FUSION 8

USER MANUAL

Blackmagic design

September 2016

Fusion 8

The world's most advanced visual effects and motion graphics solution!

Fusion 8

Fusion is the world's most advanced compositing software for visual effects artists, broadcast and motion graphic designers and 3D animators. With over 25 years of development, Fusion has been used on over 1000 major Hollywood blockbuster feature films! Fusion features an easy and powerful node based interface so you can construct complex effects simply by connecting various types of processing together. That's super easy and extremely fast! You get a massive range of features and effects included, so you can create exciting broadcast graphics, television commercials, dramatic title sequences and even major feature film visual effects!

Fusion 8

Fusion 8

Fusion 8

Fusion 8

Fusion 8

Fusion 8

Fusion 8

Fusion 8

Fusion 8

Getting Started

1

Getting Started

Introduction to Fusion

Fusion is the world's most advanced visual effects and motion graphics compositing application for Mac OS X, Windows and Linux. It has been used for compositing on over 1,000 Hollywood blockbuster feature films and prime time television shows, including Avengers: Age of Ultron, The Hunger Games, Breaking Bad, and Downton Abbey, as well as hundreds of commercials, music videos and game productions.

Compositing is the art of blending imagery from multiple sources to create a new whole. The job of the visual effects artist or motion graphics designer is to take all of the discrete elements that make up a shot, and turn them into a single image that matches the artistic vision of the production. To accomplish this, you use techniques as diverse as color correction, green screen replacement, tracking, particles, blur and glow effects.

Node based compositing was once only available to large visual effects studios creating the most sophisticated visual effects in the world. Now, with Fusion, the efficiency and flexibility of powerful node based composting is available to everyone. Fusion uses its unique node based compositing environment with over 200 image processing tools to improve the compositing workflow. There is also a deep and powerful Macro system for combining tools together as a way for you to create your own plug-ins without needing to know how to write code.

Fusion is a resolution independent compositing application. That means it can use images of varying resolutions within a single composition but more importantly it normalizes the tool controls not matter what the resolution of the image maybe. For example, a .4 blur results in the same effect in HD, UltraHD or 8K. Similarly, Transform controls are normalized so the center of the image is always .5 irrespective of the resolution.

There are two versions of Fusion available. The free version of Fusion for Mac OS X and Windows can be downloaded from the Blackmagic Design web site. It contains all the features and tools for a single artist or small studio to produce incredible Ultra HD visual effects and motion graphics. Larger facilities can choose to purchase Fusion Studio to access more advanced features, including higher resolution output up to 32K by 32K, third party OFX plug-in support, network based distributed rendering, as well as 3D Stereoscopic and Optical Flow analysis tools.

If you are new to Fusion, we recommend reading through the Getting Started section of this manual. It will introduce you to the user interface and take you step by step through an example workflow that will get you up and running quickly by covering the essential concepts.

System Requirements

Before installing Fusion, you should review the following system requirements to make sure you are getting the best possible performance.

Hardware Requirements

Processor

Intel and AMD processors; Core i7 or AMD multi-core processor minimum.

Graphics

OpenGL and OpenCL capable graphics cards with 2GB memory minimum. 3D scenes that make heavy use of textures will require more graphics memory.

RAM

8GB of memory minimum; 16GB or greater is strongly recommended.

USB Port

1 open USB port for Fusion Studio hardware Licensing lock.

Hard Drives

1GB of hard disk space for installation.

Tip: Fusion will work on MacBook Pro, or Windows laptops with Nvidia or AMD graphics.

Software Requirements

Operating Systems

Mac OS X Yosemite, Microsoft Windows 10 64-bit, Microsoft Windows 8 64-bit, Linux Centos 6.x and 7.x 64bit.

QuickTime

Apple QuickTime required to read and write MOV.

Keyboard and Mouse Conventions

You can use Fusion on Mac OS X, Windows or Linux platforms. Commands that are platform specific have been documented whenever possible. This section summaries the differences.

In most cases Mac OS X and Windows keyboard commands are identical, with the Command and Ctrl keys being interchanged.

The Mac OS X Delete key is isolated below the F12 key and is the equivalent of the Backspace key on Windows.

It is highly recommended that you use a three button mouse for contextual menus and navigational controls. Although you can use a combination of keyboard short cuts to mimic the three button mouse functionality, you will find you are much more productive investing in a three button mouse.

In this documentation we assume a three button mouse is present and we refer to the buttons as follows:

Mouse Button	Documentation Reference
Left mouse button	Click
Middle mouse button	Middle click
Right mouse button	Right click

Working with a Stylus

Fusion can be used with a graphics tablet and stylus. The stylus can often come in handy when working with the Paint tool or rotoscoping. To make your stylus work efficiently with Fusion, you can set the pen tip to be your left mouse click, the top thumb button to work with a right mouse click and the lower thumb button to work as the middle mouse click.

Using Blackmagic Design Capture and Playback Hardware
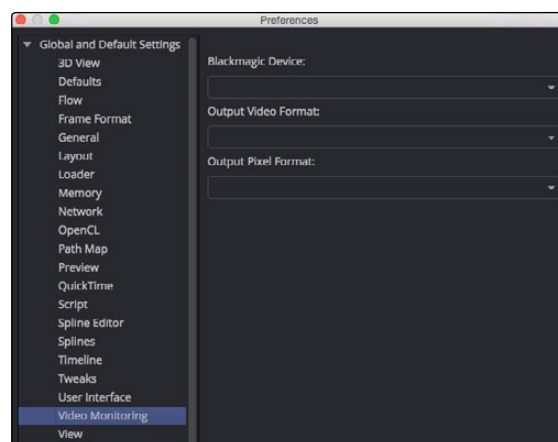
Blackmagic Design capture and playback products like DeckLink PCIe cards, can be used for monitoring your work on an HD, Ultra HD or DCI 4K display. The Video Monitoring Preferences allow you to select the type of video card you are using, the output resolution and the pixel format.

User Interface Overview

Once you install Fusion and open the application, a new, empty project opens. The projects you create in Fusion are called compositions or "comps." When you start working on a comp, you'll use five main areas.
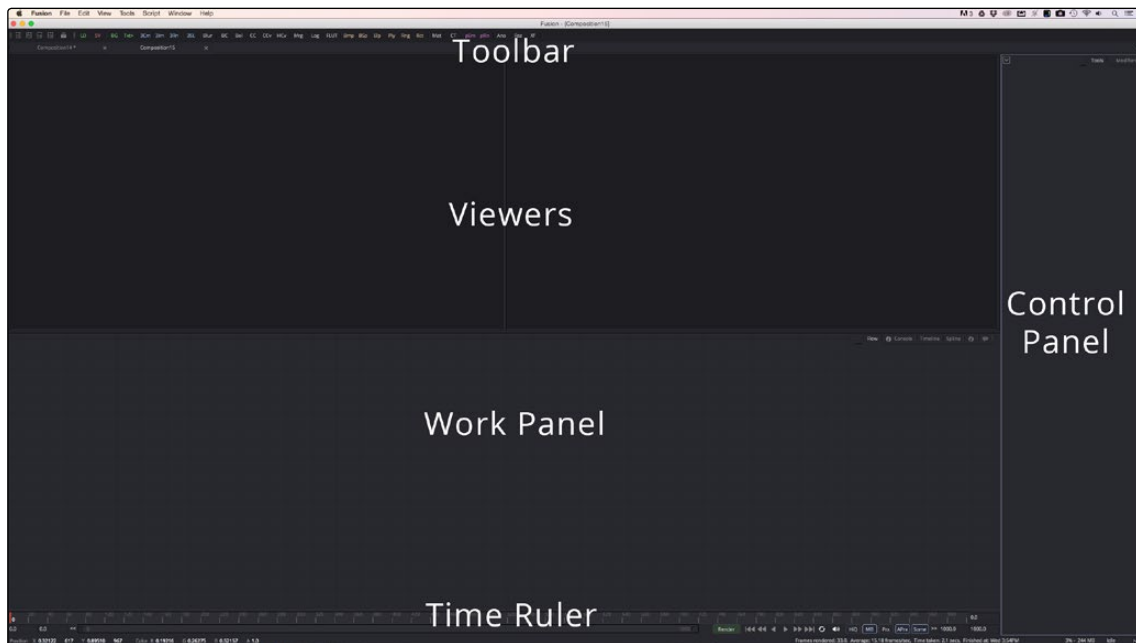
Viewers, Work Panel, Control Panel, Time Ruler, Tools Bar

Viewers

There are two identical Viewers in Fusion that are displayed above the Work Panel. The viewers can show 2D images from your composition as well as any 3D environment that you set up. You can load images into either viewer or use one viewer for images and the other viewer to display waveforms, histograms and vectorscopes. You can create floating Viewers if additional viewers are needed, or use an external video monitor by installing a DeckLink card.

Dual viewers are accelerated using graphics hardware

Work Panel

The Work Panel switches between the three main editing panels that you use to construct your composition in Fusion. The three main panels are the Flow Node Editor, the Timeline Editor and the Spline Editor. There are also tabs to switch to a scripting console, a comments windows and a chat window.

The Work Panel uses the tabs in the upper right to switch to different editors

Flow Node Editor

The Flow Node Editor is the default Work Panel view and is the primary work area you use to construct your compositions. You use the Flow Node Editor to create and display the structure of your composition much the way you use a layer list to put together elements in a graphics application. Instead of rigid layers, your composition is viewed as a schematic or flow chart where each image, mask, filter effect, and graphic, as well any other element in your comp, is represented by a tile called a Tool. The Flow Node Editor is used to connect tools together to modify how images are processed. By default tools in the Flow connect from left to right, but can be set to flow from top to bottom, or all directions simultaneously. The direction of the Flow is not related to how the Flow is processed. For example, the downstream nodes do not have to be below the upstream nodes because Fusion follows the arrows not the flow. The schematic view offers a lot of flexibility for connecting tools, including the ability to branch a tool's output into separate image processing streams and collapse entire groups of tiles to view a single thumbnail image of several tools' outputs.
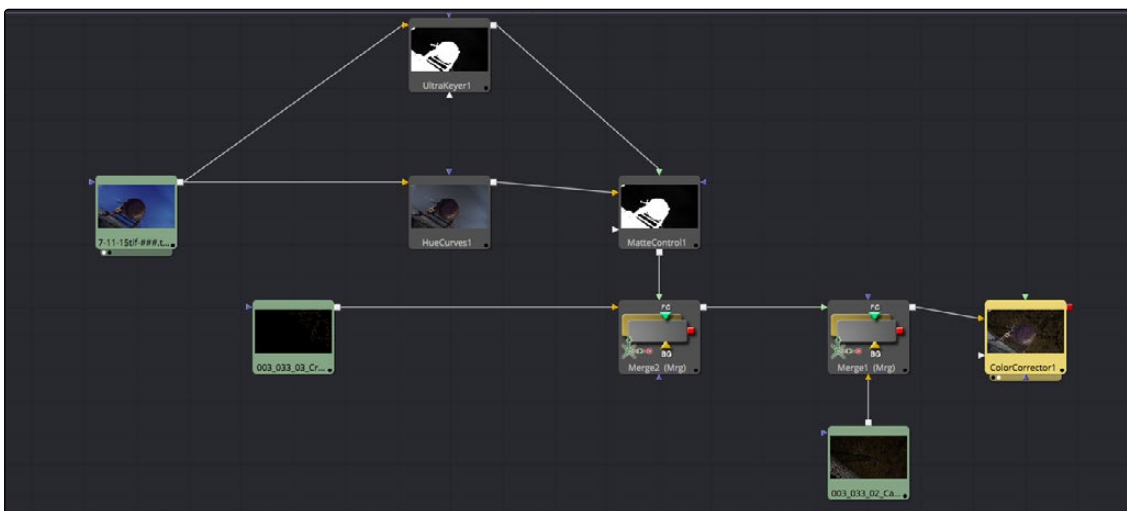
The Flow Node Editor with tools connected to create a composition

Control Panel

The Control Panel is a tabbed panel on the right side of the Fusion window that you use to modify a tool's parameters. When a tool is selected in the Flow Node Editor, its parameters and settings appear in the Control Panel.

The Control Panel shows parameters from a selected tool

Time Ruler

The Time Ruler, located beneath the Work Panel, includes transport controls for playback as well as for setting composition options like proxy level and motion blur. The Time Ruler's playhead determines the frame that is displayed in the viewers.

The Time Ruler uses a playhead to move to a specific frame

Toolbar

The Toolbar, located along the top of the Fusion window, gives you quick access to commonly used tools and functions. The toolbar is divided into two customizable sections, one for preset layouts, and one for Tools, Macros and Scripts. Each toolbar section can be detached and moved to other parts of the interface, and additional toolbars can also be created as needed.

The customizable Toolbar can show commonly used tools and functions

Timeline Editor

The Timeline Editor is used for adjusting the timing of elements in a comp. You can use the Timeline to trim and extend clips, to adjust the timing of an animation spline, or to trim the duration of an image processing effect. Unlike layer based applications, the Timeline does not determine the layer order in your composition, only the Flow Node Editor does that.

The Timeline Editor is used to adjust the timing of elements and animations

Spline Editor

The Spline Editor has three main areas: the Graph view, Tree view and Toolbar. The Graph view shows the animation spline along two axes. The horizontal axis represents time and the vertical axis represents the spline's value. The Tree view is where you decide what splines are visible in the Graph view. Checkboxes beside each name are used to show whether that spline is currently displayed in the Graph view and whether the spline is enabled for editing. The toolbar across the top of the Spline Editor is used to quickly access the most common operations applied to splines. The thin red bar that runs vertically through the Spline Editor represents the playhead in the comp.
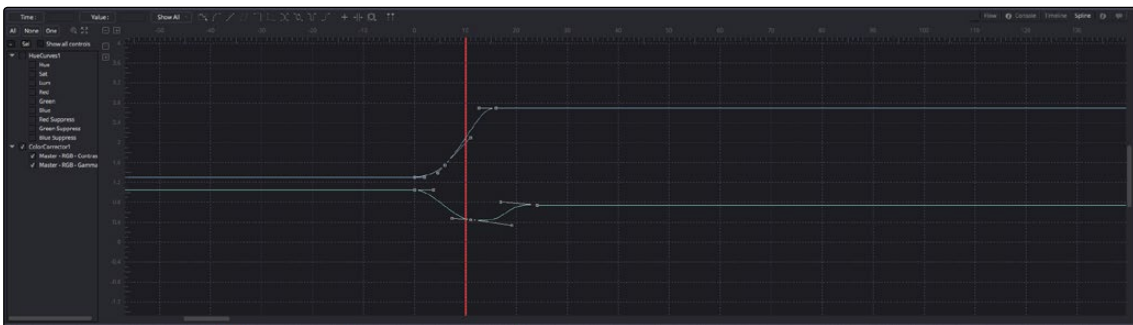
The Spline Editor is divided into the Graphics view, Tree view and Toolbar

Navigating the Interface

This section explains how you can adjust the position and scale of the main panels in the Fusion interface. The methods described here generally apply to all parts of Fusion.

The various panels may also provide additional methods for panning and scaling. We'll describe those in more detail in the chapter dedicated to that panel. Examples of these include the Flow's Navigator window and the Spline Editor's axis labels.

Panning and Scaling

As your comp grows in size, elements that you want to work on will often be off screen. The Flow Node Editor contains a Navigator window, which can be used to select an area to view in a large comp. You can also pan using the mouse.
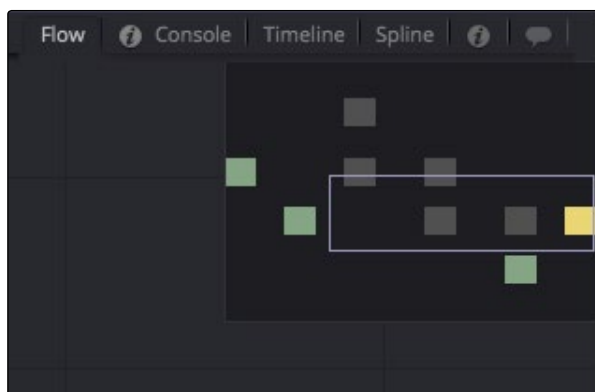
The Navigator window for accessing
offscreen parameters or tools

To pan a panel do one of the following:

→ Hold the middle mouse button, then drag in any direction.

→ Hold Command–Shift (Mac OS X) or Ctrl–Shift (Windows), then drag in the panel.

→ On Mac OS X, swipe in any direction using the Apple Mouse.

To smoothly adjust the scale of a panel:

→ Holding the left and middle mouse buttons down, drag in the panel. Moving the pointer to the left, the panel will zoom out. Moving it to the right will zoom in.

→ On Mac OS X, hold Command and swipe up to zoom in or swipe down to zoom out.

If you want to scale a panel in steps, you can use the mouse or keyboard short cuts.

To scale a panel in steps do one of the following:

→ Hold the middle mouse button down, then click the left mouse button to zoom in.

→ Hold the middle mouse button down, then click the right mouse button to zoom out.

→ Press + on the numeric keypad to zoom in.

→ Press - on the numeric keypad to zoom out.

To scale a panel using the contextual menu do the following:

→ Press the right mouse button while the pointer is positioned over the Viewer or Work Panel, then choose Scale from the Viewer's contextual menu.
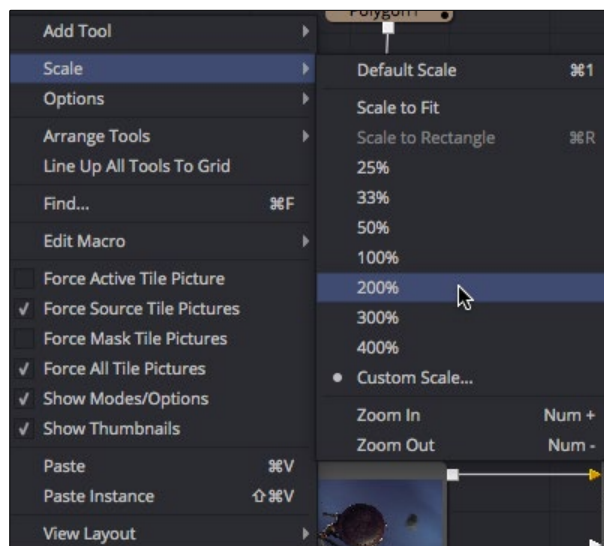
A contextual menu can be used for choosing a Scale percentage

Changing the Layout

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 edges of a panel will resize it

The Layout toolbar provides four preset panel layouts. The first two switch between a single viewer and dual viewers, leaving the Control Panel full size along the entire right side of the Fusion window. The second two layouts also switch between single and dual viewers, but they expand the viewers by making the Control Panel half height, aligning it to the bottom of the display views. To toggle between these modes, click on the small arrow button in the top left corner of the tool controls.The display views will adjust their width to compensate automatically.
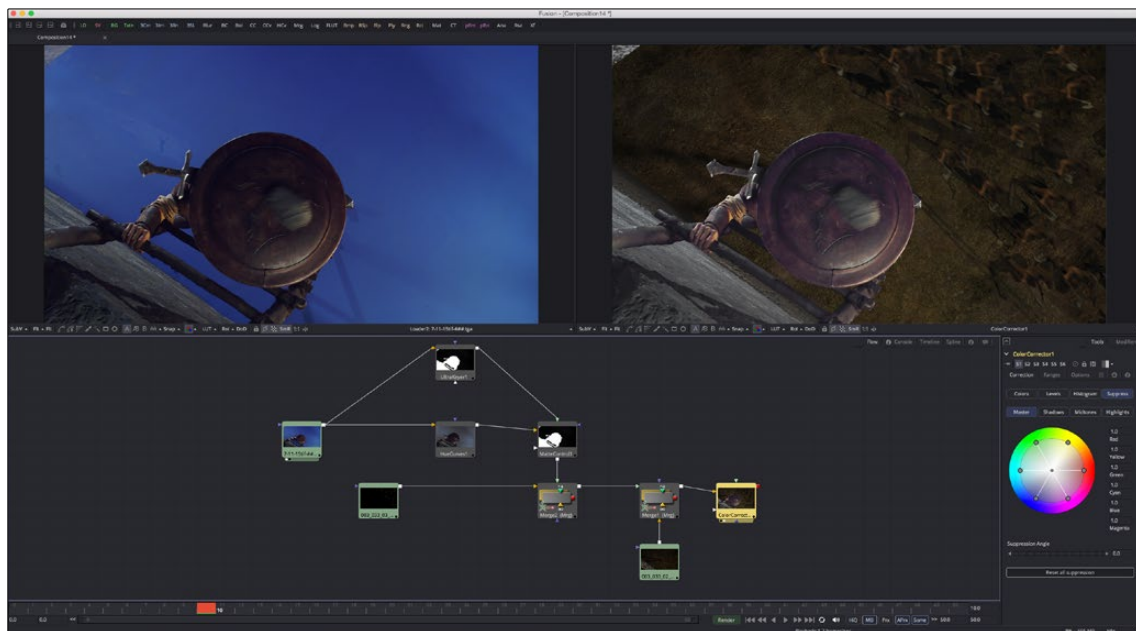
The Control Panel at half heigh with dual viewers

Pressing F4 on the keyboard will collapse all interface panels except the one under the mouse pointer, allowing it to take over almost the entire screen. This is incredibly helpful when you want to see greater detail in the Viewer or move tiles in the Flow graph in a large comp. The F4 key is a toggle, pressing it again restores the previous layout.

Tabbed Layout

Panels in Fusion can be arranged to better fit your personal workflow. They can be arranged both in the main window, as well as within floating views.

To add a tab to an existing panel do the following:

1. Right click in the upper right corner of any viewer or tab.

2. From the Add View menu choose the desired view from the list.
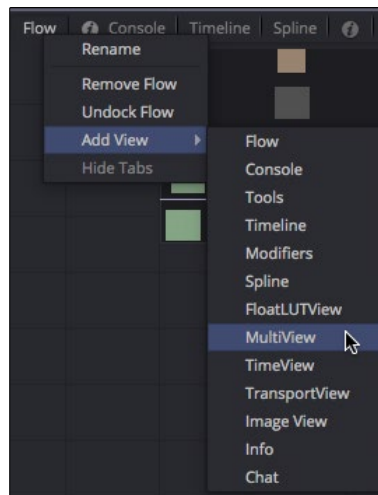
The Add View is accessed by right clicking
in the corner of a viewer or panel

You can change the order of tabs by dragging them with the cursor to the desired position. When you drag a tab out of its original position, the View Arrange menu will appear with nine drop zones. Five drop zones are centered in the panel the cursor is hovering over, and four drop zones are along the outer perimeter of the Fusion interface.
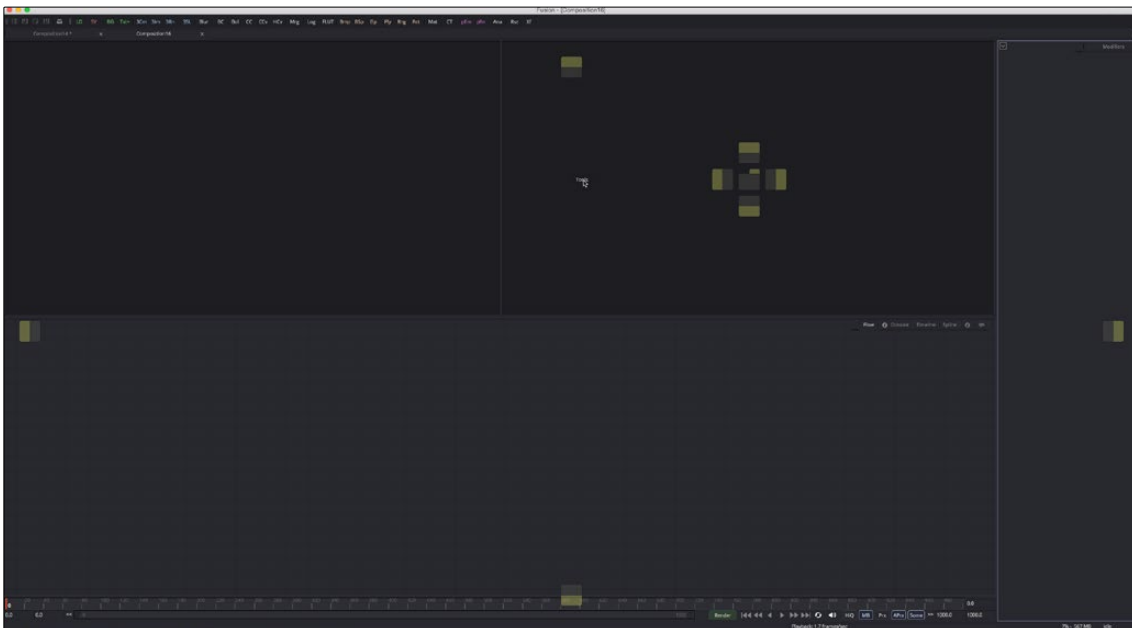
The View Arrange menu

If you use the top, left, right or bottom drop zones, the tab splits the panel the pointer is over and adds the tab into a panel of its own. Using the center drop zone adds a new tab in the current set of tabs in the panel, rather than splitting it. The perimeter drop zones split the Fusion window and adds the tab at the respective position.

You can choose Window > View Layout to return to the default layout or right click in an empty upper right corner area of the any panel to bring up the View Layout menu to Remove or Undock a tab. You can also select to Show or Hide tabs here.

The View Layout menu accessed from the menu bar

About Compositions

When you first launch Fusion, a new empty comp is opened. You can begin adding images to the comp immediately or use the Settings to define the resolution of your project. Since Fusion is resolution independent, comps do not ask you to define a resolution or frame rate. Those settings are determined in the Globals Frame Format preferences. When you save a comp the file extension .comp will automatically be added to the end of the file name you choose. Compositions that have not been saved or have unsaved changes will display an asterisk (*) next to the composition's name in the title bar.

To create a new composition do one of the following:

→ Choose File > New.

→ On your keyboard press Command – N (Mac OS X) or Ctrl – N (Windows).

To save a composition do one of the following:

→ Choose File > Save.

→ Choose File > Save As to save the composition with a new name.

→ On your keyboard press Command – S (Mac OS X) or Ctrl– S (Windows).

Unsaved compositions have an asterisk next to their name

To open an existing composition do one of the following

→ Select File > Open.

→ Press Command–O (Mac OS X) or Ctrl – O (Windows).

→ Drag a Composition file from a file browser or Finder window into an empty Fusion panel.

→ Drag a Composition file from the Bin into an empty Fusion panel.

→ Double click on a Composition file in a Windows file browser or Mac OS X Finder window.

To close an existing Composition do one of the following

→ Choose File > Close.

→ Click on the 'Close' button in the top right of the composition's window on Windows or the top left on Mac OS X.

Auto Save

You can set Fusion to automatically save the current composition to a temporary file at pre-set intervals. This helps to protect you from loss of work due to a power loss, software issues or accidental closure. An auto save file does not overwrite the currently saved composition. A file with the same name is created in the same folder as the composition, but with the extension .autosave. If an autosave file is present when you open a composition, a dialog is displayed allowing you to open either the autosaved or original version.

To enable Auto Save for new compositions

1. Choose File > Preferences.

2. Select the Global category, then select General.

3. In the Auto Save section, click the Auto Save check box to enable it.

The Auto Save settings are located in the General Global preferences

The Composition File Format

When you save a composition, it is saved as a plain text file which allows it to be opened and edited using any standard text editing application. Since Fusion comps are readable text files it is easier to integrate them with other software like asset management software.

Tip: Popular word processors on Mac OS X and Windows generally save additional formatting information when they save documents. This makes them inappropriate choices for editing Fusion Composition files. Instead, use a plain text editor like SciTE or the Mac OS X TextEdit program set to plain text mode.

Most of the time, the advantages of saving a composition as a text format far outweigh the disadvantages, but there is one scenario where it may be desirable to save a composition as a compressed binary file instead of plain text. A composition that contains a very large number of animated spline masks can sometimes become abnormally large, with sizes in excess of 100MB. In these cases you can enable the option to save compositions in a compressed binary format in the Global General preferences.

Working with Multiple Compositions

You can have multiple compositions open simultaneously. A list of the currently open compositions are seen in the Window menu, allowing you to switch between them. Having multiple compositions open at once makes it possible to copy/paste between them or to compare parts of one composition with another.

To switch to another open composition do one of the following:

→ Choose Window > composition name.

→ Press Ctrl – Tab (Windows).

Example Workflow

If you are new to Fusion or even new to compositing, this brief example is designed for you. It takes you step by step through a short but complete project, beginning with loading images and ending with a final render. In a single example workflow we'll only cover a small fraction of Fusion's functionality, but the goal of this tutorial is to teach you the core concepts of how Fusion works. Although it is only intended as an example, after reading through it you'll then be able to navigate around the user interface and explore the specifics of Fusion at your own pace. To begin we'll start with an empty composition and load in a clip.

Loading Clips

1. Open Fusion using the Mac OS X Dock or the Windows Start menu.

    Fusion opens with an empty project called a composition, or comp for short. Next you'll want to load in source media.

2. Choose Tools > I/O > Loader.
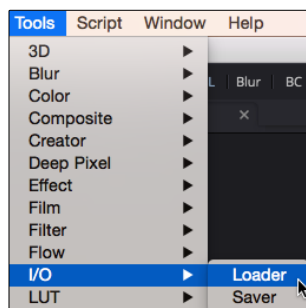
To read in clips use the Loader tool

The Loader is the primary way you get source media into a comp. After choosing the Loader tool from the Tools menu an Open dialog appears allowing you to navigate to the location of your source media on your hard drives.

3. Navigate your hard drives to locate the source media, select it, then click Open.
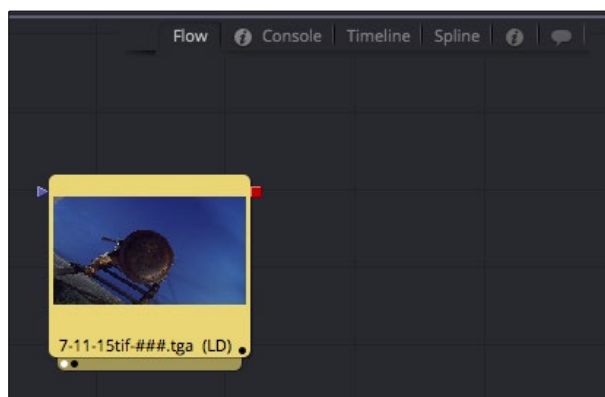
A tile represents any tool you add to the Flow Node Editor

The dialog closes and the Loader tool with the selected source media is represented by a tile in the Flow Node Editor. Every piece of media you bring into Fusion, as well as any filter effect or compositing operation, will be displayed as a tool in the Flow Node Editor.

Viewing Clips

The two viewers at the top of the Fusion windows are identical and can be used to view any tool.

1. To view a tool, select it in the Flow Node Editor.

2. Press the 1 key on the keyboard to have it displayed in the Viewer on the left
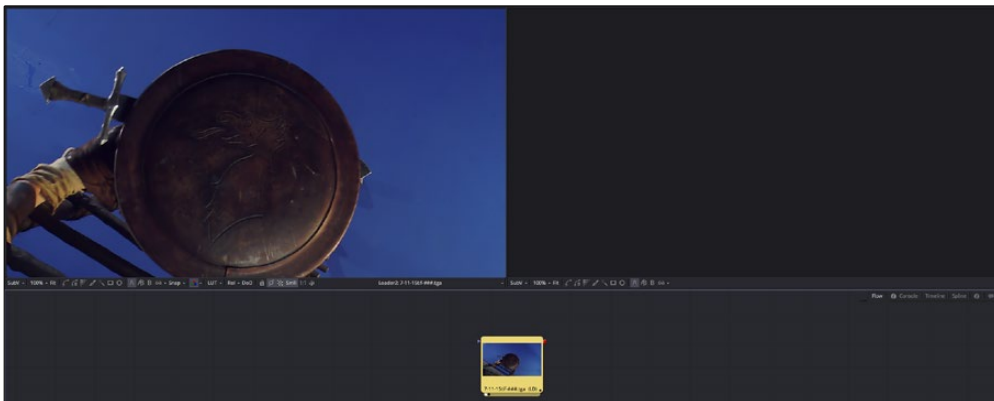
Pressing 1 displays the output of the selected tool in the Viewer on the left

A small white icon is displayed at the bottom of the tool to indicate you are viewing the tool on the left Viewer.

3. Press the 2 key on the keyboard to have it displayed in the Viewer on the right
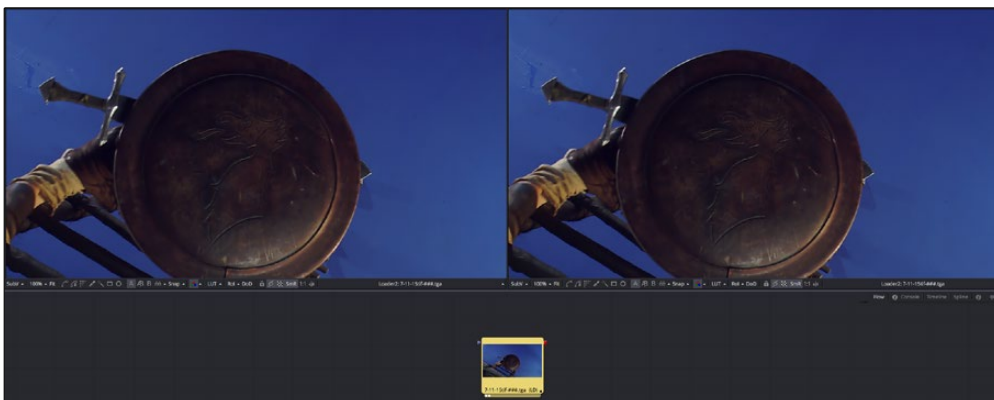
Two white icons on a tool indicate on which viewer the tool is displayed

A second white icon is displayed to indicate you are viewing this tool on both the left and right viewers. The keyboard short cuts are toggles, allowing you to view and remove tools from a viewer quickly.

4. Press the 2 key to remove the tool from the Viewer on the right

You can pan and zoom in the Viewer to see different areas of an image with more detail.

5. To zoom in on an image in the Viewer, press the middle and left mouse buttons, then drag right. On Mac OS X, hold command and swipe down to zoom in using the Apple Mouse.

The Viewer smoothly zooms in on the area where your cursor begins dragging.

6. Click the Fit button under the Viewer to see the entire image.
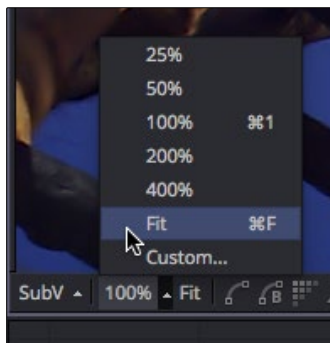
Buttons under each viewer allow you to zoom in and out or fit the entire image in the Viewer

Playing Clips

The Time Ruler at the bottom of the Fusion window is used to play or scrub through the frames in your composition. The Time Ruler defaults to a duration of 1,000 frames.

1. To shorten or lengthen the duration of a comp, type the desired number of frames in the Global End number field.

The Global End number field sets the end of your composition

2. Click anywhere on the Time Ruler to jump to that frame and have it displayed in the Viewer.

Click in the Time Ruler to move to a specific frame

Transport controls allow you to play your composition forward or reverse, and step through frame by frame.

3. Click the Play button or press the spacebar to play the composition. Press the spacebar again to stop.

The Transport controls in the Time Ruler

Tip: Compositions can contain hundreds of tools with high resolution source media. Playback is not always done in real time, depending on the size of the comp and whether or not the frames are cached in memory.

4. To move to the beginning of the comp, click the start of the Time Ruler to jump to the first frame.

Adding and Removing Tools

The Flow Node Editor works by connecting tools together, sending the image from one tool to another for a specific process to be performed.

1. Select a tool in the Flow Node Editor that you want to apply another tool to, such as a filter effect.

2. Choose Tools > Color > Brightness/Contrast to add a brightness contrast filter to the selected tool in the Flow Node Editor.
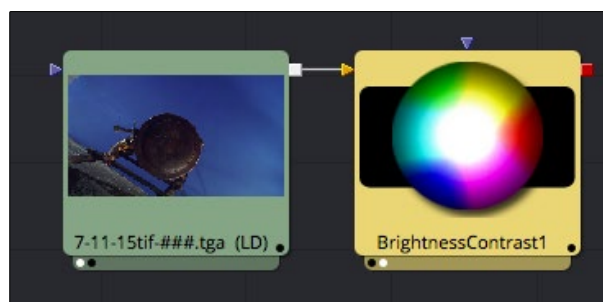
Two tools connected in the Flow Node Editor

Tools attach together via "connections," seen as lines that connect one tool to another. Image data flows from left to right, by default, although the tools can be rearranged. Tools can also be manually connected if nothing is selected in the Flow Node Editor.

3. Click in an empty area of the Flow Node Editor to deselect all the tools.

4. Choose Tools > Matte > Ultra Keyer.
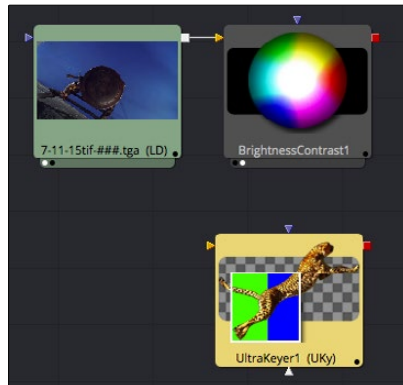
Tools can be added to the Flow Node Editor unattached to other tools

The Ultra Keyer is added to the Flow Node Editor, but it is not connected to any other tool. The Ultra Keyer is used to extract mattes from images using blue screen or green screen backgrounds as we have here.

The Ultra Keyer converts blue or green areas of an image into transparent areas

5. To manually connect tools, drag from the red square output of one tool to the yellow triangle input of another

Drag from the output to the input to connect tools

Tip: If you find it difficult to connect tools, you can hold the right mouse button and drag to the center of the Tool tile. A pop up menu will appear with the different inputs. You can now click the input you want to connect.

6. To view the newly connected tool, select it, then press 2 on the keyboard.

The Loader image displayed in the left Viewer,
the Ultra Keyer displayed in the right Viewer

The original image is displayed in the left Viewer while the right Viewer displays the output of the Ultra Keyer.

You can disconnect tools in the Flow Node Editor if you want to reconnect them someplace else.

7. To disconnect a tool, position the pointer over the connection closest to the input.

Position the pointer over the connection near the Input to disconnect it

Half the connection closest to the input highlights.

8. Click the connection to remove it.

If you find you no longer need a tool in the composition you can remove it completely.

9. To remove a tool from the Flow Node Editor, select it, then press the Delete key (Mac OS X) or the Del (Windows). You can also right click over the selected tool, then choose Delete from the contextual menu.

Choose Delete from the contextual menu to remove a tool from the comp

Making Adjustments to Tools

Every tool has parameters you can change to modify their results. Accessing a tool's parameters is done through the Control Panel.

1. Select a tool in the Flow Node Editor.
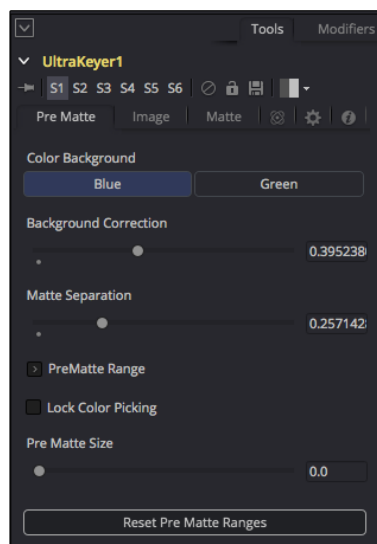
Parameters for the Ultra Keyer in the Control Panel

The parameters for the selected tool are displayed in the Control Panel. The Ultra Keyer tool has a number of parameters for turning a blue or green background into a transparent area.

2. Click the Blue button so the parameters operate on the image's blue screen background.

   When extracting a matte from a blue screen or green screen image, you are not concerned with the colors of the image. You are only concerned with the matte or alpha of the image.

3. Click the Channels button to view the alpha channel.
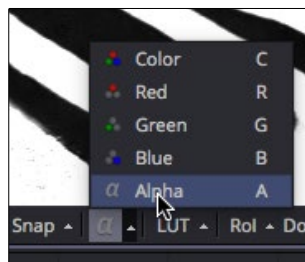
Viewing the alpha or matte channel

   When viewing a matte or alpha channel, areas that are white are fully opaque, areas that are black are fully transparent. Gray areas are varying degrees of transparency. The goal is to have a clean solid white for the areas you want to keep in the image and a clean solid black for the areas you want fully transparent.

4. Drag the Matte Separation slider to the right until you see erosion to the white areas of your matte.
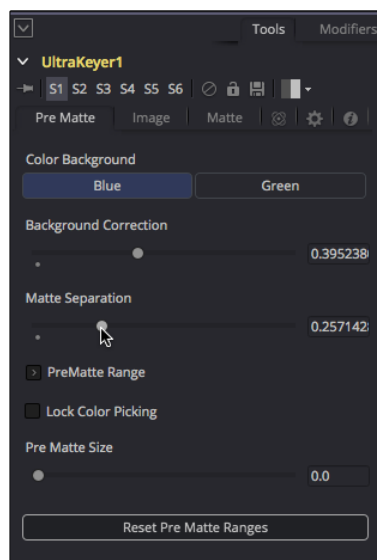
View the alpha channel after using the Matte Separation slide

   The Matte Separation slider should be able to get your blacks nice and dark.

   The Ultra Keyers, like many tools' controls, are divided into tabs to organize parameters.

5. Click the Matte tab in the Control Panel.
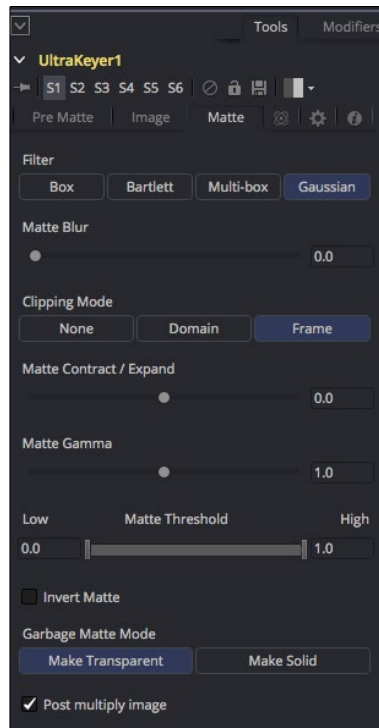
The Matte tab in the Ultra Keyer's controls

The Matte tab has many controls for modifying the matte.

6. To clean the white region of our matte, drag the High Threshold parameter of the Matte Threshold until the white region is solid.

The Matte Threshold pushes gray areas below a certain
point to black and above a certain point to white

Tip: Often on very difficult blue screen or green screen clips, using multiple Ultra Keyers to address different areas (e.g., foreground vs. edges or different regions of unevenly lit blue or green screens) can produce better results than attempting to do it with a single tool.

At some point you need to play through your entire clip to make sure your adjustments work for all the frames you need to key.

7. Click in the Time Ruler to display the first frame in your comp.

8. Press spacebar to play through the clip, then press spacebar to stop.

Drawing Masks

Often in blue or green screen shots there are areas of the frame which are not covered by the blue or green screen and they are not part of the set. They are unwanted objects in the frame. The easiest way to remove objects like these is to create a spline based mask. Creating a mask that either surrounds unwanted objects, often called a garbage mask, or a mask that surrounds your foreground subject, makes using blue and green screen keyers easier, by allowing you to focus on a smaller area.

1. Hold the middle and left mouse buttons, then drag left to zoom out on the image in the Viewer
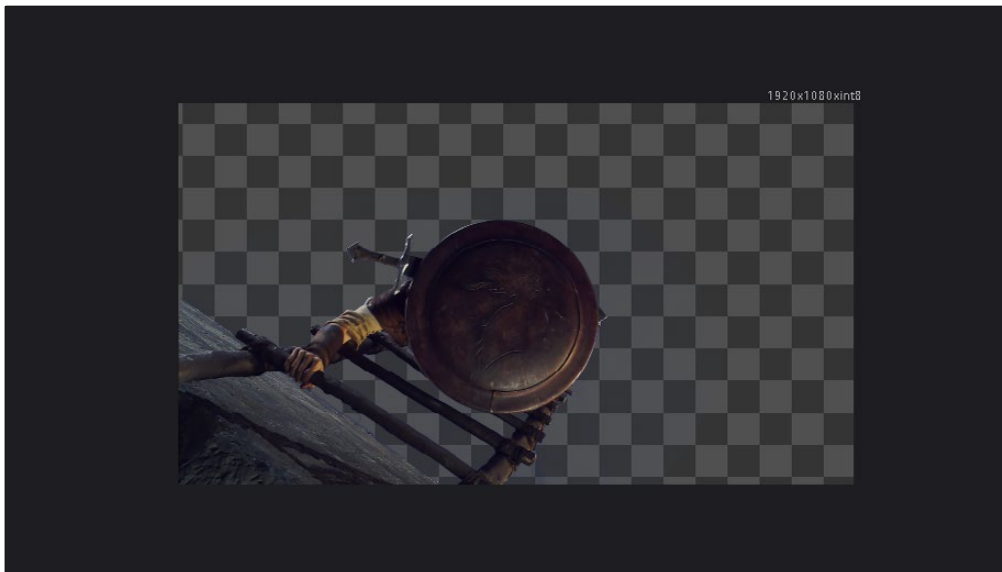
The Viewer zoomed out, to provide room for drawing a mask

Zooming out on an image can give you more room to draw a mask that needs to go outside the frame boundary.

2. Click in an empty area of the Flow Node Editor to deselect any tools.

3. Choose Tools > Mask > Polygon.

    A Polygon tool is added to the Flow Node Editor, unconnected to any other tool. The Polygon drawing tools are displayed along the right side of the Viewer.
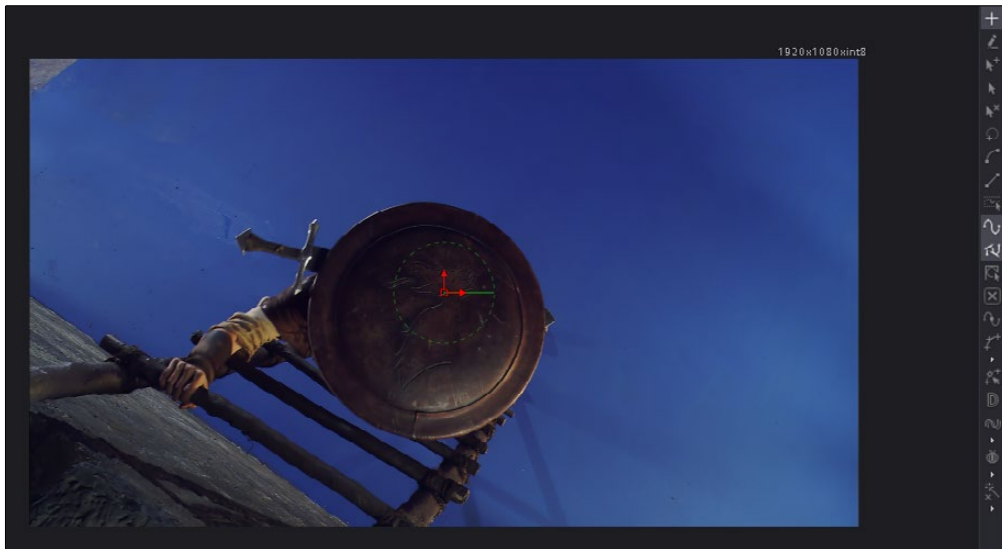
Viewing the Ultra Keyer while the Polygon tool is selected

In Fusion, you can view any tool in the Flow Node Editor while selecting a different tool to perform adjustments. Here, we are viewing the Ultra Keyer mask, but selecting the Polygon tool so we can use its drawing tools.

4. To draw a polygon spline, click where you want the spline to begin, adding an anchor point at that location.

5. Position the pointer to set another anchor point, then click to draw a line between the two points.

The default Polygon tool draws straight lines between points

6. To create a curved corner, position the pointer where you want the next anchor point to be placed, then press the mouse and drag out to set the slope of the curve. The more you drag away from the point, the more exaggerated the curve will be.

Click and drag to create a curved corner

7. When you are ready to close the mask, place the pointer over the first anchor point you added, then click to close the spline.

The close pointer is displayed when positioned
over the first anchor point in a spline

With the spline drawn you can connect it to a mask input on a tool. Depending on the type of mask input and the tool you connect it to, the spline can be used in different ways.

8. Drag the red output of the polygon tool to the white garbage matte input of the Ultra Keyer.
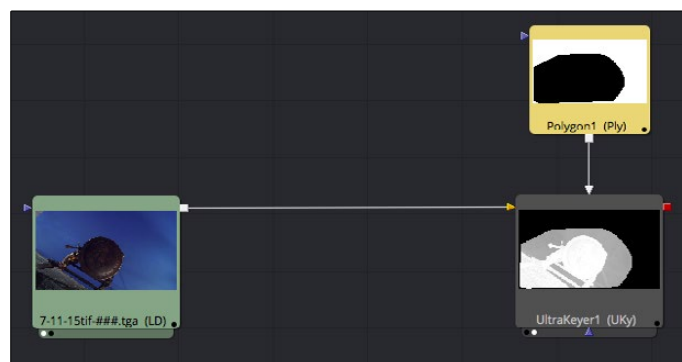
Using the garbage mask input of the Ultra Keyer to remove unwanted objects in a frame

The object is removed using the spline and garbage matte input.

Branching Tools

With the matte completed, it is time to deal with the color of our foreground image.

The foreground image has areas of blue around the edges. This is called spill. It's a result of light reflected off the blue screen onto the foreground. There are dedicated parameters in the Ultra Keyer to deal with spill colors, but it gives you more control to deal with matte creation using one tool and color balancing with another. To do this, you can have the output of your Loader go in two or more different directions, connecting to different tools. This is called branching.

1.   Click in an empty area of the Flow Node Editor to deselect any tools.

     Tip: For learning purposes, we are deselecting all the tools in the Flow Node Editor so we can connect them manually when we add new tools. As you become more proficient with Fusion, you'll become familiar with which tools to select in the Flow to have them connect automatically.

2.   Choose Tools > Color > Hue Curves.

     A Hue Curves tool is added to the Flow Node Editor, unconnected to any other tool.

3.   Drag from the output of the Loader to the yellow input of the Hue Curves tool.
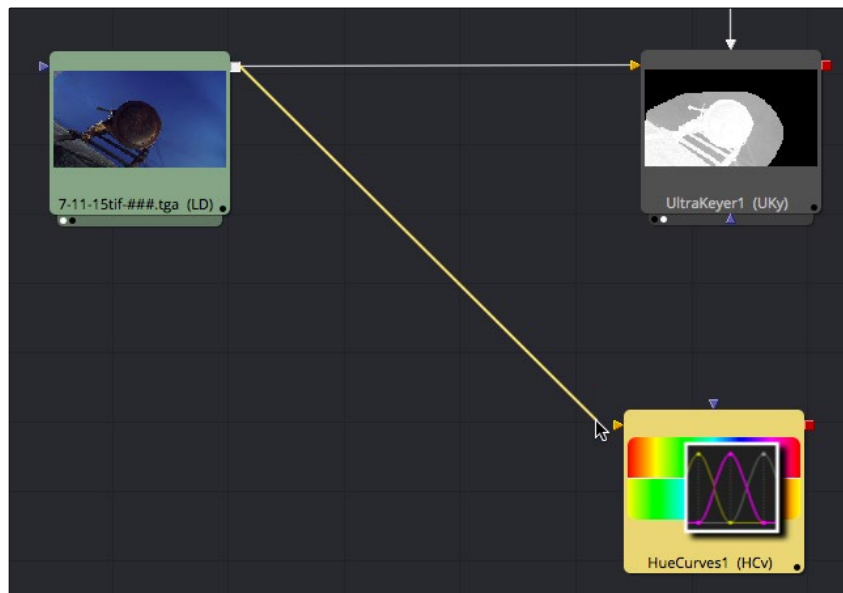
Branching eliminates the need for duplicating layers

Branching allows you to split the output of a tool and perform different and independent operations on the different branches.
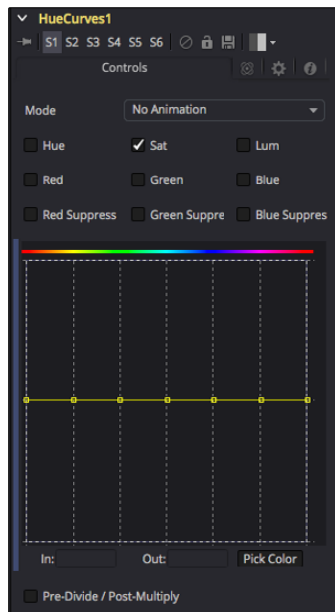
Hue Curves can selectively lower the saturation
of a specific hue, such as blue or green

4. Select the Hue Curves tool to display its parameters in the Control Panel.

5. Press 1 to display the results of the Hue Curves tool in the left Viewer.

   The Hue Curves tool displays a graph in the Control Panel. The horizontal axis represents the various hues as shown above the graph. The vertical axis increases or decreases the saturation.

6. Drag the control point below the blue hue, down to decrease the saturation.

Adjusting the blue or green hue can be used to remove spill

There are many ways to remove blue spill. Using the Hue curves to decrease a very narrow slice of blue saturation is just one way.

Combining a Foreground with an Alpha Channel

The RGB part of our image is done; now we have to combine it with the alpha channel we created earlier using the Ultra Keyer. To combine an alpha channel with an RGB image you use a Matte Control tool. The alpha channel connects to the background and the RGB image connects to the foreground.

1. Select the Hue Curves tool in the Flow Node Editor, then choose Tools > Matte > Matte Control.
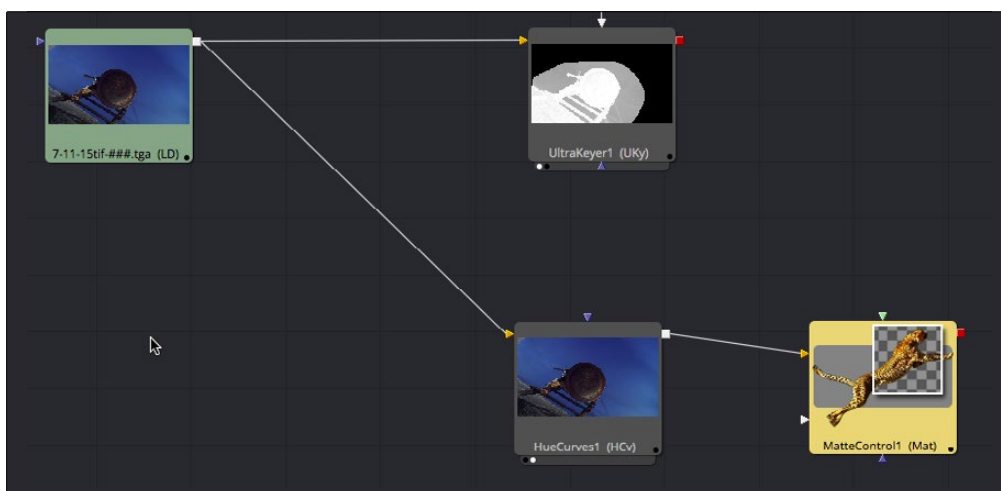
Hue Curves connected to the background input of the Matte Control

Since the Hue Curves was selected, it is automatically connected to the background input of the Matte Control.

2. Press the right mouse button over the Ultra Keyer red output and drag to the center of the Matte Control tile.

3. When you release the mouse, choose Foreground from the pop up menu. This will connect the output of the Hue Curves to the Foreground of the Matte Tool.
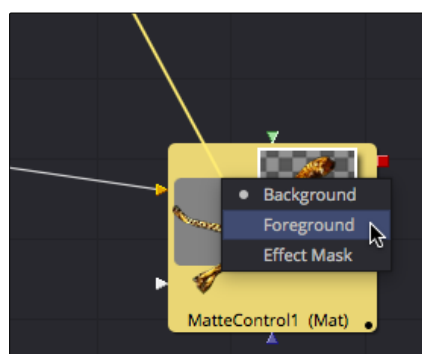
The popup menu on the Matte Control tile

4. With the Matte tool selected, choose Combine Alpha from the Matte Combine menu in the Control Panel
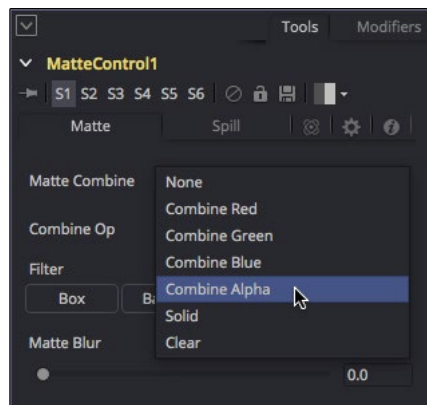
The Matte Control parameters allow you to combine alpha channels

Now, the RGB image is combined with the alpha from our Ultra Keyer. The transparent area shows a checkerboard background which we can now replace with a real background for the shot.

The checkerboard background shows where the transparency is located

Compositing Clips in the Flow Node Editor

Compositing in Fusion is done with two clips at a time; a foreground and a background. To combine these two clips you use a Merge tool. The Merge tool is the main compositing tool in Fusion.

1. Add another Loader to read in a second clip that you will use as the background for the composite.
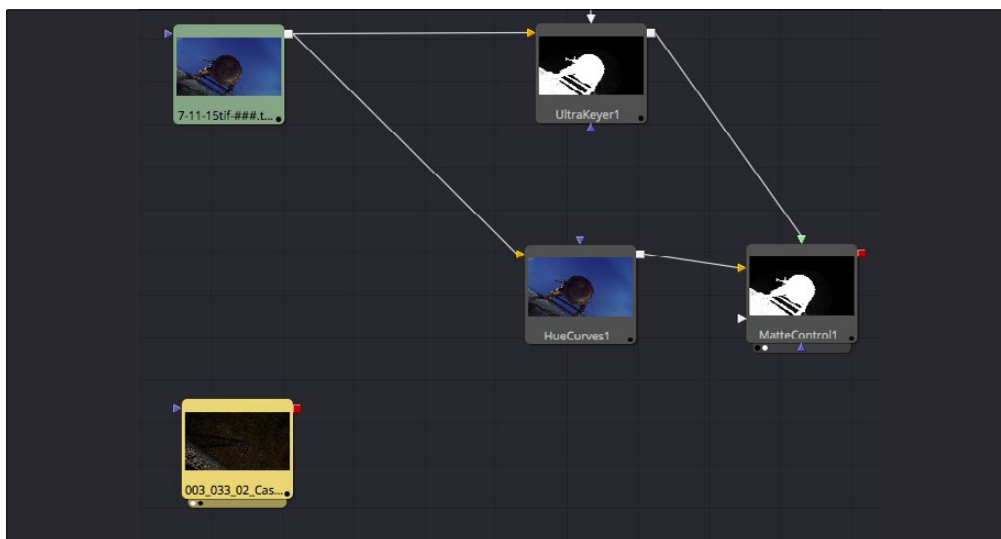
Unconnected Loaders have no affect on a comp

The background clip is placed on the Flow Node Editor. Since it isn't connected to any other tool it has no effect on your comp.

2. Again, click in an empty area of the Flow Node Editor to deselect any tool in the Flow Node Editor.

3. Choose Tools > Composite > Merge.

The Merge tool is the main compositing tool in Fusion

The Merge tool has three inputs: one for the background, one for the foreground and one for a mask.

4. Drag the output of your Background loader into the yellow background input of the Merge tool.

5. Drag the output of the Matte Control tool into the green foreground input of the Merge tool.

Drag to connect the background and foreground images to the Merge tool

Tip: If the background and foreground images are of different resolutions, the background determines the output resolution of the Merge tool. To change the output resolution of the Merge tool, add a Resize tool between the Background loader and the Merge tool input, then scale the background to the resolution you want.

6. The Merge tool blends two images. It also has controls for standard Apply modes and transparency.

The foreground with alpha combined with the background clip

To build up composites with additional images, you connect one Merge tool to another.

7.   Select the Background loader in the Flow Node Editor, then choose Tools > Composite > Merge.
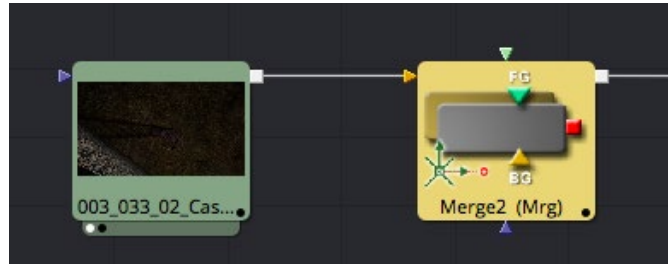
New tools are automatically connected to any selected tool in the Flow Node Editor

Instead of manually having to connect tools, new tools that are added automatically connect to the selected tool in the Flow Node Editor. The selected tool connects to the background input in the case of a Merge tool. This new merge tool provides an open foreground input between your original background and the blue screen image.
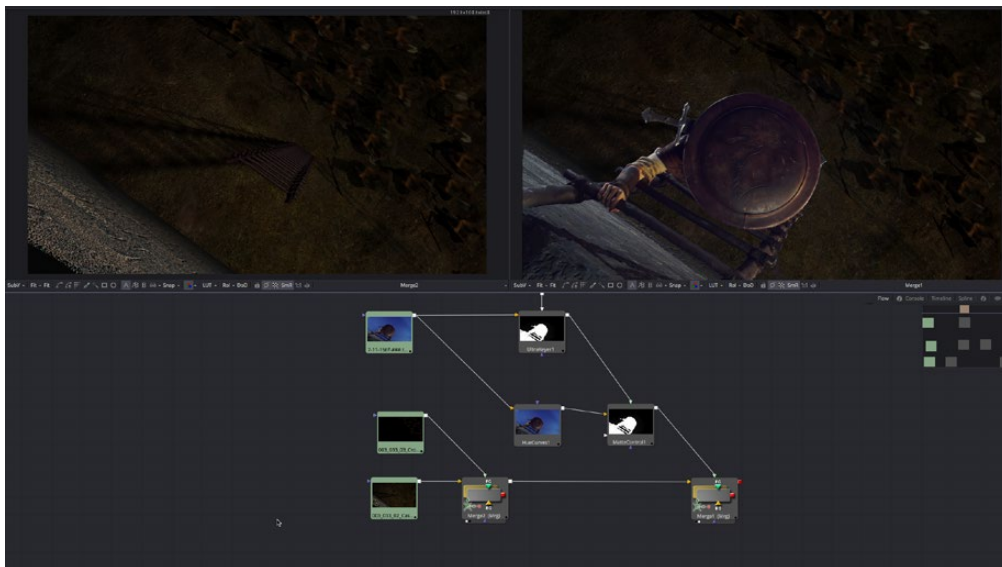
The results of two Merge tools displayed on the left and right Viewers

8.   Drag the output of a third Loader tool to the foreground input of the new Merge tool.

If you need to adjust any other tool in the Flow Node Editor, you can select it at any time and make parameter adjustments while viewing the results of the last tool in the Flow Node Editor.

Exporting a Movie

To get your composition rendered out to disk you can set a render range, then connect a Saver tool as the last tool in your composition.

1. In the Time Ruler, enter the frame number where you want to begin rendering. Often this is frame 0, but not always.

The Render Start field in the Time Ruler

2. Enter the last frame you want to render.

The Render End field in the Time Ruler

Once you set the render range, the frames are shaded in the Time Ruler as a visual reference.

3. In the Flow Node Editor, select the tool that displays the images and appearance you want as your final render.

    Tip: You can have multiple Saver tools connected to different parts of the Flow Node Editor to batch render various elements.

4. Choose Tools > I/O > Saver, enter a name and the location to save your rendered file(s), then click OK.
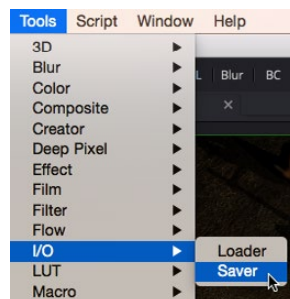
After adding a Saver tool, choose a location and name

Tip: When naming the Saver file, adding a File Extension (e.g., My Comp.mov, My Comp.tif) will automatically configure the Control Panel for your selected file format.

5. With the Saver tool selected in the Flow Node Editor, check the Control Panel to ensure the correct settings.
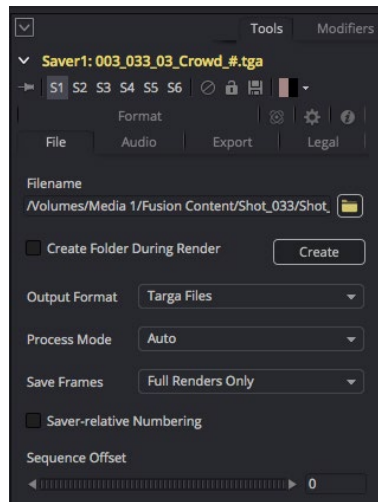
The Control Panel for the Saver tool

6. In the Time Ruler, click the Render button, then click Start Render in the render settings window.
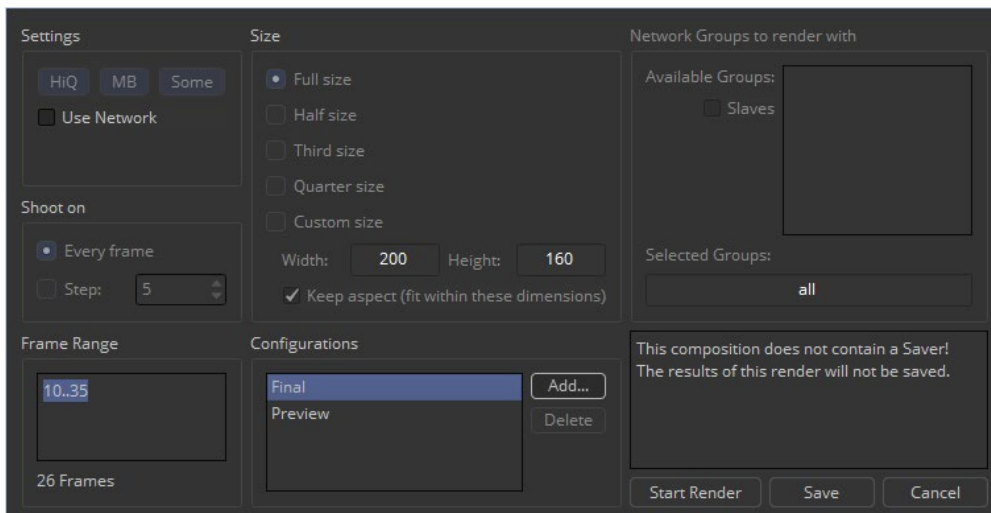
The Render Settings window

Your comp will begin rendering. The Time Ruler and Viewer will step frame by frame through the comp as it renders and a completion dialog will appear when the rendering is done.

Preparing Projects and Images

2

Preparing Projects and Images

Overview

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

Setting up a Project

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

The first thing you do when starting on a new project is to set the preferences to match the intended final output format. The preferences are organized into separate groups: one for global preferences and one for the preferences of the currently opened compositions.

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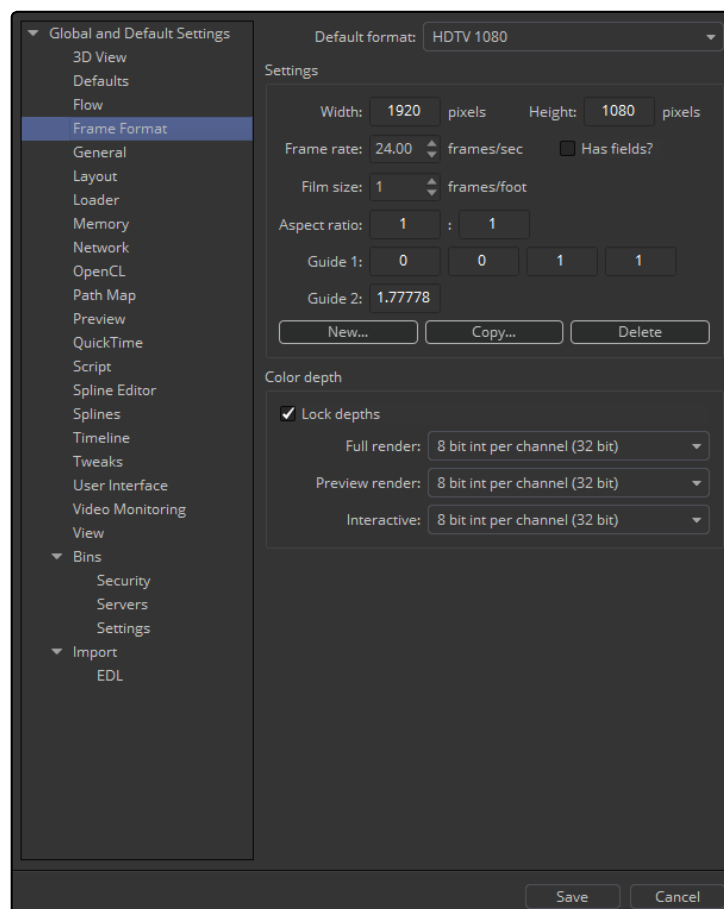
Frame Format preference for images generated in Fusion

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 File > Preferences.

2. Click the Global and new Comp disclosure triangle in the sidebar to open the Globals group.

3. Select the Frame Format category to display its options.

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

The preferences are described in greater detail in Chapter 16.

Loading Images

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

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

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

→ Choose Tools > I/O > Loader.

→ Right click over the Flow Node editor, then choose Add Tool > I/O > Loader.

→ Drag a file from a Windows Explorer file browser or Mac OS X Finder window into the Flow Node Editor.

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

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

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

Color Bit Depths

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

Generally, 8-bit integer color processing is the lowest bit depth you'll come across for video formats. 8-bit images come from older or consumer grade video equipment. If you try to perform any significant gamma or color correction on 8-bit images, you can often see more visible banding.

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

This sounds perfectly fine until you realize that many digital cameras from Blackmagic and others record in formats that can capture over-range values with shadow areas below 0 and specular highlights above 1.

The 16-bit float color depth sacrifices a small amount of the precision from standard 16-bit integer color depth to allow storage of color values less than 0 and greater than 1. 16-bit float, sometimes called half-float, is most often found in the OpenEXR format, and contains more than enough dynamic range for most film and television purposes, yet requires significantly less memory and processing time than is required for full float, 32-bit images.
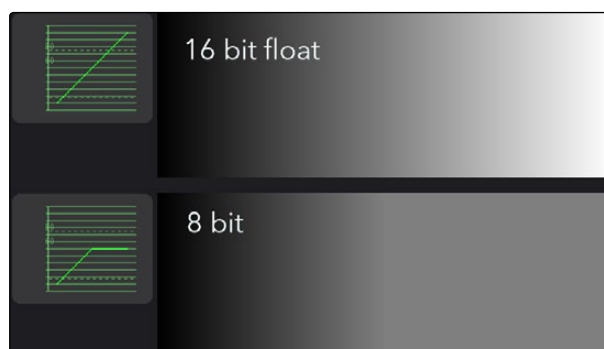
Preserving over-range values allows you to change exposure while maintaining highlights

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

Setting Color Depth

Fusion automatically uses the color depth that makes the most sense for the file format of an image. For example, if you read in a TGA file from disk, then the color depth for the Loader will be set to 8-bits per channel. The TGA format is an 8-bit format, so loading the image at a greater color depth would generally be wasteful. If a 16-bit TIFF is loaded, the color depth will be 16-bits. Loading a Cineon or DPX file defaults to 32-bit float, whereas OpenEXR generally defaults to 16-bit float. However, you can override the automatic format color depth using the settings found in the Import tab of the tool's Control Panel. The Loader's Control Panel, as well as the Control Panel for images generated in Fusion (i.e., text, gradients, fractals and others), have buttons for 8-bit, 16-bit integer, 16-bit float and 32-bit float.

The Loader's Control Panel Color Bit Depth settings

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

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

To improve performance as you work on your comp, you can set the Interactive and Preview depth to 8-bits per channel, while final renders can be set to 16-bit integer. However, if your final render output will be 16-bit float or 32-bit float, you should not use the integer options for the interactive setting. The final results may look significantly different from interactive previews that are set to integer options.

The Frame Format Color Bit Depth settings

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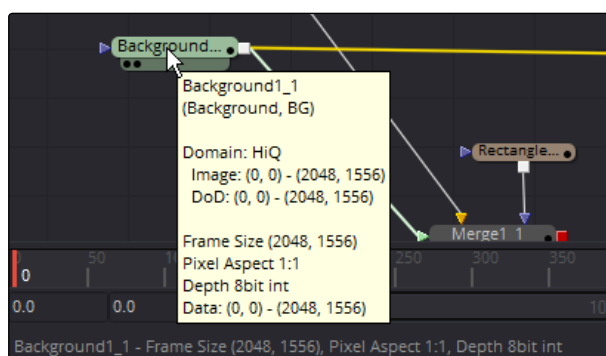
Hover over a tool to view its Color Bit Depth setting

Tip: When working with images that use 10-bit dynamic range or greater, like Blackmagic Cinema DNG files, set the bit depth in the Control tab to 16-bit float or 32-bit float. This will preserve highlight detail as you composite.

Combining Images with Different Color Depths

You can combine images with different color depths in a single composition. When images of different color depths are combined, the image from the foreground input of the tool will be adjusted to match the color depth of the background.

Advantages of Floating Point Processing

There are two major advantages to floating point processing that make the additional RAM requirements and longer render times worth your while. The first benefit is that floating point values are more accurate than integer. The second benefit is the preservation of shadow and highlight values that go beyond the normal tonal range.

Greater Accuracy

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

Imagine that the gain of that pixel is halved using a Color Correction tool. Now the pixel's red value is half of 75, or 37.5. You cannot, however, store decimal or fractional values in 8-bits, so you must round that value to 37. Now, double the brightness of the pixel with another Color Correction tool. The pixel's red value is 37 multiplied by 2, or 74. We lost a full value of precision due to integer

rounding on a very simple example. This is a problem that can result in visible banding over several color corrections. Similar problems come up when merging images together, or transforming them. The more operations that are applied to an image, the more color precision is lost to rounding.

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

Because floating point processing allows decimal or fractional values for each pixel, it is not required to round off the values of the pixel to the closest integer. As a result, color precision remains virtually perfect, regardless of how many operations are applied to an image.

Accessing Extended Highlights and Shadows

More and more productions are capturing out of range images thanks to cost effective cinema cameras like the Blackmagic URSA. These cameras capture very high dynamic range RAW images and maintain color detail even in heavily over or under exposed frames. The extended white color detail can also give very nice, natural results when blurred, glowed, color corrected, or even just when faded or dissolved. While it is possible to work with these RAW images using integer data, to do that you must clip all the extended range values, losing all detail in the highlights and shadows. Float processing makes working with logarithmic RAW images considerably easier by preserving highlight and shadow detail.

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

When processing floating point colors, pixel values greater than white or darker than black are maintained. There is no value clipping. The pixel is still shown in the display view as pure red, but if float processing is used instead of 8-bit, the second operation where the gain was halved would have restored the pixel to its original value of 200.

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

Detecting Extended Highlight and Shadow Values

Although floating point processing preserves extended values below 0 and greater than 1, also called out of range values, the viewer still displays them as black or white. This can make it difficult for you to determine the overall dynamic range of an image.

To discover if there are out of range values in a viewed image do the following:

→ Click on the Show Full Color Range button in the Viewer toolbar.
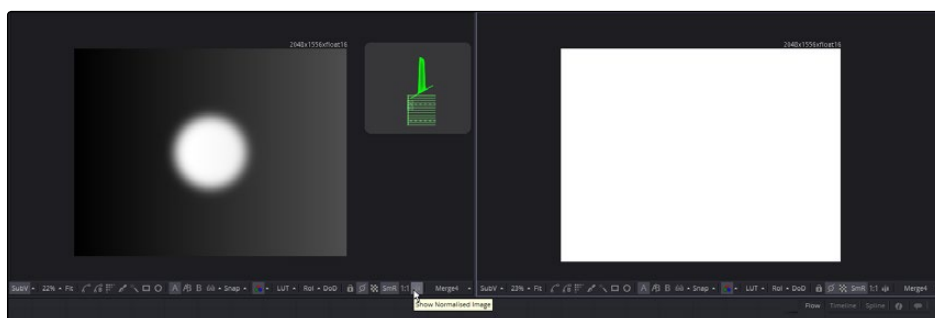
Use the Show Full Color Range button to detect out of range images

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

The 3D Histogram view type can also be helpful in visualizing out of range colors in an image. See Chapter 5, Viewers, for details.

Clipping Out of Range Values

When processing in floating point, there may be situations where the out of range values in an image need to be clipped. The Brightness/Contrast tool provides checkboxes that can be used to clip out of range values to 0 or 1.

For example, there may be files that contain out of range alpha values. Since the alpha channel represents the opacity of a pixel, it makes little sense to be more than completely transparent or more than fully opaque, and compositing such an image may lead to unexpected results. To easily clip alpha values below 0 and above 1, add a Brightness/Contrast tool set to Clip Back and Clip White, with only the Alpha checkbox selected.

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

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

Converting images to Linear Color Space

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

A nonlinear, logarithmic color conversion

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

A linear color conversion

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

To convert an image into linear color space

1. In the Flow Node Editor, select the Loader for the image

2. In the Control Panel, click the Import tab

3. Click the disclosure triangle for the Source Gamma Space

4. Enable the Remove Curve check box.
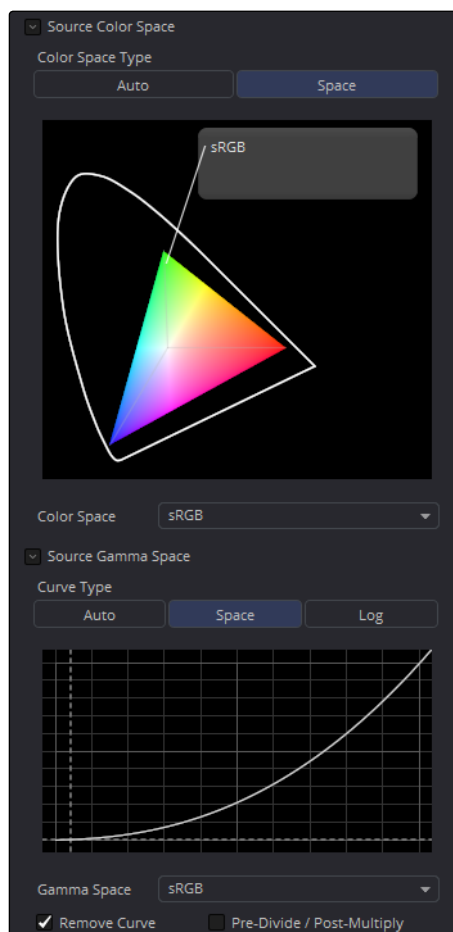
The Source Gamma Space settings in a Loader

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

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

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

To apply a Gamut LUT to the Viewer do the following

1. Click the Viewer LUT pop up menu under the Viewer.

2. Select Gamut View LUT from the menu.

3. Click the LUT button under the Viewer to enable the Gamut View LUT

4. Click the View LUT pop up menu again and choose Edit.

5. Choose sRGB as the Output Space
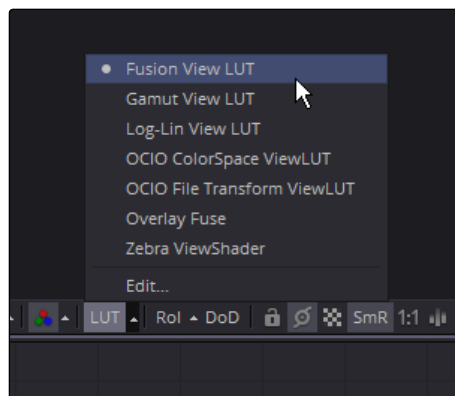
The Viewer LUT pop-up menu

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

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

To Save the Gamut LUT set up as the default Viewer set up do the following:

→ Right click in the Viewer, then choose Settings > Save Defaults.

You now have a linear color space working environment for the Viewer and one Loader. When you are ready to render you must output the images back to the appropriate color space for the file type you are saving.

To set the output render color space:

1. Connect a Save tool to the last tool in your comp

2. With the Saver tool selected, click the Export tab in the Control Panel

3. Click the disclosure arrow to expand the Output Gamma Space parameters.

4. Click the Space button, then from the Gamma Space menu choose the appropriate color space for your file type.

5. Click the Apply Curve check box too apply the color space setting.
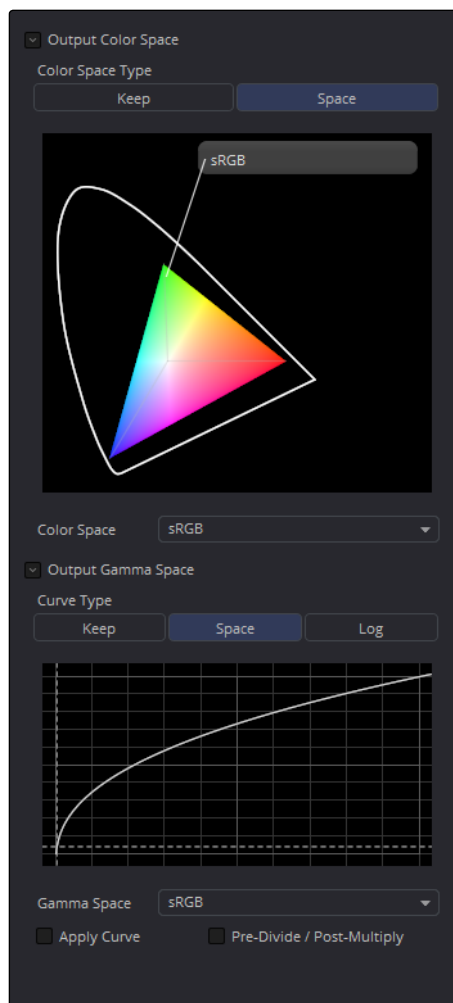
The Saver tool's Gamma Space settings

Working with Blackmagic Cinema DNG Images

CinemaDNG is an open format capable of high resolution raw image data with a wide dynamic range, and is one of the formats recorded by Blackmagic cameras when you shoot in raw mode.

The BMD setting decodes the Cinema DNG file into a log encoded color space that remaps the raw data into an approximation of the Log–C standard. The Control panel for a Cinema DNG Loader will contain some processing controls for the RAW conversion.

To convert an Cinema DNG file into linear color space

1. In the Flow Node Editor, select the Loader for the image

2. In the Control Panel, click the Format tab

3. Click the disclosure triangle for the Source Gamma Space

4. Enable the Remove curve check box.
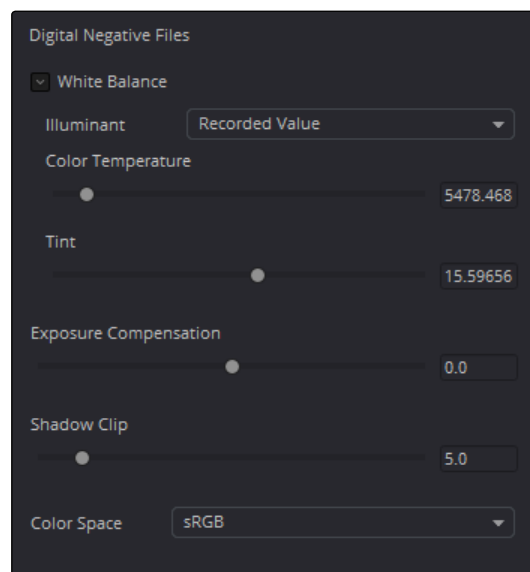
The Control Panel for a Cinema DNG file

Flow Node Editor

3

Flow Node Editor

Flow Node Editor Overview

The Flow Node Editor with connected tools

What is Node Based Compositing?

The Flow Node Editor is the heart of Fusion's compositing interface. It uses a schematic or flow chart structure to build a composition instead of layers in a layer list. Each image and effect that you add to your composition is called a Tool, which is represented as a tile in the Flow Node Editor. You draw connections between tools to arrange the order that image processing operations are applied. Just by viewing the connections between tools, you can get a clear overview of the process each image goes through to be composited.

Using a Schematic view, sometimes called a Tree view, has tremendous flexibility for compositing, including the ability to branch a tool's output into separate branches and eliminate the need to precompose or nest layers.

This flow shows a Loader to read the image, connected to a Resize tool, connected to a Saver to render the image to disk

This chapter discusses how to work in the Flow Node Editor, add and connect tools together, and other basic operations you can perform on tools in the Flow Node Editor.

Navigating the Flow Node Editor

The Flow Node Editor is one of the default panels whenever you create a new composition or when an existing comp is opened. The Work Panel that contains the Flow Node Editor can be switched to display the Spline Editor, Timeline Editor and Console, as needed. In order for you to work effectively in the Flow Node Editor, you need to be able to pan and zoom in on areas of a composition.

To display the flow when it is not currently visible, do one of the following:

→ Select the Flow tab from the top right corner of the work area.

→ Press the F5 key on the keyboard.

The tabs in the Work Panel

To zoom in and out on an area of a flow, do the following:

→ Hold the left and middle mouse buttons, then drag to zoom in and out.

To pan in the Flow Node Editor, do the following:

→ Hold the middle mouse button, then drag to pan.

→ Drag the scrollbars at the right and left sides of the Flow Node Editor panel.

→ Press Command-Shift (Mac OS X) or Ctrl-Shift (Windows), then drag to pan.

If a tool that is not visible in the Flow Node Editor becomes selected, either by performing a Find or selecting the header in the Control Panel, the Flow Node Editor will automatically pan to display the tool's tile in the visible area.

The Flow Navigator

Another useful way to pan around the Flow Node Editor is to use the Flow Navigator. The Flow Navigator is a small rectangular overview in the upper right corner of the Flow 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 Flow Navigator when you are zoomed in on a flow and want to pan around a composition.
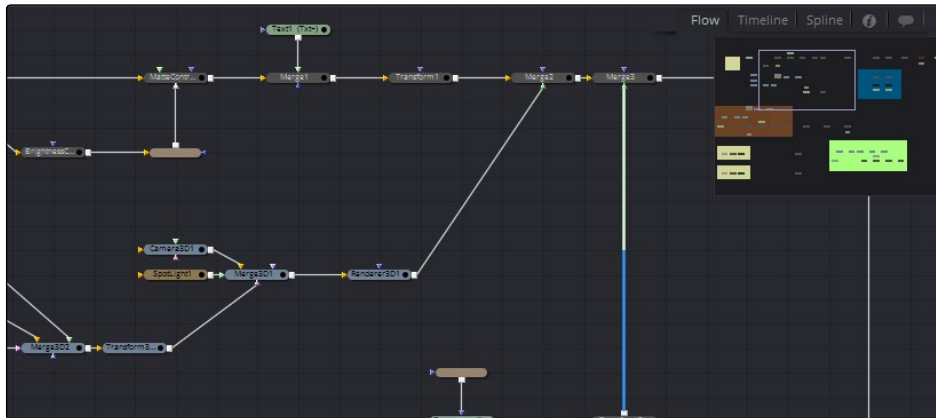
The Flow Navigator

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

→ Right click in an empty area of the Flow Node Editor, then choose Options > Show Navigator.

→ Press the V key.

To have the Flow Navigator display automatically when needed:

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

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

→ Right click in the Flow Navigator, then choose Reset Size.

To pan the flow using the Flow Navigator, do the following:

→ Drag in the Flow Navigator to change what portion of the flow is visible.

Flow Node Editor Tool Tiles

When you want to load images from disk, save images to disk, merge multiple images together, and apply effects like blurs, glows or color corrections, you must add a tool to the Flow Node Editor.

Once a tool is added to the Flow Node Editor it is represented by a Tool tile.
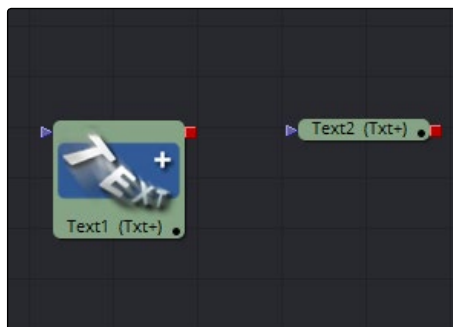
A Tool tile in the Flow Node Editor

The tiles can be displayed as a small rectangle or as a larger square. The rectangular form displays the tool's name in the center of the tile, while the square form shows either the tool's icon or a thumbnail of the image it is outputting. You can get detailed information about the tool and the data it is processing by hovering your cursor over it in the Flow Node Editor.

Adding Tools to the Flow Node Editor

You can add tools to the Flow Node Editor in a number of different ways depending on what is the easiest place for you to access them and the type of tool you are adding. The most obvious way to add a tool is to use the Tools menu in the menu bar since it contains every tool in Fusion, organized into general categories. When you select a tool from any submenu in the Tools menu it is immediately added to the Flow Node Editor. If another tool is selected in the flow, the selected tool's output is automatically connected to the new tool's input. If no tool is selected in the flow, the new tool is added, unconnected to any other tool. The Tools menu is the most obvious way, but not necessarily the fastest or most productive way, to add a tool to the flow. There are a number of other ways that, once you learn them, may better fit your workflow.

The Tools menu

Adding Tools from the Toolbar

If you find yourself constantly using the same tools over and over in different comps, you can use the Toolbar, located above the Viewers, to display tools of your choice for easy access. The tools are displayed as buttons. Clicking on a button adds that tool to the composition exactly as if the tool had been selected from the Tools menu. You can also drag the button into the flow to position the tool manually. Dropping the dragged tool over the connection between tools will automatically insert the tool into the connection.

The Toolbar

The buttons use an abbreviated form of each tool's name (e.g., BC for Brightness/Contrast), but if you can't remember what some of the abbreviations stand for, you can position the cursor over any toolbar button and a tooltip will display the full name of that tool.

Toolbar buttons can also be dragged into the Viewers, causing them to be added to the flow immediately after the currently viewed tool, regardless of the currently active tool.

In some cases you may want to add a tool from the Toolbar in order to replace an existing tool in the flow. For instance, if you have a Brightness/Contrast tool already connected in the Flow Node Editor, but you want to use the Color Curves instead, you can remove the Brightness Contrast and replace it with the Color Curves in one step.

To replace a tool in the flow with a tool from the Toolbar, do the following:

1. Drag a tool from the Toolbar directly over the tool in the Flow Node Editor that you want replaced.

2. When the mouse button is released, click OK in the dialog to confirm the replacement.
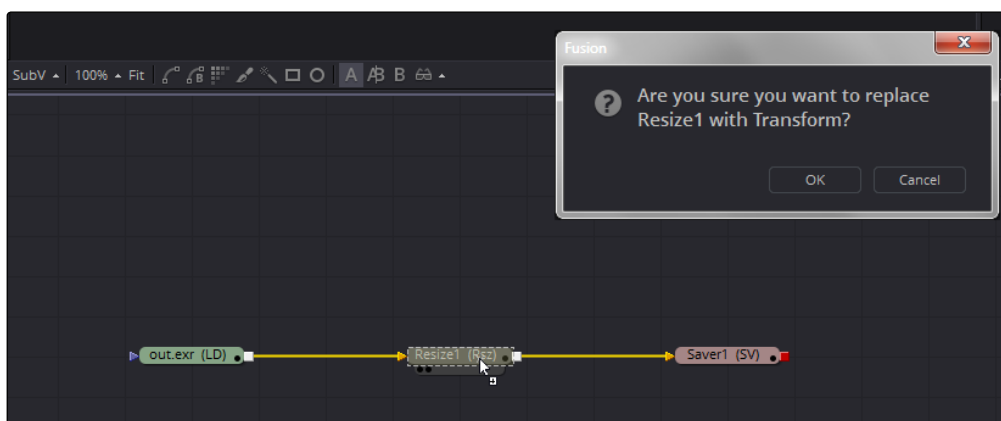
Dragging a tool from the Toolbar to replace an existing tool.

When you replace one tool with another, the settings that are identical between the two tools are carried over. For example, replacing a Transform tool with a Merge will copy the existing center and angle values from the Transform to the Merge.

For details on customizing the toolbars, refer to the Toolbar chapter later in this manual.

Adding Tools from the Contextual Menu

There are different submenus you can use to add tools to a flow when you right click over a tool, over an empty area of the Flow Node Editor, or over a connection line also called a pipe. This can be a very convenient way to add a tool since your cursor is often in the Flow Node Editor selecting, moving and connecting tools.
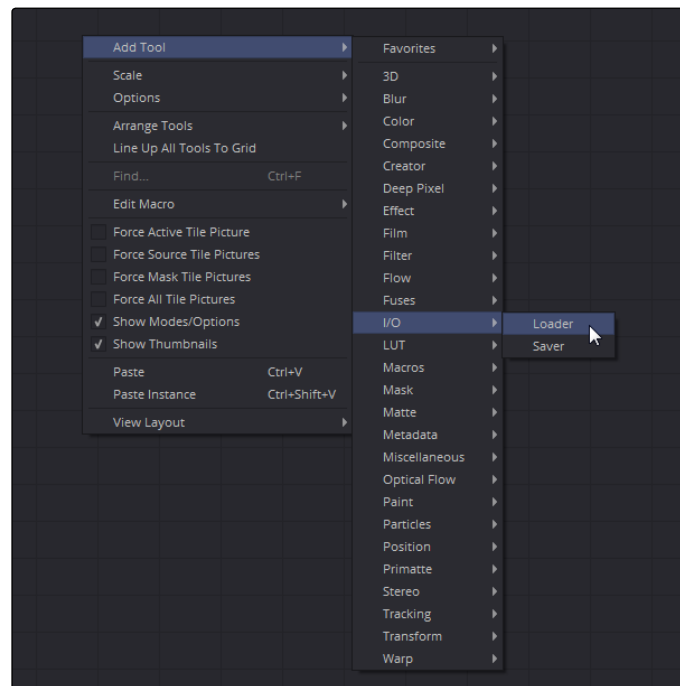
Adding tools from the contextual menu

To Add a tool, do the following:

1. Right click in an empty area of the Flow Node Editor.

2. Choose a tool from the Add Tool submenu.

Selecting a tool from the Add Tool submenu will add the tool at the position of the cursor. If the cursor is positioned over a connection between two tools, the new tool is inserted between the existing tools.

When you right click over a tool in the Flow Node Editor, you have a choice to either Insert a Tool or Replace a Tool.

To Insert a Tool, do the following:

1. Right click over a tool in the Flow Node Editor.

2. Choose a tool from the Insert submenu.

Selecting a tool from the Insert submenu will insert the new tool after the selected tool.

To Replace a Tool, do the following:

1. Right click over a tool in the Flow Node Editor.

2. Choose a tool from the Replace submenu.
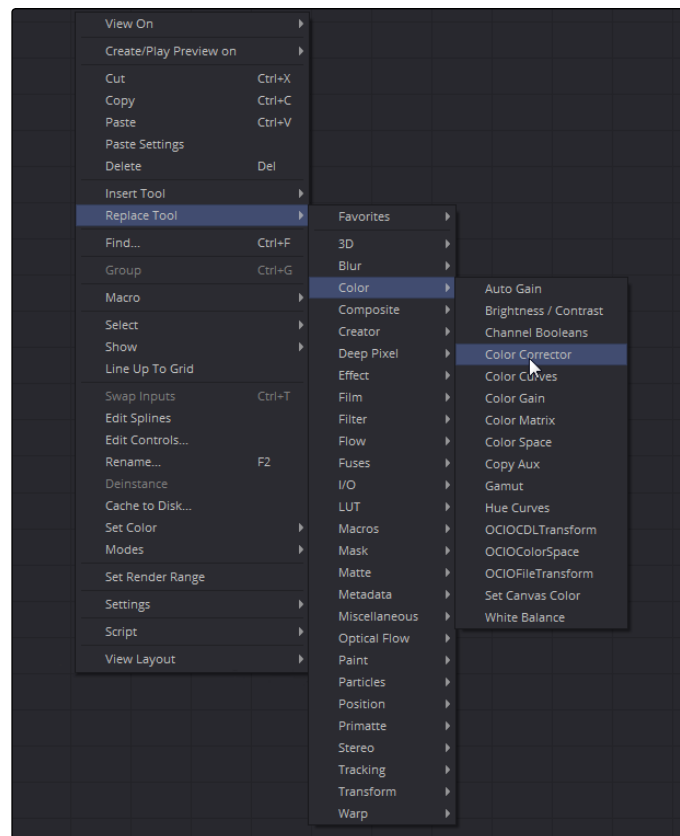
Replacing tools from the contextual menu

Selecting a tool from the Replace menu will replace the selected tool with a new one.

If the new tool has controls with the same name as the tool you are replacing, those settings will be retained. For example, replacing a Transform tool with a DVE will preserve the values of the center, angle and size controls.

Adding Tools from the Bins

The Bins are an extremely useful tool in Fusion that we cover in more detail in the Bins chapter of this manual. However, in terms of adding tools, it is as simple as drag and drop.

You open a bin by choosing File > Bins or pressing Command-B (Mac OS X) or Ctrl-B (Windows).

The bin contains different categories of items you can add to a flow. By selecting the Tools category you can drag and drop a tool directly into the flow.

Tools in a bin

Adding Tools from a File Browser

If you drag a clip or image sequence from an Mac OS X or Windows file browser into the Flow Node Editor, it automatically adds a Loader and reads in the file. If multiple clips or image sequences are dragged into the flow, a separate loader is added for each file and they are merged together.
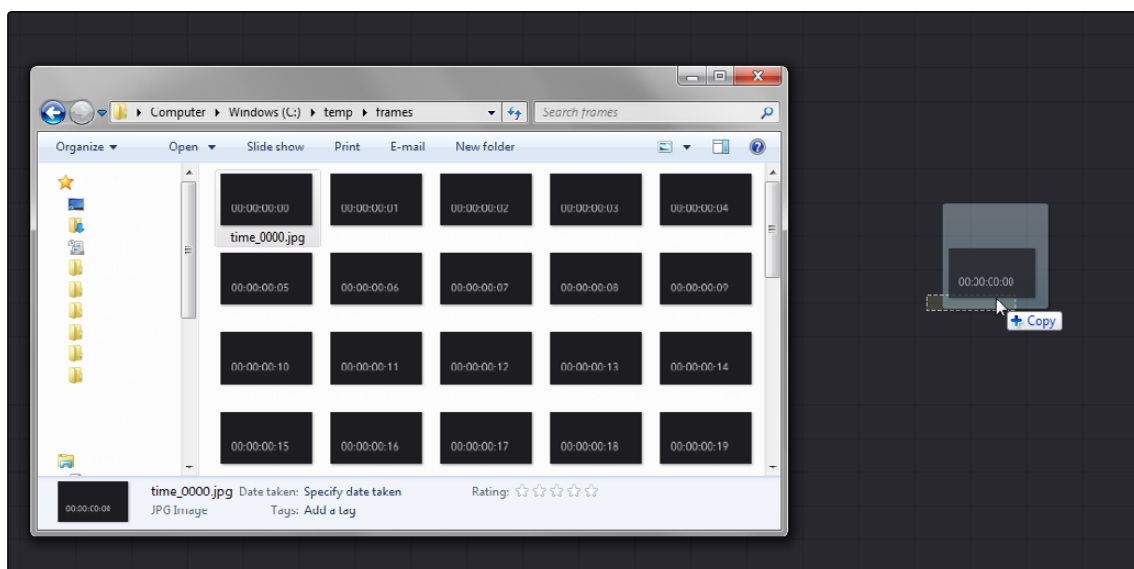
Dragging from a file browser

You can choose to drag in any file from an image sequence and the entire sequence will be read in as a single clip. However, to read in a single frame from an image sequence, hold Shift and drag the frame you want into the Flow Node Editor. That frame will be read in as a still. This can be useful when loading a number of images taken with a digital still camera.

Tool settings (.settings) files can also be dragged from the file browser into the Flow Node Editor. Dragging a settings file will add the tool saved in that file and configured it according to the saved settings.

Deleting Tools

If you want to delete a tool from the Flow Node Editor, select the tool, then press Delete (Mac OS X) or Backspace (Windows) on the keyboard. The tool is removed from the Flow Node Editor. As an alternative, you can right click on a tool and then select Delete from the context menu.

Viewing Tools

There are a number of different ways to display a tool in a viewer, starting with the Tool Tile View indicators. The View indicators are displayed under each tile. By default, there are two indicators, representing the two Viewers in the Fusion window. The View indicator on the left is used for the Viewer on the left, while the View indicator on the right is used for the Viewer on the right. If a new viewer is added, an additional indicator will appear on each tile. Clicking an indicator will turn it white and display that tool in the appropriate viewer. Clicking it again removes it from the Viewer and removes the white highlight.

A Viewer indicator enabled for the left Viewer a
nd disabled for the right Viewer

Tip: If a tool is not currently displaying View indicators, position the cursor over the tile to make the indicators visible and available for clicking.

You can use keyboard short cuts to display the selected tool in a viewer. Each viewer is assigned a numeric keyboard shortcut from 1-9. Pressing that number on the keyboard will display the selected tool in the appropriate display view.

To view a tool using keyboard short cuts, do one of the following:

→ Select a tool in the Flow Node Editor, then press 1 to display it in the left Viewer.

→ Select a tool in the Flow Node Editor, then press 2 to display it in the right Viewer.

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

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

Drag and Drop into a Viewer

If the View indicators are too small of a target for you to click on reliably and you are not a keyboard oriented person, another way to display a tool in a viewer is to drag it there. This offers a quick explicit way to assign a tool to a viewer. As you drag, the tool will appear to move at first, but will snap back into its original location once the cursor leaves the area of the Flow Node Editor. After you release the mouse button over the Viewer, the tool will be displayed.
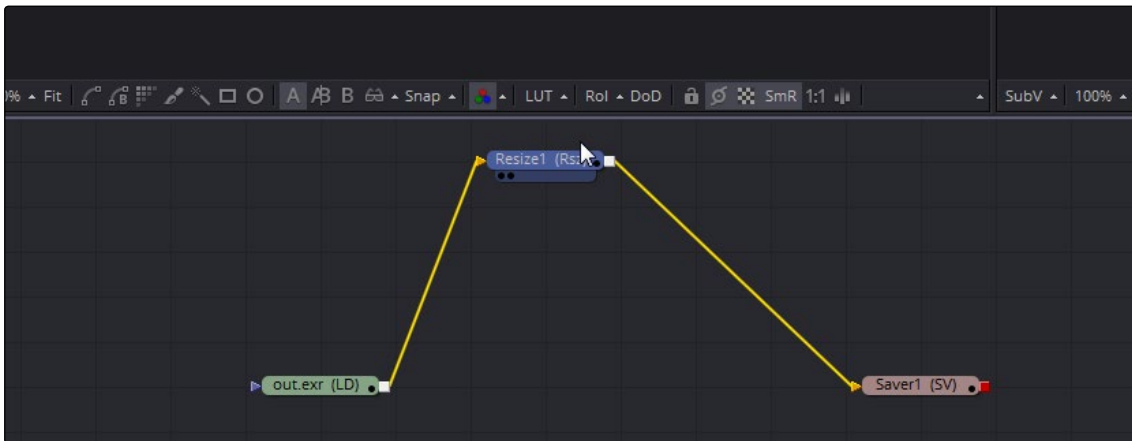
Dragging a tool into a viewer

The last way you can view a tool in a viewer is to use the tool's contextual menu.

To view a tool using the tool's contextual menu, do one of the following:

→ Right click over a tool in the Flow Node Editor, then choose View On > Left to display the tool on the left Viewer.

→ Right click over a tool in the Flow Node Editor, then choose View On > Right to display the tool on the right Viewer.
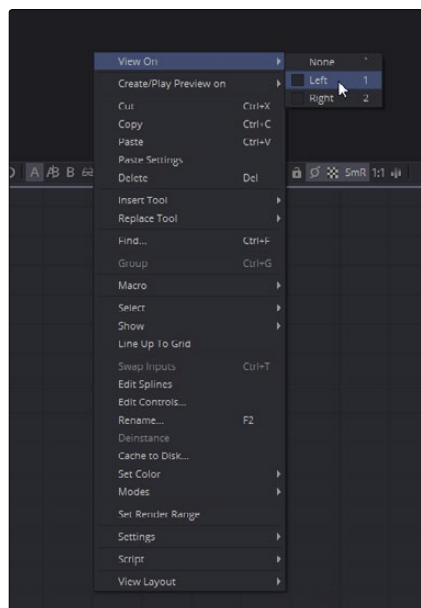
Choosing a viewer from the contextual menu

Connecting Tools in the Flow

Once you have added tools to the Flow Node Editor, you must connect them together in order for them to perform any operation. Each Tool tile displays small colored knots around the edges of the tile. The arrow shaped knots represent one or more inputs, and the grey or red square knots represent the tool's processed output, of which there is always only one.

Blur tool with Input, Output and Mask knot

By connecting outputs to inputs you begin to build a flow from one tool to the next. For example, to connect an image to a Color Corrector tool, you click on the output knot on the Loader tool and drag a connection to the input knot of the Color Corrector. A thin line, called a pipe, will be drawn between these two tools to represent that connection.

Connecting a pipe from an output to an input

The order in which you draw the connection pipe is not important; you can just as easily drag in the reverse direction from an input knot to an output knot and get the same results.

Whichever way you draw the connection, once the connection is made, multi-channel and single channel image data, 3D scene and material data, or particle information is carried from one tool to the other. The exact type of data required by each input and produced by an output is described in the Tool Reference documentation for the specific tool.

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

TIp: Rather than remembering the different knot types press the right mouse button, hold Option and Alt and drag from the output of a tool 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.

Upstream and Downstream Tools

Since tools can be positioned anywhere in the Flow Node Editor, you often refer to tools as being upstream and downstream of each other. If you have a selected tool, any other tools that directly or indirectly connect to the input of the selected tool are considered to be positioned upstream. Any other tools that directly or indirectly connect to the output of the selected tool are said to be positioned downstream. This is an important distinction to make because, unlike layer based systems, the visual positioning of the file or effect in the Flow Node Editor has no bearing on the composite's outcome. Only whether tools are upstream or downstream of each other impacts the composite.
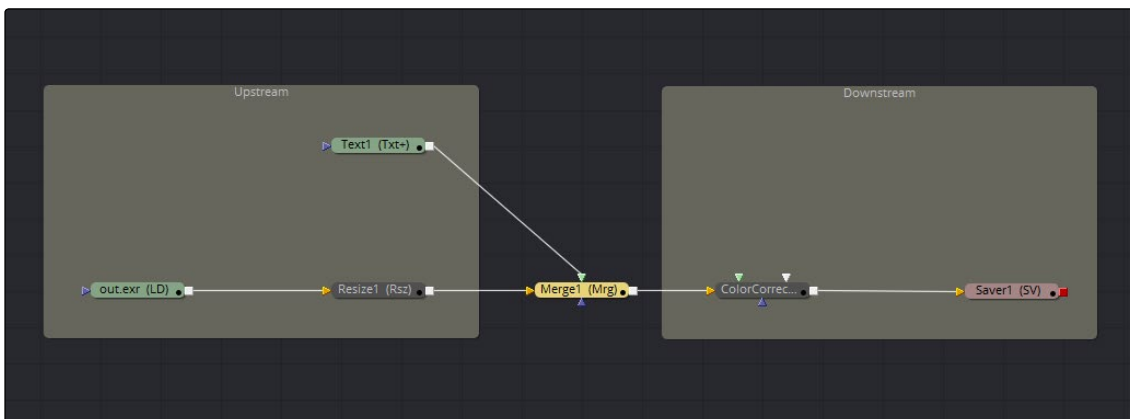
Tools upstream and downstream of the selected Merge tool

Disconnecting Tools

As often as you connect tools, you need to be able to disconnect them as well. This allows you to rearrange the order of tools in a flow. The pipe that connects two tools is divided into two halves: the output half and the input half.

To disconnect the input or output of a tool, do one of the following:

→ Click on the input half of the connection between two tools.

→ Click on the input arrow, then drag to pull the connection away from the tool.

Click the pipe nearest to the input to disconnect it

To disconnect both the input and output of a tool, do the following:

→ Hold Shift and drag the tool away from the connections.

You don't have to disconnect an input only to replace it with the output from another tool.

If you drag a connection to an input that already has a connection, you will simultaneously disconnect the tool and reconnect it with the output you are dragging.

Tracing Pipes Through the Flow

Several visual cues exist to help you follow the pipes through the flow.

The cursor positioned over a tool, highlights the connected pipes

Positioning the cursor over a specific tool will cause the pipes connected to that tool to become highlighted.

Branching

A tool's input can only have one connection from one output, but a tool's output can be connected to inputs on multiple tools. Splitting the output of a tool so it connects to inputs on multiple tools is called Branching. This allows you to have one image that is processed in different ways and then combined back together later on in the flow.
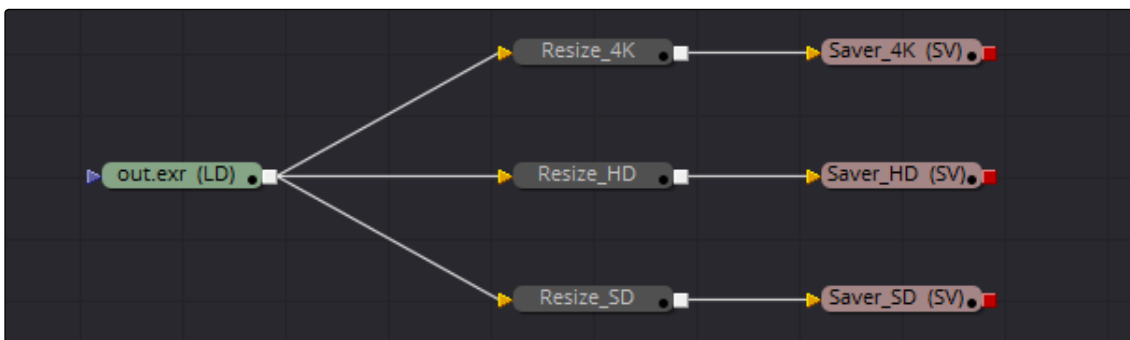
A Loader tool branched to three output

Connection Types

You can change the flexibly of the pipes used to connect tools together. The default uses linear lines that are drawn straight between two knots. This can cause pipes to overlap tiles, which some people may feel interferes with the view of the flow.
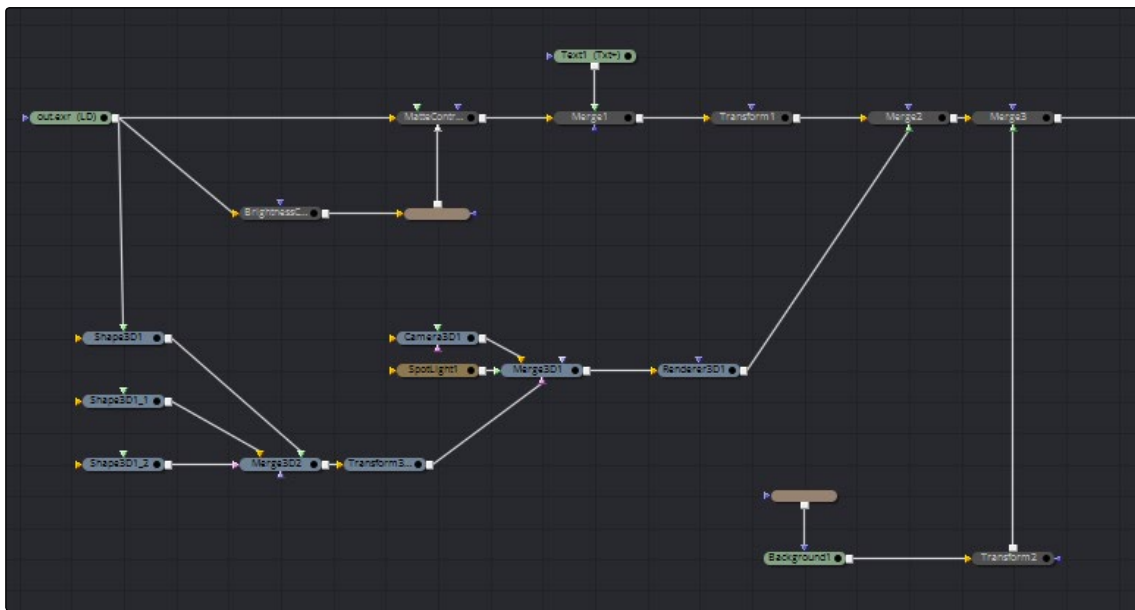
The default linear pipes style

You can switch the linear connection style with one that allows you to used right angles to avoid overlapping tiles.

To use connection pipes with right angles, do the following:

1. Right click in an empty area of the Flow Node Editor

2. From the contextual menu choose Options > Orthogonal Pipes
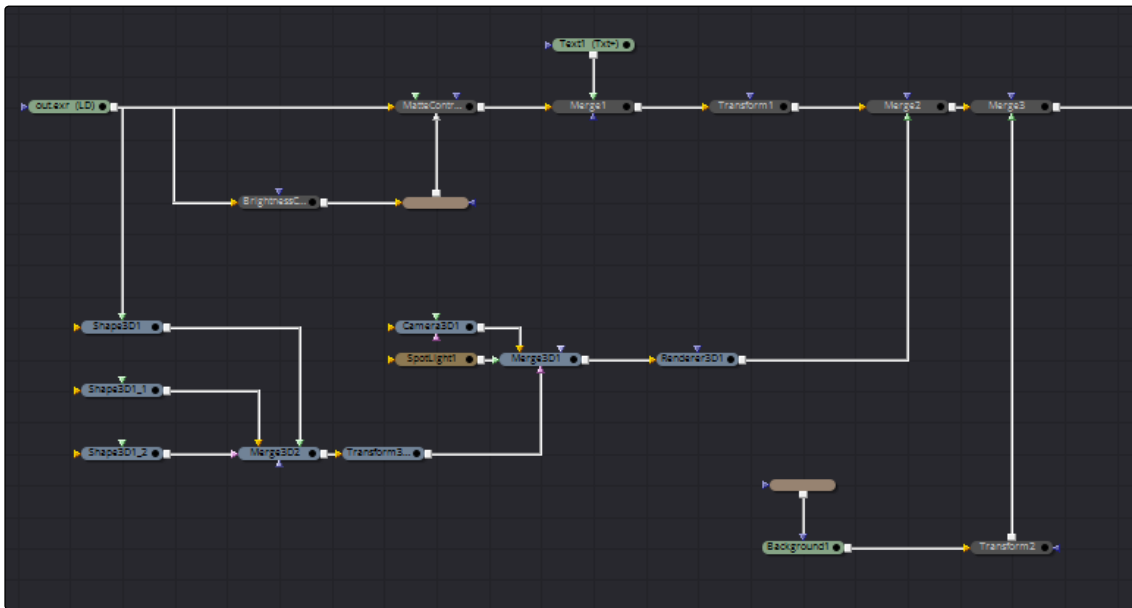
The Orthogonal pipes style

The connections now use right angles to connect one tool to the next. You can even add your own routers to a default straight pipe or orthogonal pipe to force a right angle and route it however you like.

To add a router to an Orthogonal connection pipe, do the following:

→ Option-click (Mac OS X) or Alt-Click (Windows) on a connection pipe where you want to add a router.
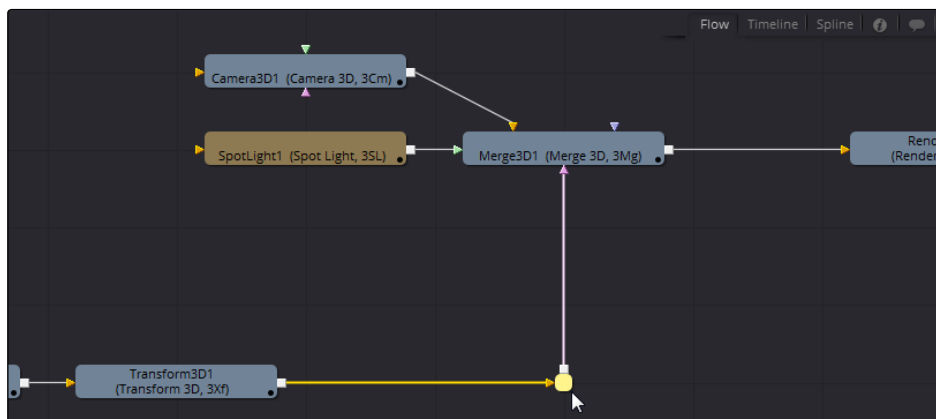
A router added to a pipe

To return to Linear connection pipes, do the following:

1   Right click in an empty area of the Flow Node Editor.

2   From the contextual menu choose Options > Direct Pipes.

Any routers you've added will remain on those connections.

To remove a router, do one of the following:

→ Select the router you want to delete, then press Command-Delete (Mac OS X) or Delete (Windows) on the keyboard.

→ Right click over the router you want to delete, then choose Delete from the contextual menu.

Inserting Tools

When building a comp, you don't always build it by tagging tools onto other tools consecutively. You often add them in a different order than the one they end up in for the final comp. That's why being able to insert a tool between two connected tools can be done in a few different ways.

To insert a tool from the Tools menu, do the following:

1. Select the tool in the flow that the new tool will come after.
2. Select the tool from the Tools menus.

The new tool will be inserted immediately after the selected tool.

To insert a tool from the Toolbar, do the following:

1. Select the tool in the flow that the new tool will come after.
2. Click a tool in the Toolbar button.

The new tool is inserted immediately after the selected tool.

To insert a tool from a Bin, do the following:

1. Drag a tool from a bin and position the cursor over the connection where you want the tool inserted.
2. When the connection pipe is highlighted, release the mouse button.

This will cause the input of the inserted tool to be connected to the upstream tool and the output of the inserted tool to be connected to the downstream tool.

To insert a tool that already exists in the flow, do the following:

1. Hold Shift and drag the tool over the connection pipe.
2. Once the pipe highlights, release the mouse button.
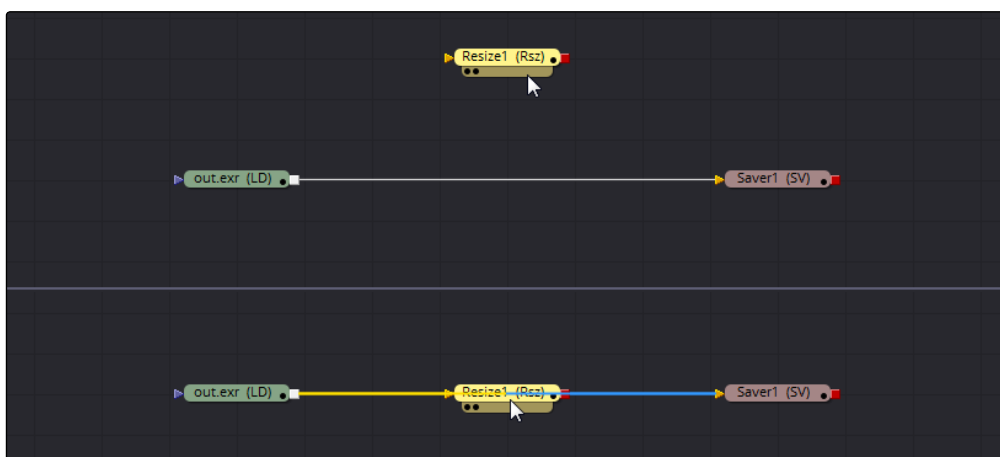
Hold Shift to insert a tool

The new tool is inserted between the tools previously connected by the pipe.

Organizing the Flow

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

Selecting Tools

In order to move, copy, view or modify a tool's parameters you first need to select the tool. Selecting a tool can be done by clicking directly on a tool's tile, on the title bar of the Tool's Control window, or even by clicking on the tool's name in the Spline and Timeline Editors. When you select a tool using any of these methods, the selected tool is known as the Active tool and is highlighted bright yellow. Only one tool can be active at a time, although multiple tools can be selected by dragging a rectangle around a group of tools you want. The tools within the selection group will highlight blue. You can then add or remove tools from the selection group by pressing Command (Mac OS X) or Ctrl (Windows) while clicking Tool tiles.
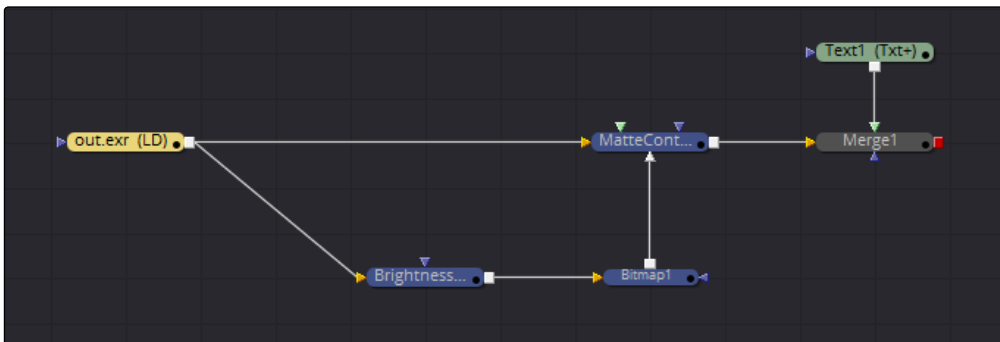
Four tools are selected, one of which is active in yellow and two other tools are not selected

Moving Tools

Moving tools is another basic function you'll use everyday when using Fusion. You will move tools around the Flow Node Editor to better arrange them for connecting and to improve visual organization. Dragging a tile or multiple tiles moves them to a new location. The order or location of tiles in the Flow Node Editor does not impact the results of your composite; it is purely for your own tool management.

Cut, Copy and Paste Tools

Standard text editing operations like cut, copy and paste are available for tools in the Flow Node Editor. You can use them to temporarily remove tools from the Flow Node Editor, copy tools with the same parameter settings and paste copied tools in different locations.

To Copy a selected tool in the Flow Node Editor, do one of the following:

→ Right click over the tool and choose Copy from the contextual menu.

→ Choose Edit > Copy from the Edit menu.

→ Click the Copy icon in the Toolbar.

→ Press Command-C (Mac OS X) or Ctrl-C (Windows).

View On	▶
Create/Play Preview on	▶
Cut	Ctrl+X
Copy	Ctrl+C
Paste	Ctrl+V
Paste Settings	
Delete	Del
Insert Tool	▶
Replace Tool	▶
Find...	Ctrl+F
Group	Ctrl+G
Macro	▶
Select	▶
Show	▶
Line Up To Grid	
Swap Inputs	Ctrl+T
Edit Splines	
Edit Controls...	
Rename...	F2
Deinstance	
Cache to Disk...	
Set Color	▶
Modes	▶
Set Render Range	
Settings	▶
Script	▶
View Layout	▶
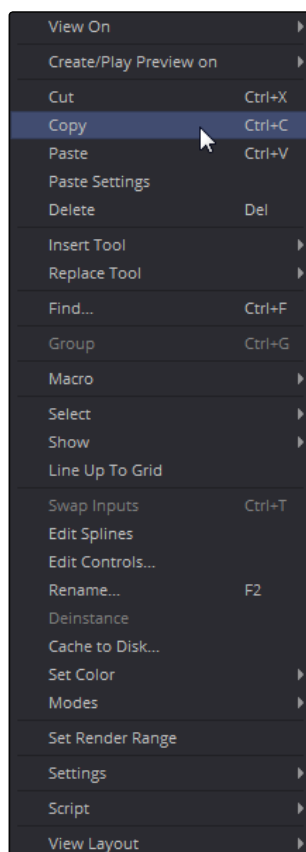
The Copy command in the tool's contextual menu

To Cut a selected tool in the Flow Node Editor, do one of the following:

→ Right click over the tool and choose Cut from the contextual menu.

→ Choose Edit > Cut from the Edit menu.

→ Click the Cut icon in the Toolbar.

→ Press Command-X (Mac OS X) or Ctrl-X (Windows).

When you paste a tool into the Flow Node Editor, a copy of the last tool(s) cut or copied into the clipboard is placed in the Flow Node Editor.

To Paste a Cut or Copied tool in the Flow Node Editor, do one of the following:

→ Right click in the Flow Node Editor and choose Paste from the contextual menu.

→ Right click over a tool and choose Paste from the Tool context menu.

→ Choose Edit > Paste from the Edit menu.

→ Click the Paste icon in the Toolbar.

→ Press Command-V (Mac OS X) or Ctrl-V (Windows).
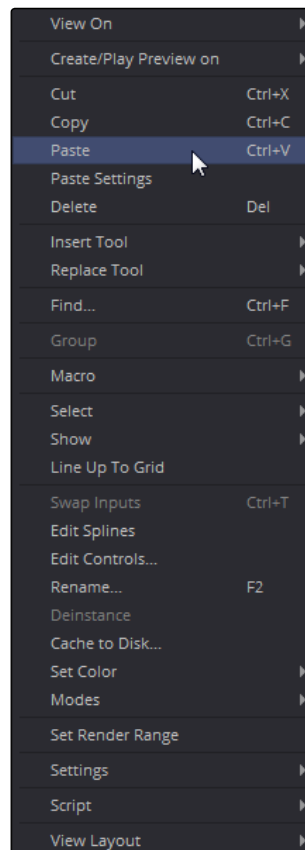
The Paste command in the tool's contextual menu

The location of the pasted tools in the Flow Node Editor varies a lot, depending on how the tools are pasted and whether any tools are selected at the time. The best way to control where the tools appear is to either click in the flow and then use the Command or Ctrl-V, or to use the contextual menus. In either case, the tools will be placed at the location of the pointer when the mouse is clicked.

If you want to replace a tool that is already in the flow with the contents of the clipboard, right click directly on the tool you want replaced and then choose Paste from the tool's contextual menu. You can also choose to insert a tool from the clipboard between two other tools. Right click directly on the pipe between the two tools and select Paste from the context menu.

When you paste a Loader or Creator (a tool that does not have any inputs), the pasted tool automatically gets composited with the active tool using a Merge tool. This can save you a few extra steps, but some effects artists and designers may prefer to perform the merges manually, so this behavior can be changed using the Default Preferences panel in the Global preferences.

Pasting Tool Settings

Instead of pasting a tool, you can choose to paste just the parameter settings for a tool. This can be useful if you've set parameters for one image like the center point in a Merge tool and want to copy that position onto another image that uses a Transform.
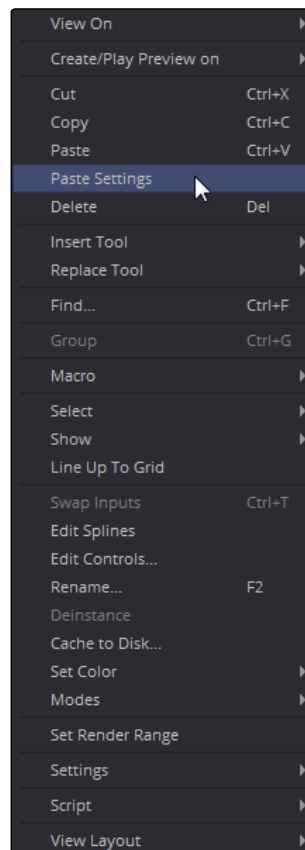
The Paste Settings command in the tool's contextual menu

To Paste settings from one tool to another, do the following:

1. Select the tool that has the correct settings and choose Copy from the Edit menu or press Command-C or Ctrl-C.

2. Right click over the destination tool, then choose Paste Settings from the contextual menu.

When copying the settings from one type of tool to another another type of tool, only the settings that match between two tools will be copied. For example, a common usage is to match animated centers in a Transform to the centers in a Mask tool.

There are common parameters on all Image tools, like Blend and which channel is being processed. Tools of similar classes also have more matching parameters.

Copying and Pasting from a Text Editor

When we covered the Fusion Composition format in Chapter 1, it was explained that the format is not binary but a simple text format. The implications of that may not have been obvious, but one benefit example is clear when you start dealing with tools. Tools can be copied from the flow and pasted directly into a text editor. This will display the tools' parameters in text format just as it would have been saved to disk. You can then edit the text and paste it back into the flow, either as a new tool, or to replace the settings of an existing tool. This is a very easy way to pass specific Tool settings back and forth between artists who may not be in the same room, city or country.

Renaming Tools

Fusion automatically assigns a name and a number to each tool added to a composition. For example, the first Blur tool added to a composition will be called Blur1, the second one will be Blur2 and so on. Although initially helpful, on larger compositions you may decide you need a more descriptive name to make it easier to identify similar tools.

You can rename a tool by right clicking over the tool and choosing Rename from the contextual menu or by pressing the F2 key. If multiple tools are selected, a dialog will appear asking for a name for each tool.

The Tool Rename dialog

Since Fusion can use scripting languages, the names of tools have some limitations. When you rename a tool, only use alphanumeric characters and do not use any spaces. Also, you cannot start a tool name with a number. If you accidentally create a name that doesn't exactly follow the guidelines, Fusion will automatically replace spaces with _ and strip away invalid characters.

Lining Up Tools

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

To have all the tools centered on the nearest horizontal and vertical grid line, do the following:

1. Right click in an empty section of the flow to display the flow's contextual menu.

2. Choose Line Up All Tools to Grid.
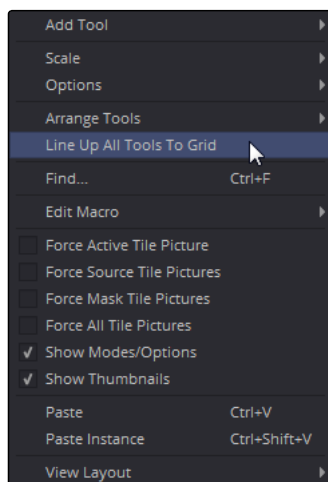
The Line Up All Tools to Grid menu command

All the tools in the Flow Node Editor will shift up or down and left or right to align their centers along a grid line up.

To have only selected tools centered on the nearest horizontal and vertical grid line, do the following:

1. Right click over a selected tool to display the tool's contextual menu.

2. Choose Line Up to Grid.

The tools that are not currently selected will be left as they are.

Snapping Tools to the Grid

As you move tools in the Flow Node Editor, you can have them snap to the grid, making it much easier to align tools manually. Enabling the Arrange Tools > To Grid option in the flow's context menu will enable this feature.

To have tools snap to the grid lines as you drag them, do the following:

1. Right click over an empty area of the Flow Node Editor.

2. Choose Arrange Tools > To Grid from the contextual menu.

Tip: You can set the Arrange Tool > To Grid option as the default for new compositions by setting the Preferences > General > Arrange To Grid checkbox to On.

Groups

When you work on complex visual effects, flows can become sprawling and unwieldy, so grouping items together can help you better organize all the tiles and pipes. Groups are containers on the flow that can hold multiple tools, similar to the way a folder on your Desktop holds multiple files. There is no limit to the number of tools that can be contained within a group and you can even create sub groups within a group.

To create a Group, do the following:

1. Select the tools you want grouped together.

2. Right click over one of the selected tools and choose Group from the contextual menu or press Command-G (Mac OS X) or Ctrl-G (Windows).
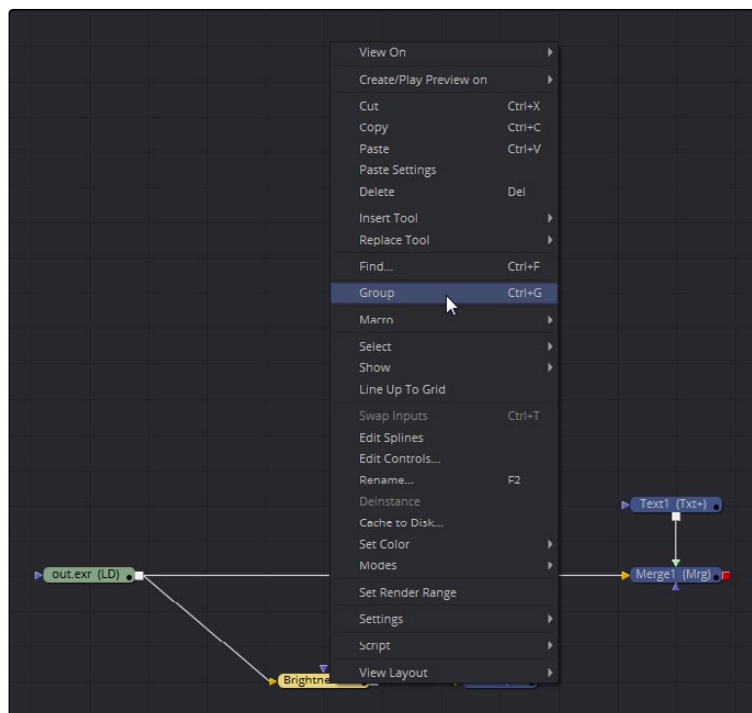
Selecting the Group menu command

The selected tools are collapsed into the group, which is displayed as a single tile in the Flow Node Editor. The Group tile can have inputs and outputs, depending on the connections of the tools within the group. The Group tile only displays inputs for tools that are already connected to tools outside the group. Unconnected inputs inside the group will not have an Input knot displayed on the Group tile.

Deleting a group is no different from deleting any other tile in the Flow Node Editor. Press Command-Delete (Mac OS X) or Delete (Windows) and Group, and the tools contained within that group are removed from the flow.

Expanding and Collapsing Groups

A collapsed group is represented by one tile on the flow. If you want to modify any of the tools inside the group, you can open it by clicking the small Group icon to the right of the Group's name in the tile or pressing Command-E (Mac OS X) or Ctrl-E (Windows).

The Group Icon for expanding the group

When you open a group, a subflow that contains the tools within the group is displayed. The subflow is a separate area that can be scaled, resized and panned independently of the main flow. You can adjust any tool and even add, insert and delete tools in the group while it is expanded. When you are ready to collapse the group again, click on the minimize icon at the top right corner of the subflow or use the keyboard short cut.
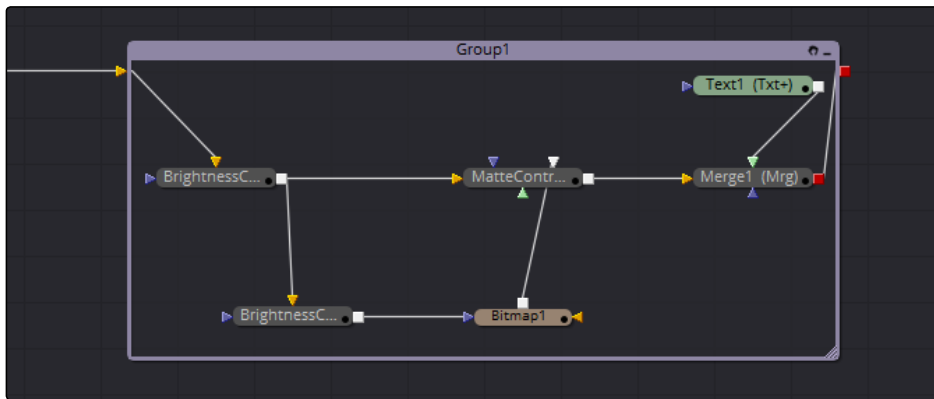
The subflow of an expanded group

Panning and Scaling in the Subflow

You can Pan and Scale the subflow using the same mouse buttons you use when panning and scaling the main flow. During sessions when you are working with an expanded group and simultaneously making changes in the main flow, you may want to prevent the expanded group from being panned or scaled. Clicking the small Hand icon in the title bar of the group attempting to pan and scale in the subflow, will pan and scale the main flow. Clicking the Hand icon again will enable the subflow to pan and scale.

The Hand icon is used to lock the
subflow from panning and scaling

Ungrouping Tools

If you no longer have the need for the group, or just find it easier to have constant access to all the tools in the group, you can remove the group without deleting the tools in it.

To Ungroup tools, do the following:

1. Right click on the group.

2. Choose Ungroup from the context menu.
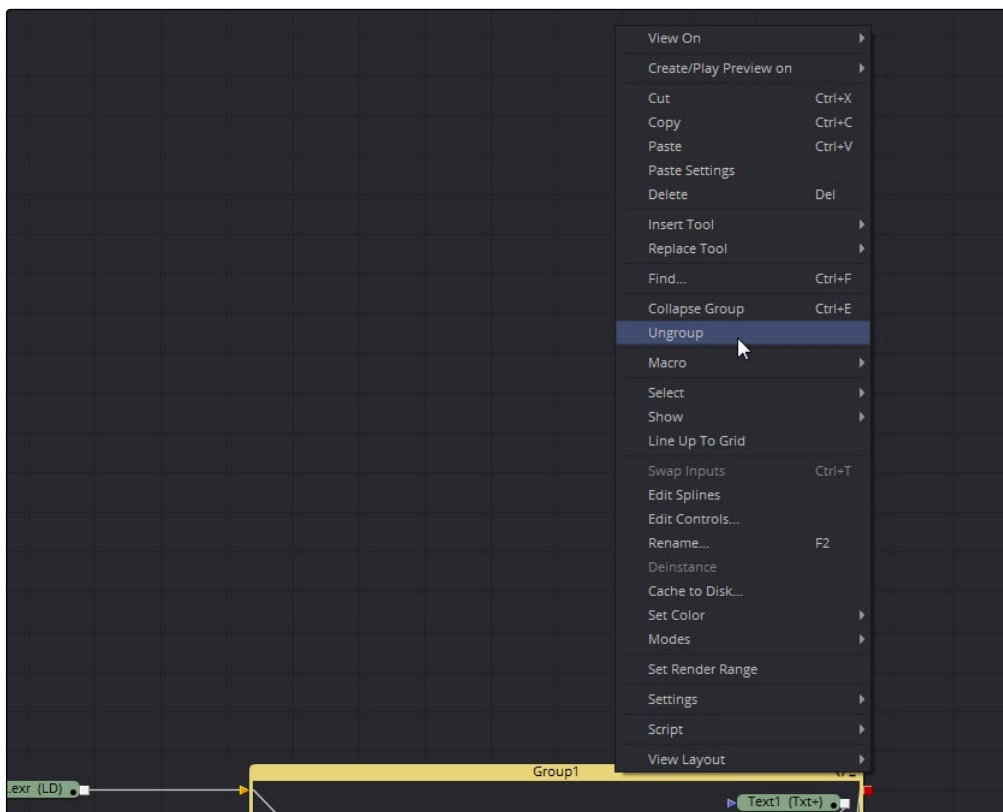
The Ungroup menu command

The tools inside the group are placed back in the main flow.

Saving and Reusing Groups

One of the best features of a group is that the group and its settings can be saved for later use in other projects. Groups can either be saved by dragging the group from the Flow Node Editor into a bin or by right clicking the group and selecting Settings > Save As from the contextual menu.

To reuse the group, you can drag it from a bin or from your computer's file browser directly into the Flow Node Editor. In this way, you will create a new tile on the flow that has all the same tools inside it. Alternatively, if you already have a group with the same tools in it, you can load the Settings by right clicking the group in the Flow Node Editor and selecting Settings > Load from the contextual menu.

An example for where you might want to Save and Load groups settings is in a studio with two or more artists. A lead artist in your studio can set up the master comp and create a group specifically for keying green screen. 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.

Instancing Tools

There are times when two tools need to have the exact same settings at all times. For example, when you are making identical color corrections to two or more images, you don't want to adjust one color correction tool and then manually adjust the other color correction tools to match. It's a hassle to do that once, but a considerable waste of time if you are constantly adjusting the color correction parameters over the course of the project. Normally, when a tool is copied and pasted into the flow, changes made to the original tool are not made to the pasted tools. There are ways to publish controls on the original tool and connect them to matching controls on the pasted tools, but this becomes prohibitively complex and again, time consuming on tools with a lot of controls.

Instanced tools share the same controls so that a change to one tool affects all of the tools in the instance.
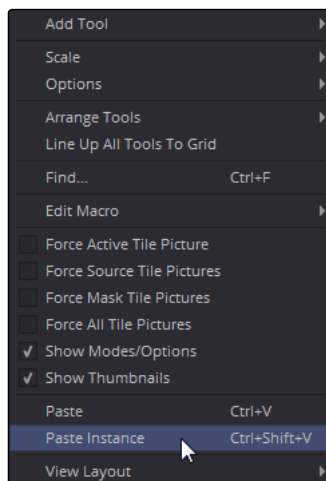
The Instance menu command

To create an Instance, do the following:

1. Right click over a tool you want to Instance.

2. From the contextual menu, choose Copy or press Command-C (Mac OS X) or Ctrl-C (Windows).

3. Right click in an empty area of the Flow Node Editor.

4. From the Contextual menu, choose Paste Instance or press Command-Shift-V (Mac OS X) or Ctrl-Shift-V (Windows).
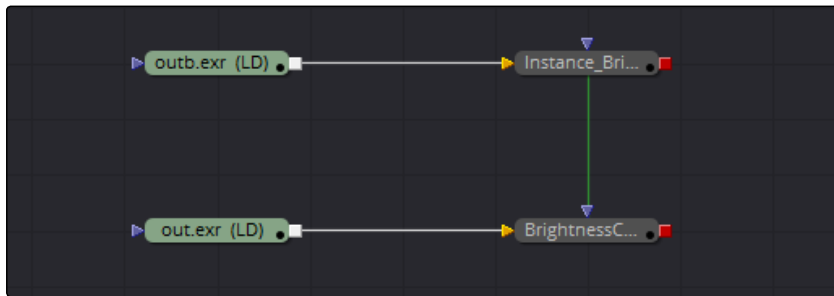
A green Instance pipe

When a flow contains Instanced Tools, a green connection pipe links the original tool with its Instances. In addition, a green outline surrounds the edit boxes in the Control Panel of an instanced tool. You can disable the green connection pipe to reduce clutter in the Flow Node Editor.
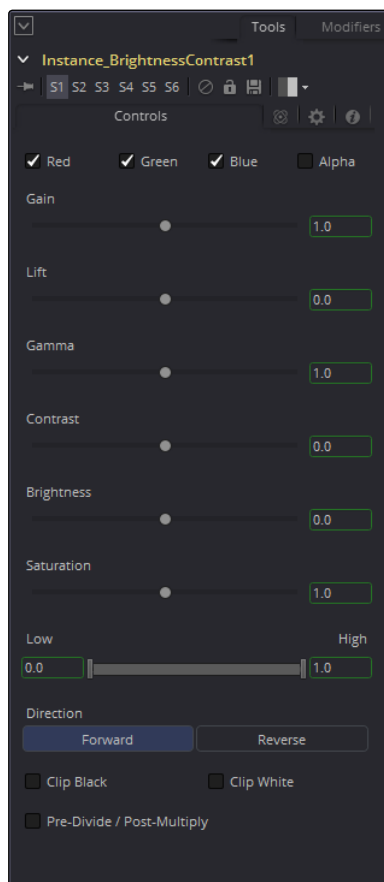
Green outlines in the Control Panel for instanced parameters

To hide green instanced connection pipes in the flow, do the following:

1. Right click in an empty area of the flow.

2. Choose Options > Show Instance Links from the contextual menu.

This submenu acts as a toggle to enable and disable the visibility of connection pipes for Instances.

If you need to remove the actual link between an Instance and the original tool, you can De-instance the tool.

To delete the Instance link from a tool, do the following:

1. Right click on an instanced tool.

2. Choose De-instance from the contextual menu.

Once you de-instance a tool you cannot re-instance it, but you can undo the operation.

If the Undo is gone because you have quit Fusion and reopened it, you can only recreate an instance by copying the original and pasting the instance again.

De-Instancing and Re-Instancing a Control

Instancing comes with the incredible flexibility to enable and disable the link on a parameter by parameter basis. By default the entire Tool is instanced, so any change to any parameter is passed on to the instanced tools. Occasionally, you will come across one parameters that you want to adjust independently from the other instances. Instead of de-instancing the entire tool, you can de-instance a single parameter.

To De-instance a single parameter, do the following:

1. Right click on the parameter's name in the Control Panel.

2. Choose De-instance from the contextual menu.
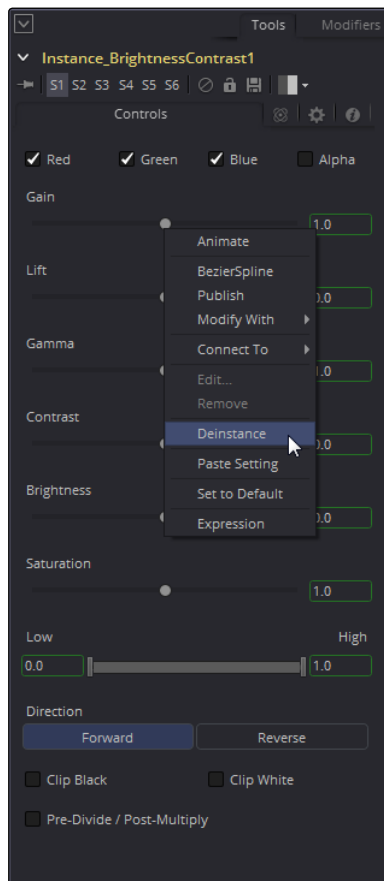
De-instancing a parameter in the Control Panel

The green outline around the parameter's edit box will disappear.

To Re-instance a single parameter do the following:

1. Right click on the parameter's name in the Control Panel.

2. Choose Reinstance from the contextual menu.

Displaying Tile Pictures

An enormously helpful aid in visualizing the operation of a flow is to display the Tool tiles with pictures of thumbnail image or a tool's icon.

To turn on tile pictures for all the tools in the Flow Node Editor, do the following:

→ Right click in an empty area of the Flow Node Editor, then choose Force All Tile Pictures from the contextual menu.
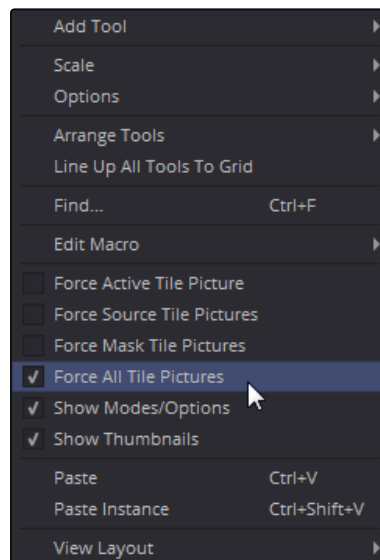
The Force All Tile Pictures menu command

Showing all tile pictures enlarges the tiles and takes more space in the Flow Node Editor. This has a benefit but also forces you to pan and zoom more often. You may decide you only want some tools to show thumbnails based on the type of tool or selection state of a tool. You would first disable the Force All Tile Pictures option in the contextual menu, then select any of the following options.

To turn on tile pictures for one or more tools, do the following:

1. Select the tool(s) in the Flow Node Editor.

2. Right click on one of the selected tools, then choose Show > Show Tile Pictures from the contextual menu.

Choosing Show Tile Pictures will enable pictures for the selected tiles. Those pictures will remain visible even if the tool is selected or not. This menu option is a toggle, selecting the tool(s) again and choosing Show Tile Pictures will hide the pictures.

You can also choose to have tools that are currently selected show pictures. With this option, pictures will temporarily be displayed for as long as the tool is the selected tool. Once you deselect the tool, the pictures will be hidden.

To make only selected tools display a tile picture, do the following:

→ Right click in an empty area of the Flow Node Editor, then choose Force Active Tile Pictures from the contextual menu.

To display tile pictures for tools that generate images (e.g., Loaders, Text, Background), do the following:

→ Right click in an empty area of the Flow Node Editor, then choose Force Source Tile Pictures from the contextual menu.

To display only tile pictures for Mask tools, do the following:

→ Right click in an empty area of the Flow Node Editor, then choose Force Mask Tile Pictures from the contextual menu.

To use Tool icons instead of rendered thumbnails, do the following:

→ Right click in an empty area of the Flow Node Editor, then deselect Show Thumbnails from the contextual menu.

As you add more and more tools to a flow, you'll notice that some tools never display a Tile thumbnail. In these cases, the default icon for that tool is displayed instead of an image.

Tools from the Particle and 3D categories fall into this group. They only show icons, for the most part. The exceptions are the pRender tool in the particles category and the Render 3D tool. These two tools will display a rendered thumbnail if you have the menu options set for Thumbnails to be displayed.

Whether some tools display thumbnails or not is more situational. Some transformation tools are able to concatenate their result, passing the actual processing downstream to another tool later in the flow. These tools do not actually process the image, so they do not produce a thumbnail.

In situations where the Loader is not reading in a clip or the clip is trimmed in the Timeline to be out of range, it can cause the tool not to process the image, so it will not produce a rendered Thumbnail image. Also, tools that have been set to Pass Through mode are disabled and do not display a rendered Thumbnail image.

Finding Tools

Modern visual effects shots require detailed work that can often result in Flow graphs that include several hundred tools. If you had to rely on your eyesight alone to locate a tool, you would be panning around the Flow Node Editor for quite a while. You can quickly locate tools in the flow using the Find dialog.
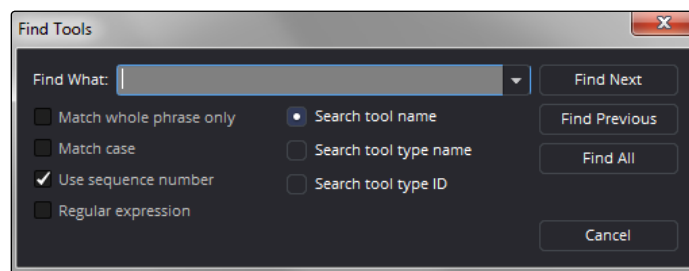
The Find dialog

To search for a tool in the Flow Node Editor, do the following:

1. Right click in an empty area of the Flow Node Editor.

2. Choose Find in the contextual menu or press Command-F (Mac OS X) or Ctrl-F (Windows).

3. In the Find dialog, enter a search term, then click Find Next or Find Previous.

Clicking Find Next or Find Previous locates and selects one tool in the flow that has a name matching the search term. You can click the Find Next or Previous button again to locate another tool that matches the search term. To narrow down your matches, the Find dialog includes several options that can be used to filter the search criteria even further. You can also click the Find All button to select all of the tools in the flow that match the entered name. Finding all the tools of a particular type can be very useful if you want, for example, to disable all Resize tools. Find All will select all of the tools based on the search term and you can temporarily disable them by pressing the short cut for Bypass, Command-P (Mac OS X) or Ctrl-P (Windows).

Tool Settings

When a tool is added to the Flow Node Editor, its parameters are set to default values. These default values are saved to a Defaults folder on your hard drive, based on the the Globals > Path Map > Defaults preference.

If you find yourself always adjusting the parameters of a tool as soon as you add it to a flow, you can override the default settings with your own custom defaults by right clicking on the tool and selecting Settings > Save Default.
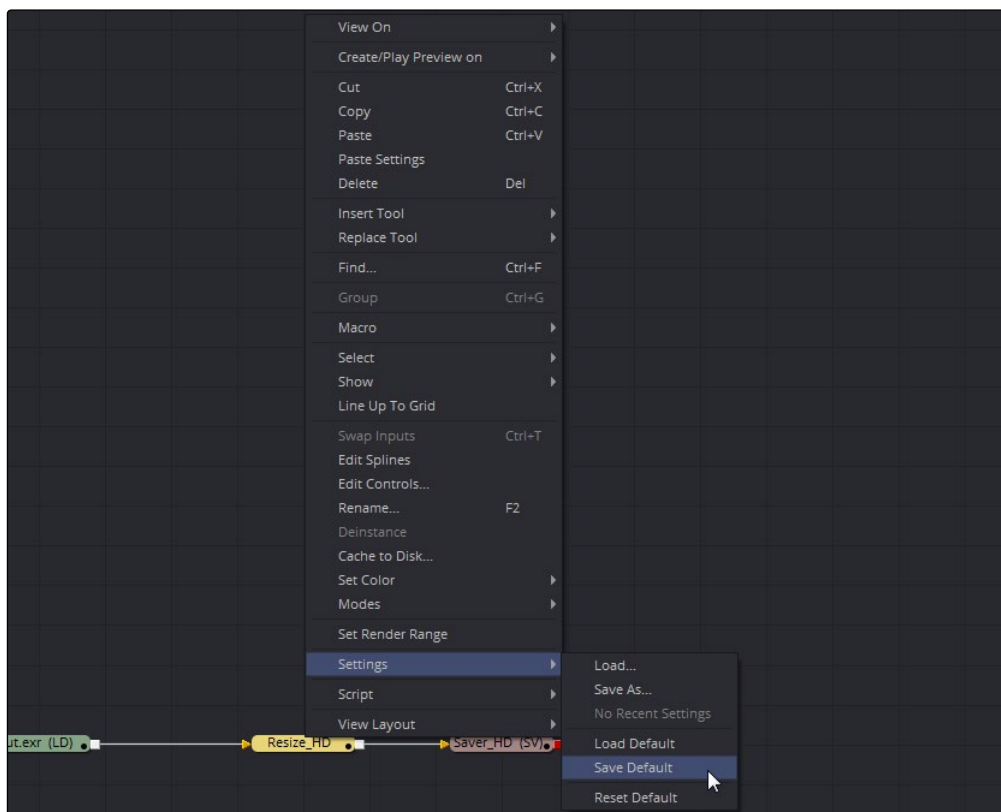
The Save Default submenu

The settings for all tools are saved using a template for the name. The name template is INTERNALNAME_PUBLICNAME.settings. For example, the default setting for a Blur tool would be called Blur_Blur.settings. This is partly to ensure that plug-in tools with identical names to Fusion tools do not overwrite the defaults for Fusion tools.

To remove a saved default setting, do one of the following:

→ Delete the .settings file from the Defaults folder.

→ Right click on the tool and select Settings > Reset Default.

Once you change parameter values for a tool using the Control Panel, those values can be saved as Settings for the tool and reused at a later time. You can load the settings to the toolbars or bins and retrieve them from the Tool context menu.

To save settings for one or more tools, do one of the following:

→ Right click on a tool, then choose Settings > Save As from the contextual menu.

→ Drag the tool into a bin.

A Save File dialog appears, allowing you to name 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.

When you drag to a bin, the setting is still saved based on the Save File dialog, but a shortcut for the saved settings is added into the bin.

To Load a saved setting for one or more tools, do one of the following:

1. Select the Tool you want the settings applied to.

2. Right click over the Tool and choose Settings > Load from the tool's contextual menu.

Saved settings can also be used to create new tools by dragging the settings file into the flow from the toolbars, bins or a standard file browser.

Custom Bin Icon for Saved Settings

You can make it easier to identify different versions of settings that are saved to a bin by creating your own custom icons.

To create a custom icon for a saved setting, do the following:

1. Create a 32-bit bitmap (.bmp) with exactly the same name as the .settings file.

2. Save the .bmp file to the same folder where you saved the setting.

The bitmap will be used as the icon for the file in the bin. Although you can create the bitmap at any size, it is recommended you keep it small for best performance. The largest Bin icon size should be created at 390x300 pixels.

Flow and Tools Options

You can access two contextual menus in the Flow Node Editor. One contextual menu is for the Flow and one is for the Tools in the flow. Both of these contextual menus hold options for customizing the appearance and operation of the flow and tools.

Flow Options

Right clicking in an empty area of the Flow 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 flow looks and behaves.
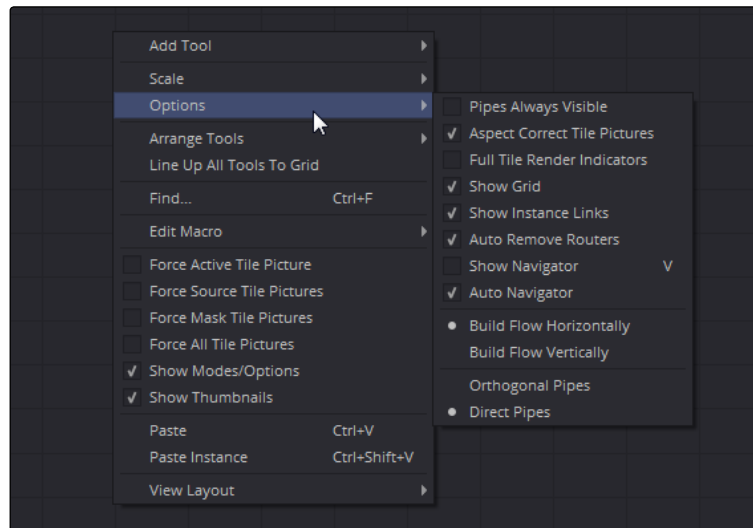
The Flow Options contextual submenu

Options > Pipes Always Visible

Enabling this option causes a pipe to cross over a tool instead of beneath it, sometimes making it easier to follow the pipe's path.

Options > Aspect Correct Tile Pictures

Aspect Correct Tile Pictures forces the display of tile pictures to be aspect corrected, which is slower but visually more accurate. This option is enabled by default.

Options > Full Tile Render Indicators

Enabling this option causes the tile to flash green when rendering, which makes it easier to identify which tool is processing in a large, complex flow.

Options > Show Grid

This option can be used to enable or disable the flow's background grid.

Options > Show Instance Links

When enabled, the flow draws a green pipe between an Instanced tool and its parent.

Options > Auto Remove Routers

If Routers are disconnected from a tool, they are automatically deleted from the flow. This option is enabled by default to eliminate the need to delete orphaned routers.

Options > Show Navigator

Enabling this option displays a small overview window of the entire flow in the Flow Node Editor's top right corner. For more information see the Navigator section in this chapter.

Options > Auto Navigator

The Navigator will automatically appear when any tool is outside the visible flow area and go off when all tools are within the flow area. For more information see the Navigator section in this chapter.

Options > Build Flow Vertically/Horizontally

Flows can either be built horizontally from left to right or vertically from top to bottom. Enabling one of these options determines whether new tools are added beneath the current tool or to the right of the current tool.

Options > Orthogonal/Direct Pipes

Use these two options to decide if pipes between tools are drawn as Direct (straight) lines or Orthogonal (bent) lines.

Tool Contextual Menu

Right clicking on a tool displays that tool's contextual menu. The options available from this menu are described below, unless they have already been discussed in detail earlier in the chapter.

View On

This option provides a submenu of available viewers or external video monitors. Choosing Viewer determines where the tool's output will be displayed.

Create/Play Preview On

You can play a preview of the tool's output by choosing one of the available viewers from the submenu. 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.
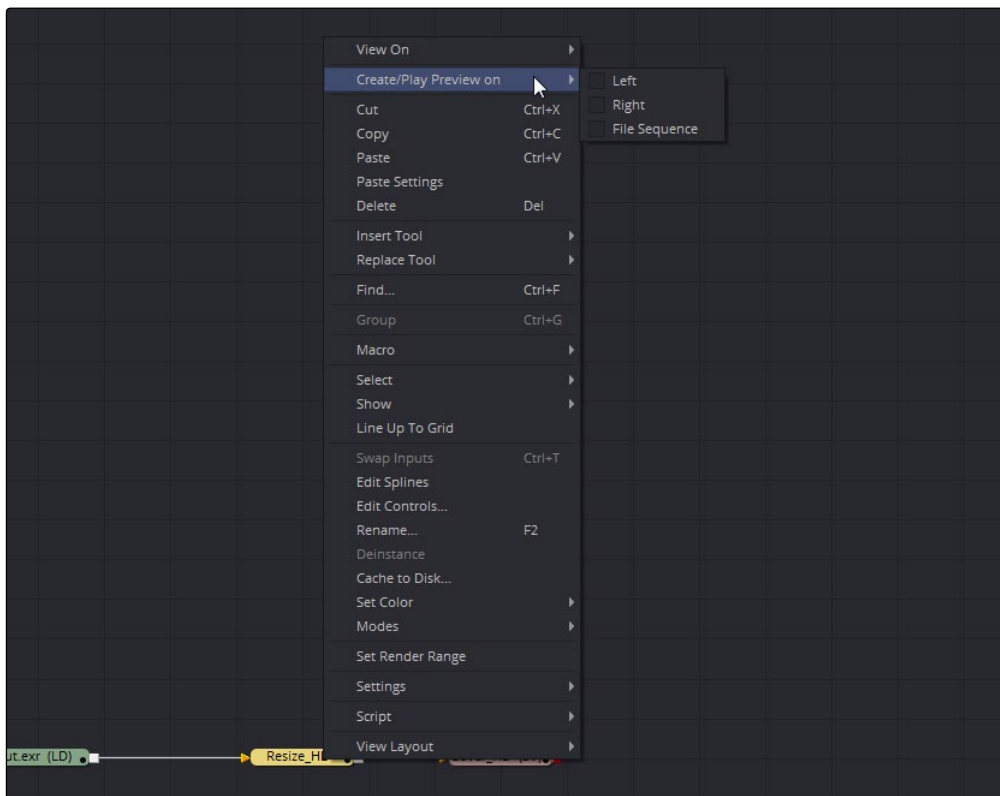
The Create/Play Preview On menu command

Cut, Copy and Paste

Cut, Copy and Paste the tool to or from the clipboard. See the Cut, Copy and Paste Tools section in this chapter for more information.

Paste Settings

Replace the active tool's settings with settings from a tool in the clipboard. See Copying and Pasting Tools for more information.

Delete

Delete the selected tool(s) from the flow. See Deleting Tools for more information.

Group/UnGroup

This command collapses the selected tools into a group or uncollapses an existing group. See Groups for more information.

Macro

Use this menu to open an existing macro for editing or to select a macro from another folder to edit.

Select

The Select submenu can be used to select the tool immediately downstream (connected to the output) or upstream (connected to the input) of the active tool. The tool is added to the current selection. It does not replace it. If there are multiple tools connected to the downstream output, all of them will be selected.

Show

The options in this submenu enable and disable the display of additional information on a tool tile, such as whether or not a tile displays a thumbnail picture, its name, or icons representing cached and locked states. See the Displaying Tile Pictures section in this chapter for more information.

Swap Inputs

Only available on a multiple input tool, such as Merge, Merge 3D and the Dissolve tool, this menu command swaps the Primary and Secondary inputs, such as the foreground and background of a Merge tool. You can also use the keyboard shortcut Command-W (Mac OS X) or Ctrl-W (Windows) to swap the inputs of the selected tool.

Edit Splines

This menu command switches the Work Panel to the Spline Editor with the splines for the selected tool(s) displayed and enabled for editing.

Rename

The Rename command displays a dialog where you can change the name of the selected tool(s).

Cache To Disk

This menu command displays a dialog used to create and manage the Disk Cache. The disk cache is covered in more detail in the Network Rendering chapter later in this manual.

The Cache to Disk dialog

Set Color

By default, the tile display color for all tools in the flow is light gray. You can change the color of any tile or text in the tile depending on which option is selected in the Set Color submenu.

Modes > Show Controls

If Show Controls is enabled, the active tool will display its parameters in the Control Panel. When disabled, the parameters will no longer appear in the Control Panel. In addition, its onscreen controls will no longer appear in the Viewers. This option is enabled by default.

Modes > Pass Through

A tool that has been passed through will not process its images or masks when rendered. The tool will simply be ignored.

Mode > Locked

Locking a tool will prevent changes from being made to its parameters in the Control Panel.

Modes > Update

When this option is enabled, any change to the tool will cause it to re-render. When this option is disabled, the tool will still allow changes to its parameters but it will not process an image until the Update option is enabled. If you have a large comp, disabling this option temporarily can allow you to quickly make a number of parameter adjustments on different tools without waiting for the flow to render. This option is enabled by default.

Modes > Force Cache

When this option is enabled, the tool's output for the current frame has an extremely high cache priority, essentially forcing it to stay cached in memory.

Cluster

This option causes the output of the tool to be pre-rendered automatically by one or more machines. This option is covered in greater detail in the Network Rendering chapter found later in this manual.

Set Render Range

The Set the Render Range menu command limits the render range to the length of the selected tool. This is useful when, for example, you want to set the render range to match the duration of a specific loader.

Tile Properties and Selection States

Certain tiles will change color to indicate special properties or selection states. The following is a list of the colors and their context.

A Text tool

Green tool tiles represent source tools, either Loaders or Creators.

An Active Brightness Contrast tool

Yellow tool tiles represent the actively selected tool. This overrides all other tool coloring.

A Selected Brightness Contrast tool

A Blue tool tiles represent a selected tool.

A Camera 3D tool

A Light Blue tool tiles represent a 3D tool.

A pEmitter (Particle Emitter) tool

Purple tool tiles represent Particle tools.

A Saver tool

Pink tool tiles represent Savers.

A failed Merge tool

Red tool tiles represent a tool that failed to render.

A passed through Brightness Contrast tool

A tool that has been passed through (disabled) will have its color darkened.

Mode Icons

When a special mode or option applied to to a tool, a small icon is applied to the lower right corner of the tile.

Cache to Disk

A tool that has been cached to disk with the icon in the lower right

Locked

A tool that has been locked with the icon in the lower right

Tooltips and the Status Bar

In even simple flows, it is very easy to forget some essential information about the clips you have read in to your comp. You can position the cursor over any tool to display a Tooltip containing that tool's name, aspect, resolution and color depth. Those same details are simultaneously displayed in the Status bar in the lower left corner of the Fusion window.

Status bar located in the lower left corner of the Fusion window

Control Panel

4

Control Panel

Control Panel Overview

Every tool added to a flow in Fusion has a set of controls which determines how that tool behaves. When a tool is selected, the controls are displayed on the right side of the screen in the Control Panel.
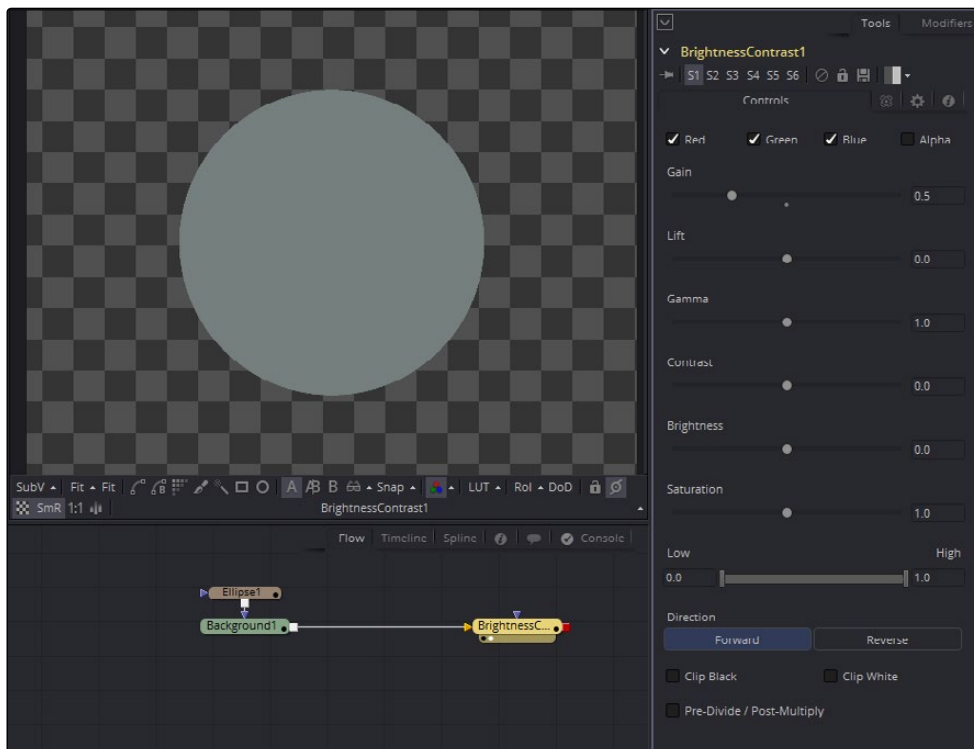
Control Panel displays the Brightness Contrast controls

To display a tool's controls in the Control Panel, do the following:

→ Select the tool from the Flow, Timeline header or Spline Editor header.

If multiple tools are selected, the Control Panel is divided into sections for all the selected tools.

The controls for a tool can be hidden so they never appear, even if the tool is selected. This can be useful in preventing accidental changes by you or other compositors who may be working on the comp.

To hide the controls for a tool, do the following:

1. Right click on the tool in the flow, or on the control header.
2. Choose Modes > Show Controls from the contextual menu.

To re-enable the display of the controls for a tool, do the following:

1. Right click on the tool in the flow.
2. Choose Modes > Show Controls from the contextual menu.

The Control Panel Interface

The Control Panel includes a small title bar at the top called the Control Header, which displays the tool's name and a button to close or open the controls.

The Control Header with the Close/Open button on the left side

More than one tool can be selected at a time and their controls are displayed in the Control Panel divided by a header for each selected tool.

Two tools in the Control Panel

Maximizing the Control Panel

A small downward arrow height button can be found at the top left of the Control Panel. This button is used to toggle the Control Panel between maximized height and half height, depending on how much room you want to provide to the Viewers.

The Maximize button on the
left side of the Control Panel

In maximized height mode, the Control Panel is aligned with the top of the window, taking up the entire right side of the Fusion window for the display of controls.
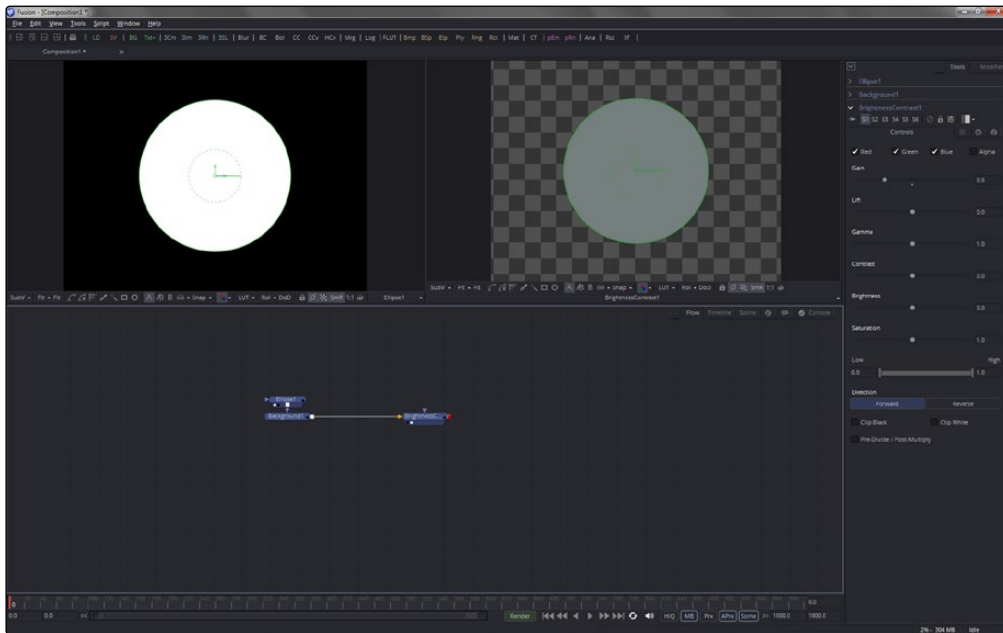
The Maximize mode

In half height mode, the top of the Control Panel is aligned with the top of the Flow Node Editor, expanding the horizontal space for the Viewers.
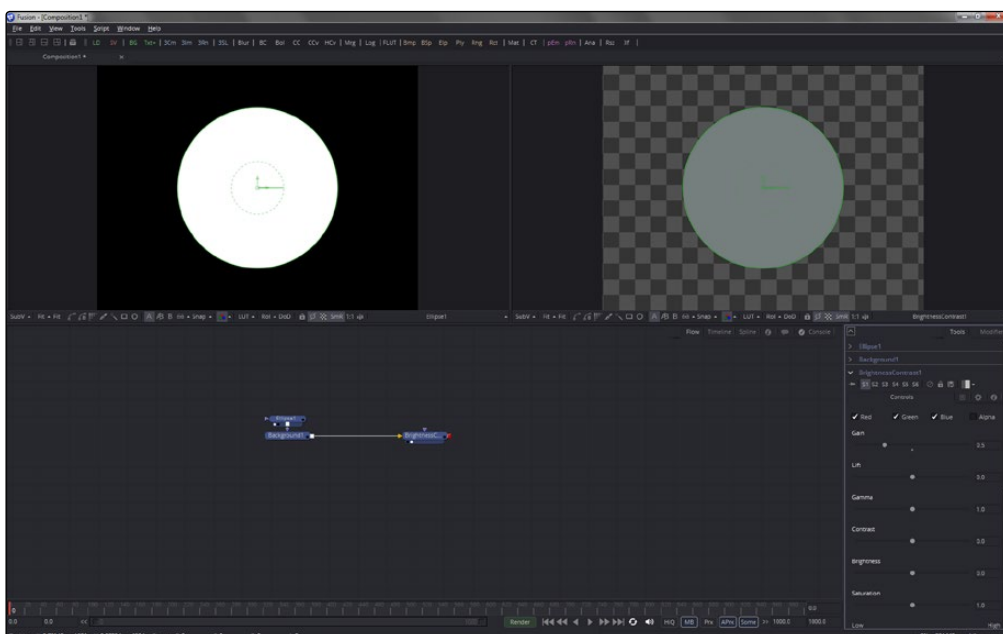
The Regular mode

Scaling and Panning the Control Window

You can drag the vertical splitter between the Control Panel and the Work Panel or Viewers to increase the overall width of the panel. Most controls will resize together with the panel. This allows for greater precision with sliders, angles and other tools.
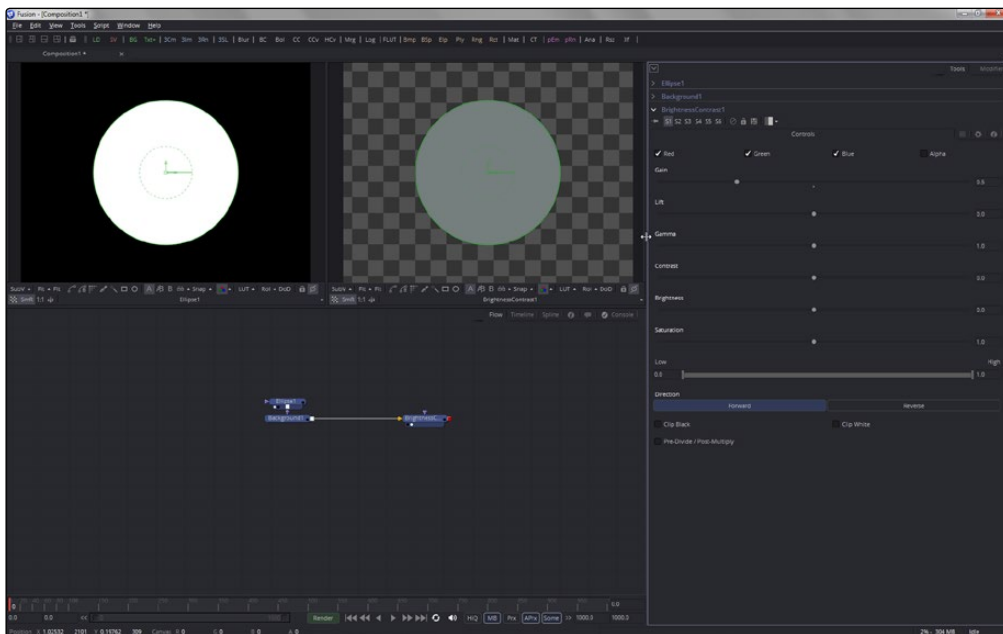
Drag the divider to increase or decrease the Control Panel width

The controls have a minimum width of approximately 280 pixels. They will begin to be cropped rather than be resized if you decrease the panel size beyond that point.

The Control Panel can be panned up or down to reveal portions that have scrolled out of the visible portion of the screen.

To pan the Control Panel up or down, do one of the following:

→ On the right side of the Control Panel, drag the scrollbar up or down.

→ Press the middle mouse button and drag up or down inside the Control Panel.

Control Panel Display Preferences

By default you only see selected tools in the Control Panel and only the active tool is expanded to show controls. You can, however, choose to change this behavior by choosing Fusion > Preferences (Mac OS X) or File > Preferences (Windows) and opening the User Interface preferences. In the User Interface preferences are two checkboxes that manage the display of controls.
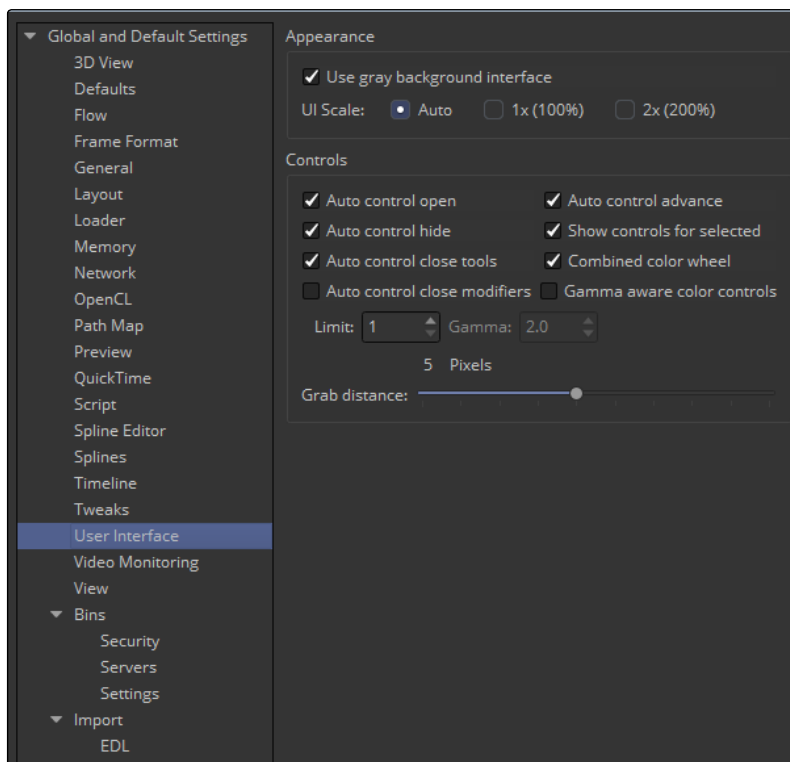
Control preferences in the General category.

Auto Control Hide

When the Auto Control Hide option is enabled, the only controls visible are for the currently selected tool. All other controls are removed from the Control Panel to reduce clutter. This is enabled by default. When this checkbox is disabled once a tool is selected its controls always remain in the Control Panel.

Auto Control Close Tools

When the Auto Control Close checkbox is enabled, the Control Panel for the active selected tool is the only expanded controls. The controls for every other selected tool are automatically collapsed in the Control Panel. This checkbox defaults to enabled.

Selecting and Viewing with the Control Header

Selecting with the Control Header

When you have multiple tools selected, you can make a tool the active tool by clicking its control header in the Control Panel. As the actively selected tool, the control header changes to a yellow color, the other tools are removed from the Control panel and any viewer displaying that tool's output will then show the tool's onscreen controls.

Viewing a Tool with the Control Header

You can view a tool by dragging the control header into one of the Viewers. To make a preview in one of the Viewers, hold Option (Mac OS X) or Alt (Windows) and drag the control header into the Viewer.

Viewing a Tool's Splines with the Control Header

If you want to view the animated curves for a tool in the Spline Editor, you can add them by dragging the control header from the Control Panel into the Spline Editor. The animated splines attached to the tool will automatically be displayed for editing.

Types of Controls

Although a few tools use customized interface elements, the vast majority of tools are modified using a mix of sliders, angle wheels and checkboxes.

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 (Mac OS X) or Ctrl (Windows) while clicking on the gutter will adjust the values in smaller increments. Holding Shift while clicking will adjust the value in larger increments.
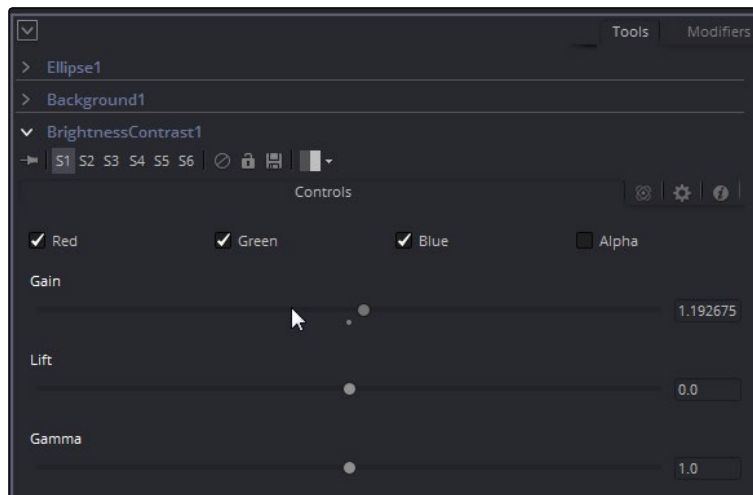
Hold Command or Ctrl while clicking in the gutter to move in smaller increments

Once you click directly on a slider handle you can make changes to its value using the left and right arrow keys. The Command/Ctrl and Shift keys can again be used to modify the value in larger or smaller increments.

While slider controls use a minimum and maximum value range, entering a value in the edit box outside that range will often expand the range of the slider to accommodate the new value. For example, it is possible to enter 500 in a Blur Size control, even though the Blur Size sliders default maximum value is 100. The slider will automatically adjust its maximum displayed value to allow entry of these larger values.

If the slider has been altered from its default value, a small circular indicator will appear below the gutter. Clicking on this circle will reset the slider to its default.
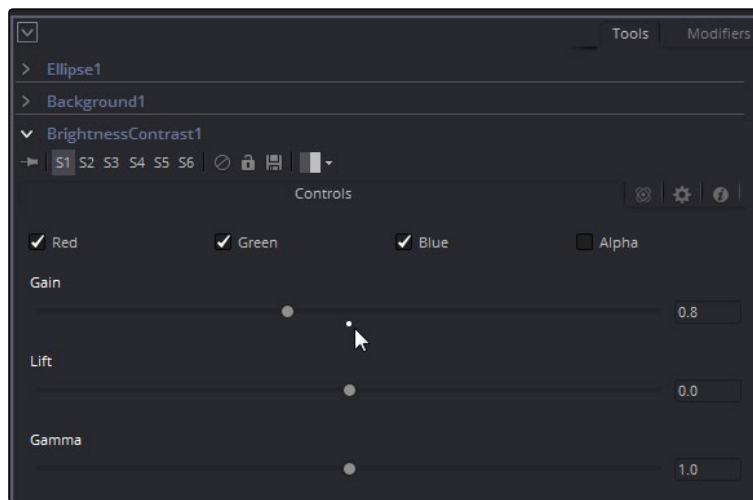
The Default indicator can be clicked to reset the slider

Thumbwheel

A Thumbwheel control is identical to a slider except it does not have a maximum or minimum value. To make an adjustment you drag the center portion left or right or by entering a value directly into the edit box. Thumbwheel controls are typically used on angle parameters, although they do have other uses as well.

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

You can use the arrowheads at either end of the control to fine tune your adjustments. Once the thumbwheel has been selected either by dragging or using the arrow keys, you can use the left and right arrows on your keyboard to further adjust the values. As with the slider control, the Command/Ctrl and Shift keys can be used to increase or decrease the change in value in smaller or larger increments.

If the thumbwheel has been altered from its default value, a small circular indicator will appear below above the thumbwheel. Clicking on this circle will reset the thumbwheel to its default.

Range Controls

The Range Controls are actually two separate controls, one for setting the Low Range value and one for the High Range value. To adjust the values, you 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 (Mac OS X) or Ctrl (Windows) and dragging either end of the Range bar. You find Range controls on parameters that require a high and low threshold, like the Matte Control, Chroma Keyer and Ultra Keyer tools.

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 color channels

Drop down Lists

Drop down lists 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 Control Panel.

Drop down list of fonts in the Text tool

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.

The Color Depth button array

Color Wheel and Picker

The Color Picker is displayed wherever a parameter requires a color as its value, such as the Fill or Outline color in the Text+ tool. The selected color is shown in a swatch below the Pick button. The swatch has two halves; the left half shows the color and the right half shows the color overlaid on a checkerboard background to preview transparency.

The color swatch with transparency preview

The color picker is extremely flexible and has four different techniques for selecting and displaying colors.

Tip: Color an be represented by 0-1, 0.255, or 0 - 65000 by setting the range you want in the Preferences > General panel.

Mac OS X and Windows Color Tools

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

Mac OS X Colors panel
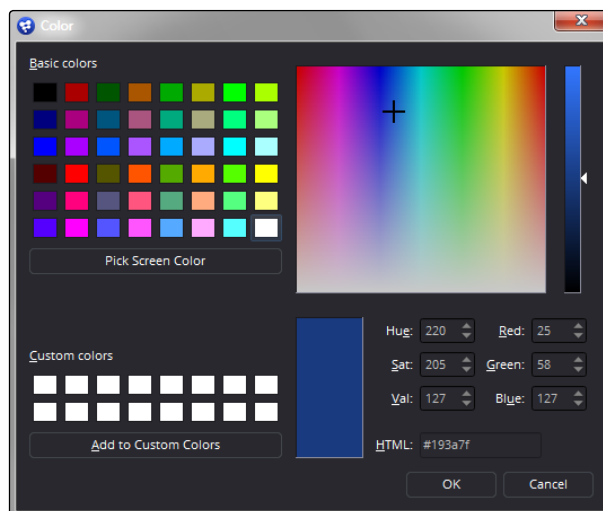
Windows Color dialog

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 Mac OS X, or the color palette on Windows. However you choose your color, you must click OK for the selection to be applied.

The Color Wheel

You also have access to the built in Color Wheel by clicking the small disclosure arrow above the Pick button. The built in color wheel includes a luminance bar for Grayscale values as well as edit boxes for Hue, Saturation and Luminance values.
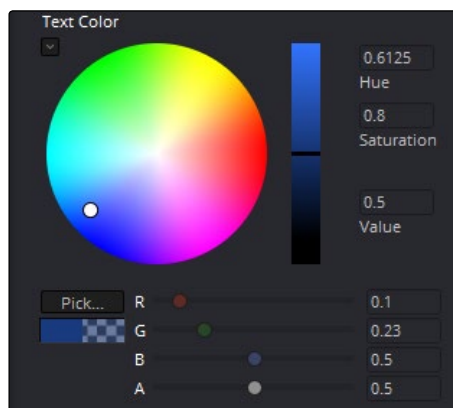
The Color Wheel in the Text+ tool

Picking Colors from an Image

If you are trying to match the color from an image in the Viewer, you can press the cursor down over the Pick button, then drag the cursor into the Viewer. The cursor will change to an eye dropper and a popup swatch will appear above the cursor with the color you are hovering over and its values. When you are over the color you want, release the mouse button to set the color.

The eye dropper with color swatch

The color picker normally selects from a single pixel in the image, but you can adjust the size of the selection by dragging into the Viewer with the eye dropper, then holding Command (Mac OS X) or Ctrl (Windows) and dragging out a rectangle for the sample side you want. The size change applies to all color pickers until the size is changed again.

Gradients

The Gradient Control bar is used to create a gradual blend between colors. The gradient bar displays a preview of the colors used from start to end. By default there are two triangular color stops, one on the left that determines the start color and one on the right that determines the end color.
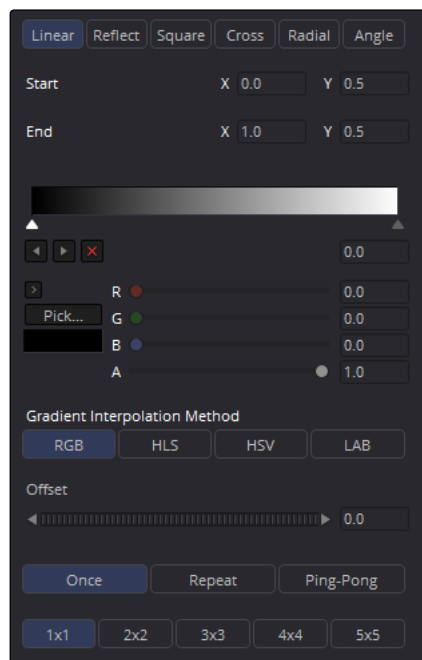
The default Gradient controls

Gradient Type

The Gradient Type button array is used to select the form used to draw the gradient. Linear draws the gradient along a straight line from the starting color stop to the ending color stop.

Linear Gradient

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

Reflect Gradient

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

Square Gradient

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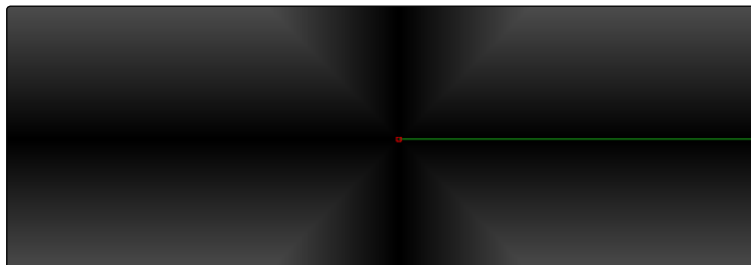
Cross Gradient

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

Radial Gradient

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

Angle Gradient

Start and End Position

The Start and End Position controls have a set of X and Y edit boxes which are useful for fine tuning the start and end position of the gradient. The position settings are also represented by two crosshair on screen controls in the Viewer, which may be more practical for initial positioning.

Gradient Colors Bar

The Gradient Colors bar is used to select the blending colors for the gradient. The default two color stops set the start and end colors. You can change the colors used in the gradient by selecting the color stop, then using the eye dropper or color wheel to set the new color.

You can add, move, copy and delete colors from the gradient using the colors bar.

To add a color stop to the gradient colors bar, do the following:

1. Click anywhere along the bottom of the gradient colors bar.
2. Use the eye dropper or color wheel to set the color for the color stop.

To move a color stop on the colors bar, do the following:

→ Drag a color stop left or right along the gradient color bar.

To copy a color stop on the colors bar, do the following:

→ Hold Command (Mac OS X) or Ctrl (Windows) while you drag a color stop.

To delete a color stop from the colors bar, do one of the following:

→ Drag the color stop up past the gradient bar.

→ Select the color stop, then click the red X button to delete it.

Gradient Interpolation Method

The Gradient Interpolation Method button array lets you select what color space is used to calculate the colors between color stops.

The Gradient Interpolation Method button array

Offset

When you adjust the Offset control, the position of the gradient is moved relative to the start and end markers. This control is most useful when used in conjunction with the repeat and ping-pong modes described below.

Once/Repeat/Ping-Pong

These three buttons are used to set the behavior of the gradient when the Offset control scrolls the gradient past its start and end positions.The Once button is the default behavior, which keeps the color continuous for offset. Repeat loops around to the start color when the offset goes beyond the end color. Ping pong repeats the color pattern in reverse.

1x1, 2x2, 3x3, 4x4, 5x5

These buttons control the amount of sub-pixel precision used when the edges of the gradient become visible in Repeat mode, or when the gradient is animated. Higher settings will take significantly longer to render but be more precise.

Gradient Contextual menu

Gradients have their own context 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 tool in the flow. To learn more about Modifiers, read Chapter 14 Introduction to Controls and Customization later in this manual.

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 Control Panel Modifiers tab.

To attach a Modifier, do the following:

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.

To learn more about Modifiers, read Chapter 14 Introduction to Controls and Customization later in this manual.

The Common Controls Tab

The Common Controls tab is represented by a radioactive icon and can be found in the Control Panel for almost every tool, including third party plug-ins. This tab includes standard controls that normally appear for every tool.
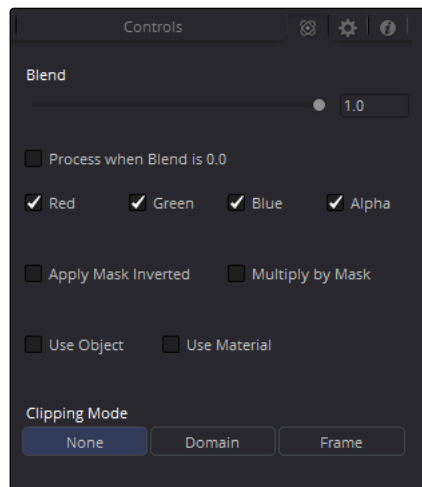
The Common Tab in the Control Panel

The following controls are found in the Common Controls tab, although some controls may not appear on all tools. For example, Motion Blur settings have no purpose in the Cineon Log tool.

Blend

The Blend control is found in all tools, except the Loader and Creator tools. It is used to blend between the tool's unaltered image input and the tool's final modified output image. When the blend value is 0.0, the outgoing image is identical to the incoming. Normally, this will cause the tool to skip processing entirely, copying the input straight to the output. The default for this tool is 1.0, meaning the tool will output the modified image 100 %.

Process When Blend is 0.0

This checkbox forces the tool 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 tools or third party plug-ins that store values from one frame to the next. If this checkbox is disabled on tools that operate in this manner, the tool will skip being processed when the Blend is set to 0, producing incorrect results on subsequent frames.

Red/Green/Blue/Alpha Channel Checkboxes

The four checkboxes for Red, Green, Blue and Alpha are used when you want to limit the effect of a tool to specific color channels. This filter is often applied after the tool has been processed.

For example, in the graphic below, the Transform operation is limited to the channels selected (Green, Blue and Alpha)

Results of having the Red checkbox disabled on a Transform tool

There are some exceptions, such as tools where deselecting a channel causes the tool to skip processing that channel entirely. Tools that operate this way will generally include an identical set of Red, Green, Blue and Alpha checkboxes on another tab in the tool. In these cases, the Common Control channel checkboxes are instanced to the channel boxes found elsewhere in the tool.

Apply Mask Inverted

When the Apply Mask Inverted checkbox is enabled, the Mask channel for the tool is inverted. The Mask channel is the combined result of all masks applied to a tool.

Tip: The Apply Mask Inverted check box only operates on Effects Masks, not on Garbage Masks.

Multiply By Mask

Selecting this option will cause the RGB values of the masked image to be multiplied by the Mask channel's values. This will cause all pixels of the image not included in the mask (i.e., those set to 0) to become black. This creates what is often called a premultiplied image.

Use Object/Use Material (For Masking)

Some 3D software can render to file formats that support additional channels. Notably, the Open EXR file format supports object and material ID channels, which can be used as a mask for the effect. This checkbox determines whether the channels will be used if they are available. The specific Material ID or Object ID affected is chosen using the next set of controls.

Pick Controls

The Pick Controls are only displayed once the Use Object or Use Material checkbox is enabled. These controls select which ID is used to create a mask from the Object or Material channels saved in the image. You use the Pick button to grab IDs from the image in the Viewer, the same way you use the color picker to select a color. The image or sequence must have been rendered from a 3D software package with those channels included.

Correct Edges

The Correct Edges check box is only displayed once the Use Object or Use Material checkbox is enabled. When the Correct Edges checkbox is enabled, the Coverage and Background Color channels are used to separate and improve the effect around the edge of the object. When disabled (or no Coverage or Background Color channels are available), aliasing may occur on the edge of the mask.

Motion Blur

The Motion Blur checkbox enables and disables the rendering of motion blur for the tool. When this checkbox is enabled, the tool's predicted motion is used to produce the blur caused by the virtual camera's shutter. When the control is disabled, no motion blur is created.

Quality

The Quality slider is only displayed when the Motion Blur checkbox is enabled. The slider 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 increase the render time.

Shutter Angle

The Shutter Angle slider is only displayed when the Motion Blur checkbox is enabled. 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

The Center Bias slider is only displayed when the Motion Blur checkbox is enabled. This slider modifies the position of the center of the motion blur. Adjusting the value allows for the creation of trail type effects.

Sample Spread

The Sample Spread slider is only displayed when the Motion Blur checkbox is enabled. Adjusting this control modifies the weighting given to each sample. This affects the brightness of the samples set with the Quality slider.

The Scripting Tab

The Scripting tab is present on every tool and is represented by a red gear icon. The Scripting tab contains several edit boxes used to add scripts that process when the tool is rendering. For more details on the contents of this tab, please consult the Scripting documentation.
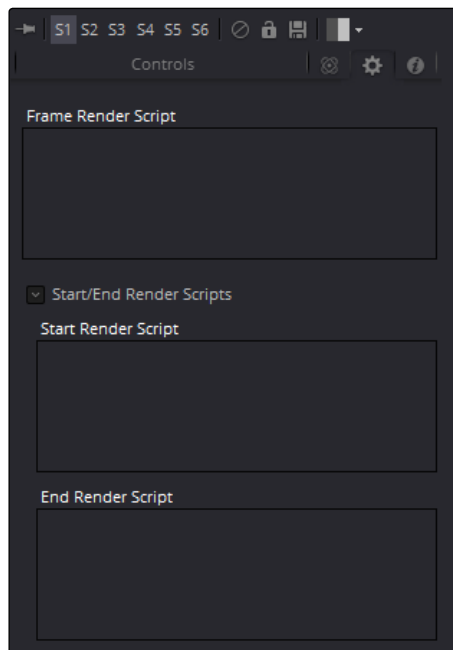
The Scripting tab

The Comments Tab

The Comments tab is found on every tool and can be identified by a small blue letter I icon. It contains a single text control that is used to add comments and notes to the tool. To enter text you click in the entry area and begin typing.
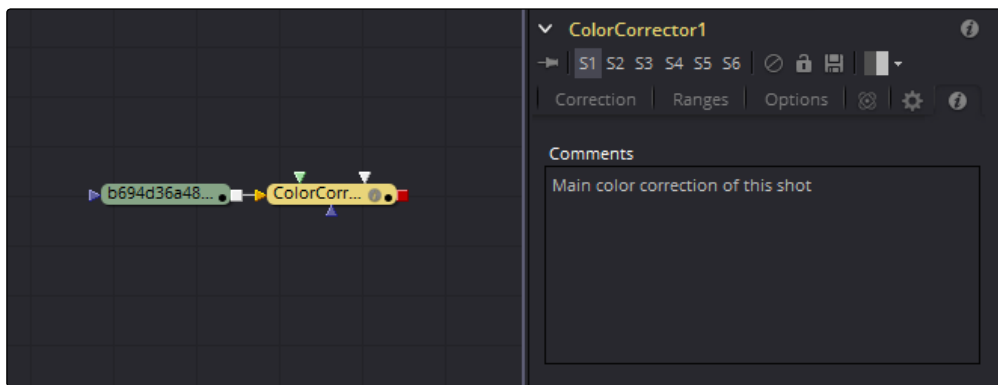
The Comments tab

When a note is added to a tool, the comments icon appears in the control header and can be seen in a tool's tooltip when the cursor is placed over the tool in the flow. The contents of the Comments tab can be animated over time, if required.

Additional controls appear under this tab if the tool is a Loader. See the Creator Tools chapter in the Tools Reference manual for more details.

Animating a Parameter

Almost every parameter can be animated over time. To animate a parameter you attach it to an animation curve.

To animate a Parameter, do the following:

1. Right click on the parameter you want to animate.

2. In the contextual menu choose Animate.

When you attach a parameter to an animation curve, a keyframe is set at the current frame using the current parameter value.
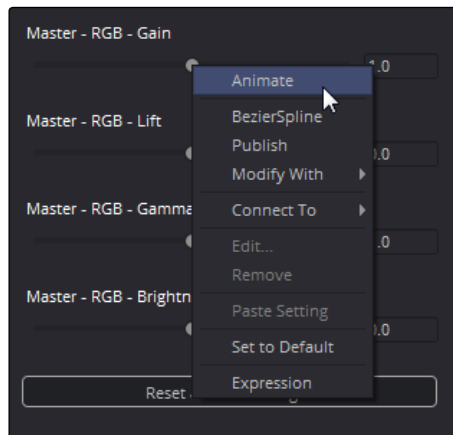
The Animate command in the contextual menu

The edit boxes of animated parameters are highlighted in blue or green and Animation Indicator bars appear to the left of each animated control. Blue indicates that the parameter is animated. Green indicates that a keyframe exists at the current position in time.

Gain and Lift parameters with animated Value fields and Indicator bars

To remove a keyframed spline from a parameter, do the following:

1. Right click on the parameter.

2. Choose Remove Bezier Spline from the context menu.

Tip: If you change the default spline type from Bezier, the contextual menu will display the name of the current spline type.

Setting Keyframes in the Control Panel

To animate a parameter you must set at least two different values on different frames so the parameter can interpolate between them. To set values on an animated parameter you use Keyframes. The first keyframe is automatically added when you attach the animation curve from the contextual menu. Additional keyframes are automatically added whenever you change the value of the parameter. If there is a keyframe already at the current time, the keyframe is updated, otherwise, a new keyframe is created.

If you want to add a keyframe without changing the value of the parameter, right click on the parameter and choose Set Key from the contextual menu. This is useful when you want the animation to hold at specific value for a length of time.
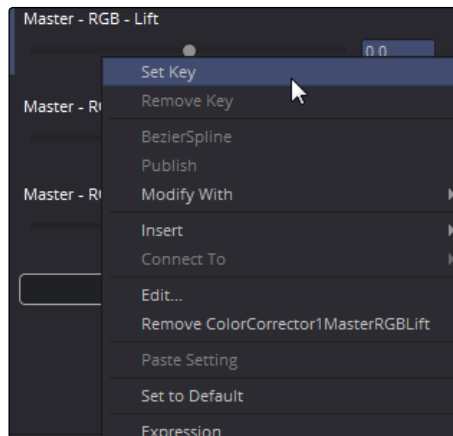
The Set Key menu command adds a keyframe
without changing the parameter's value.

Attaching to an Existing Animation Curve

Multiple parameters can be connected to the same animation curve. This can be an invaluable time saver if you are identically animating different parameters in a tool.

To connect a second parameter to the same animation curve, do the following:

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.
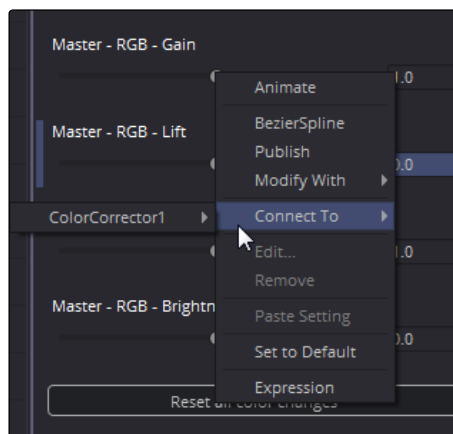
The Connect To submenu connecting a Gain parameter

Connecting Parameters without Animation

It is often useful to connect two parameters together even without an animation curve. 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 second parameter.

To Publish and Connect parameters, do the following:

1. Right click on the parameter you want to publish.

2. Choose Publish from the contextual menu.

3. Right click on the second parameter you want to attach.

4. In the contextual menu, hover over the Connect To submenu.

5. In the Connect To submenu, choose the name of the animated parameter.
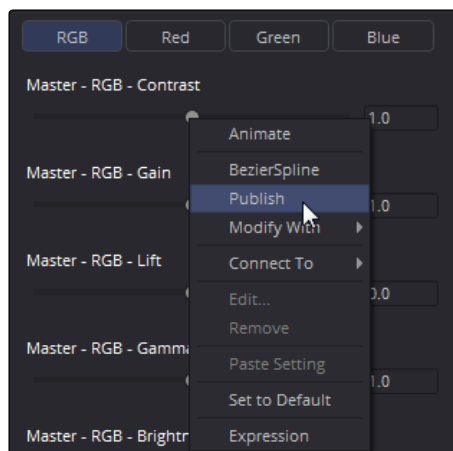
The Publish contextual menu

Pickwhipping

You can also use simple expressions to link two parameters together. By using simple Expressions and Pickwhipping, values can be connected and combined visually without the need to publish a value first.

To link two parameters using a Pickwhip, do the following:

1. In the Edit box, type =.

2. Press Return (Mac OS X) or Enter (Windows).

3. Drag from the Add button next to the expression field to the target parameter.

4. Adjust the target parameter to adjust the original parameter.
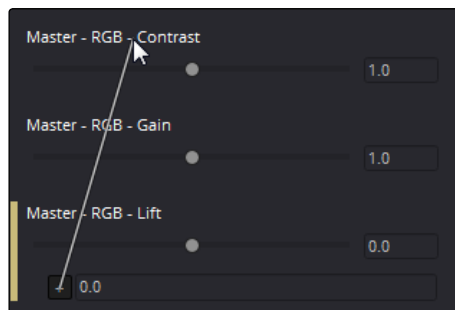
Pickwhipping between two parameters

Tip: Disabling the Auto Control Close tool's General preference, then selecting two tools in the Flow Node Editor will allow you to pickwhip two parameters from different tools

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

For more information on Pickwhipping and Expressions, read Chapter 14 Introduction to Controls and Customization later in this manual.

Contextual Menus

Tool Contextual Menu

To display the Tool Context Menu from the Control Panel, right click on the control header. The tool's contextual menu includes the same menu options that are accessed by right clicking on a tool in the flow. See Chapter 3 on the Flow Node Editor for more information on these options.

Control Contextual Menu

The contextual menu for individual parameters is accessed by right clicking over a slider, thumbwheel, range control, button array or other control type. For example, right clicking on a slider will provide the slider's contextual menu, with options to animate the control or add additional modifiers. Many of these options were described in this chapter.

Viewers

5

Viewers

Viewer Overview

The Viewers display 2D images, but they also show any 3D environment in the composition as well as offer waveforms, histograms and vectorscopes to analyze the image.
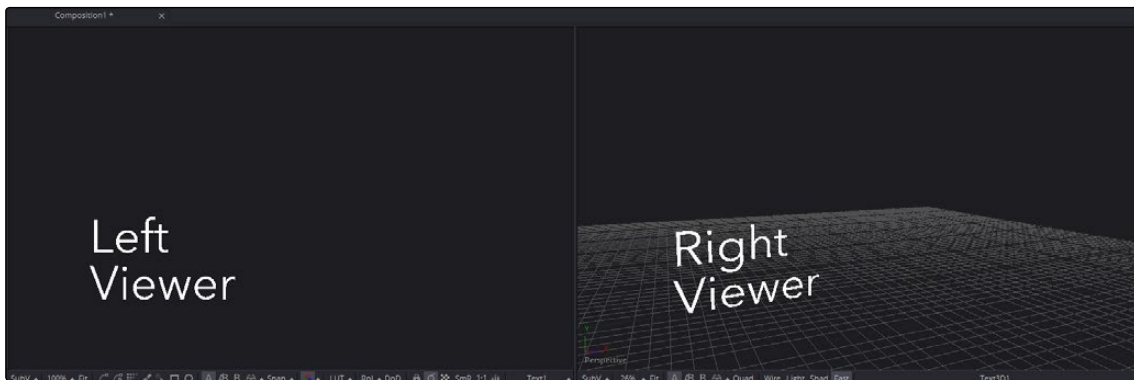
Viewers side by side, along the top half of the window

By default there are two viewers positioned side by side across the top of the window. Additional floating viewers can be created, or an external video monitor can be used if supported hardware is present on the computer, such as a Blackmagic Design DeckLink.

The viewers are OpenGL-accelerated and rely on the presence of a good OpenGL-accelerated display device.

Types of Viewers

Although the primary viewers are the Left and Right Viewers that are always available above the Flow Node Editor, these viewers are not the only viewers available to you. You can open additional viewers as well as HD, UltraHD and DCI 4K monitors to evaluate your comps.

Floating Viewers

You can open an unlimited number of Floating Viewers which are ideal for dual monitor configurations. Floating viewers can be set to full screen mode to make the best use of screen real estate. To open a new floating viewer, choose Window > New Image View from the menu bar. If this is an option you want to use regularly, you can save the position and configuration of the new viewer using the layout preferences.
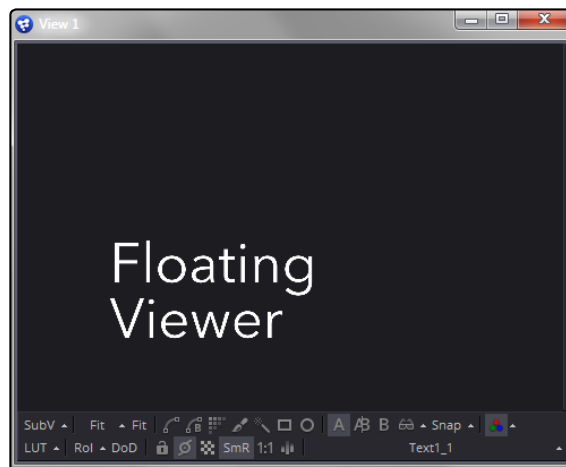
Floating Viewer

HD, Ultra HD and DCI 4K Monitors

If video hardware, like the Blackmagic Design DeckLink 4K Extreme, is present in the computer, you can display images on external SD, HD, UltraHD or full frame DCI 4K monitors. You can configure your Blackmagic Design DeckLink hardware by going to Fusion > Preferences (Mac OS X) or File > Preferences (Windows) and selecting the Video Monitoring preferences.
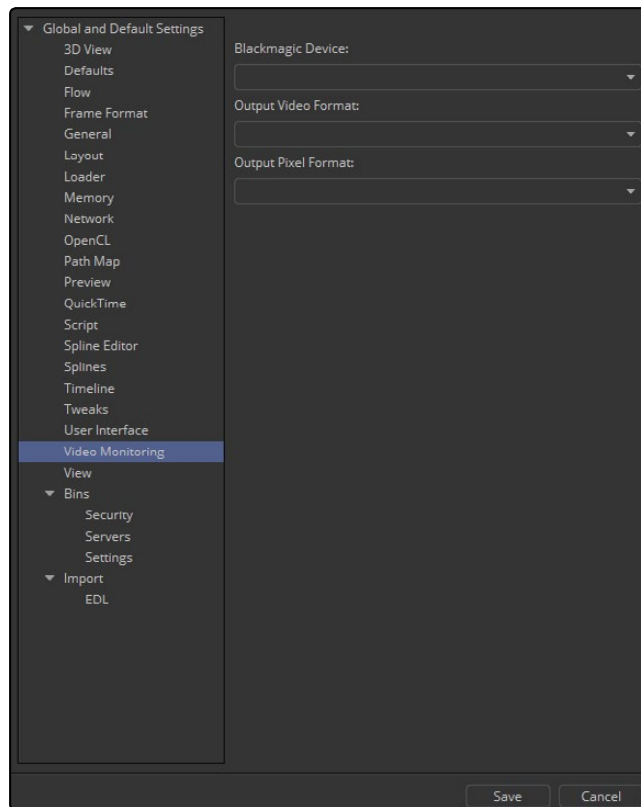
The Blackmagic DeckLink card preferences

The monitors cannot be used to manipulate preview controls like center cross hairs or masks, but they are extremely valuable for evaluating the images in the same format in which they will be output, and for determining color fidelity and accuracy with a properly calibrated monitor.

Displaying a Tool in the Viewer

There are a number of different ways to display a tool in the Viewer, starting with the Tool Tile View Indicators. The View indicators are displayed under each tile. By default, there are two indicators, representing the two Viewers in the Fusion window. The View indicator on the left is used for the Viewer on the left, while the View indicator on the right is used for the Viewer on the right. If a new viewer is added, an additional indicator will appear on each tile. Clicking an indicator will turn it white and display that tool in the appropriate viewer. Clicking it again removes it from the viewer and removes the white highlight.

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

Tip: If a tool is not currently displaying View indicators, position the cursor over the tile to make the indicators visible and available for clicking.

The name of the currently viewed tool is displayed on the horizontal toolbar at the bottom of the Viewer.

The name of the tool currently displayed is located at the bottom of the Viewer

You can use keyboard shortcuts to display the selected tool in a viewer. Each viewer is assigned a numeric keyboard shortcut from 1-9. Pressing that number on the keyboard will display the selected tool in the appropriate Viewer.

To view a tool using keyboard shortcuts, do one of the following:

→ Selected a tool in the Flow Node Editor, then press 1 to display it in the left Viewer.

→ Selected a tool in the Flow Node Editor, then press 2 to display it in the right Viewer.

You can also drag the tool from the Flow Node Editor into a viewer. The tool will appear to move at first but will snap back into its original location once the cursor leaves the area of the Flow Node Editor.

To view a tool by dragging, do the following:

1   Drag the tool into the Viewer.

2   Release the mouse button over the Viewer.

To view a tool using the Control Panel, do the following:

→ Drag the tool's control header from the Control Panel into a viewer.

If you have an image that is stored on your computer's hard drive, you can use the viewers for quick previews before adding it to a flow.

To view an image stored in a file browser, do one of the following:

→ Drag the image from the operating system's file browser into a viewer.

→ Right click over a viewer, then choose Load Preview.

This contextual menu will load the clip as a preview with transport controls at the bottom of the Control Panel to play, review, and loop the sequence.

The transport controls at the bottom of the Control Panel

When you drag an image into a viewer, whether it be from the Control Panel, the Flow Node Editor or the file browser, holding Option (Mac OS X) or Alt (Windows) while dragging will create a Flipbook Preview in that viewer. The Render Settings dialog will be displayed, and upon accepting the settings, the tool will be rendered and the resulting frames stored in RAM for fast playback on that viewer. This is similar to selecting Create Preview On from the tool's contextual menu. You can also hold Shift-Option (Mac OS X) or Shift-Alt (Windows) to skip the dialog and immediately start rendering the preview with the last used settings.

Tools can also be loaded into a viewer using the different contextual menus.

To load a tool into a viewer using the tool's contextual menu, do the following:

1. Right click on a tool's tile in the Flow Node Editor.

2. Choose View On, then in the submenu choose the viewer to use.
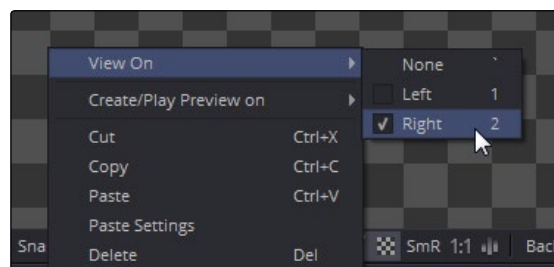
The Tool contextual menu

To load a tool into a viewer using the Control header contextual menu, do the following:

1. Right click on the tool's Control header in the Control Panel.

2. Choose View On, then in the submenu choose the viewer to use.
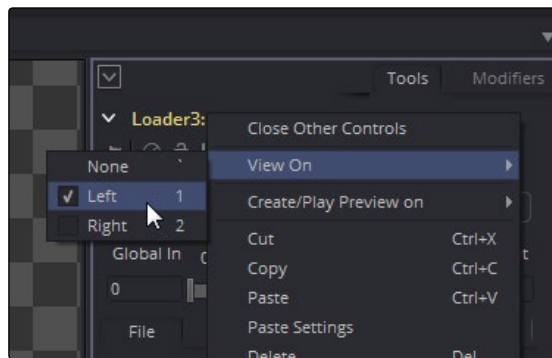
The Control header's contextual menu

To load a tool into a viewer using the Timeline header contextual menu, do the following:

1. Right click on the tool's name in the Timeline header area to the left of the Timeline.

2. From the contextual menu choose View On, then from the submenu choose the viewer.

Removing Images from the Viewer

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

Position and Layout

When you resize and change the layout of viewers in the comp, that configuration is always saved with the comp. So each time you open the comp, the size and layout is remembered. You can prevent this behavior by disabling the Recall Layout checkbox in the Global Layout preferences.

If you want all new compositions to open with a certain Viewer layout, you can configure the layout of the two primary viewers then use the Grab Document Layout button in the Global Layout preferences to remember the layout for any new compositions. To save the position and size of floating viewers, you use the Grab Program Layout button. Finally, if you want to have the floating viewers opened automatically when you open Fusion, enable the Create Floating Views checkbox.

Layout Toolbar

The Layout Toolbar at the top of the Fusion window can be used to configure the Viewers and panels using four preset layouts.

The Layout Toolbar

First Button

Use this option to maximize a single viewer (Viewer 1), aligning the Control Panel with the top of the screen.

Second Button

Use this option for dual viewers, aligning the Control Panel with the top of the screen. This is the default layout.

Third Button

Use this option to maximize a single viewer (Viewer 1), aligning the Control Panel with the top of the Flow Node Editor.

Fourth Button

Use this option for dual viewers, aligning the Control Panel with the top of the Flow Node Editor.

Splitter Bars

You can change the relative sizes of the left and right viewers using the horizontal splitter bar that runs between them.
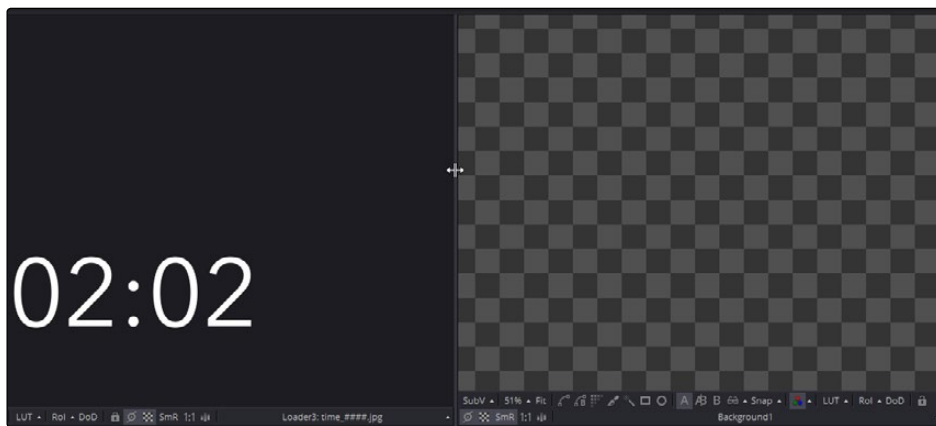
The Viewer Splitter bar

Drag the splitter to increase or decrease the amount of space used by one viewer. The adjacent viewer will adjust to accommodate the new layout.

The amount of vertical space available for both viewers can be adjusted by dragging the horizontal scroll bar between the Viewers and the Work area below them.

Tip: You can use the height button above the Control Panel to shrink the Control Panel to half height, giving the Viewers additional horizontal space. To learn more about the Control Panel, read Chapter 4.

Fullscreen Viewers

Floating viewers are not affected by the buttons in the Layout toolbar. You adjust the size of a floating viewer by dragging any corner of the Viewer window, or by double clicking on the Viewer's titlebar. Additionally, a floating viewer can be scaled to fullscreen by right clicking in the Viewer and choosing Global Options > Full Screen from the contextual menu.

You can maximize one of the fixed viewers to full screen by placing the cursor over the Viewer and pressing the F4 key. To restore a full screen viewer to its original layout position, press F4 again. This shortcut can also be used to maximize the Work area below the Viewers.

Tip: The Full screen option is only available in Fusion Studio

Panning in the Viewer

The image in the Viewer can be panned so you can focus your work a specific area of the frame. To pan, position the cursor in the Viewer, hold the middle mouse button, then drag the mouse in the direction you want to pan. If your mouse does not include a middle mouse button, hold Command-Shift (Mac OS X) or Ctrl-Shift (Windows) while dragging.

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.

Tip: While viewing a 3D scene, zooming in will decrease the sensitivity of panning, to allow finer control when navigating around small images and scenes. Correspondingly, pan sensitivity will increase when zooming out. If you need to move the point of view faster, zoom out first or reset the scale with Command-1 (Mac OS X) or Ctrl-1 (Windows).

Scaling in the Viewer

There are a number of ways to scale an image or 3D scene once it is displayed in a viewer.

You can use the mouse, the keyboard, the contextual menu or the Viewer toolbar.

Tip: There is a lower limit to the scale of a 3D scene. Continuing to zoom in past this limit will instead move ("dolly") the point of view forwards. The mouse wheel will move forward slowly, and the keyboard will move more quickly.

Using the Mouse

To scale the image in the Viewer by dragging the mouse, do the following:

1. Place the cursor over the Viewer.
2. Hold the middle and left mouse buttons.
3. Drag the mouse to the left or right.

The image smoothly scales, centered on the starting position of the cursor. Dragging to the right zooms in, while dragging to the left zooms out.

To scale the image in the Viewer by clicking on the mouse, do the following:

1. Place the cursor over the Viewer.
2. Hold the middle mouse button.
3. Click once using the left mouse button to zoom in or the right button to zoom out.

The scaling uses a fixed amount, centered on the position of the cursor.

To scale the image in the Viewer by rolling the mouse wheel, do the following:

1. Place the cursor over the Viewer.
2. Press Command (Mac OS X) or Ctrl (Windows) and roll the mouse wheel up or down.

Each step on the wheel increases or decreases the scale by a fixed amount, centered on the cursor.

Using the Keyboard

Using the keyboard, you can change the scale in a viewer by a fixed value. The + or - keys on the numeric keypad can be used on the active viewer to zoom in and out. Pressing Command-1 (Mac OS X) or Ctrl-1 (Windows) will set the image or scene back to its default of 100% scale, and will recenter it as well. If the image is larger than the Viewer, which in most cases it is, pressing Command-F (Mac OS X) or Ctrl-F (Windows) will fit the image or scene in the Viewer's frame.

Using the Contextual Menu

The benefit of using the Viewer's contextual menu for scaling is the number of preset scale settings that can be selected to scale the Viewer by a specific amount. You can also quickly scale the image to fit in the Viewer's frame.
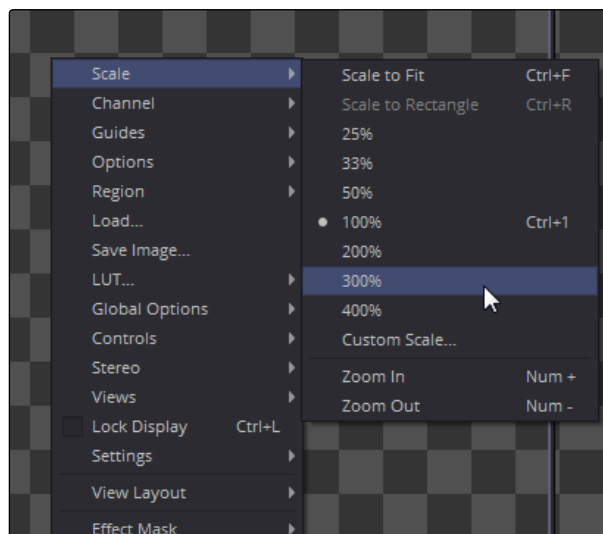
The Scale submenu in the Viewer's contextual menu

Using the Viewer Toolbar

The Viewer Toolbar contains a Scale menu similar to the one found in the contextual menu. Click on the Scale menu to bring up a list of several preset scales to choose from, or click on the Fit button beside it to fit the image or scene to the Viewer window.
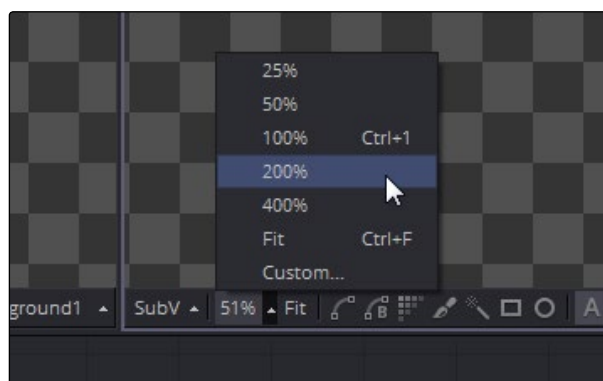
The Scale menu and Fit button in the Viewer toolbar

Making Previews

In addition to displaying the output of a tool, the viewers are also used to display flipbook previews and do file playback of image sequences and AVI or MOV movie files straight from disk.

RAM Flipbook Previews

To watch a clip or the output of a tool play back in the Viewer, you can use the RAM flipbook. RAM Flipbook Previews are preview renders that exist entirely within RAM and allow you to watch a tool's output at playback speed. Flipbook previews are created by holding Option (Mac OS X) or Alt (Windows) while dragging a tool into the Viewer, or by selecting Create/Play Preview from the tool's contextual menu.

Either way you invoke the flipbook preview, a Preview Render dialog will open where you can change the quality, resolution and motion blur settings of the flipbook preview.
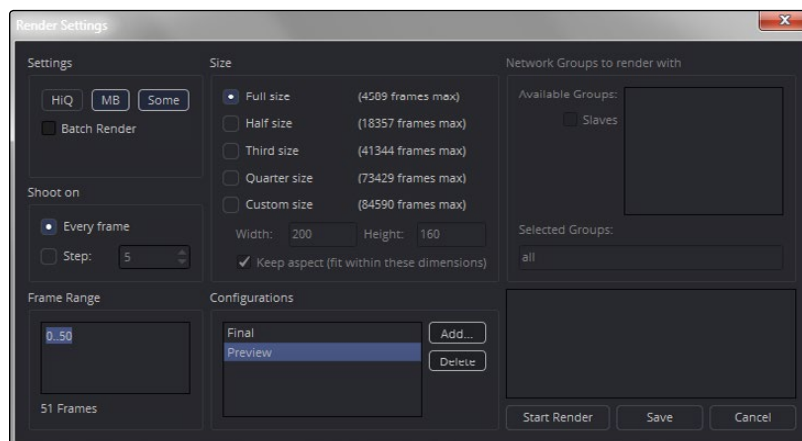
The Flipbook Preview Render dialog

If you want to bypass the dialog and use the default setting or the settings that were chosen last, hold Shift-Option (Mac OS X) or Shift-Alt (windows) while you drag the tool into the Viewer. The settings dialog will not show up and rendering the preview will start right away.

Settings

The Settings section of the Preview Render dialog include three buttons that determine the overall quality and appearance of your flipbook preview. These buttons also have a significant impact on render times.

The Settings buttons in the Flipbook Preview Render dialog

HiQ

When enabled, this settings renders the preview in full image quality. If you need to see what the final output of a tool 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 tool 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 tools specifically needed to produce the image of the tool you're previewing are rendered.

Size

Since RAM flipbook previews use RAM, it is incredibly 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.
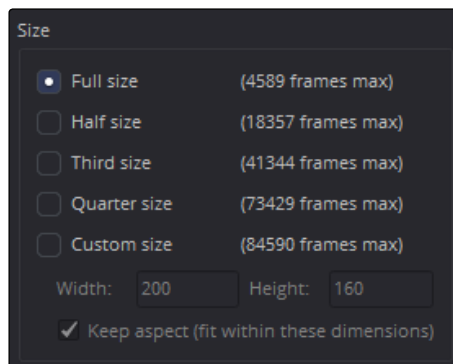
The Size settings with the number of frames that will fit into RAM

Network

You can use other computers on your network running render only versions of Fusion or full versions of Fusion Studio to help render frames for the flipbook preview. The available computers running on the network are displayed in the Network Groups pane. See the Network Rendering chapter for more information.
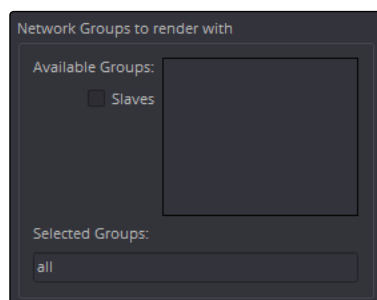
The Network Groups pane displays additional computers available for rendering

Shoot On

Sometimes you may not want to render every single frame, but every second, third or fourth to save render time and get faster feedback. You can use the Step parameter to determine the interval at which frames are rendered.

The Shoot On parameters are used for skipping frames during render

Frame Range

This field uses 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.

The Frame Range determines the starting and ending frames for the flipbook preview

Configurations

Once you're done setting your personal preview configuration, it can be saved permanently for later use by clicking Add and then typing in a new name for your personal settings.

The Configurations area can be used to save your settings for future flipbook previews

Playing Back the Preview

Once the preview is done rendering, the flipbook is loaded into the Viewer. A frame counter appears in the corner of the Viewer, and a Play/Pause button is added to the right of the Viewer's bottom toolbar. A more complete set of playback controls appear in the bottom right corner of the Fusion window beneath the Control Panel.

The flipbook preview with a Play/Pause button in the Viewer's toolbar.

Once you've watched the flipbook, you can remove it from the Viewer and free up your RAM by right clicking in the Viewer and choosing Remove Preview from the contextual menu.
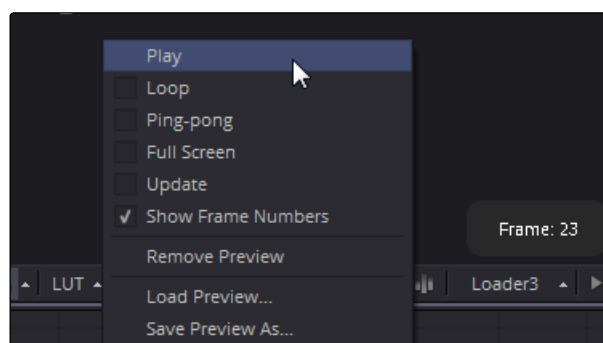
The Flipbook Preview options in the Viewer's contextual menu

With the flipbook preview still loaded in the Viewer, there are other options in the Viewer's contextual menu to enable looped or ping-pong playback, as well as enter Fullscreen mode, which allows you to focus on the clip without the rest of Fusion's interface. After you've viewed the clip, to exit Fullscreen mode you deselect Fullscreen from the contextual menu, or press the Esc key.

It is also possible to "trim" the flipbook in order to skip playback of the first or last few frames. This can be done using the scrub bar in the transport controls that appear at the bottom right of Fusion's window. Hold Command (Mac OS X) or Ctrl (Windows) and drag the the start or end of the scrub bar to set or change the frames to be played. Right clicking on the scrub bar will bring up a contextual menu where you can reset the In and Out points.
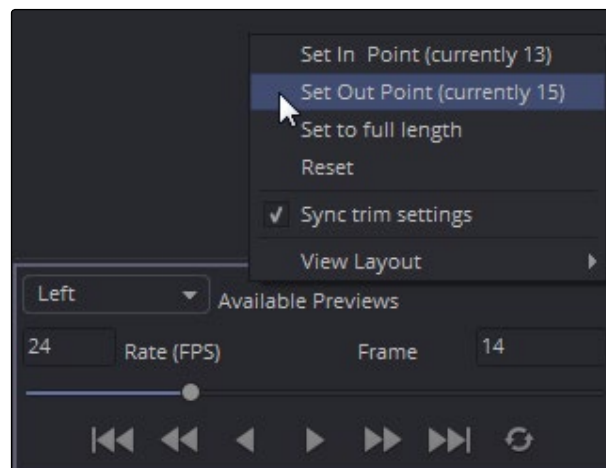
The transport controls can be used to trim frames
from the start and end of the flipbook

Updating a Preview

An option that is designed for the interactive frame by frame work of rotocoping and paint is to set a Preview to Update from its contextual menu. When active, any frames that are modified on the previewed tool are automatically updated in the preview's playback. This lets you reserve the RAM to do playback. You can keep it playing on a loop or ping-pong while you work in another viewer.

Mouse and Keyboard Shortcuts

While the flipbook preview is loaded into a viewer you can play or scrub through it using only the mouse or the keyboard.

To play back a flipbook using the mouse, do the following:
> → Double click in the Viewer to start playback.

To scrub through a flipbook using the mouse, do the following:

→ Hold the right mouse button and drag the mouse left or right to scrub through frames.

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

→ Press spacebar to start or stop playback.

→ Hold Shift and press 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 (Mac OS X) or Ctrl (Windows) left arrow key to jump to the first frame.

→ Press Command (Mac OS X) or Ctrl (Windows) right arrow key to jump to the last frame.

Tip: The mouse and keyboard shortcuts work in Fullscreen mode as well.

File Sequence Previews

If you want to play the file sequence or movie file read in by a Loader tool you don't have to create a flipbook preview, you can use the transport controls that appear in the bottom of the Control Panel.

After pressing the Play button, the images are displayed as quickly as possible in the Viewer. By default, the left Viewer is used, but that default can be changed using the Preview preferences. By changing these preferences you can also play the file sequence to installed video hardware, but this may decrease playback performance.

As with flipbooks, you can trim the In and Out points for image sequence previews, using the scrub bar in the transport controls. By default, trimming done in the scrub bar while a Loader tool is selected will be mirrored in the Loader's Trim In/Out controls. This can be disabled using the scrub bar's contextual menu.

While the file sequence preview is enabled, normal viewing of tools to that viewer is disabled. When you are ready to return to displaying tools in the Viewer, select another tool or empty area of the Flow Node Editor.

Preview Controls

When it comes to adjusting images, the Control Panel provides very precise numerical values, but sometimes visually positioning an element using on screen controls can get you where you want to go with less tweaking. The Viewers show preview controls for manipulating the parameters of the currently selected tool. Common preview 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 are determined only by the tools that are selected, not by the tool 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 tools are selected, the controls for all of those tools are shown.

Enabling/Disabling Preview Controls

The preview controls for a viewer can be disabled so they do not interfere with viewing the image.

The Show Controls button in the Viewer toolbar is
used to show and hide preview controls

To disable the preview controls, do one of the following:

→ Click the Show Controls button in the View toolbar.

→ Right click in the Viewer and choose Show Controls in the Options submenu.

→ Press Command (Mac OS X) or Ctrl (Windows) -K on the keyboard.

Tip: Some tools, like Masks, allow disabling of their preview controls on a per-tool basis because you often use multiple polygon tools to organize and animate masks.
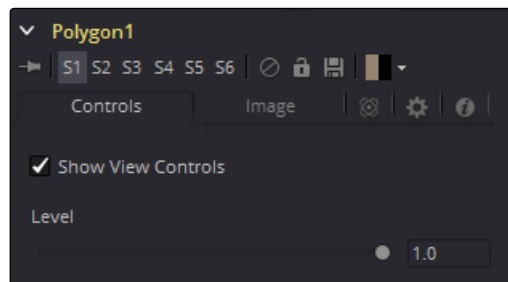
You can disable some tools, like the
Polygon tool, on a per tool basis.

Making Fine Adjustments to Controls

If you want the visual guidance of preview controls with the precision of the Control Panel, you can use a few keyboard shortcuts.

The arrow keys can be used to adjust the position of the preview control along the vertical or horizontal axis by small steps.

Holding Command (Mac OS X) or Ctrl (Windows) while using the arrow keys reduces the scale of each step by a factor of ten. Holding Shift increases the scale of each step by a factor of ten.

Toolbars

There are two toolbars in the Viewer: a horizontal toolbar called the View Toolbar, and a vertical toolbar called the Tool Toolbar. Both are used for showing buttons that are specific to the selected control.

Viewer Toolbar

The Viewer Toolbar runs across the bottom of the Viewer, providing access to many of the most commonly used Viewer related settings. Most of the buttons found on this toolbar are described in detail throughout this chapter.

The Viewer toolbar located at the bottom of the Viewer

This toolbar is dynamic, meaning the buttons that are displayed may change depending on the type of viewer. A 3D Viewer shows a set of buttons for controlling the display of lights and shadows, for example, but these buttons are not needed in a 2D viewer since there are no lights or shadows in 2D. A tooltip can be displayed by hovering the cursor over any button.

Tool Specific Toolbars

In addition to the Viewer toolbar, a vertical Tool toolbar is displayed along the side of the Viewer, depending on the selected tool. Text, masks, paths, paint strokes and the 3D environment are some of the tools that make use of this toolbar.
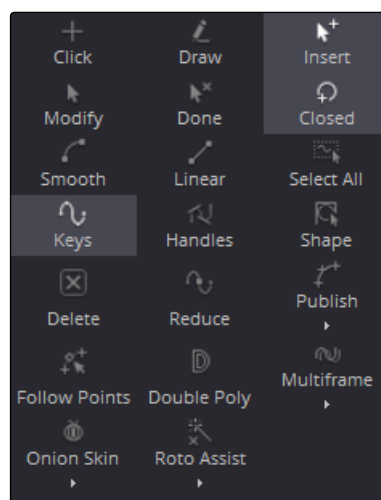
The Tool Toolbar for masks

You can customize the size and type of buttons displayed in this toolbar by right clicking on the toolbar to bring up the contextual menu.

Customizing the Toolbars

If you want to change the size of the buttons and icons in the toolbars, these settings are available in the context menu.
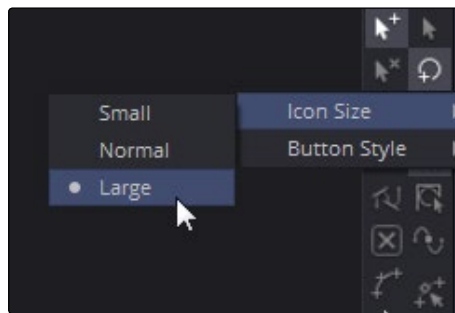
The contextual menu for the Tool toolbar

A/B Buffers

Each viewer has two buffers, each of which can contain images. Each buffer can be considered a complete and separate viewer within the same viewer pane. The default is to always use the A buffer, so when you click the Viewer indicator on a tool, the image loads into the A buffer on whichever viewer you have chosen.

Tip: Buffers can be set to different display types. For instance, you can have an A buffer showing the image with a B buffer showing a histogram.

Switching Between Buffers

Toolbar

The most common way to switch between buffers is to click the Buffer buttons located on the Viewer toolbar. Clicking the button labeled A or B switches to the associated buffer.

The Viewer Buffer buttons are used to switch between buffers

Keyboard

There are keyboard equivalents for the Buffer buttons. The comma (,) key selects the active viewer to the A buffer. The period (.) key sets the active viewer to the B buffer. If you want to see a split display of the two buffers in the Viewer, press the forward slash (/) key.

Split Wipe Images

The A,B buffers can be used to enable a split screen display, where the two buffers can be compared on screen. The displays are separated by a red splitter bar. The splitter bar can be moved and rotated in any direction or angle by dragging on the splitter. Dragging it while holding the Shift key will "snap" it to the nearest 45-degree angle.
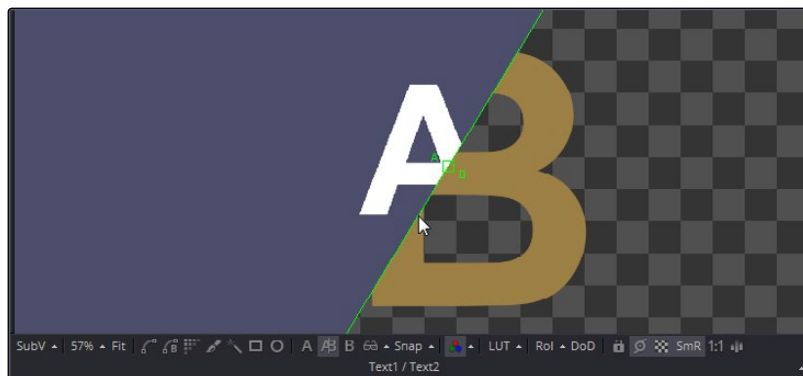
The splitter bar can be adjusted for comparing different areas of the A and B images

You enable the split wipe by clicking the A|B button, then disable it by switching to either the A or B buffer. While the split wipe is enabled, the Viewer still has the same functionality as if you are viewing just a single buffer. New tools can still be dragged and dropped into either buffer to change the image displayed. Panning or zooming one buffer will perform the same pan or zoom on the other buffer to keep them in sync.

Different channels, view LUTs, or other display options can be selected for each buffer individually. This allows easy comparison of channels, LUTs, etc. while in Split Wipe mode, or by quickly switching between the A and B buffers with the hotkeys.

Occasionally, you will have either zoomed in or panned so far from the splitter that it will no longer be visible in the view. Holding Command-Option (Mac OS X) or Ctrl-Alt (Windows) and clicking anywhere in the image will cause the splitter to jump to the current cursor.

Subviews

A Subview is a small viewer that docks within the main Viewer. It is often used to show different visual information about the image.
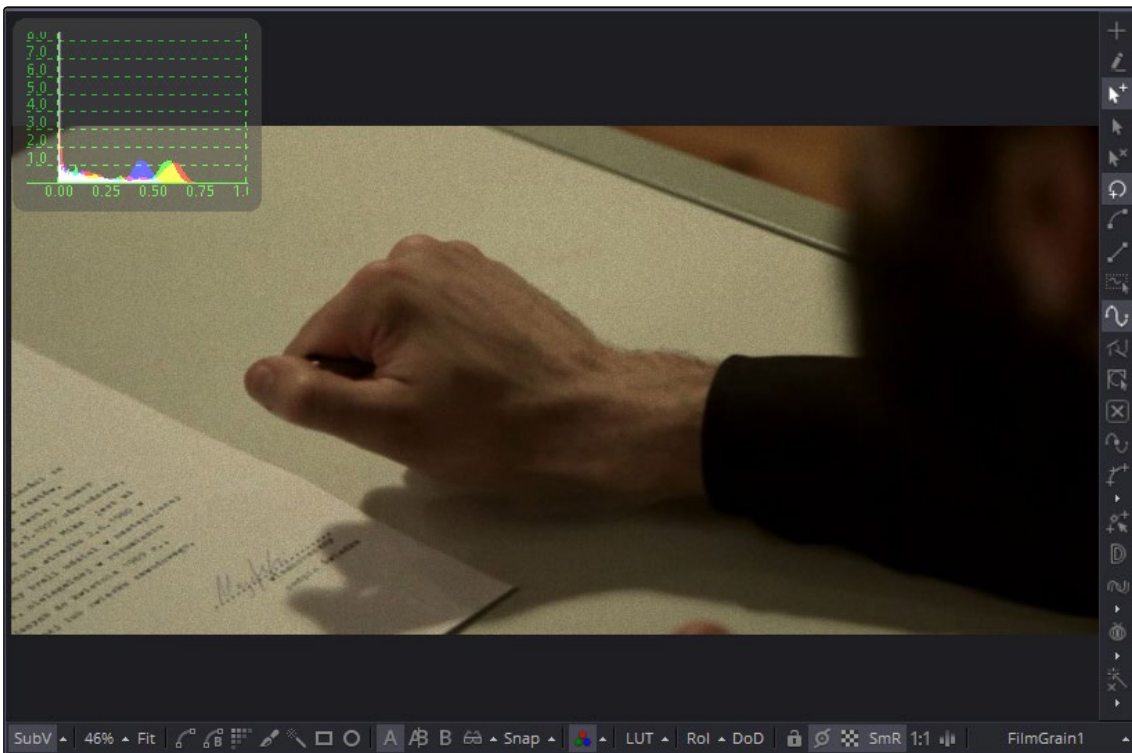
The Subview menu with the Histogram subview displayed

For example, the RGB channels can be viewed in the main Viewer, while the alpha channel displayed in a Subview. For the most part, the Subview is a fully functional mini viewer, with its own contextual menu and options. It responds to the same keyboard shortcuts and navigation controls as any other viewer. However, there are several view types that are designed for use only in the Subview, including the Navigator, Magnifier, Color Inspector and Image Info.

Displaying/Hiding The Subview

To enable the Subview, click on the Subview button in the View toolbar, or select Views > Subview > Enabled from the context menu. The V key acts as a toggle turing the Subview on and off in the active viewer.

The Subview button in the Viewer toolbar

Changing the Subview Type

The Subview button enables and disables the default Histogram subview or the last subview chosen. To change which subview type is displayed, you click the small arrow to the right of the Subview button to get the popup menu or, you can right click in a subview to bring up the Subview's contextual menu. The Subview popup menu and contextual menu show all the available subview types. Once you choose from the list, the selection will be displayed and the Subview button can then be used to enabled and disable the selected subview type.

Swapping the Subview with the Main View

Pressing Shift-V on the keyboard will swap the main Viewer and the Subview. You can also choose Views > SubView > Swap from the Viewer contextual menu. Certain view types like the Color Inspector and Magnifier can only be used in the subview. In these cases, the swap will not take place.

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 tool and some of the measurement viewers are only available as subviews. Below is detailed information about the different view types available.

2D Viewer

The 2D Viewer is the default type for showing images. When used as a subview, a different tool than the one used in the main Viewer can be displayed by dragging the tool into the Subview. This is the only Subview type that is not just a different view of the the same tool in the main Viewer.

The Subview used as another viewer

3D Image Viewer

The 3D Image viewer is available when viewing a tool from the 3D category.
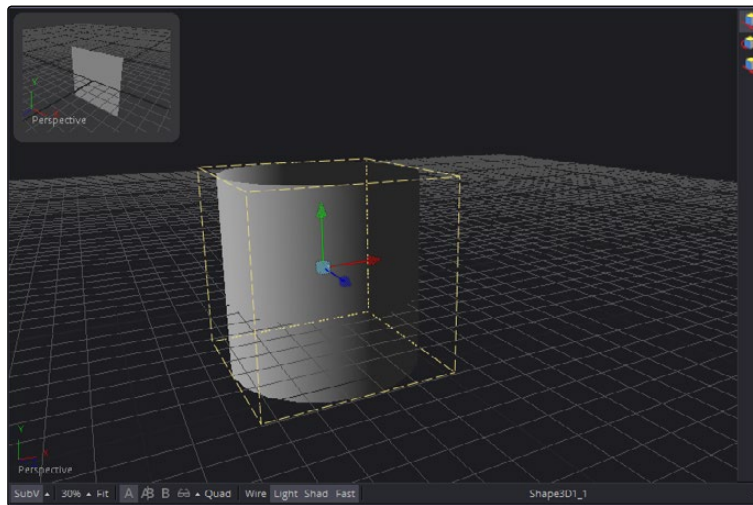
A 3D image viewer as a subview

Histogram

The Histogram is an analysis tool that can be used to identify problems with contrast and the dynamic range in an image. The graph shows the frequency distribution of colors in the image, including out of range colors in floating point images. The horizontal axis shows the colors from shadows to highlights. The vertical axis shows the number of pixels in the image that occur at each level.

The Histogram viewer will also display Gradient information. You can use the From Image and Perturb modifiers to output gradients. If you need to see the gradient represented in a histogram, drag the modifier's titlebar into the Viewer.

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

3D Histogram

This more advanced histogram type shows the color distribution in an image within a 3D cube. One advantage to a 3D Histogram is that it can accurately represent the out of range colors commonly found in floating point and high dynamic range images. It can also be used to look at vector images like position, normal, velocity, etc. Navigate the 3D cube like any other 3D view, holding Option (Mac OS X) or Alt (Windows) with the middle mouse button and dragging to rotate.

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

Vectorscope

The Vectorscope duplicates the behavior of a specific type of video test equipment, displaying a circular graph that helps to visualize the intensity of chrominance signals.

The Vectorscope viewer type for evaluating chrominance in an image

Waveform

The Waveform duplicates the behavior of a specific type of video test equipment, displaying a line or bar graph that helps to visualize the voltage or luminance of a broadcast signal.

The Waveform viewer type for evaluating luminance in an image

Navigator

The Navigator can only be used in a subview. It provides a small overview of the entire image, with a rectangle that indicates the portion of the image that is actually visible in the main Viewer. This is useful when zooming in on an image in the main view.

The Navigator subview for panning the image while zoomed in

Magnifier

The Magnifier can only be used in a subview. It shows a zoomed in version of the pixels under the cursor in the main Viewer.

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

Image Info

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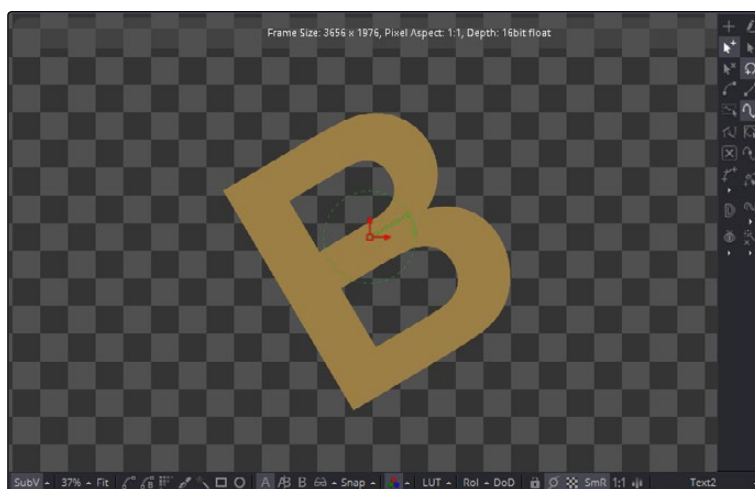
The Image Info subview for viewing size, pixel aspect and color depth information

Color Inspector

The Color Inspector can only be used in a subview. The color inspector shows information about the color channels of the pixel under the cursor. It will show all channels present, even the auxiliary channels such as Z buffer, XYZ normals and UV mapping channels.

The Color Inspector subview for evaluating colors in an image

Metadata

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

The Metadata subview for viewing embedded metadata

Choosing Color Channels

When compositing, you 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 separate color, alpha and depth channels that comprise the image.

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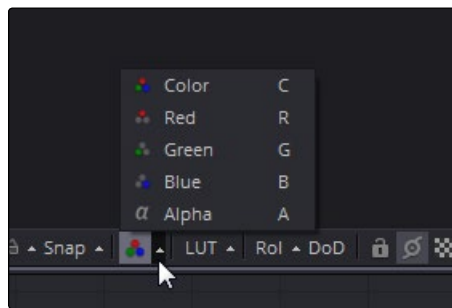
The Channel toolbar button for switching between color, alpha and depth channels

Clicking directly on the button will switch between full RGB color and alpha, or the currently viewed channel. The triangle on the right side of the toolbar button reveals the list of channels.

The following hotkeys will also toggle the channel displayed in the active view:

→ 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 Non Color Channels

The Viewers support RGBA and Z channels using keyboard shortcuts, but they support other channels as well. File formats, such as Open EXR, often include auxiliary image data that helps provide more control when compositing rendered 3D images. To view auxiliary image data in a viewer, click the arrow to the right of the RGB button to display the pop up menu or right click in the Viewer and choose Channels. From the Channels submenu select the auxiliary channel you want to view.

The 3D Viewer

Building a composite in 3D space has different requirements than traditional 2D compositing. When a tool 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 OpenGL-accelerated 3D environment.

For more information on 3D controls, read Chapter 12, Working in 3D.

Rotating, Panning and Scaling the View

For the most part, navigation in the 3D Viewer is identical to the navigation in the 2D viewer. Panning and zooming are identical but the 3D Viewer also allows you to rotate the Viewer around the scene.

To rotate in the Viewer, do one of the following:

→ Hold Option (Mac OS X) or Alt (Windows) and drag left or right using the middle mouse button.

→ Hold the middle and right mouse buttons together while dragging left or right.

The rotation is centered on the middle of the view.

Tip: The rotation controls can be used with the 3D Histogram subview as well.

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 tool, 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.

The Wireframe button in the 3D Viewer toolbar

To display 3D geometry in wireframe, do one of the following:

→ Click the Wire button in the Viewer toolbar.

→ Right click in the 3D Viewer and choose 3D Options > Wireframe from the contextual menu.

Changing 3D Viewpoints

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. 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 the 3D scene from the front, top, left or right side, collectively called Orthographic views. If you have a camera or spotlight in your scene, you can also switch to see the scene from the point of view of those objects.

To change the 3D Viewpoints, right click in the Viewer and choose the Camera submenu from the contextual menu. The choices include Perspective, Front, Top, Left and Right.

The Camera submenu in the 3D Viewpoint contextual menu

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.

There are many times you will have used the default perspective view to set up your scene but now want that point of view to become your main camera's point of view. You can easily position and orient a camera, light or other object to match the current point of view shown in the Viewer using the Copy PoV to command.

To copy the point of view in the Viewer to a Camera, do the following:

1. Set up the 3D Viewer with the point of view you want.

2. Add a camera to your 3D scene.

3. Right click in the 3D Viewer to bring up the contextual menu

4. Choose Camera > Copy PoV to > camera3D#

In the example steps above, 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 tool are not taken into account.

As you switch between view points, you are able to keep track of which one is currently displayed using the text label in the bottom corner of the Viewer. Right clicking on the label or the axis control above it acts as a shortcut to the Camera submenu, allowing you to easily choose another viewpoint.
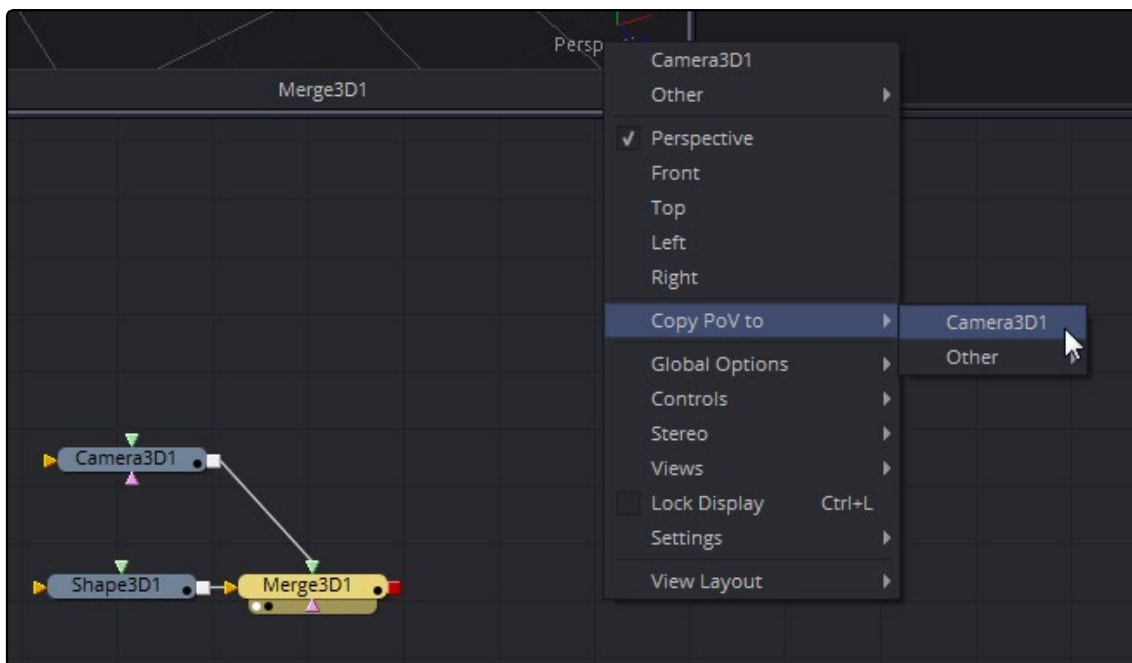
The Viewpoint label and Axis control in a 3D Viewer

Lighting

Before you add lights to a 3D scene, default lighting is provided. This basic lighting allows you to see the shading on objects without requiring you to add, and set up lights visible light icons in the viewer. To see the effects of the default light on the scene, click the Light button in the toolbar.

The Light button, enables and disables the lights in a 3D scene.

When you are ready to add your own lighting to a scene you can choose Tools > 3D > Light and then select a light type from the submenu. The lights you add must be connected to a Merge 3D tool and the Merge 3D tool, or a tool further downstream must be displayed in the 3D Viewer in order to see the effects of the lights. Once you connect a light to a Merge 3D tool, the default lighting is disabled.
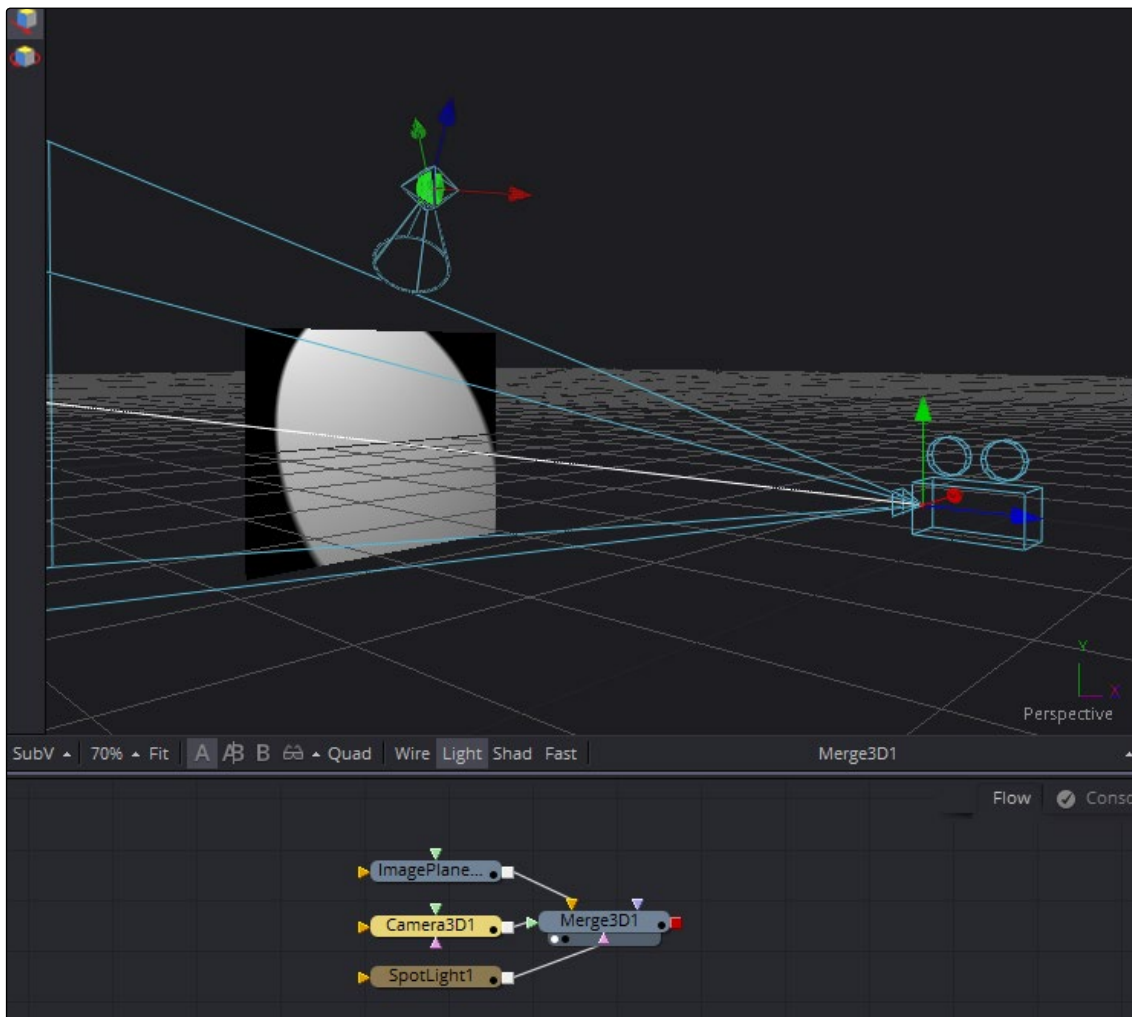
The Merge 3D tool with an image plane, camera and light connected to it

Tip: Loading a Light tool by itself into a viewer will show only an empty scene, with nothing illuminated. To see the effects of lights, view the Merge 3D that connects to the light.

Shadows

Similar to lights, the default 3D Viewer has Shadows turned off. To see shadows in the 3D Viewer, you must click the Shad button on the toolbar or right click in the Viewer and choose 3D Option > Shadows.
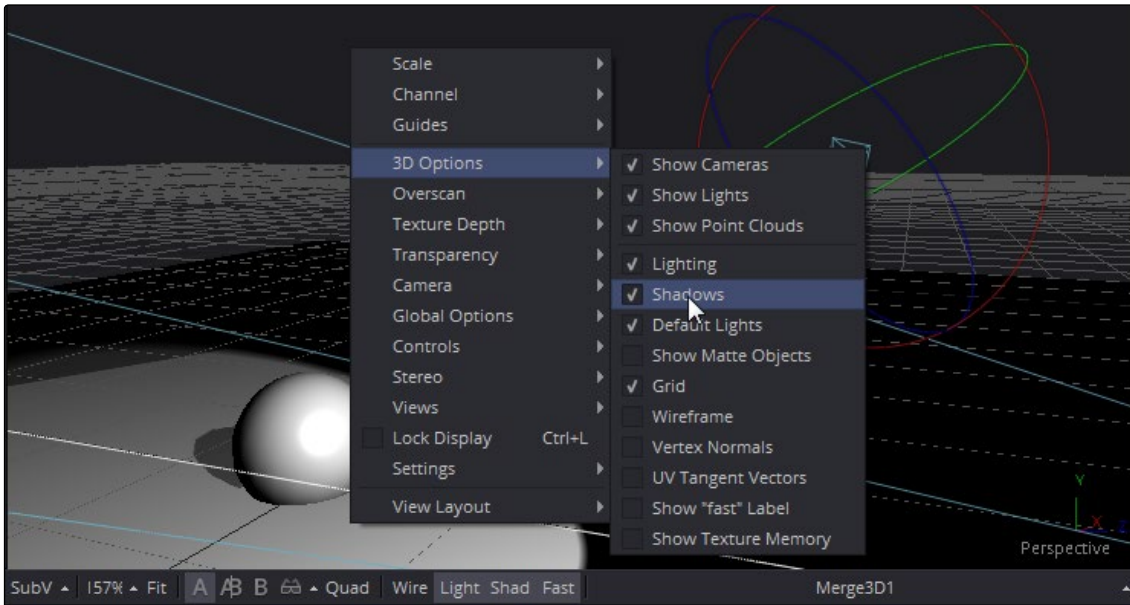
The Shadows command in the 3D options contextual submenu

Enabling shadows will automatically turn on lighting, if it is not already turned on. The shadows shown in the display view will be hard edged. Soft shadows are available in the software renderer.

Transparency

Image planes and 3D objects are obscured by other objects in the scene, based on their X, Y and Z position coordinates in 3D space. The process of figuring out which polygons are hidden and which are shown based on those coordinates is called Z-buffering. This process is extremely fast, but not always accurate when dealing with multiple transparent layers in the scene.

Although Z-buffering is the default, you also have another option for those more complex 3D scenes. The Sorted method can be significantly slower in some scenes, but provides accurate results no matter how many layers of transparency are in the scene.

The default behavior in the Viewer is to use Z-buffering, but if your scene requires the Sorted method, you can click the Fast button in the Viewer toolbar to disable Z-Buffering, or right click in the 3D Viewer and choose Full Sort or Quick Sort from the Transparency submenu. Quick Sort reorders the polygons in the scene serially, from back to front to get a more accurate rendering of transparency. Full Sort renders every polygon in Z order to get the most accurate rendering of transparency.

The Fast button in the 2D Viewer toolbar

Grid

The 3D Viewer includes a Grid that is used to help provide a point of reference in the 3D scene. 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 and small squares of 0.25 units each. To disable the display of the grid, right click and select Grid from the 3D Options sub context menu.

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 in the Viewer and choose 3D Options > Vertex Normals from the contextual menu.
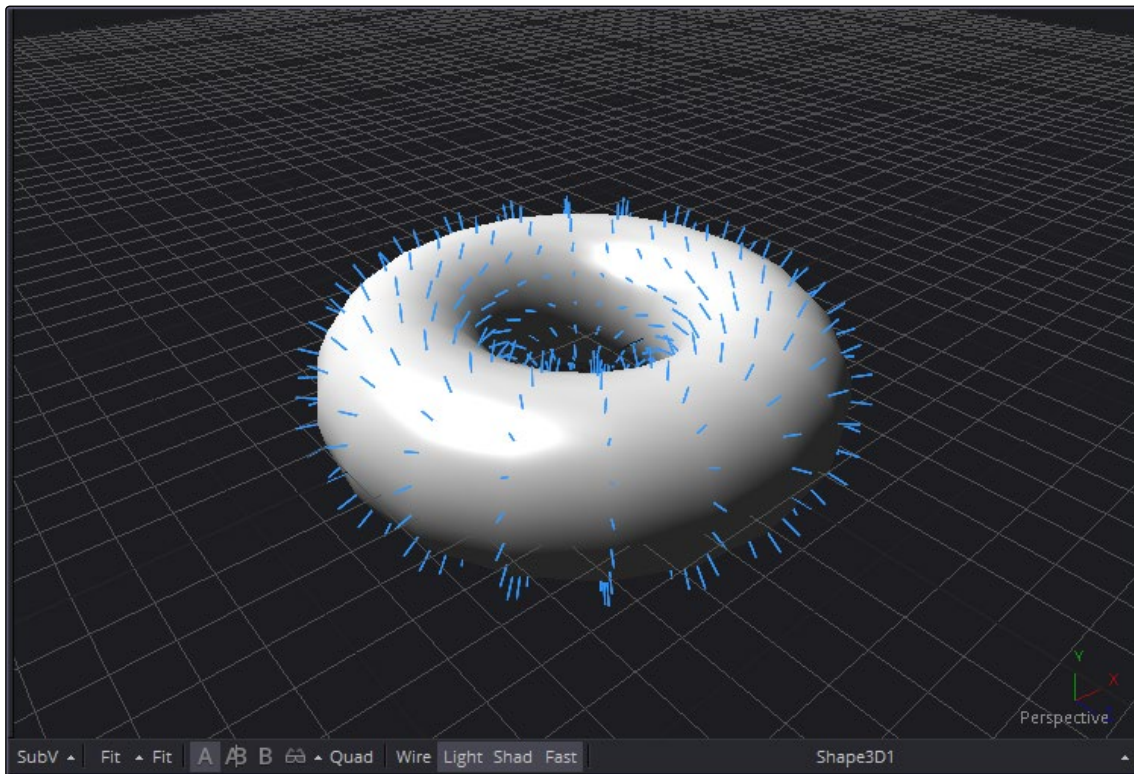
The Normals viewed in a 3D scene

Quad View

3D compositing requires you to view the scene from different points of view. You can switch the Viewer to different points of view or you can divide the Viewer to show four different angles at one time. This is called a Quad view.
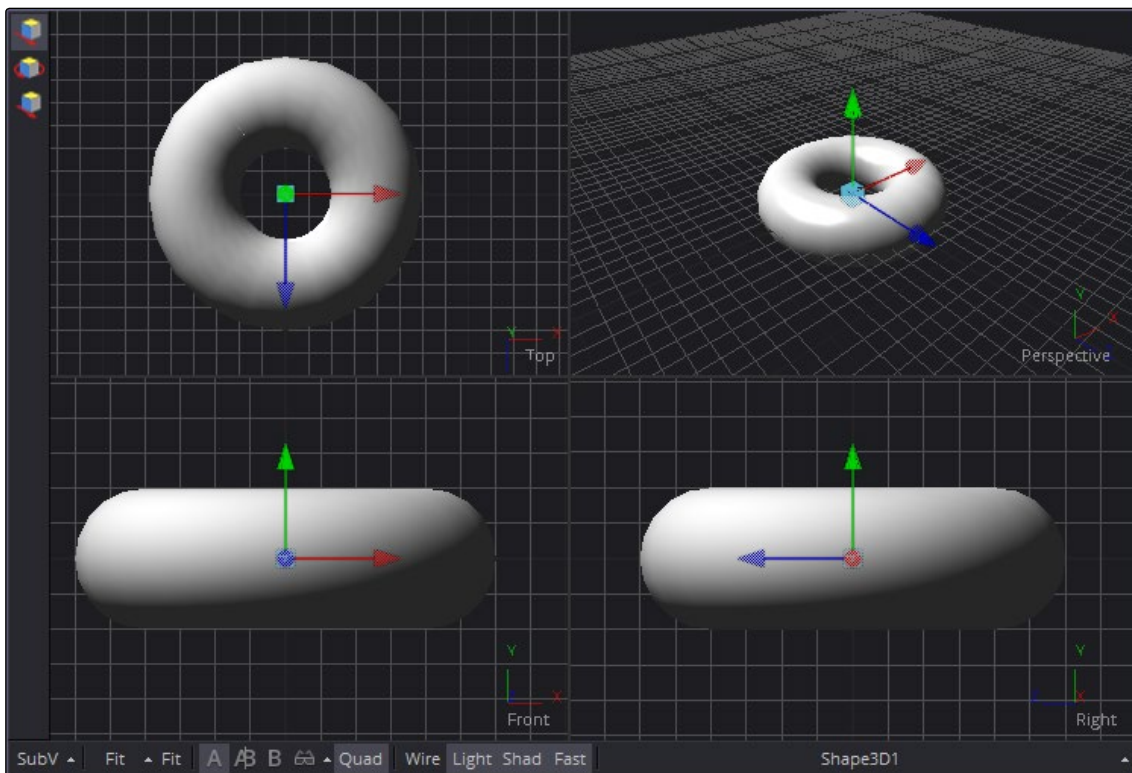
A Quad view of a 3D scene

Although there are four parts to the Quad view, all four views are of the same scene. When using a Quad view to display a 3D environment, you choose between displaying the front, left, top, bottom and perspective orthographic views, as well as the view through any cameras or spotlights present in the scene.

Quad views can also be used with a 2D scene, each part can be used to show a different view type. For example, one quadrant can show the image while the other quadrants show the alpha channel, vectorscope and histogram.

To toggle the display of the Quad view, do one of the following:

→ Right click in the Viewer, then choose Views > Quad View from the contextual menu.

→ Press Command (Mac OS X) or Ctrl (Windows) -Q on the keyboard.

→ Click the Quad button on the 3D Viewer toolbar.

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 in the Viewer, then choose Views > Quad Layouts.

3. Choose a layout from the Quad Layouts submenu.
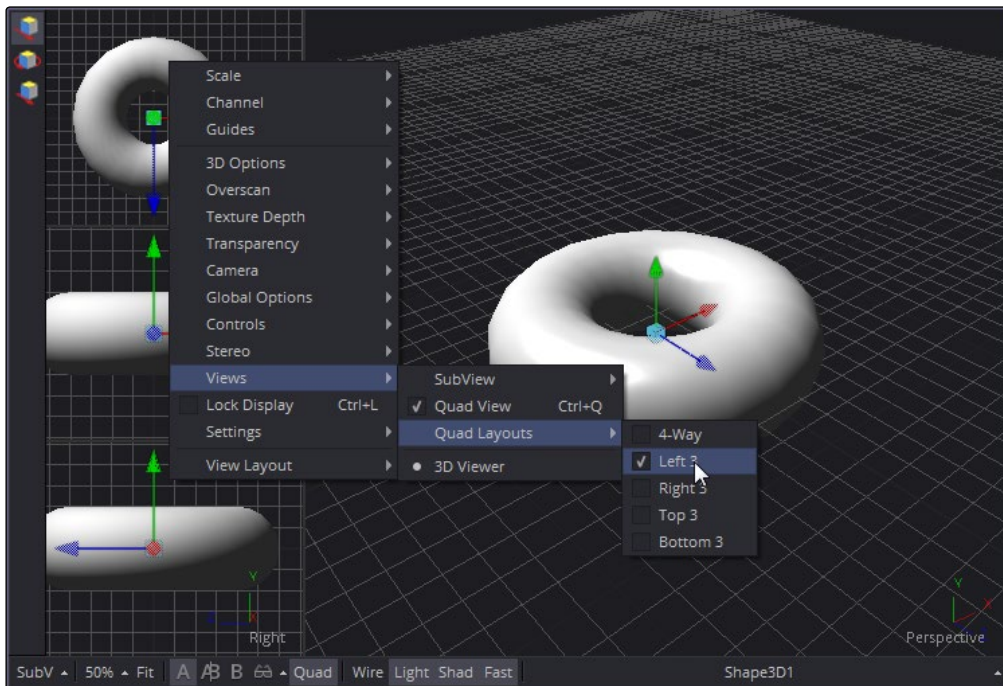
A Quad view layouts submenu

Effect Masks

Blurs, color correction and warps are typically applied to the entire image. However, there are times when it helps to be able to restrict the effect of a tool to a specified portion of an image. Effect masks can be polylines, basic primitive shapes, paint strokes or bitmaps from other tools that define a shape, which you can use as the boundary for effects.

The Polyline, B-Spline, Bitmap, Paint, Wand
and Shape Effect Mask buttons

To add an effect mask to the currently displayed tool, do one of the following:

→ Click on the button in the Viewer toolbar for the desired mask type.

→ Right click in the Viewer, then choose Effect Mask from the contextual menu.

The toolbar buttons can only be used to add effect masks. Garbage and pre-masks are added via the Viewer's contextual menu. For more information, read Chapter 10, Polylines.

Guides

Guides are used to help you compose elements within a boundary or along the center vertical and horizontal axis. There are four commonly used guides that can be displayed in the viewers: Monitor Safety, Safe Title, Center and Film.
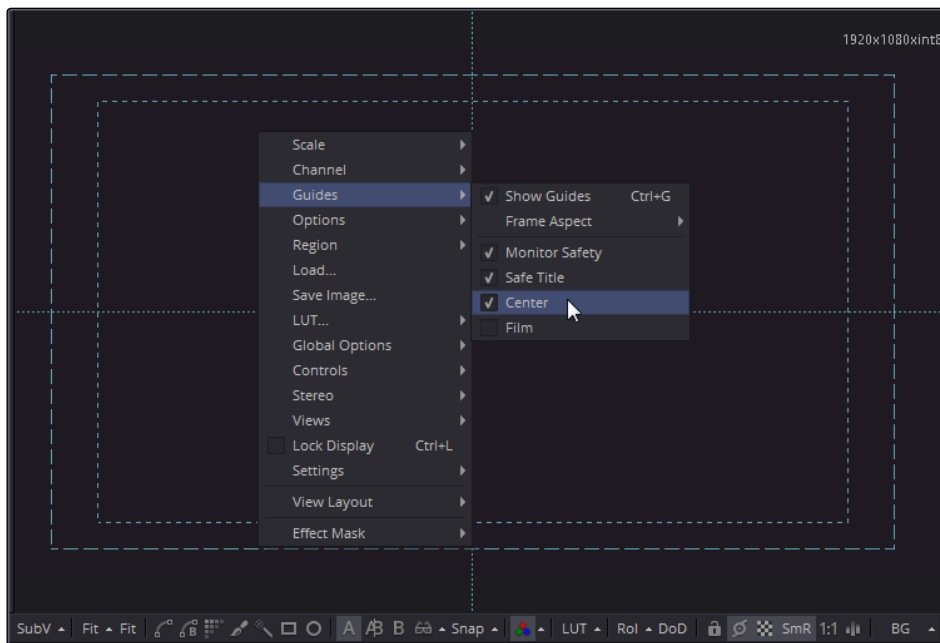
The Guides submenu in the Viewer's contextual menu

To enable or disable Guides, do one of the following:

Right click in the Viewer, then choose Guides > Show Guides from the contextual menu.

Press Command (Mac OS X) or Ctrl (Windows) -G to enable or disable the guides in the Viewer.

Using the Guides submenu you can choose to show or hide any of the four guides in the Viewer. Each guide is used for a specific purpose.

Monitor Safety

Monitor Safety indicates the safe action area viewable on most monitors and TV screens.

Safe Title

Safe Title indicates the safe area for titles viewable on all TV and monitor screens.

Center

Center shows a cross hair for the Center point and X- and Y-axis of the view.

Film

Some frame formats include film guides preset for you, whereas some will require customization. The film guides can be customized in the Preferences > Frame Format window.

In the Frame Format window, there are two film guides. Guide 1 contains four text boxes. These boxes specify the offset from the edges of the image for the left, top, right and bottom guides, in 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

Guide Frame Aspect

To find out how the current image compares to a specified frame format, you can choose from a list of frame aspect ratios in the Guides submenu of the Viewer's contextual menu. The frame aspect is usually set to Default. This forces the frame aspect to the same resolution as the image displayed in the view.

The Guide > Frame Aspect submenu in the Viewer's contextual menu.

When the frame aspect is set using the Guides > Frame Aspect context menu, the guides will represent the exact boundaries of the specified format and any image area outside of that will be dark gray.

Managing Look Up Tables (LUTs)

Look Up Tables, or LUTs, are used to help match the appearance of a viewer to its eventual output destination. It is essentially a color correction tool that affects only the Viewer, not the image data. There are two basic ways that LUTs can calculate color transformations: the first is a simple 1D LUT, while the second is a more sophisticated 3D LUT.

1D 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.

3D LUT

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 Look Up Table (LUT) is a table of values used to transform the color and luminance of an image. A 1D LUT uses a two column table for input color and output color, while a 3D LUT uses more of a matrix. A LUT is primarily used to correct for variances in the monitor or the source color space of the image. You can choose to apply a LUT to all the viewers or apply individual LUTs to each 2D viewer.

Image LUTs

Image LUTs are applied to each viewer, including separate LUTs for the A and B buffers. These LUTs can only be applied to 2D images and not 3D scenes. Image LUTs are routinely used to get from one scene referred colorspace to another. They are useful when you want to know how your image will look on a different device or standard. They're often tool or task specific.

Buffer LUTs

The Buffer LUT is applied to all of the viewers, regardless of contents, including 3D scenes, 3D materials, and subview types. Only one buffer LUT can be applied. If a 2D image is being displayed with an image LUT applied, then the buffer LUT is applied to the result of the image LUT. Buffer LUTs are typically used to account for the particular monitor you are sitting in front of. For instance, getting from DCI-P3 projector setting to an sRGB monitor setting.

When dealing with nonlinear files from many of today's digital cinema cameras, a modern workflow would be to put everything in linear from the Loader's Control Panel, then perform your color correction and compositing in the Flow Node Editor. At the end apply an image LUT or buffer LUT that matches the output device profile

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

Types of LUTs

Aside from the industry standard 1D and 3D LUTs, other types of LUTs are supported, including script based fuse tool LUTs, and macros assembled from standard tools. 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 tools they contain.

Fusion View LUT

The Fusion View LUT is the default and most frequently used LUT type. It provides an RGBA curve that can be used to assign IN/OUT value pairs. This control is identical to that provided by the Color Curve tool.

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 tool.

Macro LUTs

Any macro tool can also be used as a LUT, simply by saving the macro's .setting file to one of the directories specified by the LUTs Path Map found in the Globals > Path Map preferences.

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. To learn more about Macros read Chapter 14, Scripting, Macros and Expressions.

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. To learn more about fuses read Chapter 14, Scripting, Macros and Expressions.

Applying Image LUTs

The LUT button in the Viewer toolbar is used to toggle the image LUT on and off. The default LUT type is a standard Fusion view LUT, which is initially flat, so when you enable it, there will be no apparent change to the image. To select a different LUT type, click on the arrow to the right of the button and a LUT menu will appear with a list of all known LUT types, including macros and fuses.

The LUT menu can also be found as a submenu in the Viewers's contextual menu.
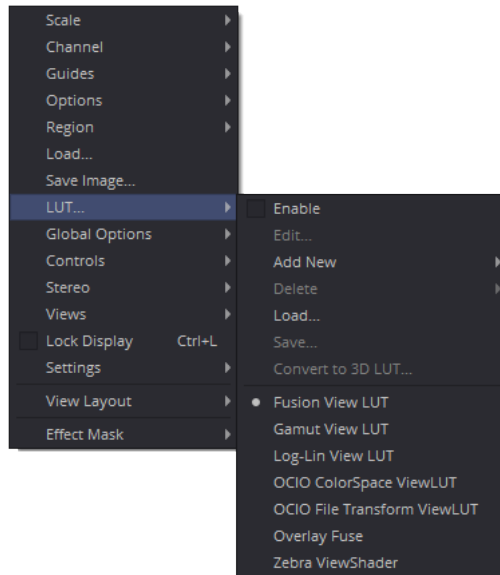
The LUT popup menu, which can be used to change the image LUT applied to the Viewer

To remove the LUT, click the LUT button again.

Applying Buffer LUTs

The buffer LUT is set for all viewers by right clicking in the Viewer and selecting the Global Options > Buffer LUT > Enable option. The Buffer LUT submenu contains a list of all known LUT types, including macros and fuses.
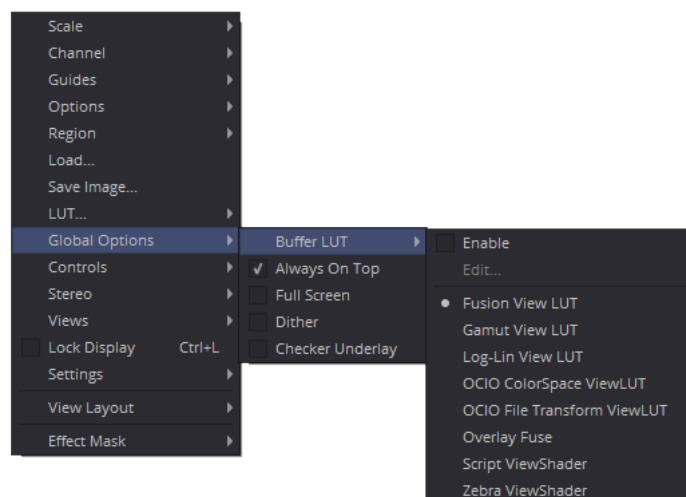
The Buffer LUT option found in the Viewer's contextual menu.

Buffer LUTs are often useful for applying monitor corrections, which do not usually change between projects.

To remove a Buffer LUT, choose Global Options > Buffer LUT > Enable.

Opening the LUT Editor

The Viewers are the primary area where color correction is assessed, so it is crucial that they give an accurate representation of what the content will look like when it is played back for an audience. The LUT Editor allows you to modify the separate color channels to match the gamma and color characteristics of your eventual playback device.

To open the Look Up Table Editor, do one of the following:

1. Click the LUT button in the Viewer toolbar to enable it.

2. Choose Edit from the LUT popup menu in the Viewer toolbar.

    or

3. Right click in the Viewer, then choose LUT > Edit from the contextual menu.

The LUT Editor for the default Fusion View LUT.

As with the Color Curves tool, 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 Gamma and Alpha Gamma sliders are used to duplicate the gamma values of the eventual output device. Video monitors, for example, commonly have a gamma of 1.7, while computer monitors can range anywhere from 1.6 to 2.2. Alpha Gamma is only applied when viewing the alpha channel of an image, or when viewing masks.

LUT Processing Order

In elaborate workflows, facilities can have multiple LUTs applied in a row before the image is seen. The order of these is important since each LUT delivers different outputs. For instance, for a Cineon file in Log color space you may often apply 3 LUTs. First a Log->Lin conversion, followed by a Fusion View LUT to apply a color calibration and a third one to correct it for display on an sRGB monitor or replace the last with a 3D DCP LUT if you are viewing on a projector.

When you select a tool 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 preview 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 in 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 in 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 view 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 tools in Fusion to create macros, which can then be saved and used as a LUT.

LUT Settings

The most straight forward way to save a LUT you have created using the Fusion View LUT Editor, is to use the LUT > Save menu found in the Viewer contextual menu. The settings are saved as an ASCII file with the extension .viewlut in the LUTs folder. Any files with this extension found in that folder will appear in the Image LUT menus, for ease of loading. You can also load the settings that are not found in the menu by choosing LUT > Load from the Viewer's contextual menu.

Using Viewer Settings

If you've modified a LUT, choosing Settings > Save New from the Viewer's contextual menu will save all the Viewer's settings, including all LUT curves and gain/gamma values. You can save these under different names, and each settings file can be reloaded at any time by choosing Settings > File name from the Viewer's contextual menu. Choosing Save Default from the same menu will make these settings the standard for all new comps.

Using LUT Curves

The View LUT edit window can be used to import and export LUT curves.You can export the LUT curves as either ASCII or Saved format. The ASCII (.alut) file format is useful for sharing LUT curves with other software, whereas the Saved (.lut) file format is preferred for Fusion, as it is more compact, accurate and allows further editing. For details on the ASCII LUT file format, see the ASCII Import Appendix.

To export a LUT, do the following:

1. Click on the LUT button in the Viewer toolbar to enable it.

2. Click on the LUT popup menu in the Viewer toolbar, then choose Edit.

3. Right click on the LUT Curve Editor, then choose Export LUT.

4. Select a LUT format at the bottom of the file browser window.

5. Enter a name for the LUT and click Save.

The Import LUT option will load LUT files back into the Curve Editor, or alternatively, if the file has been saved in Fusion's LUTs folder (as specified in Fusion > Preferences > Global > Path Maps > LUTS on Mac OS X or File > Preferences > Global > Path Maps > LUTS on Windows), it will appear in the LUT popup menu list. This is one way to move LUTs between viewers or to and from the Color Curves tool or any other LUT Editor in Fusion.

LUT Files

Any supported LUT files in the LUTs folder can be used by choosing it from the LUT popup menu, or the Viewer's contextual menu. This includes 1D and 3D LUTs such as Fusion's .lut, .alut and .alut3 formats, as well as .shlut, .look, .3dl and .itx formats. This is a convenient way to access standard format LUT files for different projects.

Settings and Macros

Since LUTs are a form of color correction, you can also use any tool, macro or group of tools as a Viewer LUT.

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

1. Select the tool, group or macros.

2. Right click over the selected tool, 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 tools to be used as a Viewer LUT. This is the most flexible approach, but is also the slowest. The LUT tools must be rendered solely on the CPU, whereas the other methods are GPU-accelerated.

Setting a Default LUT

The default LUT applied when a new composition is created can be assigned in Preferences > Globals > View panel. Clicking the Enable Display LUT checkbox allows you to select a LUT from the Display LUT plug-ins list.
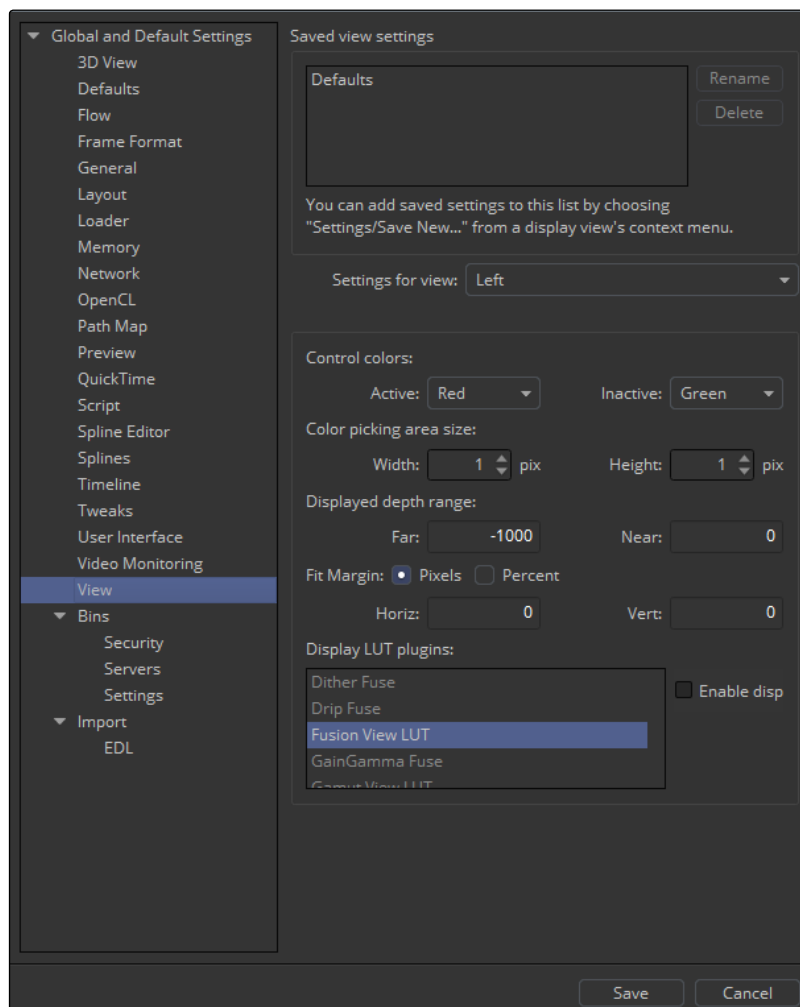
The LUT Default settings found in the Globals View preferences.

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.

For details on using the Layout and Viewer preferences, read Chapter 15, Preferences.

Viewer Settings

It is often preferable to switch between entirely different viewer configurations while working. For example, while keying, the image maybe in the main Viewer and the alpha channel may be in a subview. Viewer settings toward the end of a project may consist of the histogram, vectorscope and waveform, as well as the image in a view set to Quad view.

Fusion provides the ability to quickly load and save Viewer settings to help reduce the amount of effort required to change from one configuration to another.

To Save a Viewer setting, do the following:

1. Right click over the Viewer you want to save.

2. From the contextual menu, choose Setting > Save New.

3. Enter a name for the settings and click Save.

To Load a Viewer setting, do the following:

1. Right click over the Viewer you want to load a setting into.

2. From the contextual menu, choose Settings > Settings name.

Loading and Saving Defaults for a Viewer

The Viewer can save new defaults and be returned to its defaults using the Load Defaults and the Save Defaults options in the Settings portion of the View contextual menu.

Viewer Options

The Options contextual submenu contains several ways you can customize the look and behavior of the Viewer. Several of the options also have buttons in the View toolbar as well.
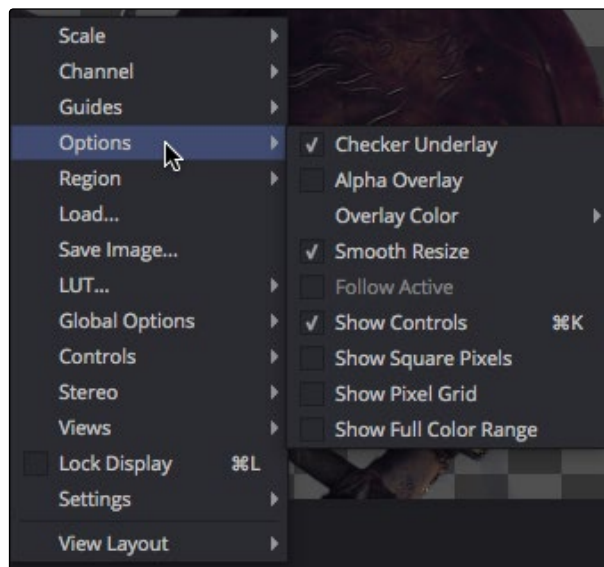
The Viewer's Options submenu

Locking the Viewer

If you lock the Viewer you can prevent it from updating. The tool still processes and the new image is queued for display in the Viewer, but until you unlock it, the Viewer does not update. You can lock the Viewer using the Lock button in the Viewer toolbar or using the Command (Mac OS X) or Ctrl (Windows) -L keyboard shortcut. By default the Viewer is unlocked.

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.

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. You can disable this option by clicking on the Checker button in the 2D Viewer toolbar to replace the checkerboard with a black.

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 when you zoom in at 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.

Follow Active

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

Show Controls

When preview 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 the Command (Mac OS X) or Ctrl (Window) -K keyboard shortcut or by clicking on the Show Controls button in the View 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 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.

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. Toggle this option with the Grid toolbar button. The default is Off.

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.

Viewer Labels

Each viewer has a label in the bottom right which indicates the name of the tool currently loaded in that viewer. Clicking on the label will select the viewed tool in the Flow Node Editor.

The Viewer label shows the currently active tool's name

Controls List

The arrow button to the far right of the label displays a popup list of the preview controls for the selected tools. The currently active control is marked on the menu. The menu lets you choose a different control for onscreen use. This can come in handy when several controls are in the exact same overlapping position on the screen.

Pressing Tab or Shift-Tab will cycle through the preview controls, activating each in turn. The Controls submenu in the Viewers contextual menu has further options for selecting, pinning, hiding and scaling preview controls.
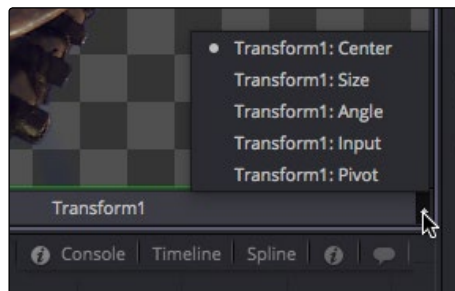
The Controls list for a Transform tool

Status Bar Information

The status bar at the bottom of the Fusion window will provide the exact RGBA and Z values for the pixel beneath the cursor in the viewers. Additional information about the X and Y coordinates of the cursor and the exact pixel position are also displayed in the Status bar.

The Status bar showing coordinates and color information

Time Ruler

6

Time Ruler

Time Ruler Overview

The Time Ruler, located beneath the Work Panel, includes transport controls for playback as well as for setting composition options like proxy level and motion blur. The Time Ruler's playhead determines the frame that is displayed in the Viewers.

The Time Ruler uses a playhead to move to a specific frame

Current Time

The Current Time edit box indicates which frame of the project is currently displayed. A red vertical line called a Playhead, is used to mark the same time in the Ruler. You can click anywhere in the Ruler to move the Playhead to that time or enter a new frame number in the Current Time edit box.

The Current Time field shows the frame that is currently displayed

When tools are displayed in the Viewers it is the current time that is rendered and displayed.

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

→ Press [to move 1 frame backward.

→ Press] to move 1 frame forward.

→ Press Shift- [to move to the project's Global Start Frame.

→ Press Shift] to move to the project's Global End Frame.

→ Press Option- [(Mac OS X) or Alt - [(Windows) to jump to the active tool's Global In.

→ Press Option-] (Mac OS X) or Alt -] (Windows) to jump to the active tool's Global Out.

Tip: If the currently active tool is animated, the Option- [or] or Alt- [or] keys will jump to the previous or next keyframe. If no tools are active, then this shortcut jumps to the next time with a keyframe on any of the animated controls in the composition.

Keyframe Ticks

When animated tools are selected, green vertical ticks on the Time Ruler will indicate the keyframe positions.

The Time Ruler with keyframe tick marks.

Cache Indicator

When assembling a flow, images are rendered for display in the Viewers. These images are stored in a RAM Cache so they can be displayed instantly at a later time without having to re-render them. You do not have to do anything since caching happens automatically. When the size of the cache reaches the limits of available memory, lower priority frames are eliminated to make space for new images.

The frames cached for a tool are indicated by a green line under that frame in the Time Ruler. The indicator only displays the cache regions for the actively selected tool in the flow.

The green line indicates frames that have been cached for playback

When you change the comp quality or proxy scale, the cache is not discarded; it is marked with a red line instead. This lets you know that if you return to the prior quality setting or proxy scale, the cache is kept. However, if you render those frames using the new quality or proxy scale, then the old cache will be overwritten.

In some cases, changing the Quality setting will not cause the Cache Indicator to turn red. When the cached frames are at the HiQ setting and you turn HiQ off, the green frames won't turn red. The HiQ cached frames will be used even though the HiQ setting has been disabled.

Ranges

The Time Ruler uses two different frame ranges, one for the duration of the entire project and one to determine what to render or what to cache in memory for previews

Global Start and End Range

The Global Start and End Range sets the total duration of a project. You cannot move the current time outside the global range. Coming from a color grading and editing application like DaVinci Resolve, the Global Start and End range is the equivalent of the Timeline duration.

The easiest way to set the Global Start and End range of a comp is to drag a tool from the Flow Node Editor to the Time Rule. This will set the Global Start and End range to the duration of the tool. You can also type in the frame numbers you want to use for the Global Start and End range using the edit boxes at the far left and right ends of the Time Ruler.

The Global Start and End Range edit boxes

Render Start and End Range

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

The easiest way to set the Render range is to drag a tool from the Flow Node Editor to the Time Ruler. This will set the Render Start and Render End range to the duration of the tool. You can also drag the ends of the shaded range in the Time Ruler, press Command (Mac OS X) or Ctrl (Windows) and drag a new range, set the range from the Time Ruler's contextual menu, or type the numbers into the inner left and right edit boxes.

The Render Start and End Range edit boxes

Visible Range Slider

The Visible Range slider found beneath the Time Ruler indicates what frames in the composition are currently shown in the Time Ruler. While this is typically left at the default setting to show the entire global range, on longer comps it can be desirable to show only a small portion of the composition's global range. Dragging the left and right edges of the Visible Range slider can adjust the visible range to focus in on a specific area of the composition in the Time Ruler.

Drag the left or right edge of the Visible Range
slider to zoom in on a range in the Time Ruler

Tip: Holding the middle mouse button and dragging in the Time Ruler will pan the visible range. This can be handy when you're zoomed in and need to move ahead to the next groups of frames.

Render Button

The Render Button opens the composition's Render Settings dialog, where you can confirm the render options and start rendering the composition to disk. Holding Shift while clicking the Render button will skip the dialog and use the default values for the rendering with full resolution, high quality and motion blur enabled.
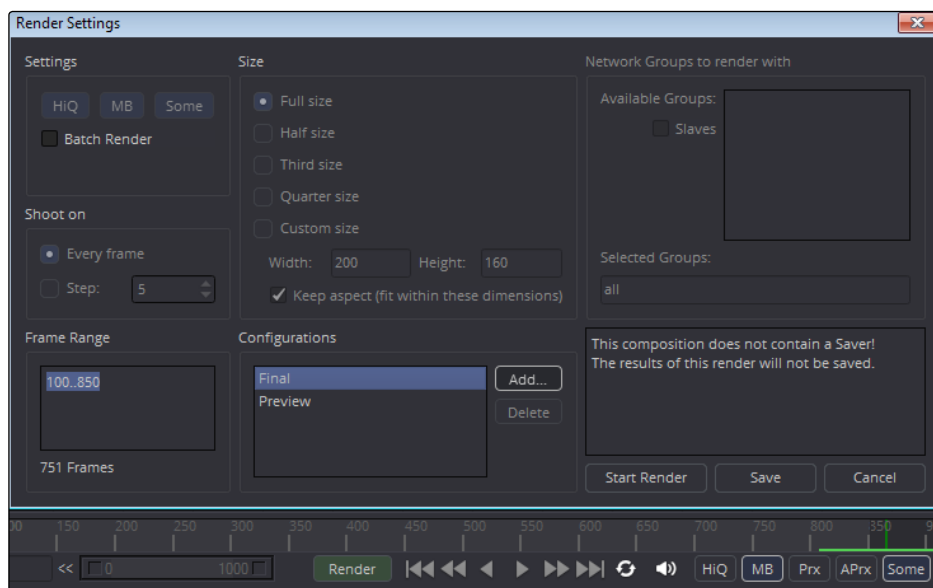
The Render Settings dialog brought up by clicking the Render button

Transport Controls

The Transport Controls in the Time Ruler can be used to play the composition forward or in reverse, to advance current time to the next or previous frame, or to jump current time to the start and end of the Render range. In this context, playback means the flow at the current frame will render into the Viewer, then advance the current time to the next frame (or previous frame) and repeat. There is no guarantee of real time playback at the project's full frame rate, unless the frames are already cached in memory.

The transport controls

Right clicking on either the frame advance or previous frame button will display a contextual menu that can be used to set a frame increment for that button. This frame increment setting will be used when you click the frame advance buttons or use the keyboard shortcuts. This is often useful for rotoscoping on fixed increments, or inspecting a frame one field at a time.
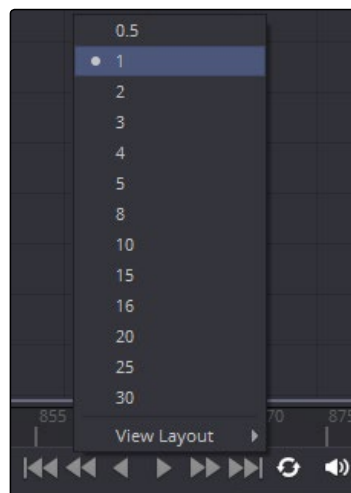
Right click to change the frame advance increment

Tip: Changing the frame increments on any transport button changes it for all transport buttons, including playback.

The loop button is a tri-state button that determines whether playback will stop when the playhead reaches the render start or end, loop back and start again, ping pong back and forth or stop completely.

Audio Mute

The Audio button is a toggle that turns on and off audio playback. Additionally, if you right click on this button, a contextual menu will allow you to select a .wav file that will play along with the composition. The same menu allows you to assign an offset to the audio playback so you can adjust when it begins.

Right click the Audio button to select an audio file to play with the composition

Composition Quality Options

The Composition Quality options consist of two buttons. The High Quality button and the Motion Blur button.

The Composition Quality options in the Time Ruler

High Quality

When working interactively to create a flow, the quality of the images in the displays is often less important than the speed with which they are shown. As a result, you have the option to render images for better interactivity or at final render quality. When you disable the High Quality (HiQ) setting, more complex and time consuming operations such as area sampling, anti aliasing and interpolation are skipped to produce results more rapidly. Enabling HiQ mode will force a full quality render identical to that produced during a final render.

Motion Blur

When the Motion Blur button in the toolbar is toggled off, the motion blur options enabled on any of the tools in the composition will be ignored. This can result in significantly quicker renders and previews, since the sub frames required to produce motion blur will no longer be calculated.

Proxy and Auto Proxy

Proxy

For increased performance, the Proxy (Prx) button in the Time Ruler reduces the resolution of the images that are processed. Enabling proxy mode causes only one out of every x pixels to be processed, instead of processing every pixel. The value of x is decided by right clicking on the Proxy button and selecting a proxy ratio from the button's contextual menu.

The Proxy Ratio menu

If a lower resolution version of media is provided in the Loader's Proxy Filename control, then this footage will be loaded when the Proxy option is enabled, rather than the full resolution media.

Auto Proxy

The Auto Proxy button enables the image to degrade only while adjustments are made to controls, snapping back to normal resolution when the control is released.

For example, while dragging a center control, the image may only process at a third of the resolution so that updates to the screen can keep up with adjustments. Once the center control is released, the image returns to full quality to confirm the change.

You can set the auto proxy ratio by right clicking on the Auto Proxy button and selecting it from the contextual menu.

The Auto Proxy Ratio menu

Selective Updates

Fusion normally renders only the tools that are needed to display the images in the views. The Selective Update button in the Time Ruler opens the Render Settings dialog to select the mode used during final and network renders.

There are three modes of selective update that are cycled through by clicking the Selective Update button in the Time Ruler.

The Selective Update button set to Some by default

Some

During final renders, the Some setting causes only tools that directly contribute to the output to be rendered. This is the default mode of selective update called Some, because only some of the tools in the flow are rendered.

None

The None setting will prevent rendering of all tools in the flow, which can be handy for making a large number of changes to a slow composition. When the Selective Update button is set to None, the Flow Node Editor, Timeline Editor and the Spline Editor will be highlighted with a red border to indicate that the Tools are not being updated to reflect changes taking place in the composition.

All

The All setting forces all of the tools in the flow to render. This is primarily used when you want to update all of the thumbnails in the flow.

Changing the Display Format

The Current Time edit box and the markers in the Time Ruler typically count in frames, however this can optionally be configured to display in either SMPTE timecode or in Feet+Frames for film. To change the current format choose View > Show Timecode. Once the Show Timecode option is enabled, you can use View > Timecode Format to select SMPTE or Feet+Frames as the display.
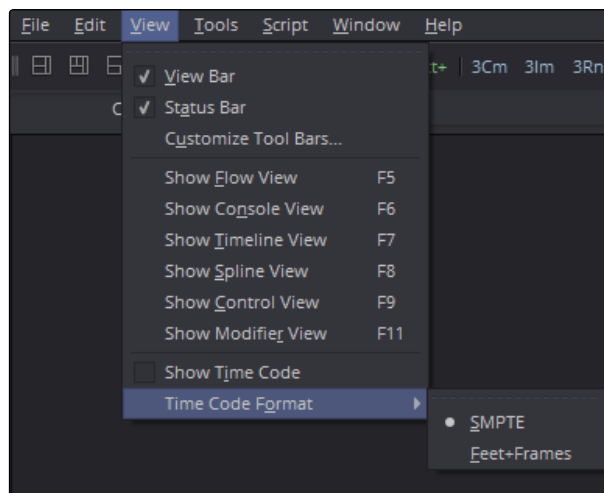
The View menu Timecode Format submenu

The options in the composition's Frame Format preferences panel determine the frame rate used for timecode, or the number of frames found in a foot of film.

Tip: Many edit boxes can use mathematical expressions. For example, typing 2.0 + 4.0 in most edit boxes will result in the value 6.0. The Current Time edit box SMPTE timecode uses the + symbol as a separator symbol, so this edit box will not correctly evaluate mathematical expressions that use the + symbol, even when the display format is set to Frames mode.

Entering Frame Numbers

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

Timeline Editor

7

Timeline Editor

Timeline Editor Overview

The Timeline Editor is used to adjust the timing of elements in a project. You can trim and extend clips, adjust the timing of an animation spline, or trim the duration of a tool. You can freely rearrange the order of tools in the Timeline without affecting the layering order of your comp. All compositing operations are handled in the Flow Node Editor, while the Timeline Editor manages the timing of your composition.
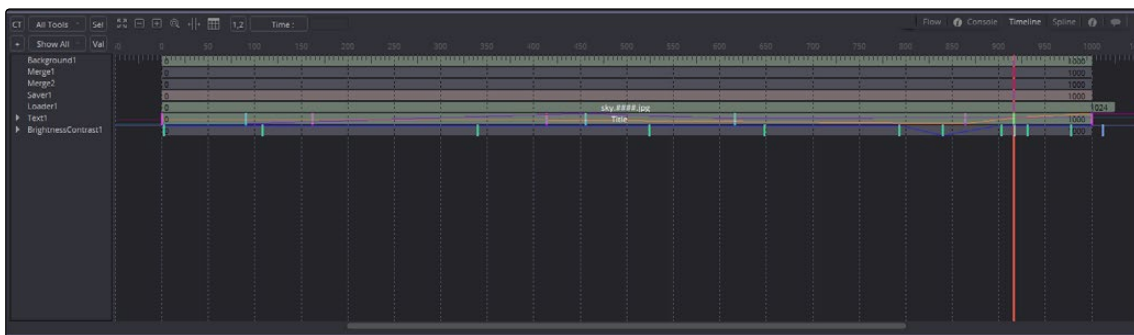
The Timeline Editor

To show the Timeline Editor, do one of the following:

→ Click the Timeline tab located at the top of the Work area.

→ Press F7 on the keyboard.

Timeline Tracks

The Timeline tracks are the main part of the Timeline interface. The horizontal axis represents the duration or time for each tool and it is usually labeled with frame numbers. The segments that run along the horizontal axis represent the time range of each tool and use the same color properties as the Tool tiles in the flow.

The Timeline tracks

The Timeline Header

The Timeline header area on the left side of the Timeline is a hierarchical list of the tools in the flow. Each tool is displayed by name and can be opened further, revealing a track for each keyframe animation, modifier and mask attached to that tool. The Timeline header acts as a label and selection tool for the segments in the Timeline.

The Brightness/Contrast tool in the Timeline Header area

Collapse/Open All

A quick way to open all the tracks and keyframe splines is to use the Collapse/Open All button. Located above the Timeline header, the button is labeled with a + or a - symbol, depending on the current state of the tracks.

The Collapse/Open All button set to Open All

Clicking this button when the label is a minus (-) sign will collapse the list of tools in the header. When the label is a plus (+) symbol, all tool categories and specific tools in the header will be opened revealing any splines created from animation.

The Playhead

The Playhead is a vertical bar that runs through the Timeline view to indicate the position of the current frame or time.

The Playhead

Dragging the playhead left or right can change the current frame of the project and ultimately what is displayed in the Viewer. As current time is changed, the frame number is displayed in the Time Ruler to assist with precise positioning.

To position the playhead, do one of the following:

1. Hold Command-Option (Mac OS X) or Ctrl-Alt (Windows).
2. Click in the Timeline view where you want the playhead to be positioned.

    Or

1. Position the cursor over the playhead.
2. Drag the playhead to a new frame in the Timeline view.

Timeline Toolbar

The Timeline toolbar sits above the Timeline view, providing access to various filters, scaling modes and the Spreadsheet.

The Timeline toolbar

Spreadsheet

The Spreadsheet view offers an alternate way of entering time and keyframes in a project. The Spreadsheet view is similar to a table with rows and columns of cells that contain the numeric values of the selected tools, including trimming and keyframes.
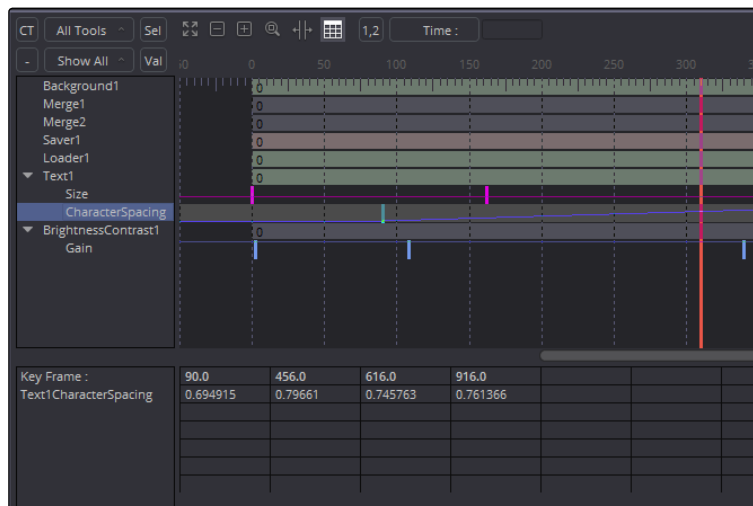
The Spreadsheet view

Scaling and Panning the Timeline

Similar to the Flow Node Editor, learning to scale and pan in the Timeline will allow you to make adjustments more efficiently.

To pan the Timeline view, do one of the following:

→ Hold the middle mouse button and drag to pan in all four directions.

→ Drag the scrollbars at the bottom or along the sides of the Timeline.

The scrollbars at the bottom and sides of the Timeline can be used to pan.

To scale the Timeline view from the toolbar, do one of the following:

→ Click the the Zoom Out and Zoom In icon in the toolbar to scale the Timeline in or out.

→ Click the Zoom to Rectangle button, then drag a rectangle in the Timeline.

→ Click the Fit button so that all currently visible tools are scaled to fit within the view.

The Fit, Zoom Out, Zoom In and Zoom to Rectangle buttons in the toolbar

To scale the Timeline view using the mouse, do the following:

1. Position the cursor over the Timeline view.

2. Hold the middle mouse button.

3. Click the left mouse button to zoom in or the right button to zoom out.

To scale the Timeline view using the keyboard, do one of the following:

Press the + and - keys on the numeric keypad to scale in and out.

Press Command-R (Mac OS X) or Ctrl-R (Windows) to enable scale to rectangle mode.

To scale the Timeline view using the Axis labels, do the following:

1. Position the cursor over the frame labels of the horizontal axis.

2. Drag right to zoom in and left to zoom out.
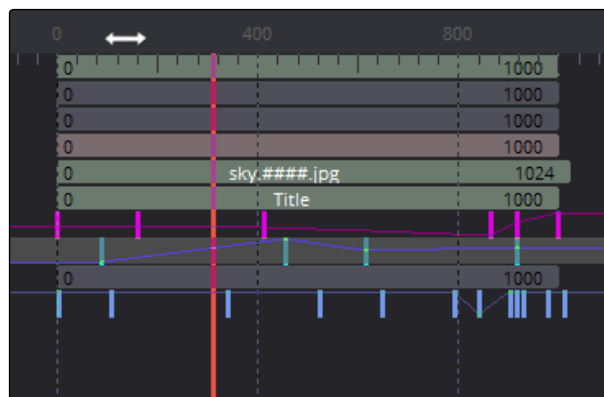
Dragging in the Axis label area will zoom in and out of the Timeline

Working with Tools in the Timeline

Most of the work in the Timeline has to do with aligning clips and adjusting keyframes. Whether you are changing a single clip or aligning multiple elements, you must know how to position the cursor in the right location to move timeline segments in the direction you want them to go.

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

→ Click the tool's name in the header.

→ Click the tool's segment in the Timeline.

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

→ Hold Command (Mac OS X) or Ctrl (Windows), then click a segment.

→ Hold Shift and click a segment to select all contiguous segments.

To remove a segment from the selecting, do the following:

Hold Command (Mac OS X) or Ctrl (Windows), then click a selected segment.

Tip: Selecting a tool's name from the Timeline header also selects the tool's tile in the flow, with its controls displayed in the Control Panel.

Sliding Segments in the Timeline

To slide the position of a segment in the Timeline, drag on the tool's segment in the Timeline. The cursor will resemble a bar with two arrows pointing in either direction.

The Move cursor

This will change where, on the Project's global time range, the clip begins and ends.

Trimming segments

In a traditional editing sense, trimming is the process of removing or adding frames to the start or end of a clip. Trimming a Loader tool works in this same way. It will modify the clip start and clip end of the loaded sequence or movie file. Trimming the segments that represent effects tools in the Timeline is a bit different. When trimming an effect tool segment, it will modify the range of that tool's effect. Outside of the trimmed region, the tool will behave as if it were passed through. This can be used to optimize performance in a flow.

Imagine a clip, for example, that is 100 frames long and is connected to a tool that animates a defocus from frames 80 -100. There is little to no point in processing the defocus tool between frames 0 - 79. Trimming the defocus segment to start at frame 80 in the Timeline will effectively prevent it from rendering and consuming memory or processor time until needed.

To trim a segment in the Timeline, do the following:
> → Drag on either end of the tool's segment in the Timeline.

The cursor will change 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 (Mac OS X) or Ctrl (Windows) while you drag beyond the first or last of the segment in the Timeline.

A Loader with freeze frames created at the first and last of the clip

Working with Keyframes in the Timeline

The Timeline also shows any keyframes set on a tool. When the keyframes are collapsed, they are drawn over the tool's segment. Clicking on the disclosure icon to the left of the tool's name in the header will expand the display so that each keyframes parameter has its own track in the Timeline.

Click the disclosure icon to reveal keyframes for the tool

Keyframes for the splines are normally represented as vertical colored bars, making them easier to manipulate.

Changing Keyframe Time

Whether the keyframes are expanded in the Timeline or not, are always available so you can adjust the timing on the segments or on the splines. Moving keyframes can be done using the mouse or the Time Edit box for more precise control.
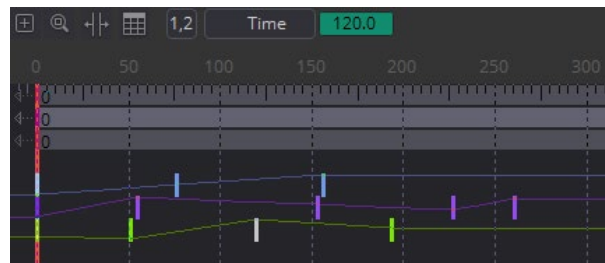
Drag to position keyframes.

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

→ Drag the keyframe to a new position.

The exact position of the dragged keyframe in time is displayed to the right of the Spline toolbar, as well as in the status bar.

You can also select multiple keyframes using a selection box, by holding Command or Ctrl and clicking on the desired keyframes or holding shift and selecting a contiguous range of keyframes. Even keyframes from several splines can be selected and moved together.

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

→ Select a keyframe, then enter a new frame number in the Time Edit box.

→ Select a keyframe(s), click the Time button to switch to Time Offset mode, then enter a frame offset.

→ Select a keyframe(s), click the Time button twice to switch to T Scale mode, then enter a frame offset.

The Time button can switch to Time Offset or Time Scale for moving keyframes

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

→ Select the name of the tool in the header area.

→ Type a new frame number in the Keyframe box to move.

The Spreadsheet's Keyframe box shows the frame number for the keyframe

Duplicating Spline Keyframes

Keyframes and multiple keyframe selections can be duplicated on the same spline or onto different splines. This can save you time if you need to repeat keyframes at another time on the same segment or even just have keyframes timed similarly on two different timeline segments.

To duplicate keyframes, do the following:

1. Select the keyframe(s) to be duplicated.

2. Hold Command (Mac OS X) or Ctrl (Windows) and drag the keyframes to a position.

While you drag, the cursor will change from the No symbol to a sub selection cursor until you release the mouse button and place the keyframes.

Duplicating keyframes

Time Stretching

Time Stretching allows 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 time stretching, the tracks in the Timeline must be open to show the splines.

To enable keyframe time stretching once the splines are open, do one of the following:

→ Click on the Time Stretching button in the toolbar.

→ Right click over a parameter's name in the Timeline header, then choose Time Stretch from the contextual menu.

→ Right click over a keyframe and choose Time Stretch from the contextual menu.

The Time Stretch button

If more than one keyframe is selected on the spline, the keyframes will be surrounded by a purple rectangle to identify which keyframes will be stretched or squashed. If no keyframes are selected, you can drag a rectangle to set the boundaries of the bars. The time stretcher can affect more than one spline at a time.
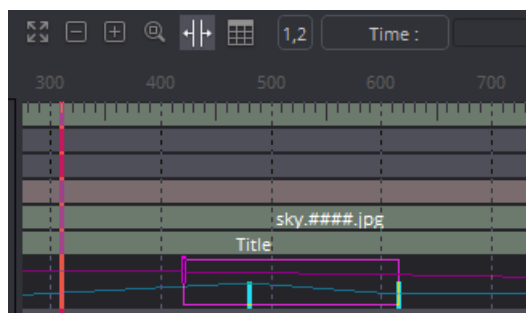
The Time Stretch selection rectangle

Once the selection is made, you can drag on the left or right edge of the purple rectangle to stretch or shrink the timing of the spline segments within the rectangle.

To disable time stretching, click on the toolbar button again or re-select Time Stretch from the contextual menu.

Showing Keyframe Values

If parameters are animated on a tool, colored tick marks in the tool's timeline segment indicate where keyframes are placed. If the tracks and splines are open on a parameter, the Val button above the header area will reveal text boxes beneath each keyframe on the spline. The boxes show the keyframe's current value and allow you to edit the value by typing in a new number.

Keyframes with value displayed in the Timeline

Timeline Filters

When a comp grows to includes hundreds of tools, locating items can quickly become difficult in the Timeline. Timeline filters can be applied to sift out tools that are not necessary to the current operation. The Global Timeline preferences include a number of pre-made filters that you can enable, or create new ones as needed.

To apply a Timeline filter, do the following:

1. Click the Filters button above the Timeline header to reveal the Filters menu
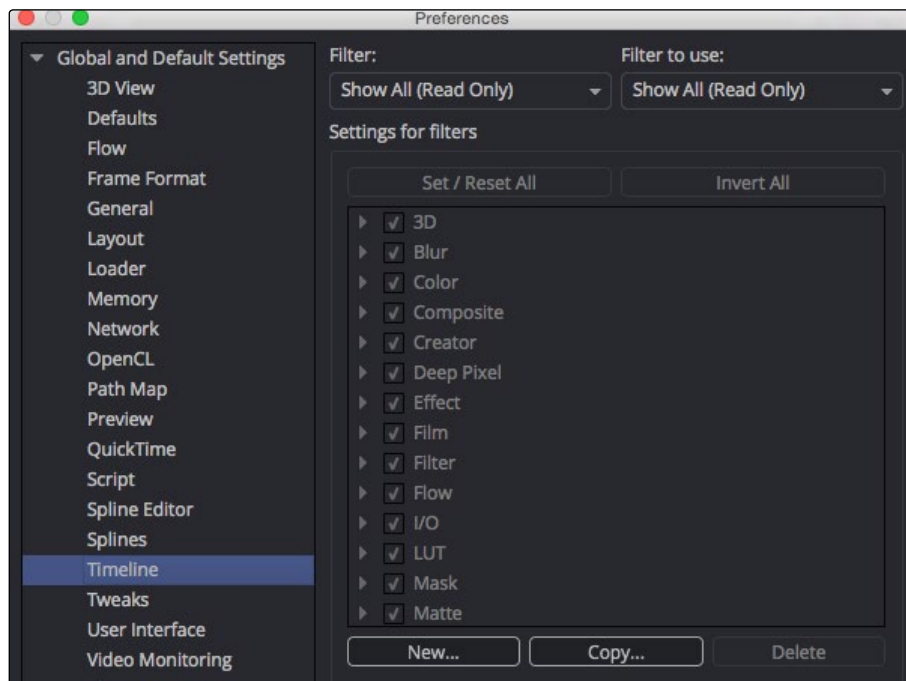
The Global Timeline preferences for enabling filters

2. Choose the desired filter from the list to apply it to the Timeline.
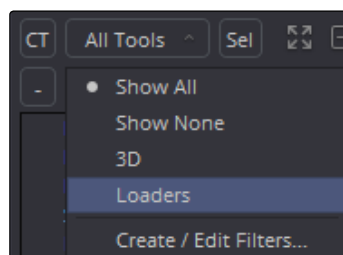
The Timeline Filters list

Creating Filters

You create new Timeline filters in the Timeline Preferences dialog by choosing Fusion > Preferences (Mac OS X) or File > Preferences (Windows) and then selecting Global > Timeline in the sidebar. You can also open the Timeline preferences by clicking the Filter button above the Timeline header and choosing Create/Edit Filters.

Once in the Timeline Preferences dialog, clicking the New button will allow you to name and create a new filter with the default set to all the tools being enabled.

Add and Remove Tools from Filters

The Timeline Preferences dialog includes a list of all the tools you can choose to include in your custom filter. Clicking the disclosure button next to a Tool category opens it to reveal the specific Tool types. If you enable the checkbox next to a tool, that tool will not be filtered out when you select the custom filter from the Filter list in the Timeline.

Enable each tool you want to keep in the Timeline when the filter is selected

The Invert All button and Set/Reset All button can apply global changes to all of the checkboxes, toggling the selected states as described.

When you are done configuring the custom filter, click the Save button to close the preferences dialog and save the filter.

Removing Filters

To disable the currently applied filter, choose Show All from the Filter pop up menu in the Timeline. You can permanently remove a filter from the list by returning to the Timeline preferences, choose the filter from list at the top of the dialog, then clicking the Delete button.
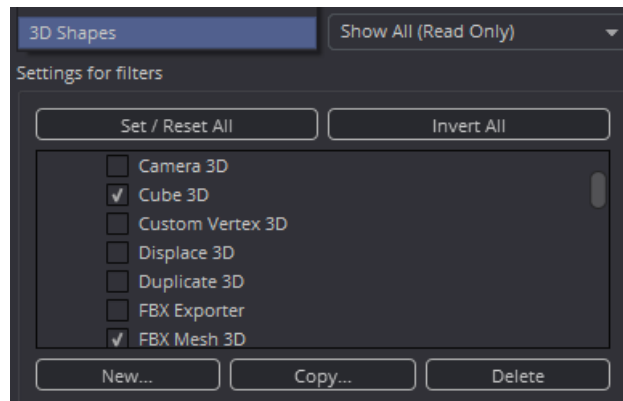
Clicking Delete permanently removes the custom filter from the Filter list

Current Time Mode

While not part of the Filter menu, the CT (Current Time) button can filter the Timeline. When you are working on a complex composite with many elements that come and go throughout the comp, clicking the CT button can filter the Timeline to only display the segments that are under the Playhe ad.

The CT button will filter the Timeline to show only tools that are at the Playhead position

Selected Filtering

Another button located in the toolbar that can help clear the Timeline of segments that maybe temporarily in the way is the Select or Sel button. Located to the right of the CT button, the Sel button will filter out all the segments from the Timeline except the segment of the selected tools in the Timeline header. Clicking the Sel button again will return to displaying all the tools.

The Sel button filters the Timeline to display only the tool you select in the Timeline header

Tip: When the Sel button is enabled, you can select tools in the Flow Node Editor or Control Editor and they will be the only tools displayed in the Timeline.

Sorting in the Timeline

There are three ways you can change the order in which the tools are displayed from top to bottom in the Timeline. You can drag them in any order you like, use the Preset Sort menu that sorts them based on certain criteria or use the Order Selection button to click them in the order you want them to appear.

Reordering Tools in the Header

Although reordering how tools in the Timeline are stacked does not impact how they are composited, it can be helpful when comparing segments and keyframes. To re-order the tools listed in the Timeline header, click on the name of the tool and drag it to a new position in the list.

Drag to re-order tools in the Timeline header

The Sort menu

The Timeline toolbar has a Sort menu to determine the order in which the tools are listed in the Timeline header. Once a Sort mode is set from the menu, the segments in the Timeline are rearranged according to the sort you select. Setting the menu back to All Tools will display them in a linear order, scanning the flow from left to right and top to bottom. This is the default setting.

The Timeline Sort Order menu

Hierarchy

Hierarchy sorting will display tools starting with the most background element at the top of the header, through to the most foreground at the bottom, following the connections of the tools in the flow.

Reverse

Reverse does the opposite of Hierarchy, working backward from the last tool in the flow toward the most background source tool.

Names

The Names menu item uses alphabetical order of the tools starting at the top, with the beginning of the alphabet.

Start

The Start setting uses orders the tools based on their starting point in the comp. Tools that start earlier in the Global project time are listed at the top of the header, while tools that start later are at the bottom.

Animated

The Animated sort order is not really a sorting method at all. It is more of a filter that restricts the Timeline to showing animated tools only. This is an excellent mode to use when adjusting the timing of animations on several tools at once.

Order Selection button

Order Selection is a quick way to pick the order in which the tools should appear by clicking them. Once you click the Order Selection button to enable it, the Timeline header Tool names will highlight. Click the Tool names in the order you want them to appear from top to bottom. Once you disable the Order Selection button, the Tool names in the header and the corresponding tracks will be rearranged based on your selection. You can reset the order using the Sort menu.

Click the Order Selection button, then click the tools how you want them to appear

Guides

Guides are markers to help identify important frames in a project. They may indicate the frame where a dragon appears on screen, the moment that someone passes through a portal in the image, or any other important event in a composite.

Guides are created in the Time Axis label. They have the appearance of a small teal triangle and a teal line that extends from them vertically through the graph view.
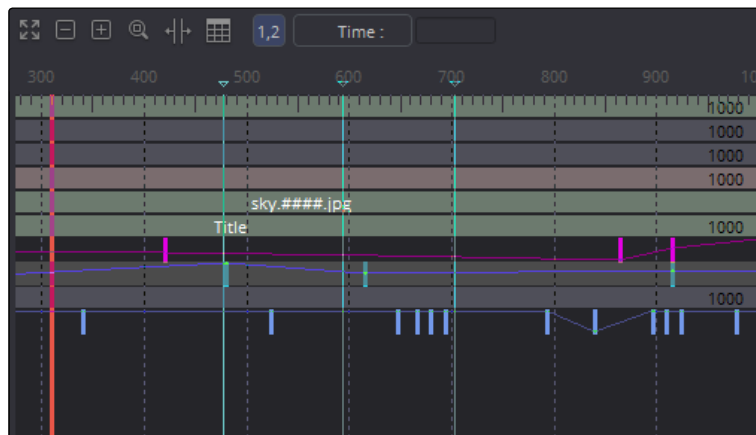
A Guide being moved in the Timeline

To create a guide, do the following:

→ Right-click in the bottom half of the Time Axis labels.
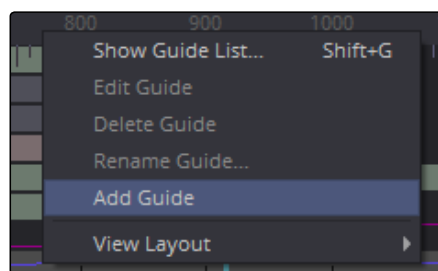
Right-click in the bottom half of the Time axis to add a guide

Working With Guides

The most important attribute of a guide is its position. For it to add value it needs to be placed on the frame you intended it to be on. Hovering the cursor over a guide displays a tooltip with its current frame position. If it is on the wrong frame, you can drag it along the Timeline to reposition it.

Renaming Guides

By default, a guide uses the frame number it is on as its only name, but you can give it a more descriptive name to go along with the frame number, making it easier to identify. To rename a guide, right click at the top of the guide and choose Rename Guide from the contextual menu. Enter a name in the dialog and click OK.

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

Show Guide List

Guides can be used to jump to specific locations in a comp using the Guide List. If you right click at the top of a guide to bring up the contextual menu, you can choose Show Guide List or press Shift-G to display the Guide List dialog. The Guide List shows all the current guides in the composition, listed according to their frame number with any custom name you've given them. If you double click a guide's name from the list, the Playhead jumps to the guide's location.

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

There are a pair of checkboxes beside the names of each guide. One is for the Spline Editor and one is for the Timeline. By default, guides are shown in both the Spline Editor and Timeline Editor, but you can deselect the appropriate checkbox to hide the guides in that view.

The Guide List is a dialog and will remain on top of the main window until you close it.

Deleting Guides

You can delete a guide by dragging it up beyond the axis labels and releasing the mouse. You can also use the Guide's contextual menu to choose Delete Guide.

Autosnap

To help with precisely positioning spline keyframes and the start and end of segments as you drag in the Timeline, you can have them snap to a field, frame or to guides. Autosnap is an option that is accessed through the Timeline's contextual menu. There are two submenu options for autosnapping. One option controls the snapping behavior when you drag keyframes, control points or the starting and ending edges of segments. The other option controls the snapping behavior of guides.

Autosnap Points

When you drag Keyframes or the edges of segments, often you want them to fall on a specific frame. Autosnap restricts the placement of keyframes and segment edges to frame boundaries by default, but you have other options found in the contextual menu. To configure autosnapping on keyframes and segment edges, right click over the Timeline and choose Options > Autosnap Points from the contextual menu. This will display the Autosnap Points submenu with options for the snapping behavior.
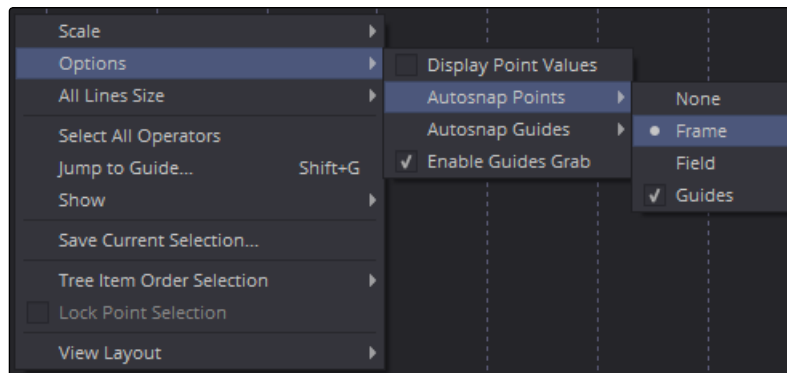
The Autosnap Points submenu

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.

Guides

When enabled, the keyframes and segment edges snap to guides in the Timeline.

Autosnap Guides

When you click to create a new guide, the default behavior is that it will snap to the closest frame. If you reposition the guide, it also snaps to the nearest frame as you drag.

This behavior can be changed in the Timeline's contextual menu by choosing from the Options > Autosnap Guides submenu.

The Autosnap Guides submenu

None

Guides can be placed anywhere with subframe accuracy.

Frame

Frame forces all guides to snap to the nearest frame.

Field

Field forces all guides to snap to the nearest field.

The Spreadsheet Editor

The Spreadsheet editor is a separate panel that can be displayed beneath the Timeline Editor. It is used to compactly show the numeric values of the keyframes for selected tools, using a table with rows and columns.
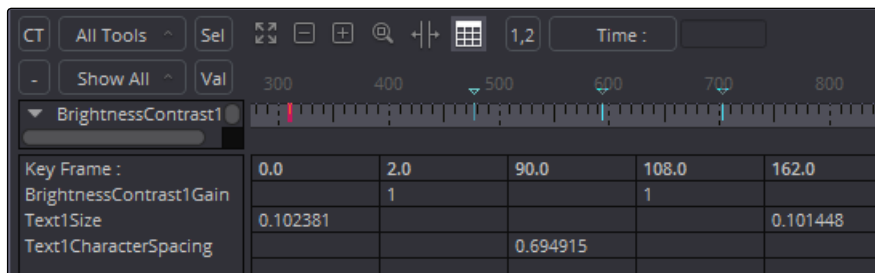
The Spreadsheet editor

To reveal the Spreadsheet editor, click on the Spreadsheet button in the toolbar. The spreadsheet will split the Work Area panel and appear below the Timeline interface.
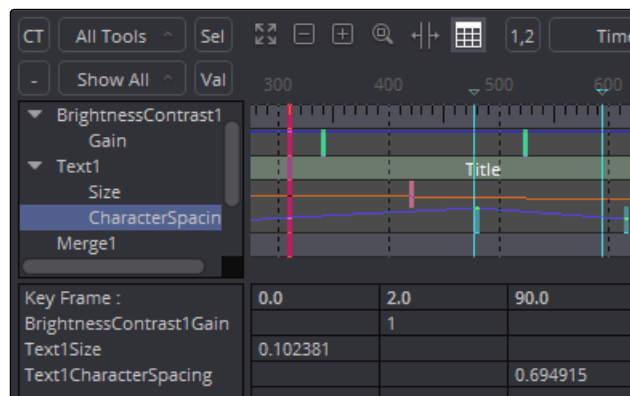
The Spreadsheet button in the toolbar

Selecting a Tool to Edit

To display a tool's timing in the spreadsheet, select the tool in the Timeline header. The Start and End points of the selected tool will appear in the keyframe's line of the spreadsheet.

To edit an animation applied to a parameter in the Spreadsheet editor, select the Keyframed parameter in the Timeline header. The keyframe row includes a box for each frame number that contains a keyframe. The value of the keyframe is displayed in the cell below the frame number. Clicking on a cell allows you to change the frame number the keyframe is on or the parameter's value for that keyframe.

0.0	2.0	90.0
	1	
0.102381		
		255.0256

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 Tools to Edit

Multiple splines and tools can be edited together in the spreadsheet. By default, selecting a new parameter in the Timeline header will replace the parameter and keyframes currently listed in the Spreadsheet editor. Holding Command or Ctrl, you can click on additional parameters on different tools to add to the spreadsheet.

Tip: You can use the Tab or Shift Tab keys to move between cells in the Spreadsheet editor.

Changing the Timeline Appearance

There are a few ways you can change the appearance of the Timeline to better fit your needs.

Timeline Tool Color Coding

The same colors used by tiles in the flow are also used in the Timeline Editor.

See Chapter 3, Flow Node Editor, for details on which Tool colors represent which categories.

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.

To increase or decrease the height of a Timeline segment, right click on the segment in the Timeline view and choose a size from the Line Size submenu.
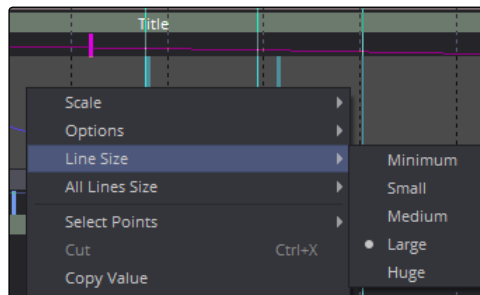
Timeline contextual menu's Line Size submenu for changing segment height

All Lines Size

The All Lines Size submenu includes the same choices as the Line Size submenu, but it will affect all segments currently visible in the Timeline.

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 choose Options > Display Point Values to switch the appearance of keyframes.

Keyframes with display point values enabled

Timeline Contextual Menu

The Timeline contextual menu includes most of the options found in the Tool contextual menu.

To open the Timeline context menu, right click on the Name in the Timeline header or right click on the tool's segment in the Timeline.

For complete documentation of the options in the Tool context menu, see Chapter 3, Flow Node Editor.

Spline Editor

8

Spline Editor

Spline Editor Overview

The Spline Editor is the main area where animation is manipulated and refined. It is primarily used 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 a keyframe is set 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.

On the left side of the Spline Editor, the Spline header shows a list of tools and their parameters that are keyframed. The majority of the panel is taken up by a graph area that displays the splines for those keyframed parameters. The horizontal axis of the graph area represents time while the vertical axis represents the parameter's value. By manipulating the keyframes and splines using different interpolation methods, you can dramatically change the style of animation.
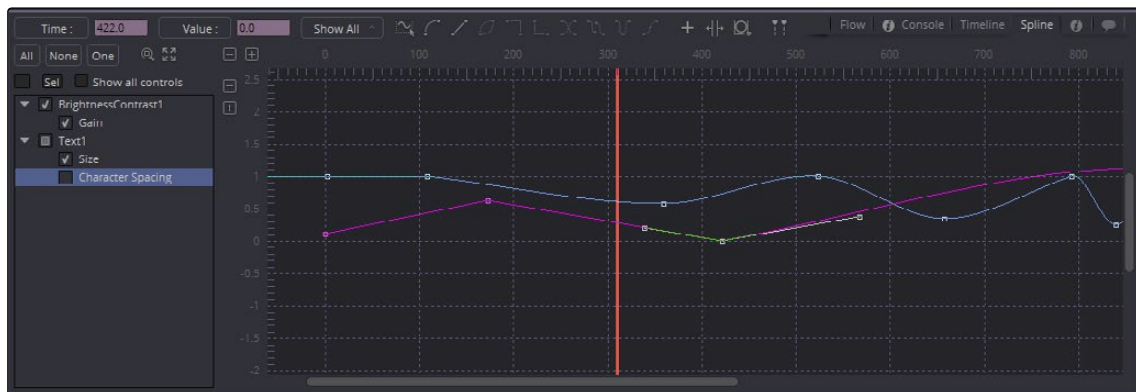
The Spline Editor

Spline Editor Interface

The Spline Editor Interface has three main working areas: the Graph, the Header and the Toolbar. All three of these areas as described below in detail.

Spline Graph

The Spline graph is the largest area of the interface. It is here that the actual animation spline is displayed for editing. 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 and the frame that is displayed in the Viewers.
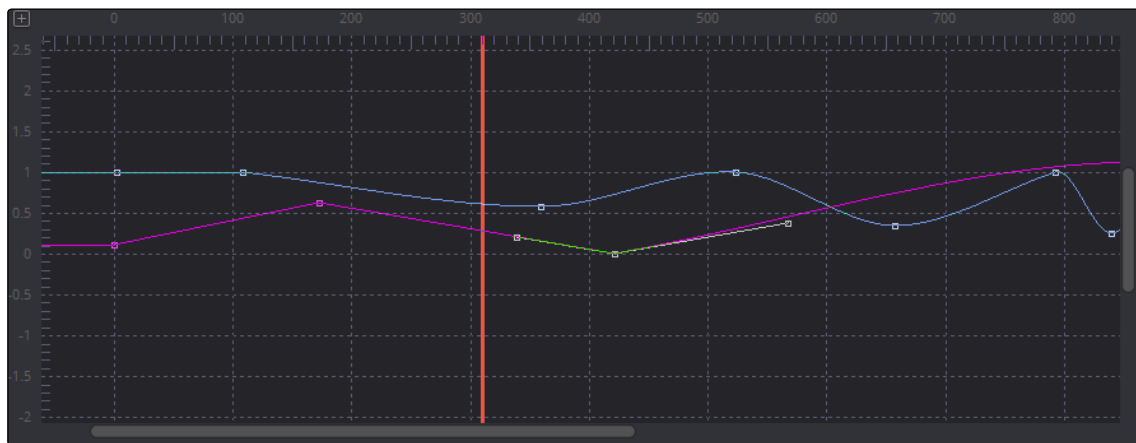
The Spline graph

Spline Header

The Spline header displays a list of each tool with animated parameters in the project. Each tool has a disclosure arrow that can be used to reveal its animated parameters. Clicking the checkbox next to a parameter's name displays the parameter's spline in the Spline graph.
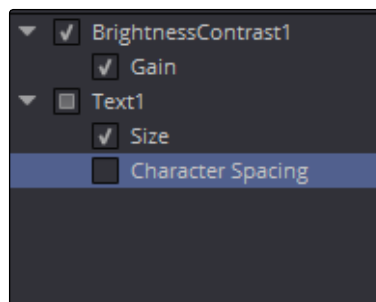
The Spline header

Toolbar

The Toolbar across the top of the Spline Editor represents the most common operations applied to a spline, or splines, in the Graph. Although you can access many of them from the Graph's contextual menu, the buttons provide a shortcut.

The operation of each button is described later in this chapter.

The Spline toolbar

Viewing, Scaling and Panning the Spline Editor

It is sometimes necessary to adjust the view of the Graph within the Spline Editor, just as it is in the Timeline Editor and Flow Node Editor. You often need to scale and pan around to focus on an area of animation that needs work. Below are a number of different ways this can be accomplished.

To view the Spline Editor, do one of the following:

→ Click the Spline Editor tab located at the top of the Work Panel.

→ Press F8 to switch to the Spline Editor.

→ Right click on a tool within the Flow Node or Timeline Editor or a parameter in the Control Panel and select Edit Splines from the contextual menu.

The Spline Editor tab

To scale the Spline Graph using the Toolbar buttons, do one of the following:

→ Click the Scale buttons located in the top left corner of the Graph.

→ Click the Zoom to Fit button to fit all the currently active splines within the Graph view.

→ Click the Zoom rectangle button, then draw a rectangle in the Graph to define the area you want to fill the view.

Toolbar Scale buttons

To scale the Spline Graph using the axis labels, do the following:

→ Place the cursor over the horizontal or vertical axis label and drag to resize the Graph view on the corresponding axis.
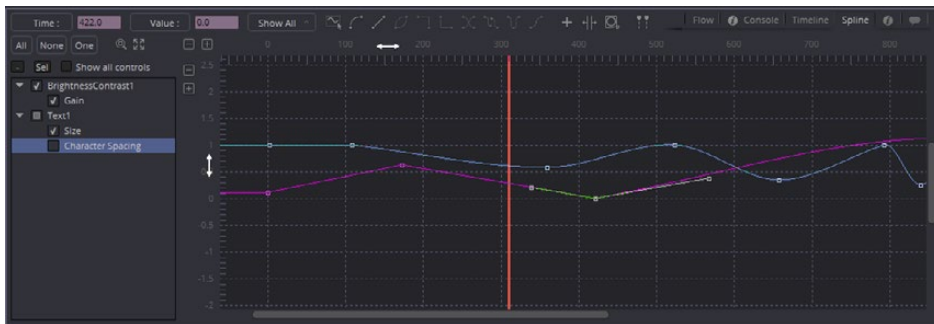
Dragging in the horizontal axis label will scale the graph horizontally

To scale the Spline Graph using the mouse, do the following:

→ Place the cursor over the Graph, hold the middle mouse button, then click once on the left mouse button to zoom in and once on the right button to zoom out.

To scale the Spline Graph using the keyboard, do one of the following:

→ Press the + and - keys on the numeric keypad.

→ Press Command-R (Mac OS X) or Ctrl-R (Windows) to enable scale to rectangle mode, then drag of a rectangle around the area of the Graph you want to fill the view.

To pan the Spline Editor, do one of the following:

→ Hold the middle mouse button and drag up, down, left or right.

→ Drag the scrollbars on the right or bottom of the Graph.

Dragging in the horizontal scroll bar will pan the graph horizontally

In general, scaling and panning the Spline Editor works the same as it does in all navigable parts of the Fusion interface. See Chapter 1, Getting Started, for more information.

Creating Animation Splines

When you add keyframes to a parameter, an animation spline is created in the Spline Editor. The default animation spline uses the standard Bezier spline found in most drawing and graphics software, but you have other spline interpolation options as well.

To create and view a default spline animation, do the following:

1. In the Control Panel, right click over the parameter to be animated.

2. Choose Animate from the contextual menu.

3. Right click over the parameter again and choose Edit from the contextual menu.

Selecting Animate from the context menu will animate the parameter using the default Bezier spline type. Instead of choosing Animate in the Control Panel contextual menu, other spline types can be used from the Modify With contextual submenu.
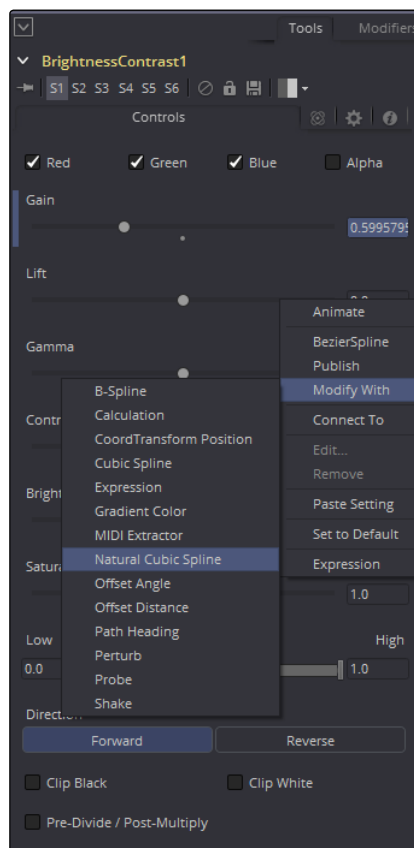
The Modify With contextual submenu

You can change the default spline type by changing the Default Animate option in the Globals > Defaults panel of the Preferences.

Modify with > Bezier Spline

Choosing Bezier Spline from the Modify With contextual submenu is the same as choosing Animate, unless you have changed the default spline type in your preferences. Bezier splines use three points for each keyframe on the spline to determine the smoothness of the curve. The first point is the actual keyframe or control point, representing the value at a given time. The other two points are direction handles that determine how smoothly the curve for the segments leading in and out of the keyframe are drawn. Bezier splines are powerful because they allow you to create combinations of curves and straight lines.

Modify with > B-Spline

B-splines are simpler and easier to control than Bezier splines because they use a single control point that pulls the spline in its direction, producing a smooth curve when you set the interpolation to smooth. The spline is shaped by the control point but it doesn't necessarily go through it. Unlike Bezier splines, B-splines do not have direction handles, so the control point determines the keyframe value as well as the smoothness of the curve.

Modify with > Natural Cubic Spline

Natural Cubic Splines are very similar to Bezier splines, especially when the interpolation between points is linear. When the interpolation is smooth, the handles for Cubic splines are locked and invisible, and always set to make the smoothest possible curve.

Other choices in this menu allow you to use Expressions and Modifiers for animation. You can learn more about Modifiers in the Tool Reference Manual. You can learn more about Expressions in Chapter 14, Introduction to Controls and Customization.

Displaying Splines for Editing

The Spline header on the left of the Spline Editor is a hierarchical list. The names of any tools in the flow that have an animated control are displayed. Clicking the disclosure arrow on a tool reveals all of the names of the animated parameters on that tool. Clicking directly on the parameter name in the Spline header will activate that spline for display and editing.
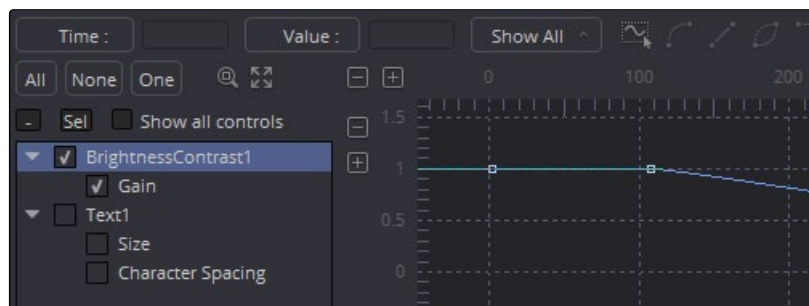
The Spline header with Brightness Contrast tool and its animated Gain parameter

Tool Status Checkbox

To the left of each parameter's name in the header is a checkbox that indicates the spline's status. When you select a parameter name, the checkbox becomes Active. Selecting a different parameter disables the first parameter when the other becomes Active. There are three selection modes for each checkbox: Active, Viewed and Disabled.

Active

When the checkbox is enabled with a check mark, that spline is displayed in the Graph and is Active. This spline can then be edited.

Viewed

When the checkbox is enabled but filled with a gray box, the spline is visible in the Graph but is Read only. It cannot be edited.

Disabled

When the checkbox is empty, the spline is not visible in the Graph and cannot be edited.

Three parameters for Brightness Contrast in three different modes: Active, Viewed and Disabled

You can click on a checkbox to change its mode even if the parameter name is not selected. This allows you to make multiple parameters Active, Viewed or Disabled. Changing the state of the checkbox at the Tool level will set the mode for all splines attached to that tool.

Selection States

There are three Selection States that determine how the list of parameters in the header behave when a checkbox or parameter name is selected to activate a spline. These buttons are located in the Spline toolbar, just above the header. They are labeled All, None and One. Additional options control what tools are displayed in the header. These are controlled by the Sel button and the Show All control checkbox.

All

Clicking on this button activates all splines for editing.

None

All spline checkboxes will be set to Disabled.

One

When the One button is selected, only one spline in the header is active and visible at a time. Clicking another spline's checkbox will set it to Active and all other splines will be disabled.

The Selection State buttons used to activate and deactivate splines in the graph

Show Selected Only

Another button located above the Spline header that can help clear the Graph of splines that maybe temporarily in the way is the Select or Sel button. The Sel button will filter out all the splines from the Spline Editor except the spline of the selected tools in the header. Clicking the Sel button again will return to displaying all the tools.

The Selection State buttons used to activate and deactivate splines in the graph

The Sel button filters the Spline Editor to display only the tool you select in the header

Tip: When the SEL button is enabled, you can select tools in the Flow Node Editor or Control Panel and they will be the only tools displayed in the Spline Editor.

Show All Control

By default, only animated parameters are displayed in the header. Enabling this checkbox will display all tools and parameters that are in the flow. Activating one of the controls will automatically add an animation spline to it.

With a large number of tools displayed, which themselves might have a large number of parameters, you may end up with a cluttered header that's slow to work with. You can more efficiently use the Show All control option when the Sel option is enabled to limit the number of tools and parameters displayed.
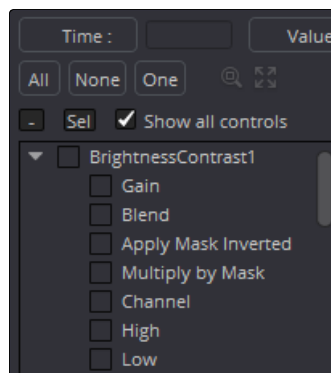
Options to show only selected parameters or all tools and parameters that are in the flow

Removing and Deleting Animation Splines

To remove an animation spline from a parameter, right click on the parameter in the Control Panel and choose Remove [tool parameter's name] from the contextual menu. Removing a spline from a control will only delete the spline if no other tools in the composition are connected to the same spline at that time.
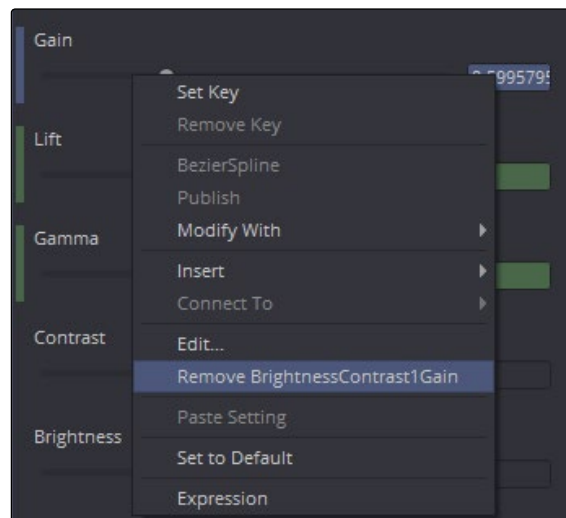
The Remove menu for a Brightness Contrast Gain animation

Working With Splines

Renaming Splines

The default name for a spline is the name of the parameter it animates. It may be useful to assign more descriptive names to a spline, especially if a composition is shared between several artists.

To rename a spline, do the following:

1. Right click on the name of the spline in the header.
2. Choose Rename Spline from the contextual menu.
3. In the dialog that opens, enter a name and click OK.

The Rename Spline contextual menu

Changing Spline Colors

Each spline in the Graph is assigned a different color so it is easier to identify when many splines are visible at once. When the spline is Active, the checkbox in the header will also be drawn in the same color as the spline..

To change the color of the spline, do the following:

1. Right click on the name of the spline in the header and choose Change Color from the contextual menu.
2. Chose the new color from the dialog box that appears and click OK.

Scaling Splines in the Graph

There are six Scale options in the contextual menus that allow you to scale the splines to make them easier to see within the Graph.
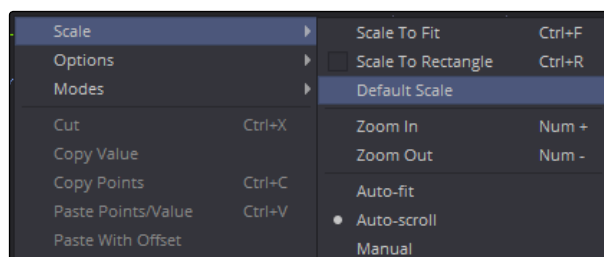
The Scale contextual submenu

To scale Splines in the Graph using the Scale contextual, do one of the following:

→ Choose Scale to Fit to fit splines into the view.

→ Choose Scale to Rectangle to zoom the view in on the contents of a drawn rectangle.

→ Choose Default Scale to reset the scaling to the default values.

→ Choose Zoom In/Zoom Out to incrementally scale or shrink the content of the Graph.

→ Choose Auto Fit to dynamically perform a Scale to Fit when new splines get selected or added.

→ Choose Auto Scroll to fit the view in X direction.

Spline Options

There are three options in the contextual menus that ensure the splines stay within the visible Graph area as you work.
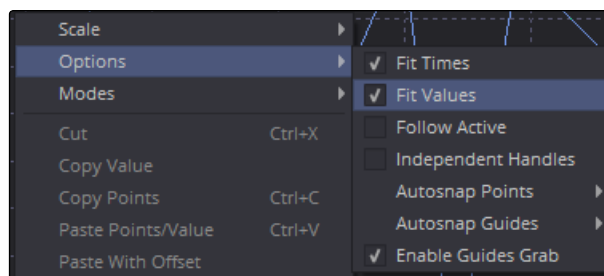
The Options contextual submenu

Fit Times

When enabled, the Graph automatically scales to fit the X-axis when a spline is selected. 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.

Fit Values

When enabled, the Graph automatically scales to fit the Y-axis when a spline is selected. 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.

Follow Active

This option causes all splines connected to the currently active tool to be made visible and ready for editing whenever a tool is selected.

Show All Inputs

The Show All Inputs option shows all the parameters on all tools in the Spline header, regardless of whether they are animated or not. This mode is handy when animating a bunch of parameters quickly, but it can be cumbersome with large compositions.

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, do the following:

1. Right click over a parameter in the header or in the Graph.

2. From the contextual menu, choose Save Current Selection.

3. Name the Selection group in the dialog, then click OK.
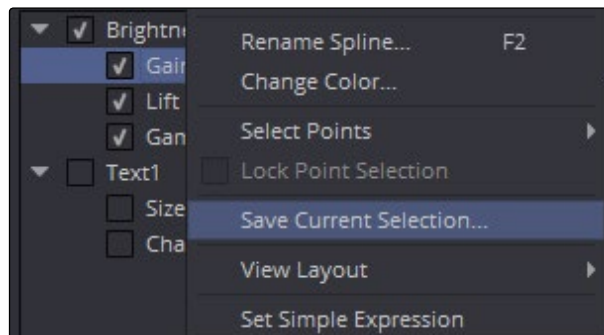
The Save Current Selection menu

To reapply the selection group, right click over any parameter in the header and choose the Selection group by name from the Set Selection submenu. Other contextual menu options will allow Selection groups to be renamed or deleted.

Working With Keyframes

Once a spline is displayed in the Graph and enabled for editing, the spline's keyframes can be manipulated in a variety of ways. No matter how you want to manipulate a keyframe, you have to be able to select it first.

To select keyframes on a visible spline, do one of the following:
→ Click directly on the keyframe.
→ Drag a rectangle around multiple keyframes.

Dragging a rectangle allows you to select all the keyframes within an area

To add a keyframe to the current selection, do the following:
→ Hold Command (Mac OS X) or Ctrl (Windows), then click on the point to add. This will add the keyframe to the selection.

To select all keyframes from the active splines, do one of the following:
→ Click in an empty area of the Graph, then press Command-A (Mac OS X) or Ctrl-A (Windows).
→ Right click over the Graph and choose Select Point > All from the Spline Editor's contextual menu.

Moving Keyframes

Keyframes can be moved with the cursor, keyboard or using the Time and Value Edit controls in the toolbar.

Tip: Although extremely flexible, keyframes on the same spline cannot be made to overlap. If a key exists at frame 2 and frame 5, the key at frame 5 cannot be moved to frame 1.

To move a keyframe with the cursor, do one of the following:
→ Drag the selected keyframe to its new position in the spline. If more than one point is selected, all selected points will be moved with the one selected.
→ Hold Option (Mac OS X) or Alt (Windows) and drag the point to constrain its motion to the first axis you drag.

To move a keyframe with the keyboard, do one of the following:

→ Select a keyframe, then use the arrow keys to move the keyframe left, right, up or down.

→ Select a keyframe, then hold Shift and press the arrows keys to move the keyframe in larger increments.

Tip: When using the arrow keys to move a keyframe, the scale of the Graph determines the amount of movement applied to each keypress. The closer the zoom in, the finer the adjustment.

Time and Value Editors

The Time and Value buttons are used to move the position and the parameter value of a keyframe by entering a number into the edit box for each button. Each button can switch between three modes that are helpful in modifying the time and value of a keyframe in three precise but distinct ways. The default mode for each button is used to enter the explicit frame number or parameter value at which you want the keyframe to be set. The other modes can be used to offset and scale the keyframe's position or value.
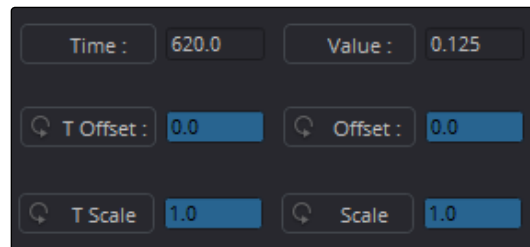
The Time and Value Edit boxes

Time Button

The Time button is used to modify the current time of the selected keyframe using a Time mode, Time Offset and Time Scale. The three modes are cycled through each time you click on the Time button.

Time

When the Time button is set to Time, the edit box shows the current position of the selected keyframe in frames. Typing a new frame number in the edit box will move the selected keyframe to the specified frame. If no keyframes are selected, or if multiple keyframes are selected, the edit box is empty and nothing can be entered into it.

Time Offset

Clicking the Time button switches it to Time Offset mode. The edit box is used to move the selected keyframe based on its relative position. An offset of either positive or negative values can be entered. For example, entering an offset of 2 will move a selected keyframe from frame 10 to 12.

Time Scale

Clicking the button when it is set to Time Offset switches it to Scale mode. The selected keyframe's position will be scaled based on the position of the playhead. For example, if a keyframe is set at frame 10 and the playhead is at frame 5, entering a scale of 2 will move the keyframe 10 frames forward from the playhead's position, to frame 15.

Value Button

The Value button is used to modify the selected keyframe's parameter value using a Value mode, Value Offset and Value Scale. The three modes are cycled through each time you click on the Value button.

Value

When the Value button is set to Value, the edit box shows the value of the currently selected keyframe. Entering a new number into the edit box will change the value of the keyframe. If more than one keyframe is selected, entering a new value will set all the selected keyframes at that value.

Value Offset

Clicking the Value button switches it to Offset mode. Entering a new value in the edit box will cause the selected keyframe's values to be moved by the amount specified. For example, entering a value of -2 will change a value from 10 to 8.

Value Scale

Clicking the button when it is set to Value Offset switches it to Scale mode. Entering a new value will cause the selected keyframe's values to be scaled or multiplied by the specified amount. For example, entering a value of 0.5 will change a keyframe's value from 10 to 5.

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.
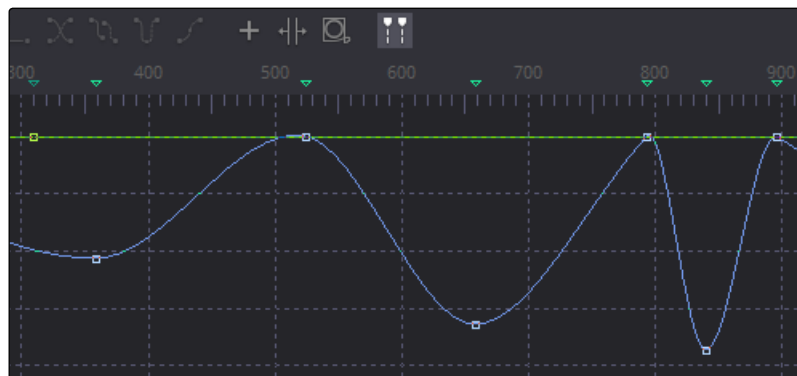
The Key Markers button in the toolbar with show keyframes in the horizontal axis

Adding New Keyframes

Once one keyframe is added to a parameter, additional keyframes are automatically added whenever you make a change to the parameter's value. For example, if the strength of an animated glow is changed at frame 15, a keyframe with the new value will be set at frame 15.

The easiest ways to explicitly add keyframes to a spline is either by clicking directly on the spline or using Shift-K to add it at the current playhead position. For special situations, the Spline Editor contextual menu includes a number of other ways to set keyframes.

Set Key

Selecting Set Key from the contextual menu adds a new keyframe along the active spline wherever the position of the cursor is when the contextual menu is displayed.
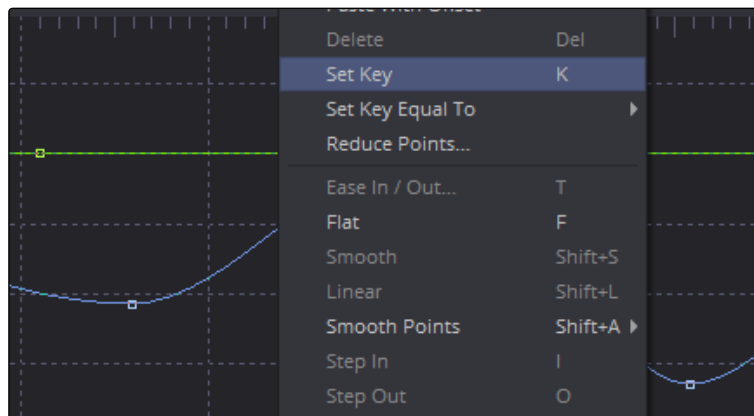
The Set Key menu

Set Key Equal to

This option allows a new keyframe to be added with a value equal to the next or the previous keyframe. Easily add linear sections in your spline with no value change.

Deleting Keyframes

To delete a keyframe or keyframes, simply select the keyframes and press Command-Delete (Mac OS X) or Delete (Windows). This will only remove the keyframes. It will not remove the spline, even if all of the keyframes are deleted.

Copying and Pasting Keyframes

The Keyframe values from one spline can be copied to a different area on the same spline or from one spline to another.

To copy Keyframe values, do one of the following:

→ Select the desired keyframes, then press Command-C (Mac OS X) or Ctrl-C (Windows).

→ Select the desired keyframes, then choose Copy Value from the Graph's contextual menu.

→ Hold Command (Mac OS X) or Ctrl (Windows) and drag the points where you want them.

To paste Keyframes points, do one of the following:

→ Place the cursor where you want the points pasted, then choose Paste from the Graph's contextual menu.

→ Place the cursor where you want the points pasted, then press Command-V (Mac OS X) or Ctrl-V (Windows).

Reshaping Splines

The shape of a spline impacts an animation by changing the interpolation between keyframe values. A Bezier spline is incredibly flexible for making the adjustments to a spline's shape because it uses a pair of direction handles to shape the spline as it passes through the keyframe or control point. These handles can only be seen when the keyframe is selected. Dragging on direction handles will adjust the slope of the segments passing through the spline.

A Bezier Control point and direction handles

By default, the two direction 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 direction handles separately for a more pronounced curve or effect.

To move a direction handle independently of the other, do the following:

1. Select the control point to be modified.

2. Hold Command (Mac OS X) or Ctrl (Windows), then drag on the direction handle.

Holding Command or Ctrl allows you to move a handle independently.

To make the Spline Editor treat all Bezier handles as independent, right click in the Graph and choose Independent Handles from the Options contextual submenu.

Although manipulating the handles on a Bezier spline gives you a lot of flexibility, you can take a few short cuts to place the direction handles in commonly used handle positions. These options can be found in the Graph's contextual menu as well as in the Spline Editor's toolbar.

Linear

A Linear segment will effectively take the shortest route, which is a straight line between two keyframes. The rate of change stays the same throughout the animation so it produces a fairly rigid, mechanical style. To make the selected keyframe(s) linear, press the L key or click the Linear button in the toolbar.
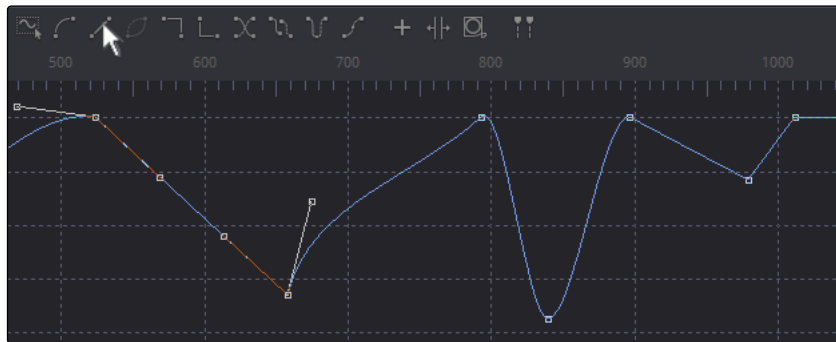
The Linear button in the toolbar

Smooth

A smoothed segment will provide a gentle keyframe transition in and out of the keyframe, by slightly extending the direction handles on the curve. It slows down the animation as you pass through the keyframe. To smooth the selected keyframe(s), press the S key or click the toolbar's Smooth button

The Smooth button in the toolbar

Step In/Out

On occasion, it is not desirable to have any interpolation between two keyframes. Instead, the value of one keyframe may need to be held until another keyframe is set. For these cases, use Step In or Step Out mode.

Setting Step In on a keyframe will hold the value of the keyframe steady, then switch values when it reaches the next keyframe. Step Out immediately switches to the value of the upcoming keyframe.

To set Step In or Step Out, do one of the following:

→ Select at least two keyframes, then click the Step In or Step Out buttons on the toolbar.

→ Select at least two keyframes, then right click and choosing the appropriate option from the contextual menu.

→ Select at least two keyframes, then press I or O on the keyboard.

The Step In/Out buttons in the toolbar

Looping Splines

It is often useful to repeat a given section of an animation, either infinitely or for a specified number of times, such as is required to create a strobing light, or a spinning wheel. There are a number of ways to repeat a selected segment in the Spline Editor, accessible in the Graph's contextual menu and in the toolbar.

The various Loop buttons in the toolbar

Set Loop

If you want to repeat or loop a selected spline segment, you first select the two or more keyframes that define the section you want looped. You can then either select Set Loop from the Graph's contextual menu or click on the Set Loop button in the toolbar. The selected section of the spline will appear as a dotted line and it will repeated 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 that were originally used to create the loop. Simply select one of the originating key points, make any necessary modifications and the looped segment will update. To remove the loop, select the keyframes you used to create the loop, then click the Loop button in the toolbar.

Relative Loop

The Relative Loop mode repeats the segment as with Loop, where 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

Ping-Pong

The Ping-Pong Loop mode repeats the selected segment, where each successive loop is reversed so that the segment plays through, reverses itself and repeats. Ping-Pong is often used for pendulum like animations where an object has to swing back and forth in a repetitive seamless motion.

A Ping-Pong section in the Graph

Set Pre-loop

The Set Pre-loop option contains the same options for looping as the Loop option, except that the selected segment is repeated backward in time rather than forward. The Pre-loop option is found in the Graph's contextual menu.

Duplicate Repeating Splines

Repeated splines are like looped splines, except that the selected segment is only repeated 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 Repeat modes are available through the Repeat submenu in the Graph contextual menu. Selecting any of the modes in the submenu will open a dialog where the number of times the segment is to be copied can be entered.

The Repeat contextual submenu

Reversing Splines

Reverse can be used to invert a segment of an animation spline. To apply reverse, choose a group of points in a spline and select Reverse from the contextual menu or use the V hotkey. The keyframes are reversed in the Graph and points immediately surrounding the reversed selection adjust of needed.

The Reverse button in the toolbar

Time Stretching

Time Stretching allows for a selected group or 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, then choose Modes > Time Stretching from the Graph's contextual menu or click on 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 will surround the outer keyframes with two vertical purple bars. Click drag on the purple 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.
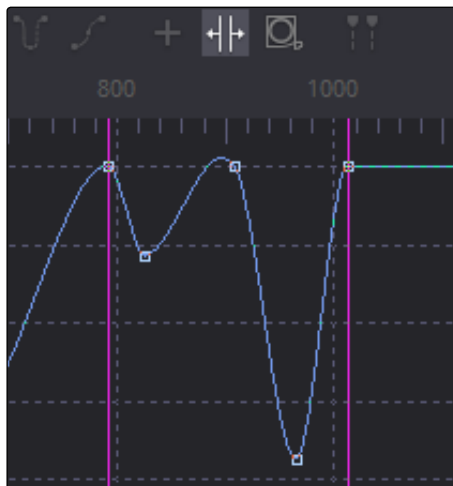
The Time Stretch bars in the Graph

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, then choose Modes > Shape Box from the contextual menu.

→ Select a group of keyframes, then click the Shape Box button in the toolbar.

→ Select a group of keyframes, then press Shift-B to enable and disable the Shape Box mode.

The Shape Box button in the toolbar

A purple rectangle will be drawn around any points selected when the mode is enabled. To scale, skew or stretch the spline, drag on any of the control points located around the box. Dragging on the box edges will move the keyframes.
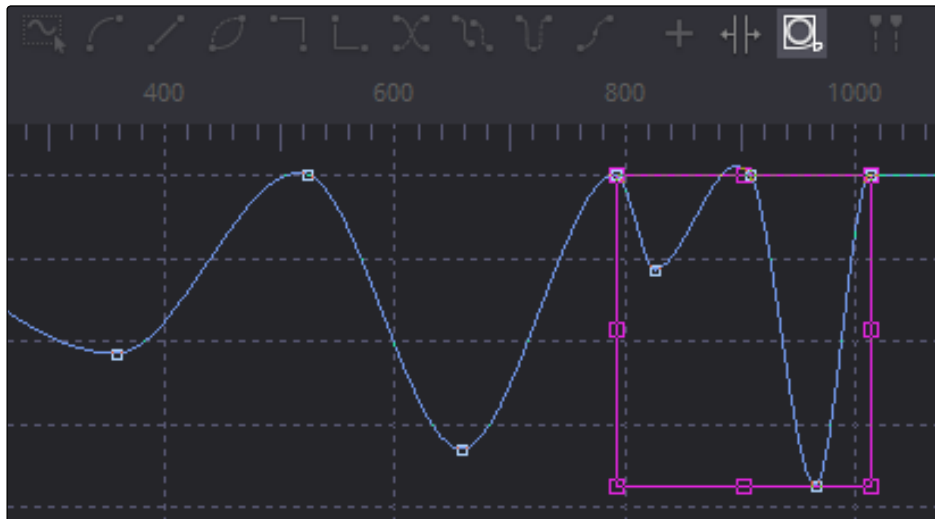
The Shape Box in the Graph

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 Bezier direction handles attached to selected keyframes, you can use the Spline Ease dialog. The dialog can be accessed by selecting a keyframe in the Graph, then choosing Edit > Ease In/Out from the Graph's contextual menu or by pressing T on the keyboard.

The appearance of the toolbar changes to present you the Ease In/Out controls. You can use the sliders to adjust the length of the direction handles or enter a value in the edit boxes.

The Ease In/Out controls in the toolbar

Clicking the Lock In/Out button will collapse the two sliders into one, so any adjustments apply to both direction handles.

Reducing Points

When there are too many control points, too close together on a spline, you can Reduce Points to decrease the number of them, making it easier to modify the points remaining. The overall shape of the spline is maintained as closely a possible, while eliminating redundant points from the path.

To reduce the number of points on a spline, do the following:

1. Select the range of keyframes you want to thin out.
2. Select Reduce Points from the contextual menu
3. In the Reduce Points dialog, 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

A composition can easily contain dozens, if not hundreds, of animation curves, making it difficult to locate a specific spline. You can choose to filter the splines shown in the Spline Editor by displaying selected tools only or creating a filter to show only certain tools.

A filter allows the tools shown in the Spline Editor and Timeline to be limited to specific types and categories of tools. For example, you can make a filter that shows only Creator tools and Loaders, or one that only shows Color Corrector and Brightness/Contrast tools.

To apply a filter in the Spline Editor, do the following:

1. Click the Filters button above the toolbar to reveal the Filters menu.

2. Choose the desired filter from the list to apply it to the Spline Editor.
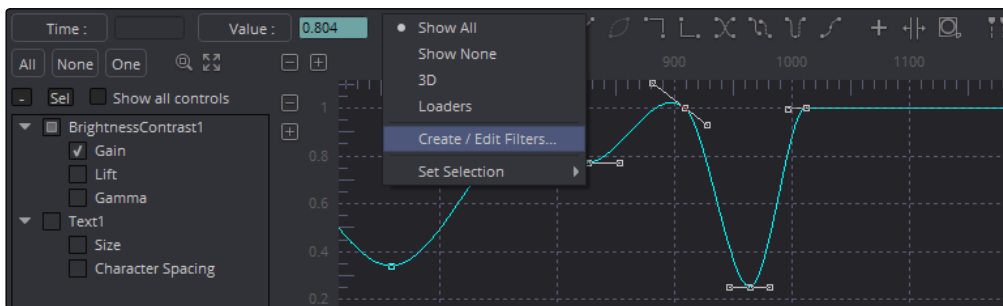
The Filter button in the toolbar, labeled Show All by default

Filters are applied simultaneously to the Spline Editor and Timeline.

Creating Filters

You create new Spline Editor filters in the Global Timeline Preferences dialog by choosing Fusion > Preferences (Mac OS X) or File > Preferences (Windows) and then selecting Global > Timeline in the sidebar. You can also open the Timeline preferences by clicking the Filter button above the Spline header and choosing Create/Edit Filters.

Once in the Timeline Preferences dialog, clicking the New button will allow you to name and create a new filter with the default set to all the tools being enabled.

Add and Remove Tools from Filters

The Timeline Preferences dialog includes a list of all the tools you can choose to include in your custom filter. Clicking the disclosure button next to a Tool category opens it to reveal the specific Tool types. If you enable the checkbox next to a tool, that tool will not be filtered out when you select the custom filter from the Filter list in the Timeline.
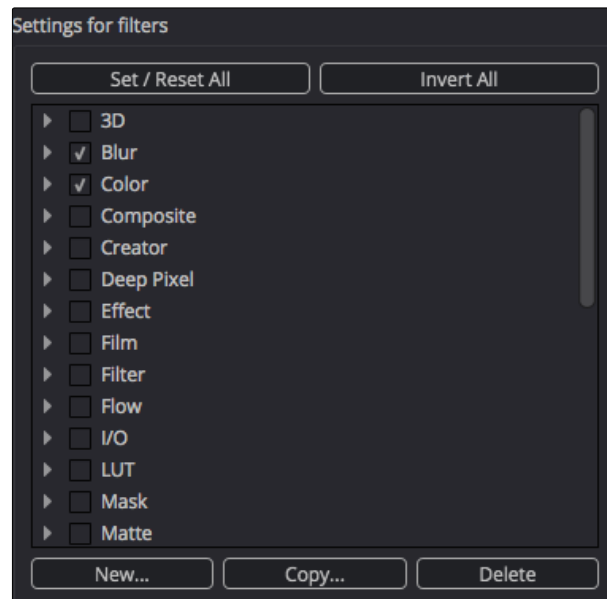
Enable each tool you want to keep in the Spline Editor when the filter is selected

The Invert All button and Set/Reset All button can apply global changes to all of the checkboxes, toggling the selected states as described.

When you are done configuring the custom filter, click the Save button to close the preferences dialog and save the filter.

Removing a Filter

To remove a filter and show all the tools in the Spline Editor again, click on the Filter button and choose Show All from the pop up menu.

Guides

Guides are markers to help identify important frames in a project. They are created by clicking in the Time axis label of the Spline Editor. Markers appear as an aqua line that extends vertically through the Graph.
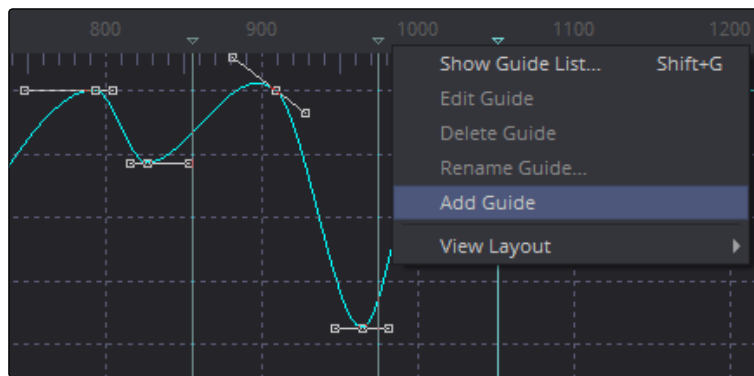
Click in the Time axis label to create a guide

Enable Guides Grab

With this option on, guides can also be moved by clicking at them in the Spline view. When it is off, guides can only be moved by clicking directly in the Guides area, which gives some protection against accidentally moving the guides.

To find out more about Guides, read Chapter 7, Timeline Editor.

Autosnap

to a field, frame or to guides. Autosnap is an option that is accessed through the Graph's contextual menu. There are two submenu options for Autosnapping. One option controls the snapping behavior when you drag keyframes, while the other option controls the snapping behavior of guides.

Right clicking in the Graph and choosing Options > Autosnap displays a submenu of Autosnap choices.

To find out more about AutoSnapping Points and Guides, read Chapter 7, Timeline Editor.

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 re-use, or to transfer tracking, path and animation data from one application to another.

Exported splines are assigned the file extension .spl for easier to identification.

To export a spline, do the following:

1. Select a spline in the Spline Editor.

2. Right click on the spline in the Graph to display the contextual menu.

3. Select Export from the contextual menu.

4. Choose from three format options in the submenu.

5. Enter a name and location in the file browser dialog, then click Save.

Tip: Exporting a spline gives you three options. You can export the Samples, Key Points or All Points.

To import an ASCII spline, do the following:

1. Animate the control.

2. Right click on the animated spline in the Graph.

3. From the contextual menu choose Import Spline.

4. In the File Browser dialog, select the spline curve .spl file, then click Open.

Importing a new curve will replace any existing animation on the selected spline.

Display Controls in the Spline View

The Spline view can also be used to display and modify any tool's splines, like in the case of a Hue Curves tool. Sometimes the screen real estate for a spline based tool like Color Curves or Hue Curves, or spline based values like Size Over Life in a Particle Emitter is limiting. The Spline view offers an excellent working environment in this case.

To edit the splines of Color curves or Hue curves in the Spline Editor, do the following:

1. Select the Color or Hue Curves tool in the flow.

2. In the Control Panel, right click over the header.

3. Choose Edit Splines from the contextual menu.

Bins

9

Bins

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.
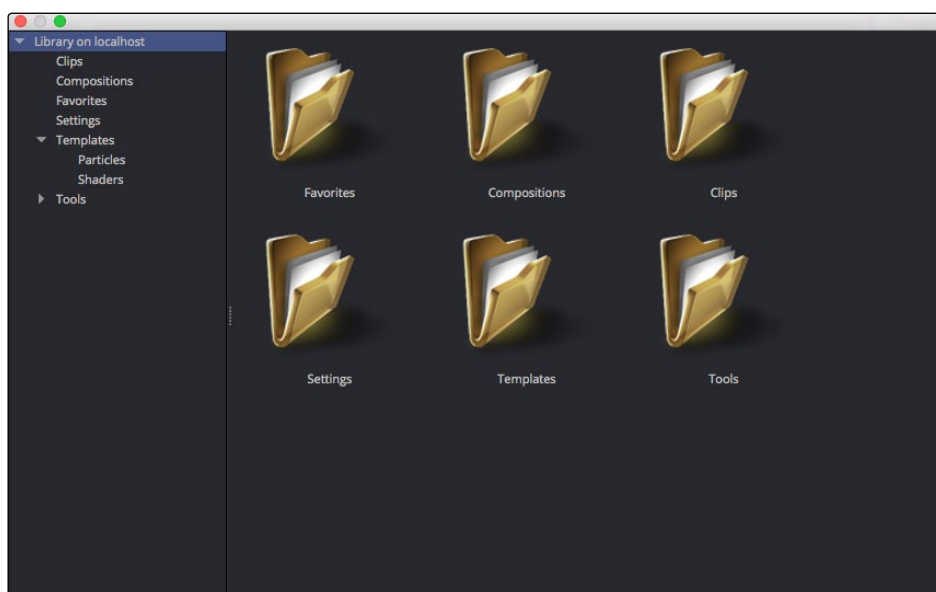
The Bins window

Bins Interface

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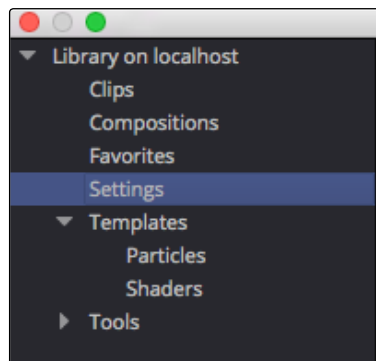
The Bins sidebar

The sidebar organizes content into bins or folders using a hierarchical list view. These folders can be organized however suits your workflow, but standard folders are provided for Clips, Compositions, Favorites, Settings and Tools. Parent folders contain sub folders that hold the content. For instance, the Tools bin is parent folder to all the categories of Tools. Parent folders can be identified by the disclosure arrow to the left of the name.

When you select a folder from the sidebar, the contents of the folder are displayed in the Contents panel as thumbnail icons.
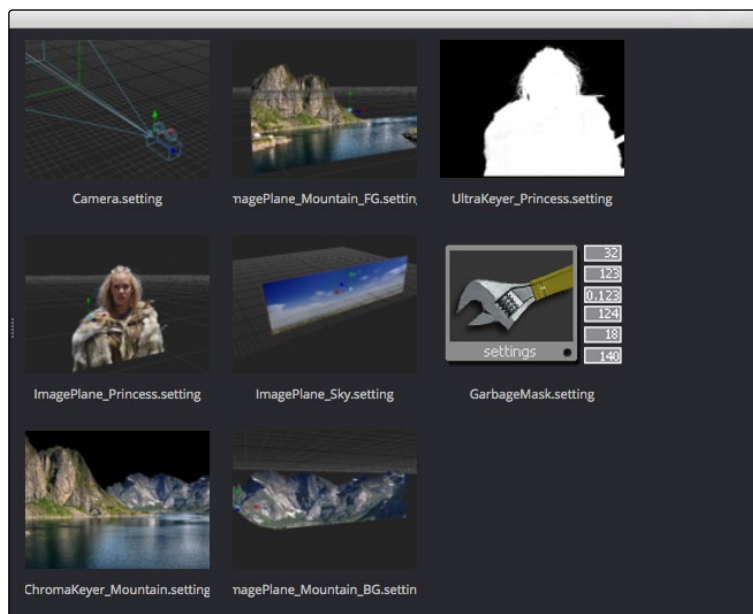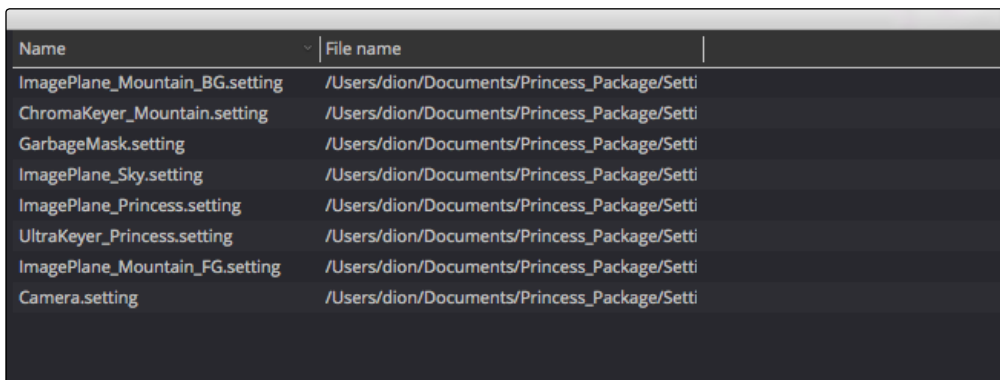
The Bins icon view

A contextual menu is used to access most of a bin's features. You show the contextual menu by right clicking in an empty area in the Contents panel. Right clicking on an item will show the same menu with additional options for renaming, playing or deleting the item.

One use of the contextual menu is to switch between viewing the contents as thumbnail icons or in a list.

To view a bin's contents in List view, do the following:

1. Right click in an empty area of the Contents panel.

2. From the contextual menu, choose View > Details.

The Bins List view

Clicking on the heading of a column in List view will sort the list in ascending order and clicking it again will reverse the sort order.

Each bin in the sidebar can be set to List view or Icon view independently of each other. So while you may have some bins you want to see as a list, others maybe easier to view as icons.

The icons can be adjusted to small, medium, large or huge by 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.

Organizing Bins

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

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

→ Choose File > Bins from the menu bar.

→ Click the Bins button in the main toolbar.

→ Press Command-B (Mac OS X) or Crtl-B (Windows).

The Bins button in the toolbar

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

1. In the sidebar, select the parent folder under which the new folder will be listed.

2. Right click in an empty area of the Contents panel.

3. From the contextual menu, choose New Folder.

4. Enter a name for the new folder, then click OK on the dialog.
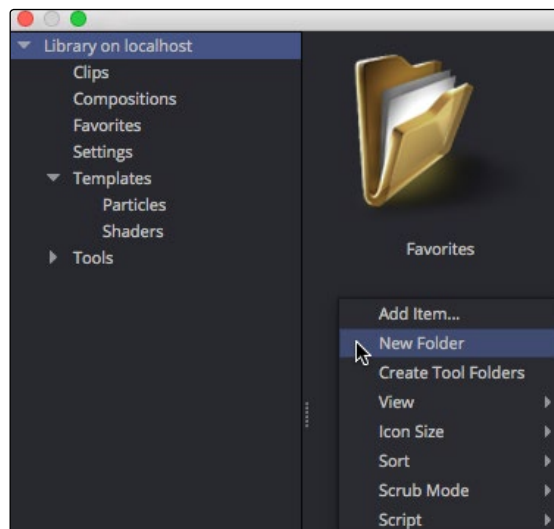
The New Folder menu item

To rename a bin folder, do the following:

1. Right click on the folder icon in the Contents panel.

2. Choose Rename from the contextual menu or press F2 on the keyboard.

To move a folder into or out of a parent folder, do the following:

1. Select the parent folder that contains the folder you want to move.

2. In the Contents panel, drag the folder into the sidebar where you want it moved.
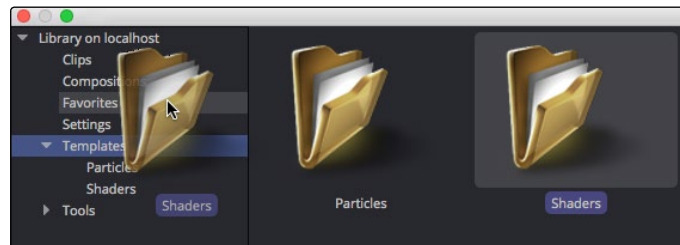
Drag a folder from the Contents panel to move it in the sidebar

When you drag a folder onto another folder in the sidebar, the folder you are dragging is nested under the targeted folder. Dragging it to the Library parent folder at the top of the sidebar will add it to the top level of the Bins window.

To remove an item from the bins, do the following:

1. Select the folder in the Contents panel.

2. Right click on the item and choose Delete from the contextual menu.

   Or

2. Press Command-Delete (Mac OS X) or Backspace (Windows).
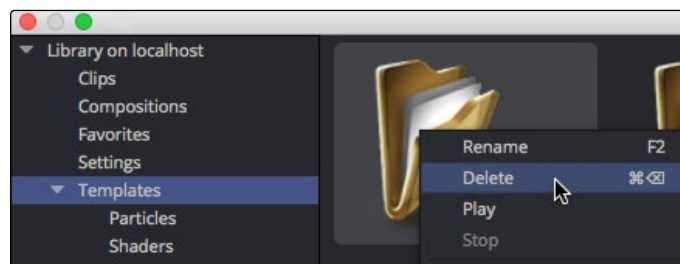
The Delete function in the contextual menu

Tip: You cannot undo removing a folder from the Bins window.

Adding and Using Content

You can add and use different types of content with bins. Fusion compositions, tools, saved tool settings, macros, tool groups and any file format that is supported in Fusion can be added to bins and then used in compositions at a later time.

To add an item to a bin, do the following:

1. Select a bin in the sidebar where you want to add the content.

2. Right click in the Contents panel.

3. Choose Add Item from the contextual menu.

4. Select the footage, files or flows in the file browser, then click Open.
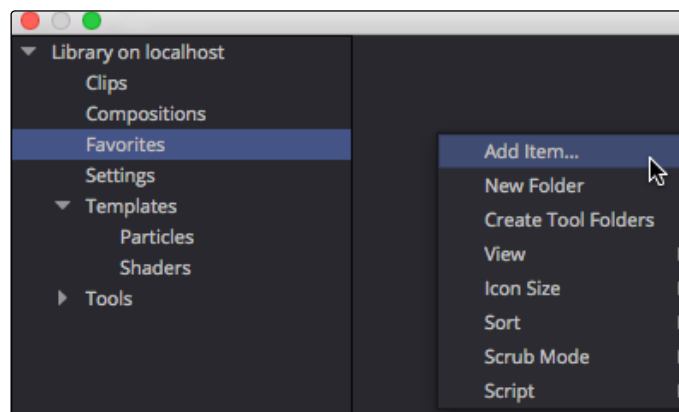
Add Item in the contextual menu

Tip: Unsupported files like PDFs can also be saved in the bin and the application that supports it will launch when you double click the file. It's a useful location for scripts and notes.

If you have an Mac OS X Finder window or Windows Navigator open, you can drag files directly into the bin as well. When adding an item to the bins, the files are not copied. A link is created between the content and the bins, but the files remain in their original location.

File Type Details

Some types of content have additional methods of being added to bins. Some additional methods are due to the file type and some because of where they are located.

Projects and Media

In addition to using the Add Item contextual menu as explained earlier, project files with the extension ".comp" and media can also be added to bins by dragging them to the Contents panel from an Mac OS X Finder or Windows file browser.

Tools

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

Tool Settings

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

Image Sequences and Stills

Image sequences are automatically identified on disk and loaded as clips rather than stills, so it is not necessary to select more than one frame from an image sequence when dragging it into a bin.

To ignore the image sequence and only import a single frame, hold Shift when you drag the frame in to a bin.

Using Content from Bins

Once you have content in your bins, you'll want to add them to a composition. In general you can either drag the content directly into the Flow Node Editor or double click it to add it to the flow, however each type of content behaves a little differently when added to a flow.

Media

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

Compositions

When a composition is added, it is opened in a new window. It is not added to the existing composition. To add a composition you must double click it in the bin to open it. Dragging a comp item onto an open composition will have no effect.

Tools, Tool Settings and Macros

When you add tools to a composition the methods you use and results you get are similar to adding tools using the Tool buttons in the toolbar. Dragging a tool allows you to place it anywhere in the flow, unconnected or, if you drag it over a pipe, inserted between two existing tools. Double clicking a tool in the bin will insert it after the Active tool in the flow. Dragging a tool from a bin into a Viewer will insert that tool after the currently viewed tool.

Working with Clips

Clips from image sequences, .MOV files and AVI files can be added to bins and previewed in different ways without having to first add the clip to a flow.

To play back a clip in a bin, do the following:

1. Right click over a clip in a bin.

2. Choose Play from the contextual menu.

The clip's thumbnail icon in the bin as well as the left Viewer are used to view the playback.

Scrubbing Footage

You can quickly preview the clip by scrubbing through it rather than playing it back.

To scrub a clip in a bin, do the following:

1. Hold the right mouse button over the clip thumbnail in the bin.

2. Drag the mouse left and right to scrub forward and backward.

The footage will appear in the left Viewer of the currently open composition.

Jog and Shuttle

You can scrub clips using one of two modes. Jog mode is the default mode. It moves the clip forward and backward as long as you are dragging the mouse. Once the mouse stops, the clip is paused.

You can choose Shuttle mode by right clicking over the clip's thumbnail in the bin, and choosing Scrub Mode > Shuttle.
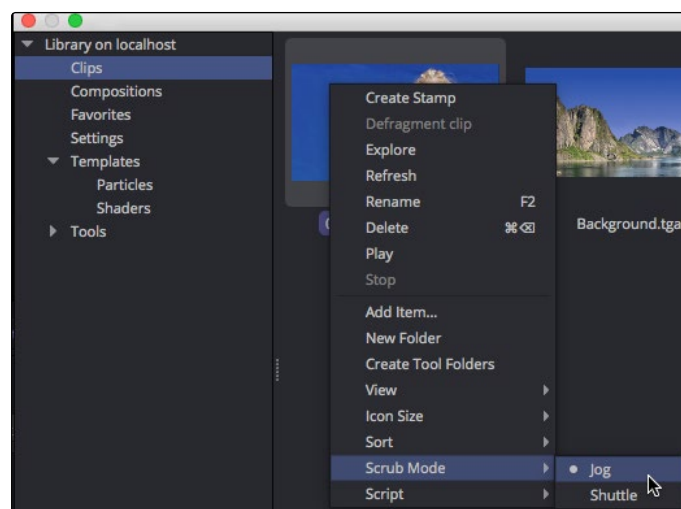
Jog and Shuttle contextual menu

Shuttle mode begins playing the clip forward or backward once you press the right mouse button and drag either left or right. The clip continues to play until the mouse button is released or you reach the end of the clip.

Stamp Files

Stamp Files are low resolution, local proxies of clips, which are used for play back of clips stored on a network server, or for very large image sequences.

To create a stamp file for a clip, do the following:

1. Right click on the clip in the bin.

2. Choose Create Stamp from the contextual menu.

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

Connecting Bins Over a Network

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

Tip: Remote bin functionality is only available in Fusion Studio.

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

1. Choose Fusion > Preferences (Mac OS X) or File > Preferences (Windows).

2. In the Preferences dialog, select Global > Bins > Servers in the list.
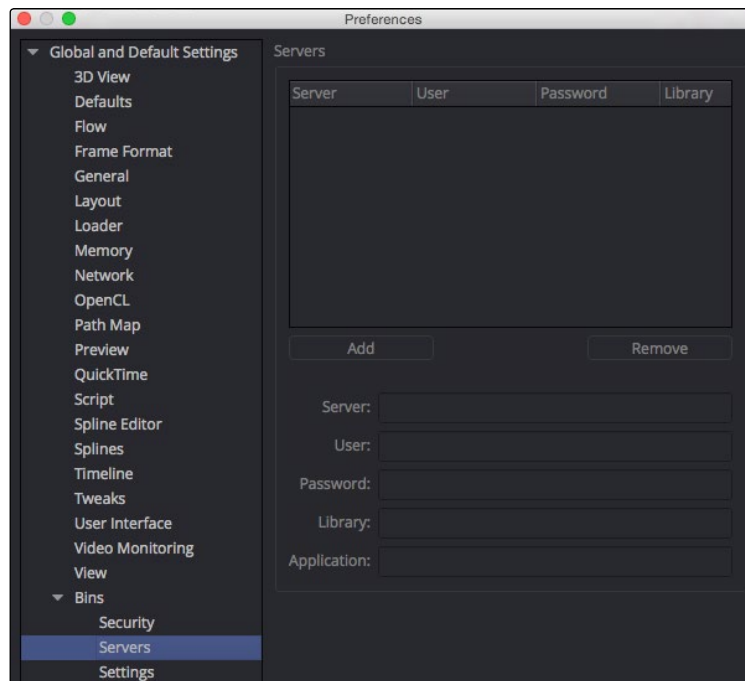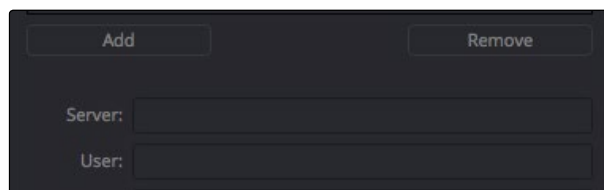
The bin servers Preferences panel

This panel shows a list of the bin servers available, with buttons below for entries to be added to or deleted from the list.

Adding a Remote Bin Entry

If you want to add a Remote bin to the list of available Remote bins, you can click the Add button in the bin servers Preferences panel. The text controls below the button will become enabled for editing. In the Server field type the system name or IP address where the bin is hosted.

Add the IP address where the Bin Server is hosted

Then add a User name and Password if one is needed to access the server.

The Library field lets you name the bins. So if you want to create a bin for individual projects, you would name it in the Library field and each project would gets its own bin.

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

Once you've finished setting up the bin server information and clicked Save in the Preferences window, you can open the bin window to test your bin server. Opening the bin 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 it requires, 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.

Polylines, Masks and Motion Paths

10

Polylines, Masks and Motion Paths

Polylines Overview

Polylines are splines that are used whenever a control is animated with a motion path or when a tool's effect is masked with a drawn shape. They are also used in the Paint and Grid Warp tools. In a more basic form, polylines are used to control the animation in the Spline Editor.

Since these splines are used for just about everything, they are extremely flexible, with a considerable amount of controls, modes and options.

This chapter offers an overview of polylines and their operation with sections on how to create paths or effect masks.

Polyline Types

You can draw polylines using B-Spline or Bezier spline types. Which you choose depends on the shape you want to make and your comfort with the spline style.

Bezier Polylines

Bezier Polylines are shapes that are composed of control points and direction handles. Several points together are used to form the overall shape of a polyline.

Bezier control point with direction handles
extended to create a smooth curve

Each control point has a pair of direction handles that are used to define the exact shape of the polyline segments passing through each control point. Adjusting the angle or length of the direction handles will affect whether that segment of the polyline is smooth or linear.

Bezier control point with direction handles
aligned to create a linear segment

Common graphics packages, such as Adobe Photoshop or Illustrator, will offer familiarity with many of the basic concepts that relate to the use of Bezier polylines.

B-Spline Polylines

A B-Spline polyline is similar to a Bezier spline, however, these polylines excel at 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 Polyline Types

Just because a shape using a B-Spline or polyline has been created does not mean that the spline type has to be used. You can convert from B-Spline to Bezier, or Bezier to B-Spline, as needed. Once you create a polyline, you can right click in the Viewer and choose Convert Bezier Spline to B-Spline or Convert B-Spline to Bezier from the spline's contextual menu. Only the appropriate option will be displayed.

When converting from one type to another, the shape is preserved. The new polyline generally has twice as many control points as the original shape to ensure the minimum change to the shape.

Although animation is preserved, the conversion process will not yield perfect results every time. It is a good idea to review the animation after you convert spline types.

Polyline Toolbars

Whenever a tool that contains one or more polylines is selected, the polyline is shown on all Viewers and the Polyline Toolbar is displayed along the side of each Viewer. The toolbar contains several buttons that make switching polyline modes and options easy to access.

The polyline toolbar

If you hover the cursor over any of the Polyline Toolbar buttons, a tooltip that describes the button's function appears. Clicking on a button will affect the currently active polyline or the selected polyline points, depending on the button.

You can change the size of the toolbar icons, add labels to the buttons, or make other adjustments to the toolbar's appearance in order to make polylines easier to use. All the options can by found by right clicking on the toolbar and selecting from the options displayed in the contextual menu.

Selecting a Polyline

It is possible to have several polylines in the Viewer at once, so it is 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.

→ To the right of the Viewer toolbar, click on the disclosure arrow to display a list of all polylines to choose from.
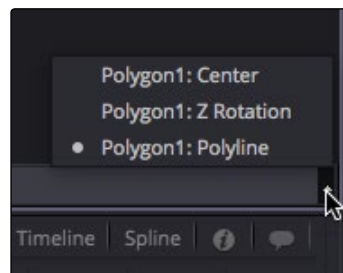
Click the Arrow icon to display a list of polylines and Control previews

Polyline Creation Modes

There are several different modes available from the toolbar for creating and modifying polylines. The specific mode used when a polyline is first added will depend on whether it is used as a path or a mask.

Each of the modes is described in more detail below.

Click Append

This mode is the default mode for mask creation. It's used to quickly define the rough shape of the mask, after which you switch to Insert and Modify mode to refine the mask further.

The Click Append toolbar button (Shift-C)

To create a mask using the Click Append mode, do the following:

1. Select Click Append from the toolbar or press Shift-C.

2. Click the cursor where you want to start the shape.

3. Move and click the cursor to append a point to the last one.

4. To close the shape, place the mouse pointer over the first point created and click when the cursor changes shape.

When a shape is closed, the polyline is automatically switched to Insert and Modify mode.

Although the Click Append mode is rarely used with paths, it can be helpful when you know the overall shape of a motion path but you don't know the timing yet.

Tip: Holding Shift while you draw a mask constrains subsequent points to 45 degree angles relative to the previous point. This can be very helpful when drawing regular geometry.

Insert and Modify

Masks, which are created in Click Append mode, automatically switch to Insert and Modify mode when the mask shape is closed. You can also manually switch to this mode by clicking the Insert and Modify button in the toolbar or using the Shift-I keyboard short cut. This mode makes it easier to add additional points and refine the shape of the mask. Dragging the control points or direction handles modifies existing points on the polyline.

The Insert Modify toolbar button (Shift-I)

Insert and Modify mode is also the default mode for creating motion paths. A new control point is automatically added to the end of the polyline, extending or refining the path, any time a parameter that is animated with a motion path is moved.

Draw Append

The Draw Append mode creates a freehand polyline shape directly on the Viewer, like drawing with a pencil or a paintbrush. This mode is ideal to use in conjunction with a tablet and for the creation of garbage mattes and effect masks around complex shapes.

The Draw Append toolbar button (Shift-D)

Protection Modes

In addition to the modes used to create a polyline, two other modes are used to protect the points from further changes after they have been created.

Modify Only

Modify Only mode allows existing points on the polyline to be modified, but new points may not be added to the shape.

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 cursor over the first point created, then click on the point.

→ Press Shift-O on the keyboard.

→ Click the Close button on the polyline toolbar.

→ Choose Closed from the polyline's contextual menu.

The Close toolbar button (Shift-O)

All of these options are toggles that can also be used to open a closed polygon.

Selecting and Adjusting Polylines

To create the shape you need for a mask or a motion path, you need to know how to manipulate the splines. In Fusion there are a number of simple techniques for selecting, moving and smoothing a spline, but also more complex adjustment techniques for scale, skewing and twisting a spline.

Polyline Points Selection

To select one or more control points on a polyline, do one of the following:

→ Click directly on the control points.

→ Lasso around the points.

To add or remove points from the current selection, do one of the following:

→ Hold the Shift to select a continuous range of points.

→ Hold Command (Mac OS X) or Ctrl (Windows) and click each control point you want to add or remove.

→ Press Command-A or Ctrl-A to select all of the points on the active polyline.

Tip: Once a control point is selected, you can press Page Down or Page Up on the keyboard to select the next control point in a clockwise or counter clockwise rotation. This can be very helpful when control points are very close to each other.

Moving Polyline Points

The selected polyline points can be moved using either the keyboard or the mouse.

To move selected control points using the cursor, do one of the following:

→ Drag on the selected points anywhere in the Viewer.

→ Hold Shift while dragging to restrict movement to a single axis.

→ Hold Option (Mac OS X) or Alt (Windows) and drag anywhere in the Viewer to move the selected control point.

To move selected control points using the keyboard, do one of the following:

→ Press the left, right, up or down arrow keys on the keyboard.

→ Hold Command-Arrow keys (Mac OS X) or Ctrl-Arrow Keys (Windows) to move in smaller increments.

→ Hold Shift-Arrow keys to move in larger increments.

Smoothing a Polyline Segment

If you want to shape the polyline and control its slope, you can choose to smooth a spline segment by adjusting the Bezier direction handles.

To smooth the selected points on an active polyline, do one of the following:

→ Press Shift-S.

→ Click the Smooth button on the Polyline toolbar.

→ Choose Smooth from the polyline's contextual menu.

The Smooth button in the toolbar (Shift-S)

Linearizing a Polyline Segment

To make certain that a polyline segment is perfectly straight, that segment needs to be linearized. A linear segment aligns the Bezier direction handles with the segment and therefore has no curvatures. The segment is always drawn in a straight line between two points on the polyline.

To linearize the selected points on an active polyline, do one of the following:

→ Press Shift-L.

→ Click the Linear button on the polyline's toolbar.

→ Choose Linear from the polyline's contextual menu.

The Linear button in the toolbar (Shift-L)

Transforming Individual or Multiple Points

Select the points to be transformed, then do one of the following:

→ Hold T and drag to twist.

→ Hold S and drag to scale.

→ Hold X and drag to scale horizontally only.

→ Hold Y and drag to scale vertically only.

→ Hold O and drag to offset the points perpendicular to the tangent.

The position of the cursor when the transformation begins becomes the center used for the transformation.

Deleting Selected Points

You can delete a selected point or group of points by pressing the Delete (Mac OS X) or Backspace (Windows), choosing Delete from the contextual menu or by clicking the Del button in the toolbar. The shape of the polyline changes to reflect the removal of these points.
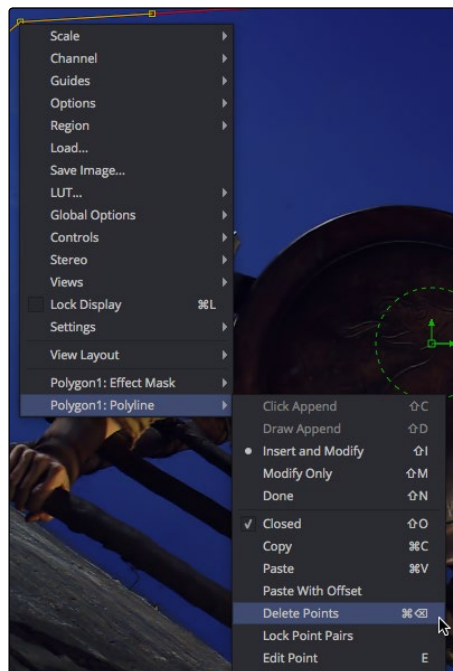
The polyline contextual menu has a Delete command

Tip: Deleting all of the points in a polyline does not delete the polyline itself. To delete a polyline, delete the tool or modifier that created the polyline.

Bezier Handles

For Bezier polylines, each control point has two direction handles that adjust the slope of a curve through the control point. These direction handles only appear when the point is selected.

Dragging a direction handle makes adjustments to the curve of the segment that emerges from the control point. The direction handle on the opposing side of the control point will also move to maintain the relationship between these two handles.

To break the relationship between direction handles and adjust one independently, hold Command (Mac OS X) or Ctrl (Windows) while dragging a handle. Subsequent changes will maintain the relationship, unless the Command or Ctrl is held during each adjustment.

Hold Command or Ctrl to adjust one handle independently

If you want to adjust the length of a handle without changing the angle, hold Shift while moving a direction handle.

Edit Points

The Edit Points dialog can be used to reposition control points using precise X and Y coordinates. Pressing the E key on the keyboard will bring up the dialog and allow you to reposition one or more selected control points.
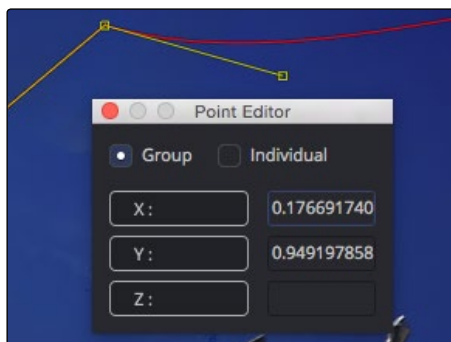
The Edit Points dialog can be used to position control points

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 Edit Points dialog with Offset values

If you are not sure of the exact value, you can also perform mathematical equations in the dialog box. For example, typing 1.0-5 will move the point to 0.5 along the given axis.

Reduce Points

When freehand drawing a polyline or an editable paint stroke, the spline is often created using more control points than you need to efficiently make the shape. If you choose Reduce Points from the polyline's contextual menu or toolbar, a dialog box will open allowing you to decrease the number of points used to create the polyline.

The Reduce Points button in the toolbar

The overall shape will be maintained while eliminating redundant control points from the path. When the value is 100, no points are removed from the spline. As you drag the slider to the left you reduce the number of points in the path.

Shape Box

If you have a polyline shape or a group of control points you want to scale, stretch, squish, skew or move, you can use the Shape Box to easily perform these operations.

To enable the shape box, do one of the following:

→ Click the Shape Box toolbar button.

→ Choose Shape Box from the contextual menu.

→ Press Shift-B.

The Shape Box button in the polyline toolbar

If there are selected points on the polyline when the Shape Box mode is enabled, the shape box is drawn around those points. Otherwise, you can drag the shape box around the area of control points you want to include.

If you want to freely resize the shape box horizontally and vertically, you can drag a corner handle. Dragging a handle on the side of the shape box resizes the polyline along a specific axis.

Dragging a side handle resizes along a specific axis

Holding Command or Ctrl while dragging a shape box handle will apply adjustments from the center of the shape box, constraining the transformation to the existing proportions of the shape box. Holding Shift while dragging a corner handle affects only that handle, allowing skewed and non-uniform transformations.

Hold shift while dragging a corner to perform non-uniform transformations

Showing and Hiding Points and Handles

It is often difficult to identify individual points when they are placed closely together. You can choose to display both points and their direction handles, just points, or just handles. These display mode options are selected using the Show Key Points and Show Handles toolbar buttons, or from the polyline's context menu.

The Show Key Point and Show Handles
buttons in the toolbar

You use these options to simplify the screen display when adjusting control points placed closely together and to avoid accidentally modifying controls and handles that are adjacent to the intended target.

Stop Rendering

While points along the polyline are being moved, the results are rendered to the Viewer to provide constant interactive feedback. Although extremely useful, there are situations where this can be distracting and can slow down performance on a complex effect. To disable this behavior so renders only happen when the points stop moving, you can toggle the Stop Rendering button in the toolbar or select this option from the polyline contextual menu.
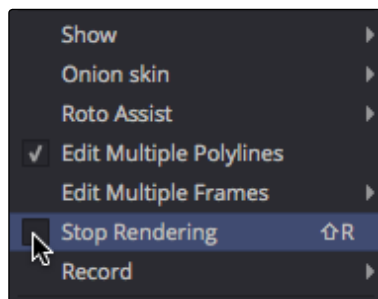
Click the Stop Rendering menu in the
polyline contextual menu to limit rendering

Importing and Exporting Polylines

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

Native Format

To save a polyline shape in Fusion's native ASCII format, you right click on the header of the Mask tool in the Control Panel and select Settings > Save As from the contextual menu. Provide a name and path for the saved file and select OK to write a file with the .setting extension. This file will save the shape of a mask or path, and any animation splines applied to its points or controls.

To load the saved setting back into Fusion, you first create a new polyline of the same type, then select Settings > Load from the mask's context menu or drag the .setting file directly into the Flow.

If you want to move a polyline from one comp to another, you can also copy the tool to the clipboard, open your second comp and paste it from the clipboard into the new comp.

Roto Assist

You can enable the Roto Assist button in the toolbar when you begin drawing your shape to have points snap to the closest edge as you draw the shape. The points that have snapped to an edge are indicated by a cyan outline. There are three main Roto Assist options.

The Roto Assist options in the toolbar

→ Multiple Points: Allows adding multiple points along an entire edge with a single click instead of having to add each point individually.

→ Distance 8: Defines the range within which searching for an edge will take place.

→ Reset: Used for resetting the snap attribute of the snapped points. After resetting, the points will become unavailable for tracking.

Mask Types

Typically, a tool applies its effect to every pixel in an image. Masks are single channel images that restrict an effect to specific regions of an image. Masks can be created using primitive shapes, such as circles and rectangles, complex polyline shapes, or by extracting channels in another image. Multiple masks of various types can be combined to produce complex and powerful results. In Fusion there are three ways you can apply a mask to tools. You can use an effects matte, garbage mask and pre-masking.

Masks are created and animated using a combination of primitive shape masks or more complex polyline shapes. Any number of masks of any type can be combined together on a single tool. Two masks can combine together to create a more complex shape or one mask can subtract from another to cut holes into the mask channel.

Effects Mask

Effect masks are by far the most common type of mask applied to tools. They are applied post effect, which means the tool's result is calculated first and then the effects mask is used to limit the result.

A Polygon tool used as an effects mask

Although many tools support effects masking, there are a few where this type of mask does not apply, notably Savers, Time tools, and Resize, Scale and Crop tools.

Garbage Mattes

Garbage Mattes are a special type of effect mask that are unique to the keying tools (difference keyer, ultra keyer, chroma keyer, matte and luma keyer). Garbage matting is used to isolate portions of the image and to explicitly discard or include the area defined by the mask. It's used to remove rigging from blue and green screen shots and to isolate a section from an image not suited to chroma or color keying.
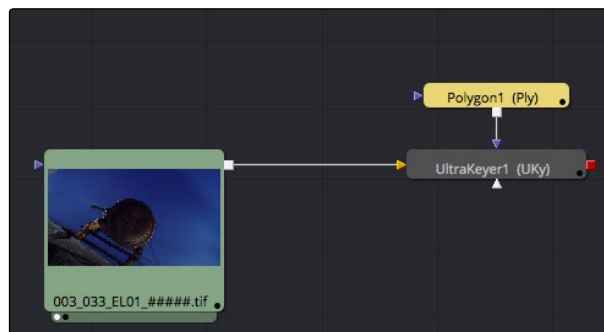
A garbage matte added to the ultra keyer

The option of creating a garbage matte is available from the contextual menu when one of the keying tools is selected. The various garbage matte types and shapes are identical to effect masks.
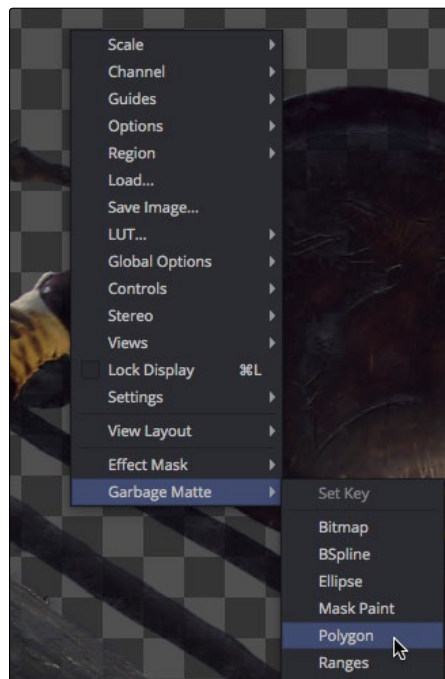
Selecting a garbage matte from the contextual men

You choose whether a garbage matte is applied to a keying tool as opaque or transparent using the garbage matte's drop down menu in the Control Panel for the key.

Garbage Matte setting in the keyer's controls

Pre-Masking

Unlike effect masks, a pre-mask is used by the tool before the effect is applied. This usually causes the tool to render more quickly and to produce a more realistic result.

In the case of the Highlight and the Glow tools, 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 tool, 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.

Mask Tools

Masks are like other creator tools in the flow, with the exception that they create a single channel image rather than a full RGBA image. Most of these Mask tools are also located under the Viewers, with the exception of Ranges Mask and Triangle Mask.

The Mask tools are described in the Tool Reference Manual.

Polygon Mask

Polygon Masks are user created Bezier shapes. This is the most common type of polyline and the basic workhorse of rotoscoping.

B-Spline Masks

B-Spline Masks are user created shapes made with polylines that are drawn using the B-splines. They behave identically to polyline shapes when linear, but when smoothed, the control points influence the shape through tension and weight. This generally produces smoother shapes while requiring fewer control points.

Bitmap Masks

The Bitmap Mask allows images from the flow to act as masks for tools 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 tools.

Wand Mask

A Wand Mask provides a crosshair that can be positioned in the image. The color of the pixel under the crosshair is used to create a mask, where every contiguous pixel of a similar color is also included in the mask. This type of mask is ideal for secondary color correction.

Ellipse, Rectangle and Triangle Masks

These are primitive shape masks. See Rectangle, Ellipse, and Triangle in the Tool Reference.

Ranges Mask

Similar to Bitmap Mask, the Ranges Mask allows images from the flow to act as masks for tools 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.

Animating Polyline Masks

When Polygon or B-Spline Masks are added in the flow, they are automatically ready to be animated. All you have to do to animate a mask is move the playhead to a new frame and make the required changes to the mask. A new keyframe is added in the Spline Editor and Timeline Editor. This one keyframe controls the position of all of the points that comprise the mask for that frame. The polygon or B-Spline is automatically morphed between the two keyframes. 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 Control Panel for the mask, right click in the bottom of the panel where it says Right Click Here For Shape Animation. From the contextual menu choose Remove Bezier Spline. If you decide you need to animate the mask at a later time, right click in the same area again and choose Animate.

Adding and Removing Points from an Animated Mask

When adding points to an animated mask, the new point is fit into the shape at all keyframes. Deleting a point removes that point from all keyframes in the animated mask.

Published Points

Although you can rapidly animate the entire shape of a polyline, the Spline Editor and Timeline only display one keyframe for the entire shape at any given frame. The default setting does not allow a specific point's animation to be controlled explicitly. If you are working on a complex mask that requires more precise timing or interpolation for control points, you can expose individual control points on a polyline by publishing them.

Publishing a point on a polyline removes that point from the standard animation spline. Once removed, this point will not be connected to paths, modifiers, expressions, or trackers that are connected to the main polyline.

Publishing a Point

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.

The Publish Points button in the polyline toolbar

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 Control panel

The onscreen control indicates published points on the polyline by drawing the point as a much larger point. 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

Publish Point to Path

When a point is published, any animation already applied to that point is removed. However, if you need to keep the animation, you can use the Publish Point to Path option. This polyline contextual menu option publishes the selected points and converts their existing animation to a path. You can also use the Publish Point to Path button in the Polyline toolbar.
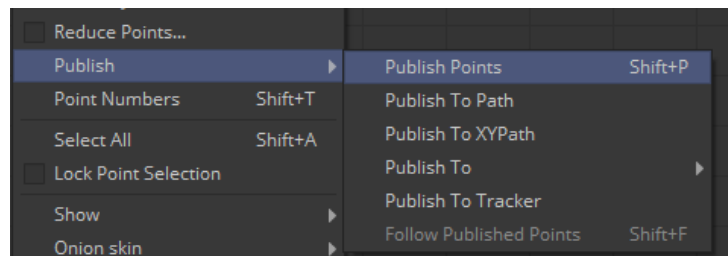
The Published Point to Path contextual menu option

Follow Published Points

There are times when you will need to have points in a mask that are between two published points, follow the motion of the published points but still maintain their relative offset and shape. For this reason, points in a polyline mask can be set to Follow Published Points using the polyline's contextual menu.

When a point of an effect mask is set to follow points, the point will be drawn as a diamond shape rather than a small box.

A control point set to Follow Points

When this mode is enabled, the new following points will maintain their position relative to the motion of any published points in the mask, 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.

Double Polylines

The standard soft edge control available in all mask tools softens the entire mask equally. However, there are times, particularly with a lot of motion blur, when softening part of the curve while keeping other portions of the curve sharp is often required.

This form of softness is called Non-Uniform Softness, which is accomplished by converting the shape from a single polyline to a double polyline. The double polyline is composed of two shapes, an inner and an outer shape. The inner shape is the original shape from the single polyline, whereas the outer shape is used to determine the spread of the softness. The further the outer shape gets from the inner shape, the softer that segment of the shape becomes.

A double polyline uses an inner and
outer shape for non-uniform softness

Converting a Single Polyline to a Double Polyline

To convert a mask into a double polyline, click the Double Polyline button in the polyline toolbar or right click in the Viewer and select Make Outer Polyline from the mask's contextual menu. The shape will be converted into an inner and an outer polyline spline. Both polylines start with exactly the same shape as the original single polyline. This keeps the mask sharp to start with and allows any animation that may have already been applied to the shape to remain.

Make Double Polyline button

The control points on the outer shape are automatically parented to their matching points on the inner shape. This means that any changes made to the inner shape will also be made to the outer shape. The relationship is one way. Adjustments to the outer shape can be made without affecting the inner shape.

A dashed line drawn between the points indicates the relationship between the points on the inner and outer shapes.

Adding Softness to a Segment

The outer shape is drawn using a green dashed line instead of a solid line to help distinguish it from the inner shape. If you want to select the outer shape, use the Tab key to cycle between the onscreen controls until the dashed outline is visible, or you can select the outer polyline using the contextual menu's Controls > Select menu.

Once the outer polyline is selected, you can drag any of the points away from the inner polyline to add some softness to the mask.

Tip: Press Shift-A to select all the points on a shape then hold O and drag to offset the points from the inner shape. This gives you a starting point to edit the falloff.

The further the outer shape segment is from the inner shape, the larger the falloff will be in that area.

Adding Additional Points to the Shape

It is not necessary for every point on the inner shape to have a match on the outer shape, or vice versa. You can add additional control points to refine the shape of either shape.

Each polyline stores its animation separately; however, if a point is adjusted on the inner shape that is parented to a point on the outer shape, a keyframe will be set for both splines. Adjusting a parented point on the outer shape only sets a keyframe for the outer shape's spline. If a point that is not parented is adjusted, it will only set a keyframe on the relevant spline. You can disable this behavior entirely for this polyline by disabling "Follow Inner Polyline" from the contextual menu.

Locking/Unlocking Point Pairs

If you want to parent additional control points you can select the points, right click in the Viewer and choose Lock Point Pairs from the contextual menu for either spline. This will cause the selected point on the outer shape to become parented to the selected point on the inner shape.

Any animation already applied to either point is preserved when the points become parented.

To unlock a point so it is no longer parented, select the point, right click in the Viewer and deselect Lock Point Pairs from the contextual menu.

Motion Paths

Motion Paths are polylines that define the movement for positional controls of layers, effects and masks. You add motion paths to coordinate controls, such as a Merge tool's center or published polyline points on a mask. Coordinate controls are represented onscreen with a crosshair or an X.

A Center Offset onscreen control

It is not possible to add a motion path to a single value control, like a blur strength or merge angle.

Types of Motion Paths

There are three types of Motion paths: Poly paths, XY paths and 3D motion paths for 3D scenes. A Poly path uses two splines to record the path, one for shape of the path and a displacement spline for the timing along the path.

The XY path type employs a spline for the X position of the point and another for the Y position. The XY paths are explained in detail toward the end of this chapter.

The Poly path is the default type for motion paths and most documentation in this chapter assumes that this type is used.

Poly Paths

To create a Poly path, do the following:

1. Position the playhead on the frame where the motion will begin.

2. Position the control for the layer, effect or mask at its starting position.

3. Right click on the onscreen control in the Viewer.

4. Choose Path from the contextual menu for that control.

    The control's value boxes will turn green to indicate that the parameter is now animated and on a keyframe. Any changes made to the control will cause a new keyframe to be created.

5. Move the playhead to a new frame.

6. Drag the onscreen control or adjust the Offset or Center values in the Control panel.

A keyframe is automatically created on the motion path and a polyline is drawn from the original keyframe to the new one. The position of the control is interpolated between the two keyframes.

You can continue adding points by moving the playhead and adjusting the control until the entire motion path is created.

Upon completion, set the polyline to Insert and Modify mode by pressing Command-I (Mac OS X) or Ctrl-I (Windows) on the keyboard or pressing the Modify button on the toolbar.

Do not 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.

Path Modifier

A Path is a Modifier tool. When the path is added to a control, controls for the path are visible in the Modifiers tab of the Tool Control window.
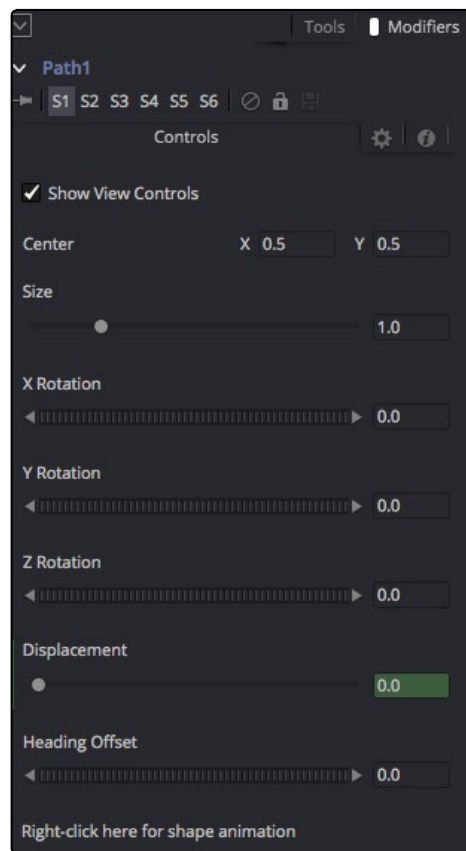
The path's controls in the Modifier tab of the Control panel

Displacement Splines and Path Timing

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

Displacement splines are used to control the speed of the movement along a path. To slow down, speed up, stop or even reverse the motion of the control along the path, adjust the points for the path's displacement in the Spline Editor.

A displacement value of 0.0 in the Spline Editor indicates that the control is at the very beginning of a path. A value of 1.0 indicates that the control is positioned at the end of the path.

Types of Points

Displacement paths are composed of Locked and Unlocked points. Whether a point is locked or not is determined by how it was added to the polyline. Locked points on the motion path will have an associated point on the displacement spline. Unlocked points will not have a corresponding point on the displacement spline. Each has a distinct behavior, as described below.

Locked Points

Locked points are the motion path equivalents of keyframes. They are created by changing the playhead position and moving the animated control. These points indicate that the animated control must be in the specified position at the specified frame.

The locked points are displayed as larger sized hollow squares in the view. Each locked key has an associated point on the path's displacement in the Spline Editor.
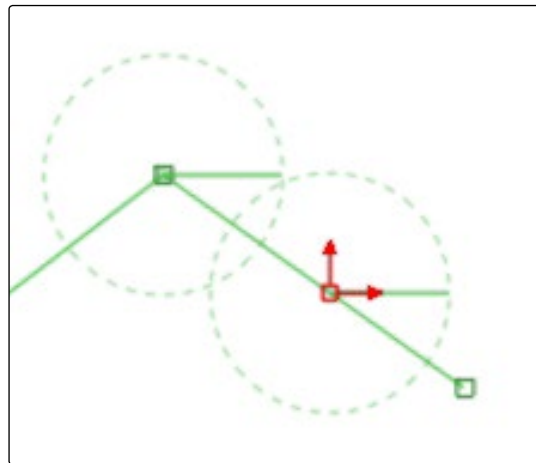
A locked point in the Viewer

Deleting a locked point from the motion path will change the overall timing of the motion.

Unlocked Points

Unlocked points are created when additional points are added to the motion path while in Insert and Modify mode. These points are used to adjust the overall shape of the motion path, without directly affecting the timing of the motion.

Unlocked points do not have corresponding points on the path's displacement spline. They are represented in the Viewer as smaller, solid square points.

To experience the difference between locked and unlocked points, do the following:

1. Add a Text tool to the flow and type a word in the Text tool.

2. Position the text's center in the upper left corner of the frame.
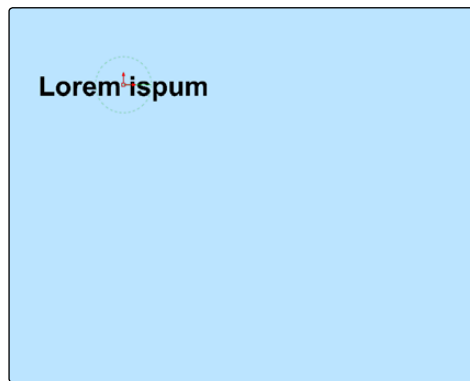
Text placed in the upper left corner of the frame

3. Set the Playhead at frame 0.

4. In the Viewer, right click on the text center and choose Animate from the contextual menu.

   This creates the first locked point of the path.

5. Position the playhead on frame 45.

6. Move the text center to the center of the screen.

Moving the playhead and repositioning the text adds a locked point

This sets the second locked point.

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

The path's displacement spline

At a value of 0.0, the control will be located at the beginning of the path. When the value of the displacement spline is 1.0, the control is located at the end of the path.

8. Select the keyframe at frame 45 in the displacement spline and drag it to frame 50.

The motion path is now 50 frames long, without making any changes to the motion path's shape.

If you try to change this point's value from 1.0 to 0.75 it cannot be done because the point is the last in the animation, so the value must be 1.0 in the displacement spline.

9. Position the playhead on frame 100 and move the text center to the bottom right corner of the screen.
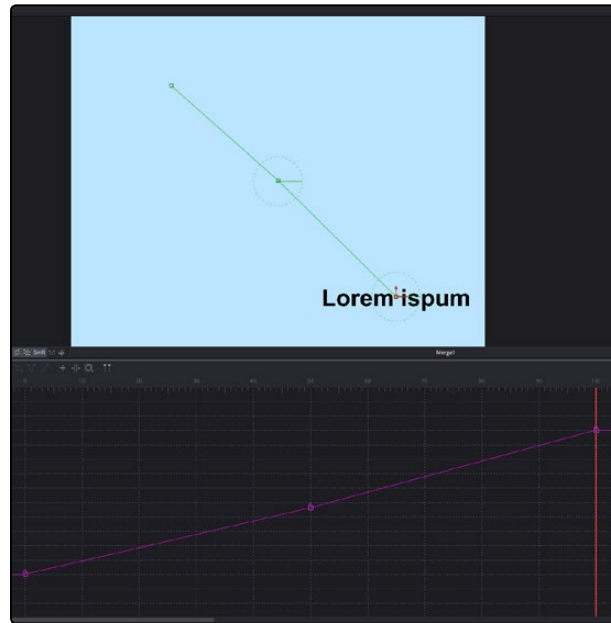
Position the playhead at 100 and drag to reposition the text, creating another locked point

This will create an additional locked point and set a new ending for the path.

10. Select the motion path spline by using the Tab key to cycle controls until the path is selected.

Currently, the path is in Click Append mode.

11. To insert some points, click the Insert and Modify button in the toolbar.

Click the Insert and Modify button in the toolbar

12. Click on the path and create two new points: one part-way between the first and the second points and the other part-way between the second and the third.

Add two new points to the path

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

Adjust the path using the newly added points

The two points just added are not present in the motion path's displacement spline. These are unlocked points, used to shape the motion but unrelated to the timing of the path. This behavior can be used to make the layer's motion pause briefly.

14. Select the point at frame 50 in the displacement spline and hold down Command (Mac OS X) or Ctrl (Windows) while dragging it to frame 65. The point is copied at frame 65.

Create a pause in the motion by copying locked points

15. Select the point at frame 0 and at frame 50 and press Shift-S while in the Spline Editor to smooth the displacement curve.

    This will cause the motion to accelerate slightly at the beginning and slow to a stop at the center.

16. Render a preview of frames 0 to 100 to see the results so far.

17. Remove the preview when done.

18. In the Viewer, delete the unlocked point added between the first point and the middle point. The spline segment of the motion path will become linear.

    Deleting this point did not change the amount of time it takes the spline to reach the center of the screen, only the path it takes to get there.

19. Step through a few of the frames in the scene to confirm this.

20. Now delete the locked point in the center of the screen.

    Removing this point changes the timing of the animation rather profoundly because the key point in the displacement spline is also removed.

Knowing the difference between locked and unlocked points gives you independent control over the spatial and temporal aspects of motion paths.

Locking and Unlocking Points

You can change an unlocked point into a locked point, and vice versa, by selecting the point(s) and choosing the Lock point option from the contextual menu.

XY Paths

The XY path type uses a separate spline for the position along the X-axis, as it does for the position along the Y-axis.

To animate a coordinate control using an XY path, do the following:

→ Right click on the onscreen control and choose Control name > Modify With > XY Path from the contextual menu.
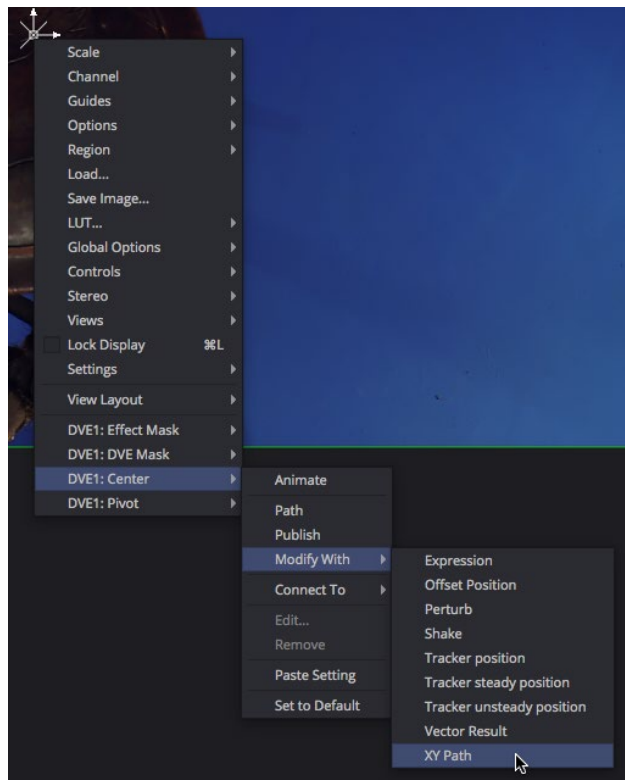
Choosing XY Path for a DVE's center

At first glance, XY paths work like displacement paths. To create the path, position the playhead and drag the onscreen control where you want it. Position the playhead again and move the onscreen control to its new position. The difference is that the control points are only there for spatial positioning. There are no locked points for controlling temporal positioning.

The Spline Editor for the XY path displays the X and Y channel splines. Changes to the controls position will be keyframed on these splines. The advantage to the XY path is that you can explicitly set an XY coordinate at a specific time for more control.

Tip: XY path and Poly path can be converted between each other from the contextual menu. This gives you the ability to change methods to suit your current needs without having to redo animation.

Switching Default Paths

If you want to change the Default Path type to XY path you can choose Fusion > Preferences > Globals (Mac OS X) or File > Preferences > Globals (Windows) and select the Defaults category. In the Point With pop up menu, choose XY path. The next time Animate is selected from a coordinate control's contextual menu, an XY path will be used instead of a displacement path.
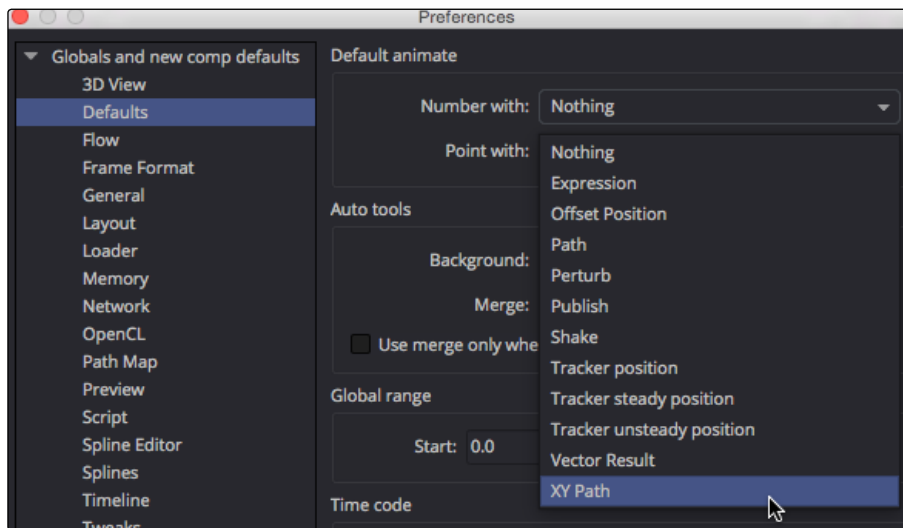
The Point With pop up menu is used to change the default motion path type

General Path Techniques

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.

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. This would require two paths. The bee would be connected to a first path, which would be a figure eight. This first path's center would be connected to another path defining the forward motion of the bee through the scene.

Copying and Pasting Motion Paths

It is possible to copy an entire motion path to the clipboard and then paste it onto another tool or path or composition.

To copy a motion path, do the following:

→ In the Control Panel's Modifier tab, right click on the path's control header and choose Copy from the contextual menu.

To cut a motion path out of a tool, do the following:

→ In the Control Panel, right click on the path's control header and choose Cut from the contextual menu.

To paste the copied path over another path, do the following:

→ In the Control Panel, right click on the path's control header and choose Paste from the contextual menu.

The settings of the old motion path will be overwritten with the settings of the one in the clipboard.

Removing Motion Paths

To remove a control from a motion path, right click on the motion path's control header and select Delete from the motion path's context menu. Alternately, right click on the animated control and select Remove "Name of Modifier."

Recording

You can animate both of the control's spatial and temporal information at the same time using the Record mode. This is useful when both position and speed are crucial to achieve the desired result.

Right click on the desired path to display its context menu and select Record from the contextual menu. This displays a submenu of available data that may be recorded.

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

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

Tracking

11

Tracking

Tracking Overview

Tracking is one of the most powerful automation tools available to a compositor. A tracker can follow a pattern of pixels in an image over time, generating a motion path from the pattern's movement. This information is then used to perform a variety of tasks, including image stabilization, matching the movement of one image to another and driving the positions of points on a mask.

Image with tracking patterns

Detailed reference for each of the controls in the tracker can be found in the Tool Reference Manual. This chapter provides more general information on how to use a tracker.

Basic Workflow

The basic approach to tracking has two steps. A pattern is selected and tracked, then the results are applied to one of several operations using the tracking data. The layout of the Tracker tool's tabs in the Control Panel reflects the workflow. The Tracker tab is where you select the patterns and perform the track. The Operations tab is where you decide how the tracking data is used.
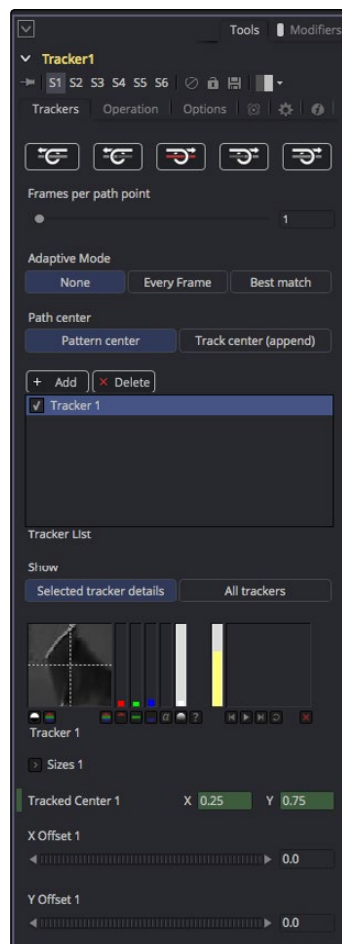
The Tracker Control Panel

Tracking Uses

There are four primary applications for the data that results from tracking.

Stabilizing

You can use one or more trackers to remove all of the motion from the sequence or to smooth out vibration and shakiness. When you use a single tracker pattern to stabilize, you only stabilize the X and Y position. Using multiple patterns together, you are able to stabilize position, rotation and scaling.

Match Moving

The reverse of stabilizing is Match Moving, which detects position, rotation and scaling in an image sequence using one or more patterns. Instead of removing that motion, it is applied to another image that matches the movement so that the two images can be composited together.

Corner Positioning

Corner Positioning tracks four patterns which are then used to map the four corners of a new foreground into the background. This technique is generally used for sign or TV screen replacement.

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.

Connecting Parameters to a Tracker

The position of a tracker pattern can also be used to drive the position of a center or axis control. For example, once you track the face of a suspect on a surveillance tape, you then connect the center of a mask on a Blur tool to the result in order to obscure the features.

As an example, to connect a Mask tool's Center parameters to a tracker, do the following:

1. In the Control panel, right click over the Mask's Center Parameter label.

2. From the contextual menu, choose Connect To > Tracker 1: Offset Position.
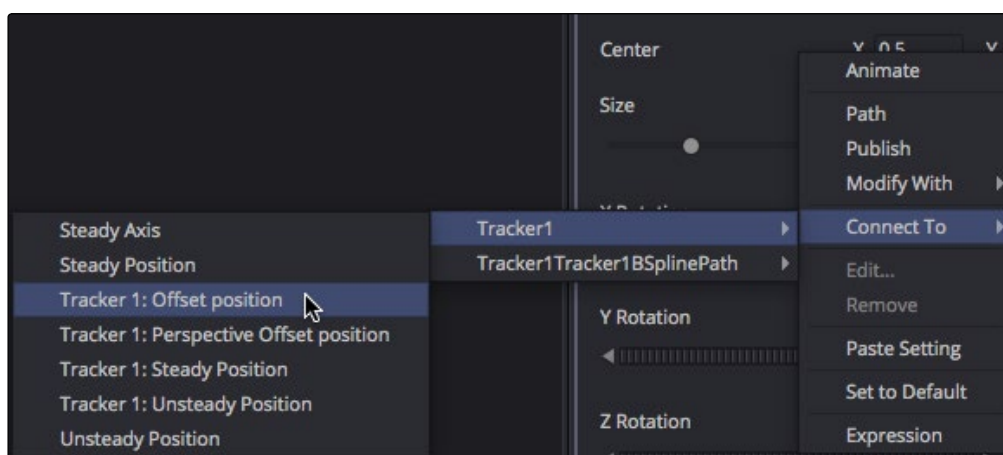
The Connect To submenu for connecting a track to a parameter

What Is a Pattern?

A Pattern is the region of pixels that are selected for tracking in an image. The pattern is represented in the Viewer as a solid rectangle when the Tracker tool is active. A single Tracker tool can have many patterns defined. Each pattern will produce its own path.

A pattern in the Viewer

When you add a Tracker tool to the flow, you start with one pattern displayed in the Viewer as a small rectangle. When the cursor is placed over the pattern rectangle, the control expands and two rectangles appear. The outer has a dashed line and the inner has a solid line. The outer rectangle is the search area and the inner rectangle is the pattern.

If you need to select a new pattern, you can move it by clicking on the small box in the top left of the pattern rectangle.

To move a pattern in the Viewer, drag it from the upper left corner

While moving the pattern rectangle, an overlay pop-up will appear. The overlay shows a zoomed version of the pixels contained within the rectangle to help you precisely position the pattern.

The pattern rectangle can also be resized by dragging on the edges of the rectangle.

Search Area

When the cursor is over the onscreen control for a pattern, a second rectangle with a dotted border will appear surrounding the pattern. This outer rectangle represents the search area.

When moving from one frame to another while tracking, the tracker will search a region surrounding 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 it will take to track. 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. It is important to take the speed and direction of the motion into consideration when setting the search area.

Resize the search area by dragging on
the edges of the dotted outline

Tip: The size of the search area can be animated over time. This is useful for times when the track pattern accelerates or decelerates rapidly part of the way through the scene.

Adding and Removing Patterns

Although each Tracker tool starts with only a one pattern, a single tool is capable of tracking multiple patterns. Additional patterns can be added by clicking on the Add button immediately above the Tracker List control.

To add an additional tracker, click the
Add button in the Control Panel

Multiple patterns are useful when stabilizing and corner or perspective positioning. They also help to keep the flow from becoming cluttered by collecting into a single tool what would otherwise require several tools.

If you need to delete a pattern, select the pattern in the Tracker list and click the Delete button.

Enabling/Suspending/Disabling Patterns

To the left of each pattern's name in the Tracker list is a checkbox that can be used to set the pattern into one of three states. The checkbox is 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 tools and the data is available for stabilization and corner positioning.

An Enabled checkbox

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 tools. The data is also available for advanced tracking modes like stabilization and corner positioning.

A Suspended checkbox

Disabled

A Disabled tracker does not create a path when tracking is performed. Its data is not available to other tools or for advanced tracking operations like stabilization and corner positioning.

A Disabled checkbox

Pattern Selection

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 tool controls.

Channel bars indicate which channel is selected for tracking

You can override the automatic channel selection by clicking the buttons beneath the bars for each channel to determine the channel used for tracking.

You can choose any one of the color channels, the luminance channels or the alpha channel to track a pattern.

When choosing a channel, the goal is to choose the cleanest, highest contrast channel for use in the track. Channels that contain large amounts of grain or noise should be avoided. Bright objects against dark backgrounds often track best using the luminance channel.

Selecting Patterns for Stabilization

Selecting patterns for stabilization can be a tricky business. The location of the pattern, when it is selected, is used to determine precisely how the image will be stabilized. At least two patterns are required to correct for rotation; using three patterns will correct for scaling and more will usually improve the quality of the solution.

Try not to select just any potentially valid pattern in the sequence, as some patterns will make the solution worse rather than better. To help with your selection, use the following guidelines when selecting patterns for stabilization.

→ Locate patterns at the same relative depth in the image. Objects further in the background will move in greater amounts compared to objects in the foreground due to perspective distortion. This can confuse the stabilization calculations, which do not compensate for depth.

→ Locate patterns that are fixed in position relative to each other. Patterns should not be capable of moving with reference to each other. The four corners of a sign would be excellent candidates, while the faces of two different people in the scene would be extremely poor choices for patterns.

Tracking the Pattern

Once the patterns are selected, you can begin tracking to create paths that represent their position for each frame. Each pattern enabled in the Track list will produce its own path.

Before initiating the Tracker, you'll need to set a render range that is usually a range of frames during which the pattern is visible in the frame. This maybe the entire clip of a small portion.

To learn how to set the Render Range, read Chapter 6, The Time Ruler.

Once the render range is set, you can use on any of the Tracking Transport buttons at the top of the Control Panel to start tracking. Once tracking has started you cannot work in the flow until it has completed.

The Tracking Transport buttons

To begin tracking, do one of the following:

> → Click the Track Reverse button to track from the end of the render range.

> → Click Track Backward from Current Frame, to track backward from the current playhead position.

> → Click the Track Forward button to track from the start of the render range.

> → Click Track Forward from Current Frame, to track forward from the current playhead position.

Pattern tracking will stop automatically when it reaches the end of the render range (or the start when tracking backward), but you can also interrupt it and stop tracking at anytime.

To stop tracking, do one of the following:

> → Click the Stop Tracking button in the Tracker transports.

> → Click Stop Render at the bottom of the Fusion window.

> → Press the Escape key.

When tracking is complete, the path will be connected to the pattern. The path from that pattern can now be connected to another tool or used for more advanced operations like stabilization and corner positioning.

Using The Pattern Flipbooks

Each pattern has a pair of thumbnail windows shown in the Control Panel. The left window shows the pattern that is selected.
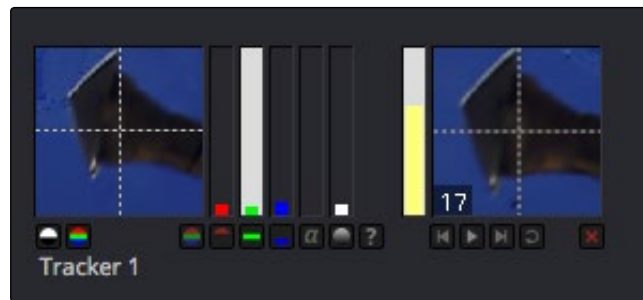
The Tracker Pattern Selection and Flipbook thumbnails

The right window is updated during the track to show the actual pattern acquired on each frame. Each pattern adds these images to a Flipbook and once the render is complete you can play back the Thumbnail Pattern flipbook to evaluate the accuracy of the tracked path.

Adaptive Pattern Tracking

Even the most ideal pattern will usually undergo shifts in profile, lighting conditions and other variables. These can adversely affect pattern recognition to the point that a pattern becomes unusable. The Tracker offers three modes of pattern acquisition during tracking that can help to correct these conditions. The modes can be set using the Adaptive Mode buttons in the Control Panel.

The Adaptive Mode buttons

None

When the Adaptive mode is set to None, the pattern within the rectangle is acquired when the pattern is selected, and that becomes the only pattern used during the track.

Every Frame

When Every Frame is chosen, the pattern within the rectangle is acquired when the pattern is selected, then re-acquired at each frame. The pattern found at frame 1 is used in the search on frame 2, the pattern found on frame 2 is used to search frame 3, and so on. This method helps the Tracker adapt to changing conditions in the pattern.

Every Frame tracking is slower and can be prone to drifting from sub-pixel shifts in the pattern from frame to frame. Its use is therefore not recommended unless other methods fail.

Best Match Tracking

Best Match tracking works in much the same way as Every Frame tracking, however, it will not re-acquire the pattern if the difference between the original pattern and the new one is too great. This helps to prevent cases where transient changes in the image cause the Tracker to become confused.

As a comparison between the two Adaptive modes, if a shadow passes over the tracker point, the Every Frame tracking mode may start tracking the shadow instead of the desired pattern. The Best Match mode would detect that the change from the previous frame's pattern was too extreme and would not grab a new pattern from that frame.

The Adaptive mode is applied to all active patterns while tracking. If you only want some patterns to use the Adaptive mode, disable all other patterns in the list before tracking.

Dealing With Obscured Patterns

Often, an otherwise ideal pattern can be temporarily obscured 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. You can then set the track path to Insert and Modify mode and insert points in the non-tracked range to compensate for any non-linear motion in the tracked pattern.

The Insert and Modify button can be used to insert new points on the path

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 button

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, change the playhead to the last frame that tracked successfully.

2. Click the Track Center (Append) button in the Control Panel.

3. Reposition the Pattern selector to the new pattern.

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 perspective and axial rotation will disturb the accuracy of the tracked result.

Tracker Offsets

Often, the pattern selected for tracking is not exactly the same as the pattern required for the track. For example, the pattern to be tracked is a button on an actor's sleeve but 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 would be necessary to make use of the Tracker Offset control.
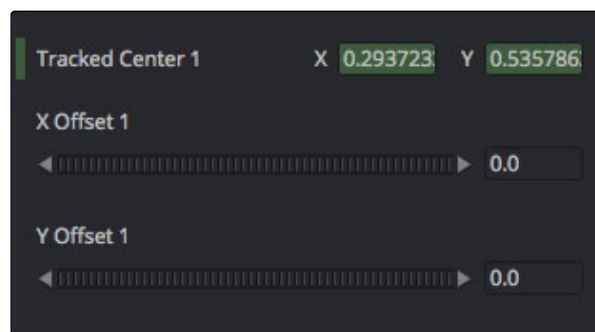
The Tracker Offset controls in the Control Panel

The tracker offset allows for constant or animated positional offsets to be created relative to the actual tracker's pattern center. You can position an offset from the pattern's tracked position using the X Offset and Y Offset controls in the Tracker's Pattern controls.

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 enabled you to reposition the path while keeping the tracker pattern in place.

The Tracker Offset button

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.

Match Move Operations

When the Tracker is set to Match Move in the Operations tab, it is capable of stabilizing footage to completely remove motion from the scene or smooth existing motion. It can also apply the motion from one clip to another, basically matching the movement and stabilizing one shot with reference to another. Here are some common scenarios for stabilization that are handled when the tracker is set to Match Move.

    → A sequence that should be steady has vibrations or undesirable movement.

    → A sequence that requires a smooth camera move suffers from jarring.

    → A static CG element must be believably added to a moving sequence.

    → Two sequences with different motions must be composited together.

The Tracker Operation tab Match Move button

Stabilizing Motion

Stabilizing Motion completely removes the appearance of motion from the image. The motion from frame to frame is calculated, and the contents of the frame are transformed to return the image to a reference position. This position can be either the start or end of the sequence or a manually selected frame from the sequence.

Stabilization can correct for position with as little as one pattern. Two or more patterns are required to correct for rotation or scaling within the image.

When the Match Move button is selected in the Tracker's Operation tab, the Tracker can use the data from its patterns for stabilization. Only the controls that are applicable for stabilization operations will appear in the Operation tab.

Several of the stabilization controls are always available, collected under the Match Move Settings disclosure button. These controls are available at all times because the Steady and Unsteady positions of a tracker are always published. This makes them available for connection by other controls, even when the Tracker's operation is not set to match moving.
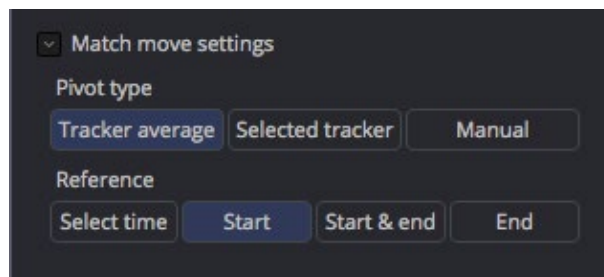
The Match Move settings

Merge

The Merge buttons determine to what input knot on the Tracker tool the stabilization data is applied. When stabilizing an image to remove all motion, or smooth the motion, the Merge button should be set to BG Only.

Edges

The Edges buttons determine whether the edges of an image that leave the visible frame are cropped, duplicated or wrapped when the stabilization is applied. Wrapping edges is often desirable for some methods of match moving, although rarely when stabilizing the image for any other purpose. These controls are described in greater detail in the Tool References for the Tracker later in this manual.

Position/Rotation/Scaling

Use the Position, Rotation and Scaling checkboxes to select what aspects of the motion are corrected.

Pivot Type

The Pivot Type for the stabilization is used to calculate the axis of rotation and scaling calculations. This is usually the average of the combined pattern centers but may be changed to the position of a single tracker or a manually selected position.

Reference

The Reference controls establish whether the image is stabilized to the first frame in the sequence, the last frame or to a manually selected frame. Any deviation from this reference by the tracked patterns is transformed back to this ideal frame.

As a general rule, when tracking to remove all motion from a clip, set the Merge mode to BG Only, the Pivot type to Tracker Average or Selected Tracker and the Reference control to Start, End or Select Time.

Smoothing Motion

When confronted with an image sequence with erratic or jerky camera motion, instead of trying to remove all movement from the shot, you often need to preserve the original camera movement while losing the erratic motion.

The Start & End reference option is designed for this technique. Instead of stabilizing to a reference frame, the tracked path is simplified. The position of each pattern is evaluated from the start of the path and the end of the path along with intervening points. The result is smooth motion that replaces the existing unsteady move.

The Reference Intermediate Points slider is
displayed when Start & End is selected

To 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.

Matching Movement in Clips

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.

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 tools, such as transform or merge, to a tracker's outputs. The other method is to stabilize an image by trying to remove all motion but instead of setting the Merge buttons to BG Only, set it to FG Over BG, FG Only, or in rare occasions, BG Over FG.

Set the Merge buttons to BG Only, FG Over BG or BG Over FG

When using this Merge buttons, you connect a foreground image to the Tracker tool's Input knot in the flow.
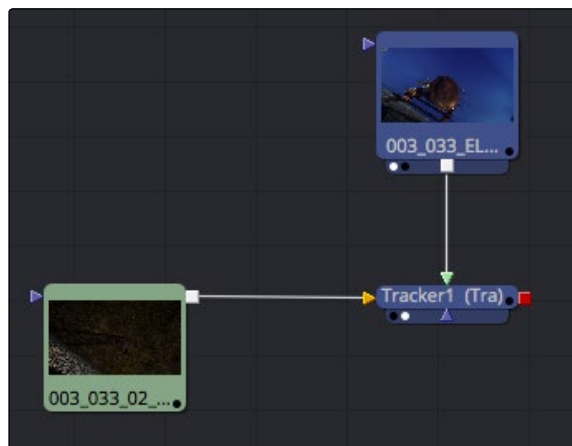
Connect a foreground image to the Tracker's foreground input

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 tools 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.
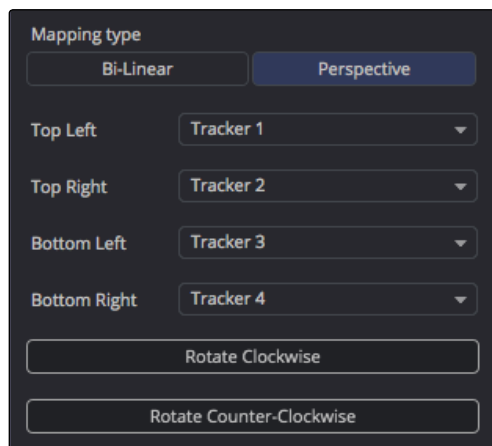
The Corner Positioning operation in the Control Panel

The Corner Positioning operation of the Tracker requires the presence of a minimum of four patterns. If this operation mode is selected and there are not four patterns set up in the Tracker already, additional patterns will automatically be added to bring the total up to four.

When this mode is enabled, a set of drop down boxes will appear to select which tracker relates to each corner of the rectangle. It has no effect when the Merge control buttons are set to BG Only.

Perspective Positioning Operations

The Perspective Positioning operation is used to remove perspective from a foreground image or apply the perspective from one sequence to another.

The Perspective Positioning operation in the Control panel

The Perspective Positioning operation of the Tracker requires the presence of a minimum of four patterns. If this operation mode is selected and there are not four patterns set up in the Tracker already, additional patterns will automatically be added to bring the total up to four.

When this mode is enabled, a set of drop down boxes will appear to select which tracker relates to each corner of the rectangle. It has no effect when the Merge control buttons are set to BG Only.

Tracker Outputs

One of the most common applications for a tracked pattern is using the tracked position or path to drive the position of another tool's parameters. For example, tracking an eye in order to color correct the eye to blue using an effect mask. You start off by tracking the eye, then create a color corrector with the desired settings. You create a mask in the shape of the eye and connect the tracker's position to the Center of the mask.

In addition to the path, each pattern in a tracker outputs five values for use as connections that are available for use by other tools.

You connect a tool's position parameters to a tracker by selecting the connection type from the controls contextual menu (for example, Transform 1: Center > Connect To > Tracker 1 > Unsteady Position).

There are five connection types available to connect to a position parameter in another tool.
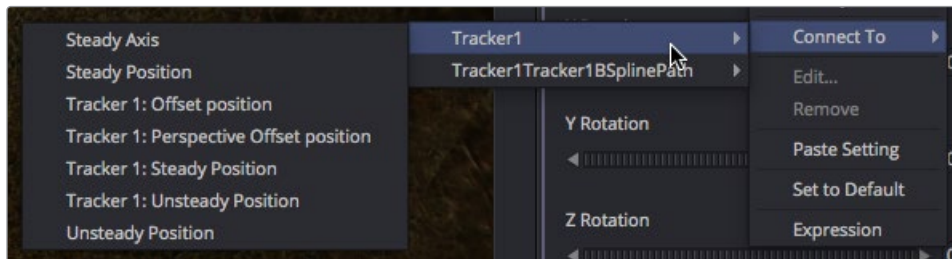
The Connect To submenu

Steady Position

Steady Position can be used to stabilize footage in both X and/or Y to remove camera shake and other unwanted movement. The connection inverts the output of the tracked pattern's motion. When you connect a Center parameter to the Steady Position of the tracker, it will be placed at 0.5/0.5 (the center of the screen) by default at frame 1. You can change this using the Reference mode in the Tracker's Operation tab.

Steady Angle

The Steady Angle mode can be used to stabilize footage in both X and/or Y to remove camera shake and other unwanted movement. When you connect a control, for example the Angle of a Transform, to the Steady Angle of the Tracker, it will be placed at 0 degrees by default at frame 1. This can be changed by means of the Reference mode in the Tracker's Operation tab. From there on, the resulting motion of the Steady Angle mode will rotate into the opposite direction of the original motion.

So if the actual angle at frame 10 is 15 degrees, the result of the Steady Angle will be -15 degrees.

To use Steady Angle you need at least two tracked patterns in your tracker. With just one point you can only apply (Un)Steady Position.

Offset Position

An Offset Position is available for each single tracker in the Tracker node and refers to that single tracker only. When you connect the Center X and Y parameters to the Offset Position of the tracker, the tool'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.
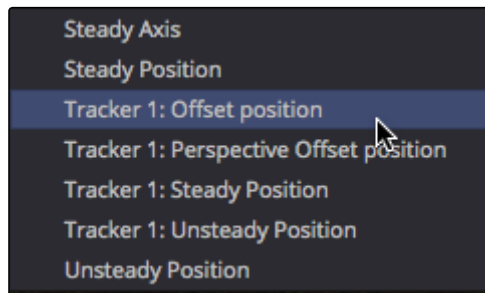
The Offset Position connection

Unsteady Position

The Unsteady Position is used to re-introduce the original movement on an image after an effect or new layer has been added. The resulting motion from Unsteady Position is basically an offset in the same direction as the original motion.

Steady Size

The Steady Size connection outputs the inverse of the tracked pattern's scale. When you connect a parameter, for example the Size of a Transform, to the Steady Size of the Tracker, it will be placed with a Size of 1 (i.e., the original size) by default at frame 1. This can be changed by means of the Reference mode in the Tracker's Operation tab. The resulting size of the Steady Size mode will then counteract the size changes of the original motion. So if the actual size at frame 10 is 1.15, the result of the Steady Size will be 1-(1.15-1)=0.85.

To use Steady Size you need at least two tracked patterns in your tracker. With just one point you can only apply (Un)Steady Position.

The connections above are output by pattern in the Tracker tool. Each tool itself also outputs a Steady Position, Angle and Size output. The values of these outputs are calculated using all of the patterns in that tracker, as configured by the Stabilize Settings controls in the Tracker's Operation tab.

As an example, to use the Connect To menu to perform a Match Move, do the following:

1. Add a Transform tool to the clip you want to Match Move.

2. Right click over the Transform's Center and choose Connect to > Steady Position.

3. Set the Transform tool's Edges mode to Mirror so that pixels will not get cropped from the image when it is stabilized.

4. Add another Transform tool to the flow after the merge.

5. Connect the new Transform's Center to the tracker's Unsteady Position.

The image will be restored to its original state with the additional effect included.

To better understand how this works, imagine a pattern that is selected at frame 1, at position 0.5, 0.5. The pattern does not move on frame 2, so its position is still 0.5, 0.5. On the third frame, it moves 10 percent of the image's width to the right. Now its position is 0.6, 0.5.

If a transform center is connected to the Steady Position output of the tracker, the Transform tool'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.

Tracker Modifier

Another technique for adding a tracker directly to a control is to add it as a modifier.
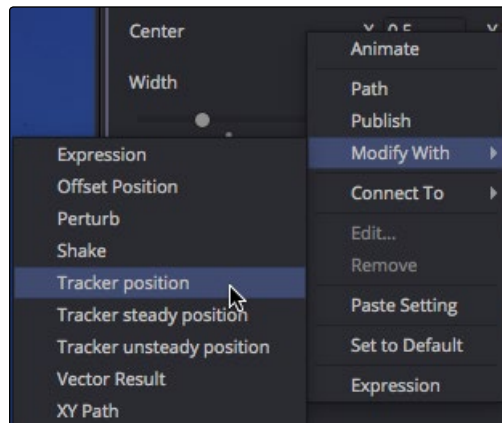
Choosing a Tracker modifier

The differences between a Tracker modifier and a Tracker tool are as follows:

→ The Tracker modifier can only track a single pattern.

→ A source image must be set for the Tracker modifier.

The Tracker modifier can only output a single value and cannot be used for complex stabilization procedures.

As an example, to apply the Tracker as a modifier, do the following:

Imagine that you needed to track an actor's eyes so that an unearthly, alien glow could be applied to the eyes.

1. Add a Glow tool.

2. Create an ellipse mask in the shape of the eye.

3. In the Control Panel, right click on the Mask's Center parameter.

4. From the contextual menu choose Modify With > Tracker Position.
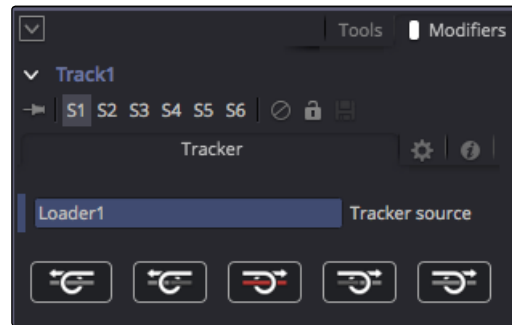
Tracker Modifier controls in the Modifier tab of the Control Panel

Choosing the Tracker from the Modify with Contextual menu adds a modifier in the Control Panel with a set of parameters almost identical to those found in the Tracker tool itself.

The default source image for the modifier is the tool immediately upstream of the tool that contains the modifier (i.e., when adding a Tracker modifier to a Glow tool with a Loader as its input, the Tracker Source input will default to the output of the Loader). You can set a different source image for the Tracker modifier by typing in the name of the tool at the top of the Control Panel Modifier tab. Alternately, drag and drop the source tool from the flow into the Text Box control or use Connect To from the contextual menu.

For detailed information on the Tracking controls please refer to the Tracking Tools section in the Tool Reference.

Working in 3D

12

Working in 3D

3D Compositing Overview

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 has been created in a true 3D modeling and rendering application.

Within Fusion's Flow Node Editor you have a complete OpenGL-accelerated 3D compositing environment, including support for geometry, point clouds and particles. This makes tasks such as the ones listed below much easier to perform.

→ Convert 2D images into image planes in 3D space.

→ Import matched cameras and point clouds from 3D matchmoving software, such as SynthEyes and PF Track.

→ Import camera, lights and materials from popular 3D applications, such as Maya, LightWave and 3Ds Max.

→ Create rough primitive geometry or import mesh geometry from FBX or Alembic scenes.

→ Create realistic surfaces using illumination models and shader compositing.

→ Render with realistic depth of field, motion blur and supersampling.

→ Use 3D particle systems.

→ Extrude and bevel text.

→ Cast shadows across geometry.

Example 3D scenes

This chapter will cover some of the tools for 3D compositing, the tasks they perform and how they can be combined to produce effective 3D composites.

The 3D Scene

Tools that work in a 3D compositing environment create a 3D scene. Examples are tools that generate primitive geometry, import geometry, and create lights and cameras, as well as tools that warp, deform or duplicate geometry. Almost all of these tools are collected together into the 3D category of tools.

When you create a new composition, you don't have to specify whether it will be 2D or 3D, because you can seamlessly combine any number of 2D and 3D scenes together. 3D geometry is added to the scene using the Image Plane 3D tool and Shape tools or imported via the FBX Mesh tool. The geometry is combined with lights and cameras using the Merge 3D tool, and rendered into 2D images using the Renderer 3D tool. The look of the geometry is defined by material tools, such as the Blinn and Phong tools, and textured with 2D images or environment maps.

When these tools are displayed , a 3D Viewer is used, allowing you to navigate and manipulate the scene in 3D. The interactive 3D Viewer is highly dependent on the computer's graphics hardware, relying on support for OpenGL. The amount of onboard memory as well as the speed and features of the Graphics Processing Unit (GPU) make a huge difference in the speed and capabilities of the 3D Viewer.

Tools that output 3D scenes cannot be connected directly to inputs that expect 2D images. For example, the output of an Image Plane 3D cannot be connected directly to the input of a Blur. A Renderer 3D must first be used to create a 2D image from the scene for the Blur tool to operate.
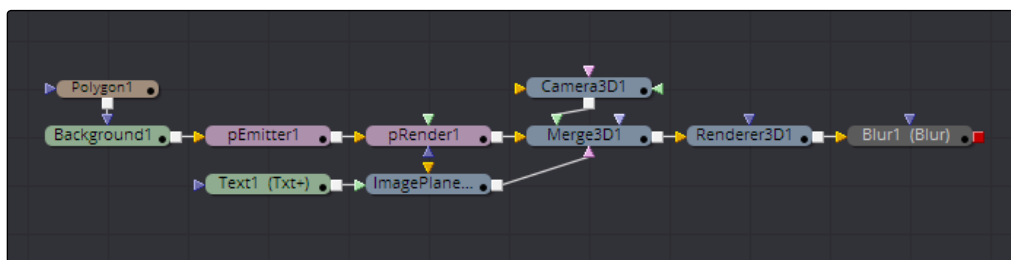
2D tools are gray, 3D tools are light blue, and Particle tools are purple

Merge 3D

Each 3D tool outputs a complete 3D scene. This is unlike most traditional 3D modeling and animation programs, where all objects reside within a global scene environment. So the scenes created by a Camera 3D tool and an image plane are separate until they are somehow combined into the same scene.

The Merge 3D takes the output of 3D tools and combines them into a single scene. Unlike the 2D Merge tool, the ordering of the elements in the scene is not restricted to just the background and foreground inputs. Instead, the Merge 3D allows for an unlimited number of inputs and combines them according to their absolute positions in 3D space.

The Merge 3D presents very few options in the Control Panel. The main Control Panel contains a single checkbox that you can enable to pass lights through to the output. Any lights present in the scene are passed onto next connected tool. This checkbox is disabled by default allowing you to light elements in the scene without worrying about how the light might affect geometry introduced further downstream. For example, you may want to apply a spotlight to brighten the wall of a building without having that spotlight spill over onto the grass or pavement at the foot of the wall.

Additionally, the Merge 3D includes a common 3D Transformation tab. These transformation parameters can be used to adjust the position, scale and rotation of all the elements combined by the merge. This applies not just to geometry but also to lighting and particles. All transformations take place around a common pivot point. This forms the basis of parenting in the 3D environment.
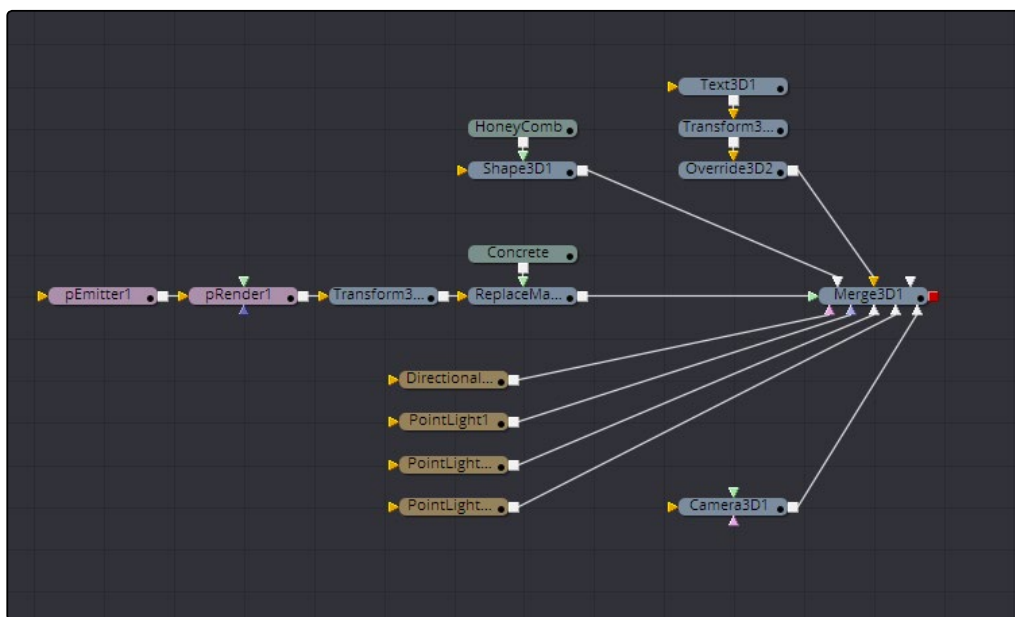
3D scene built around a Merge 3D

Tip: Most 3D primitives also provide a 3D scene input that can be used as an alternative to the Merge 3D.

Rendering 3D Scenes

A 3D scene must go through a Renderer 3D before it can be rendered into 2D images. In the Renderer 3D there are three different types of renderers by default: a GPU-accelerated OpenGL renderer which produces results almost identical to those shown in the interactive 3D views, a software renderer which is sometimes slower but supports advanced rendering techniques such as soft shadows or extremely large textures, and an OpenGL UV renderer which is used for baking textures or shading.

Renderer 3D

The Renderer 3D is placed at the end of your 3D scene.

Output of a Merge 3D going into a Renderer 3D

The Renderer 3D uses one of the cameras in the scene to produce an image. If no camera is found, a default perspective view is used. This default view rarely provides a useful angle, so most 3D scenes include at least one camera.

The image produced by the Renderer 3D can be any resolution with options for fields processing, color depth, and pixel aspect.

Software Renderer

The Software Renderer is the default renderer and is generally used to produce the final output. The software renderer is not dependent on the capabilities of the graphics hardware, ensuring that the results of the render are always the same whatever the configuration of the machine used to render.

The software renderer can easily handle textures much larger than one half the graphics hardware's maximum texture size. If you are working with images larger than 8K as textures you should choose the software renderer to get maximum quality. It's also recommended to use the software renderer if your render farm has weak graphics hardware, or no graphics hardware at all.The software renderer is required to produce soft shadows. This feature is not supported by the OpenGL renderer.
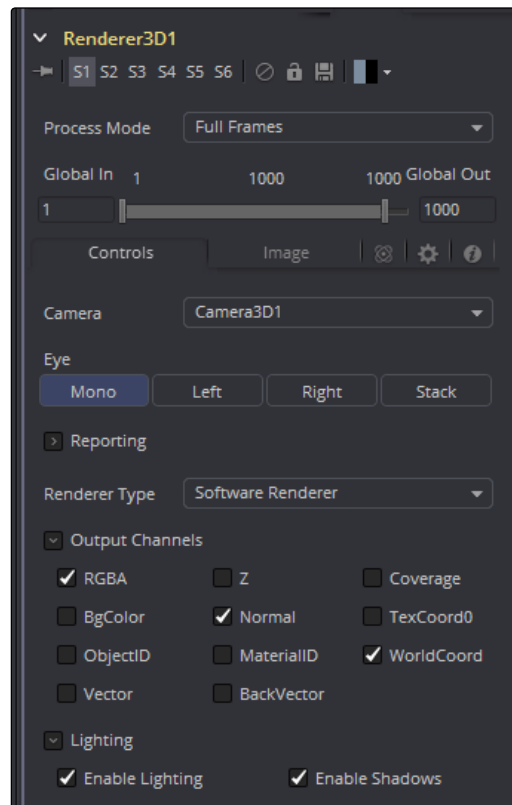
Software Renderer options in the Control Panel

OpenGL Renderer

The OpenGL Renderer takes advantage of the GPU in your computer to produce the image. The textures and geometry are uploaded to the graphics hardware, and OpenGL shaders are used to produce the image. This can produce high quality images perfect for final rendering and potentially orders of magnitude faster than the software renderer, but it does pose some limits on some rendering effects.
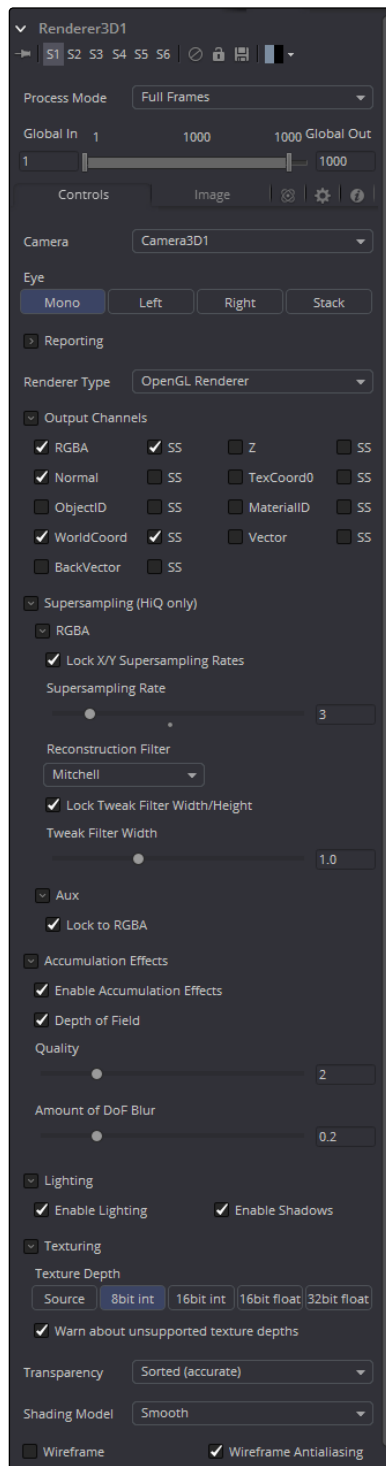
Open GL Render options in the Control Panel

OpenGL UV Renderer

The OpenGL UV Renderer is used for baking out texture projections or materials to a texture map. Baking out projections can speed up a render, but it also allow you to modify the texture with other 2D tools before applying it back onto the mesh.

For instance, you have a scene on a street corner and there's a shop sign with a phone number on it, but you want to change the numbers. If you track the scene and have standing geometry for the sign, you can project the footage onto it, do a UV render, switch the numbers around with a Paint tool, and then apply that back to the mesh with a Texture2D.

The UV renderer can also be used for retouching textures. You can combine multiple DSLR still shots of a location, project all of those onto the mesh, UV render it out, then retouch the seams and apply it back to the mesh.

You could project tracked footage of a road with cars on it, UV render out the projection from the geometry, do a temporal median filter on the frames, then map a "clean" roadway back down.

3D Viewer

Displaying a tool 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 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.
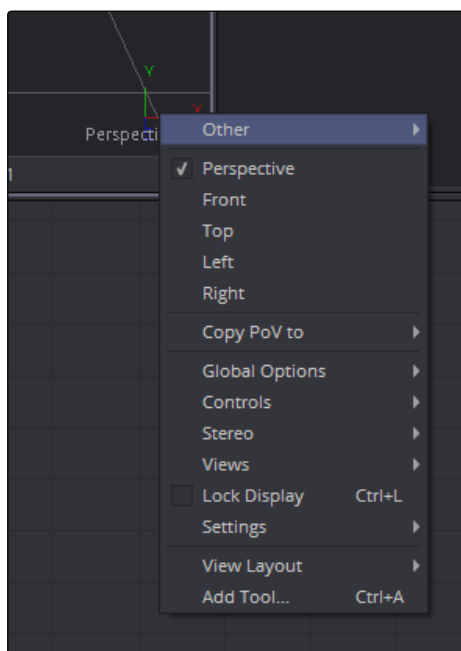
Right click over the Axis label to change the viewpoint

In addition to the usual Perspective, Front, Top, Left and Right viewpoints, the Camera submenu also includes any cameras and lights present in the scene as potential viewpoints. It is 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.

For a detailed description of the options available in the 3D Viewer, read Chapter 5, Viewers.

Navigating the 3D View

For the most part, panning and scaling of the 3D Viewer uses the same controls as the 2D Viewer.

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 the middle and left mouse buttons and drag left or right in the Viewer.

→ Hold down Command (Mac OS X) or Ctrl (Windows) and use the scroll wheel.

To rotate around the 3D Viewer, do the following:

→ Hold Option (Mac OS X) or Alt (Windows) and drag left and right in the Viewer, using the middle mouse button.

If you want to frame certain objects in the Viewer, do one of the following:

→ Press A to Fit all objects in the Viewer.

→ Press F to Fit to selection (or Fit All if nothing is selected).

→ Press C to Rotate the Viewer to look at the center of the currently selected object without moving the Viewer's position.

Furthermore, selecting a 3D tool in the flow area also selects the associated object in the 3D Viewer.

Transforming With the Viewers

When the Viewer is set to look through an object in the scene, like a camera or spotlight, the usual controls for panning and rotating the Viewer will now directly affect the position of the object.

For example, to adjust the camera position when looking through the camera, do the following:

1. Place the cursor in the Viewer.

2. Hold the middle and left mouse buttons and drag to change the camera.

It is even possible to view the scene from the perspective of a Merge3D or Transform3D tool by selecting the object from the Camera > Others menu. The same transformation techniques will then move the position of the object. This can be helpful when you are trying to orient an object in a certain direction.

Transparency Sorting

While generally the order of geometry in a 3D scene is determined by the Z-position of each object, sorting every face of every object in a large scene can take an enormous amount of time. To provide the best possible performance, a Fast Sorting mode is used in the OpenGL renderer and Viewers. This is set by right clicking in the Viewer and choosing Transparency > Z-buffer or enabling the Fast button in the Viewer toolbar. While this approach is much faster than a full sort, when objects in the scene are partially transparent it can also produce incorrect results.

The Sorted (Accurate) mode can be used to perform a more accurate sort at the expense of performance. This mode is selected from the Transparency menu of the Viewer's contextual menu. The Renderer 3D also presents a Transparency menu when the Renderer Type is set to OpenGL. Sorted mode does not support shadows in OpenGL. The software renderer will always use the Sorted (Accurate) method
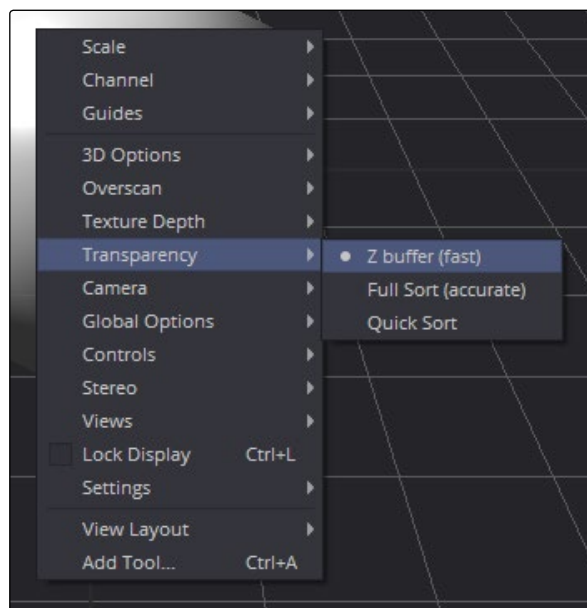
Transparency Sorting in the Viewer contextual menu

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 tool that comes from the 3D > Material category of tools, the Viewer displays a Material Viewer. This Material Viewer allows you to preview the material applied to a lit 3D sphere rendered with OpenGL by default. The type of geometry, the renderer, and the state of the lighting can all be adjusted from either the Viewer's toolbar, or the contextual menu. Each Viewer includes an A and B buffer to assist with comparing multiple materials.

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, but the view can be rotated to provide a different angle on the material, by holding Option (Mac OS X) or Alt (Windows), pressing the middle mouse button and dragging left and right.

Using the middle mouse button to drag in the Viewer will adjust the position of the light. Light can also be toggled on and off using the display view's toolbar button.

Material Viewer

Transformations

Merge3D, 3D Objects and Transform3D all have Transformation parameters that are collected together into a Transformation tab in the Control Panel. The parameters found in this tab affect how the object is positioned, rotated and scaled within the scene.
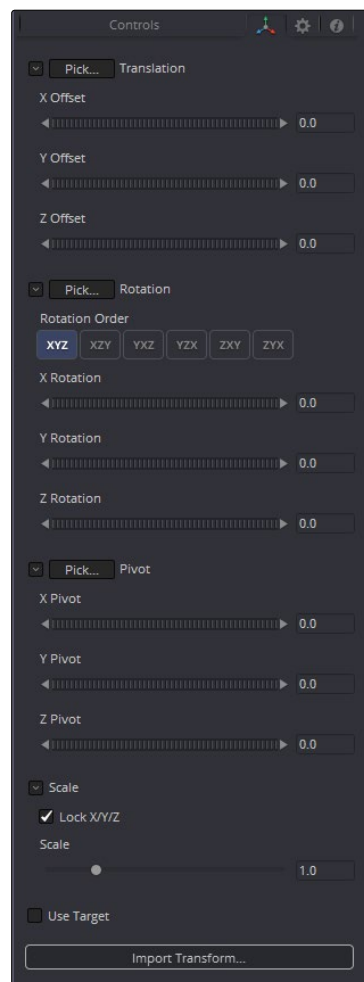
Transformation tab

The Offset 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 sliders affect its size. The same adjustments can be made in the Viewer using the transformation on screen controls.

Onscreen Transformation Controls

When an object is selected, it displays onscreen Transformation controls in the Viewers that allow you to adjust the object's position, rotation and scale. There are buttons in the Transform toolbar so you can switch the modes or use the keyboard short cuts.

To switch transformation modes, use the following keyboard short cuts:

→ Press Q for Position

→ Press W for Rotation

→ Press E for Scaling

Transformation toolbar

In all three modes, the red portion of the onscreen control indicates the object's local X-axis, while green and blue are used for the Y- and Z-axis respectively. You can drag directly on the red, green or blue portion of the onscreen control to constrain the transformation to that axis, or if you drag in the center of the onscreen control you can apply a transformation without constraint. Holding Option (Mac OS X) or Alt (Windows) and dragging in the Viewer allows you to freely translate in all three axes without clicking on a specific control.
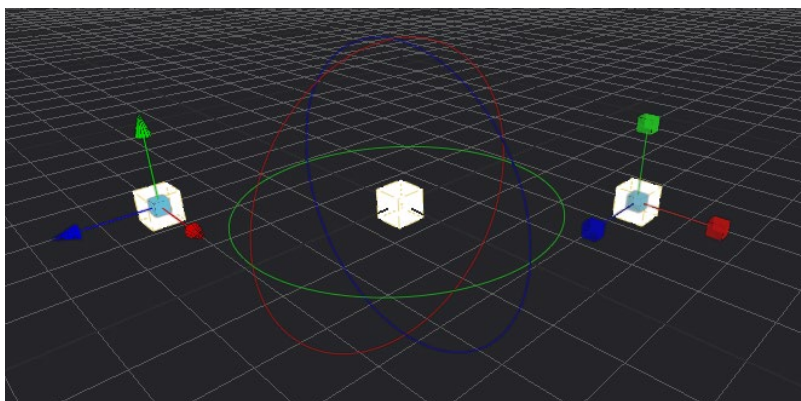
Transformation onscreen control

If the Scale's Lock XYZ checkbox is enabled in the Control Panel, only the overall scale of the object is adjusted by dragging the red or center onscreen control while the green and blue portions of the onscreen control have no affect. If you unlock the parameters you are able to scale an object along individual axes separately.

Pivot

In 3D, 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 Control Panel.
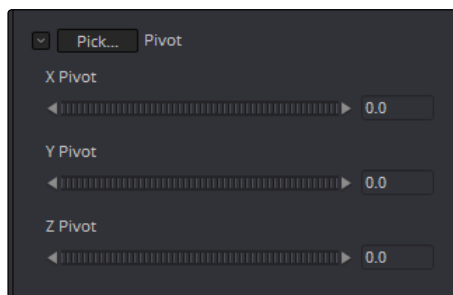
Pivot parameters in the Control Panel

Target

Targets are used to help orient one object to another or to a specific point in the scene. No matter where the object moves, it will rotate in the local coordinate system to face its target.

Enabling the Use Target checkbox will reveal another set of XYZ parameters in the Control panel and new onscreen controls for target position. A line is drawn between the target and object center in the Viewer to show the relationship between these controls.
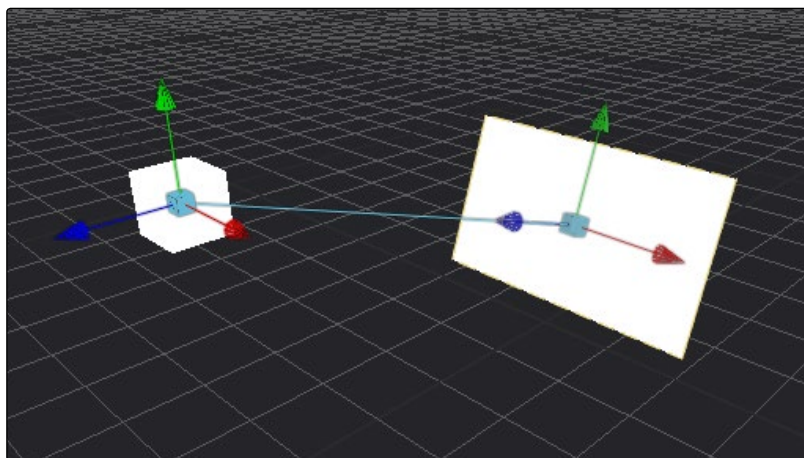
Image facing the Target onscreen controls

For example, if a spotlight is required in the scene to point at an image plane, enable the spotlight's target in the Transformation tab and connect the target's XYZ position to the image plane's XYZ position. No matter where the spotlight is moved, it will rotate to face the image plane.

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 the 3D flow. The basis for all parenting is the Merge 3D. For example, picture a scene with two spheres that are both connected to a Merge 3D. The Merge 3D can be used to rotate one sphere around the other, like the moon around the earth. Then the Merge 3D can be connected to another Merge 3D to create the earth and the moon orbiting around the sun.

One Merge 3D with two spheres parented to another Merge 3D and parenting using three connected spheres

Any transformation applied to a Merge 3D is applied to all objects connected to that merge; that includes cameras, lights, geometry and other Merge tools.

Cameras

A frequent task for compositors involves matching the cameras from live action clips or cameras from 3D applications. To accommodate this task, Fusion provides a flexible Camera 3D where you can adjust common camera controls like the Angle of View, Focal Length, Aperture and Clipping planes to set up your own camera or import camera from other applications. The Camera 3D is a virtual camera through which the 3D environment can be viewed. When adding a camera you connect it and view it through a Merge 3D.
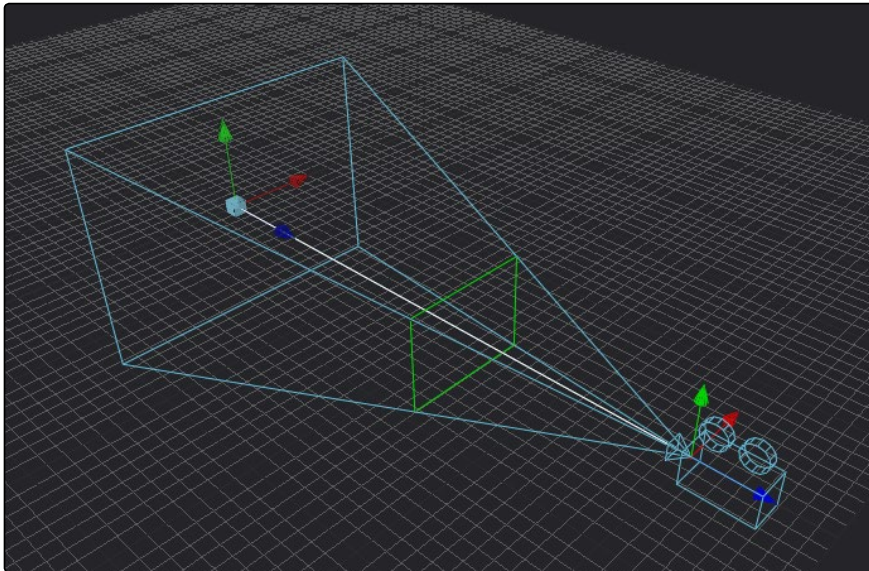
A camera displayed with on screen Transform controls in the Viewer

To view the scene through the camera, do the following:

1. Select the Merge 3D that the camera is connected to, or any tool downstream of that Merge 3D.

2. Load the selected Merge 3D or downstream tool 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 rendered by the Renderer 3D. If there is no Renderer 3D added to your scene yet, you can use Guides that represent the camera's framing. To learn more about Guides, read Chapter 5, Viewers.

Additionally, the Plane of Focus, for depth of field rendering is available in the Control Panel for the camera. To render the depth of field you must add a Renderer 3D and enable the Accumulation Effects in the OpenGL render.

Plane of Focus controls in the Camera Control Panel

To learn more about specific camera controls, read the Camera3D section of the Tool 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 Flow Node Editor.

2. At the bottom of the Control Panel, click the Import Camera button.

3. In the file browser, navigate to and select the scene that contains the camera you want to import.

A dialog box with several options will appear. When the Force Sampling checkbox is enabled, Fusion will sample each frame of the motion, regardless of the format.
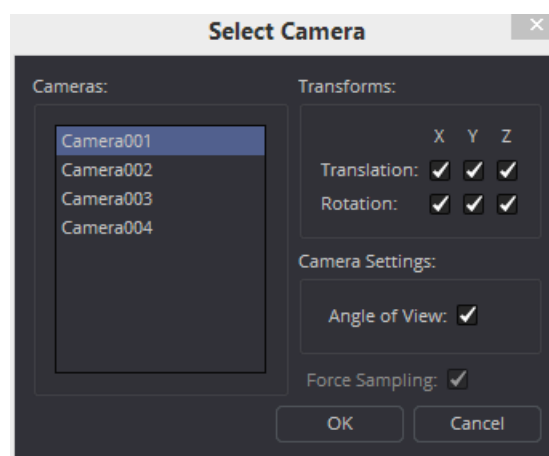
The Import Camera Control dialog

Tip: When importing parented or rigged cameras, baking the camera animation in the 3D application before importing it into Fusion often produces the more reliable results.

Lighting and Shadows

You can add light sources to a scene to form very detailed lighting environments and create atmosphere. There are four different types of lights you can use in 3D scenes: ambient, directional, point and spotlights. A scene without lights uses a default directional light which disappears once you add a light. However, lights will not be visible in the Viewer unless Lighting is enabled in the Viewer toolbar. Light will also not be visible in the Renderer3D until the Enable Lighting checkbox is checked in the Control Panel. Tools that create or load geometry also include lighting options which are used to choose whether the object is affected by light or not.

The Light button under the Viewer

The Light button in the Viewer toolbar can be used to toggle lights off and on. When lighting is disabled in either the Viewer or final renders, the image will appear to be lit by a 100% ambient light.

Ambient Light

You use Ambient light to set a base light level for the scene since it produces a general uniform illumination of the scene. Ambient light exists everywhere without appearing to come from any particular source; it cannot cast shadows and will tend to fill in shadowed areas of a scene.

Directional Light

A Directional light is composed of parallel rays that light up the entire scene from one direction, creating a wall of light. The sun is an excellent example of a directional light source.

Point Light

A Point light is a well defined light that has a small clear source, like a light bulb, and shines from that point in all directions.

Spot Light

A Spot light is an advanced point light that produces a well defined cone of light with falloff. This is the only light that produces shadows.

From L to R: Directional Light, Point Light and Spot Light

All of the Light tools display on-screen controls in the Viewer, although not all controls affect every light type. In the case of the ambient light, the position has no effect on the results. The directional light can be rotated, but position and scale will be ignored. The point light ignores rotation. Both position and rotation apply to the spotlight.

Light Hierarchies

Lights normally do not pass through a Merge. This provides a mechanism for controlling which objects are lit by which lights. For example, in the flow below, two cubes and an ambient light are combined into a merge, which is then connected to a merge with an image plane and a spotlight. The ambient light will affect only the cubes. Its light does not continue downstream to affect the image plane introduced by the next merge. The spotlight will affect the sphere and both cubes, since all three objects are in the scene by this point.
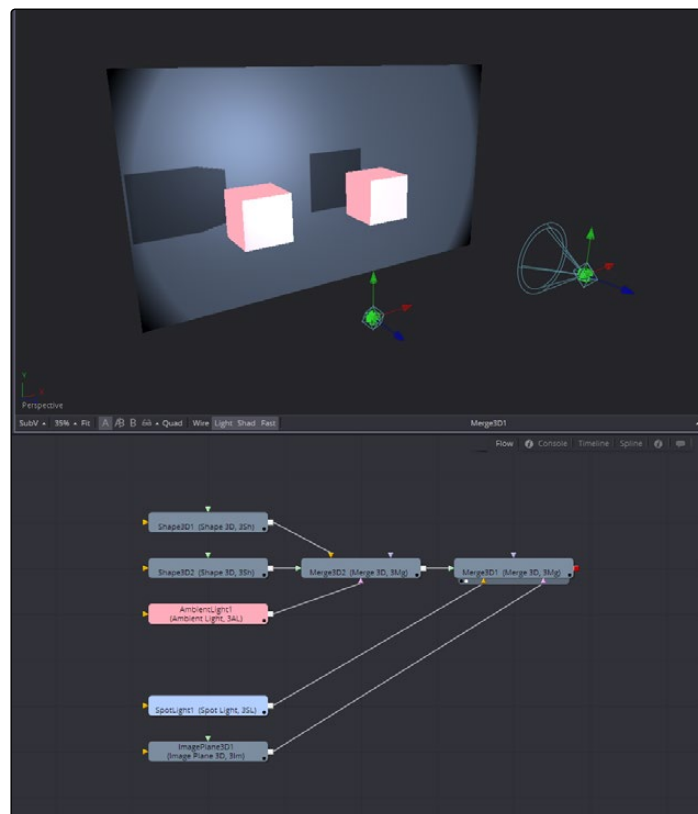
Lights can be confined to the merge they are connected to

To have lights pass through merges to affect downstream objects, click on the Merge's Pass Through Lights option.

Lighting Options

Most tools that generate geometry have additional options for lighting. These lighting options are used to determine how the object reacts to lights and shadows in the scene.

Lighting options in a Shape tool Control Panel

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 tools cast shadows by default, although these shadows will not be visible in the Viewer until shadows are enabled using the Viewer toolbar button. Shadows will not appear in the output of the Renderer 3D unless the Shadows option is enabled for that renderer. If you want to prevent a spotlight from casting shadows, you can disable the Enable Shadows checkbox in the tool's Control Panel.

Image with spotlight casting a variable soft shadow

See Spot Light in the Tool Reference Manual for a more detailed description of the shadow controls.

Shadow Maps

The 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 Control Panel.

The quality of the shadow produced depends greatly on the size of the shadow map. Larger maps generate better looking shadows, but will take longer to render. The wider the cone of the spotlight, or the more falloff in the cone, the larger the shadow map will need to be to produce useful quality results. Setting the value of the Shadow Map Size control sets the size of the depth map in pixels.

Generally, through trial and error, you'll find a point of diminishing returns where increasing the size of the shadow map no longer improves the quality of the shadow. It is not recommended to set the size of the shadow maps any larger than they need to be.

The Shadow Map Proxy control is used to set a percentage by which the shadow map is scaled for fast interactive previews, such as autoproxy and LoQ renders. A value of .4, for example, represents a 40% proxy.

Shadow Softness

By default, the spotlight generates shadows without soft edges but there are options for constant and variable soft shadows. Hard edged shadows will render significantly faster than either of the Soft Shadow options. Shadows without softness will generally appear aliased, unless the shadow map size is large enough. In many cases, softness is used to hide the aliasing rather than increasing the shadow map to preserve memory and avoid exceeding the graphics hardware capabilities.
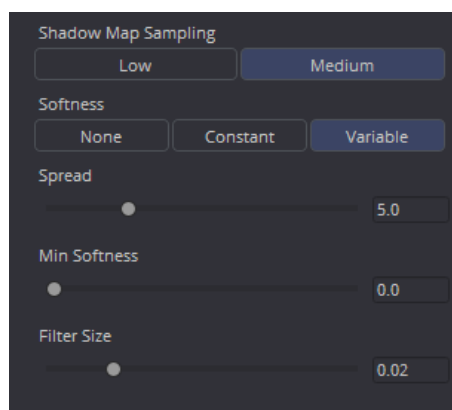
Soft Shadow controls in the Control panel

Setting the spotlight's shadow softness to None will render crisp and well defined shadows. The Constant option will generate shadows where the softness is uniform across the shadow, regardless of the shadow's distance from the casting geometry. The Variable option generates shadows that become softer as they get farther from the geometry that is casting the shadow. This is a more realistic effect, but the shadows are somewhat harder to control. When this option is selected, additional controls for adjusting the falloff of the shadow will appear, as well as sliders for the minimum and maximum softness.

Hard shadow casted by a spotlight

Selecting the Variable option reveals the Spread, Min Softness and Filter Size sliders. A side effect of the method used to produce variable softness shadows is that the size of the blur applied to the shadow map can become effectively infinite as the shadow's distance from the geometry increases. These controls are used to limit the shadow map by clipping the softness calculation to a reasonable limit.

The filter size determines where this limit is applied. Increasing the filter size increases the maximum possible softness of the shadow. Making this smaller can reduce render times but may also limit the softness of the shadow or potentially even clip it. The value is a percentage of the shadow map size.

See Spotlight in the Tool Reference Manual for additional details on these controls.

Multiplicative and Additive Bias

Shadows are essentially textures applied to objects in the scene, occasionally resulting in Z-fighting. Z-fighting results when portions of the object that should be receiving the shadows render over top of the shadow instead because they effectively exist in the same exact location in 3D space.

Top: Results of shadow map Z-fighting Bottom: Corrected shadow shown using Biasing

Biasing works 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 goal is to adjust the Multiplicative Bias until the majority of the Z-fighting is resolved, then adjust the Additive Bias 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 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 Control Panel overrides this, causing the shadow map produced by the tool to ignore transmittance entirely.

Materials and Textures

In order to render a 3D scene, the renderer must take into account the shape of the object as well as its appearance. The geometry of an object defines the shape of the object, while the material applied to the object defines its appearance. Fusion comes with a range of options for applying materials and textures to geometry, so you can give your 3D objects the surface qualities you want.

Tools that describe the geometry's response to light are called illumination models. Blinn, Cook-Torrance, Ward and Phong are the included illumination models. These tools are found in the 3D > Material category of tools.

Most materials also accept textures, which are typically 2D images. Textures are used to refine the look of an object further, by adding photorealistic details, transparency or special effect. More complex textures like bump maps, 3D textures and reflection maps are also available and located in the 3D >Texture category.

Materials can also be combined together to produce elaborate and highly detailed composite materials.

Each tool 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 tools that output a 3D material. Some of these materials provide a greater degree of control over how the geometry reacts to light, providing inputs for diffuse and specular texture maps, bump mapping and environmental maps, which mimic reflection and refraction.

Material examples from the bin

Material Components

All of the standard illumination models share certain characteristics that must be understood.

Diffuse

The Diffuse parameters of a material control the appearance of an object where light is absorbed or scattered. This diffuse color and texture are the base appearance of an object, before taking into account reflections. The opacity of an object is generally set in the diffuse component of the material.

Alpha

This parameter defines how much the object is transparent to diffuse light. It does not affect specular levels or color. However, if the value of Alpha, either from the slider or a Material input from the diffuse color, is very close to or at zero, those pixels, including the specular highlights, will be skipped and disappear.

Opacity

This parameter fades out the entire material, including the specular highlights. This value cannot be mapped; it is applied to the entire material.

Specular

The Specular parameters of a material control the highlight of an object where the light is reflected to the current viewpoint. This causes a highlight that is added to the diffuse component. The more specular a material is, the glossier it appears. Surfaces like plastics and glass tend to have white specular highlights, whereas metallic surfaces like gold have specular highlights that tend to inherit their color from the material color.

Specularity is made up of color, intensity and exponent. The specular color determines the color of light that reflects from a shiny surface. Specular intensity is how bright the highlight will be.

3 spheres L to R: diffuse only, specular only, and combined

The specular exponent controls the falloff of the specular highlight. The larger the value, the sharper the falloff and the smaller the specular component will be.

L to R: white, complimentary, and matching specular colors

Transmittance

When using the software renderer, the Transmittance parameters control how light passes through a semi-transparent material. For example, a solid blue pitcher will cast a black shadow, but one made of translucent blue plastic would cast a much lower density blue shadow. The transmittance parameters are essential to creating the appearance of stained glass.

Tip: You can adjust the opacity and transmittance of a material separately. It is possible to have a surface that is fully opaque yet transmits 100% of the light arriving upon it, so in a sense it is actually a luminous/emissive surface.

Transmissive surfaces can be further limited using the Alpha and Color Detail control.

Attenuation

The transmittance color determines how much color is passed through the object. For an object to have fully transmissive shadows, the transmittance color to must be set to RGB = (1, 1, 1), which means 100% of green, blue and red light pass through the object. Setting this color to RGB = (1, 0, 0) means that the material will transmit 100% of the red arriving at the surface but none of the green or blue light.

Alpha Detail

When this slider is set to 0, the non-zero portions of the alpha channel of the diffuse color are ignored and the opaque portions of the object casts a shadow. If it is set to 1, the alpha channel determines how dense the object casts a shadow.

Tip: The OpenGL renderer will always cast a shadow from the entire object, ignoring the alpha. 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, and 0

Color Detail

Color Detail is used to color the shadow with the object's diffuse color. Increasing the Color Detail slider from 0 to 1 brings in more of diffuse color and texture into the shadow.

Tip: The OpenGL renderer will always cast a black shadow from the entire object, ignoring the color. Only the software renderer supports color in the shadow maps.

Color Detail set to 1 and 0. The gradient is clearly 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 Control Panel 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 tools require at least one light source to affect the appearance of the object. There are four different illumination models that can be found in the Tools > 3D > Material menu.

Illumination models L to R: 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 tools that load or create geometry. Connecting any tool that outputs a material to that tool's Material Input will override the Standard material and the controls in the Material tab will be hidden. See the Common 3D Controls section of the Tool Reference Manual for a complete description of the material's options.

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 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 tools in the Flow 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 tool that outputs a material is frequently used, instead of an image, to provide other shading options. Materials passed between tools are RGBA samples; they contain no other information about the shading or textures that produced them.
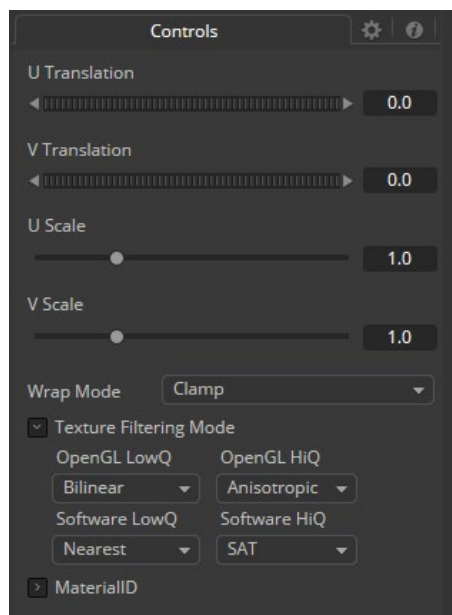
Texture 2D tool is used to translate the texture in the UV space of the object, as well as set the filtering and wrap mode

Composite Materials

Building complex materials is as easy as connecting the output of a Material tool to one of the Material inputs of another Material or Texture tool. 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 tool can be used to simulate reflections and refractions on an object. Reflections are direct bounce light that hits an object, while refractions simulate the distortion of light seen through semi-translucent surfaces.

The reflections and refractions use an environment mapping technique to produce an approximation that balances realistic results with greater rendering performance. Environment maps assume an object's environment is infinitely distant from the object and rendered into a cubic or spherical texture surrounding the object.

The Tools > 3D > Texture > Cube Map and Sphere Map tools 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 tool 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 will not respond to lights in the scene. For this reason, the Reflect is usually combined with the Blinn, Cook-Torrance, Phong or Ward tools.

Reflection

Reflection outputs a material making it possible to apply the reflection or refraction to other materials either before or after the lighting model with different effects.

A Blinn material connected to a Background Material input of the Reflect.
This causes the reflection to be added to the Blinn output.
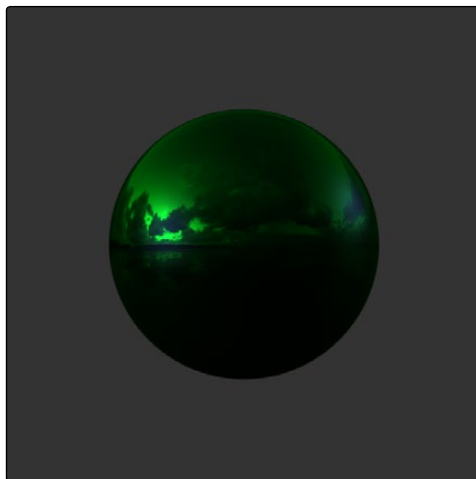
A Reflect is connected to the Diffuse Color component of the Blinn, causing the
reflection to be multiplied by the diffuse color and modulated by the lighting.

Refraction

Refraction only occurs where there is transparency in the background material, which is generally controlled through the Opacity slider and/or the alpha channel of any material or texture used for the Background Material Texture input.

The Reflect tool 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 sometimes. 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 tool. 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 Material tool

Using a Height Map

A Height map is an image where the value of a pixel represents the height. It is possible to select which color channel is used for bump creation. White means high and black means low, however it is not the value of a pixel in the height map that determines the bumpiness, but rather how the value changes in the neighborhood of a pixel.

Using a Bump Map

A bump map is an image containing normals stored in the RGB channels.

Tip: Normals are generated by 3D modeling and animation software as a way to trick the eye into seeing smooth surfaces, even though the geometry used to create the models only uses triangles to build the objects.

Normals are 3 float values (nx, ny, nz) whose components are in the range [-1, +1]. Because you can store only positive values in 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 tool 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 which expects a packed bump map or height map and will do the conversion of the bump map to work correctly.

If your bump mapping isn't appearing correctly, here are a few things to look for:

→ Make sure you have the tools 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 Hight 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 image you have.

→ Check to ensure HiQ is on. Some tools like Text+ produce an anti-aliased version in HiQ 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 tool. 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 if the projection behaves like a spotlight or an ambient light, however, alpha channels cannot be projected. Overlapping projections add together like any other light tool. 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. You can learn more about the Camera 3D and Projector 3D in the Tool Reference Manual.
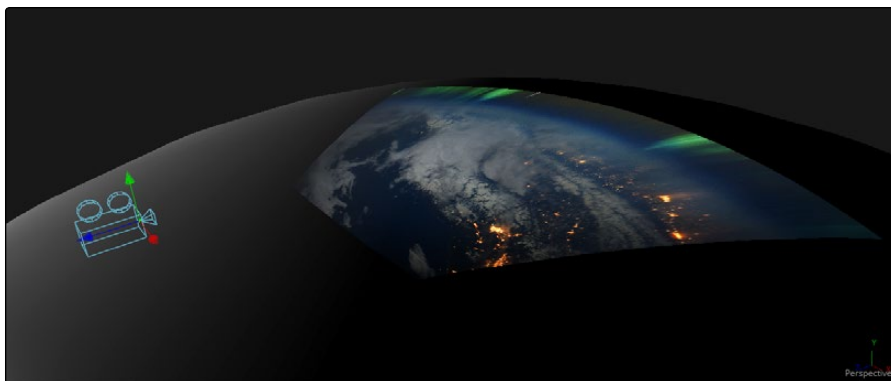
Camera tool used for a projection map

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 combined in various ways (mean, median, blend, etc.) via the Catcher tool. 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.
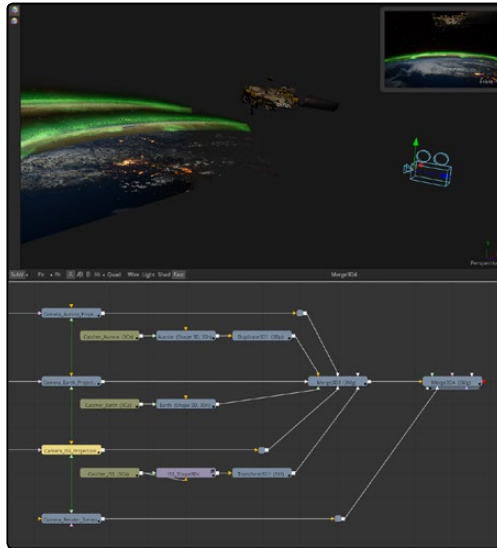
Camera projection used with a Catcher tool

Project using the UVMap Tool

This mode requires a camera and a UVMap 3D downstream of the objects to which the texture is being projected. In the Control Panel, when the UVMap Map mode is set to Camera, it gathers the information from the camera and creates new UVs for the input objects; these are used for texturing. Because the UVs are stored in the vertices of the mesh, the object must be tessellated sufficiently.
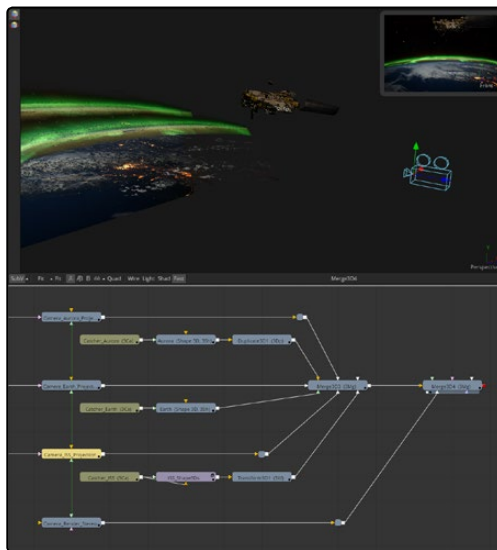
UV Map tool incorrectly tessellated (L) and correctly tessellated (R)

Textures are assigned to the object like any other texturing technique. The UVs can be locked to the vertices at a chosen frame using the Ref Time slider. This locking only works as long as vertices are not created or destroyed or reordered (e.g., projection locking will not work on particles because they get created/destroyed, a Cube 3D with its subdivision level slider animated, a Duplicate 3D tool with its time offset set to non zero so additional meshes get created over time).

Camera projection onto two image planes

Tip: Projected textures can be allowed to slide across an object. If the object moves relative to the Projector 3D, or alternatively, by grouping the two together with a Merge 3D, they can be moved as one and the texture will remain locked to the object.

Geometry

There are five tools used for creating geometry in Fusion. These tools 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 tool 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 tools is covered in more detail in the Tool Reference Manual, a summary of the 3D creation tools 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 tool used to place a 2D image into a 3D scene with an automatically scaled plane.

Shape 3D

This tool 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+ tool. This version supports beveling and extrusion but does not have support for the multi-layered shading model available from Text+.

Particles

When a pRender tool is connected to a 3D view, it will export its particles into the 3D environment. The particles are then rendered using the Renderer 3D instead of the Particle renderer. See the Particle chapter in the Tool Reference Manual for more detail.

Visibility

The Visibility options are found on most 3D source tools under the Visibility disclosure arrow. They control object visibility in the Viewers and in the final render.

A 3D geometry tool's visibility settings

Visible

If the Visibility checkbox is not selected, the object will not be visible in a Viewer, nor will it be rendered into the output image by a Renderer 3D. A non-visible object does not cast shadows. This is usually enabled by default, so objects that you create are visible in both the Viewers and final renders.

Unseen by Cameras

If the Unseen by Cameras checkbox is selected, the object will be visible in the Viewers, but invisible when viewing the scene through a camera and so the object will not be rendered into the output image by a Renderer 3D. Shadows cast by an Unseen object will still be visible.

Cull Front Face/Back Face

Use these options to cull (exclude) rendering of certain polygons in the geometry. If Cull Back Face is selected, all polygons with normals pointing away from the view will not be rendered and will not cast shadows. If Cull Front Face is selected, all polygons with normals pointing away from the view will likewise be excluded. Selecting both checkboxes has the same effect as deselecting the Visible checkbox.

Ignore Transparent Pixels in Aux Channels

For any piece of geometry, the Renderer 3D rejects transparent pixels in the auxiliary image channels. The reason this is the default is to prevent aux channels (e.g., normals, Z channel, UVs) from filling in where there should be blank space or full transparency. For example, suppose in post you want to add some fog to the rendered image. If you had fully transparent geometry in the foreground affecting the Z channel, you would get incorrect fog rendering. By deselecting this checkbox, the transparency will not be considered and all of the aux channels will be filled for all of the pixels. This could be useful if you wanted to replace texture on a 3D element that is fully transparent in certain areas with a texture that is transparent in different areas; it would be useful to have the whole object set aux channels (in particular UVs).

FBX Format

The Filmbox FBX format is a scene interchange format that facilitates the moving of 3D scene information from one application to another. Fusion's FBX format extends support to other 3D files such as Collada and OBJ.

Importing An FBX Scene

To import an entire FBX scene, you choose File > Import FBX Scene. This creates a composition with the same lights, cameras, materials and geometry found in an FBX file.
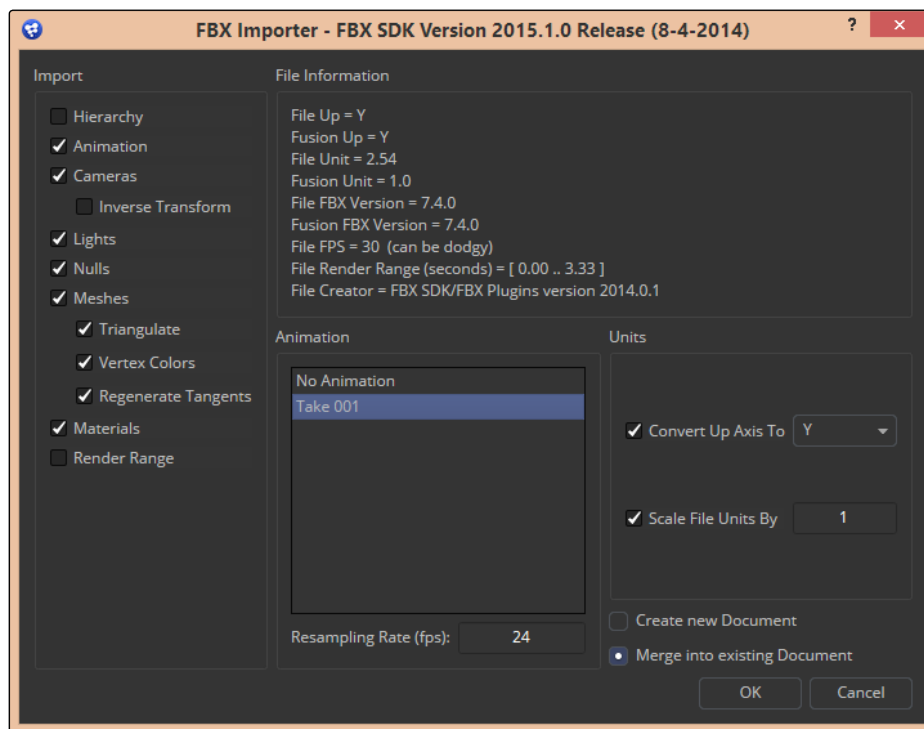
FBX Scene Import dialog

Importing FBX Geometry

The FBX Mesh tool is used to import mesh geometry from an FBX file. The first texture applied to the 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. Fusion will treat the scales 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 into Fusion as a massive 100 units.

You can use the Size slider in the FBX Mesh to reduce the scale of such files to something that matches the Fusion 3D scene.

FBX Exporter

You can export a 3D scene from Fusion to other 3D packages using the FBX Exporter tool. On render it saves geometry, cameras lights and animation into different file formats like .dae or .fbx. The animation data can be included in one file or it can be baked to sequential frames. Textures and materials cannot be exported.

Fog 3D and Softclipping

The Fog 3D tool helps to create atmospheric depth cues.

Split screen with and without fog

The Fog 3D works well with depth of field and antialiasing supported by the OpenGL renderer. Since it is not a post processing tool (like the Volume Fog tool found in the Tools > Position menu or Fog tool in Tools > Deep Pixel), it does not need additional channels like Position or Z channel color. Furthermore, it supports transparent objects.

SoftClip 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.

Matte Objects

3D Geometry uses a Matte Objects option for 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 settings

When the Matte disclosure arrow is clicked it reveals the IsMatte option which can enable another two options.

Matte settings opened and IsMatte enabled to reveal additional options

IsMatte

Located in the Controls tab for the geometry, this is the main checkbox for matte objects. When enabled, objects whose pixels fall behind the matte object's pixels in Z do not get rendered.

Opaque Alpha

When the IsMatte checkbox is enabled, the Opaque Alpha checkbox is displayed. Enabling this checkbox sets the alpha value of the matte object to 1. Otherwise the alpha, like the RGB, will be 0.

Infinite Z

When the IsMatte check box is enabled, the Infinite Z checkbox is displayed. Enabling this checkbox sets the value in the Z channel to infinite. Otherwise, the mesh will contribute normally to the Z channel.

Matte objects cannot be selected in the Viewer unless the Show Matte Objects viewer option is selected. However, it is always possible to select the matte object by selecting its tool in the flow.
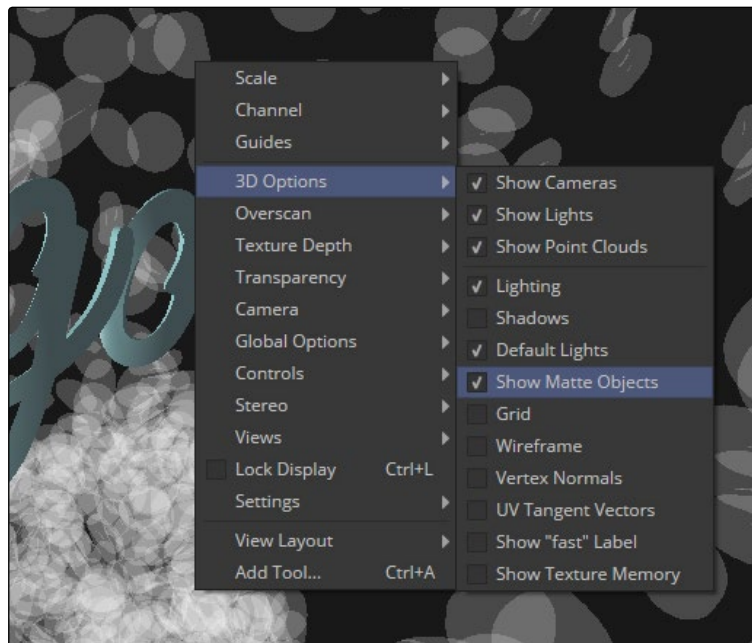
Show Matte Objects setting in the Viewer contextual menu

Material and Object IDs

Most tools in Fusion that support effect masking can use the Object ID and Material ID auxiliary channels to generate a mask. The settings used to accomplish this are found in the Common Controls tab of each tool.

Material and Object ID settings in the Common Controls tab in the Control Panel

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 ⅞ of those possible positions containing negative coordinates, the Position channel is always rendered in 32bit float.

A World Position Pass Rendering of a scene with its center at (0,0,0). Actual image is on the left.

3D Scene Input

Tools 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 Flow View if the world space Camera Position inputs are set to the correct value.. VolumeMask does not use a camera input. Tools that support the World Position Pass, located under the Position category, offer a Scene input, which can either be a 3D Camera or a 3D scene containing a camera.

There are three Position tools that can take advantage of World Position Pass data.

Tools > Position > Volume Fog

Tools > Position > Volume Mask

Tools > 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 tools to render correctly.

Without a bounding mesh to generate Position values, the fog fills in the background incorrectly.

Point Clouds

The Point Cloud tool 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 flow.
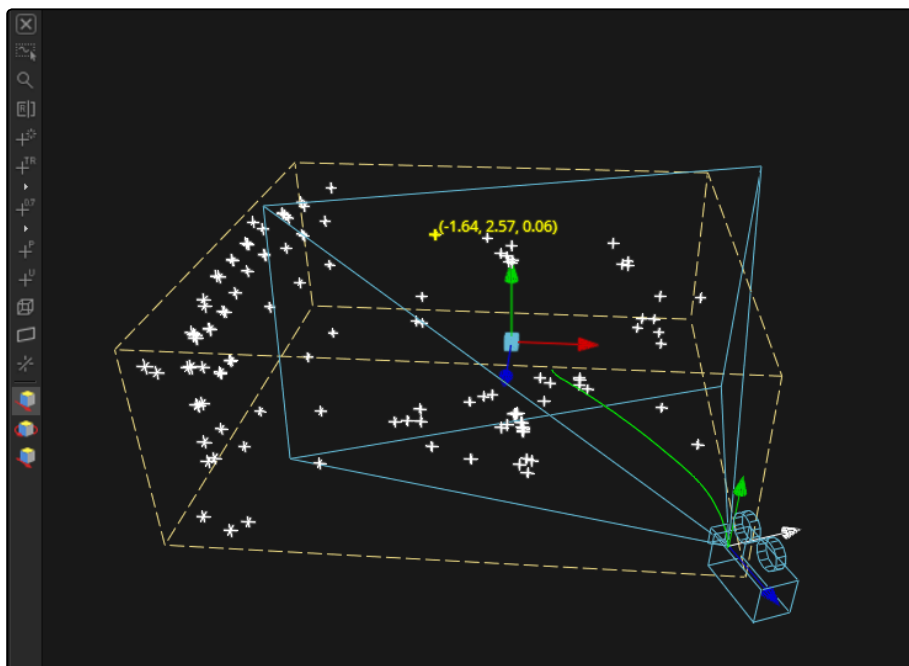
Point cloud in the Viewer

The Point Cloud tool can import point clouds written into scene files from matchmoving or 3D scanning software.

To import a point cloud, do the following:

1. Add the Point Cloud 3D to your comp.

2. Click the Import Point Cloud button in the Control panel.

3. Browse to the scene file and select the 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 in a Viewer.
2. Select Find from the Point Cloud's submenu in the contextual menu.
3. Type the name of the point and click OK.
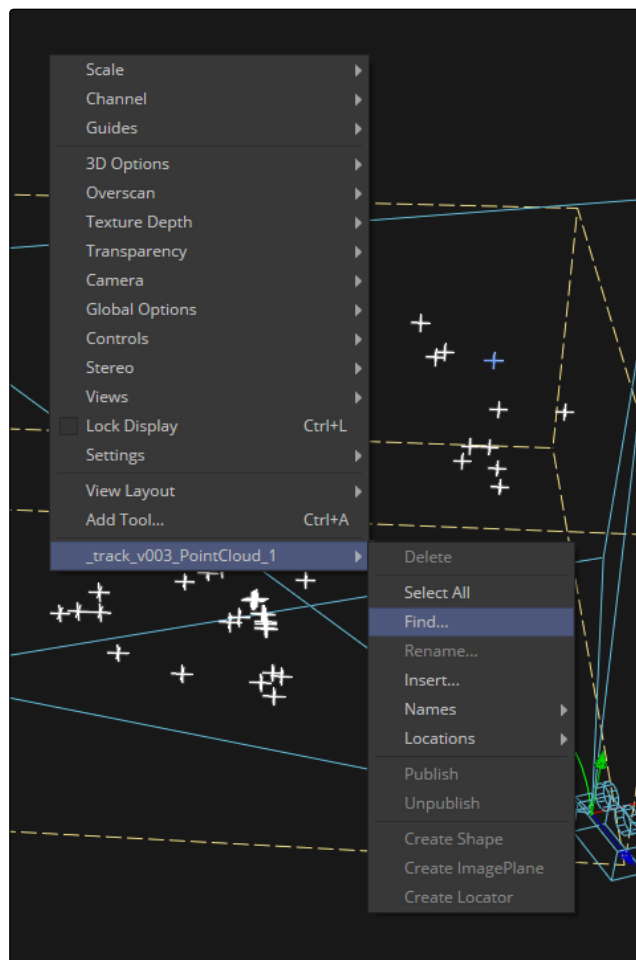
Finding a Point Cloud using the Viewer contextual menu.

If a point that matches the name entered is found, it will be selected in the point cloud and colored yellow.

Tip: The Point Cloud Find function is a case-sensitive search. A point named tracker15 will not be found if the search is Tracker15.

Renaming a Point in the Cloud

You can use the Point Cloud contextual menu to rename a selected point. This only works for a single point. A group of points cannot be renamed.

Publishing a Point

If you want to use a point's XYZ positions for connections to other controls in the scene, you can publish the point. This is useful for connecting objects to the motion of an individual tracker. To publish a point, right click over it and choose Publish from the contextual menu
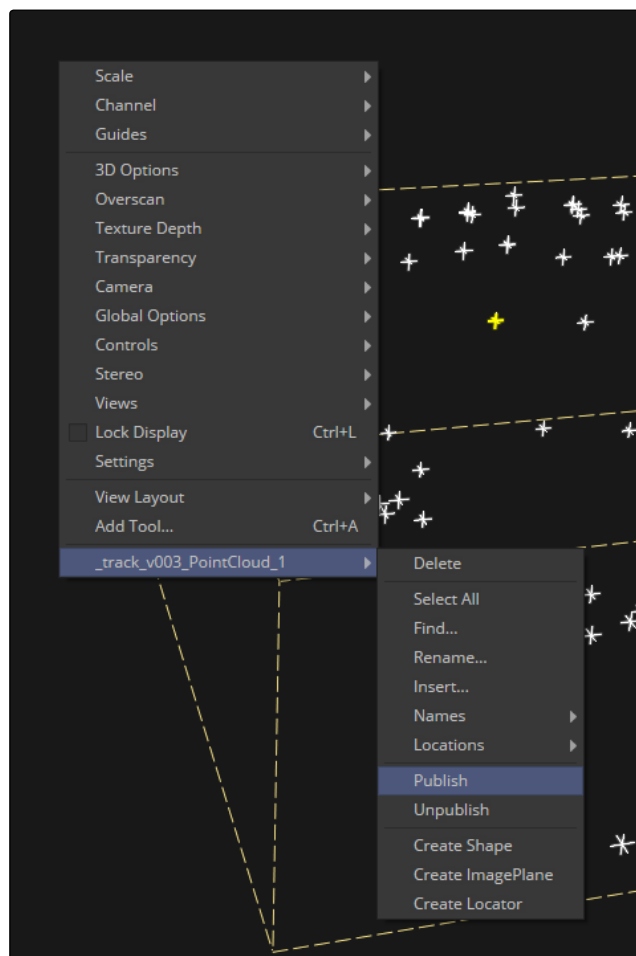
Publishing a point using the Viewer contextual menu

Auxiliary Channels

13

Auxiliary Channels

Auxiliary Channels Overview

The vast majority of computer generated images have four channels: a red, green and blue channel, and an alpha channel for transparency. 3D applications are capable of outputting more detailed data on the 3D objects in a scene. This extra information, known as the Auxiliary Channels, can contain depth information, mapping coordinates, normals orientation and a variety of other data for each pixel in the 3D rendered scene.

Consult your 3D application's documentation to determine which auxiliary channels are supported.

An RGBA 3D rendered scene

Channel Types

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

Depth

Each pixel in the Z-Depth channel contains a value, typically floating point, that represents the relative depth of that pixel in the scene. A Z-Depth channel can be used to perform depth merging. Typical merging requires one layer to be selected as the foreground and another as the background. Depth merging uses the Z-Depth value for each pixel to determine which is in front of which.

The rendered Z-Depth channel for the previous RGBA image

Most 3D applications will take the depth value from the object closest to the camera when two objects are present in the pixel.

Z-Coverage

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

Z-Coverage channel

Background RGBA

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

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

Background RGBA

Object ID

Most 3D applications are capable of assigning ID values to objects in a scene. Each pixel in the Object ID channel will be identified by that ID number, allowing for the creation of masks.

Object ID

Material ID

Most 3D applications are capable of assigning ID values to materials in a scene. Each pixel in the Material ID channel will be identified by that ID number, allowing for the creation of masks based on materials.

Material ID

UV Texture

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

UV Texture

X, Y And Z Normals

The X, Y and Z Normal channels contain information about each pixel's orientation (direction) in 3D space.

XYZ Normals

XY Vector and XY BackVector

The Vector channels indicates the pixel's motion from frame to frame. It can be used to apply motion blur to an image as a post process or to track pixels over time for retiming. The XY Vector points to the next frame, while the XY BackVector points to the previous frame.

XY Vector

XYZ Position

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

For more information on using Position channels in Fusion, read Chapter 12, 3D.

XYZ Position

XY Disparity

The XY Disparity channels indicate where each pixel's corresponding matte can be found in a stereo image. Each eye, left and right, will use this vector to point to where that pixel would be in the other eye. This can be used for adjusting stereo effects, or to mask pixels in stereo space.

XY Disparity

Image Formats

Fusion supports auxiliary channel information contained in a variety of image formats. The number of channels and methods used are different for each format.

Digital IFF (*.IFF)

The IFF image format used by Alias' Maya is capable of containing Z-Depth data in addition to the RGBA channels. Not all IFF files contain auxiliary channel information.

OpenEXR (*.exr)

The OpenEXR file format can contain an arbitrary number of additional image channels. Many renderers that will write to the OpenEXR format will allow the creation of channels that contain entirely arbitrary data. For example, a channel with specular highlights might exist in an OpenEXR. In most cases, the channel will have a custom name that can be used to map the extra channel to one of the channels recognized by Fusion.

SoftImage PIC (*.PIC, *.ZPIC and *.Z)

The PIC image format (used by SoftImage) can contain Z-Depth data in a separate file marked by the ZPIC file extension. These files must be located in the same directory as the RGBA PIC files and must use the same names. Fusion will automatically detect the presence of the additional information and load the ZPIC images along with the PIC images.

Wavefront RLA (*.RLA), 3ds Max RLA (*.RLA) and RPF (*.RPF)

These image formats are capable of containing any of the image channels mentioned above. All channels are contained within one file, including RGBA, as well as the auxiliary channels. These files are identified by the RLA or RPF file extension. Not all RLA or RPF files contain auxiliary channel information but most do. RPF files have the additional capability of storing multiple samples per pixel, so different layers of the image can be loaded for very complex depth composites.

Fusion RAW (*.RAW)

Fusion's native RAW format is able to contain all of the auxiliary channels as well as other metadata used within Fusion.

Creating Auxiliary Channels

Some tools in Fusion make auxiliary channels on their own. Renderer 3D creates many channels just as any other 3D application would. 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.

Viewing Auxiliary Channels

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

Selecting a channel from the Viewer's Toolbar

Using Auxiliary Channels

Fusion includes tools that are specifically intended to work with images that contain auxiliary channels.

Merge

In addition to regular compositing operations, Merge is capable of merging two or more images together using the Z-Depth, Z-Coverage, and BG RGBA buffer data. This is accomplished by enabling the Perform Depth Merge checkbox from the Channels tab.

Depth Blur

The Depth Blur tool is used to blur an image based on the information present in the Z-Depth. A focal point is selected from the Z-Depth values of the image and the extent of the focused region is selected using the Depth of Field control.

Fog

The Fog tool makes use of the Z-Depth to create a fog effect that is thin closer to the camera and thickens in regions farther away from the camera. You use the Pick tool to select the Z-Depth values from the image and to define the Near and Far planes of the fog's effect.

Shader

The Shader tool applies data from the RGBA, UV and the Normal channels to modify the lighting applied to objects in the image. Control is provided over specular highlights, ambient and diffuse lighting, and position of the light source. A second image can be applied as a reflection or refraction map.

SSAO

SSAO is short for Screen Space Ambient Occlusion. Ambient Occlusion is the lighting caused when a scene is surrounded by a uniform diffuse spherical light source. In the real world, light lands on surfaces from all directions, not from just a few directional lights. Ambient Occlusion captures this low frequency lighting, but it does not capture sharp shadows or specular lighting. For this reason, Ambient Occlusion is usually combined with Specular lighting to create a full lighting solution.

The SSAO tool uses the Z-Depth channel, but requires a Camera3D input.

Texture

The Texture tool uses the UV channels to apply an image from the second input as a texture. This can replace textures on a specific object when used in conjunction with the object ID or material ID masks.

Shadow

The Shadow tool can use the Z-Depth channel for a Z-Map. This allows the shadow to fall onto the shape of the objects in the image.

Vector Motion Blur

Using the forward XY Vector channels, the Vector Motion Blur tool can apply blur in the direction of the velocity, creating a motion blur effect.

Vector Distortion

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

Time Speed and Time Stretcher

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

New Eye

For stereoscopic footage, New Eye uses the Disparity channels to create new viewpoints or to transfer RGBA data from one eye to the other.

Stereo Align

For stereoscopic footage, the Disparity channels can be used by Stereo Align to warp one or both of the eyes to correct misalignment or to change the convergence plane.

Smooth Motion

Smooth Motion uses Vector and Back Vector channels to blend other channels temporally. This can remove high frequency jitter from problematic channels such as Disparity.

Volume Fog

Volume Fog is a raymarcher that uses the Position channels to determine ray termination and volume dataset placement. It can also use cameras and lights from a 3D scene to set the correct ray start point and Illumination parameters.

Volume Mask

Volume Mask uses the Position channels to set a mask in 3D space as opposed to screen space. This allows a mask to maintain perfect tracking through a camera move.

Custom Tool, Custom Vertex 3D, pCustom

The "Custom" tools can sample data from the auxiliary channels per pixel, vertex, or particle and use that for whatever processing you would like.

Lumakeyer

The Lumakeyer tool can be used to perform a key on the Z-Depth channel by selecting the Z-Depth in the channel drop down list.

Disparity to Z, Z to Disparity, Z to WorldPos

These tools use the inherent relationships between depth, position, and disparity to convert from one channel to another.

Copy Aux

The Copy Aux tool can copy auxiliary channels to RGB and then copy them back. It includes some useful options for remapping values and color depths, as well as removing auxiliary channels.

Channel Boolean

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

Using Object and Material IDs

The Object ID and Material ID auxiliary channels can be used by some tools in Fusion to generate a mask. The settings used to accomplish this are found in the Common Controls tab of the tool.
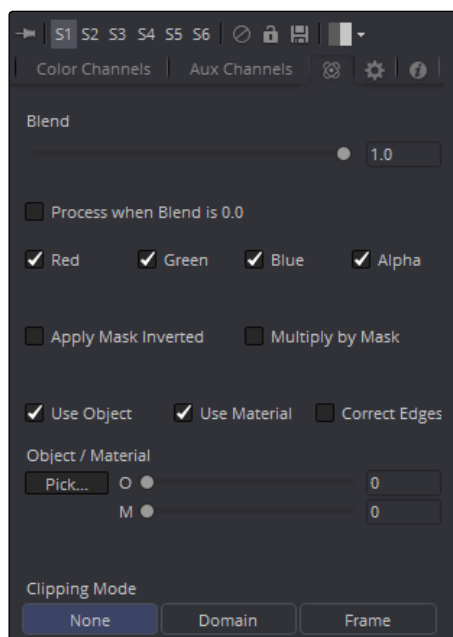
Common Controls Tab

Introduction to Controls and Customization

14

Introduction to Controls and Customization

Overview

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 acts as an introduction to a variety of advanced features, including Expressions, Scripting and Modifiers, which can help you extend the functionality and better integrate Fusion into your studio.

Modifiers

Parameters, also called Inputs, in Fusion can be controlled with Modifiers. These modifiers can be as simple as keyframe animation or linking the parameters to other tools, or can be complex expressions, procedural functions, external data, 3rd party plug-ins, or fuses.

You can add modifiers to an input through the Control Panel contextual menu or on the Preview control in the viewer.
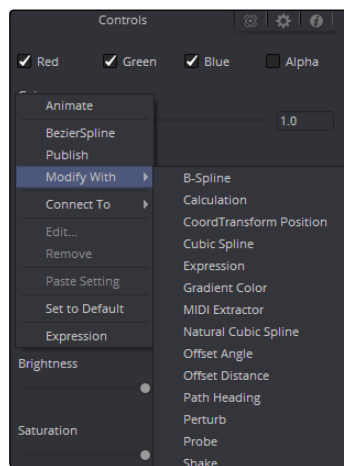
Control Panel contextual menu Modify with submenu
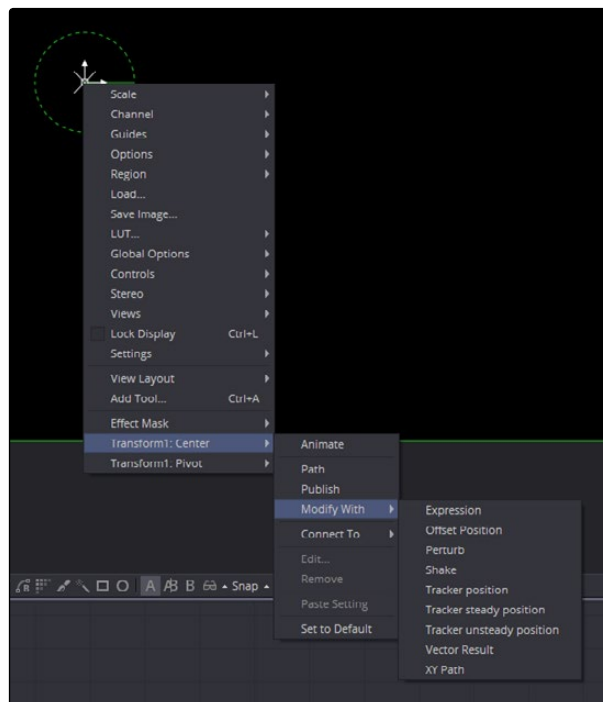
Preview Control Point contextual menu Modify with submenu

Modifiers are filtered according to the type of the input: numbers, text, polylines, gradients, points, etc.

Fusion's user interface will alert you that an input is modified by highlighting the left side of the Control Panel for that parameter to either green for spline animation or blue for other modifiers. Some parameters, like those for numbers and points, will also color the numerical entry field.

Animated parameter

Modified parameter

Modifiers with additional UI will be displayed in the Modifiers tab of the Control Panel. When a selected tool has modifiers applied, a marker will appear on the Modifier tab as an indication.

Modifier panel

Modifiers will appear with title bars just like tools, and those title bars can also be dragged into a Viewer to see their output.

Once a modifier has been applied to a parameter, you can connect another parameter to that modifier so that the modifier can affect both parameters. This can be handled through the Connect To contextual menu. As with modifier assignment, the list is filtered by the type of the parameter. This connection is bi-directional. Editing either parameter will cause the other parameter to change.
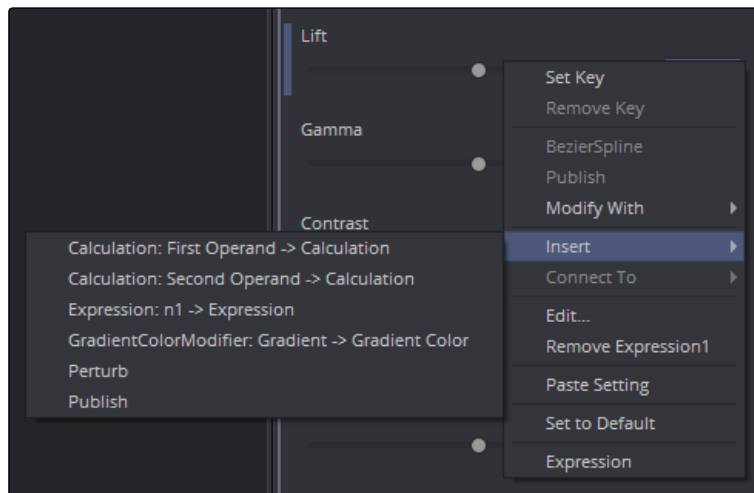
Connect To contextual menu

The Publish modifier does nothing by itself, but it allows you to connect parameters together without animation or another modifier.

Modifiers can be connected to each other and branched, just like any other tool in Fusion. For example, the Calculation modifier outputs a Number, but has two Number parameters, both of which can have modifiers added to them. If you want to insert a modifier between the existing modifier and the modified parameter, use the Insert submenu of the parameter's contextual menu.
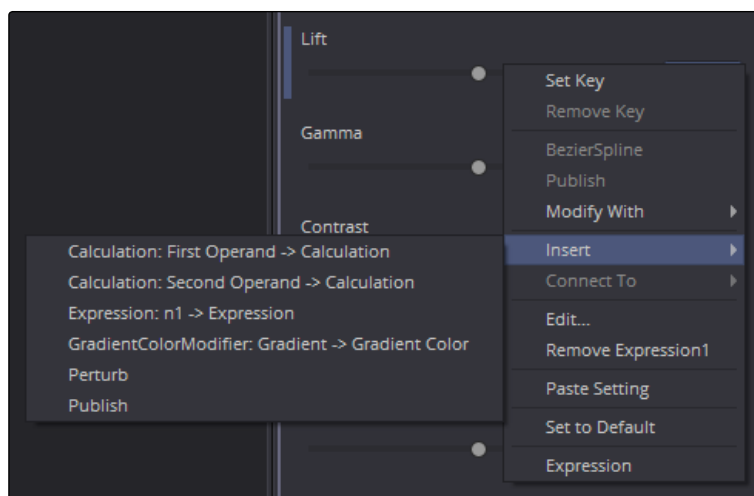
Insert submenu

For a complete listing of built-in modifiers, see the Tool Reference Manual.

SimpleExpressions

SimpleExpressions are a special type of script that can be placed alongside the parameter it is controlling. These are useful for setting simple calculations, building uni-directional parameter connections, or a combination of both.

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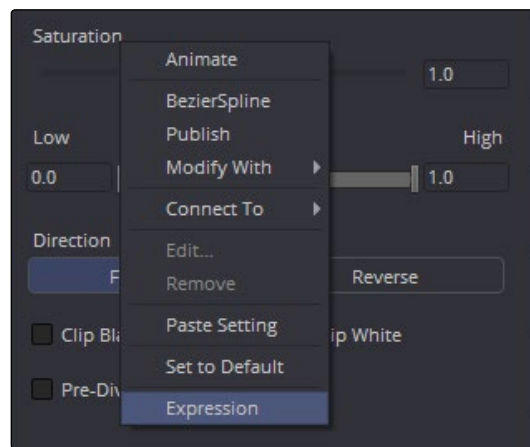
Control contextual menu

A text entry box will appear below the control, and a yellow indicator will appear to the left of the control. The current value of the parameter will be entered into the text box.

Parameter with SimpleExpression applied

Inside this text box you can enter one line scripts in Lua with some Fusion specific shorthand.

Some examples of SimpleExpressions and their syntax:

`tiwme`

This returns the current frame number.

`Merge1.Blend`

This returns the value of another input, Blend, from another tool, Merge1.

`Merge1:GetValue("Blend", time-5)`

This returns the the value from another input, but sampled at a different frame, in this case 5 frames before the current one.

```
sin(time/20)/2+.5
```

This returns a sine wave between 0 and 1. If you look in the Spline Editor, you can see the values plotted out over time. This is a good way to check how your SimpleExpression evaluates over time.

Sine wave in Spline Editor

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

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

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

This returns 0 if the image connected to the current tool's Input is tagged with sRGB colorspace. When no other tool name is supplied, the expression assumes the Input is coming from the current tool. It is equivalent to self.Input. The Input in most, but not all, Fusion tools is the main image input shown in the Flow Node Editor as an orange triangle. Images have members that can you can read, like Depth, Width, Metadata, etc.

```
Point(Text1.Center.X, Text1.Center.Y-.1)
```

Unlike the previous examples, this returns a Point, not a Number. Point inputs use two members, X and Y. In this example, the Point returned is 1/10 of the image height below the Text1's Center. This can be useful for making unidirectional parameter links, like offsetting one Text from another. Another way to get the same result would be:
```
Text1.Center - Point(0,.1)
```

```
Text("Colorspace: "..(Merge1.Background.Metadata.ColorSpaceID )
```

This SimpleExpression returns Text instead of a Number or Point. The string inside the quotes is concatenated with the metadata string, perhaps returning:
```
Colorspace: sRGB
```

```
Text("Rendered "..os.date("%b %d, %Y").. " at "..os.date("%H:%M").."\n
on the computer "..os.getenv("COMPUTERNAME").. " running "..os.
getenv("OS").."\nfrom the comp "..ToUNC(comp.Filename))
```

This returns a much larger Text, perhaps:

```
Rendered Nov 12, 2015 at 15:43
on the computer Rn309 running Windows_NT
from the comp \\SRVR\Proj\Am109\SlateGenerator_A01.comp
```

The OS library can pull various information about the computer and environment. Here os.date with different formatting arguments gets the date and time. Any environment variable can be read by os.getenv, in this case the computer name and the operating system. Various attributes from the comp can be accessed with the comp variable, like the file name, expressed as a UNC path. To get a new line in the Text, the escape sequence \n is used. When working with long SimpleExpressions, it may be helpful to drag the Tool Control panel out to make it wider or to copy/paste from a text editor or the Console.

For more details about writing SimpleExpressions, see the Fuse Guide, Scripting Guide, and the official Lua documentation.

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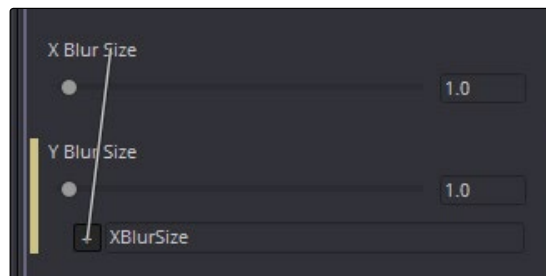
Pickwhipping with Tool Control ID already in the box

SimpleExpressions can also be created and edited in the Spline Editor. Right click on the parameter in the Spline Editor and select Set Expression from the contextual menu. The SimpleExpression will be plotted in the Spline Editor, allowing you to see the result over time.

A quick way of setting a SimpleExpression in a Number or Point input is to type = into the text box to the left of the number already there.

UserControls

The user interface for each tool 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 SimpleExpressions and Macros, but it can be done for usability and aesthetic reasons for favorites and presets.

UserControls are most easily added through the UserControls tool script.
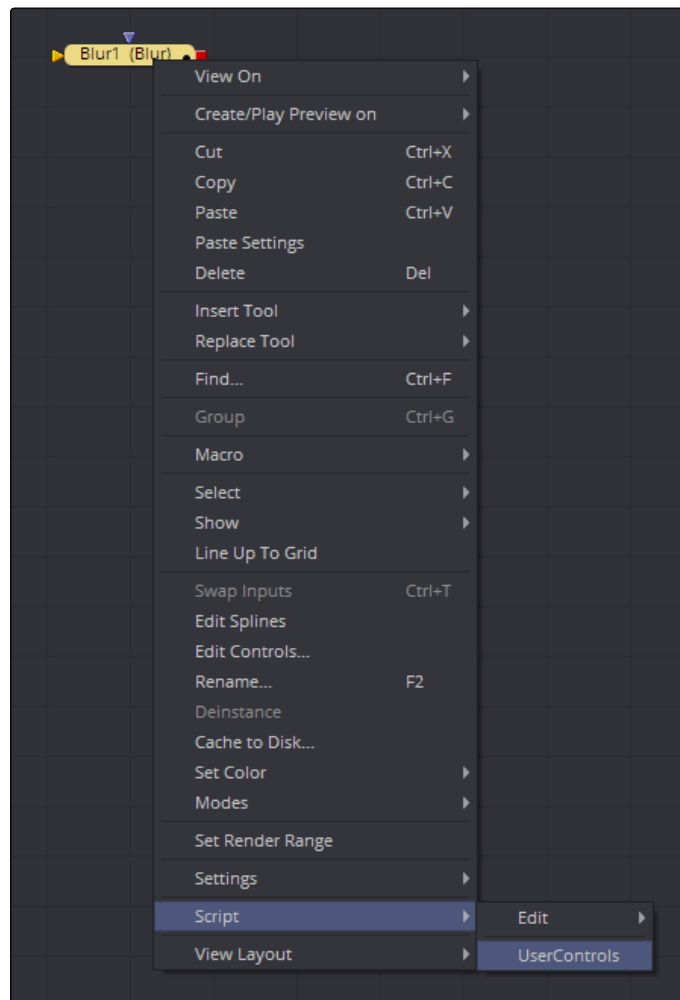
UserControls in tool contextual menu

A new window will appear, titled Add Control.
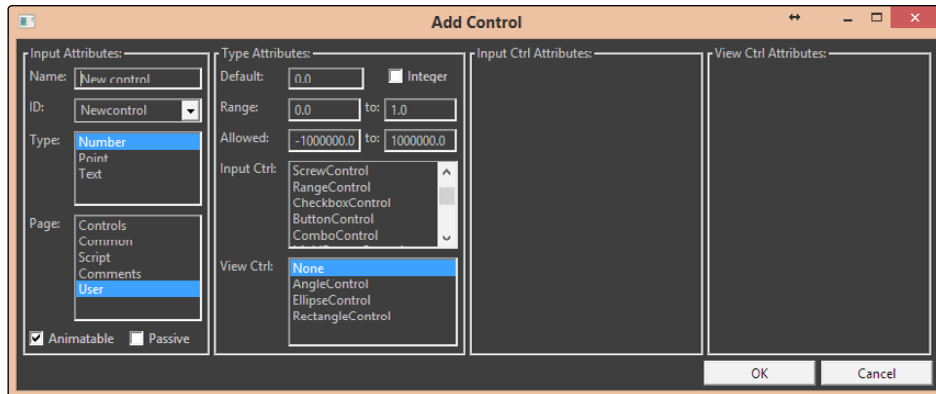
Add Control window

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 tool control, and the View Ctrl Attributes box contains settings for the preview control, if any.

All changes made using UserControls are stored in the tool instance itself, so they can be copy/ pasted, saved to a setting, added to the Bins, or added to your favorites.

Consider this example, a customization of the controls for DirectionalBlur:
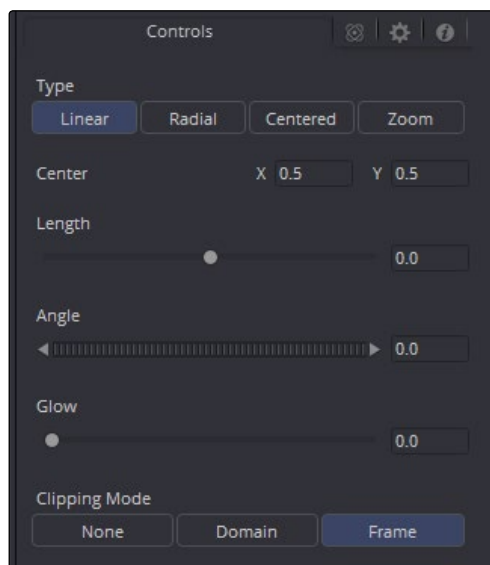
A default DirectionalBlur

Let's say we wanted a more intuitive way of controlling a linear blur than using the Length and Angle sliders. We could use the Center input control, along with its preview control, to set an angle and distance from directly within the Viewer. We'll do that using SimpleExpressions for Length and Angle:

```
sqrt(((Center.X-.5)*(Input.XScale))^2+((Center.Y-.5)*(Input.
YScale)*(Input.Height/Input. Width))^2) and atan2(.5-
Center.Y , .5-Center.X) * 180 / pi respectively.
```
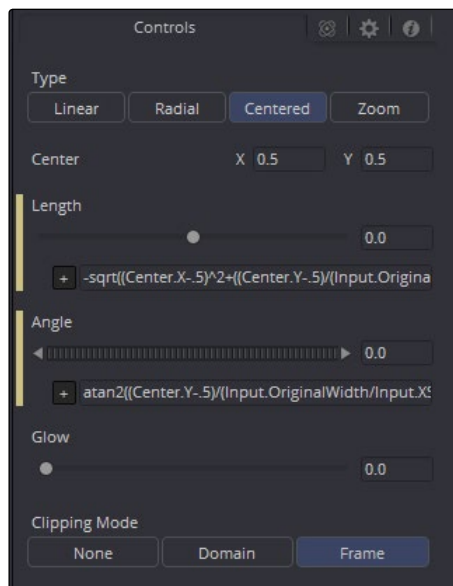

DirectionalBlur controlled by the Center's position

This functions fine, but the controls are confusing. The Center doesn't function as the center anymore, it should be named "Blur Vector" instead. The controls for the Length and Angle aren't meant to be edited, so they should be hidden away, and we're only doing a linear blur, so we don't need the buttons for Radial or Zoom. We just need to choose between Linear and Centered.

For the first task, let's rename the Center. From the Add Control window, select the Center from the ID list. A dialog will appear asking if you would like to Replace, Hide, or Change ID. We'll chose Replace. Now we are editing the Center input. We'll change the Name to Blur Vector, set the Type to Point, and the Page to Controls, which is the first tab where the controls are normally. Press OK, and our new input will appear on our tool in the flow. The ID of the control is still Center, so our SimpleExpressions did not change.

To hide the Length and Angle, we'll run the UserControls script again. This time when we select the Length and Angle IDs, we'll choose Hide in the dialog. Press OK for each.

Finally, to change the options available in the Type, we have two options. We can hide the buttons and use a checkbox instead, or we can change the MultiButton from 4 entries to 2. Let's try both.

To add the checkbox, run UserControls again, but this time instead of selecting an existing ID, we'll type Centered into the Name. This will set the name and the ID of our input to Centered. The Type is set to Number and the Page is set to Controls. Now in the Type Attributes, set the Input Ctrl to be CheckboxControl. Press OK, and now we have our checkbox. To make the new control affect the Type, add a SimpleExpression to the Type:

`iif(Centered==1, 2, 0).`

Once that's done, we can use the UserControls to hide the Type control.

To make a new MultiButton, run the UserControl script, and add a new control ID, TypeNew. You can set the Name to be Type, as the Names do not need to be unique, just the IDs. Set the Type to Number, the Page to Controls, and the Input Ctrl to MultiButtonControl. In the Input Ctrl Attributes, we can enter the names of our buttons. Let's do Linear and Centered. Type them in and hit Add for each. Press OK, and we have our new buttons with the unneeded options removed. To make this new control affect the original Type, add a SimpleExpression to the Type:

`iif(TypeNew==0, 0, 2).`

Once that's done, we can use the UserControls to hide the original Type control.
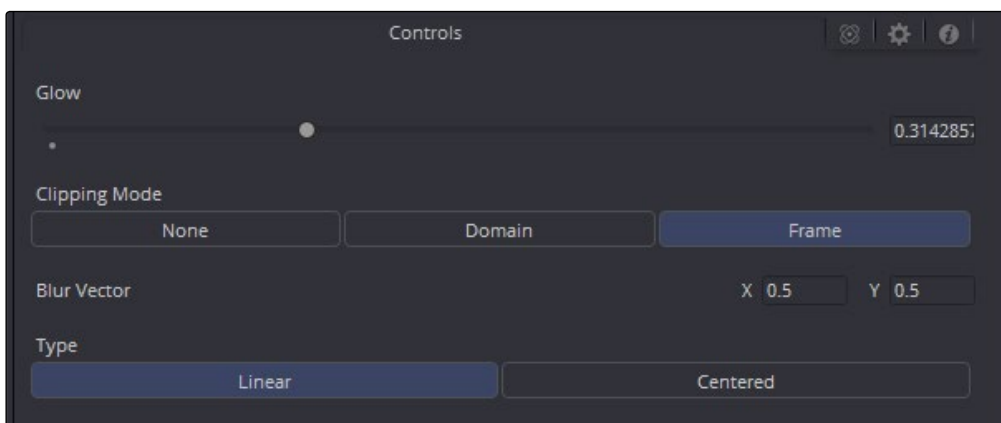
DirectionalBlurs with UserControls applied

For more information, see the Reference Manual.

Macros and Groups

Some effects aren't built with one tool, but with a series of operations, sometimes in complex branches with interconnected parameter controls. Fusion provides many individual tools but gives users the ability to bundle them up into Macros and Groups. These bundles reduce visual clutter, ensure proper user interaction, and improve productivity.

Macros and Groups are similar in function but differ slightly in how they are presented to the user and how they are generally created. Groups can be thought of as being quick ways of organizing a comp by reducing the visual complexity of a flow. Macros, however, take longer to create but are easier to re-use in other comps.

Groups are created by selecting the tools you want to include and pressing Command-G (Mac OS X) or Ctrl-G (Windows) or selecting Group from the selected tools' contextual menu. This will collapse the selected tiles in the flow down into a single group tile. Groups are represented in the flow as purple tiles with special buttons on them. Pressing the Expand/Collapse button or Command-E (Mac OS X) or Ctrl-E (Windows) will open up the Group into a small subflow where you can see and edit the contents or close it back down to a compact tile. If you no longer want the group, you can Ungroup it from the Group's contextual menu.
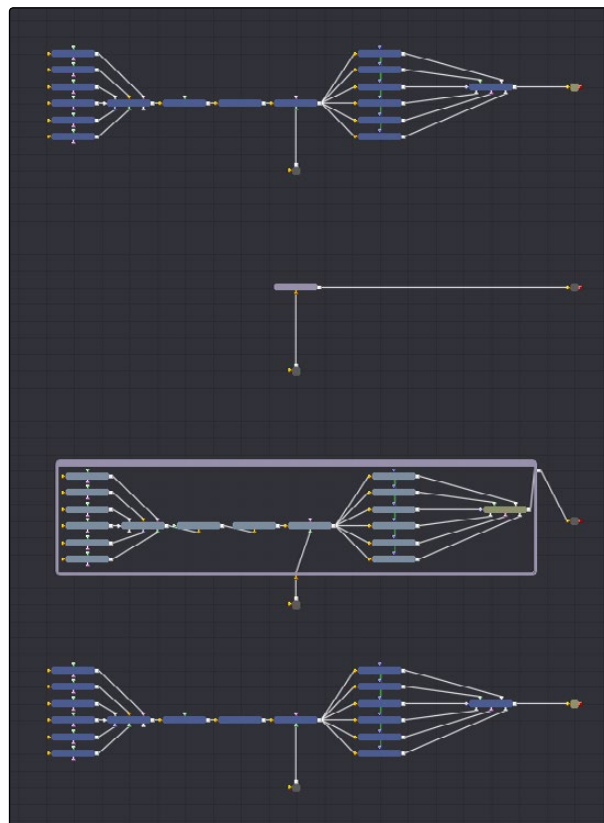
Selected tiles, Collapsed group, Expanded group, Ungrouped

Macros are created by selecting the tools you want to include and selecting Macro > Create Macro… from the selected tools' contextual menu. A Macro Editor window will appear, allowing you to name the macro, select which inputs and outputs to expose to the user, and what settings those should have including names, defaults, and ranges. Pressing close will prompt you to save the macro.
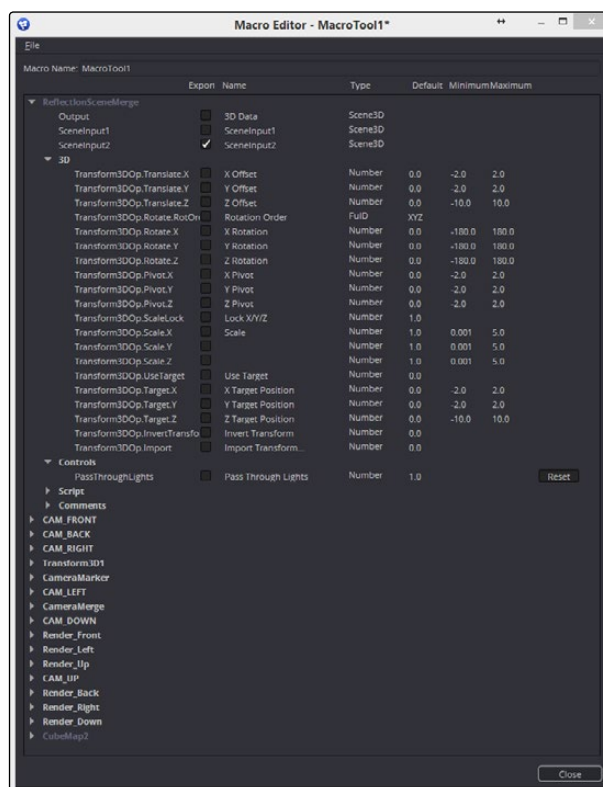
Macro Editor window

The macros can be added to the flow just like any other tool, using the AddTool window or from the Tools > Macros menu. You can access the Macro Editor to load existing macros for editing from the File menu.

Groups and Macros are somewhat interchangeable beyond this point as groups can have exposed controls added to them, just like macros, and can be loaded from the Tools > Macros menu if their .setting is saved in the Macros folder. For more information, see the Reference Manual.

Macro Examples

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

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. You could batch dailies, conform a program, generate multiple output files for delivery, rearrange tools in a flow, submit renders to a render farm, manage local data caches, log artist time, request approvals, update webpages and databases, send emails, build comps from other applications, manage comp exports to other applications, even do cloud processing of images from a web interface.

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.

Interactive scripts can be executed at various contexts in Fusion.

The most common scripts:

→ 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 commandline scripting as well and network rendering Job and Slave scripting.

FusionScript also forms the basis for Fuses and ViewShaders, special scripting based plug-ins for tools and viewers that can be used in both Fusion and Fusion Studio.

While Intool Scripts and SimpleExpressions exist in the .comp or .setting file where they were created, most FusionScripts will have the .lua or .scriptlib extension and need to be copied to the correct folders for Fusion to see them. The root folders for scripts are defined by the Script: pathmap, which defaults to AllDocs:Scripts and Fusion:Scripts.

Type of FusionScript	Extension	Location
Scriptlib	.scriptlib	:Scripts
Event Suite	.scriptlib	:Scripts
Bin Scripts	.lua	:Scripts\Bin
Comp Scripts	.lua	:Scripts\Comp
Utility Scripts	.lua	:Scripts\Utility
Tool Scripts	.lua	:Scripts\Tool
Job Scripts	.lua	:Scripts\Job
Slave Scripts	.lua	:Scripts\Slave
Hotkey Scripts	.lua	:Scripts\HotkeyScripts
Startup Scripts	.lua	:Comp
Fuses	.fuse	:Fuses
ViewShaders	.fuse	:Fuses\ViewShaders
Commandline	.lua	Anywhere

Commandline scripts can be run from anywhere, since .lua can be registered to FuScript.exe.

FuScript.exe can also be run as a service.

External scripting is done from another application by starting or connecting to an active instance of Fusion locally or on a network.

For more information, see the Scripting Guide and the Scripting Reference.

The Console

The Console view provides a place to get log and error messages, read FusionScript outputs, or input FusionScripts directly. It can be accessed with the F6 hotkey or clicking the Console tab. The Console tab will display an icon when new content is produced. The icon will identify if it is an error, log, or script output.
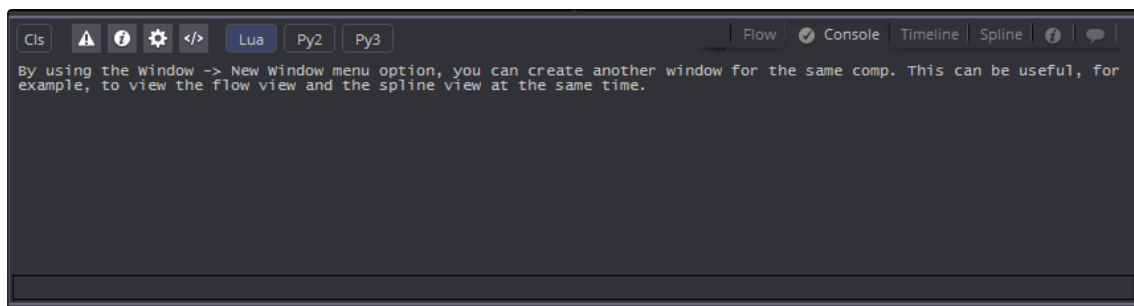
The Console

At the top of the Console is a Toolbar. The button on the far left is Clear Screen and wipes the Console contents away. The next four buttons toggle the visibility of Error messages, Log messages, Script messages, and Input echoing. Showing only a specific type of message can help you find what you are looking for in the Console. The next three buttons let you choose the input script language. Lua 5.1 is the default and is installed with Fusion. Python 2.7 and Python 3.3 require that you have the appropriate Python environment already installed on your computer. Since scripts in the Console are executed immediately, you can switch between input languages at any time.

At the bottom of the Console is the Entry Field. You can type your scripts here and execute them in the current comp context. Scripts are entered one line at a time and are executed immediately.

Useful shortcuts are available for input in the Console. A single equals sign = is the same as print(). Two equals signs == is the same as dump(), printing out the contents of a table provided.

Pressing the up arrow will walk through previous input entries. The history of these inputs is saved in the .comp file.

Some examples of Lua scripts you can input:

```
=tan(pi/3)^2
```
The Console is a convenient place to do calculations.

```
comp:GetAttrs()
```
The Console uses the composition context, so we can refer to the current comp as "comp." The current instance of Fusion can be referred to as "fu."

```
Help(composition)
```
The Console is an easy way to get help with FusionScript.

```
comp.CurrentTime=comp:GetNextKeyTime(comp.CurrentTime)
```

The Console is also useful for testing out FusionScript snippets, making sure that it works as expected. This particular line will skip to the next keyframe or range start/stop for any tool in the comp. This doesn't make much sense to use in the Console, but you could bind it to a hotkey, perhaps Shft+Ctrl+]. Now you can rapidly skip through your keyframes. To jog backwards, bind:

```
comp.CurrentTime=comp:GetNextPrevTime(comp.CurrentTime) to Shft+Ctrl+[.

for i,v in ipairs(fu:GetRegSummary(CT_Tool)) do if v.REGS_OpIconString
then print(v.REGS_OpIconString.." = ".. v.REGS_Name) end end
```

Because scripts are executed immediately after you hit return, typical loop formatting won't work. Doing the loop in one line, however, does.

```
= composition:AskUser("Motion Blur Quality For Selected Tools",
{{"Quality:", "Slider", Min = 0, Default = 10 ,Integer = true}}) for i,tool
in pairs(comp:GetToolList(true)) do tool.Quality = ask["Quality:"] end
```

Even though scripts are executed one line at a time, you can still input them one at a time. If there is an error in your script, you will get a message immediately following the input in the Console.

When you create a new comp or open an existing one, a Fusion tip will be placed in the Console. Typing "tip" into the Entry Field will display another tip.

Tip: The font used in the Console can be scaled using Command-Mousewheel or Ctrl-Mousewheel.

Fuses

A special form of FusionScript is the Fuse, which is a tool built using Lua that acts much like a plug-in created from the C++ SDK, but without needing to be precompiled. This allows it to be used in both Fusion and Fusion Studio, but it also allows you to create nearly any effect you can conceive of.

Lua is a very fast scripting language, but on high resolution images, executing pixel by pixel, even it would be slow compared to a real C++ plug-in. Fortunately, there are special enhancements for Fuses that greatly speed up the execution so that fully interactive rendering can occur.

→ Fusion API. Many functions available to Fuses come directly from the SDK, including resizing, compositing, transformations, blurring, vector painting, color matrixes, sampling, and much more. These not only speed up common operations, but they reduce the amount of code your Fuse needs to contain.

→ LuaJIT. Instead of interpreting the Lua script whenever the Fuse is rendered, the LuaJIT (Just In Time) compiler interprets the script once and compiles it to a highly optimized form for your computer that takes less memory and runs many times faster. LuaJIT also includes the ability to use C via an FFI.

→ OpenCL. Since most image processing is highly data parallel, a GPU can render some operations much more quickly than a CPU can. This not only speeds up the current operation, it also frees the CPU for other tasks, like rendering other tools in the flow. OpenCL is a standardized framework that allows parallel processing on a variety of devices from many manufacturers. You can mix OpenCL with Lua to create Fuses that split up processing according to the specific demands of the tool. OpenCL also has interoperability with OpenGL, allowing the image output from an OpenCL fuse to be displayed in the Viewer directly from GPU memory, improving interactivity by reducing memory bottlenecks.

→ Cg. Viewshader Fuses leverage the Cg language to create optimized OpenGL shaders that are applied in the Viewers. These shaders can effortlessly provide custom lookups and viewing effects run entirely on the GPU.

The examples below are not designed as a complete tutorial, but they serve as examples of the relative simplicity of the Fuse structure. At the most basic level, Fuses need to contain 3 things.

→ FuRegisterClass, which says what the Fuse is called, what type of Fuse it is, where it should go in the UI, etc.

→ Create, which says what to do when the Fuse is created. This typically adds all the inputs and outputs, like buttons and sliders and such, as well as the inputs in the Flow view.

→ Process, which is what actually gets run every time a render is requested for the output.

If you wish to actually try these Fuses out, save them as .fuse files in the Fuses: directory. You will need to restart Fusion if it is running, as Fuses get registered when Fusion launches. Once they are loaded, anything other than the FuRegisterClass portion can be refreshed without restarting Fusion, so making iterations on your tools is very fast.

The first example of a Fuse adjusts the contrast of an input. It uses two methods; one does a multithreaded iteration over the pixels of the image, the other applies a color matrix transformation to the image.

```
FuRegisterClass("UGFuseContrast", CT_Tool, {
    REGS_Name = "UG Fuse Contrast",
    REGS_Category = "Fuses\\Examples\\User Guide",
    REGS_OpIconString = "UGFC",
    REGS_OpDescription = "Example Fuse from the User Guide",
    })


function Create()
    InContrast = self:AddInput("Contrast", "Contrast", {
    LINKID_DataType = "Number",
    INPID_InputControl = "SliderControl",
    INP_MaxScale = 2.0,
    INP_MinScale = -2.0,
    INP_Default = 0.0,
    })

    InMethod = self:AddInput("Method", "Method", {
    LINKID_DataType = "FuID",
    INPID_InputControl = "MultiButtonIDControl",
    INP_Default = 0.0,
    INP_Integer = true,
    { MBTNC_AddButton = "MultiProcessPixels", MBTNCID_AddID
    = "MultiProcessPixels", MBTNCD_ButtonWidth = .5},

    { MBTNC_AddButton = "ColorMatrix", MBTNCID_AddID = "ColorMatrix",
    MBTNCD_ButtonWidth = .5},
    })

    InImage = self:AddInput("Input", "Input", {
    LINKID_DataType = "Image",
    LINK_Main = 1,
    })
```

```
    OutImage = self:AddOutput("Output", "Output", {
    LINKID_DataType = "Image",
    LINK_Main = 1,
    })
end
function Process(req)
    local img = InImage:GetValue(req)
    local contrast = InContrast:GetValue(req).Value + 1
    local method = InMethod:GetValue(req).Value

    if contrast == 0 then
    OutImage:Set(req, img)
    else
    if method == "MultiProcessPixels" then
        local out = Image({IMG_Like = img})
        out:MultiProcessPixels(nil, { Contrast = contrast },
        0,0, img.Width, img.Height, img, function (x,y,p)
            p.R = ((p.R - 0.5) * Contrast) + 0.5
            p.G = ((p.G - 0.5) * Contrast) + 0.5
            p.B = ((p.B - 0.5) * Contrast) + 0.5
            return p
        end)
        OutImage:Set(req, out)

    else
        local m = ColorMatrix()
        m:Offset(-0.5, -0.5, -0.5)
        m:Scale(contrast, contrast, contrast)
        m:Offset(0.5, 0.5, 0.5)
        out = img:ApplyMatrixOf(m, {})
        OutImage:Set(req, out)
    end
    end
end
```

This next example of a Fuse generates a new grayscale texture image using an input image as a seed. It uses the built in Blur method to simplify and speed up the code.

```
FuRegisterClass("UserGuideFuseTexture", CT_Tool, {
    REGS_Name = "UG Fuse Texture",
    REGS_OpIconString = "UGFT",
    REGS_Category = "Fuses\\Examples\\User Guide",
    REGS_OpDescription = "Example Fuse from the User Guide",
    REG_NoMotionBlurCtrls = true,
    })
function Create()
    InScale = self:AddInput("Scale", "Scale", {
    LINKS_Name = "Scale",
    LINKID_DataType = "Number",
    INPID_InputControl = "SliderControl",
    INP_MaxScale = 100.0,
    INP_Default = 2.0,
    })
    InIterations = self:AddInput("Iterations", "Iterations", {
    LINKS_Name = "Iterations",
    LINKID_DataType = "Number",
    INPID_InputControl = "SliderControl",
    INP_MaxScale = 100,
    INP_Default = 25,
    INP_Integer = true,
    })
    InImage = self:AddInput("Input", "Input", {
            LINKID_DataType = "Image",
            LINK_Main = 1,
            })
    OutImage = self:AddOutput("Output", "Output", {
    LINKID_DataType = "Image",
    LINK_Main = 1,
    })
end
```

```
function Process(req)
    local img = InImage:GetValue(req)
    local scale = InScale:GetValue(req).Value / 720
    local iterations = InIterations:GetValue(req).Value
    local result = Image({
            IMG_Like = img,
            IMG_CopyChannels = false,
            { IMG_Channel = "Alpha" },
            IMG_Depth = IMDP_8bitInt,
            })
    result = result:ChannelOpOf("Copy", img, { A = "Fg.YUV.Y" } )
    if iterations > 0.5 then

    for i = 0, iterations - 1 do
    result:Blur(result, {
            BLUR_Type = "Gaussian",
            BLUR_XSize = scale,
            BLUR_YSize = scale,
            BLUR_Normalize = 1.0,
            })
    result:Blur(result, {
            BLUR_Type = "Sharpen",
            BLUR_XSize = scale,
            BLUR_YSize = scale,
            BLUR_Normalize = 1.0,
            })
self:SetProgress(i/iterations)
    end
end
OutImage:Set(req, result)
end
```

For proper Fuse tutorials, see the Fuse Guide.

Final and Preview Renders

15

Final and Preview Renders

Rendering Overview

Since most composites happen over time, it is very important to be able to view a project or a selected clip at the intended frame rate, not just a single frame at a time. Although this can often be done by rendering the final output to disk, Fusion allows you to create previews using full resolution, floating point playback at your full frame rate.

When the intention is to evaluate an effect before your final output, Fusion offers several different ways to create and playback previews.

RAM Cache Previewers

When working with a composite, the frames rendered for interactive display are cached into memory, up to a preset memory limit. Priority is given to maintaining caches for tools that are currently displayed in a Viewer. Other tools may also cache their result, depending on available memory and the time cost of the tool, among other factors.

Pressing the Play Forward or Play Backward buttons in the Time Ruler will attempt to play back the frames in the current render range at the rate configured in the comp's frame format preferences. For example, if the frame format is set to HDTV 720, the frame rate attempted will be 60.

Transport controls in the Time Ruler

The composition will, by default, render only the tools that are required to produce the image from the tools that are currently displayed in the Viewer.

If a cache already exists for the tools currently being played, the composition will generally play at the full frame rate. The actual frame rate achieved during playback will be displayed in the status bar during playback. Frames without a cache will need to render first, but as they render they will add themselves to the cache, and the next time they play through it will be at the full frame rate.

The green line in the Time Ruler shows regions of cached frames

You can read more about RAM Flipbook Previews in Chapter 5, Viewers.

File Sequence Previews

File Sequence and Disk Previews are previews of file sequences on disk, as in the files used by a tool that has been cached to disk. A File Sequence preview plays a series of files from disk as quickly as they can be loaded from disk and formatted for use in the Viewer.

Both Loaders and Savers, which are tools that refer to files stored on disk, are able to read and write these types of previews. From the Loader or Saver you can create Flipbooks or .fb files, which are optimized for fast playback, or create .raw files, which are slower but contain more information.

When a Loader or Saver is the selected tool in the flow, transport controls appear beneath the tool controls. Tools that have a pre-rendered or populated disk cache will also reveal transport controls when they are selected.
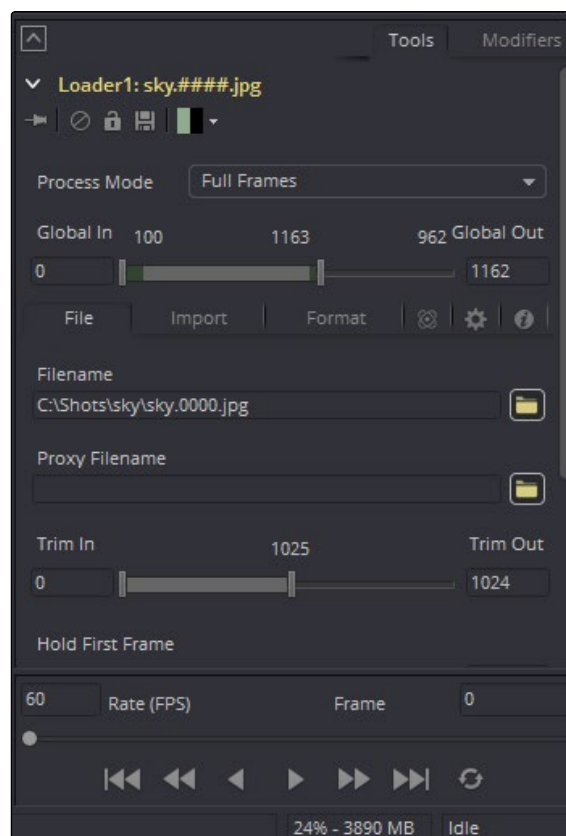
Transport controls under the Control Panel for a Loader

Clicking on any of the transport buttons will begin the file sequence playback. The left Viewer is normally used to display these previews, although the target display can be changed using the Preview preferences.

The target frame rate is chosen from the Frame Format preferences, but it can be changed using the RATE (FPS) control in the transports. The faster the drives are able to provide footage, the faster frames are displayed.

Tip: Frames played through file sequence playback are not spooled into the interactive cache.

While the file sequence playback is active, the Viewer cannot be used for normal image viewing.

To create a file sequence preview for a Loader or Saver, do the following:

1. Right click on the header of a Loader or Saver tool in the Control Panel, then choose Create Preview on > File Sequence from the contextual menu.

2. In the dialog, click Start Render to create the file sequence preview.

3. Use the transport control under the Control Panel to play the clip back.

To disable the playback of the file sequence preview, do the following:

→ De-select the tool that has enabled the file sequence preview to turn the display of the preview off.

Flipbook Previews

Flipbook Previews are stored entirely in RAM. These previews can play back at frame rates limited only by the refresh rate of your monitor, which is typically as much as 100/120 fps. The frame rate is artificially limited by the rate configured in the frame format preferences.

When a flipbook preview is available, transport controls will appear beneath the Control Panel.

To create a flipbook preview, do one of the following:

→ Right click on a tool and select Create > Play Preview from the contextual menu.

→ Hold Option (Mac OS X) or Alt (Windows) and drag the image into the Viewer.
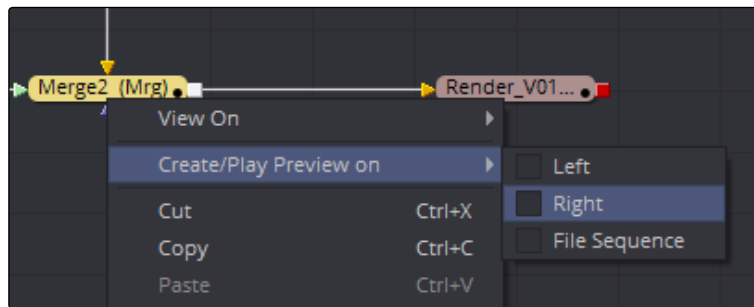
A tool's Contextual menu for creating a flipbook preview

This method can be used to create previews to video display hardware when using the Blackmagic Design DeckLink card.

To remove a flipbook preview from the Viewer do the following:

→ Right click in the Viewer and choose Remove Preview from the contextual menu.

Resizing Flipbooks

Scaling and panning a flipbook uses the same keyboard shortcuts and mouse commands used to scale and pan a tool that is displayed in the views. See Chapter 5, Viewers, for information on scaling and panning.

Saving Flipbooks

A flipbook preview can be saved to disk for later use by right clicking on the Preview and choosing Save As from the contextual menu. Flipbooks are saved to disk with the extension .FB. They can be reloaded in a Loader or Viewer that contain a flipbook at a later time for playback.

Memory Usage

Since flipbooks are stored in memory, they consume available memory used for image processing and tool caches. The amount of available memory determines the number of frames that can play back smoothly.

When creating a file preview, the Render Settings dialog will display the number of estimated frames that will fit within working memory. Rendering more than this number of frames will likely exceed available memory and cause the system to use the swap file.
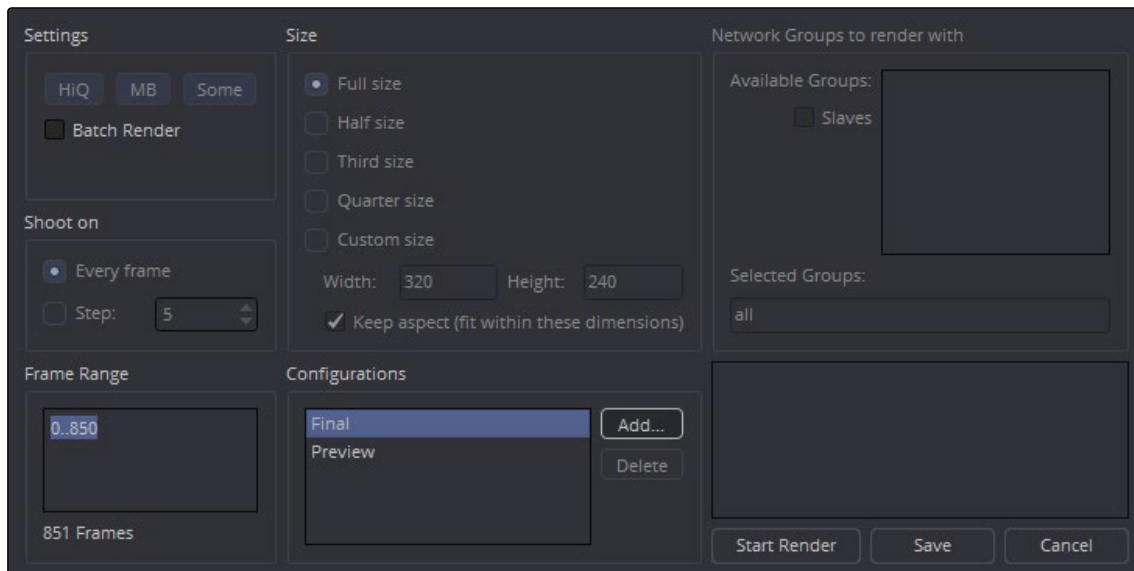
Render Setting dialog

Final Renders

When all is said and done, the composite has to be rendered to disk. This is accomplished by adding a Saver tool to the flow and pointing it to a directory on disk or over the network.

Any number of Savers can be present at any point in a flow, making it easy to save intermediate composites, as well as final results in any resolution or format with a single pass.
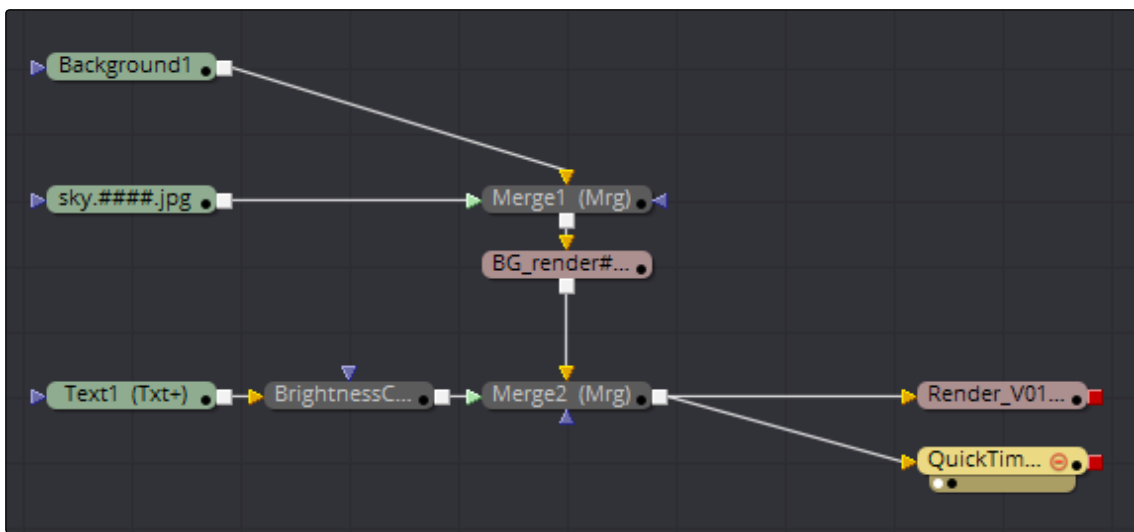
Multiple Savers used in a composition

Any tool that has a valid high quality, full resolution frame in the memory or disk cache will not need to be rendered. The cached result will be used instead.

Starting a Final Render

You can click the green Render button in the Time Ruler, choose Start Render from the File menu or press F12 to start a render. When a render is started, the Render Settings dialog will appear.

Click the Render button to begin a final render

This dialog is used to configure the rendering options. Clicking on the Start Render button on the dialog will begin the actual render.

Fusion can continue to be used while rendering, although performance will be affected by the ongoing render. You can open new compositions and work on them, but any attempt to change controls or settings in the rendering composition will display a dialog asking if you would like to abort the current render.

Tip: To bypass the Render Settings dialog when starting a render, hold Shift down while clicking on the Start Render button.

Stopping a Final Render

When a render is in progress, the green Render button in the Time Ruler turns into a brown Stop Render button. To stop the render at the current frame, click the Stop button. A dialog will appear asking for confirmation.

Click the Stop button to cancel a render

You can also stop the render by pressing the ESC (escape) key.

Tip: To pause the render, click the Stop button or ESC, but do not confirm the dialog. As long as the dialog is open, no rendering will take place. To resume, click No.

Render Status and Statistics

While a render is in progress, the status bar on the bottom right of the interface will display some statistics about the current render. The time taken to render the last frame, the average time per frame, the amount of time remaining, and the estimated time and day of completion are shown.

Status display of a render

When the render is complete, final statistics for the time taken to render and the average time to render each frame will be displayed in a dialog box, as well as in the console.
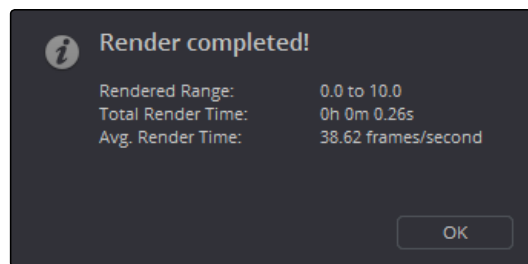
Dialog displayed at the end of a render

Troubleshooting Renders

A render can fail to start or complete for several reasons. Generally, an error message in the console window will indicate a possible cause of the error. The following are common errors that can cause a render to fail.

→ Missing frame in a file sequence used by a Loader.

→ Background input of a Merge tool could not be found.

→ Saver path is invalid.

→ Ran out of drive space on the destination.

Memory Optimization

The Preferences > Memory section has two options that affect memory use during rendering.

These options can be adjusted during a render to tune memory performance as a render progresses.
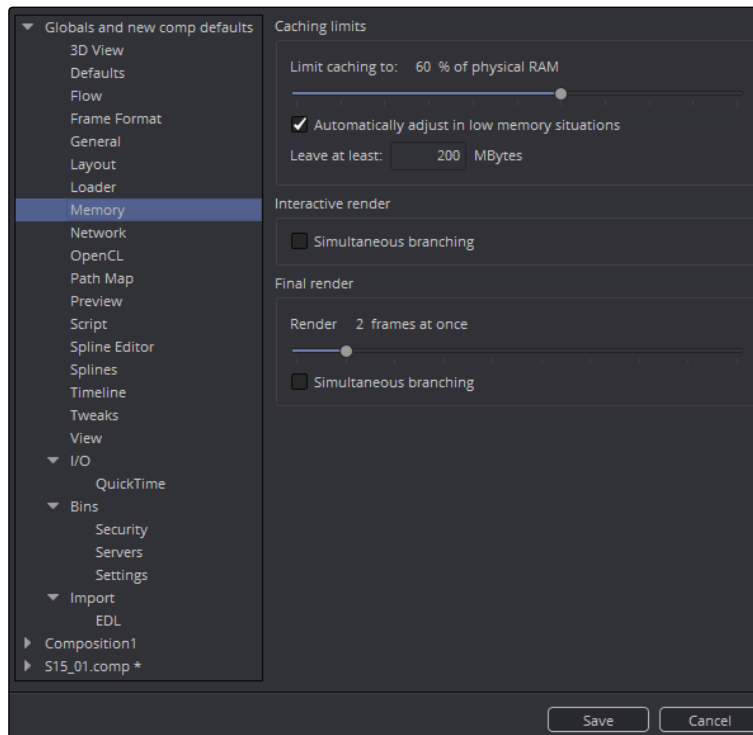
Memory Preferences

Simultaneous Branching

Enabling this Preference checkbox will render more than one branch of the flow at a time. When enabled, this option can make the most efficient use of the system's processors on a multiprocessor system, provided enough memory is available. If the memory used during rendering exceeds physical memory, disable this checkbox. Leaving this disabled significantly reduces memory usage, which is why it defaults to unchecked.

Render 'X' Frames at Once

This slider is primarily used if you are using tools that are not multithreaded or are GPU tools. It allows multiple frames to render; for instance, a 12 CPU system can render 12 frames of non-multithreaded operators efficiently. Since some tools use GPU processing, this allows other frames to do non-GPU tasks, increasing occupancy on all processors, both CPU and GPU. Decrease this value only if memory is running out and disabling Simultaneous Branching did not solve the problem.

The Render Settings Dialog

The Render Settings dialog generally appears whenever the Render button is clicked or whenever a Preview render is initiated.

The dialog provides control over the quality, range and size of the rendered frames. Features like step rendering, non-contiguous render ranges and low quality, proxied final renders are available from this dialog.
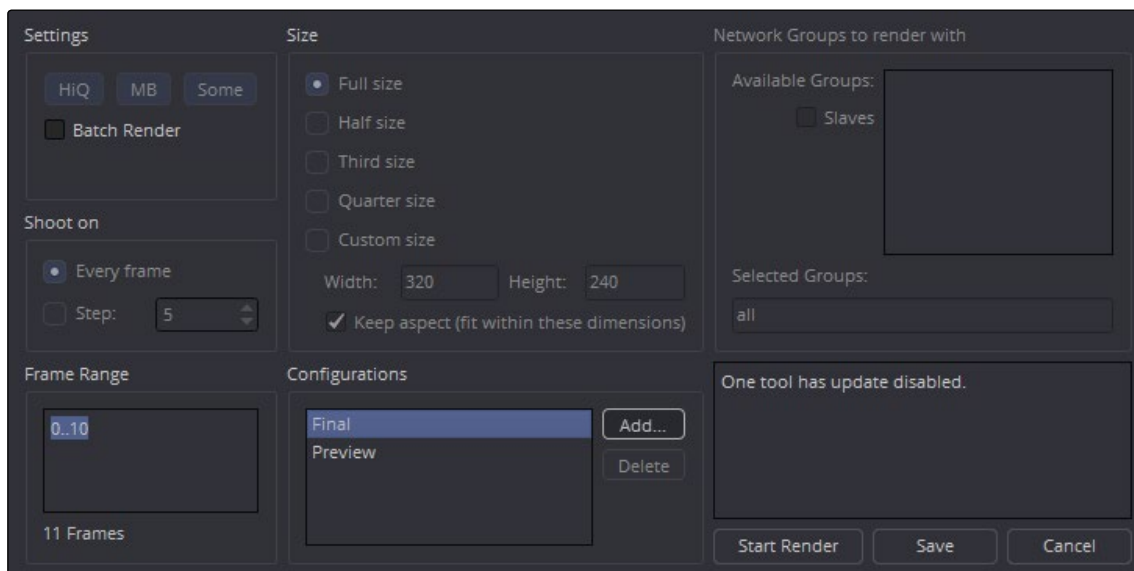
Render Settings dialog

Clicking the Render button defaults the dialog to the Final configuration, which renders every frame of the current render range at full quality, full size, with motion blur. These options are disabled unless the configuration is changed from Final.

High Quality

With the HiQ checkbox selected, the rendered image will be produced with the best possible quality. De-selecting the HiQ checkbox will cause Fusion to use faster, lower quality methods for merges, transforms, blurs and other operations. This is the same as the HiQ button in the Time Ruler.

Some/All

This option is similar to the Some/None/All option displayed in the Time Ruler. When the composition is set to Some, it will only render the tools that are required to produce output required by a Saver or other Sink tools in the flow. Orphaned branches in the composition will not render. This is a forced option for the final render settings.

The All option renders all branches in the flow regardless of whether they terminate in a Saver tool. There are very few situations where this would be required and it is mostly there for legacy support. A different configuration than Final render must be configured to enable this option.

Use Network/Batch Render

Selecting this option submits the render task to the Render Master configured in the network preferences. When selected, the Network Classes portion of the dialog will be enabled, allowing you to choose the group of Render nodes used to render the flow.

Network rendering is only available in Fusion Studio.

Size

The Size section provides control over whether the full image is rendered or if the image is processed at the selected proxy resolution.

When rendering a flipbook preview, a rough estimate of the number of frames that can be held in memory without paging to disk is displayed to the right of each size option.

Frame Range

This section is used to set the range of frames rendered. It can accept a start and end Frame Range separated by two periods, individual frames separated by a comma, or a combination of both methods. For example:

> 1..100
> 1,5,10
> 1..100,110,120,121..150

Although dashes (-) can also be used to separate the frame numbers in a range, such as 1-50, it can cause confusion when rendering frames with a negative value (i.e., Fusion allows a render range of -50..100).

Shoot On

By default, every frame in the render range is rendered. For quick previews, it is often useful to render only one out of every few frames for an overview of the motion or effect. This technique is known as Shoot On. To perform a Shoot On render, click the Step option in the Render Settings dialog and enter the step interval desired in the spinner box beneath the control.
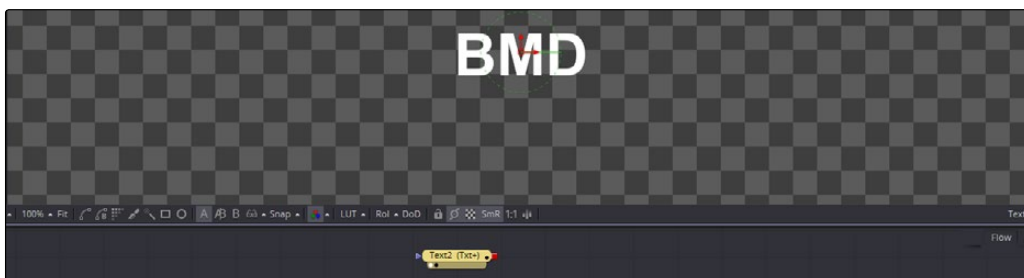
When this option is enabled for a Preview render, the frame rate of the preview will be modified so that playback of the preview will still take the same amount of time as it would if all of the frames had been rendered. For example, if the Shoot On option is set to a step of 5 frames for a RAM preview, and the render range is set to 1..100, the render will render frame 1,

6, 11, 16 until frame 100. When the preview is re-played, frame 1 will be held for five frames, frame 6 will be held for another 5 frames, and so on until the last frame is reached.

Configurations

It is possible to store the current render options for future use. Clicking the Add button in the Render Settings dialog will allow you to name current render settings. That name will then appear in the Configurations list.
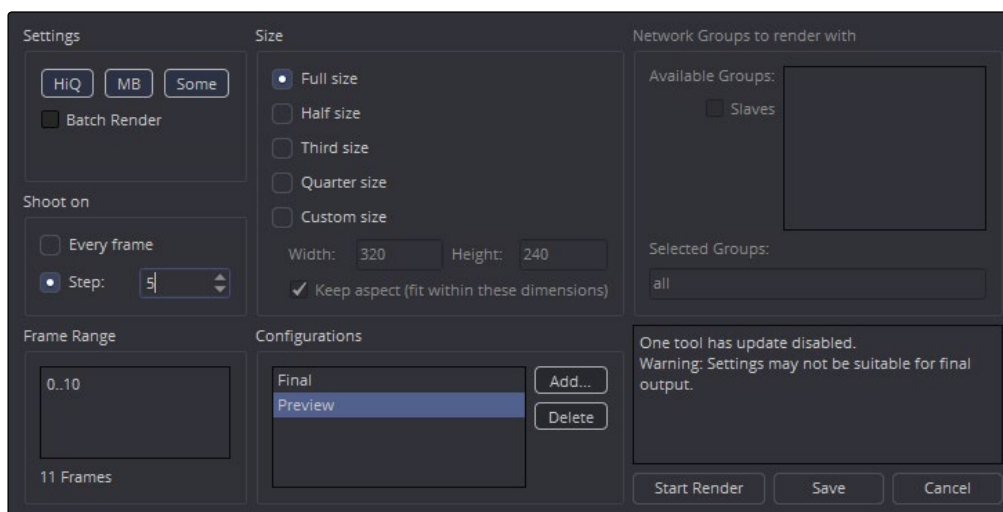
Click the Add button to save the current settings to the Configurations list

Status

The Status section displays any warnings and informational messages related to the render. For example, warnings for compositions that do not contain Savers, or that have passed through tools, will appear here.

Bypassing the Render Settings Dialog

It is possible to skip or bypass the Render Settings dialog by using whatever settings were last applied to the render. To bypass the Render settings for a Preview render, hold the Shift-Option (Mac OS X) or Shift-Alt (Windows) and drag the tool into the Viewer. To bypass the Render settings for a Final render, hold Shift and click on the Start Render button. To reverse this behavior so that the Render Settings dialog does not appear by default and will only be displayed if the Shift key is held down, use the option in the Preferences > General dialog.

RoI and DoD

Fusion uses advanced techniques called Region of Interest (RoI) and Domain of Definition (DoD) to dramatically improve rendering.

Domain Definition

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, tools will no longer be required to render portions of the image that do not get affected by the tool. This helps the renderer to optimize its performance. Secondly, Fusion can now keep track of, and apply a tools effect to, pixels that lie outside the visible portion of the image.

A small amount of text in a larger frame

For example, consider the above 320x240 pixel image, which shows the output of a Text+ tool rendered against a transparent black background. The text only occupies a portion of the pixels in the image. Without Domain of Definition, you would be required to process every pixel in the image. 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 the same image with the DoD outlined.

Text outlined by DoD

For the most part, the DoD is calculated automatically and without the need for manual intervention. For example, all of the tools in the Creator category automatically generate the correct DoD. For tools 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 tools, the DoD can often be much smaller or larger.

Media loaded from disk will typically also have the DoD default to the full image width. The exception is media stored in the .raw or OpenEXR format. These formats are capable of storing the data window of the image, and Fusion will apply this as the DoD when loading the image and will write out the DoD from the Savers.

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 tool.

Tools will shrink, expand or move the DoD as they apply their effect to the image. This causes the DoD to change from tool to tool. The current DoD for a tool is indicated in the tooltip that appears when the cursor hovers over the Tool tile in the flow. It is also visible in the Viewer when the DoD button is enabled.
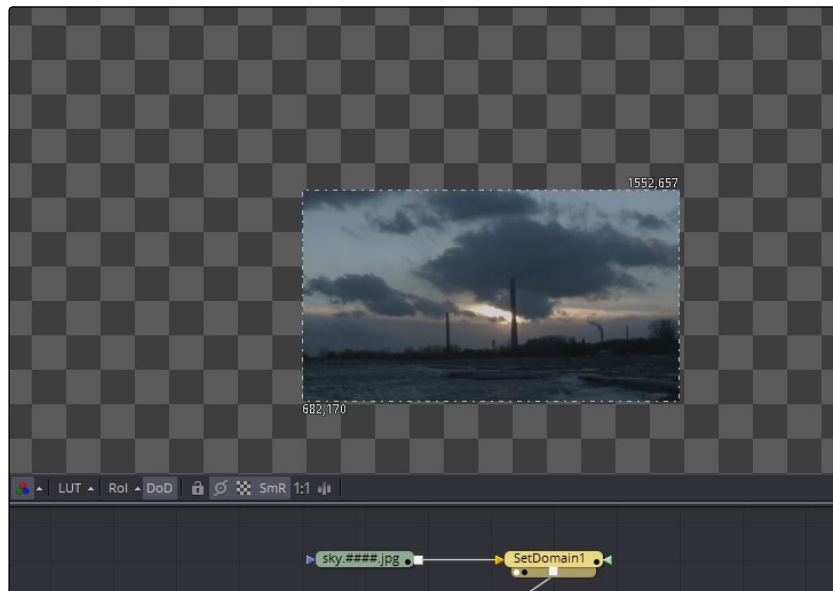
The DoD is shown as two XY coordinates indicating the corners
of an axis aligned bounding box (in pixels).

It is also possible to set the DoD for an image manually using the Tools > Miscellaneous > Auto Domain. This can be useful when dealing with pre-created media that does not occupy the full image dimensions. For example, a rendering of a 3D character that walks toward the camera will frequently only occupy a portion of the image. The Auto Domain tool can be used to animate a DoD that covers the character and ignores the rest of the image.

Region of Interest

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

You can use the RoI to restrict rendering to a small region of the image. This can significantly improve performance when working with very high resolution or complex compositions.

The RoI button in the 2D Viewer toolbar can be used to turn the RoI on or off for that Viewer. Clicking on the small arrow next to the RoI toolbar button will display a list of options to Show, Lock, Set, Reset and Auto. These ROI options are also found in the Viewer 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.

Dragging on the edges of the rectangle allows you to adjust the size of the RoI. Dragging on the corners allow you to adjust both dimensions at once. To move the RoI without adjusting its dimensions, drag on the small circle found at the top left corner of the RoI rectangle.

A faster way to set the RoI is to choose Set from the ROI pop up menu or right click over the Viewer and choose Region > Set and then drag in the Viewer to set a RoI rectangle directly.

Choosing Auto from the menu will set the region to whichever pixels are visible at the current zoom/pan level in the Viewer.

While the RoI is active, Fusion will only request rendering of the pixels inside the region when it displays an image in that Viewer. Flipbook previews created in that Viewer will also respect the current RoI. Saver tools will always use the full image dimensions when writing the image to disk, ignoring any RoI set in the views.

The RoI improves not only rendering speed and memory use, it can also reduce file I/O, since the Loader tool will only attempt to load pixels from within the RoI. 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.

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

A color correction with an RoI applied when Auto Proxy mode is on

A color correction with an RoI applied when Auto Proxy mode is off

To reset the RoI to the full width and height of the image select Region > Reset from the context menu or the Toolbar button menu, or disable ROI which will also reset it.

Tip: Image Overlay will defeat RoI, forcing renders of pixels for the entire image.

Preferences

16

Preferences

Preferences Overview

The Preferences window provides a wide variety of optional settings available for you to configure Fusion's behavior to better suit your working environment. These settings are accessed via the Preferences dialog. The Preferences dialog can be opened from the File menu at the top of the interface.

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

→ On Mac OS X, Choose Fusion > Preferences.

→ On Windows, Choose File > Preferences.

The Preferences window is divided into a sidebar that appears on the left, which is used to select which of the settings panels appear on the right. There are two levels of preferences: Global and Composition. The Global preferences are used to set options that are specific to Fusion's overall behavior as well as defaults for each new composition. The composition preferences can further modify the currently open composition without affecting the Global prefs, or any other flow that is open but not displayed.

Categories of Preferences

The first entry in the Preferences sidebar is assigned to the Global preferences. Clicking the disclosure arrows for the Global preferences reveals the following sections.

3D View

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

Default

The Default preference is used to select default behavior for a variety of options, such as animation, global range, timecode display, and automatic tool merging.

Flow

You use the Flow preference to set many of the same options found in the Flow contextual menu, like settings for Tile picture, the Navigator and pipe style.

Frame Format

The Frame Format preference is used to create new frame formats as well as select the default image height and width when adding new creator tools like Background and Text+. You also set the frame rate for playback.

General

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

Layout

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

Loader

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

Memory

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

Network

Only available in Fusion Studio, the Network rendering preference is used to configure options, such as selecting a render master, email notification and whether the machine can be used as a render slave.

OpenCL

The Open CL preference contains options for selecting OpenCL devices, enabling caching, and debugging OpenCL devices and tools.

Path Map

Path Map preference is used to configure virtual file path names used by Loaders and Savers as well as the folders used by Fusion to locate comps, macros, scripts, tool settings, disk caches and more.

Preview

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

QuickTime

This section lets you configure QuickTime codecs.

Script

The preference for Scripting includes a field for passwords used to execute scripts externally, programs to use for editing scripts, and the default Python version to use.

Spline Editor

The Spline Editor preference allows you to set various spline options for Autosnap behavior, handles, markers and more.

Splines

Options for the handling and smoothing of animation splines, Tracker path defaults, onion-skinning, roto assist and more are found in the Splines preference.

Timeline

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

Tweaks

Tweaks preference handles miscellaneous settings for modifying the behavior when loading frames over the network and queue/network rendering.

User Interface

These Preferences set the appearance of the user interface window and how the Control panel is displayed.

Video Monitoring

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

View

The View preference is used to manage settings for Viewers, including default control colors Z-depth channel viewing ranges, default LUTs, padding for fit zoom, and more.

Bins

There are three panels as part of the Bins preference: a Security panel where you set users and passwords for serving the local bins; a Servers panel used to select which remote Bin servers are connected; and a Settings panel for stamp rendering. These are only available in Fusion Studio.

Import

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

Preferences In Depth

3D View Preferences

The 3D View preferences contain settings over various defaults in the 3D Viewers including grids, default ambient light set up and Stereoscopic views.
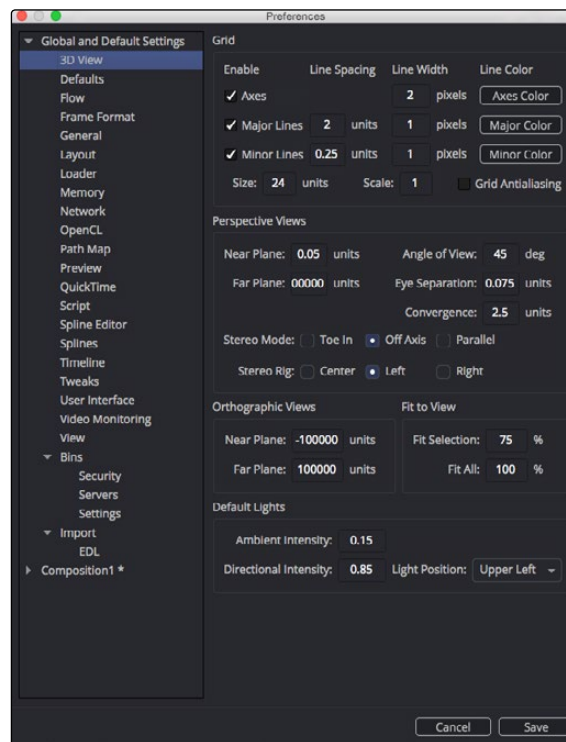
The 3D View preferences

Grid

The Grid section of the 3D View preference configures how the grid in 3D Viewers are drawn.

→ **Grid Antialiasing:** Depending on the graphics drivers installed in your computer, the antialiased lines of the grid may sort incorrectly in the 3D Viewer. Disabling this checkbox will disable antialiasing of the grid lines, correcting the issue. To turn the grid off completely, use the 3D View > Right click > Grid option, rather than preferences.

→ **Size:** Increasing the Size value will increase the number of grid lines drawn. The units used for the spacing between grid lines are not defined in Fusion. A 'unit' is whatever you want it to be.

→ **Scale:** Adjusting the overall scaling factor for the grid is useful, for example, if the area of the grid appears too small compared to the size of your geometry.

PerspectiveViews

The Perspective Views section handles the appearance of the perspective view in both a normal and stereoscopic project.

→ Near Plane/Far Plane: These values set the nearest and 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 define the defaults when stereo is turned on in the 3D Viewer.

Fit to View

The Fit to View section has two value fields that manage how much empty space is left around objects in the Viewer when the F key is pressed.

→ Fit Selection: Fit Selection determines the empty space when one or more objects are selected and the F key is pressed.

→ Fit All: Fit All determines the empty space when you press F with no objects selected.

Default Lights

These three settings control the default light set up in the 3D Viewer.

The default ambient light is used when lighting is turned on and you have not added light to the scene. The default light moves with the camera, so if the default light is set to 'upper left,' the default light appears to come from the upper left side of the image/camera.

Defaults

The choices made here are used to determine Fusion's behavior when new tools are added to the flow and when controls are animated.
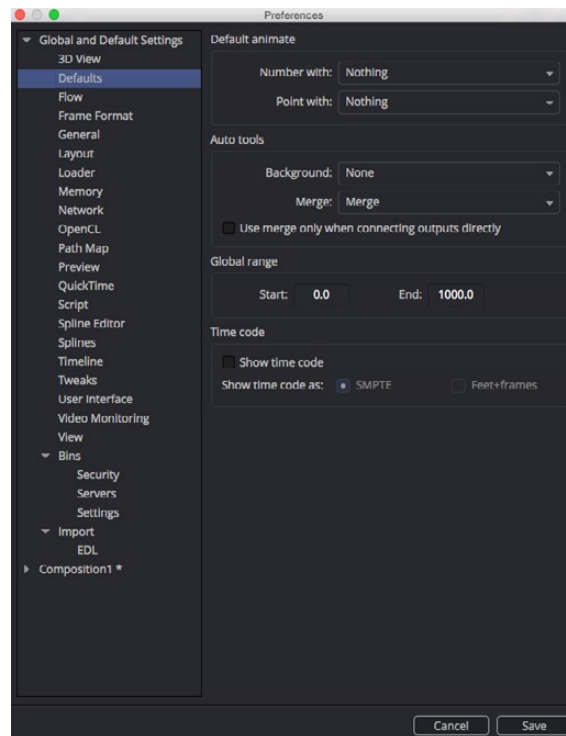
The Defaults preferences

Default Animate

The Default Animate section is used to change the type of modifier attached to a parameter when the Animate option is selected from its contextual menu. The default option is Nothing, which uses a Bezier spline to animate numeric controls and a path modifier for positional controls.

→ **Number with and Point with:** Drop down lists are used to select a different modifier for the new default. For example, change the default type used to animate positions by setting the Point with drop down menu to XY Path.

Choices shown in this menu come from installed modifiers that are valid for that type of parameter. These include third party plug-in modifiers, as well as native modifiers installed with Fusion.

Auto Tools

The Auto Tools section determines which 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 wide variety of available 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 Flow area uses a standard Merge. Other valid options for this are Anaglyph, Channel Booleans and Dissolve.

→ **Use Merge Only When Connecting Outputs Directly:** When this option is active, Merges are not automatically added when you drag multiple clips from the Finder or Windows Explorer onto the Fl ow area.

Global Range

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

Time Code

You use this option to determine if new compositions will default to showing SMPTE Time Code or frames (Feet + Frames) to represent time.

Flow

Many of the same options found in the Flow contextual menu, like settings for Tile picture, the Navigator and pipe style, are found in this category.
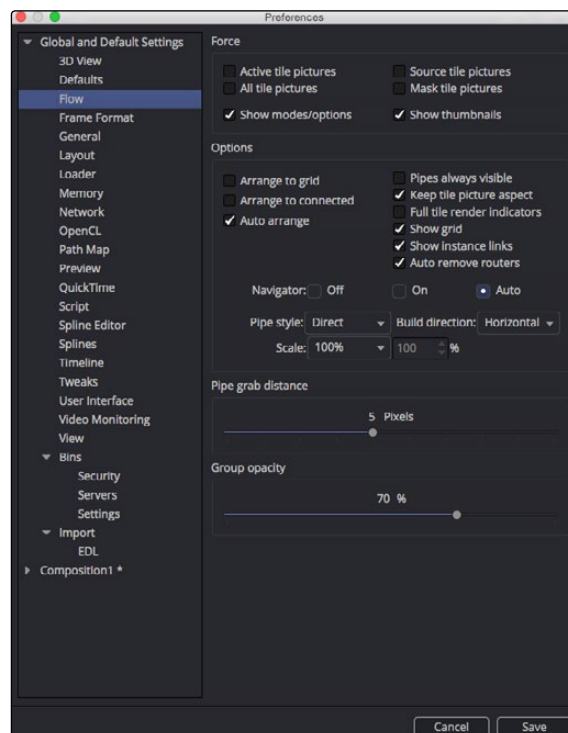
The Flow preferences

Force

The Force section can set the default to display pictures in certain tool tiles on the flow, rather than showing plane tiles. The Active checkbox sets pictures for the actively selected tool, the All checkbox enables pictures for all tiles, and the Source and Mask checkbox enables tile pictures for just Source and Mask tools.

When All is enabled, the picture shown will either be a thumbnail of the image rendered by the tool if the tool has rendered, or if the Show Thumbnails option is disabled, the tool's default icon is used. Concatenated transforms will also show a default icon.

→ Show Modes/Options: Enabling this option will display icons in the tool tile depicting various states, like Disk Caching or Locked.

→ Show Thumbnails: When this checkbox is selected, tool tiles set to show tile pictures will display the rendered output of the tool. When the checkbox is cleared, the default icon for the tool is used instead.

Options

The Options section has a number of settings that control or aid in the layout and alignment of tools in the flow.

→ Arrange to Grid: This enables a new flow'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 flow 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 Flow Editor's background grid.

→ Auto Remove Routers: Pipe Routers or "elbow nodes" in the flow are considered to be '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 connecting pipes between tools will be drawn over top of the tool tiles.

→ Keep Tile Picture Aspect: Enabling this option forces tool tile thumbnail pictures to preserve the aspect of the original image in the thumbnail.

→ Full Tile Render Indicators: Enabling this checkbox will cause the entire tile to change color when it is rendering. This can make it easier to identify which tools are rendering in a large composition. The coloring itself will form a progress bar to alert you to how close slower tools are to finishing their renders.

→ Show Instance Links: This option can be 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 your entire composition, which you can use to quickly navigate to different parts of the flow while you are zoomed in. The checkboxes in this section determine when the Navigator is displayed, if at all.

On: The Flow Navigator will always be visible.

Off: The Flow Navigator will always be hidden.

→ Auto: The Flow Navigator will only be visible when the Flow Editor's contents exceed the currently visible Flow 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 flow, Build Direction controls whether tools are organized horizontally or vertically.

→ Scale: The Scale menu allows you to select the default zoom level of the Flow Editor when a new flow is created.

Pipe Grab Distance

This Pipe Grab Distance slider allows you to choose how close the cursor must be (in pixels) to the pipes in the flow when selecting them.

Group Opacity

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

Frame Format

Frame Format preferences select the format settings for the interface. The settings in this preferences dialog are described in greater detail in the Frame Formats chapter of the manual.
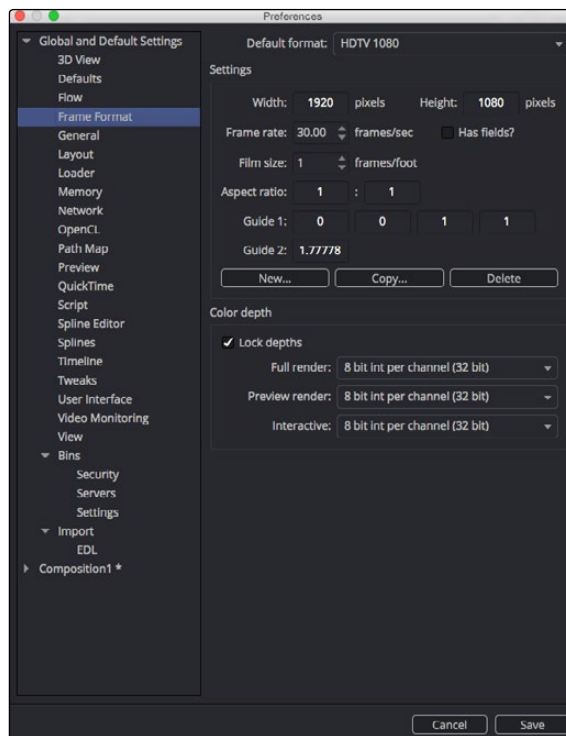
The Frame Format preferences

Default Format

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

Use the edit boxes to change any of the default settings. When creating a new setting, press the New button and enter a name for the setting in the dialog box that appears and enter the parameters.

Settings

The Settings section defines the format that is selected in the Default Format menu. You can modify an existing format or create a new one.

→ Width/Height: When creating a new format for the menu or modifying an existing menu item, you specify the Width or Height in pixels of the format using these fields.

→ Frame Rate: Enter or view the frames per second played by the format. This sets the default Frame Rate for previews and final renders from the Saver tool. It also sets the playback for the comp itself, as well as the frame to timecode conversion for tools with temporal inputs.

→ **Has Fields:** When this checkbox is enabled, any Creator or Loader tool added to the flow will be in Fields process mode.

→ **Film Size:** This field is used to define how many frames are found in one foot of film. The value is used to calculate the display of time code in Feet + Frames mode.

→ **Aspect Ratio:** These two fields set the pixel Aspect Ratio of the chosen frame format.

→ **Guide 1:** The four fields for Guide 1 define the left, top, right and bottom guide positions for the custom guides in the Viewer. To change the position of a guide, enter a value from 0 to 1. The bottom left corner is always 0/0, the top right corner is always 1/1. If the entered value's aspect does not conform to the frame format as defined by the Width and Height parameters, an additional guide is displayed on screen. The dotted line represents the image aspect centered about Guide 1's Center values.

→ **Guide 2:** This setting determines the image aspect ratio in respect to the entire frame format width and height. Values higher than 1 cause the height to decrease relative to the width. Values smaller than 1 cause height to increase relative to width.

→ **New Button:** You use the New button to create a new default setting in the drop down menu. Once you click the button you type a name for the setting in the dialog box that appears.

→ **Copy Button:** The Copy button copies the current setting to create a new one for customization.

→ **Delete Button:** The Delete button will remove the current setting from the default drop down list.

Color Depth

The three menus in the Color Depth section are used to select the color mode for processing preview renders, interactive renders and full (final) renders. Processing images at 8-bit is the lowest color depth and is sufficient for almost all video and multimedia tasks. 16-bit color has much higher color fidelity but uses more system resources. 16-bit and 32-bit float per channel uses even more system resources and is best for digital film and HDR rendered images.

Generally, these options are ignored by the composition unless a Loader or Creator tool's Color Depth control is set to Default.

General

The sections contained in the General preferences affect the behavior of the Control Panel as well as some other user interface elements.
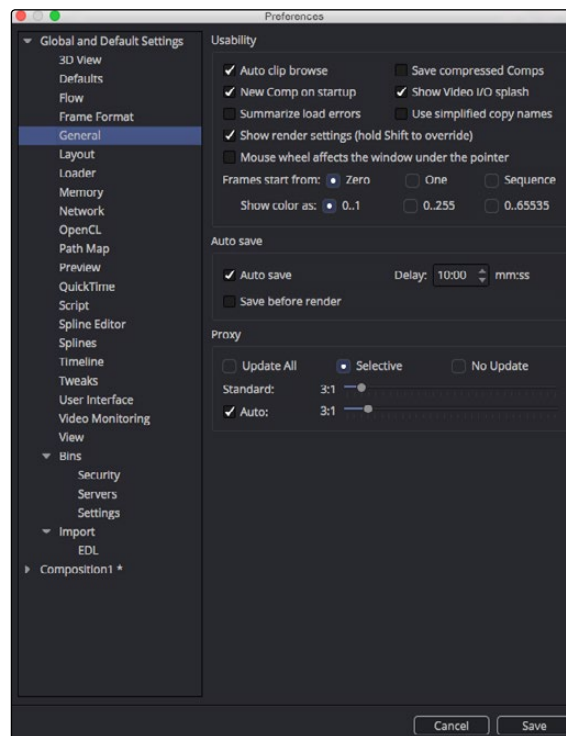
The General preferences

Usability

Usability has a number of Project, Flow and User Interface settings that can make the application easier to work with, depending on your workflow.

→ Auto Clip Browse: When this checkbox is enabled, the File Browser is automatically displayed when a new Loader or Saver is added to the flow.

→ New Comp on Startup: When checked, a new, empty flow is created each time the application is launched.

→ Summarize Load Errors: When loading 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 dialogue being presented for every missing tool.

→ **Save Compressed Comps:** This option enables the saving of compressed flows, rather than ASCII based text files. Compressed flows take up less space on disk, although they may take a moment longer to load. Flows containing complex spline animation and many paint strokes can grow into tens of megabytes when this option is disabled. However, compressed comps cannot be edited with a text editor unless saved again as uncompressed.

→ **Show Video I/O Splash:** This toggles whether the Splash image will be displayed over the video display hardware.

→ **Use Simplified Copy Names:** This option reduces the occurrence of underscores in tool names when copying.

→ **Show Render Settings:** When this checkbox is selected, the Render Settings dialog will be displayed every time a render is started. Holding Shift while starting a render will prevent the display of the dialog for that session, using whatever settings were applied during the last render. Disabling this option reverses this behavior.

→ **Mouse Wheel Affects the Window Under the Pointer:** Normally the Mouse wheel or Apple Mouse swiping works in the currently active window. With this option enabled it will work in the window underneath the cursor, so you don't have to click into a window first to make it active.

→ **Frames Start From:** This designates the starting frame number for clip times in the Loader and its Clip list.

→ **Show Color As:** This setting determines the numeric scale used to represent colors. The available options are Normalized (0 to 1), 8-bit (0 to 255) and 16-bit (0 to 65,535). This does not affect the actual processing or quality of the image, but it can make the mental math sometimes used to figure out adjustments a bit easier.

Controls

The Controls section includes settings for how the Control Panel is displayed.

→ **Auto Control Close:** When enabled, multiple tool headers will be visible, but only one tool header will be opened to display parameters. When disabled, any number of tools may be opened to display parameters at the same time. This setting has no effect unless the Auto Control Hide checkbox is disabled.

→ **Auto Control Hide:** When enabled, only the parameters for the currently active tool will be made visible. Otherwise, all tool headers will be visible and displayed based on the Auto Control Close setting.

→ **Auto Control Advance:** If the Auto Control Advanced checkbox is enabled, the Tab key and Return/Enter key will cause the keyboard focus to advance to the next edit box within a tool's control. When disabled, Return/Enter will cause the value entered to be accepted, but the keyboard focus will remain in the same edit box of the control. The Tab key can still be used to advance the keyboard focus.

→ **Show Controls for Selected:** When this option is disabled, only the active tool's parameters are shown tool in the Control Panel. 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 Control Panel, enabling this checkbox will show one color wheel with buttons to switch between the master, shadow, midtones and highlight channels. Otherwise four color wheels are displayed in the Control Panel.

→ **Grab Distance:** This slider ranges from 1 to 10 and defaults to 5. It designates the active area around the cursor and can be modified if you have difficulties in selecting points for modification in paths and spline curves. Smaller values represent more accurate selection.

Auto Save

When enabled, comps are automatically saved to a backup file at regular intervals defined by the Delay setting. If a backup file is found when attempting to open the comp, you are presented with the choice of loading either the backup or the original.

If the backup comp is opened from the location set in the Path Map preference, saving the backup will overwrite the original file. If the backup file is closed without saving, it is deleted without affecting the original file.

→ **Save Before Render:** When enabled, the comp is automatically saved before a preview or final render is started.

→ **Delay:** This preference is used to set the interval between Auto Saves. The interval is set using mm:ss notation, so entering 10 causes an Auto Save to occur every 10 seconds where as entering 10:00 causes and Auto Save every 10 minutes.

Proxy

→ **Update All, Selective, No Update:** The Update mode button is located in the Time Ruler. You can use this preference to determine the default mode for all new comps. Selective is the usual default. It renders only the 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 beside the Render button and selecting the desired value from the contextual menu

Layout

These controls are used to control the Layout, Size and Position of various elements of Fusion's interface at startup or when a comp is created.
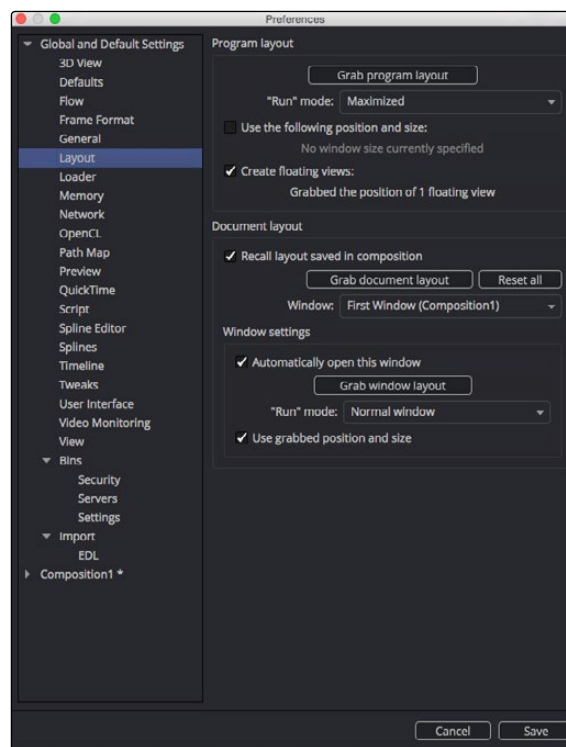
The Layout preferences

There are a lot of options, but in practice, you simply organize the interface the way you prefer it on startup and when a new composition is created, then open this preferences panel and click on the three buttons to grab the Program layout, the Document layout and the Window layout.

Program Layout

→ Grab Program Layout: Pressing this button stores the program's overall current position and size.

→ Run Mode: This menu is used to select the program's default mode at startup. You choose between a Maximized program, a Minimized program, or a Normal program display.

→ Use the Following Position and Size: When checked, the values stored when Grab Program Layout was selected will be used when starting Fusion.

→ Create Floating Views: When checked, the values of the floating display Viewers will be recorded when the Grab Program Layout button is used.

Document Layout

→ **Recall Layout Saved In Composition:** When checked, all Document layout settings in the controls below will be recalled when a saved composition is loaded.

→ **Grab Document Layout:** Pressing this button stores the entire interface setup, including all of the internal positions and sizes of working areas.

→ **Window:** When multiple windows on the same composition are used, this menu is used to select the window to which the Window settings will apply.

Window Settings

→ **Automatically Open This Window:** When checked, the selected window will automatically be opened for new flows.

→ **Grab Window Layout:** Pressing this button stores the size and position of the selected window.

→ **Run Mode:** Select the default run mode for the selected window. You choose between a Maximized window, a Minimized window, or a Normal window display.

→ **Use Grabbed Position and Size:** When checked, the selected window will be created using the stored position and size.

Loader

Using the Loader preference, you can set options for the Default Loader's depth and aspect ratio as well as define the local and network cache settings.
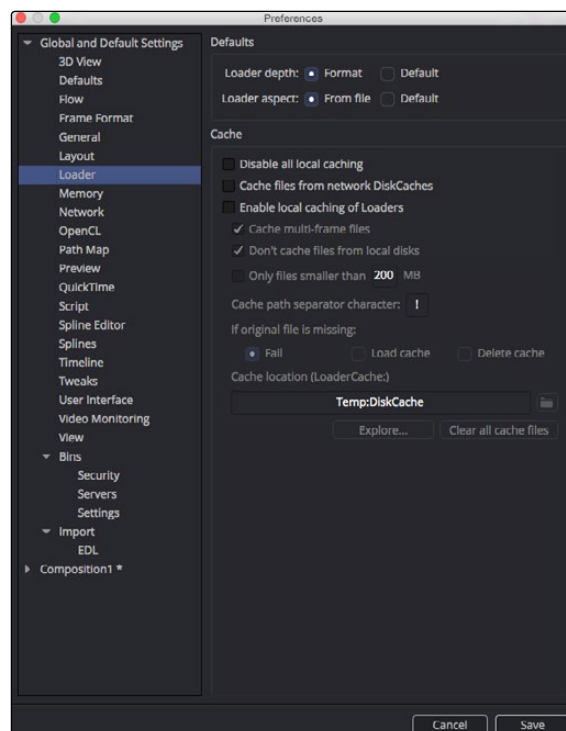
The Loader preferences

Defaults

The Defaults section includes two settings to determine how color depth and aspect ratio are handled for Loaders.

→ Loader Depth: The Loader Depth defines how color bit depth is handled when adding a Loader. Choosing Format means that the correct bit depth is automatically selected, depending on the file format and the information in the file's header. Choosing Default sets the bit depth to the value specified in the Frame Format preferences.

Cache

→ Disable All Local Caching: This setting disables local caching.

→ Cache Files from Network DiskCaches: If a tool has Disk Caching enabled, and the disk cache files are stored remotely on the network, then enabling this option will use a local copy of those cache files, similarly to the local cache on a networked Loader.

→ Enable Local Caching of Loaders: Files will be copied into the LoaderCache path set below or in the Path Maps preferences.

→ Cache Multi-Frame Files: Files like AVI or QuickTime will be copied into the LoaderCache path. This may take some time if the file is large.

→ Don't Cache Files from Local Disks: Files that do not sit on a network drive will not be copied into the LoaderCache path. You can disable this option if you have, for example, a fast SSD cache drive and want to use it for local files as well, to speed up file access while working interactively.

→ Only Files Smaller Than xxx MB.: Files larger than the value set here will not be copied into the LoaderCache path.

→ Cache Path Separator Character: When Enable Local Caching of Loaders is enabled, you can use this setting to re-build the path of the original files in LoaderCache. For instance, given the default "!" character, the original path X\Project\MyShots\Shot0815\ will be translated into X!Project!MyShots!Shot0815! in the LoaderCache path. Other separator characters may be used, including the "\" character, which will use subdirectories in LoaderCache: X\Project\ MyShots\Shot0815\.

→ If Original File is Missing: This setting provides three options to determine the caching behavior when the original files can't be found. The Fail option behaves exactly as the Default Loader in Fusion. The Loader will not process, which may cause the render to halt. The Load Cache option loads the cache even though no original file is present. The Delete Cache option clears missing files from the cache.

→ Cache Location: For convenience, this is a copy of the LoaderCache path set in the Path Maps preferences.

→ Explore: This button opens the LoaderCache path in the Mac OS X Finder window or a Windows Explorer window.

→ Clear All Cache Files: This button deletes all cached files present in the LoaderCache path.

Memory

Occasionally, it will be necessary to adjust the Memory Preferences in order to make the best use of available memory on the computer. For example, some people prefer a higher cache memory for faster interactive work, but for final renders the cache memory is often reduced, so there's more memory available for simultaneous processing of tools or multiple frames being rendered at once.
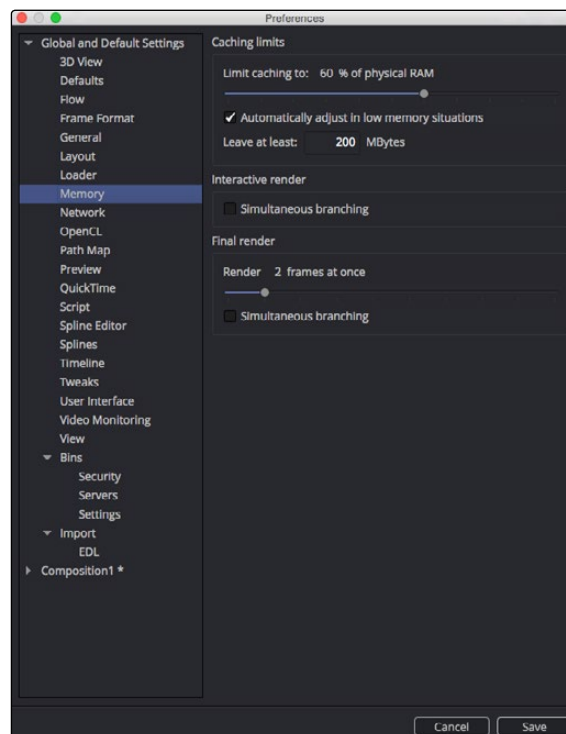
The Memory preferences

Caching Limits

→ **Limit Caching To:** This slider is used to set the percentage of available memory used for the interactive tool cache. Available memory refers to the amount of memory installed on the system, not the amount currently available.

When the interactive cache reaches the limit defined in this setting, it starts to remove lower priority frames in the cache to clear space for new frames.

→ **Automatically Adjust In Low Memory Situations:** This checkbox will set the caching to adjust when memory is low. The console will display any cache purges.

→ **Leave At Least X MBytes:** This setting is used to set the hard limit for memory usage. No matter what the setting of the Cache Limit, this setting determines the amount of physical memory available for use by other applications. Normally, this value should not be smaller than 25 MBytes.

Interactive Render

→ Simultaneous Branching: When checked, more than one tool will be processed at the same time. Disable this checkbox if you are running out of memory frequently.

Final Render

These settings apply to memory usage during a rendering session, either preview or final, with no effect during an interactive session.

→ Render Slider: This slider adjusts the number of frames that are rendered at the same time.

→ Simultaneous Branching: When checked, more than one branch of a flow will be rendered at the same time. If you are running low on memory, turn this off to increase rendering performance.

Network

These preferences are used to set up and control Fusion's network rendering. The majority of settings are found in the Render Manager dialog. These settings are only available in Fusion Studio.
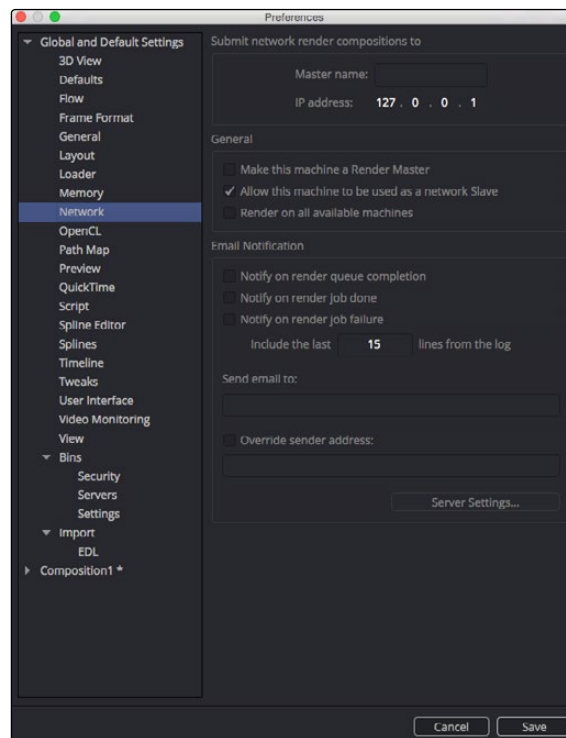
The Network preferences

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, this control may be pre-filled, but it may be uneditable. 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.

General

→ Make This Machine a Render Master: When enabled, Fusion will accept network render compositions from other computers and manage the render. It does not necessarily mean that this computer will be directly involved in the render, but it will submit the job to the render nodes that are 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. De-select it to prevent other people from submitting flows to this computer.

→ Render on All Available Machines: Enable this checkbox to ignore groups and priorities configured in the Render Manager. Compositions submitted from this computer for network rendering will always be assigned to every available slave.

Email Notification

→ Notify Options: These checkboxes cause emails to be sent when certain render events take place. The available events are Queue completion, Job done and Job failure.

→ Send Email to: Enter the address or addresses to which notifications should be sent. You separate multiple addresses with a semi colon.

→ 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 details on these settings and clustering, refer to the Network Rendering chapter for Fusion Studio.

OpenCL

The Open CL preference contains options for executing and debugging OpenCL devices and tools.
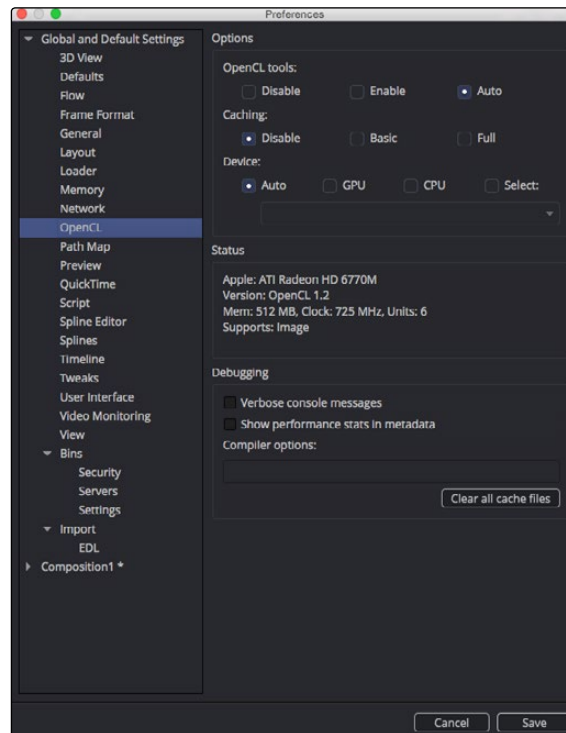
The OpenCL preferences

Options

→ **OpenCL tools:** This preference has three settings: Disable, Enable and Auto. When set to Disable, no OpenCL tools or plug-ins will be used. Fuses may still require OpenCL. If Enable is selected, OpenCL will be available for tools and plug-ins, if appropriate drivers are installed.

→ **Caching:** This preference enables write back cache for OpenCL images.

→ **Device:** The Device settings determine where the OpenCL code is executed. When GPU is selected all OpenCL code will be rendered on the graphics card, if drivers are available. The CPU setting renders OpenCL code on the CPU, if drivers are available.

→ The Auto setting gives priority to GPU rendering, however, if it is unavailable, the vendor's platform default will be used. Currently both the AMD and CPU options require either the AMD Catalyst 10.10 Accelerated Parallel Processing (APP) technology Edition driver, or the ATI Stream SDK 2.1, or later, to be installed. The Select setting allow you to chose the device explicitly. This is only available to Fusion Studio users.

Debugging

→ Verbose Console Messages: Enabling this option causes information to be shown in the Console. For example, Startup Logs, Compiler Warnings and Messages.

→ Performance Stats in Metadata: This setting adds some timing information to the output image's metadata. This can be seen with the Metadata view or accessed using the Metadata tools in Fusion.

→ Compiler Options: The Compiler Options allow specific command-line arguments to be given to the OpenCL compiler.

→ Clear Cache Files: Enabling this option deletes the contents of the AllDataOCLCache folder, where compiled OpenCL binaries are cached.

Path Maps

Path Maps are virtual paths used to replace segments of filepaths with variables. For example, define the path 'movie_x' as actually being in X\Shows\Movie_X. Using this example, Fusion would understand the path 'movie_x\scene_5\ scan.000.cin' as actually being X:\Shows\Movie_X\scene_5\scan.000.cin.
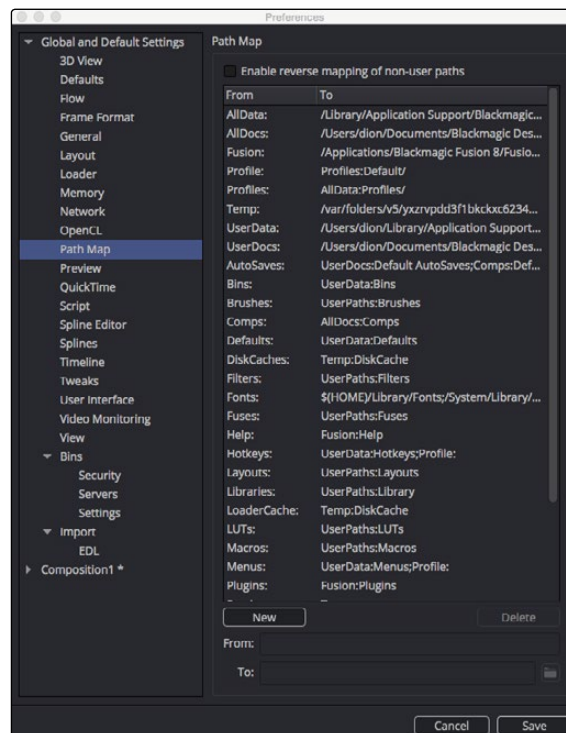
The PathMaps preferences

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

Also, path maps are used in render farms to bypass the different filename conventions.

→ **Built in Path Maps:** There are several built in path maps. Comp refers to the folder where the current composition is saved. Temp refers to the system's temporary folder. Fusion refers to the folder where Fusion is installed.

If the Reverse Pathmap option is enabled, Fusion will automatically attempt to replace portions of filenames that match known path maps. For example, a composition saved as C\Comps\example1.comp would automatically replace the path to c:\comps\example.jpg with Comp:\example.jpg.

→ **Global and Composition Path Maps:** Both Global and Composition preferences show a Path Maps preferences panel. The Global preferences are applied to all compositions, while Composition path maps are saved with the composition and only apply to it.

Composition path maps will override Global path maps with the same name. The built-in Comp.Path Map refers to the default Composition folder when used in a Global path map.

→ **Nesting Path Maps:** You can use a path map in a path map's definition, provided it has been defined first. For example, define a path map called 'Episode' that maps to x\shows\ Episode1. Then create path maps called Renders and Stills that mapped to Episode\ Renders_v1 and Episode\Stills_v1.

→ **Creating a Path Map:** To create a path map, click on the New button and enter the name of the path map in the From field below. Enter the value of the path map in the To: field.

→ **Deleting a Path Map:** To delete a path map, select it from the list and click on the delete button.

→ **Enable Reverse Mapping of Paths Preferences:** When this checkbox is selected, the built-in path maps for entries in the Paths preferences will be taken into account when applying mapping to existing filenames.

Preview

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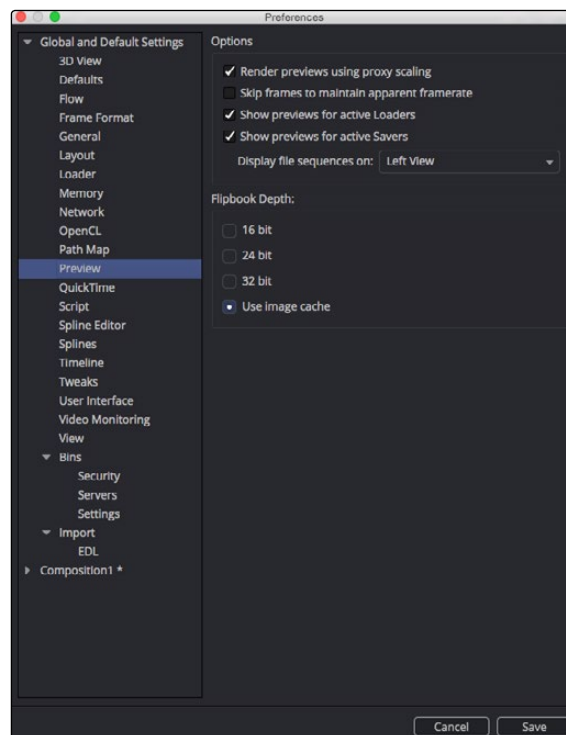
The Preview preferences

Options

→ Render Previews Using Proxy Scaling: When checked, this option scales the images down to the preview size for the Loader and Creator tools. This causes much faster rendering. If this control is disabled, frames will be rendered at full size, 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 if the preview playback controls below the Control Panel are shown when a Loader with a valid file is activated.

→ Show Previews for Active Savers: This setting determines if the preview playback controls below the Control Panel are shown when a Saver with a valid file is activated.

→ Display Files 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

QuickTime Video Compression Options
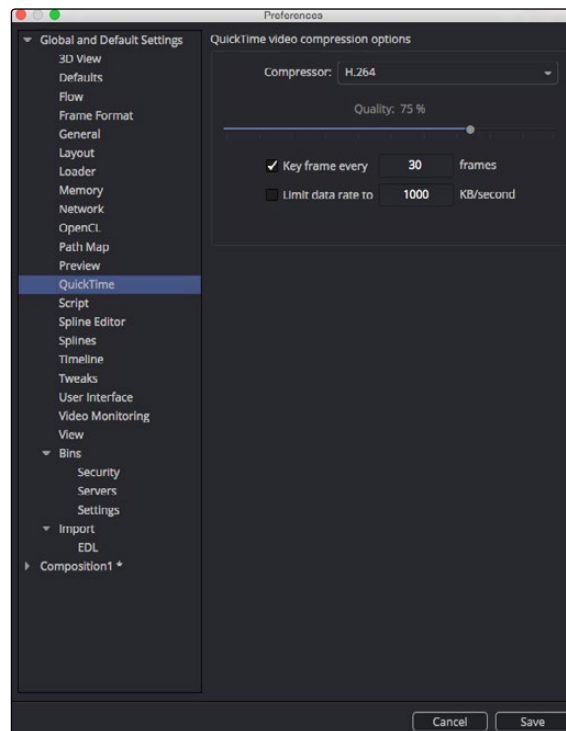
The QuickTime Preferences

Script

The preferences for Scripting include a field for passwords used to execute scripts from the command line and programs for use when editing scripts.
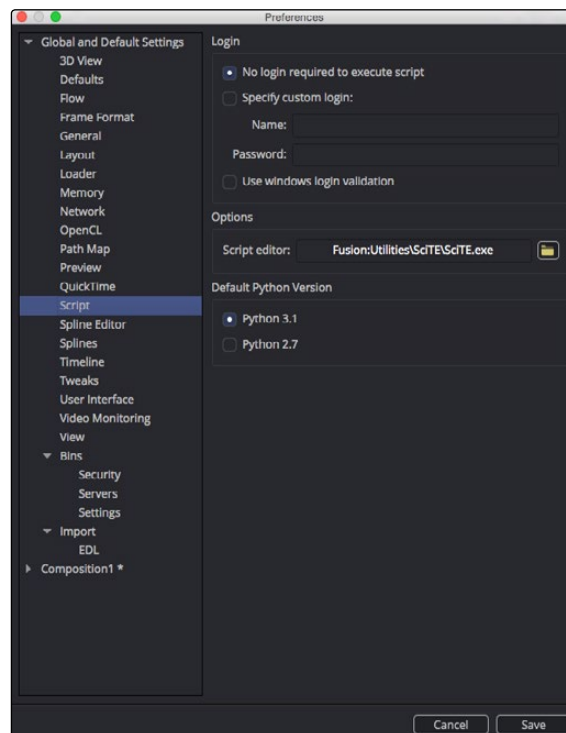
The Script preferences

Login

→ No Login Required to Execute Script: If a username and password is assigned, Fusion will refuse to process incoming external script commands (from FusionScript, for example), unless the Script first logs onto the workstation. This only affects scripts that are executed from the command line, or scripts that attempt to control remote copies of Fusion. Scripts executed from within the interface do not need to log in regardless of this setting. See the Scripting documentation for details.

Options

→ Script Editor: Use this preference to select an external editor for scripts (used when selecting Scripts > Edit from the menu).

Spline Editor

The Spline Editor preference allows you to set various spline options for Autosnap behavior, handles, markers and more.
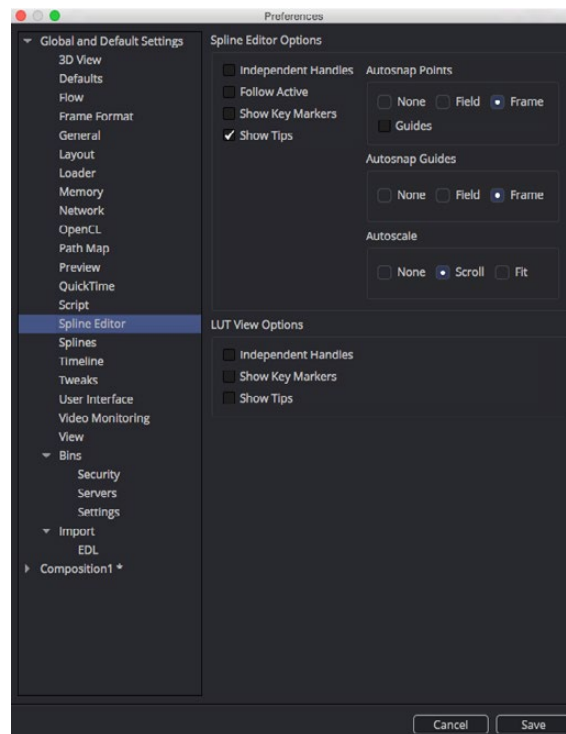
The Spine Editor preferences

Spline Editor Options

Independent Handles: Enabling this option allows the In or Out direction handle on newly created keyframes to be moved independently without affecting the other.

> → Follow Active: The Spline Editor focuses on the currently active tool.

> → Show Key Markers: Small colored triangles will be displayed at the top of the Spline Editor Time Ruler to indicate keyframes on active splines. The colors of the triangles match the colors of the splines.

> → Show Tips: Toggles if tooltips are displayed or not.

> → Autosnap points: When moving points in the Spline Editor, these will snap to the Fields or Frames or can be moved freely.

> → Guides: When moving points in the Spline Editor, these will snap to Guides as well.

> → Autosnap Guides: When moving or creating Guides these will snap to the Fields or Frames or can be moved freely.

→ **Autoscale:** Keeps the Spline View scales intact on changing the editable spline content of the view.

→ **Scroll:** Scrolls horizontally and vertically to show all or most of the spline points on changing the editable spline content of the view.

→ **Fit:** Zooms to fit all points within a Spline view, if necessary, on changing the editable spline content of the view.

LUT View Options

→ **Independent Handles:** Enabling this option allows the In or Out direction handle on newly created keyframes to be moved independently without affecting the other.

→ **Show Key Markers:** Small colored triangles will be displayed at the top of the Spline Editor Time Ruler to indicate keyframes on active splines. The colors of the triangles match the colors of the splines.

→ **Show Tips:** Toggles if tooltips are displayed or not.

Splines

Options for the handling and smoothing of animation splines, tracker path defaults and more are found in the Splines preferences.
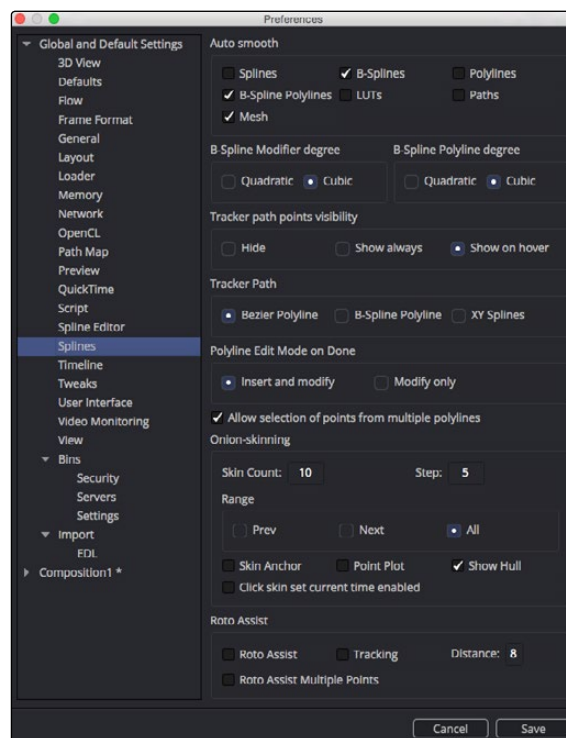
The Splines preferences

Autosmooth

Automatically smooths out any newly created points or keyframes on the splines selected in this section. This is valid for both animation splines as well as masks, LUTs, paths and meshes.

B-Spline Modifier Degree

This setting determines the degree used for interpolating B-Splines when used as modifiers. Cubic B-Splines determine a segment through 2 control points between the anchor points, and Quadratic B-Splines determine a segment through 1 control point between the anchor points

B-Spline Polyline Degree

This setting is like the one above, but applies to polylines, like those used for masks.

→ **Tracker Path Points Visibility:** This setting controls the control points on tracker paths. You have three options. You can show them, hide them or show them when your cursor hovers over the path, which is the default behavior.

→ **Tracker Path:** The default tracker creates a Bezier style spline paths. Two other options in this setting allow you to choose B-Spline or XY Spline paths.

→ **Polyline Edit Mode on Done:** This setting determines the Polyline tool that is selected in the Toolbar once you complete the drawing of a Polyline.

Onion Skinning

The Onion Skinning settings determine the number of frames displayed while rotoscoping, allowing you to preview and compare a range of frames. You can also adjust if the preview frames only from frame prior to the current frame, after the current frames, or split between the two.

Timeline

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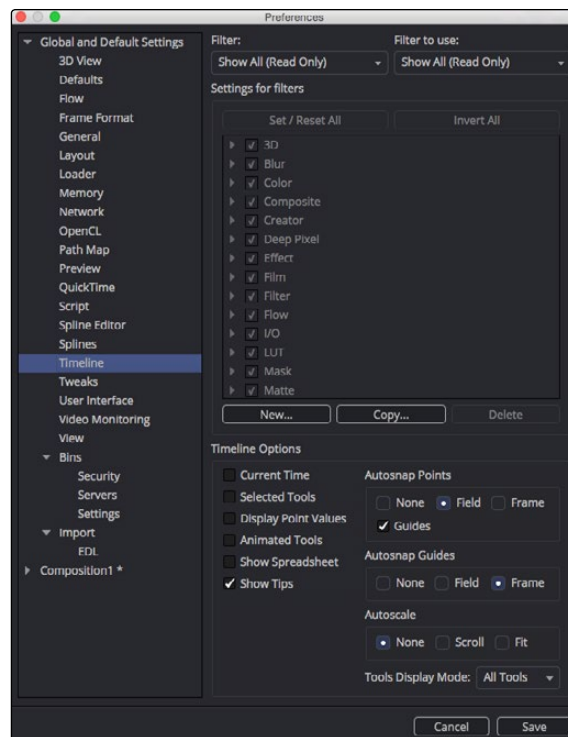
The Timeline preferences

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 actually sets the filter in the Timeline.

Settings for Filters

To create a new filter and define its settings, you first click the New button. Creating a new filter will allow you to select any of the tools that the filter will contain. The Settings for Filters area includes a list of all the tools. Only tools that are checked will appear in the Timeline or Spline Editor when the filter is selected. You can also create a copy of the filter using the Copy button or remove a filter from the list by clicking the Delete button.

Timeline Options

The Timeline Options configure which options in the Timeline are enabled by default. A series of checkboxes correspond to buttons located in the Timeline, allowing you to determine the states of those button at the time a new comp is created. For more information on the Timeline functions read Chapter 7, The Timeline Editor.

→ Autosnap points: When moving points in the Timeline Editor, the points will snap to the Fields, Frames or can be moved freely.

→ Guides: When moving points in the Timeline Editor, the point will snap to the Guides that are placed in the Timeline.

→ Autosnap Guides: When moving or creating Guides, the Guides will snap to the Fields, Frames or can be moved freely.

→ Autoscale: Keeps the Timeline scales intact while changing the editable spline content of the view. When set to scroll, the Timeline scrolls horizontally and vertically to show all or most of the spline points when changing the editable spline content of the view. When set to Fit, the Timeline zooms to fit all points within the view, if necessary.

→ Tools Display Mode: This menu controls the default sort order or the tools displayed in the Timeline. The default can be changed using the Sort order menu in the Timeline.

Tweaks

The Tweaks preferences handle a collection of settings for fine tuning I/O, memory and graphics hardware behavior.
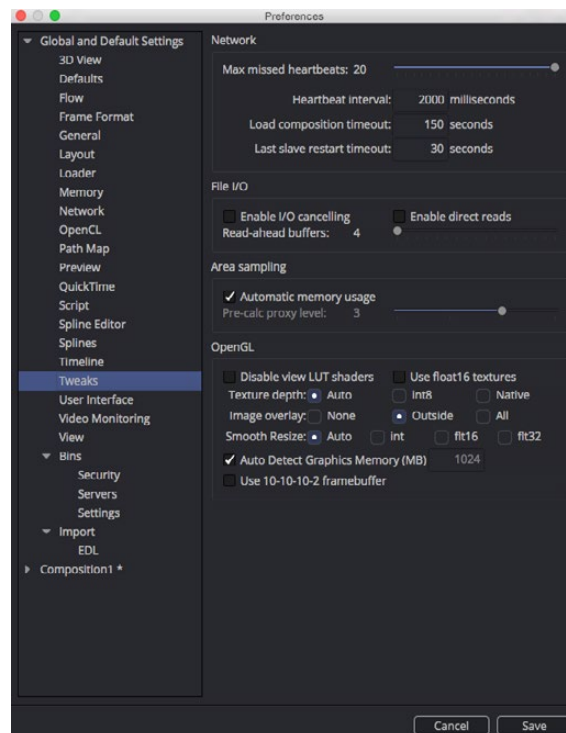
The Tweaks preferences

File I/O

> → I/O Canceling: This option enables a feature of the operating system that allows queued operations to be canceled when the function that requested them is stopped. This can improve the responsiveness, particularly when loading large images over a network.

> Enabling this option will specifically affect performance while loading and accessing TIFF, VideoPump and other formats that perform a large amount of seeking.

> This option has not been tested with every hardware and OS configuration, so it is recommended to enable it only after you have thoroughly tested your hardware and OS configuration using drive loads from both local disks and network shares.

> → Enable Direct Reads: Enabling this checkbox uses a more efficient method when loading a large chunk of contiguous data into memory by reducing I/O operations. Not every operating system employs this ability, so it may produce unknown behavior.

> → Read Ahead Buffers: This slider determines the number of 64K buffers that are use to read ahead in a file I/O operation. The more buffers, the more efficient loading frames from disk will be, but the less responsive it will be to changes that require disk access interactively.

Area Sampling

> → Automatic Memory Usage: This checkbox determines how Area Sampling uses available memory. Area Sampling is used for Merges and Transforms. When the checkbox is enabled (default), Fusion will detect available RAM when processing the 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 flows that deal with images larger than 4K, it may be desirable to override the automatic scaling and fix the proxy scale manually. This can preserve RAM for future operations.

> → Pre-Calc Proxy Level: De-selecting the Automatic Memory will enable the Pre-Calc Proxy Scale slider. Higher values will use less RAM, but take much longer to render.

Open GL

This section controls how Fusion makes use of your graphics card. Most settings may be left as they are, but since OpenGL hardware varies widely in capabilities and different driver revisions can sometimes introduce bugs, these tweaks can be useful if you are experiencing unwanted behavior.

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

→ **Use Float16 Textures:** If your graphics hardware supports 16-bit floating-point textures, activating this option will force int16 and float32 images to be uploaded to the Viewer as float16 instead, which may improve playback performance.

→ **Texture Depth:** Defines in what depth images are uploaded to the Viewer.

→ **Auto:** The Auto option (recommended) lets Fusion choose the best balance of performance and capability.

→ **int8:** Similar to the Use Float16 Textures switch, this option can be used to force images to be uploaded to the Display View as int8, which can be faster but gives less range for View LUT correction.

→ **Native:** The Native option uploads images at their native depth, so no conversion is done.

→ **Image Overlay:** The Image Overlay is a Viewer control used with Merge and Transform 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, this setting will render Viewers with 10 bits per primary accuracy, instead of 8 bits. Banding is greatly reduced when displaying 3D renders or images deeper than 8-bit. This is only available in Fusion Studio.

User Interface

These Preferences set the appearance of the user interface window and how the Control panel is displayed.
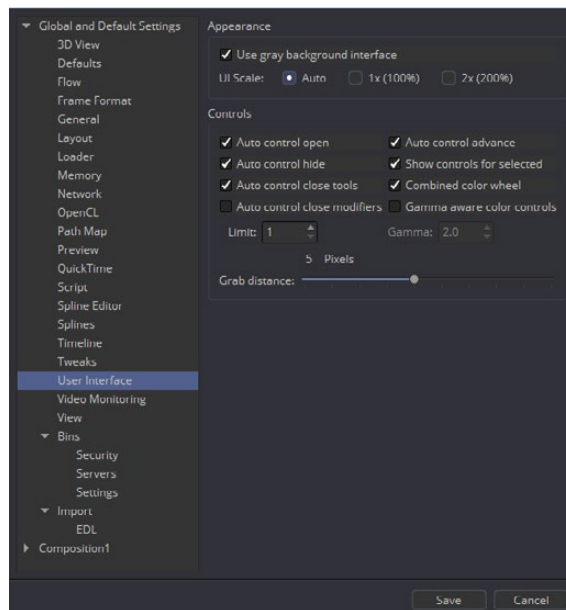
The User Interface preferences

View

The View preference is used to manage settings and default controls for Viewers.
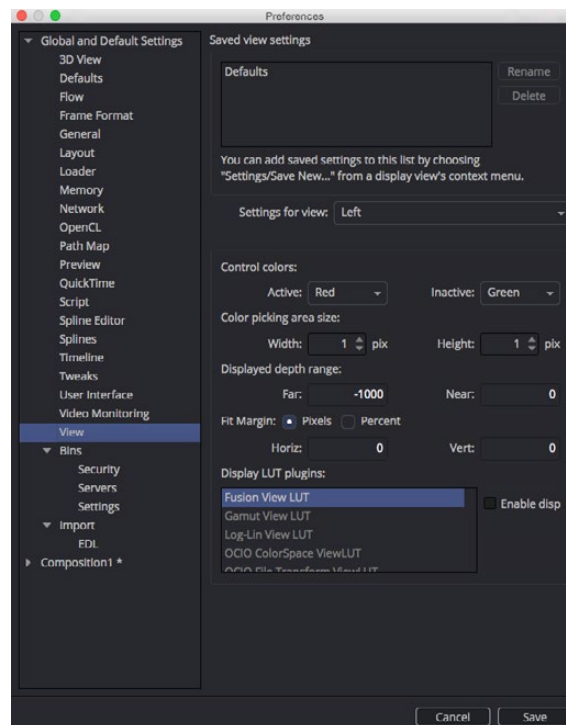
The View preferences

Saved View Settings

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

Settings for View

Each Viewer has its own preferences. The Settings for View drop down menu is used to select the Viewer for which you want to configure the settings.

Control Colors

The Control Colors setting allows you to determine the color of the Active/Inactive onscreen control.

Color Picking Area Size

You can use these width/height controls to set the number of pixels sampled when using the color picker in the Viewers.

Fit Margin

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

Displayed Depth Range

The Displayed Depth Range setting controls the view normalization of the Z-Channel.

Display LUT Plug-ins

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

Video Monitoring

The controls here are used to configure the settings of Blackmagic Video hardware.
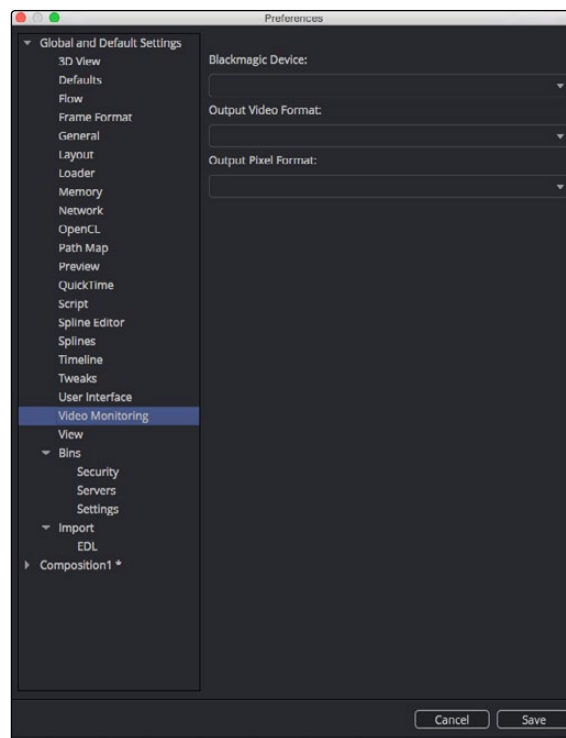
The Video Monitoring preferences

> → Compressor: This drop down list displays the QuickTime codecs available from your system. Fusion tests each codec when the program is started.
>
> 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.

→ 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.

→ 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.

Bins/Security

This preference is used to manage the Bin users and their permission. Bin Server and Remote Bins are only available in Fusion Studio.
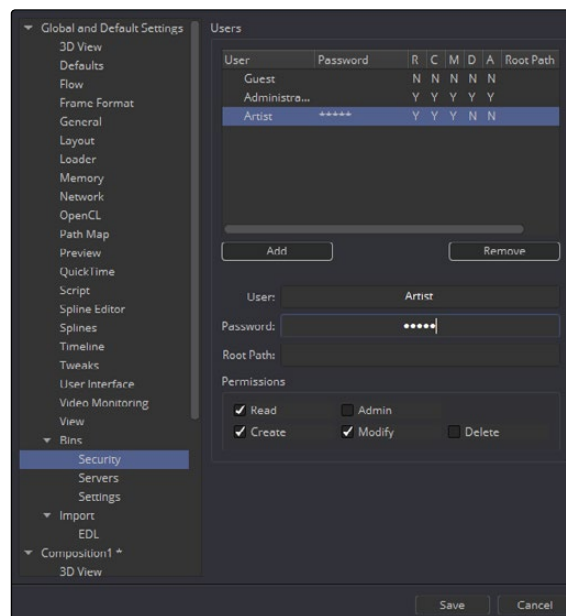
The Bins Security preferences

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 user name and password.

→ Remove: Click on this button to remove the selected entry.

User

This editable field shows the user name for the selected Bin Server item. If the user name is unknown, try Guest with no password.

Password

Use this field to enter the password for the Bin user entered in the Users list.

Permissions

The administrator can set up different permission types for users.

→ Read: This will allow the user to have read only permission for the bins.

→ Create: This will allow the user to create new bins.

→ Admin: This gives the user full control over the bins system.

→ Modify: This allows the user to modify existing bins.

→ Delete: This allows the user to remove bins.

Bins/Server

This preference is used to add Bin Servers to the list of bins Fusion will display in the Bins dialog. Bin Server and Remote Bins are only available in Fusion Studio.
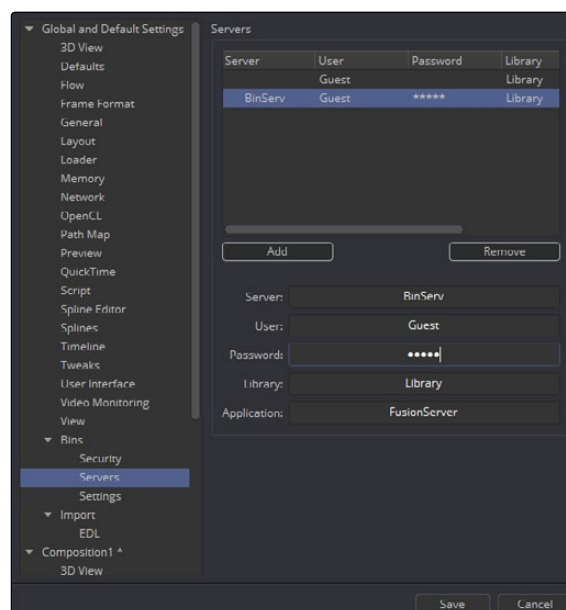
The Bin Server preferences

Servers

This dialog lists the servers that are currently in the connection list. You can select one of the entries to edit its settings.

→ Add: Use this button to add a new server to the list.

→ Remove: Click on this button to remove the selected entry.

Server

This editable field shows the name or IP address of the server for the selected entry in the list.

User

This editable dialog shows the username for the selected Bin Server item.

Password

Use this field to enter the password for the server entered in the Server list.

Library

The Library field lets you name the bins. So if you want to create a bin for individual projects you would name it in the Library field and each project would gets its own bin.

Application

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

Bins/Settings

This preference is used to control the default behavior of bins.
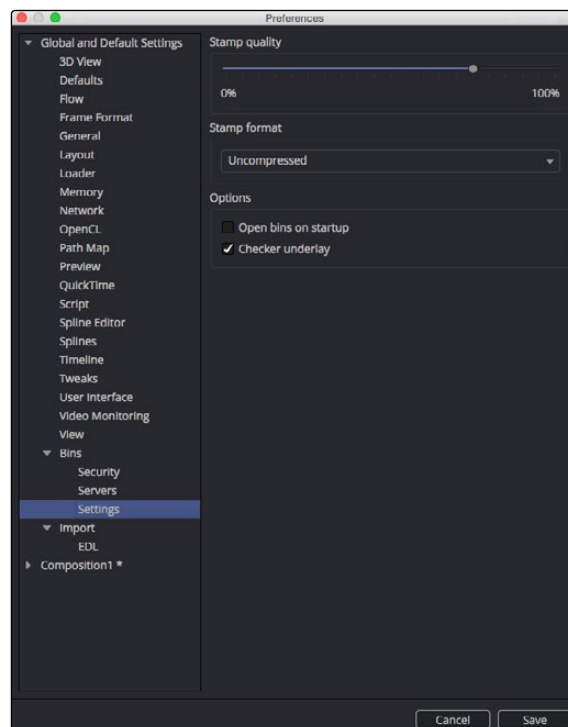
The Bins Settings preferences

Stamp Quality

The Stamp Quality is a percentage slider that determines the compression ratio used for Stamp file creation. Higher values offer better quality but take up more space.

Stamp Format

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

Options

→ Open bins on startup: When Open bins on startup is checked, the bins will open automatically when Fusion is launched.

→ Checker Underlay: When the Check Underlay is enabled, a checkerboard background is used for clips with alpha channels. When disabled, a gray background matching the Bin window is used as the clip's background.

EDL Import

The EDL Import options are used to determine how compositions are created from imported EDL files.
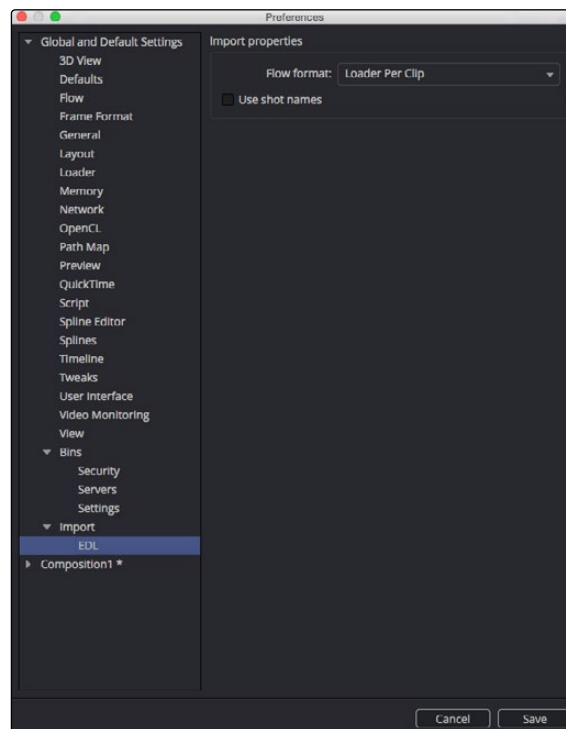
The EDL Import preferences

Flow Format

This drop down menu provides three options that determine how the flow is constructed for the imported EDL file.

→ Loader Per Clip: A Loader will be created for each clip in the EDL file.

→ A-B Roll: A Dissolve tool will be created with the settings imported from the EDL file.

→ Loader Per Transition: A Loader will be created with the settings representing the settings imported from the EDL file (Loader with a Clip list created).

Use Shot Names

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

Steroscopic and Optical Flow

17

Steroscopic and Optical Flow

Overview

Fusion Studio includes 3D stereoscopic and optical flow-based tools, which partially overlap in functionality.

Stereoscopic comp displayed in the Viewers

Stereoscopic Overview

All stereoscopic features are fully integrated into Fusion's 3D-environment. Stereoscopic images can be created using a single camera, which supports eye separation and convergence distance, and a Renderer 3D for the virtual left and right eye. It is also possible to combine two different cameras for a stereo camera rig.

Stereoscopic tools 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 tools can also be used for creating depth maps.

Optical Flow Overview

Optical Flow analyzes the motion in a clip and generates motion vectors between neighboring frames. It generates X and Y vectors from the previous frame to the current frame (Back Vectors) and to the next frame in sequence (Forward Vectors). Once calculated, Optical Flow data can be used to create smooth slow motion and variable retiming of clips, repair missing frames and even correct disparity in stereo 3D clips.

The OpticalFlow tool is used to analyze a clip and generate motion vector fields. These motion vectors are then used by additional tools to apply functions to images.

Toolset Overview

Optical Flow Tools

→ Optical Flow > OpticalFlow tool analyzes motion between neighboring frames in a sequence to generate motion vectors

→ Miscellaneous > TimeSpeed tool retimes a clip at a constant speed using Flow Interpolation mode

→ Miscellaneous > TimeStretcher tool retimes a clip at variable speeds using Flow Interpolation mode

→ Optical Flow > RepairFrame tool generates a new frame using the motion vectors between two neighboring frames

→ Optical Flow > SmoothMotion tool smoothes the color or aux channels using motion vectors

→ Optical Flow > Tween tool interpolates between two non-sequential images to generate a new frame

→ Color > CopyAux tool copies Aux channels including motion vectors, into RGBA more efficiently than Channel Booleans

Stereoscopic Tools

→ Stereo > Anaglyph tool for combing stereo images to create a single anaglyph image for viewing

→ Stereo > Combiner tool stacks a separate stereo images into a single stacked pair, so they can be processed together

→ Stereo > Disparity tool generates disparity between Left/Right images

→ Stereo > DisparityToZ tool converts disparity to Z depth

→ Stereo > Global Align tool shifts each stereo eye manually to do basic alignment of stereo images

→ Stereo > NewEye tool replaces left and/or right eye with interpolated eyes

→ Stereo > Splitter tool separates a stacked stereo image into to left and right images

→ Stereo > StereoAlign tool for adjusting vertical alignment, convergence, eye separation

→ Stereo > ZToDisparity tool converts Z depth to disparity

Working with Aux Deep Channels

Fusion images can contain channels other than RGBA color, called Aux Deep channels. Stereo Disparity and OpticalFlow deal directly with Auxiliary deep channels.

Below is a complete list of Aux channels supported in Fusion Studio:

→ RGBA: These are the standard colors

→ Z: The eyespace Z coordinate is almost always negative since in eyespace Fusion's camera sits at (0, 0, 0) looking down the Z-axis, Z values start at Z = 0 at the camera focal point and progressively become more negative for objects deeper in the scene

→ Coverage: The percentage of the pixel that is covered by the frontmost pixel, used for antialiased z-compositing

→ Object ID: These are user-assigned integers to meshes

→ Material ID: These are user-assigned integers to materials

→ Texture Coords: Normalized texture coordinates stored as (u, v) pairs

→ Normal Vector: Normal vector (nx, ny, nz) where the components are typically in the range [-1, +1]

→ Background Color: The color of the pixel if the frontmost layer were removed, used for antialiased z-compositing

→ Vector: The forward motion vector is an offset (vx, vy) that compares every pixel's position in one frame to the same pixel's position in the next frame

→ Back Vector: The backward motion vector is an offset (vx, vy) that compares every pixel's position in one frame to the same pixel's position in the previous frame

→ World Position: The position (wx, wy, wz) of the pixel in world coordinates

→ Disparity: An offset (dx, dy) that maps a pixel in the Left > Right or Right > Left frames

These extra channels are used by some Fusion tools. 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 tools (e.g., BrightnessContrast) have options on their common controls tab to limit their processing by MaterialID and ObjectID

→ The Fog and DepthBlur tools make use of the Z channel.

→ The Texture tool makes use of the TexCoord channel

→ The Shader tool makes use of the Normal channel.

There are a couple of ways to retrieve or generate those extra channels. For example:

→ The Renderer3D tool is capable of generating most of these channels.

→ The OpticalFlow tool generates the Vector and BackVector channels, then TimeStretcher and TimeSpeed can make use of these channels.

→ The Disparity tool generates the Disparity channels and then DisparityToZ, NewEye, and StereoAlign tools 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 renderfarm, and save results into an OpenEXR sequence. The Optical Flow toolset is designed around four types of tools that either generate, destroy, pass through or construct the motion vectors:

OpticalFlow
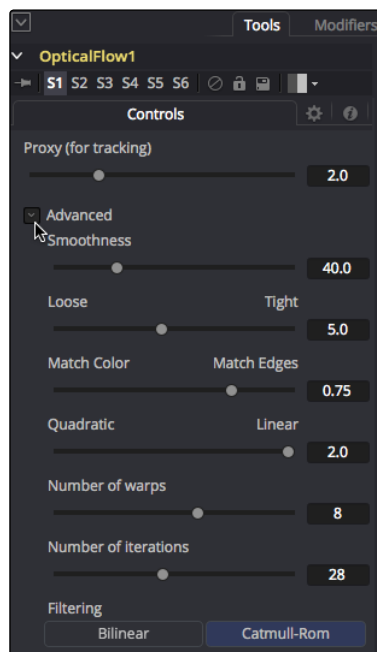
Optical Flow controls

The Optical Flow tool 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
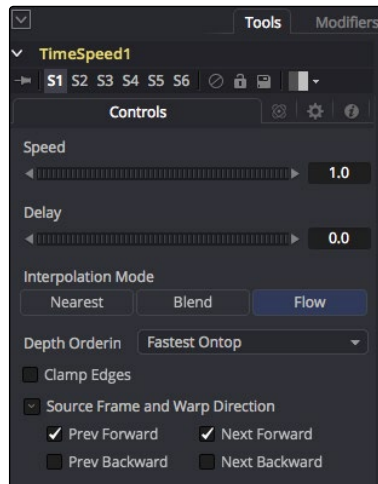
TimeSpeed Controls with Flow Mode

You can create smooth constant or variable slow motion effects using the TimeSpeed or TimeStertcher tool. When optical flow motion vectors are available in the Aux channel of an image, enabling Flow Mode in the TimeSpeed or TimeStretcher Interpolation settings will take advantage of the Vector and Back Vector channels. For the Flow Mode to work there must be either an upstream OpticalFlow tool generating the hidden channels or an OpenEXR Loader bringing these channels in. These tools use the Vector/BackVector data to do interpolation on the motion channel and then destroy the data on output since the input Vector/BackVector channels are invalid. For more information on TimeSpeed or TimeStretcher see the Tools Reference guide.

SmoothMotion
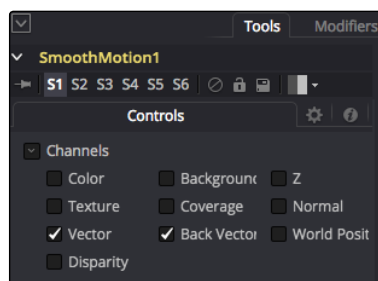
SmoothMotion Controls with Flow Mode

SmoothMotion can be used to smooth the Vector and BackVector channels or smooth the disparity in a stereo 3D clip. This tool passes through, modifies, or generates new aux channels, but does not destroy them. To learn more about SmoothMotion see the Tool Reference guide.

RepairFrame, Tween
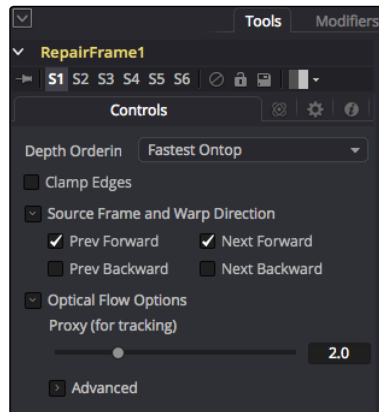
RepairFrame Controls

The Tween and Repair Frame tools are different from standard optical flow tools 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, etc. Because these tools 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 each frame, do their processing, and then destroy the flow channels. This being the case, these tools are computationally expensive. For more information on Tween or Repair Frame see the Tools Reference guide.

Stereoscopic workflows

Disparity is the difference between the left and right image. The Disparity map is used by tools to align and massage the stereo pair of images.
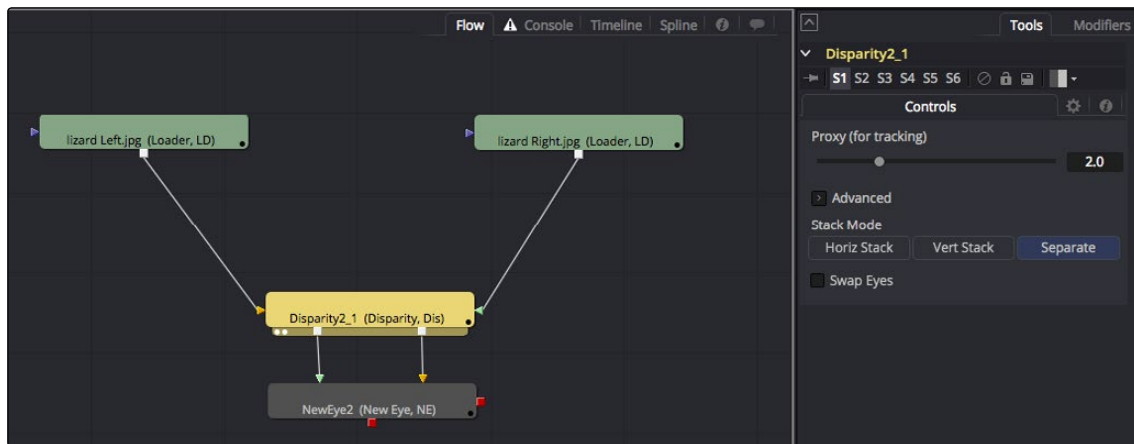
The Disparity tool analyzes a stereo pair of images and generates an X&Y disparity map.

The workflow is to load a left and right stereo image pair and process those in the Disparity tool. Once the Disparity map is generated, other tools can process the images.

Tip: When connecting Stereo paris in the Flow, make sure that the left and right images are connected to the left and right inputs of the Disparity tool.

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 renderfarm, and save them into an EXR sequence.

The toolset is designed around this philosophy.

Stereo Camera

There are two ways to setup a stereoscopic camera. The common way is to simply add a Camera 3D and adjust the eye separation and convergence distance parameters.
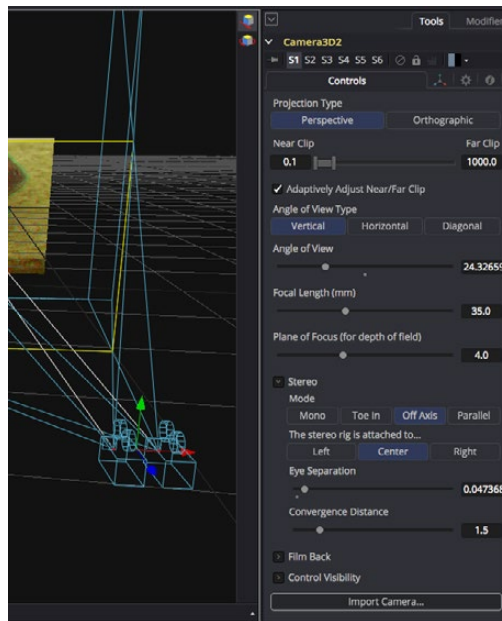
Stereoscopic cameras can be done with a
single camera or two connected cameras

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 tool, it is possible to assign different textures per eye.
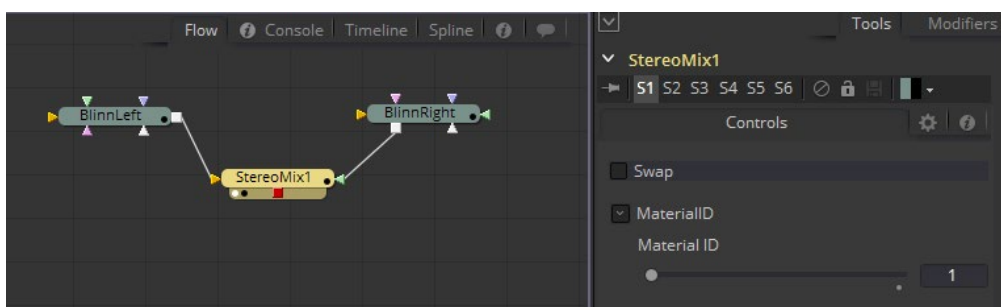
Material Viewer showing stereoscopic material

Disparity
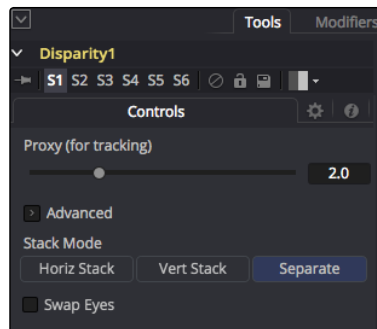
Disparity Controls

The Disparity tool 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
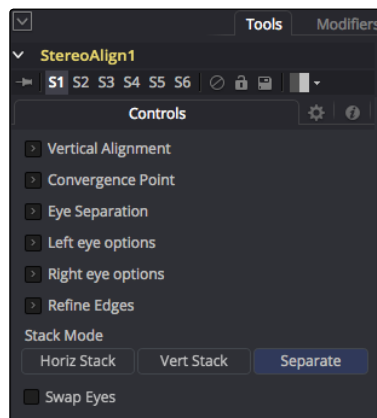
Stereo Align Controls

These use and destroy the Disparity channel to do interpolation on the color channel.

The hidden channels are destroyed in the process because, after the tools have been applied, the original Disparity channels would be invalid.

For these tools to work there must be either an upstream Disparity tool generating the hidden channels or an OpenEXR Loader bringing these channels in.

DisparityToZ, ZToDisparity
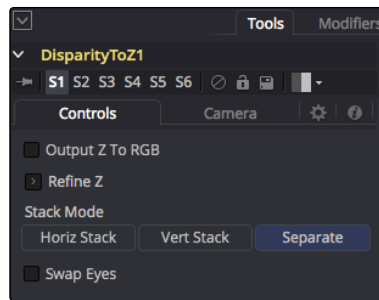
DisparityToZ Controls

These tools 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

Separate and Stack buttons

Stereo tools 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 flow for the left and right eyes. As a consequence, you will see Stereo tools 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 tools. In Stack mode the respective Right eye inputs and outputs are hidden.

Setting up Stereo in the Flow Editor

The disparity generation is the first operation to happen . This can be configured in the Flow editor in two different ways.
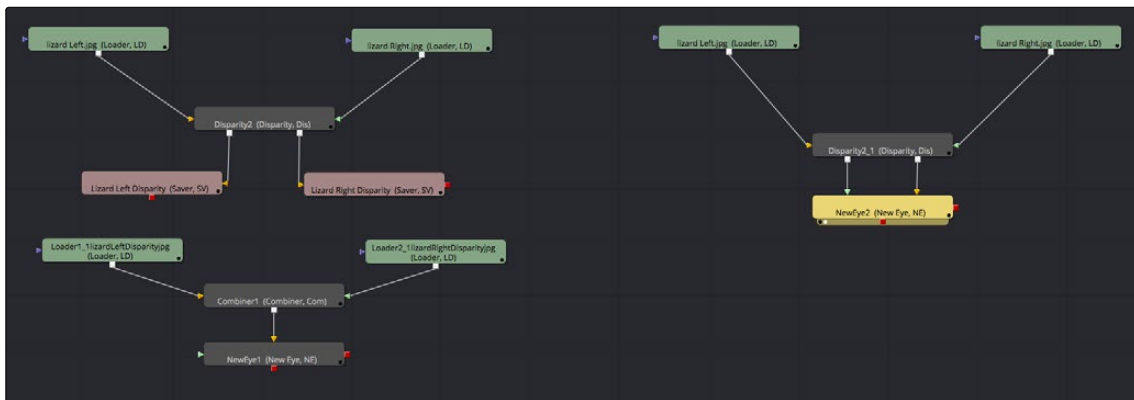
Two stereoscopic workflows.

In the above example the workflow on the right takes the left and right eye, generates the Disparity, 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 tools and used to create the NewEye images.

By using rendernodes to compute the disparity first, the later processing of the creative operations can be a much faster and interactive experience.

Although not shown in the above diagram, it is usually a good idea to color correct the right eye to be similar to the left eye before disparity generation as this helps with the disparity-tracking algorithm. The color matching does not need to be perfect, for example, it can be accomplished using the "Match" option in a Color Corrector's histogram options.

About the Disparity Channel

The Disparity channel stores the displacement vectors that match pixels in one eye to the other eye. The left image's Disparity channel will contain vectors that map left>right and the right image's Disparity channel will contain vectors that map right>left. For example:

```
(xleft, yleft) + (Dleft. x, Dleft. y) -> (xright, yright) (xright, yright) +
(Dright. x, Dright. y) -> (xleft, yleft)
```

You would expect for non-occluded pixels that Dleft = -Dright, although, due to the disparity generation algorithm, this is only an approximate equality.

Note that disparity stores both X and Y values because rarely are left/right images perfectly registered in Y, even when taken through a carefully setup camera rig.

Both Disparity and Optical Flow values are stored as un-normalized pixel shifts. In particular, note that this breaks from Fusion's resolution independent convention. After much consideration, this convention was chosen so the user wouldn't have to worry about rescaling the Disparity/Flow values when cropping an image or working out scale factors when importing/exporting these channels to other applications. Because the Flow and Disparity channels store things in pixel shifts, this can cause problems with Proxy and AutoProxy. The convention that Fusion follows is that, for proxied images, these channels store unscaled pixel shifts valid for the full sized image. So if you wish to access the disparity values in a script or via a probe, you need to remember to always scale them by (image. Width/image. OriginalWidth, image. Height/image. OriginalHeight).

Viewing of Disparity and Vector Channels

Aux channels can be displayed directly in the Viewers through the Channel toolbar button's pop up menu.
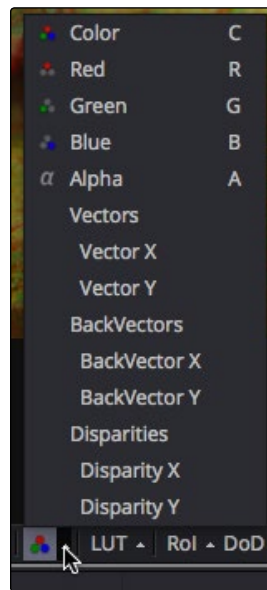
Viewer's Channel toolbar button pop up menu.

The CopyAux tool tis used to copy those channels directly into the RGB channels for viewing or further processing. The advantage of using the CopyAux tool 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 tool has the option to kill off aux channels and keep only the current RGB channels, freeing up valuable memory so you can cache more frames.

Tip: Although you can use the Channel Booleans to copy any aux channel into RGBA, it involves a few additional clicks when compared to CopyAux.

One thing to be aware of is that aux channels tend to consume a lot of memory. A float32 1080p image containing just RGBA uses about 32MB of memory but, with all the aux channels enabled, it consumes around 200MB of memory.

Stereo and Optical Flow Best Practices

How you create your composition, the images you are using and the type of shot you are working on can all have an impact on the success of the Disparity generation and Optical Flow analysis. Below we'll take a look at some of the situations to be aware of and how you can avoid some pitfalls when dealing with optical flow.

Semi-Transparent Objects

The optical flow and disparity generation algorithms Fusion uses assume there is only one layer per pixel when tracking pixels from frame to frame. In particular, transparent objects and motion blur will cause problems. For example, a shot flying through the clouds with the semi-transparent clouds in the foreground and a distant landscape background will confuse the of/stereo algorithms as they do not recognize overlapping objects with different motions. Usually the optical flow will end up tracking regions of one object or the other. If the transparent object and the background are near the same depth and consequently have the same disparity, then it is not a problem

Motion Blur

Motion blur is also a serious problem for the reason explained in the previous point. The Disparity and Optical Flow algorithms are unsure whether or not to assign a pixel in the motion blur to the moving object or the background pixel. Because the algorithms used are global in nature, not only the vectors on the motion blur will be wrong, but it will confuse the algorithm on regions close to the motion blur.

Depth of Field

Depth of field is also another problem related to the above two problems. The problem occurs when you have a defocused foreground object over a background object that is moving (OF case) or shifts between L/R (stereo disparity case). The blurred edges will confuse the tracking since they can't figure out the edges are actually two separate objects.

Where to Calculate Disparity and Optical Flow?

Where you chose to generate optical flow or disparity in your composition can drastically affect the results.

For example, if you have composited a lens flare in, it is better to compute OpticalFlow/Disparity before that, since the semi-transparent lens flare will confuse the tracking algorithms.

If you are color correcting the left/right eyes to match or for deflickering, it is better to apply the OpticalFlow/Disparity after since it will be easier for the tracking algorithm to find matches if the color matches between frames.

If you are removing lens distortion, think carefully about whether you want to do it before or after disparity computation. If you do it after, your disparity map will also act as a lens distortion map, combining the two effects as one.

As a general rule of thumb, it is best to use OpticalFlow/Disparity before any compositing operations except an initial color matching correction and a lens distortion removal.

Cropping the Source

As a general tip, if you are cropping your input images down for any reason, it is probably better to compute the optical flow or disparity before the crop and then afterward crop the flow/disparity along with the color.

The reason is that flow/disparity matching works well when there is common pixel data to match in both frames, but when there are pixels that show up in just one frame (or one eye), then the Disparity/OpticalFlow tools have to make a guess and fill in the data. The biggest occlusions going from L <–> R are usually pixels along the L/R edges of the images that get moved outside. This is similar for optical flow when you have a moving camera.

Another thing to be aware of are black borders around the edges of your frames, which you should crop away.

Tools with Multiple Outputs

Many of the stereo tools 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 tool to the view, it will always display the left output. There is no way to view the right output without connecting another tool like BC (BrightnessContrast) to the right output and viewing that.

Picking from Aux Channels

Some tools, like StereoAlign allow one to drag pick from the Z or Disparity auxiliary channels. You must pick from a tool upstream of the StereoAlign, not from the output of the StereoAlign. If you try to pick a disparity from the output of a StereoAlign tool, 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 tool 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 tools would be a user trying to pick a gradient color for a Background tool from a view showing the Background tool itself (you are trying to pick a color for a tool from its own output).

Another issue that you need to be aware of, is which eye you are picking. To avoid problems, it's a good idea to always pick from the left eye. The reason is that the Disparity channels for the left and right eyes are different and when you pick from a horizontal/vertical stereo stack, Fusion has no way of knowing whether you picked the disparity value from the left or right eye.

The above are not hard and fast rules; rather they are guidelines to prevent foot shootings. If you understood the above reasoning fully, you'll realize there are exceptions, like picking disparity from the left output of DisparityToZ and Z from the left/right output of ZToDisparity, where everything is okay.

Vector and Disparity Channels

The Vector and BackVector channels store the forward and reverse optical flow.

The Vector channel might be better named "forward vector" or "forward flow" since the name "Vector" to describe a channel is "not technically correct," as the more mathematically-inclined user might recognize that all of the channels except the scalar channels Z/ID are technically "vector" channels. A frames Vector aux channel will store the flow forward from the current frame to the next frame in the sequence and the BackVector aux channel will store the flow backward from the current frame to the previous frame. If either the previous or next frames do not exist (either not on disk or the global range of a loader does not allow OpticalFlow to access them), Fusion will fill the corresponding channels with zeros (transparent black).

The Disparity channel stores the displacement vectors that match pixels in one eye to the other eye. The left image's Disparity channel will contain vectors that map left>right and the right image's Disparity channel will contain vectors that map right>left.

For example:

```
(xleft, yleft) + (Dleft. x, Dleft. y) -> (xright, yright) (xright,
yright) + (Dright. x, Dright. y) -> (xleft, yleft)
```

You would expect for non-occluded pixels that Dleft = -Dright, although, due to the disparity generation algorithm, this is only an approximate equality. Note that disparity stores both X and Y values because rarely are left/right images perfectly registered in Y, even when taken through a carefully setup camera rig.

Disparity and Optical Flow values are stored as un-normalized pixel shifts. In particular, note that this breaks from Fusion's resolution independent convention. After much consideration, this convention was chosen so the user wouldn't have to worry about rescaling the Disparity/Flow values when cropping an image or working out scale factors when importing/exporting these channels to other applications. Because the Flow and Disparity channels store things in pixel shifts, this can cause problems with Proxy and AutoProxy. The convention that Fusion follows is that, for proxied images, these channels store unscaled pixel shifts valid for the full sized image. So if you wish to access the disparity values in a script or via a probe, you need to remember to always scale them by (image. Width/image. OriginalWidth, image. Height/image. OriginalHeight).

When using Vector and BackVector aux channels remember that all tools expect these aux channels to be filled with the flow between sequential frames.

More precisely, if you have sequence of 3 frames A, B, C then:

B	Vector will contain the flow B>C
B	BackVector will contain the flow B>A
A	Vector will contain the flow A>B
A	BackVector is written with zeros as there is no frame before A
C	Vector is written with zeros as there is no frame D to flow C>D
C	BackVector will contain the flow C>B

When working with these channels, it is the user's responsibility to follow these rules (or for clever users to abandon them). Tools like TimeStretcher will not function correctly since they still expect them to contain flow forward/back by 1 frame.

Note: Currently DoD/RoI is not supported for all Fusion tools.

Network Rendering and Licensing

18

Network Rendering and Licensing

Network Rendering

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

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

Render Node Installation (Windows)

1. Locate the Render Node installer in the Fusion Studio program folder:

   → C:\Program Files\Blackmagic Design\Fusion\

   Note: During Fusion Studio installation a different location may have been selected.

2. Copy the Render Node installer Blackmagic_Fusion_Render_Node_[version].msi, to each computer on the network that you want to perform rendering operations.

3. Install the Render Node.

Render Node Installation (OS X)

1. Locate the Render Node installer in the Fusion Studio Application folder:

   → Applications\Blackmagic Fusion 8\

   Note: During Fusion Studio installation a different location may have been selected.

2. Copy the Render Node installer Blackmagic_Fusion_Render_Node, to each computer on the network that you want to perform rendering operations.

3. Install the Render Node.

By default the Render Node application will be added to Startup folder. Each time you login the Render Node application will run automatically. To disable the Render Node application from starting up automatically, press the "Advanced" button in the setup dialog as shown below:
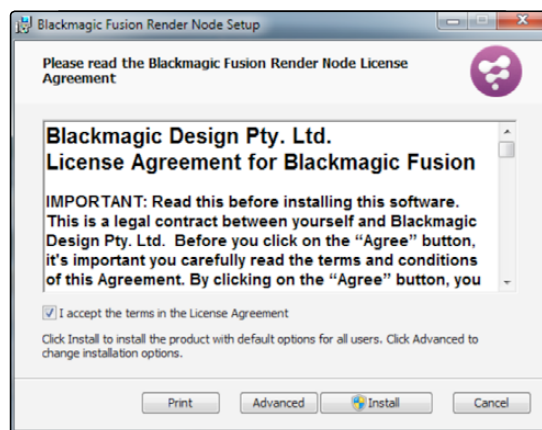
Then, uncheck the option "Automatically start at login". After the installation is complete, Render Node application is launched.
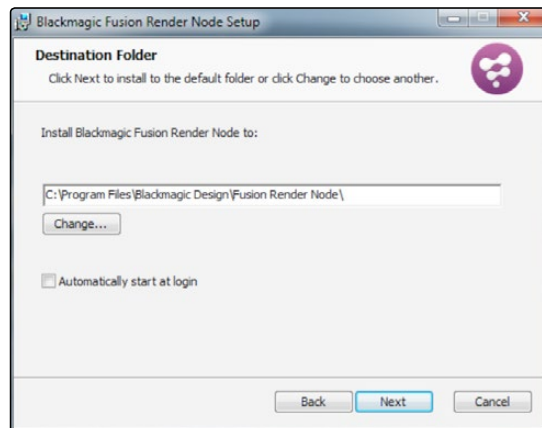
To disable the launching uncheck in the last page of the installer.
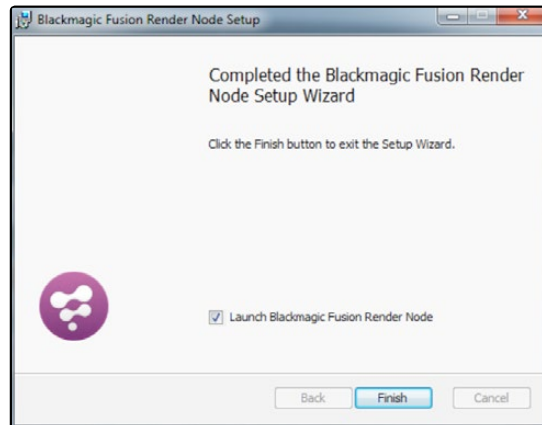
When Render Node is running, this icon is visible in the system tray.

Using a Third-Party Render Manager

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

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

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

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

Setting Up Network Rendering

What is the Render Master?

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

Any copy of Fusion or the Render Node can act as a Render Master. The rendermanager application can also be used as a standalone render master.

Run a standalone Render Manager to use a dedicated Render Manager without consuming a license of Fusion or the Render node.
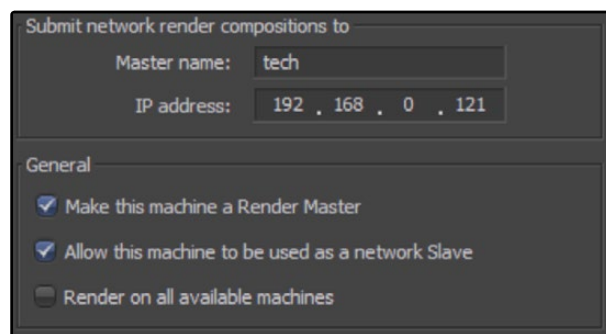
Rendermanager

Acting as a render master has no significant impact on render performance. The system resources consumed are insignificant. Many smaller facilities use one of the render nodes as a Render Master, while freelancers with just one license often use their own interactive license of Fusion as a Render Master. The larger and more complex a render farm gets, the more likely it is that a standalone Render Master will be required.

The standalone Render Manager is available at www.blackmagicdesign.com/support or in a folder on the original Fusion installation disk. To run the manager, simply copy the rendermanager application and FusionScript.dll into a directory and run rendermanager. It is usually a good idea to make the rendermanager a startup application for the chosen computer.

Setting Up the Render Master

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

For Fusion

Select Fusion > Preferences (OS X) or File > Preferences (Windows) and open the Global Network Preferences window.

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

For Render Node And Rendermanager

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

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

Preparing Render Slaves

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

For Fusion

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

For Render Node and Rendermanager

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

Choosing the Render Master Used By Slaves

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

The Render Manager

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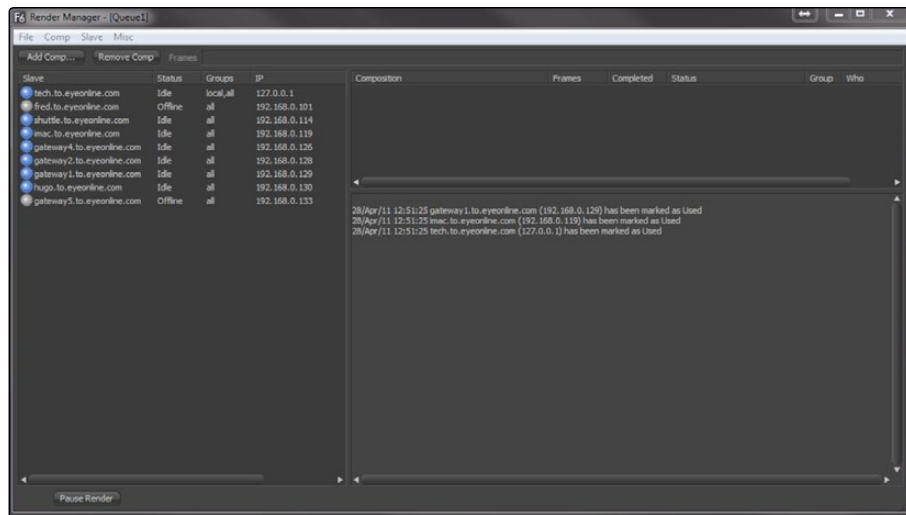
Opening the Render Manager Dialog

For Fusion

Select Render Manager from the File menu or press Command-M (OS X) or Ctrl-M (Windows) to open the Render Manager Dialog.

For Render node and rendermanager

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

Adding Computers to The Slave List

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

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

Scanning for Slaves

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

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

Manually Adding Slaves

To manually add a slave to the Slave List, select Add Slave from the Slave menu on Windows or from the Slave List contextual menu on OS X.

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

Removing Computers from the Slave List

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

Loading and Saving Slave Lists

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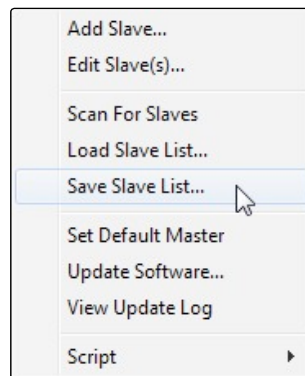
The Render Queue

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

Adding Compositions to the Queue

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

From the Render Manager

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

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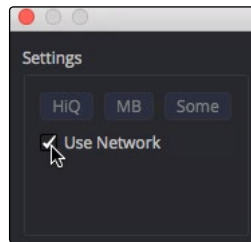
When starting a preview or final render, selecting the Use Network checkbox from the Render Settings dialog will add the composition to the end of the current queue in the Render Master's queue manager. The Render Master used is the one configured in the network preferences of the Fusion that is submitting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

From the File Explorer

Add one or more compositions to the queue by dragging the .comp file into the Render Manager's queuelist from the Finder (OS X) or File Explorer (Windows).

From the Command Line

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

Removing Compositions from the Queue

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

Saving and Loading Queues

It may be useful to save a Queuelist for later re-use. The current Queuelist is normally saved as queue.slv in the Fusion

→ Queue directory. To save the current queue with a new name, select Save Queue from the File menu. To load a queue from disk, select Load Queue from the File menu.

Re-Ordering the Queue

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

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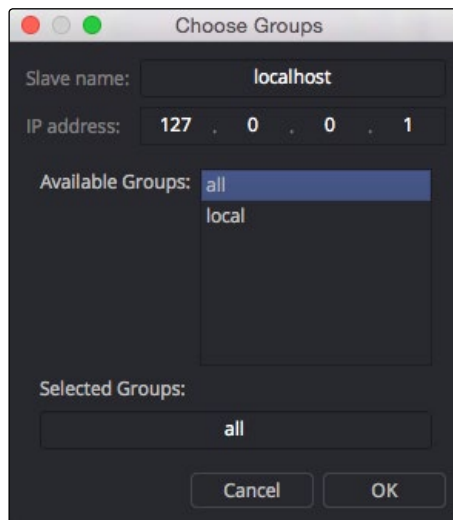
Slaves can be configured into Groups, which are then used when submitting compositions. For example, imagine that there are five Render Slaves. All of the slaves are members of the group All, but two of them have more memory than the other slaves so these two slaves are also added to a group called Hi-Mem.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

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

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

Assigning Slaves to Groups

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

Assigning Compositions to Groups

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

From Render Settings Dialog

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

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

From the Render Manager

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

Using Multiple Groups

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

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

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

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

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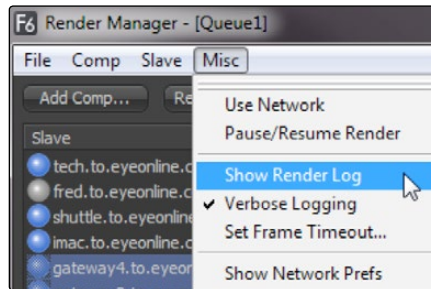
The Render Log is displayed beneath the Queuelist in the Render Manager. The text shown in this window displays the Render Manager's activities, including which frame is assigned to what slave, which slaves have loaded the compositions in the queue, together with statistics for each render after completion.

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

Preparing a composition for Network Rendering

Paths

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

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

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

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

Using UNC Names

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

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

Using Mapped Drives

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

Fonts

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

Plugins

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

Submitting a Composition

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

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

Flipbook Previews

Fusio Studio is able to use slaves to accelerate the production of Flipbook Previews, allowing for lightning fast previews. Frames for previews that are not network rendered are rendered directly into memory. Select the Use Network checkbox and the slaves will render the preview frames to the directory set in the Preferences Global>Path>Preview Renders. This directory should be valid for all

slaves that will participate in the network render. The default value is Temp\, which is a virtual path pointing to the system's default temp folder. This will need to be changed before network rendered previews can function.Once the preview render is completed, the frames that are produced by each slave are spooled into memory on the local workstation. As each frame is copied into memory, it is deleted from disk.

Disk Cache

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

Clustering

Clustering is a feature that allows the memory cache for a tool to be filled by remote workstations during an interactive session. When the Cluster mode is enabled on a tool, Fusion will attempt to maintain a cache for the tool in memory, using Remote Slaves on the network to assist with filling the cache.

By default, the clustering slaves will attempt to fill the cache by rendering the frames to either side of the current time. For example, if the current frame is 20, the slaves will fill the RAM cache for that tool between frames 10 and 30 in the background.

If a value for one of the controls on the clustered tool or one of its sources is changed, the cache will invalidate and the slaves will re-fill the cache.

Configuring Slaves for Clustering

Select the slaves to be used for clustering in the Render Master's Slave List and add them to a new group. Generally, the clustering group is called Cluster for clarity, although any group can be used for clustering. Once the group is created and slaves are assigned to that group, close the Render Manager on the master and open Preferences on the workstation from which the compositions will be clustered. Locate the Cluster panel in the Global preferences. Select the Slave group or groups to be used to render cluster jobs.
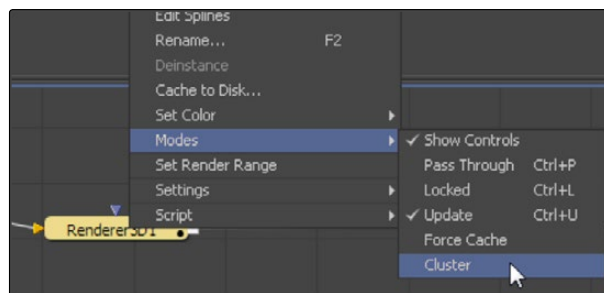
Select from the other options in this dialog to choose whether clustering will fill the cache for side frames around the current time or whether it will render the entire render range.

Enabling Clustering on a Tool

To enable clustering for a tool, right-click on the tool and select Modes>Cluster from the contextual menu that appears.
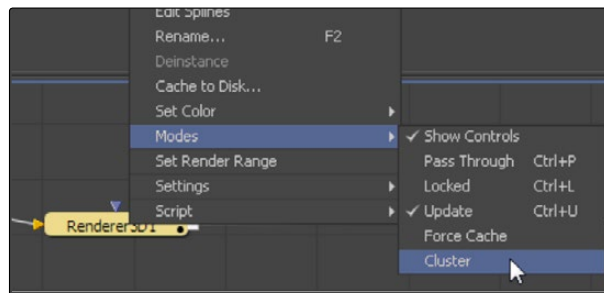
If the cluster does not start working immediately, change the current frame or modify a setting on the tool, or one of its source tools, to trigger the cluster for the first time.Once the cluster has been triggered, it will always work to keep the cache filled. The cluster will detect changes that affect the cache of the clustered tool and fill the cache automatically, whether or not that tool is currently selected or viewed.

Removing Clustering from a Tool

To stop clustering for a tool, right-click on the tool and select Modes>Cluster again from the contextual menu.

Remote Administration

Managing a large farm of Render Slaves can become a time-consuming task. Several capabilities are available from the Render Master to assist with Remote Administration of Render Slaves.

Installing Updates Remotely

It is possible to update one or more slaves from the Render Manager. To update remote slaves with a new version of Fusion, select the slaves to be updated and display the slave's contextual menu. Choose Update Software from the menu that appears and select the network directory that contains the files required for the update.

Before an update is applied remotely, create a directory with the update files. The update directory should be an uncompressed copy of a standard installation for either Fusion or the Render Node. Often, the easiest way to accomplish this is to download the archive version of the update and extract it into an empty directory.The update directory should be created in a directory accessible to all slaves via UNC or mapped drives paths, or the update will fail.
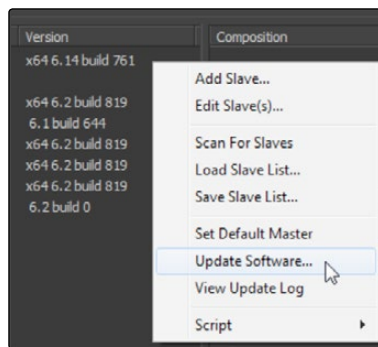
The master will instruct the remote slave to run the Overseer, application which must be installed in the Fusion directory of the remote machine. Overseer will shutdown the slave to prevent any errors with files in use. Once the slave has shutdown, the Overseer will copy the files from the update directory. The directory structure of the update directory will be duplicated on the slave. New files will be copied over and existing files will be overwritten. When the copy is complete, the Overseer will restart the slave. This will allow the slave to participate in network renders once more, completing the update.

If the update fails for any of the slaves, view an update log by right-clicking on the slave and selecting View Update Log from the contextual menu. The update log describes the history of updates for that node, with the most recent update appended to the end of the log file. Any errors that may have occurred will be detailed in the log. Updates do not have to be complete updates. Use an update directory to copy new plug-ins, scripts and macros to other systems on the network.

Notes: Fusion will not allow itself to be updated by update directories containing renderslave Nodes will not allow themselves to be updated with Fusion.

Read-only files in the update directory may not allow themselves to be copied over the network using this technique. Ensure that all files in the update directory are set with the read-only flag disabled.

The user who is logged into the Render Master workstation must have permission to write to each of the updated directories.

Setting the Render Master Remotely

Set the Render Master used by one or more slaves remotely, which can assist with remote administration. Open the Render Manager on the Master and select the slave(s) to be configured in the Slave List. Right-click and select Set Render Master from the contextual menu. Each selected slave will be set to use the current machine as their Render Master.

When Renders Fail

It is a fact of life that render queues occasionally fail. Software crashes, the power goes out or a computer is accidentally disconnected from the network. These are only a few of the things that can happen in a network. The more complex the topology of the network and the larger the compositions, the more likely some failure or another will interrupt the render in progress.If someone is monitoring the render as it progresses, this may not be too catastrophic. If no one is available to monitor the render, the risk that an entire queue may sit inactive for several hours may become a serious problem.

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

Automatic Rejoining of the Queue

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

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

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

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

FusionServer

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

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

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

Frame Timeout

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

The timeout is set per composition in the queue. To change the timeout value for a composition from the default of

60 minutes, right-click on the composition in the Render Manager's Queue list and select Set Frame Timeout from the contextual menu.
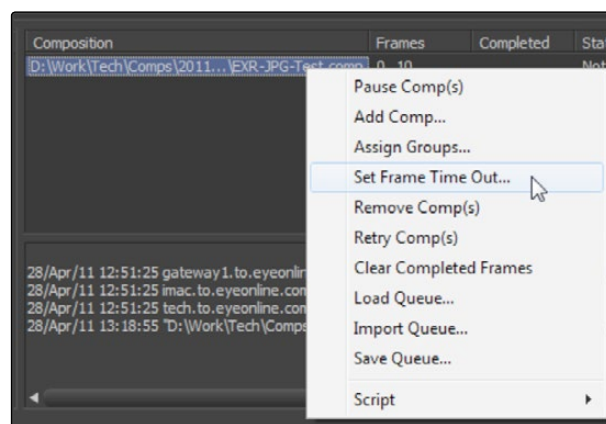
If it is expected to never have a single frame take longer than 10 minutes (for example, film or HD-sized footage is never rendered), reduce the timeout to 10 minutes instead. To change the frame timeout value, select Set Frame Timeout from the Render Manager's Misc menu.

Heartbeats

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

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

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

Managing Memory Use

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

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

Override Composition Settings

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

Render Several Frames at Once

Fusion has the ability to render multiple frames at once for increased render throughput. This slider controls how many frames are rendered simultaneously. The value displayed multiplies the memory usage (a setting of 3 requires three times as much memory as a setting of 1). Normal values are 2 or 3, although machines with a lot of memory may benefit from higher values, whereas machines with less memory may require the value to be 1.

Simultaneous Branching

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

Things You Need To Know

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

Time Stretching

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

Linear Tools

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

Saving to Multiframe Formats

Multiple machines cannot render a single AVI or QuickTime file due to the limitations of the file format itself. AVI and QuickTime movies are formats that are a single stream containing multiple frames. These frames must be written one after the other and not in the out-of-order sequence typically produced by network rendering. Always render to separate sequential file formats like EXR, DPX, JPEG and so on. Once the render is complete, a single workstation can load the image sequence in order and save to the desired compiled format.

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

Video Hardware

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

Troubleshooting

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

Check the Render Log

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

Check the Composition

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

No Slaves Could Be Found

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

The Composition Could Not Be Loaded

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

The Slave(s) Stop Responding

If a network link fails, or a slave goes down for some reason, the slave will be removed from the active list and its frames will be reassigned. If no more slaves are available, the composition will fail after a short delay (configurable in network preferences). If this happens, check the Render Log for clues as to which slaves failed and why.

The Slave(s) Failed to Render a Frame

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

Check the Slaves

Fusion's Render Manager incorporates a number of methods to ensure the reliability of network renders. Heartbeats are used to detect network or machine failures. In this event, failed slave's outstanding frames are reassigned to other slaves where possible. In rare cases, a slave may fail in a way that the heartbeat continues even though the slave is no longer processing. If a slave may have failed (although the Render Master may not have detected it) and you do not want to wait for the Frame Timeout, simply restart the Fusion workstation or Fusion Render Node that has hung. This triggers the heartbeat check, reassigns the frames on which that slave was working, and the render should continue. Heartbeats may fail if the system that is performing the render is making extremely heavy use of the Swap file or is spending an extraordinary amount of time waiting for images to become available over a badly lagged network. The solution is to provide the node with more RAM, adjust memory settings for that node or upgrade the network bandwidth.

Check the Network

At the Render Master, bring up the Network tab of the Preferences dialog box and click Scan. If a slave is not listed as running, the Render Master will not be able to contact it for network rendering. Alternatively, bring up a command prompt and ping the slaves manually. If the remote systems do not respond when they are up and running, the network is not functioning and should be examined further.

Fusion 8 Studio Network Licensing

Fusion 8 Studio introduces a simplified network licensing system via a dongle server setup.

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

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

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

Setup

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

Technical Info

→ Fusion Studio needs either a single-licence dongle on that machine or a multi-licence dongle with a spare licence on a machine anywhere on that subnet.

→ Fusion Studio can act as the license server if the dongle is plugged into that machine.

→ This works Cross Platform without restriction.

→ Render Nodes are unlimited and need a single or multi-licence dongle somewhere on the subnet.

→ Fusion Server is what serves multi-licence dongles so it has to be running (it is started automatically by Fusion).

→ By default, Fusion Server will stick around as long as Fusion is running locally or if a licence is being used by another machine (including Render Node). If nothing is using Fusion Server, it will quit in a few seconds.

→ "Fusion Server -S" (capital S) will stick around forever, until force-quit.

→ "Fusion Server -i" or "-u" will install/uninstall it as a service or daemon, running it on startup before user login.

→ Fusion Server also serves shared Bins, which can be connected to remotely.

→ Multiple dongles can be installed on a single machine and licences will add up.

→ Multiple licence servers can be present, and Studio will try all of them as needed.

→ The dongle, Fusion, and Render Node need to be on the same subnet for Auto License searching.

→ The FUSION_LICENSE_SERVER environment variable can now take a semicolon-separated list of IP addresses and host names to scan for license servers, example:
"bobs-mac.local;10.0.0.23;*;license.mystudio.com"
including * to indicate broadcast search of the local subnet.

→ Like most environment variables, you can put them in Fusion.prefs under Global.
EnvironmentVars:
fu:SetPrefs("Global.EnvironmentVars.FUSION_LICENSE_SERVER", "192.168.1.12;*")
fu:SavePrefs()

→ Multiple license servers (semicolon-separated) in FUSION_LICENSE_SERVER env var, like
fu:SetPrefs("Global.EnvironmentVars.FUSION_LICENSE_SERVER", "192.168.1.12;
192.168.10.55;*")

→ Removing a dongle or breaking the network will eventually (30-60secs) drop the licences of the Studio copies using them.

→ Upon losing its licence, Studio will start searching for another licence, locally or on a different machine.

→ If no licence can be found, Studio will pause rendering and display a modal dialog, with Retry/Exit options. Retry searches again. Exit will autosave the comp first.

→ Render Node just checks on startup for a dongle on the network but does not check afterward. This means it will not be affected by dongle removal or network issues.

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

→ TCP 1144

→ UDP 1144

→ TCP 49152-49159

→ UDP 49152-49159

Fusion Connect

19

Fusion Connect

Fusion Connect Overview

Fusion Connect is an AVX2 plugin for AVID Media Composer.

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

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

Fusion can be started automatically by the plugin if Fusion is installed on the same system, or it can be used on remote computers to modify the composition.

System Requirements

→ Supported Avid products: Media Composer 8.x

→ Supported product: Fusion 8. 1 or later

→ Memory recommendation: 8 GB minimum 16 GB or more recommended

→ Recommended Monitor Resolution: 1920 x 1080

→ Installation: Two files will be installed in your Media Composer:

                              Fusion Connect.avx

                              BlackmagicFusionConnect.lua

→ Avid's default directory: \Avid\AVX2_Plug-ins

The Effect Palette

Launching your Avid product, you will find Fusion Connect under the above effect palette category.

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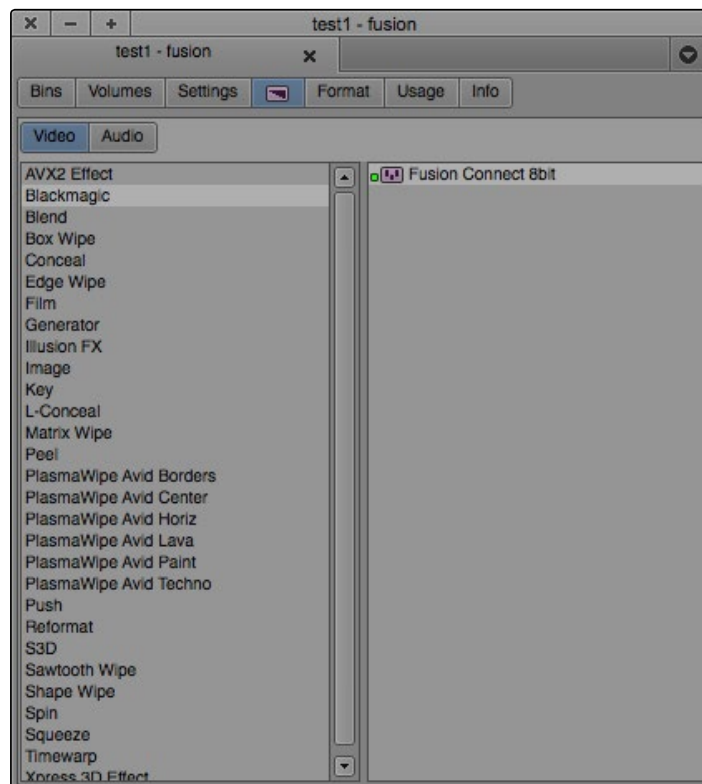
The Layer Input Dialog

When you apply the Fusion Connect AVX plug-in to a clip or a layer stack, you are presented with a layer AVX Optional Input dialogue box.
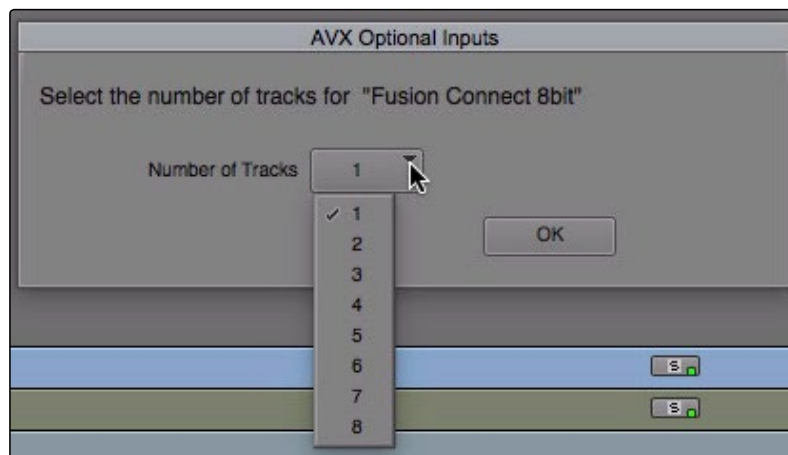
Once the layer count is selected, Fusion Connect will be applied to the timeline.

→ Select the layer count to the equal number of layers you want to ingest into Fusion.

→ Filler can be used as a layer.

→ Fusion Connect will allow a maximum of 8 layers.

Applying Fusion Connect to a transition point

If you apply Fusion Connect to a transition point, you will not receive any dialogue box, and the AVX plug-in will simply be applied to the timeline transition point.
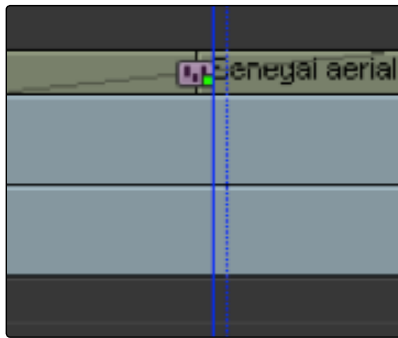
You can use the Avid dialogue boxes or smart tools to adjust the length and offset of the transition to Start, Center, End, or Custom.

The Effect Editor

Export Clips

Pressing this button, Fusion Connect will export all the associated clip(s) as sequences of frames, to provide access to them in Fusion.

Fusion Connect is always live.

Fusion Connect will save source frames each time images are displayed, scrubbed, or played back from the timeline. Depending on your Media Composer timeline settings, these interactively exported images might be half-resolution only based on Avid proxy settings.

When scrubbing around the timeline, only a few frames, namely those that are fully displayed during scrubbing, might be written to disk.

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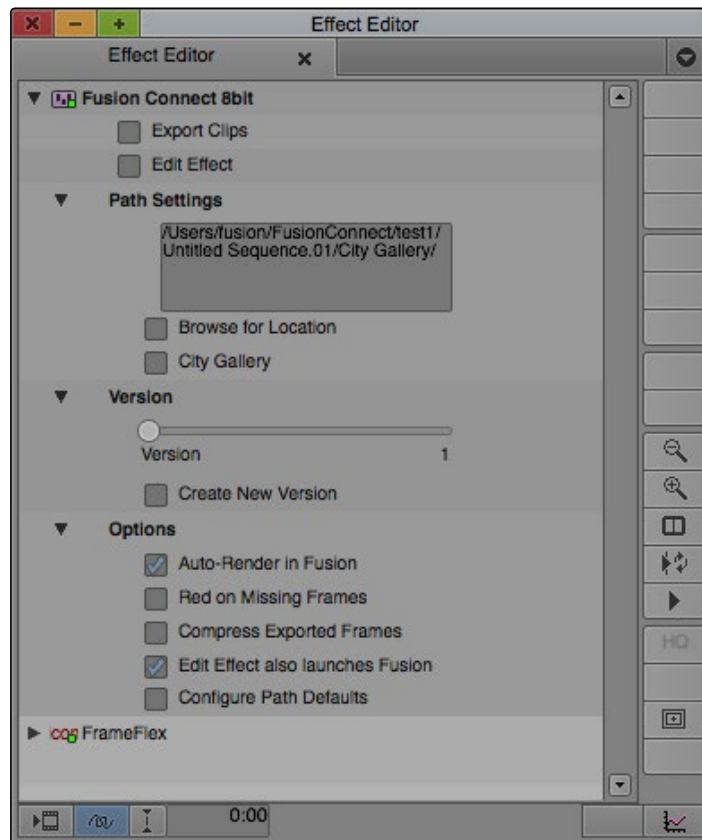
Edit Effect

This button performs three subsequent functions.

1. Creates a Fusion composition, with Loaders, Savers, and Merges (for layers) or a Dissolve (for transitions). This function is only performed the first time a comp is created when the Fusion Connect AVX2 plug-in is applied.

2. Launch Fusion (if installed on the machine), and if it is not already launched.

3. Open the Fusion comp associated with created effects.

Browse for Location

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

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

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

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

Auto Render in Fusion

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

Note this method of rendering has limitations, and can be significantly slower than rendering directly in Fusion. It is mostly used for batch rendering on the Avid timeline. Auto Render will also export the necessary media without having to manually execute the Export Clips function first.

Red on Missing Frames

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

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

Rename Shot

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

Compress Exported Frames

When enabled, this button will allow for Fusion RAW files with compression; both exported and rendered frames. This provides disk space saving because the files sizes are smaller. As with any other compression algorithm, the compression adds time to the write process of the file sequence.

Edit Effect also launches Fusion

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

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

Versioning

Create New Version

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

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

Version

This slider allows the editor to choose which version of the comp is used to populate in the timeline. It can be used to interactively switch from one version to the other in order to compare the results.

About RAW images

Fusion Connect creates a RAW file image sequence for both directions between Avid editing products and Fusion, in order to preserve all the image information. This benefits the editor in numerous ways because the images are not sitting in RAM.

The benefits include:

→ The ability to continue the editing process while an effect of treatment is rendering

→ The ability to take advantage of render slaves

→ The ability to retime footage

→ The ability to run Fusion remotely

About Color Depth

Fusion Connect derives its images directly from the RGB data within Avid MXF files. This allows the images to be codec agnostic. All RAW files from Avid that begin as 8 or 10 bit images are remapped to 16 bit float in Fusion. Rendered results from Fusion are processed in 16bit float to maintain the full color fidelity supported by Media Composer.

Export Clips Resolution Warning Overview

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

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

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

Click on the green/yellow split button or solid yellow toggle button at the bottom of the timeline.

When the button is solid green, click on the export button in the Effect Editor once again.
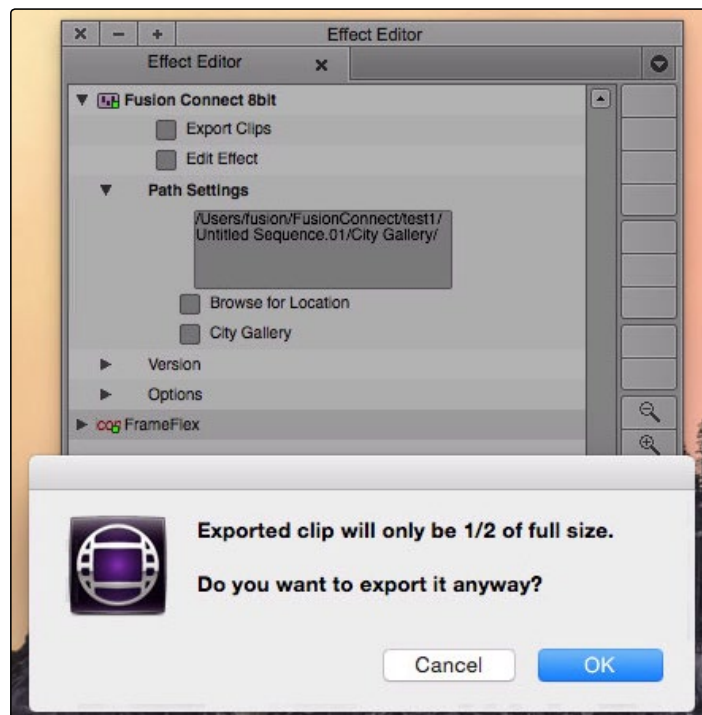
Manual vs Auto-Render

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

Manual Workflow

① Add Connection effect

---

② Click Export Clip

Creates Fusion RAW files at the given location on the media drive. These files are not deleted automatically, and manual cleanup of old files will be rquired to avoid running out of disk space.

---

③ Click Edit Effect

Create a comp file when first clicked, but will not overwrite this file when clicked again. Attempts to launch Fusion and load the comp. If Fusion is not installed locally, the comp can be accessed manually from a different machine via the network.

---

④ Edit comp in Fusion

Work your magic here.

---

⑤ Save comp in Fusion

This step is important!
Do not forget to save the comp.

---

⑥ Render comp in Fusion

You can continue working in AVID while Fusion render. Your timeline will update automatically with any frame that is succesfully rendered.

---

⑦ Render clip in AVID

Optional step, but recommended.

Auto-Render Workflow

① Add Connection effect

Ensure thatAuto-Render is ticked.

---

② Click Edit Effect

Creates Fusion RAW files as Export Clip would do. Also creates a comp file when first clicked, but will not overwrite this file when clicked again. Launches Fusion and loads the comp.

---

③ Edit comp in Fusion

Work your magic here.

---

④ Save comp in Fusion

This step is important!
Do not forget to save the comp.

---

⑤ Render clip in AVID

In this case, both the Fusion-comp and the AVID clip will be rendered simultaneously. If a full-size rendered frame is not found, the full size/depth source frames exported automatically for that time. Fusion is then instructed to start rendering from that point.
The resulting frame is loaded and returned to MC, and process is repeated for each frame thereafter.

While Auto-Render is the easier workflow, the manual approach offers faster renderings in Fusion and more control over the performance and memory usage on your system.

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

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

AVID / Fusion layer to comp relationship

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

→ a Loader representing a single clip.

→ 2 Loaders Connected to a Dissolve representing a transition.

→ 2 or more Loaders Connected to Fusion Merges representing layers.

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

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

Due to the design of the AVX2 standard, hidden embedded handles are not supported.

To add handles, increase the length of your clip to include the handle length.
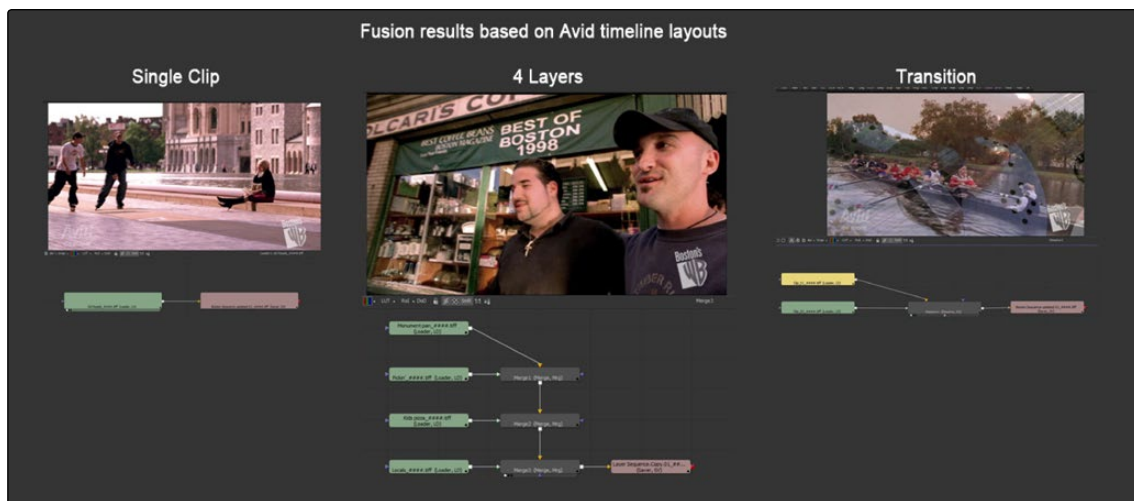
Tip – move your longest clip to the top layer and apply Fusion Connect to that clip.

Fusion / AVID Project relationship

The frame rate and image size preferences created in Avid products are adopted within Fusion's frame rate preferences.

This allows for consistency in formats formats for the round-trip process from Avid to Fusion and back to Avid.

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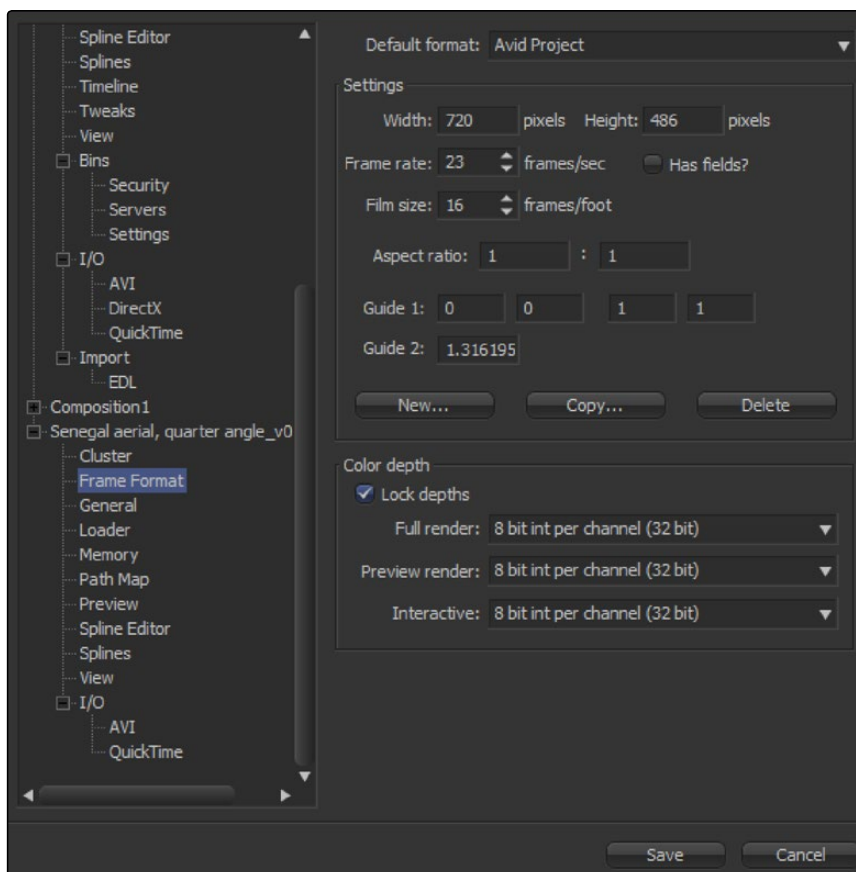
Memory Considerations between Fusion & AVID

The frame rate and image size preferences created in Avid products are adopted within Fusion's frame rate preferences.

This allows for consistency in formats formats for the round-trip process from Avid to Fusion and back to Avid.

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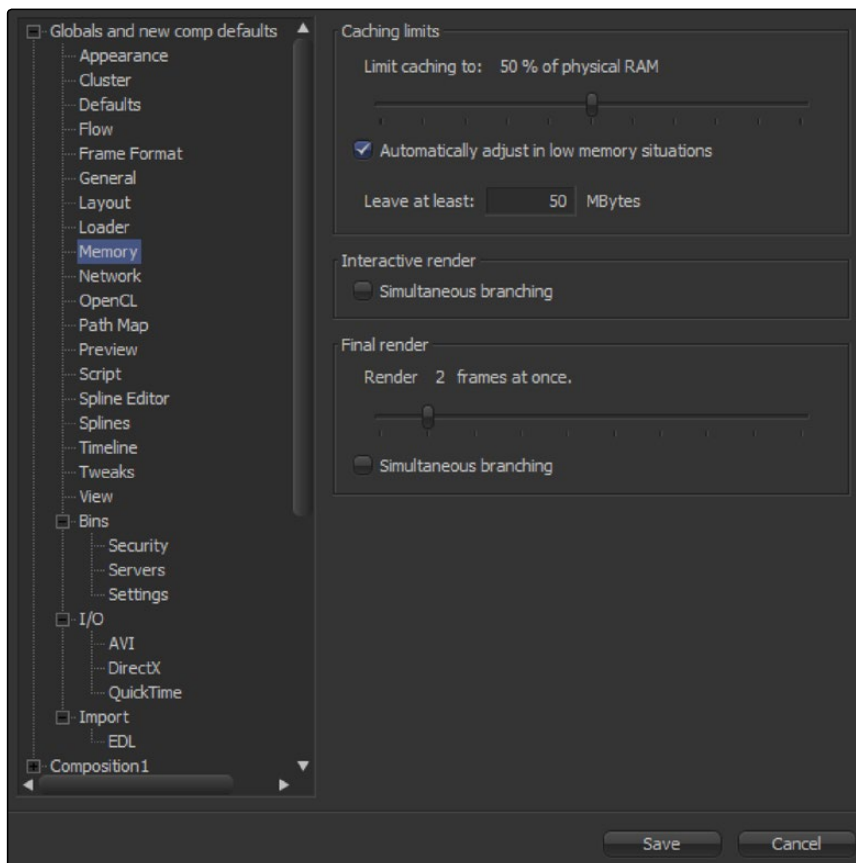
Rendering with Fusion

When you perform a render of your comp inside Fusion, the results are rendered to the output directory created by Fusion Connect during the initial application of the plug-in to the timeline. Upon rendering, you will immediately be able to see the results of the rendered Fusion comp if you switch back to Avid. Even while Fusion is rendering you can continue with the edit process on any clip except for the associated clip being rendered at the time.

Directory Structure of Fusion Connect Media

AVID ProjectName
Master directory for the AVID project.
This can be set within user defined directories.

AVID SequenceName
This directory contains the Fusion comp-files
and sub-directories, named after the AVID clip.

Bob_v01.comp
Charlie_v01.comp
Fusion composition(s) named
after the AVID clip(s).

Bob
Directory named after the AVID clip.

Charlie
Directory named after the AVID clip.

AVID
Directory named 'AVID', containing
the exported source clip(s).

Charlie
Directory named after the
exported clip.

Charlie_0000.raw
Charlie_0000.raw
Charlie_0000.raw
Image sequence named
after the exported clip.

Dave
Optional directory named after
the exported clip.

Dave_0000.raw
Dave_0000.raw
Dave_0000.raw
Optiona second source clip
for Charlie_v01.comp.

xyz_name
Optional directory named after
the exported clip.

xyz_0000.raw
xyz_0000.raw
xyz_0000.raw
Any additional source clip
for Charlie_v01 comp.

Fusion
Directory named 'Fusion'' containing
the rendered clip.

Render_v01
Directory named 'Render_' plus the
version number of the Fusio comp.

Charlie_0000.raw
Charlie_0000.raw
Charlie_0000.raw
Rendered images sequence
named after the AVID clip.

Render_v02
Directory named 'Render_' plus the
version number of the Fusion comp.

Fusion Connect creates a logical directory structure which is not affiliated with the Avid Media Files folder but rather the Fusion Connect AVX2 plug-in. Based on data gathered during the AVX application to the timeline, a logical directory hierarchy is automatically created based on Avid projects, sequences, and clips. This structure allows for multiple instances of Fusion to access the media and multiple instance of the AVX to relate to a single Fusion Comp. In a re-edit situation, media is never duplicated but over-written to avoid multiple copies of identical media.

If you apply the effect to a transition, the naming behavior might be somewhat different.

By default Media Composer refers to the two clips of a transition as "Clip_001" and "Clip_002". Based on the naming convention, Fusion Connect will create directories with matching names. If such directories already exist, because another transition has already used up "Clip_001" and "Clip_002", the numbers will be incremented automatically.

Likewise, "_001" will be added incrementally to the group directory name, if a directory of that name already exists. The corresponding comp file will be named accordingly.

You will notice that the Fusion Connect icon is a green dot (real time) effect.

If your hardware is fast enough, the results that populate the plug-in will play in real-time.

Recommended that you render the green dot effect which will force an MXF precompute to be created to guarantee real-time playback.

Advanced Project Paths

Fusion Connect AVX2 Plugin 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 Variable which are set in the Operating system. This give you the most flexible control over pathing your media and as the name depicts, you can change variables (controls) in various application environments. This useful for network storage and Fusions running on other systems.

Configuring Paths on OS X

When using Fusion Connect on OS X, the Configure Path Defaults dialogue looks like this:
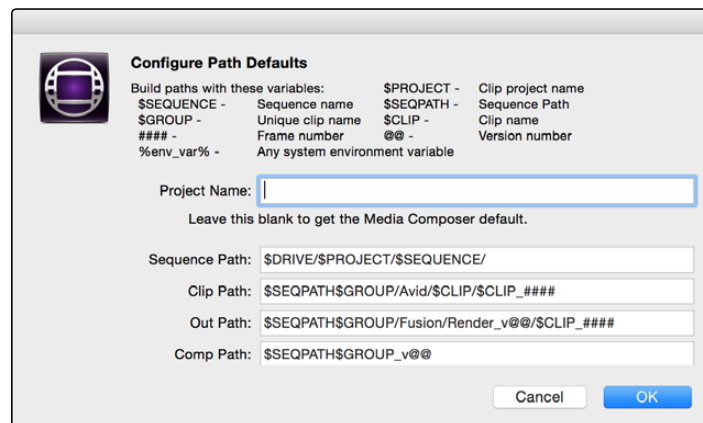
 Default paths can be configured using variables similarly as on Windows, but for added convenience it is possible to enter any desired path defaults directly into fields in the dialogue, without the need for using environment variables.

Configuring Paths on Windows

When using Fusion Connect on Windows, the Configure Path Defaults dialogue looks like this:
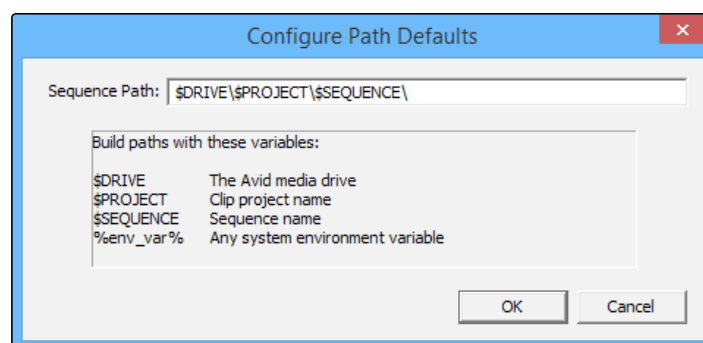
In Fusion Connect can defined the User Variables directly in the Fusion Connect plug-in. Click the Configure Path Defaults button to launch the path defaults dialogue editor. It is located in the Options section of the Fusion Connect AVX2 Plug in. You must click on the triangle to see it.

Fields and Variables

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

Field	Variable	Environment Variable	Description
Project Name	$PROJECT	CONNECT_PROJECT	Overrides the current Avid Project name.
	$DRIVE	CONNECT_DRIVE	Drive or folder for all Connect projects
	$SEQUENCE		Name of Avid sequence
Sequence Path	$SEQPATH	CONNECT_SEQUENCE_PATH	Folder for all Connect files in this sequence
	$GROUP		Unique name of this Connect instance
	$CLIP		Name of exported clip
Clip Path		CONNECT_CLIP_PATH	Folder for exported clips from Avid
Out Path		CONNECT_OUT_PATH	Folder for rendered results from Fusion
Comp Path		CONNECT_COMP_PATH	Location and name of Fusion comp file

Environment Variables

The Pathing can be set in Environment Variables of the system, so that IT managment of the Project Paths can be achieved.

Accessing the Environment Variable on Windows

The quickest way to access it is through your Windows control panel and search the word "env" without the quotes. You have a choice of editing at a user level or system level.

User Variables

Click on the link that says Edit environment variables for your account.

System Variables

Click on the link that says Edit the system environment variables.

Important Rule

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.

Definitions

→ Variable is the control title defined specifically by the application that is being controlled.

→ Value is the instructions that tell the variable what to do.

→ Variables and what they mean.

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

Important

If you type this directly in Fusion Connect's Path editor you do not have to type the Variable, just the value. You also can make modifications without having to restart the Media Composer! The only caveat is in order to clear "get rid of" a variable, you have to exit the Media Composer and clear the environment variable in the windows interface and restart the Media Composer.

Other Values you can control derived from your AVID Bin:

→ $DRIVE This will force the Directory to the Drive where the Avid Media is stored.

→ $PROJECT This will force a directory based on the Avid project name for which the media was digitized/imported or AMA Linked.

→ $SEQUENCE This will force a directory based on the Avid SEQUENCE name for which the media was digitized/imported or AMA Linked.

Here is an example of how a Variable can be setup to support project and sequence names within

your directory.
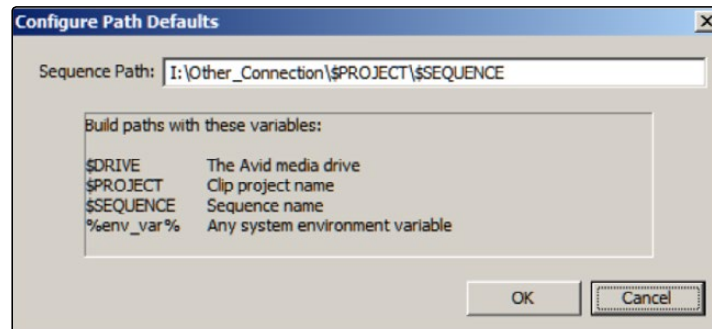
Here is the same example in the Windows environment variable editor