloff in the cone, the larger the shadow map will need to be to produce useful quality results. Setting the value of the Shadow Map Size control sets the size of the depth map in pixels.

Generally, through trial and error, you'll find a point of diminishing returns where increasing the size of the shadow map no longer improves the quality of the shadow. It is not recommended to set the size of the shadow maps any larger than they need to be.

The Shadow Map Proxy control is used to set a percentage by which the shadow map is scaled for fast interactive previews, such as Autoproxy and LoQ renders. A value of .4, for example, represents a 40% proxy.


Shadow Softness

By default, the spotlight generates shadows without soft edges, but there are options for constant and variable soft shadows. Hard-edged shadows will render significantly faster than either of the Soft Shadow options. Shadows without softness will generally appear aliased, unless the shadow map size is large enough. In many cases, softness is used to hide the aliasing rather than increasing the shadow map to preserve memory and avoid exceeding the graphics hardware capabilities.



Soft Shadow controls in the Control panel

Setting the spotlight's shadow softness to None will render crisp and well-defined shadows. The Constant option will generate shadows where the softness is uniform across the shadow, regardless of the shadow's distance from the casting geometry. The Variable option generates shadows that become softer as they get farther from the geometry that is casting the shadow. This is a more realistic effect, but the shadows are somewhat harder to control. When this option is selected, additional controls for adjusting the falloff of the shadow will appear, as well as sliders for the minimum and maximum softness.



Hard shadow cast by a spotlight


Selecting the Variable option reveals the Spread, Min Softness, and Filter Size sliders. A side effect of the method used to produce variable softness shadows is that the size of the blur applied to the shadow map can become effectively infinite as the shadow's distance from the geometry increases. These controls are used to limit the shadow map by clipping the softness calculation to a reasonable limit.

The filter size determines where this limit is applied. Increasing the filter size increases the maximum possible softness of the shadow. Making this smaller can reduce render times but may also limit the softness of the shadow or potentially even clip it. The value is a percentage of the shadow map size.

For more information, see "Spotlight" in Chapter 90, "3D Light Nodes," in the DaVinci Resolve Reference Manual or Chapter 30 in the Fusion Reference Manual.


Multiplicative and Additive Bias

Shadows are essentially textures applied to objects in the scene that occasionally result in "fighting." Z-fighting results when portions of an object that should be receiving shadows instead render over the top of the shadow because they effectively exist in the same exact location in 3D space.



Results of shadow map Z-fighting (top), and the corrected shadow shown using Biasing (bottom)

Two Biasing sliders in the Shadows group of Spotlight parameters work by adding a small depth offset to move the shadow away from the surface it is shadowing, eliminating the Z-fighting. When too little bias is added, the objects can self shadow themselves. When too much is added, the shadow can become separated from the surface.



The Multiplicative and Additive Bias sliders, and the Non-Transmissive Materials checkbox, all in the Spotlight Inspector controls

The goal is to adjust the Multiplicative Bias slider until the majority of the Z-fighting is resolved, and then adjust the Additive Bias slider to eliminate the rest. The softer the shadow, the higher the bias will probably have to be. You may even need to animate the bias to get a proper result for some particularly troublesome frames.

Force All Materials Non-Transmissive

How light passes through a semi-transparent material plays an important role in determining the appearance of the shadow an object casts. Normally, this transmittance behavior is defined in each object's Materials tab. However, selecting Force All Materials Non-Transmissive in the Spotlight Inspector overrides this, causing the shadow map produced by the node to ignore transmittance entirely.

Materials and Textures


To render a 3D scene, the renderer must take into account the shape of the object as well as its appearance. The geometry of an object defines the shape of the object, while the material applied to the object defines its appearance. Fusion provides a range of options for applying materials and textures to geometry, so you can give your 3D objects the surface qualities you want.

Nodes that describe the geometry's response to light are called illumination models. Blinn, Cook-Torrance, Ward, and Phong are the included illumination models. These nodes are found in the 3D > Material category of nodes in the Effects Library.

Most materials also accept textures, which are typically 2D images. Textures are used to refine the look of an object further, by adding photorealistic details, transparency, or special effects. More complex textures like bump maps, 3D textures, and reflection maps are also available in the 3D > Texture category.

Materials can also be combined to produce elaborate and highly detailed composite materials.

Each node that creates or loads geometry into a 3D scene also assigns a default material. The default material is the Blinn illumination model, but you can override this material using one of several nodes that output a 3D material. Some of these materials provide a greater degree of control over how the geometry reacts to light, providing inputs for diffuse and specular texture maps, bump mapping, and environmental maps, which mimic reflection and refraction.



Material examples from the bin

Material Components

All the standard illumination models share certain characteristics that must be understood.

Diffuse

The Diffuse parameters of a material control the appearance of an object where light is absorbed or scattered. This diffuse color and texture are the base appearance of an object, before taking into account reflections. The opacity of an object is generally set in the diffuse component of the material.

Alpha

The Alpha parameter defines how much the object is transparent to diffuse light. It does not affect specular levels or color. However, if the value of alpha, either from the slider or a Material input from the diffuse color, is very close to or at zero, those pixels, including the specular highlights, will be skipped and disappear.

Opacity


The Opacity parameter fades out the entire material, including the specular highlights. This value cannot be mapped; it is applied to the entire material.

Specular

The Specular parameters of a material control the highlight of an object where the light is reflected to the current viewpoint. This causes a highlight that is added to the diffuse component. The more


specular a material is, the glossier it appears. Surfaces like plastics and glass tend to have white specular highlights, whereas metallic surfaces like gold have specular highlights that tend to inherit their color from the material color.

Specularity is made up of color, intensity, and exponent. The specular color determines the color of light that reflects from a shiny surface. Specular intensity is how bright the highlight will be.



Three spheres, left to right: diffuse only, specular only, and combined

The specular exponent controls the falloff of the specular highlight. The larger the value, the sharper the falloff and the smaller the specular component will be.



Left to right: white, complimentary, and matching specular colors

Transmittance

When using the software renderer, the Transmittance parameters control how light passes through a semi-transparent material. For example, a solid blue pitcher will cast a black shadow, but one made of translucent blue plastic would cast a much lower density blue shadow. The transmittance parameters are essential to creating the appearance of stained glass.

TIP: You can adjust the opacity and transmittance of a material separately. It is possible to have a surface that is fully opaque yet transmits 100% of the light arriving upon it, so in a sense it is actually a luminous/emissive surface.

Transmissive surfaces can be further limited using the Alpha and Color Detail control.

Attenuation


The transmittance color determines how much color is passed through the object. For an object to have fully transmissive shadows, the transmittance color must be set to $RGB = (1, 1, 1)$, which means 100% of green, blue, and red light passes through the object. Setting this color to $RGB = (1, 0, 0)$ means that the material will transmit 100% of the red arriving at the surface but none of the green or blue light.

Alpha Detail


When the Alpha Detail slider is set to 0, the non-zero portions of the alpha channel of the diffuse color are ignored and the opaque portions of the object casts a shadow. If it is set to 1, the alpha channel determines how dense the object casts a shadow.

NOTE: The OpenGL renderer ignores alpha channels for shadow rendering, resulting in a shadow always being cast from the entire object. Only the software renderer supports alpha in the shadow maps.

The following examples for Alpha Detail and Color Detail cast a shadow using this image. It is a green-red gradient from left to right. The outside edges are transparent, and inside is a small semi-transparent circle.



Alpha Detail set to 1; the alpha channel determines the density of the shadow




Alpha Detail set to 0; a dense-colored shadow results

Color Detail

Color Detail is used to color the shadow with the object's diffuse color. Increasing the Color Detail slider from 0 to 1 brings in more diffuse color and texture into the shadow.

TIP: The OpenGL renderer will always cast a black shadow from the entire object, ignoring the color. Only the software renderer supports color in the shadow maps.



Color Detail set to 0; no color is visible in the shadow.

Saturation


Saturation will allow the diffuse color texture to be used to define the density of the shadow without affecting the color. This slider lets you blend between the full color and luminance only.

Transmittance and Shadows

The transmittance of an object's material plays an important role in determining the appearance of the shadow it casts. Normally, the transmittance behavior is defined in each object's Materials tab as explained above. However, selecting Force All Materials Non-Transmissive in the Spotlight Inspector overrides this, causing the shadow map produced by the spotlight to ignore transmittance entirely.

Illumination Models

Now that you understand the different components that make up a material or shader, we'll look at them more specifically. Illumination models are advanced materials for creating realistic surfaces like plastic, wood, or metal. Each illumination model has advantages and disadvantages, which make it appropriate for particular looks. An illumination model determines how a surface reacts to light, so these nodes require at least one light source to affect the appearance of the object. Four different illumination models can be found in the Nodes > 3D > Material menu.



Illumination models left to right: Standard, Blinn, Phong, Cook-Torrance, and Ward

Standard

The Standard material provides a default Blinn material with basic control over the diffuse, specular, and transmittance components. It only accepts a single texture map for the diffuse component with the alpha used for opacity. The Standard Material controls are found in the Material tab of all nodes that load or create geometry. Connecting any node that outputs a material to that node's Material Input will override the Standard material, and the controls in the Material tab will be hidden.

Blinn

The Blinn material is a general purpose material that is flexible enough to represent both metallic and dielectric surfaces. It uses the same illumination model as the Standard material, but the Blinn material allows for a greater degree of control by providing additional texture inputs for the specular color, intensity, and exponent (falloff), as well as bump map textures.

Phong

The Phong material produces the same diffuse result as Blinn, but with wider specular highlights at grazing incidence. Phong is also able to make sharper specular highlights at high exponent levels.

Cook-Torrance

The Cook-Torrance material combines the diffuse illumination model of the Blinn material with a combined microfacet and Fresnel specular model. The microfacets need not be present in the mesh or bump map; they are represented by a statistical function, Roughness, which can be mapped. The Fresnel factor attenuates the specular highlight according to the Refractive Index, which can be mapped.

Ward

The Ward material shares the same diffuse model as the others but adds anisotropic highlights, ideal for simulating brushed metal or woven surfaces, as the highlight can be elongated in the U or V directions of the mapping coordinates. Both the U and V spread functions are mappable.


This material does require properly structured UV coordinates on the meshes it is applied to.

Textures

Texture maps modify the appearance of a material on a per-pixel basis. This is done by connecting an image or other material to the inputs on the Material nodes in the Node Editor. When a 2D image is used, the UV mapping coordinates of the geometry are used to fit the image to the geometry, and when each pixel of the 3D scene is rendered, the material will modify the material input according to the value of the corresponding pixel in the map.

TIP: UV Mapping is the method used to wrap a 2D image texture onto 3D geometry. Similar to X and Y coordinates in a frame, U and V are the coordinates for textures on 3D objects.

Texture maps are used to modify various material inputs, such as diffuse color, specular color, specular exponent, specular intensity, bump map, and others. The most common uses of texture maps is the diffuse color/opacity component.



The Fast Noise texture used to control the roughness of a Cook-Torrance material

A node that outputs a material is frequently used, instead of an image, to provide other shading options. Materials passed between nodes are RGBA samples; they contain no other information about the shading or textures that produced them.




The Texture2D node is used to translate a texture in the UV space of the object, as well as set the filtering and wrap mode

Composite Materials

Building complex materials is as easy as connecting the output of a Material node to one of the Material inputs of another Material or Texture node. When a Material input is supplied just as with a 2D image, its RGBA values are used per pixel as a texture. This allows for very direct compositing of shaders.

For instance, if you want to combine an anisotropic highlight with a Blinn material, you can take the output of the Blinn, including its specular, and use it as the diffuse color of the Ward material. Or, if you do not want the output of the Blinn to be relit by the Ward material, you can use the Channel Boolean material to add the Ward material's anisotropic specular component to the Blinn material with a greater degree of control.



Combining an anisotropic highlight with a Blinn material using the Channel Boolean material

Reflections and Refractions

Environment maps can be applied with the Reflect material in the 3D > Material category. This node can be used to simulate reflections and refractions on an object. Reflections are direct-bounce light that hits an object, while refractions simulate the distortion of light seen through semi-transparent surfaces.

The reflections and refractions use an environment mapping technique to produce an approximation that balances realistic results with greater rendering performance. Environment maps assume an object's environment is infinitely distant from the object and rendered into a cubic or spherical texture surrounding the object.

The Nodes > 3D > Texture > Cube Map and Sphere Map nodes can be used to help create environment maps, applying special processing and transforms to create the cubic or spherical coordinates needed.




Sphere map example

To produce reflections with real-time interactive feedback at a quality level appropriate for production environment maps, you make some trade-offs on functionality when compared with slower but physically accurate raytraced rendering. Environment-mapped reflections and refractions do not provide self-reflection or any other kind of interaction between different objects. In particular, this infinite distance assumption means that objects cannot interact with themselves (e.g., the reflections on the handle of a teapot will not show the body of the teapot). It also means that objects using the same cube map will not inter-reflect with each other. For example, two neighboring objects would not reflect each other. A separate cube map must be rendered for each object.


The Reflect node outputs a material that can be applied to an object directly, but the material does not contain an illumination model. As a result, objects textured directly by the Reflect node will not respond to lights in the scene. For this reason, the Reflect node is usually combined with the Blinn, Cook-Torrance, Phong, or Ward nodes.

Reflection

Reflection outputs a material making it possible to apply the reflection or refraction to other materials either before or after the lighting model with different effects.



A Blinn material connected to a background material input of the Reflect. This causes the reflection to be added to the Blinn output.



A Reflect is connected to the Diffuse Color component of the Blinn, causing the reflection to be multiplied by the diffuse color and modulated by the lighting.

Refraction

Refraction occurs only where there is transparency in the background material, which is generally controlled through the Opacity slider and/or the alpha channel of any material or texture used for the Background Material Texture input. The Reflect node provides the following material inputs:


- **Background Material:** Defines both the opacity for refraction and the base color for reflection.
- **Reflection Color Material:** The environment reflection.
- **Reflection Intensity Material:** A multiplier for the reflection.
- **Refraction Tint Material:** The environment refraction.
- **Bump Map Texture:** Normal perturbing map for environment reflection/refraction vectors.

Working with reflection and refraction can be tricky. Here are some techniques to make it easier:

- Typically, use a small amount of reflection, between 0.1 and 0.3 strength. Higher values are used for surfaces like chrome.
- Bump maps can add detail to the reflections/refractions. Use the same bump map in the Illumination model shader that you combine with Reflect.
- When detailed reflections are not required, use a relatively small cube map, such as 128 x 128 pixels, and blur out the image.
- The alpha of refracted pixels is set to 1 even though the pixels are technically transparent. Refracted pixels increase their alpha by the reflection intensity.
- If the refraction is not visible even when a texture is connected to the Refraction Tint Material input, double-check the alpha/opacity values of the background material.

Bump Maps

Bump mapping helps add details and small irregularities to the surface appearance of an object. Bump mapping modifies the geometry of the object or changes its silhouette.



Split screen of a sphere—half with bump map, half without

To apply a bump map, you typically connect an image containing the bump information to the BumpMap node. The bump map is then connected to the Bump input of a Material node. There are two ways to create a bump map for a 3D material: a height map and a bump map.




Image connected to a BumpMap connected to a CookTorrance material node

Using a Height Map

A height map is an image where the value of a pixel represents the height. It is possible to select which color channel is used for bump creation. White means high and black means low; however, it is not the value of a pixel in the height map that determines the bumpiness, but rather how the value changes in the neighborhood of a pixel.

Using a Bump Map

A bump map is an image containing normals stored in the RGB channels.

TIP: Normals are generated by 3D modeling and animation software as a way to trick the eye into seeing smooth surfaces, even though the geometry used to create the models uses only triangles to build the objects.

Normals are 3 float values (nx, ny, nz) whose components are in the range $[-1, +1]$. Because you can store only positive values in Fusion's integer images, the normals are packed from the range $[-1, +1]$ to the range $[0, 1]$ by multiplying by 0.5 and adding 0.5. You can use Brightness Contrast or a Custom node to do the unpacking.

If you were to connect a bump map directly to the bump map input of a material, it will result in incorrect lighting. Fusion prevents you from doing this, however, because Fusion uses a different coordinate system for doing the lighting calculation. You first must use a BumpMap that expects a packed bump map or height map and will do the conversion of the bump map to work correctly.

If your bump mapping doesn't appear correct, here are a few things to look for:

- Make sure you have the nodes connected correctly. The height/bump map should connect into a BumpMap and then, in turn, should connect into the bump map input on a material.
- Change the precision of the height map to get less banding in the normals. For low frequency images, float32 may be needed.
- Adjust the Height scale on the BumpMap. This scales the overall effect of the bump map.
- Make sure you set the type to HeightMap or BumpMap to match the image input. Fusion cannot detect which type of image you have.
- Check to ensure High Quality is on (right-click in the transport controls bar and choose High Quality from the contextual menu). Some nodes like Text+ produce an anti-aliased version in High Quality mode that will substantially improve bump map quality.
- If you are using an imported normal map image, make sure it is packed $[0-1]$ in RGB and that it is in tangent space. The packing can be done in Fusion, but the conversion to tangent space cannot.


Projection Mapping

Projection is a technique for texturing objects using a camera or projector node. This can be useful for texturing objects with multiple layers, applying a texture across multiple separate objects, projecting background shots from the camera's viewpoint, image-based rendering techniques, and much more.

There are three ways to do projection mapping in Fusion.

Using the Projector/Camera Tool to Project Light


When lighting is enabled, a Camera 3D or Projector 3D can act as a light with all the lighting features. When Camera Projection is enabled or you use a projector, you can choose whether the projection behaves like a spotlight or an ambient light; however, alpha channels cannot be projected. Overlapping projections add together like any other light node. An internal clipping plane (at around 0.01 distance from camera) limits how close the projector or camera can get to the receivers of the projection.



Camera node used for a projection map

Project a Texture onto a Catcher Material

If you do not want to work with light sources, you can use the projector or camera as a texture projector. To work without lighting, a catcher is required in order to receive the texture and apply it to a material. Only objects using this material will receive the projection. This offers some advantages, like the projection of alpha channels, and texturing other channels like specular color or roughness. If the software renderer is used, overlapping projections can be combined in various ways (mean, median, blend, and so on) via the Catcher node. When using the OpenGL renderer, one catcher per projector is used, and the results can be combined using another material. Similar to the Light Projection technique, an internal clipping plane (at around 0.01 distance from camera) limits how close the projector/camera can get to the projection receivers.



Camera projection used with a Catcher node (example from an older version of Fusion)


Project Using the UVMap Node

This mode requires a camera and a UVMap3D node downstream of the objects to which the texture is being projected. In the Inspector, when the UVMap Map mode is set to Camera, it gathers the information from the camera and creates new UVs for the input objects, which are used for texturing. Because the UVs are stored in the vertices of the mesh, the object must be tessellated sufficiently.

Textures are assigned to the object like any other texturing technique. The UVs can be locked to the vertices at a chosen frame using the Ref Time slider. This locking only works as long as vertices are not created, destroyed, or reordered (e.g., projection locking will not work on particles because they get created/destroyed, nor will they work on a Cube3D when it's subdivision level slider animated).


TIP: Projected textures can be allowed to slide across an object. If the object moves relative to the Projector 3D, or alternatively, by grouping the two together with a Merge3D, they can be moved as one and the texture will remain locked to the object.

In the following section of a much larger composition, an image (the Loader1 node) is projected into 3D space by mapping it onto five planes (Shape3D nodes renamed ground, LeftWall, RightWall, Building, and Background), which are positioned as necessary within a Merge3D node to apply reflections onto a 3D car to be composited into that scene.



Excerpt of a composition that's projecting an image of a street scene into 3D space

The output of the Merge3D node used to assemble those planes into a scene is then fed to a UV Map node, which in conjunction with a Camera3D node correctly projects all of these planes into 3D space so they appear as they would through that camera in the scene. Prior to this UVMap projection, you can see the planes arranged in space at left, where each plane has the scene texture mapped to it. At right is the image after the UVMap projection, where you can see that the scene once again looks “normal,” with the exception of a car-shaped hole introduced to the scene.



Five planes positioning a street scene in 3D space in preparation for UV Projection (left), and the UV Map node being used to project these planes so they appear as through a camera in the scene (right)

However, this is now a 3D scene, ready for a digital car to be placed within it, receiving reflections and lighting and casting shadows into the scene as if it were there.



The new 3D scene casting reflections and lighting onto a 3D car, and receiving shadows caused by the car

Geometry

There are five nodes used for creating geometry in Fusion. These nodes can be used for a variety of purposes. For instance, the Image Plane 3D is primarily used to place image clips into a 3D scene, while the Shapes node can add additional building elements to a 3D set, and Text 3D can add three-dimensional motion graphics for title sequences and commercials. Although each node is covered in more detail in the “3D Nodes” chapter, a summary of the 3D creation nodes is provided below.

Cube 3D

The Cube 3D creates a cube with six inputs that allow mapping of different textures to each of the cube’s faces.

Image Plane 3D

The Image Plane 3D is the basic node used to place a 2D image into a 3D scene with an automatically scaled plane.

Shape 3D

This node includes several basic primitive shapes for assembling a 3D scene. It can create planes, cubes, spheres, cylinders, cones, and toruses.

Text 3D


The Text 3D is a 3D version of the Text+ node. This version supports beveling and extrusion but does not have support for the multi-layered shading model available from Text+.

Particles

When a pRender node is connected to a 3D view, it will export its particles into the 3D environment. The particles are then rendered using the Renderer3D instead of the Particle renderer. For more information, see Chapter 112, “Particle Nodes,” in the DaVinci Resolve Reference Manual or Chapter 52 in the Fusion Reference Manual.

Common Visibility Parameters

Visibility parameters are found in the Controls tab of most 3D geometry-producing nodes, exposed via a disclosure control. These parameters let you control object visibility in the viewers and in the final render.



A 3D geometry node’s visibility parameters

Visible

If the Visibility checkbox is not selected, the object will not be visible in a viewer, nor will it be rendered into the output image by a `Renderer3D`. A non-visible object does not cast shadows. This is usually enabled by default, so objects that you create are visible in both the viewers and final renders.

Unseen by Cameras

If the Unseen by Cameras checkbox is selected, the object will be visible in the viewers but invisible when viewing the scene through a camera, so the object will not be rendered into the output image by a `Renderer3D`. Shadows cast by an Unseen object will still be visible.

Cull Front Face/Back Face

Use these options to cull (exclude) rendering of certain polygons in the geometry. If Cull Back Face is selected, all polygons with normals pointing away from the view will not be rendered and will not cast shadows. If Cull Front Face is selected, all polygons with normals pointing away from the view will likewise be excluded. Selecting both checkboxes has the same effect as deselecting the Visible checkbox.

Ignore Transparent Pixels in Aux Channels


For any piece of geometry, the `Renderer3D` rejects transparent pixels in the auxiliary image channels. The reason this is the default is to prevent aux channels (e.g., normals, Z-channel, UVs) from filling in where there should be blank space or full transparency. For example, suppose in post you want to add some fog to the rendered image. If you had fully transparent geometry in the foreground affecting the Z-channel, you would get incorrect fog rendering. By deselecting this checkbox, the transparency will not be considered and all the aux channels will be filled for all the pixels. This could be useful if you wanted to replace texture on a 3D element that is fully transparent in certain areas with a texture that is transparent in different areas; it would be useful to have the whole object set aux channels (in particular UVs).

Adding FBX Models

The Filmbox FBX format is a scene interchange format that facilitates moving 3D scene information from one application to another. Fusion's FBX format support extends model import support to other 3D files such as Collada and OBJ.

Importing an FBX Scene

To import an entire FBX scene, you add an `FBXMesh3D` node to your node tree. After being prompted to choose a scene or object file, Fusion imports it to create a composition with the same lights, cameras, materials, and geometry found in an FBX file.



An imported model, via the `FBXMesh3D` node

FBX Scene Import Dialog

The FBX Mesh node is used to import mesh geometry from an FBX file. The first texture applied to a mesh will also be imported, if available.

Since different 3D applications use different units to measure their 3D scenes, the imported geometry may be enormous compared to the rest of the scene, because Fusion treats its scale of measurement as equal to its own system. For example, if your 3D application defaults to using millimeters as its scale, an object that was 100 millimeters in size will import as a massive 100 units.

You can use the Size slider in the FBX Mesh Inspector parameters to reduce the scale of such files to something that matches Fusion's 3D scene.

FBX Exporter


You can export a 3D scene from Fusion to other 3D packages using the FBX Exporter node. On render, it saves geometry, cameras lights, and animation into different file formats such as .dae or .fbx. The animation data can be included in one file, or it can be baked into sequential frames. Textures and materials cannot be exported.

Using Text3D

The Text3D node is probably the most ubiquitous node employed by motion graphics artists looking to create titles and graphics from Fusion. It's a powerful node filled with enough controls to create nearly any text effect you might need, all in three dimensions. This section seeks to get you started quickly with what the Text3D node is capable of. For more information, see Chapter 89, "3D Nodes," in the DaVinci Resolve Reference Manual or Chapter 29 in the Fusion Reference Manual.

Assembling Text Objects

Each Text3D node is a self-contained scene within which each character of text is an individual object. Because of this, the ideal way to combine numerous text objects that you might want to animate or style independently from one another is to connect as many Text3D objects as you want to be able to independently animate or style to one or more Merge3D nodes.




Merging multiple text objects to create an intricately styled scene

TIP: If you click the Text icon in the toolbar to create a Text3D node, and then you click it again while the Text3D node you just created is selected, a Merge3D node is automatically created and selected to connect the two. If you keep clicking the Text icon, more Text3D nodes will be added to the same selected Merge3D node.

Entering Text

When you select a Text3D node and open the Inspector, the Text tab shows a "Styled Text" text entry field at the very top into which you can type the text you want to appear onscreen. Below, a set of overall styling parameters are available to set the Font, Color, Size, Tracking, and so on. All styling you

do in this tab affects the entire set of text at once, which is why you need multiple text objects if you want differently styled words in the same scene.




The text entry and styling parameters in the Text tab

Near the bottom of the Text tab are the Extrusion parameters, available within a disclosure control.



The Extrusion parameters near the bottom of the Text tab

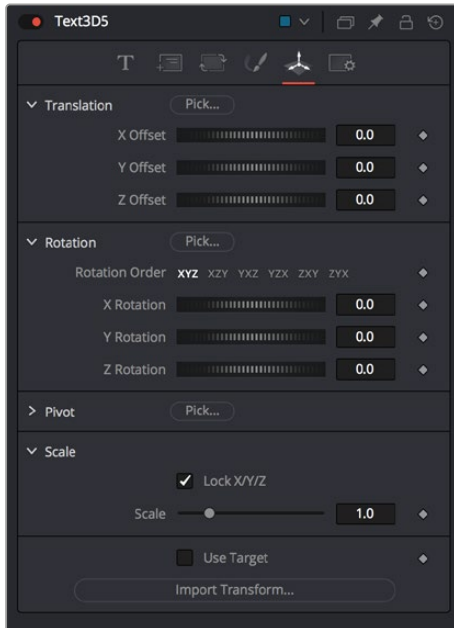
By default, all text created with the Text3D node is flat, but you can use the Extrusion Style, Extrusion Depth, and various Bevel parameters to give your text objects thickness.



Unextruded text (left), and Extruded text (right)

Positioning and Transforming Text

By default, every new Text3D node is positioned at 0, 0, 0, so when you add multiple Text3D nodes, they're all in the same place. Fortunately, every Text3D node has built-in transform controls in the Transform tab.



Text3D nodes also have Transform parameters built-in

Additionally, selecting a Text3D node exposes all the onscreen transform controls discussed elsewhere in this chapter. Using these controls, you can position and animate each text object independently.




Repositioned text objects to create a title sequence

Combining Text3D nodes using Merge3D nodes doesn't just create a scene; it also enables you to transform your text objects either singly or in groups:

- Selecting an individual Text3D node or piece of text in the viewer lets you move that one text object around by itself, independently of other objects in the scene.
- Selecting a Merge3D node exposes a transform control that affects all objects connected to that Merge3D node at once, letting you transform the entire scene.

Layout Parameters


The Layout tab presents parameters you can use to choose how text is drawn: on a straight line, a frame, a circle, or a custom spline path, along with contextual parameters that change depending on which layout you've selected (all of which can be animated).



Text using two different layouts

“Sub” Transforms

Another Transform tab (which the documentation has dubbed the “Sub” Transform tab) lets you apply a separate level of transform to either characters, words, or lines of text, which lets you create even more layout variations. For example, choosing to Transform by Words lets you change the spacing between words, rotate each word, and so on. You can apply simultaneous transforms to characters, words, and lines, so you can use all these capabilities at once if you really need to go for it. And, of course, all these parameters are animatable.



Transforming individual words in two different ways

Shading


The Shading tab lets you shade or texture a text object using standard Material controls.



Shading controls for text objects

Fog 3D and Soft Clipping

The Fog3D node helps to create atmospheric depth cues.



Split screen with and without fog

The Fog3D node works well with depth of field and antialiasing supported by the OpenGL renderer. Since it is not a post-processing node (like the VolumeFog node found in the Nodes > Position menu or Fog node in Nodes > Deep Pixel), it does not need additional channels like Position or Z-channel color. Furthermore, it supports transparent objects.

The SoftClip node uses the distance of a pixel from the viewpoint to affect opacity, allowing objects to gradually fade away when too close to the camera. This prevents objects from “popping off” should the camera pass through them. This is especially useful with particles that the camera may be passing through.


Geometry nodes such as the Shape3D node use a Matte Objects checkbox to enable masking out parts of the 3D scene. Effectively, everything that falls behind a matte object doesn't get rendered. However, matte objects can contribute information into the Z-channel and the Object ID channel, leaving all other channels at their default values. They do not remove or change any geometry; they can be thought of as a 3D garbage matte for the renderer.



Circle shape used as a Matte object to see the floor

Matte Object Parameters

Opening the Matte disclosure control reveals the Is Matte option, which when turned on enables two more options.



Matte parameters in the Shape3D node; enabling Is Matte reveals additional options

Is Matte

Located in the Controls tab for the geometry, this is the main checkbox for matte objects. When enabled, objects whose pixels fall behind the matte object's pixels in Z do not get rendered.

Opaque Alpha

When the Is Matte checkbox is enabled, the Opaque Alpha checkbox is displayed. Enabling this checkbox sets the alpha value of the matte object to 1. Otherwise the alpha, like the RGB, will be 0.


Infinite Z

When the Is Matte checkbox is enabled, the Infinite Z checkbox is displayed. Enabling this checkbox sets the value in the Z-channel to infinite. Otherwise, the mesh will contribute normally to the Z-channel.

Matte objects cannot be selected in the viewer unless you right-click in the viewer and choose 3D Options > Show Matte Objects in the contextual menu. However, it's always possible to select the matte object by selecting its node in the node tree.

Material and Object IDs

Most nodes in Fusion that support effect masking can use Object ID and Material ID auxiliary channels to generate a mask. The parameters used to accomplish this are found in the Common Controls tab of each node.



Material ID parameters in a Shape3D node's Inspector controls

The Material ID is a value assigned to identify what material is used on an object. The Object ID is roughly comparable to the Material ID, except it identifies objects and not materials.


Both the Object ID and Material ID are assigned automatically in numerical order, beginning with 1. It is possible to set the IDs to the same value for multiple objects or materials even if they are different. Override 3D offers an easy way to change the IDs for several objects. The Renderer will write the assigned values into the frame buffers during rendering, when the output channel options for these buffers are enabled. It is possible to use a value range from 0 to 65534. Empty pixels have an ID of 0, so although it is possible to assign a value of 0 manually to an object or material, it is not advisable because a value of 0 tells Fusion to set an unused ID when it renders.



Object ID for ground plane and object set to the same numeric value

World Position Pass

The World Position Pass, or WPP, is a render pass generated from 3D applications. Each pixel is assigned the XYZ position where the pixel was generated in the world coordinates. So if the face from which the pixel was derived in the scene sits at (0,0,0), the resulting pixel will have a Position value of (0,0,0). If we visualize this as RGB, the pixel will be black. If a face sits at (1,0,0) in the original scene, the resulting RGB pixel will be red. Due to the huge range of possible positions in a typical 3D scene, and 7/8 of those possible positions containing negative coordinates, the Position channel is always rendered in 32-bit float.



A World Position Pass rendering of a scene with its center at (0,0,0) The actual image is on the left


3D Scene Input

Nodes that utilize the World Position channel are located under the Position category. VolumeFog and Z to WorldPos require a camera input matching the camera that rendered the Position channels, which can either be a Camera3D or a 3D scene containing a camera. Just as in the Renderer3D, you can choose which camera to use if more than one are in the scene. The VolumeFog can render without a camera input from the Node Editor if the world space Camera Position inputs are set to the correct value. VolumeMask does not use a camera input. Nodes that support the World Position Pass, located under the Position category, offer a Scene input, which can be either a 3D Camera or a 3D scene containing a camera.

There are three Position nodes that can take advantage of World Position Pass data.

- Nodes > Position > Volume Fog
- Nodes > Position > Volume Mask
- Nodes > Position > Z to World
- The “Dark Box”

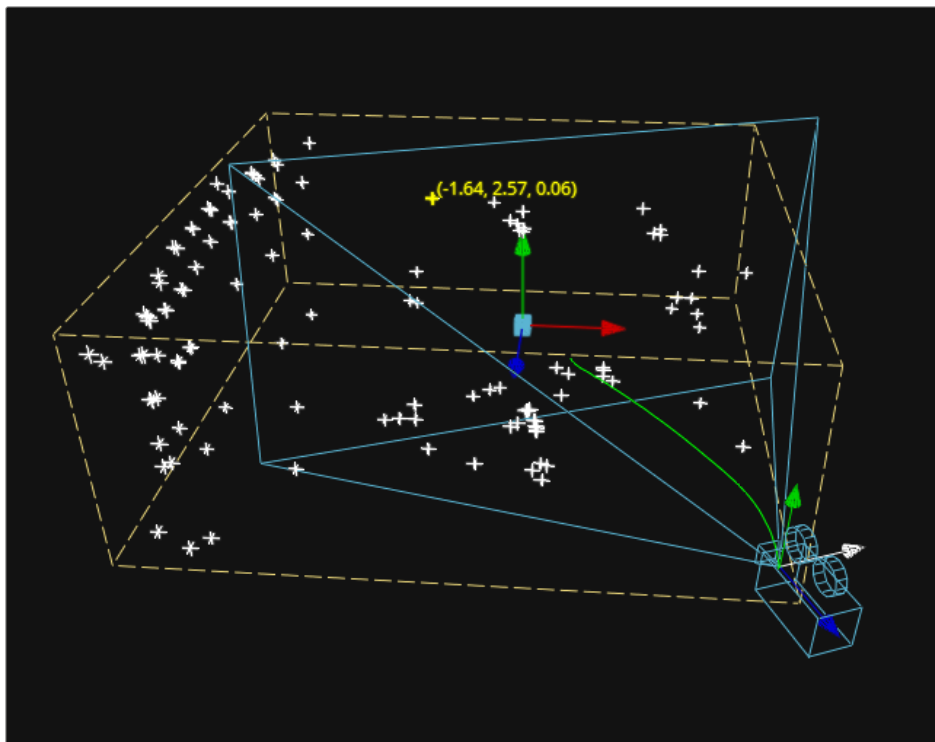
Empty regions of the render will have the Position channel incorrectly initialized to (0,0,0). To get the correct Position data, add a bounding sphere or box to your scene to create distant values and allow the Position nodes to render correctly.



Without a bounding mesh to generate Position values, the fog fills in the background incorrectly

Point Clouds

The Point Cloud node is designed to work with locator clouds generated from 3D tracking software. 3D camera tracking software, such as SynthEyes and PF Track, will often generate hundreds or even thousands of tracking points. Seeing these points in the scene and referencing their position in 3D and screen space is important to assist with lining up live action and CG, but bringing each point in as an individual Locator3D would impact performance dramatically and clutter the node tree.



Point cloud in the viewer

The Point Cloud node can import point clouds written into scene files from match moving or 3D scanning software.

To import a point cloud, do the following:

- 1 Add the PointCloud3D node to your composition.
- 2 Click the Import Point Cloud button in the Control panel.
- 3 Browse to the scene file and select a cloud to import from the scene.

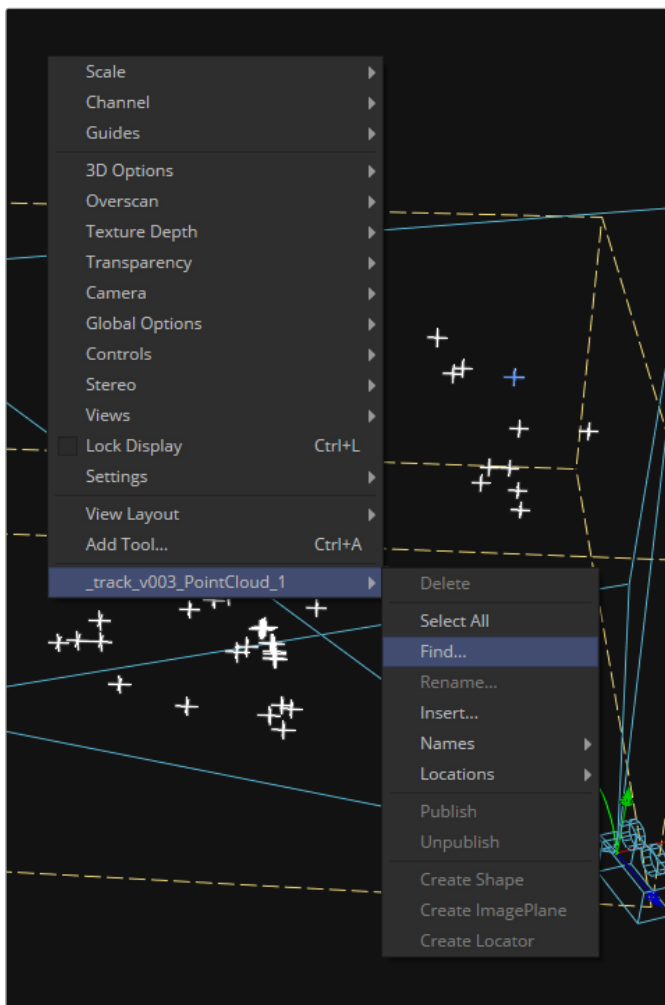
The entire point cloud is imported as one object, which is a significantly faster approach.

Finding, Naming, and Publishing Points

Many 3D trackers allow for the naming of individual tracking points, as well as setting tracking points on points of interest. The Point Cloud 3D will quickly find these points and publish them. A published point in the cloud can be used to drive the animation of other parameters.

To find a point in the point cloud, do the following:

- 1 Right-click anywhere within a viewer.
- 2 Choose Find from the Point Cloud's submenu in the contextual menu.
- 3 Type the name of the point and click OK.



Finding a point cloud using the viewer contextual menu

If a point that matches the name you entered is found, it will be selected in the point cloud and highlighted yellow.

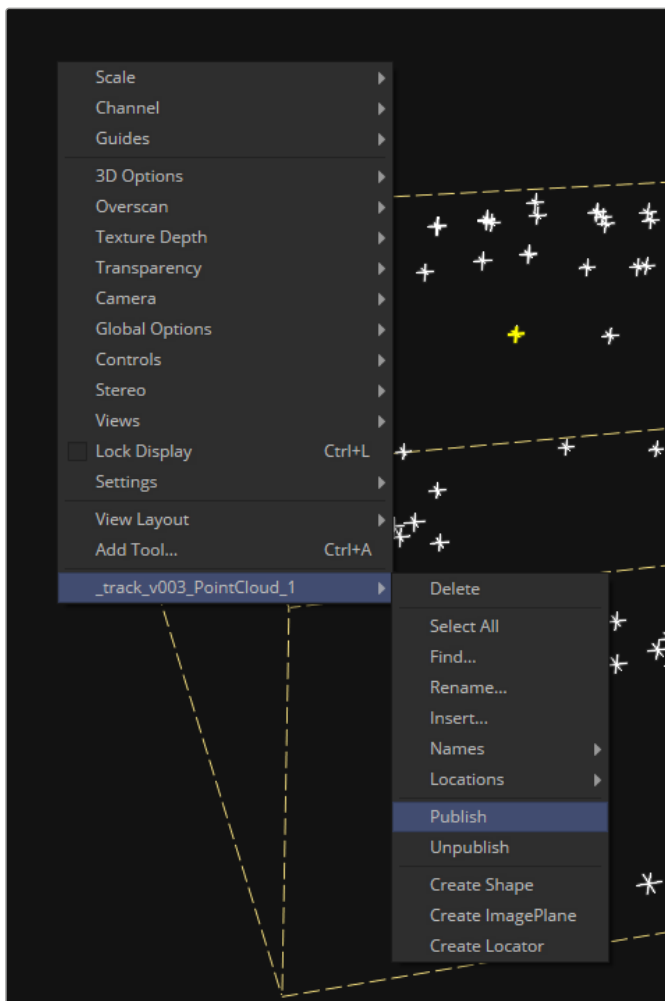
TIP: The Point Cloud Find function is a case-sensitive search. A point named “tracker15” will not be found if the search is for “Tracker15”.

Renaming a Point in the Cloud

You can use the Point Cloud contextual menu to rename a selected point. This works only for a single point. A group of points cannot be renamed.

Publishing a Point

If you want to use a point’s XYZ positions for connections to other controls in the scene, you can publish the point. This is useful for connecting objects to the motion of an individual tracker. To publish a point, right-click it and choose Publish from the contextual menu.



Publishing a point using the viewer contextual menu

Chapter 86

3D Camera Tracking

This chapter presents an overview of using the Camera Tracker node and the workflow it involves. Camera tracking is used to create a virtual camera in Fusion’s 3D environment based on the movement of a live-action camera in a clip. You can then use the virtual camera to composite 3D models, text, or 2D images into a live-action clip that has a moving camera.

For more information on other types of tracking in Fusion, see Chapter 82, “Using the Tracker Node,” in the DaVinci Resolve Reference Manual or Chapter 22 in the Fusion Reference Manual.

Contents

Introduction to Tracking	1566	Matching the Live-Action Camera	1573
What Is 3D Camera Tracking?	1566	Running the Solver	1573
How Camera Tracking Works	1566	How Do You Know When to Stop?	1574
The Camera Tracking Workflow	1567	Using Seed Frames	1575
Clips That Don’t Work Well for Camera Tracking	1567	Cleaning Up Camera Solves	1575
Outputting from the Camera Tracker	1568	Exporting a 3D Scene for Efficiency	1578
2D View	1568	Unalign the 3D Scene Transforms	1578
3D View	1569	Setting the Ground Plane	1578
Auto-Tracking in the Camera Tracker	1570	Setting the Origin	1579
Increasing Auto-Generated Tracking Points	1570	Setting the Scale	1579
Masking Out Objects	1571	Realign the Scene	1579
		Viewing the Exported Results	1580

Introduction to Tracking

Tracking is one of the most useful and essential techniques available to a compositor. It can roughly be defined as the creation of a motion path from analyzing a specific area in a clip over time. Fusion provides a variety of different tracking nodes that let you analyze different kinds of motion.

Each tracker type has its own chapter in this manual. This chapter covers the tracking techniques with the Camera Tracker node.

What Is 3D Camera Tracking?

Camera tracking is used for match moving, and it's a vital link between 2D scenes and 3D scenes, allowing compositors to integrate 3D CGI elements into live-action clips. The Camera Tracker node calculates the path of a live-action camera and generates a virtual camera in 3D space. This virtual camera is intended to be identical to the actual camera that shot the scene, not only in terms of motion but in matching the lens focal length as well. The calculated position and movement of the virtual camera is central to realistically compositing 3D elements with live action.



An example of 3D elements integrated in a live-action scene

How Camera Tracking Works

Camera tracking begins by tracking the movement of fixed features from one frame to the next. To put it another way, camera tracking algorithms follow features that are “nailed to the set.” Objects in the scene that move independently of the camera movement in the shot, such as cars driving or people walking, cause poor tracks, so masks can be used to restrict the features that are tracked in order to improve the results. Additionally, it is helpful to provide specific camera metadata, such as the sensor size and the focal length of the lens. This information guides the scene reconstruction calculation, called a solver, toward generating a more accurate virtual camera.

The Camera Tracker's purpose is to create a 3D animated camera and point cloud of the scene. A point cloud is a large group of points generated by the solver that roughly recreates the 3D positions of the tracked features in a scene. The point cloud can then be used as a guide when integrating other 2D or 3D elements alongside live-action features.


The Camera Tracking Workflow

Camera tracking has two main phases:

- 1 *Tracking*, which is the analysis of a scene.
- 2 *Solving*, which calculates the virtual 3D scene.

Once you complete these steps, an animated camera and point cloud are exported from the Inspector into a 3D composite. The Camera Tracker encompasses this complete workflow within one tool. Five tabs at the top of the Inspector are roughly laid out in the order in which you'll use them. These five tabs are:

- **Track:** Used to track a clip.
- **Camera:** Configures the basic Camera parameters.
- **Solve:** Calculates the 3D placement of the 2D tracking points and reconstructs the camera.
- **Export:** Generates a Camera 3D node, a Point Cloud node, and a 3D scene in the node tree.
- **Options:** Used to customize the look of the onscreen overlays.



The Camera Tracker tab

Clips That Don't Work Well for Camera Tracking

Even though the Camera Tracker is somewhat automatic, it sometimes needs your help. If you can identify potential issues before you even track or solve the shot, you can save yourself much time. Certain types of clips will cause more significant problems for camera tracking than others. Some are fixable, while for others you just have to admit defeat and figure out another solution. Here is a list of the types of shots to look out for, as they can be big headaches for camera tracking:

- **Lack of depth:** Camera tracking requires parallax in a clip in order to work. You must be able to identify objects further away and objects that are nearer as the camera moves. If everything is at the same distance from the camera, there is no way to calculate depth. In this case, it's better to skip the Camera Tracker node and find another solution.
- **Locked-off shots:** If the camera does not move, there is no way to calculate which objects are closer and which are nearer. Again, don't spend too much time in this situation; it is better to skip the Camera Tracker node and find another solution.

- **Tripod pans:** Similar to a locked-off shot, there is no way to calculate which objects are closer and which are nearer from a pan that remains centered on a locked-off tripod. Skip the Camera Tracker node and find another solution.
- **No detail:** Clips like green screens without tracking markers lack enough detail to track. If you are lucky enough to be involved in the shooting of these types of shots, including tracker markers makes it much easier to get a good track. Without detail, camera tracking will fail and you will need to find a more manual solution.
- **Motion blur:** Fast camera motion or slow shutter speeds can introduce motion blur, which will make it difficult to find patterns to track. It's worth trying shots like these to see if there are enough details to get a good solve, but know when give up and turn to another solution.
- **Rolling shutter:** CMOS-based cameras sometimes introduce distortion due to the shutter capturing different lines at slightly different times. This distortion can create significant problems for camera tracking. Sometimes it is possible to create motion vectors with the Optical Flow node to create new in-between frames without the wobble distortion of the rolling shutter. Then you can use the corrected image to connect to the Camera Tracker.
- **Parallax issues:** When objects at different distances in a shot overlap in the frame, the overlapping area can be misinterpreted as a corner. Having a tracker assigned to an overlapping angle like this will cause errors as the parallax starts to shift and the overlapping area slides. This can be solved in Fusion by removing that tracker before running the solver.
- **Moving objects:** It's difficult to capture a shot where objects in the clip do not move. People, cars, animals, or other object may move in and out of a shot. These objects move independent of the camera movement and must be eliminated or they will cause solving errors. You can fix these issues by masking out objects that are not “nailed to the set.” The masks are then connected to the Track Mask input on the Camera Tracker node.

TIP: Some shots that cannot be tracked using Fusion's Camera Tracker can be performed in dedicated 3D camera-tracking software like 3D Equalizer and PF Track. Camera tracking data from these applications can then be imported in the Camera3D node in Fusion.

Outputting from the Camera Tracker


Unlike most Fusion nodes, the Camera Tracker node has two outputs:

- The primary output is a 2D view used when you are setting up the Track, refining the camera, and performing your initial solve.
- There is also a 3D output used after your initial solve for viewing the camera path and point cloud in 3D space. This view can be helpful when you are refining tracks to increase the accuracy of the solve and aligning your ground plane. It can be used simultaneously with the 2D output in side-by-side views.

Note that the selection of tracks in the 2D view and their corresponding locators (in the point cloud) in the 3D view are synchronized. There are also viewer menus available in both the 2D and 3D views to give quick control of the functionality of this tool.

2D View


The 2D view is the primary display for the node. Viewing the node displays the image being tracked as well as overlay tracker markers and their motion paths. A dedicated toolbar gives you access to the common features used to track and solve a clip.



The Camera Tracker 2D output with toolbar and auto-track points

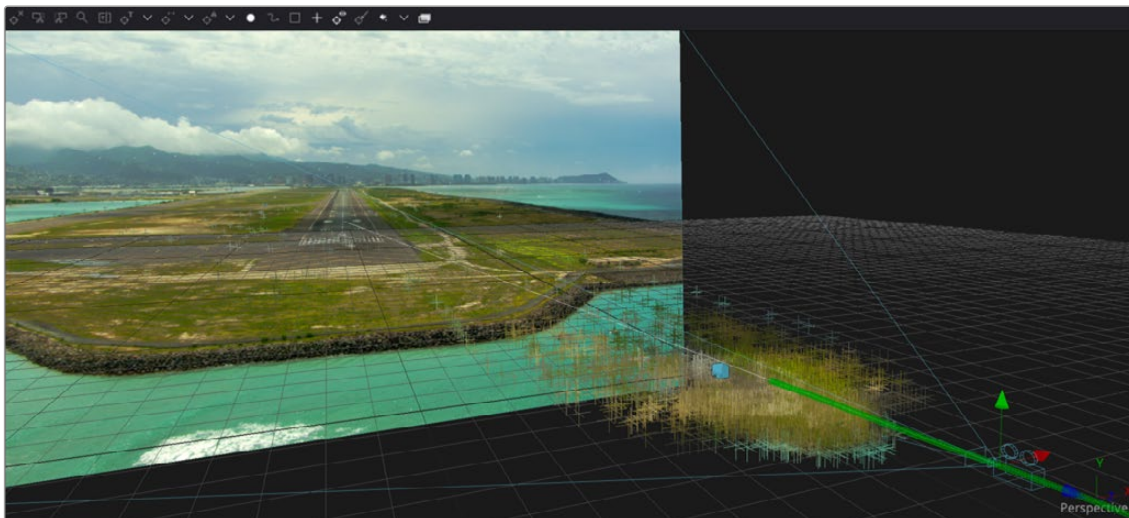
3D View

The second output of the Camera Tracker node displays a 3D scene. To view this, connect this 3D output to a 3D Transform or Merge 3D node and view that tool.



After solving, connecting the second output of the Camera Tracker node to a Merge 3D displays the point cloud in 3D

After an initial solve, the 3D output displays the point cloud and the camera, along with the image connected to it. Selecting points displays the Camera Tracker toolbar above the viewer, which gives control of various functions, such as renaming, deleting, and changing the colors of points in the point cloud.



3D output of a point cloud and a solved camera path

Auto-Tracking in the Camera Tracker

Tracking is the term used to describe the task of observing or analyzing a sequence of frames (or clip). The Camera Tracker node must take into account the movement of the source footage before it can determine the location and movement of the virtual camera. The Camera Tracker tool automatically searches for features that are high-contrast patterns within the clip and assigns trackers to those features. Having a wide distribution of tracking points across the frame and having points with long durations results in the best track.

Increasing Auto-Generated Tracking Points

Unlike the Tracker node, setting tracking points is entirely automatic in the Camera Tracker, but the Detection Threshold and Minimum Feature Separation sliders let you adjust the criteria by which tracking points are found. Lowering these parameters lets you increase the number of tracking points that will be found. This can be useful if the scene has too few points, which will prevent the solver from generating an accurate camera and point cloud. However, make these adjustments with care, since adding too many points may generate redundant trackers that slow down the entire process with minimal benefit.

Previewing Tracking Points

You can view the tracking points currently generated on the clip by viewing the Camera Tracker node and turning on the Preview AutoTrack Locations checkbox. This causes green tracking points to be displayed in the viewer as you play through the clip. Using this preview, you can decide whether you need to make adjustments to the Detection Threshold or Minimum Feature Separation to increase or decrease the number of tracking points that are found automatically.



Green tracker marks are added automatically to the features in an image

Bi-Directional Tracking

When performing a track, you can enable the Bidirectional Tracking checkbox, which first tracks forward from the start of the clip, and then tracks a second pass in reverse. This two-pass approach can potentially extend the duration of any given point by re-analyzing points initially identified in the forward pass. There is very little reason not to have this enabled unless you are very short on time. Bidirectional tracking takes longer, but it's usually worth it, and the process is reasonably quick considering the benefit.

Choosing a Tracking Algorithm

There are three available choices for the algorithm to use when tracking. The three options in the New Track Defaults section of the Inspector include:

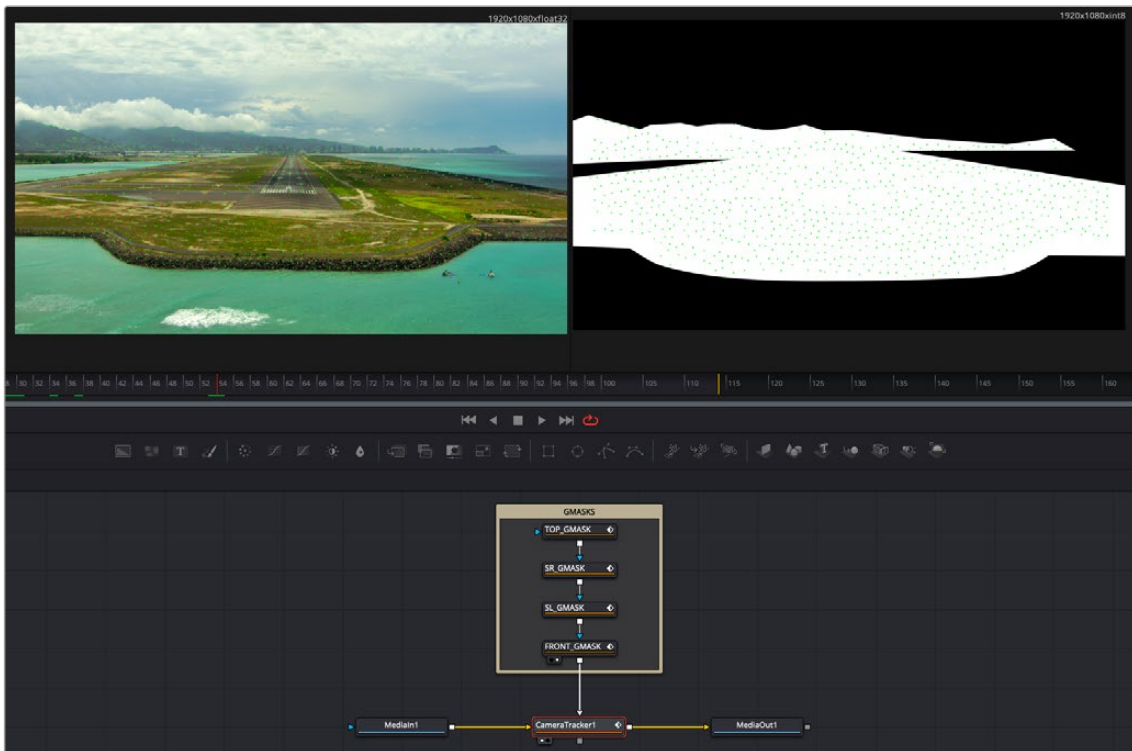
- **Optical Flow:** Usually your best choice, unless you have a great deal of criss-crossing objects in a clip.
- **Tracker:** A good second choice when Optical Flow can't be used due to motion estimation errors like criss-crossing objects.
- **Planar:** Mostly used in simpler clips, where the majority of the image consists of planar surfaces such as the facades of buildings.

Masking Out Objects

When tracking a clip, the Camera Tracker automatically generates trackers on feature details. However, not all features that stand out in a clip are appropriate for camera tracking. You only want to track features that are “nailed to the set.” In other words, objects that move independently of the camera motion, like moving cars and people, cause inaccuracies when camera tracking. You must eliminate these types of objects from the analysis.

The primary way of avoiding these problem areas is by masking. You connect a mask to the Camera Tracker node's Track Mask input to identify areas of a scene that the Camera Tracker can analyze. For example, if you have a clip of an airport runway along a shoreline, the waves of the water and moving clouds in the sky must be masked since they move independently of the camera.

When creating a mask, the fixed areas of the image to be analyzed for tracking should be encompassed in the white portion of the mask. All moving objects that need to be ignored should be encompassed in the black portion. The mask should then be attached to the Camera Tracker Track Mask input.



Masks used to omit the moving clouds and waves from being tracked by the Camera Tracker

By doing this, the tracker ignores the waves of the water and moving clouds. Unlike drawing a mask for an effect, the mask in this case does not have to be perfect. You are just trying to identify the rough area to occlude from the tracking analysis.



The original image to be tracked (left), and the occlusion mask of the clouds and water (right)

TIP: If there's a lot of motion in a shot, you can use the Tracker or Planar Tracker nodes to make your occlusion mask follow the area you want to track. Just remember that, after using the PlanarTracker or PlanarTransform node to transform your mask, you need to use a Bitmap node to turn it back into a mask that can be connected to the Camera Tracker node's Track Mask input.

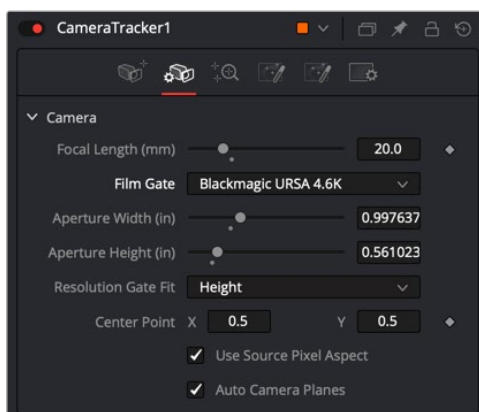
Matching the Live-Action Camera

Once you have completed tracking, the next stage of this workflow requires the controls in the Camera tab. This is where you define the actual camera used on set, primarily the film gate size and focal length. This information should have been logged on the set to make available for post-production. When using camera-original media, you can sometimes locate this information in the file metadata.

To locate camera metadata, do the following:

- If you are using DaVinci Resolve, select the MediaIn node with the camera clip, open the Metadata Editor, and view the Camera metadata preset.
- If you are using Fusion Studio, display the metadata subview from the viewer toolbar.

If the actual values are not known, try a best guess. The solver attempts to find a camera near these parameters, and it helps the solver by giving parameters as close to the live action as possible. The more accurate the information you provide, the more accurate the solver calculation. At a minimum, try to at least choose the correct camera model from the Film Gate menu. If the film gate is incorrect, the chances that the Camera Tracker correctly calculates the lens focal length become very low.



The Camera tab in the Camera Tracker tool

Unlike the Track and Solve tabs, the Camera tab does not include a button at the top of the Inspector that executes the process. There is no process to perform on the Camera tab once you configure the camera settings. After you set the camera settings to match the live-action camera, you move to the Solve tab.


Running the Solver

The next step in this workflow involves the controls found in the Solve tab. Solving is a compute-intensive process in which the Camera Tracker analyzes the currently existing tracks to create a 3D scene. It generates a virtual camera that matches the live action and a point cloud consisting of 3D locators that recreate the tracked features in 3D space. The analysis is based on parallax in the frame, which is the perception that features closer to the camera move quicker than features further away. This is much like when you look out the side window of a car and can see objects in the distance move more slowly than items near the roadside.

The trackers found in the Track phase of this workflow have a great deal to do with the success or failure of the solver, making it critical to deliver the best set of tracking points from the very start. Although the masking you create to occlude objects from being tracked helps to omit problematic tracking points, you almost always need to further filter and delete poor quality tracks in the Solver tab. That's why, from a user's point of view, solving should be thought of as an iterative process.

To solve a camera's motion:

- 1 Click the Solve button to run the solver.
- 2 Filter out and delete poor tracks.
- 3 Rerun the solver.



The Solver tab after it has run and produced an average solve error of 0.4367 pixels

How Do You Know When to Stop?

At the end of the solve process, an Average Solve Error (sometimes called a reprojection error) appears at the top of the Inspector. This is the crucial value that tells you how well the calculation has gone. A good Average Solve Error for HD content is below 1.0.

You can interpret a value of 1.0 as a pixel offset; at any given time, the track could be offset by 1 pixel. The higher the resolution, the lower the solve error should be. If you are working with 4K material, your goal should be to achieve a solve error below 0.5.

Tips for Solving Camera Motion

When solving camera movement, it's important to provide accurate live-action camera information, such as focal length and film gate size, which can significantly improve the accuracy of the camera solve. For example, if the provided focal length is too far away from the correct physical value, the solver can fail to converge, resulting in a useless solution.

Additionally, for the solver to accurately triangulate and reconstruct the camera and point cloud, it is important to have:

- A good balance of tracks across objects at different depths, with not too many tracks in the distant background or sky (these do not provide any additional perspective information to the solver).
- Tracks distributed evenly over the image and not highly clustered on a few objects or one side of the image.
- The track starts and ends staggered over time, with not too many tracks ending on the same frame.

Using Seed Frames

The solver works by first constructing a partial solution between two seed frames. These seed frames are selected automatically. However, automatic selection adds time to the process. The time is reported in the solve summary at the top of the Inspector once you run the solver. You can select your own seed frames to speed the process and potentially get a better solve on trickier clips. The solver uses these seed frames to create an intermediate solution and then extends this forward and backward for the duration of the clip.

Selecting appropriate seed frames is not necessarily recommended unless you have some experience with camera tracking. Keeping the default Auto Select Seed Frames checkbox enabled in the Solve Options section of the Solver tab selects the best frames in most cases. However, you can disable the checkbox and use the Seed Frame 1 and Seed Frame 2 slider to select frames you believe achieve better results.

When choosing seed frames, it is important to satisfy two main characteristics:

- Seed frames should have lots of tracks in common.
- Seed frames should be far apart in perspective.

Sometimes There's Nothing You Can Do

Some shots that do not have enough camera motion to triangulate feature locations cannot be reconstructed with any useful accuracy. Ensuring that a shot is camera-trackable begins on set, with proper usage of track markers, and by ensuring that camera moves have enough perspective shifts for the solver to glean useful data from.

Cleaning Up Camera Solves

Sometimes the first solve will be good enough. Other times, it may take a few hours of cleaning up tracks to get a good solve, and sometimes it is impossible. With experience, one gets a feel for which tracks should be deleted and which should be kept, and which shots will be easy, difficult, or impossible to solve.

Be aware that deleting too many tracks can cause the Average Solve Error to increase, as the solver has too little information to work with. In particular, if there are fewer than eight tracks on any frame, mathematically there is not enough information to solve the clip. However, it is strongly recommended to use a lot more than eight tracks to get a robust and accurate solve.

IMPORTANT

If you are not familiar with camera tracking, it may be tempting to try to directly edit the resulting 3D splines in the Spline Editor in order to improve a solved camera's motion path. This option should be used as an absolute last resort. It's preferable, instead, to modify the 2D tracks being fed into the solver.


How to Judge Track Accuracy

The automatic color coding of tracking markers makes deleting false or poor tracks easier. After the solver runs, each tracker is assigned a solve error color that indicates which 3D locators match their 2D tracking points well, and which match up poorly.

- **Green:** Good. Tracked very well.
- **Yellow:** Moderate confidence. Usually an acceptable track.


- **Orange:** Low Confidence. May be acceptable in some situations.
- **Red:** No Confidence. The tracks have not solved well.

Hovering the pointer over any tracking point displays a large metadata tooltip that includes the solve error for the point. For a more visual representation of the accuracy, you can enable the display of 3D locators in the viewer by clicking the Reprojection Locators button in the viewer toolbar.



After a solve, the Camera Tracker toolbar can display Reprojection locators

When the tracking points are converted into a point cloud by the solver, it creates 3D reprojection locators for each tracking point. These Reprojection locators appear as small X marks near the corresponding tracking point. The more the two objects overlap, the lower the solve error.



Reprojection locators displayed with tracking points, and tooltip

The goal when filtering the trackers is to remove all red tracker marks and keep all the green marks. Whether you decide to keep both the yellow and orange or just the yellow is more a question of how many marks you have in the clip. You produce a better solve if you retain only the yellow marks; however, if you do not have enough marks to calculate the 3D scene, you will have to keep some of the better orange marks as well.


Tips for What to Keep and What to Delete

Understanding what false tracks look like, and then manually cleaning the track data to reduce it to a precise set of clear tracks, will result in a more accurate solve. When cleaning up any track—particularly yellow and orange color coded tracks—keep the following in mind:

- Keep all tracks with motion that's completely determined by the motion of the live-action camera.
- Delete tracks on moving objects or people and tracks that have parallax issues.
- Delete tracks that are reflected in windows or water.
- Delete tracks of highlights that move over a surface.
- Delete tracks that do not do a good job of following a feature.
- Delete tracks that follow false corners created by the superposition of foreground and background layers.
- Consider deleting tracks that correspond to locators that the solver has reconstructed at an incorrect Z-depth.

Deleting Tracks

You can manually delete tracks in the viewer or use filters to select groups of tracks. When deleting tracks in the viewer, it is best to modify the viewer just a bit to see the tracks more clearly. From the Camera Tracker toolbar above the viewer, clicking the Track Trails button hides the trails of the tracking points. This cleans up the viewer to show points only, making it easier to make selections. At the right end of the toolbar, clicking the Darken Image button slightly darkens the image, again making the points stand out a bit more in the viewer.



Hiding trails and darkening the viewer can make it easier to see and select poor tracks

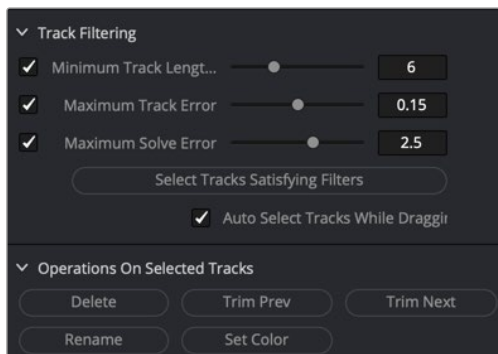
To begin deleting poor-quality tracks, you can drag a selection box around a group of tracks you want to remove and then either click the Delete Tracks button in the Camera Tracker toolbar or press Command-Delete.

You can hold down the Command key to select discontinuous tracking marks that are not near each other. If you accidentally select tracks you want to keep, continue holding the Command key and drag over the selected tracks to deselect them.

When deleting tracks, take note of the current Average Solve Error at the top of the Inspector and then rerun the solver. It is better to delete small groups of tracks and then rerun the solver than to delete one or two large sections. As mentioned previously, deleting too many tracks can have adverse effects and increase the Average Solve Error.

Using Filters to Delete Problem Tracks

The Solve tab includes filters that can be used to select groups of similar tracks by track length, track error, and solve error. These can be used to quickly select and delete poorly performing tracks that may be misleading to the resulting camera, leaving a concise list of accurate tracks.



Tracks can be selected using filters and deleted using the Operations On Selected Tracks buttons

For instance, it is generally best to run the solver using tracks with longer durations. Since shorter tracks tend to be less accurate when calculating the camera, you can remove them using the Filter section in the Inspector.

Increasing the Minimum Track Length parameter sets a threshold that each tracker must meet. Tracks falling below the threshold appear red. You can then click the Select Tracks Satisfying Filters button to select the shorter tracks and click Delete from the Options section in the Inspector.

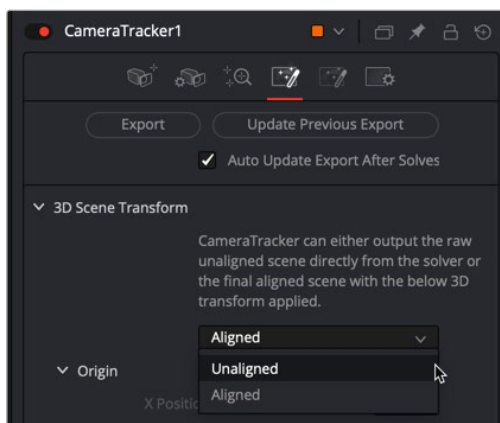
Exporting a 3D Scene for Efficiency

The Camera Tracker saves all its 2D tracks into the composition, sometimes resulting in a rather large file on disk. If you are dealing with a large clip with many 2D tracks over a long duration, the saved composition can reach over a gigabyte in size. Using a Camera Tracker node in a composition can make it cumbersome to load and operate. While it is possible to use the Camera Tracker node directly to composite via the 3D output, you'll achieve better performance by exporting. Once the quality of the solve is satisfactory, the Export tab can generate a "low memory" alternative by producing individual Camera 3D, Point Cloud, Ground Plane, and 3D Renderer nodes.

Before you can export the 3D scene, you must provide a bit more information about it. You'll do this using controls found in the Export tab. Cameras do not include tiltmeters, so clips do not contain metadata that indicates how the camera is tilted or oriented. This is critical information when recreating the virtual camera. It is also useful to determine the location for the center of this 3D scene. The Export tab provides various translation, rotation, and scale controls to set these options.

Unalign the 3D Scene Transforms

By default, the Export tab is set to Aligned in the 3D Scene Transform section. The Aligned setting locks the orientation and scale of the 3D scene to prevent accidentally changing it. So, before you can set the ground plane and origin location, you must change the Camera Tracker to be unaligned using the 3D Scene Transform menu in the Export tab. After you have gone through the Export settings and configured them how you want, you must set the menu back to Aligned before exporting.



Set the 3D Scene Transform menu to Unaligned before setting the ground plane

Setting the Ground Plane

The Camera Tracker has no idea if the camera is on its side or tilted in some way. So, it is up to you to indicate where the ground plane is in a clip. After choosing Unaligned from the 3D Scene Transform menu, you can begin identifying the ground plane.



Drag a selection box around marks that represent the ground

To set the ground plane, do the following:

- 1 Move to a frame with lots of green 3D locators where you can see a large part of the ground.
- 2 In the viewer, drag a selection box around the marks located on the ground in the clip.
- 3 In the Inspector Orientation section, click the Set from Selection button.

TIP: In some cases, the clip you are tracking may not have the ground in the frame. If necessary, you can set the Selection menu to XY, which indicates you are selecting points on a wall.

Setting the Origin

When it comes time to add and position new objects into your 3D scene, you can make it easier by setting the origin or center location.

To set the origin of the 3D scene, do the following:

- 1 Move to a frame that clearly shows the area you want to select as the center of the scene.
- 2 In the viewer, either select a single point or drag a selection box around a few marks located where you want to position the center of the 3D scene.
- 3 In the inspector Origin section, click the Set from Selection button.

Setting the Scale

The Camera Tracker has no idea of the size of the 3D scene, so the scale parameter is used to scale the scene output. This makes it possible to match the scale of two or more clips.


Realign the Scene

Before exporting the scene from the Camera Tracker, you must set the 3D scene Transform menu back to Aligned. Now you're ready to export.

Viewing the Exported Results


Clicking the Export button at the top of the Inspector creates a functional 3D scene with five new nodes automatically added to the node tree.

- Camera 3D
- Point Cloud
- Ground Plane
- Merge 3D
- Camera Tracker Renderer (3D Renderer)




Five nodes created as a result of exporting from the Camera Tracker

To work with the 3D scene, you can select the Merge 3D and load it into one of the viewers, and then select the Camera Tracker Renderer and load that into a second viewer.



Viewing the Merge 3D shows the point cloud, ground plane, and camera

When the Merge 3D is selected, a toolbar above the viewer can add 3D test geometry like an image plane or cube to verify the precision of the 3D scene and camera. You can then connect actual 3D elements into the Merge 3D as you would any manually created 3D scene. The point cloud can help align and guide the placement of objects, and the CameraTracker Renderer is a Renderer 3D node with all the same controls.



Use the point cloud to accurately place different elements into a 3D scene

At this point, there is no need for the Camera Tracker node unless you find that you need to rerun the solver. Otherwise, you can save some memory by deleting the Camera Tracker node.

Chapter 87

Particle Systems


This chapter is designed to give you a brief introduction to the creation of fully 3D particle systems, one of Fusion’s most powerful features. Once you understand these basics, for more information on each particle system node that’s available, see Chapter 112, “Particle Nodes,” in the DaVinci Resolve Reference Manual or Chapter 52 in the Fusion Reference Manual.

Contents

Introduction to Particle Systems	1583
Anatomy of a Simple Particle System	1584
Particle System Distribution	1586
Particle Nodes Explained by Type	1587
Emitters	1587
Forces	1588
Compositing	1588
Rendering	1588
Example Particle Systems	1589


Introduction to Particle Systems

Particle systems are computer simulations that use customizable rules to automatically generate and animate large numbers of elements to simulate smoke, dust, fire, leaves, sparks, or any other animated system of shapes. As Fusion is a full-featured 3D compositing environment, particle systems can be created in 2D or 3D, which makes them incredibly flexible and capable of producing all kinds of visual effects or abstract animated content for use in motion graphics.




A 3D particle system, also created entirely within Fusion

The three most fundamental nodes required for creating particle systems are found on the toolbar. As with the 3D nodes to the right, these are arranged, from left to right, in the order in which they must be connected to work, so even if you can't remember how to hook up a simple particle system, all you need to do is click the three particle system nodes from left to right to create a functional particle system.



The pEmitter, pMerge, and pRender particle system nodes available from the toolbar

However, these three nodes are only the tip of the iceberg. Opening the Particle category in the Effects Library reveals many, many particle nodes designed to work together to create increasingly complex particle interactions.




A sample of the nodes available in the Particles bin of the Effects Library

All particle nodes begin with the letter “p,” and they’re designed to work together to produce sophisticated effects from relatively simple operations and settings. The next section shows different ways particle nodes can be connected to produce different effects.


Anatomy of a Simple Particle System

The simplest particle system you can create is a pEmitter node connected to a pRender node. The pEmitter node includes the core controls for creating various kinds of particles in different ways, while the pRender node is required to render a 2D or 3D result that can be composited with other scenes within your composition.




The minimum node tree required to create a simple particle system

If your needs are more complicated, you can combine two or more pEmitter nodes using a pMerge node (the particle system version of a Merge node), to create compound particle systems where multiple types of particles combine with one another to create a result.



Compositing two pEmitter nodes to create a compound particle system, combining two kinds of particles together


If you’re trying to create particle systems with more natural effects, you can add “forces” to each emitter. These forces are essentially physics or behavioral simulations that automatically cause the particles affected by them to be animated with different kinds of motion, or to be otherwise affected by different objects within scenes.



Customizing the effect of pEmitter nodes using different forces to add complexity to the particle animation

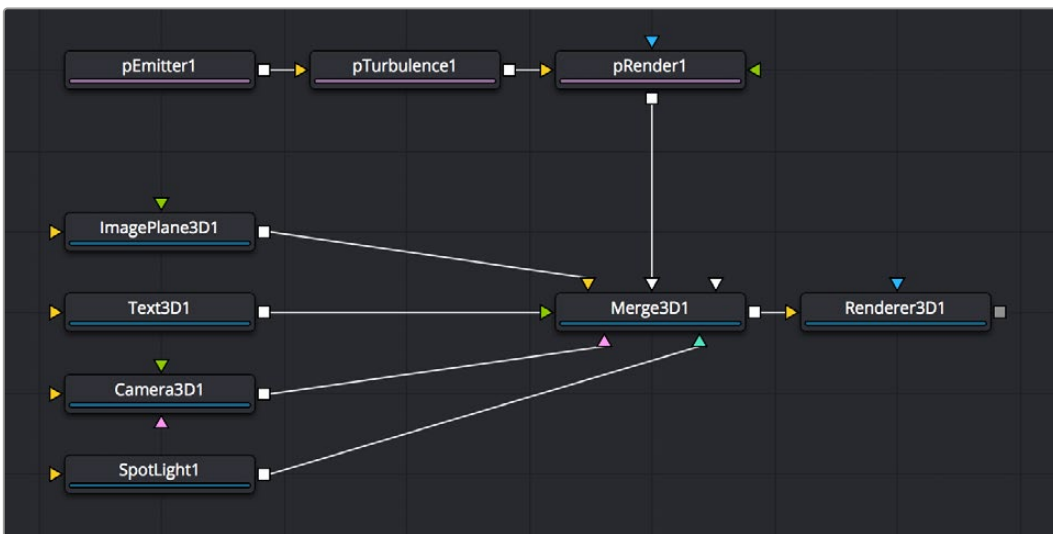
You can also attach the following types of nodes to a pEmitter node to deeply customize a particle system:

- Attach a 2D image to a pEmitter node to create highly customized particle shapes. Make sure your image has an appropriate alpha channel.
- Attach a Shape3D or other 3D geometry node to a pEmitter node to create a more specific region of emission (by setting Region to Mesh in the Region tab).




Customizing pEmitter nodes using mesh geometry to define regions and 2D images to define particle shape

The above examples assume that you'll output 2D renders to combine into the rest of a 2D composition. However, because particle systems are fully 3D, you also have the option of outputting your particle system in such a way as to be used from within other 3D scenes in your composition.




Connecting a particle system to a Merge3D node so the particles are subject to lighting and shadows within a 3D scene

The Output Mode of the pRender node, at the very top of the controls exposed in the Inspector, can be set to either 2D or 3D, depending on whether you want to combine the result of the particle system with 2D layers or with objects in a 3D scene.



Choosing whether a particle system's output is 2D or 3D in the pRender node's Inspector controls

If you connect a pRender node to a Merge3D node, the Output Mode is locked to 3D, meaning that 3D geometry is output by the pRender node for use within the Merge3D node's scene. This means that the particles can be lit, they can cast shadows, and they can interact with 3D objects within that scene.

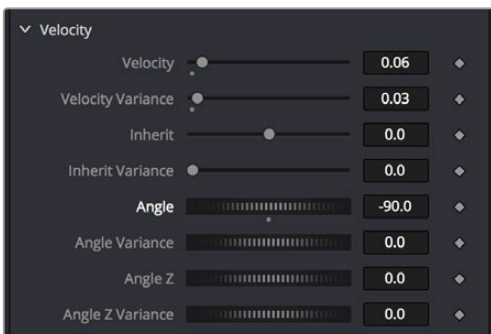


The result of using a particle system within a 3D scene

NOTE: Once you set the pRender node to either 2D or 3D and make any change to the nodes in the Inspector, you cannot change the output mode.


Particle System Distribution

To adjust the distribution of particles being emitted, select the pEmitter node to expose its controls in the Inspector, then open the Velocity controls in the Controls tab, and use the Angle, Angle Variance, Angle Z, and Angle Z Variance controls to adjust the direction and width over which particles are emitted. All these controls can be animated.



A pEmitter node's Velocity Angle and Angle Variance controls let you adjust the direction and width of particle distribution

Particle systems can be positioned and rotated by loading the pEmitter nodes that generate particles into a viewer and using the onscreen 3D position and Rotation controls provided to move the particle system around.



A pEmitter node loaded into the viewer with the rotation onscreen controls enabled

Alternatively, you can use the controls of the pEmitter's Region tab in the Inspector to adjust Translation, Rotation, and Pivot. All these controls can be animated.



A pEmitter node's Region controls open in the Inspector

Particle Nodes Explained by Type

This section introduces the four types of particle system nodes available in the Effects Library.

Emitters

pEmitter nodes are the source of all particles. Each pEmitter node can be set up to generate a single type of particle with enough customization so that you'll never create the same type of particle twice.

Along with the pRender node, this is the only other node that's absolutely required to create a particle system.

pEmitter nodes have four parameters tabs:

- **Controls:** The primary controls governing how many particles are generated (Number), how long they live (Lifespan), how fast they move (Velocity) and how widely distributed they are (Angle and Angle Variance), their rotation (Rotation Mode with X, Y, and Z controls), and whether there's spin (Spin X, Y, and Z controls). For each parameter of particle generation, there's an accompanying Variance control that lets you make that parameter less uniform and more natural by introducing random variation.
- **Sets:** This tab contains settings that affect the physics of the particles emitted by the node. These settings do not directly affect the appearance of the particles. Instead, they modify behaviors such as velocity, spin, quantity, and lifespan.
- **Style:** While the Controls tab has a simple control for choosing a color for particles, the Style tab has more comprehensive controls including color variance and Color Over Life controls. Additionally, size controls including Size Over Life, fade controls, and blur controls let you create sophisticated particle animations with a minimum of adjustments, while Merge controls give you an additional level of control over how overlapping particles combine visually. A set of controls at the bottom lets you choose how animated effects are timed.
- **Region:** The Region tab lets you choose what kind of geometric region is used to disperse particles into space and whether you're emitting particles from the region's volume or surface. The Winding Rule and Winding Ray Direction controls determine how the mesh region will handle particle creation with geometric meshes that are not completely closed, as is common in many meshes imported from external applications. Tweaking these last parameters is common when using imported mesh geometry as a region for emitting particles, since even geometry that appears closed will frequently appear to "leak" particles thanks to improperly welded vertices.

Forces

Many of the particle nodes found in the Particles bin of the Effects Library are "forces" that enhance a particle simulation by simulating the effect of various forces acting upon the particles generated by an emitter.

Some forces, including pDirectionalForce, pFlock, pFriction, pTurbulence, and pVortex, are rules that act upon particles without the need for any other input. These are simply "acts of nature" that cause particles to behave in different ways.

Other forces, such as pAvoid, pBounce, pFollow, and pKill, work in conjunction with 3D geometry in a scene such as shapes or planes to cause things to happen when a particle interacts or comes near that geometry. Note that some of the particles described previously can also use geometry to direct their actions, so these two categories of forces are not always that clear-cut.

Compositing

The pMerge node is a simple way to combine multiple emitters so that different types of particles work together to create a sophisticated result. The pMerge node has no parameters; you simply connect emitters to it, and they're automatically combined.

Rendering


The pRender node is required whether you're connecting a particle system's output to a 2D Merge node or to a Merge3D node for integration into a 3D scene. Along with the pEmitter node, this is the only other node that's absolutely required to create a particle system.

- **Controls:** The main controls that let you choose whether to output 2D or 3D image data, and whether to add blur or glow effects to the particle systems, along with a host of other details controlling how particles will be rendered.

- **Scene:** These controls let you transform the overall particle scene all at once.
- **Grid:** The grid is a helpful, non-rendering guide used to orient 2D particles in 3D space. The grid is never output in renders. The width, depth, number of lines, and grid color can be set using the controls found in this tab.
- **Image:** Controls the output of the pRender node, with controls over the process mode, resolution, and color space settings of the output.


Example Particle Systems

The Templates category in the Inspector in the Fusion page of DaVinci Resolve or in the Bins window in Fusion Studio includes over 20 different examples of particle systems creating a variety of effects. One of the best ways of learning how to create and customize particle systems is to open these and investigate how they're made.



Different particle system presets in the Templates category of the Bins window in Fusion Studio

Simply drag and drop any of the particle presets into the Node Editor, load the last node into the viewer, and you'll see how things are put together.



The Blowing Leaves preset from the Templates category

Chapter 88

Optical Flow and Stereoscopic Nodes


This chapter covers the numerous stereoscopic and optical flow-based nodes available in Fusion and their related workflows.

Contents

Overview	1591	DisparityToZ, ZToDisparity	1596
Stereoscopic Overview	1591	Separate vs. Stack	1596
Optical Flow Overview	1591	Setting Up Stereo in the Node Editor	1596
Toolset Overview	1591	About the Disparity Channel	1597
Working with Aux Deep Channels	1592	Viewing of Disparity and Vector Channels	1597
Optical Flow Workflows	1593	Stereo and Optical Flow Best Practices	1598
OpticalFlow	1593	Semi-Transparent Objects	1598
TimeSpeed, TimeStretcher	1593	Motion Blur	1598
SmoothMotion	1593	Depth of Field	1598
Repair Frame, Tween	1594	Where to Calculate Disparity and Optical Flow?	1598
Advanced Optical Flow Processing	1594	Cropping the Source	1599
Stereoscopic Workflows	1594	Nodes with Multiple Outputs	1599
Stereo Camera	1595	Picking from Aux Channels	1599
Stereo Materials	1595	Vector and Disparity Channels	1599
Disparity	1595		
NewEye, StereoAlign	1596		

Overview

Fusion includes 3D stereoscopic and optical flow-based nodes, which can work together or independently of each other to create, repair, and enhance 3D stereoscopic shots.



Stereoscopic comp displayed in the viewers.

Stereoscopic Overview

All stereoscopic features are fully integrated into Fusion's 3D environment. Stereoscopic images can be created using a single camera, which supports eye separation and convergence distance, and a Renderer 3D for the virtual left and right eye. It is also possible to combine two different cameras for a stereo camera rig.

Stereoscopic nodes can be used to solve 3D stereoscopic shooting issues, like 3D rig misalignment, image mirror polarization differences, camera timing sync issues, color alignment, convergence, and eye separation issues. The stereo nodes can also be used for creating depth maps.

Optical Flow Overview

Optical Flow analyzes the motion in a clip and generates motion vectors between neighboring frames. It generates X and Y vectors from the previous frame to the current frame (Back Vectors) and to the next frame in sequence (Forward Vectors). Once calculated, optical flow data can be used by other nodes to create smooth slow motion and variable retiming of clips, repair missing frames, and even correct disparity in stereo 3D clips.

NOTE: The stereoscopic nodes in the Fusion page work independently of the stereoscopic tools in the other DaVinci Resolve pages.

Toolset Overview

Here is an overview of the available nodes.

Optical Flow Nodes

- **Optical Flow > OpticalFlow:** Analyzes motion between neighboring frames in a sequence to generate motion vectors, which can then be used by other nodes for retiming, motion blur, and other effects.
- **Miscellaneous > TimeSpeed:** Retimes a clip at a constant speed using Flow Interpolation mode.
- **Miscellaneous > TimeStretcher:** Retimes a clip at variable speeds using Flow Interpolation mode.

- **Optical Flow > RepairFrame:** Generates a new frame using the motion vectors between two neighboring frames.
- **Optical Flow > SmoothMotion:** Smooths the color or aux channels using motion vectors.
- **Optical Flow > Tween:** Interpolates between two non-sequential images to generate a new frame.
- **Color > CopyAux:** Copies aux channels, including motion vectors, into RGBA more efficiently than Channel Booleans.

Stereoscopic Nodes

- **Stereo > Anaglyph:** Combines stereo images to create a single anaglyph image for viewing.
- **Stereo > Combiner:** sStacks a separate stereo images into a single stacked pair, so they can be processed together.
- **Stereo > Disparity:** Generates disparity between left/right images.
- **Stereo > DisparityToZ:** Converts disparity to Z-depth.
- **Stereo > Global Align:** Shifts each stereo eye manually to do basic alignment of stereo images.
- **Stereo > NewEye:** Replaces left and/or right eye with interpolated eyes.
- **Stereo > Splitter:** Separates a stacked stereo image into to left and right images.
- **Stereo > StereoAlign:** Adjusts vertical alignment, convergence, and eye separation.
- **Stereo > ZToDisparity:** Converts Z-depth to disparity.

Working with Aux Deep Channels

Certain image formats can contain channels other than RGBA color, called aux deep channels. Stereo Disparity and OpticalFlow deal directly with auxiliary deep channels.

Aux channels supported in Fusion include:

- **RGBA:** These are the standard colors.
- **Z:** The eyespace Z coordinate is almost always negative because in eyespace, Fusion's camera sits at (0, 0, 0) looking down the Z-axis. Z values start at Z = 0 at the camera focal point and progressively become more negative for objects deeper in the scene.
- **Coverage:** The percentage of the pixel covered by the frontmost pixel, used for antialiased Z-compositing.
- **Object ID:** These are user-assigned integers to meshes.
- **Material ID:** These are user-assigned integers to materials.
- **Texture Coords:** Normalized texture coordinates stored as (u, v) pairs.
- **Normal Vector:** Normal vector (nx, ny, nz) where the components are typically in the range [-1, +1].
- **Background Color:** The color of the pixel if the frontmost layer were removed, used for antialiased Z-compositing.
- **Vector:** The forward motion vector is an offset (vx, vy) that compares every pixel's position in one frame to the same pixel's position in the next frame.
- **Back Vector:** The backward motion vector is an offset (vx, vy) that compares every pixel's position in one frame to the same pixel's position in the previous frame.
- **World Position:** The position (wx, wy, wz) of the pixel in world coordinates.
- **Disparity:** An offset (dx, dy) that maps a pixel in the Left > Right or Right > Left frames.

Some extra channels are used by specific Fusion nodes.

For example:

- Merge can use the Z channel to perform a depth merge. If the Coverage and BackgroundColor channels are present, it can do a better job on antialiased edges during the Z merge.
- Most image-processing nodes (e.g., BrightnessContrast) have options on their common controls tab to limit their processing by MaterialID and ObjectID.
- The Fog and DepthBlur nodes make use of the Z channel.
- The Texture node makes use of the TexCoord channel.
- The Shader node makes use of the Normal channel.

There are a couple of ways to retrieve or generate those extra channels within Fusion.

For example:

- The Renderer3D node is capable of generating most of these channels.
- The OpticalFlow node generates the Vector and BackVector channels, and then TimeStretcher and TimeSpeed can make use of these channels.
- The Disparity node generates the Disparity channels, and then DisparityToZ, NewEye, and StereoAlign nodes can make use of the Disparity channels.
- The OpenEXR format can be used to import or export aux channels into Fusion by specifying a mapping from EXR attributes to Fusion Aux channels using CopyAux.

Optical Flow Workflows

The Optical Flow analysis is a non real-time process, and depending on your computer, the clip's resolution, and the duration of the clip, it can take some time. Because of this, the general idea is that you pre-generate the motion vectors, either by performing the analysis overnight or using a render farm, and save results into an OpenEXR sequence. The Optical Flow toolset is designed around four types of nodes that either generate, destroy, pass through, or construct the motion vectors.

OpticalFlow

The Optical Flow node generates the Vector and BackVector data. Typically, for optimal performance, you connect the Optical Flow output to a Saver to save the image as OpenEXR files with the motion vectors stored in an aux channel.

TimeSpeed, TimeStretcher

You can create smooth constant or variable slow-motion effects using the TimeSpeed or TimeStretcher nodes. When Optical Flow motion vectors are available in the aux channel of an image, enabling Flow mode in the TimeSpeed or TimeStretcher Interpolation settings will take advantage of the Vector and BackVector channels. For the Flow mode to work, there must be either an upstream OpticalFlow node generating the hidden channels or an OpenEXR Loader bringing these channels in. These nodes use the Vector/BackVector data to do interpolation on the motion channel and then destroy the data on output since the input Vector/BackVector channels are invalid. For more detail on TimeSpeed or TimeStretcher, see Chapter 109, "Miscellaneous Nodes," in the DaVinci Resolve Reference Manual and Chapter 49 in the Fusion Reference Manual.

SmoothMotion

SmoothMotion can be used to smooth the Vector and BackVector channels or smooth the disparity in a stereo 3D clip. This node passes through, modifies, or generates new aux channels, but does not destroy them.

Repair Frame, Tween

The Tween and Repair Frame nodes are different from standard optical flow nodes because they have the OpticalFlow analysis and motion vector generation built in. Tween will compare two frames and create an in-between frame, which is good for recreating a missing or flawed frame. Repair Frame will look at frames on either side of the current frame and repair scratches, dust marks, and so on. Because these nodes work with flow values between non-sequential frames, they cannot use the optical flow stored in the input image's Vector/BackVector channels, but rather must regenerate the flow of each frame, do their processing, and then destroy the flow channels. This being the case, these nodes are computationally expensive. For more detail on Tween or Repair Frame, see Chapter 110, "Optical Flow," in the DaVinci Resolve Reference Manual or Chapter 50 in the Fusion Reference Manual.


Advanced Optical Flow Processing

The Optical Flow, Repair Frame, and Tween nodes include a faster GPU-based Optical Flow algorithm. When you add the Optical Flow, Repair Frame, or Tween node to a comp, the Inspector includes a Method drop-down menu where you can choose Advanced to enable the GPU-based algorithm. This Advanced method is the same Optical Flow algorithm used in other DaVinci Resolve pages.

By choosing Classic from the Method drop-down menu in the Inspector, you can use the older CPU-based algorithm to maintain compatibility with comps created in previous versions. This method may also be better suited for some Stereo3D processing.

Stereoscopic Workflows

Disparity is the difference between the left and right image. The Disparity map is used by nodes to align and massage the stereo pair of images.



The Disparity node analyzes a stereo pair of images and generates an X&Y disparity map.

The workflow is to load a left and right stereo image pair and process those in the Disparity node. Once the Disparity map is generated, other nodes can process the images.


TIP: When connecting stereo pairs in the node tree, make sure that the left and right images are connected to the left and right inputs of the Disparity node.

Disparity generation, like Optical Flow, is computationally expensive, so the general idea is that you can pre-generate these channels, either overnight or on a render farm, and save them into an EXR sequence.

The toolset is designed around this philosophy.

Stereo Camera

There are two ways to set up a stereoscopic camera. The common way is to simply add a Camera 3D and adjust the eye separation and convergence distance parameters.




Stereoscopic cameras can be done with a single camera or two connected cameras.

The other way is to connect another camera to the RightStereoCamera input port of the Camera 3D. When viewing the scene through the original camera or rendering, the connected camera is used for creating the right-eye content.

Stereo Materials

Using the Stereo Mix material node, it is possible to assign different textures per eye.



Material viewer showing stereoscopic material.

Disparity

The Disparity node does the heavy lifting of generating disparity maps. This generates the Disparity channel and stores it in the hidden aux channels of their output image.

NewEye, StereoAlign

NewEye and StereoAlign use and destroy the Disparity channel to do interpolation on the color channel.

The hidden channels are destroyed in the process because, after the nodes have been applied, the original Disparity channels would be invalid.

For these nodes to work, there must be either an upstream Disparity node generating the hidden channels or an OpenEXR Loader bringing these channels in.

DisparityToZ, ZToDisparity

These nodes pass through, modify, or generate new aux channels, but do not destroy any.

TIP: If the colors between shots are different, use Color Corrector or Color Curves to do a global alignment first before calculating the Disparity map. Feed the image you will change into the orange input and the reference into the green input. In the Histogram section of the Color Corrector, select Match, and also select Snapshot Match Time. In the Color Curves' Reference section, select Match Reference.

Separate vs. Stack

Stereo nodes can work in Separate or Stack modes. When in Stack mode, the left/right eyes are stacked horizontally or vertically, forming one image with double width or height, respectively.

The advantage to using Stack mode is that you do not have to have duplicate branches of the Node Editor for the left and right eyes. As a consequence, you will see Stereo nodes with two inputs and two outputs labeled as "Left" and "Right."

When in Stack mode, the stack should be connected to the left eye input and the Left output should be used for connecting further nodes. In Stack mode, the respective Right eye inputs and outputs are hidden.

Setting Up Stereo in the Node Editor

The disparity generation is the first operation. This can be configured in the Node Editor in two different ways.



Two stereoscopic workflows.

In the above example, the workflow on the right takes the left and right eye, generates the disparity, and then NewEye is used to generate a new eye for the image right away.

The example on the left renders the frames with disparity to intermediate EXR images. These images are then loaded back into Stereo nodes and used to create the NewEye images.

By using Render nodes to compute the disparity first, the later processing of the creative operations can be a much faster and interactive experience.

Although not shown in the above diagram, it is usually a good idea to color correct the right eye to be similar to the left eye before disparity generation, as this helps with the disparity-tracking algorithm. The color matching does not need to be perfect—for example, it can be accomplished using the “Match” option in a Color Corrector’s histogram options.

About the Disparity Channel

The Disparity channel stores the displacement vectors that match pixels in one eye to the other eye. The left image’s Disparity channel will contain vectors that map left>right and the right image’s Disparity channel will contain vectors that map right>left. For example:

```
(xleft, yleft) + (Dleft. x, Dleft. y) -> (xright, yright) (xright, yright) + (Dright. x, Dright. y) -> (xleft, yleft)
```

You would expect for non-occluded pixels that $D_{left} = -D_{right}$, although, due to the disparity generation algorithm, this is only an approximate equality.

NOTE: Disparity stores both X and Y values because rarely are left/right images perfectly registered in Y, even when taken through a carefully set up camera rig.

Both Disparity and Optical Flow values are stored as un-normalized pixel shifts. In particular, note that this breaks from Fusion’s resolution-independent convention. After much consideration, this convention was chosen so the user wouldn’t have to worry about rescaling the Disparity/Flow values when cropping an image or working out scale factors when importing/exporting these channels to other applications. Because the Flow and Disparity channels store things in pixel shifts, this can cause problems with Proxy and AutoProxy. Fusion follows the convention that, for proxied images, these channels store unscaled pixel shifts valid for the full-sized image. So if you wish to access the Disparity values in a script or via a probe, you need to remember to always scale them by $(\text{image. Width}/\text{image. OriginalWidth}, \text{image. Height}/\text{image. OriginalHeight})$.

Viewing of Disparity and Vector Channels

Aux channels can be displayed directly in the viewers through the Channel viewer button’s menu.

The CopyAux node is used to copy those channels directly into the RGB channels for viewing or further processing. The advantage of using the CopyAux node is that it does static normalization, which reduces a lot of flicker that the viewer’s time-variant normalization causes. When viewing long sequences of aux channels, the CopyAux node has the option to kill off aux channels and keep only the current RGB channels, freeing up valuable memory so you can cache more frames.

TIP: Although you can use the Channel Booleans to copy any aux channel into RGBA, it involves a few additional clicks when compared to CopyAux.

One thing to be aware of is that aux channels tend to consume a lot of memory. A float-32 1080p image containing just RGBA uses about 32 MB of memory, but with all the aux channels enabled it consumes around 200 MB of memory.

Stereo and Optical Flow Best Practices

How you create your composition, the images you are using, and the type of shot you are working on can all have an impact on the success of the Disparity generation and Optical Flow analysis. Below, we'll look at some of the situations to be aware of and how you can avoid some pitfalls when dealing with optical flow.

Semi-Transparent Objects

The Optical Flow and Disparity generation algorithms Fusion uses assume there is only one layer per pixel when tracking pixels from frame to frame. In particular, transparent objects and motion blur will cause problems. For example, a shot flying through the clouds with the semi-transparent clouds in the foreground and a distant landscape background will confuse the Optical Flow/Stereo algorithms, as they do not recognize overlapping objects with different motions. Usually the optical flow will end up tracking regions of one object or the other. If the transparent object and the background are near the same depth and consequently have the same disparity, then it is not a problem.

Motion Blur

Motion blur is also a serious problem for the reason explained in the previous point. The Disparity and Optical Flow algorithms are unsure whether to assign a pixel in the motion blur to the moving object or the background pixel. Because the algorithms used are global in nature, not only the vectors on the motion blur will be wrong, but it will confuse the algorithm on regions close to the motion blur.

Depth of Field

Depth of field is also another problem related to the above two problems. The problem occurs when you have a defocused foreground object over a background object that is moving (Optical Flow case) or shifts between L/R (Stereo Disparity case). The blurred edges will confuse the tracking because they can't figure out that the edges are actually two separate objects.

Where to Calculate Disparity and Optical Flow?

Where you choose to generate optical flow or disparity in your composition can drastically affect the results.

For example, if you have composited a lens flare in, it is better to compute OpticalFlow/Disparity before that, since the semi-transparent lens flare will confuse the tracking algorithms.

If you are color correcting the left/right eyes to match or for deflickering, it is better to apply the OpticalFlow/Disparity afterward, since it will be easier for the tracking algorithm to find matches if the color matches between frames.

If you are removing lens distortion, think carefully about whether you want to do it before or after Disparity computation. If you do it after, your Disparity map will also act as a lens distortion map, combining the two effects as one.

As a general rule of thumb, it is best to use OpticalFlow/Disparity before any compositing operations except an initial color matching correction and a lens distortion removal.

Cropping the Source

As a general tip, if you are cropping down your input images for any reason, it is probably better to compute the optical flow or disparity before the crop and then afterward crop the flow/disparity along with the color.

The reason is that flow/disparity matching works well when there is common pixel data to match in both frames, but when there are pixels that show up in just one frame (or one eye), then the Disparity/OpticalFlow nodes must make a guess and fill in the data. The biggest occlusions going from L \leftrightarrow R are usually pixels along the L/R edges of the images that get moved outside. This is similar for optical flow when you have a moving camera.

Another thing to be aware of are black borders around the edges of your frames, which you should crop away.

Nodes with Multiple Outputs

Many of the stereo nodes in the Fusion toolset have multiple outputs. This can cause some confusion to new users. One particularly confusing thing is that when you drag a Stereo node to the view, it will always display the left output. There is no way to view the right output without connecting another node like BC (BrightnessContrast) to the right output and viewing that.

Picking from Aux Channels

Some nodes, like StereoAlign, allow one to drag pick from the Z or Disparity auxiliary channels. You must pick from a node upstream of the StereoAlign, not from the output of the StereoAlign. If you try to pick a disparity from the output of a StereoAlign node, you will get nothing because StereoAlign consumes/destroys the Disparity aux channel (and even if it did not destroy the Disparity channel, you would still be picking the wrong value since you would be picking from the aligned result).

The typical workflow for picking is:

- 1 View StereoAlign in the left view.
- 2 View the node upstream of StereoAlign in the right view.
- 3 Pick the Disparity value from the left eye in the right view.

Although this picking functionality does not operate any differently from normal picking of color channels, this issue may cause some confusion. If it helps, the analogous workflow mistake with color nodes would be a user trying to pick a gradient color for a Background node from a view showing the Background node itself (you are trying to pick a color for a node from its own output).

Another issue that you need to be aware of is which eye you are picking. To avoid problems, it's a good idea to always pick from the left eye. The reason is that the Disparity channels for the left and right eyes are different, and when you pick from a horizontal/vertical stereo stack, Fusion has no way of knowing whether you picked the Disparity value from the left or right eye.

The above are not hard and fast rules; rather, they are guidelines to prevent foot shootings. If you understood the above reasoning fully, you'll realize there are exceptions, like picking disparity from the left output of DisparityToZ and Z from the left/right output of ZToDisparity, where everything is okay.

Vector and Disparity Channels

The Vector and BackVector channels store the forward and reverse optical flow.

The Vector channel might be better named "forward vector" or "forward flow," since the name "Vector" to describe a channel is "not technically correct," as the more mathematically-inclined user might recognize that all the channels except the scalar channels Z/ID are technically "vector" channels. A frames Vector aux channel will store the flow forward from the current frame to the next frame in the sequence, and the BackVector aux channel will store the flow backward from the current frame to the

previous frame. If either the previous or next frames do not exist (either not on disk or the global range of a Loader does not allow OpticalFlow to access them), Fusion will fill the corresponding channels with zeros (transparent black).

The Disparity channel stores the displacement vectors that match pixels in one eye to the other eye. The left image's Disparity channel will contain vectors that map left > right and the right image's Disparity channel will contain vectors that map right > left.

For example:

$(x_{left}, y_{left}) + (D_{left.x}, D_{left.y}) \rightarrow (x_{right}, y_{right})$ $(x_{right}, y_{right}) + (D_{right.x}, D_{right.y}) \rightarrow (x_{left}, y_{left})$

You would expect for non-occluded pixels that $D_{left} = -D_{right}$, although due to the disparity generation algorithm, this is only an approximate equality. Note that Disparity stores both X and Y values because rarely are left/right images perfectly registered in Y, even when taken through a carefully set up camera rig.

Disparity and Optical Flow values are stored as un-normalized pixel shifts. In particular, note that this breaks from Fusion's resolution-independent convention. After much consideration, this convention was chosen so the user wouldn't have to worry about rescaling the Disparity/Flow values when cropping an image or working out scale factors when importing/exporting these channels to other applications. Because the Flow and Disparity channels store things in pixel shifts, this can cause problems with Proxy and AutoProxy. The convention that Fusion follows is that, for proxied images, these channels store unscaled pixel shifts valid for the full-sized image. So if you wish to access the disparity values in a script or via a probe, you need to remember to always scale them by $(\text{image.Width}/\text{image.OriginalWidth}, \text{image.Height}/\text{image.OriginalHeight})$.

When using Vector and BackVector aux channels, remember that all nodes expect these aux channels to be filled with the flow between sequential frames.

More precisely, if you have sequence of three frames A, B, C, then:

B	Vector will contain the flow B>C
B	BackVector will contain the flow B>A
A	Vector will contain the flow A>B
A	BackVector is written with zeros as there is no frame before A
C	Vector is written with zeros as there is no frame D to flow C>D
C	BackVector will contain the flow C>B

When working with these channels, it is the user's responsibility to follow these rules (or for clever users to abandon them). Nodes like TimeStretcher will not function correctly since they still expect the channels to contain flow forward/back by 1 frame.

NOTE: Currently DoD/RoI is not supported for all Fusion nodes.