

FUSION 7

USER MANUAL



October 2014 Edition

1

Getting Started	1		
What is Fusion?	2		
Some Key Features	2		
System Requirements	3		
Hardware Requirements and Recommendations	3		
Software Requirements	3		
Compositions	4		
Creating a New Composition	4		
Saving a Composition	4		
Opening a Composition	4		
Closing a Composition	5		
Auto Save	5		
The Composition File Format	6		
Multiple Document Interface	6		
Tabbed Layout	7		
Tool Basics	9		
Tool Tiles	9		
Tools	9		
Inputs and Outputs	10		
Tool Controls	10		
Adding Tools	11		
Using the Menus	11		
Using the Toolbars	11		
Using the Context Menu	12		
Using the Bins	12		
Using the File Browser	13		
Selecting Tools	14		
Deleting Tools	14		
Moving Tools	15		
Connecting Tools	15		
Disconnecting Tools	15		
Replacing an Input	15		
Branching	15		
Copy, Cut And Paste	16		
Copying Tools	16		
Cutting Tools	16		
Pasting Tools	16		
Pasting Tool Settings	17		
Using a Text Editor	17		
Renaming Tools	17		
Tool Defaults and Settings	18		
Saving New Defaults	18		
Saving Tool Settings	18		
Loading Saved Settings	19		
Custom Bin Icon for Saved Settings	19		
Domain of Definition	20		
RoI and DoD	20		
Sampling Outside the DoD	22		
		Region of Interest	23
		The 3D Scene	25
		Working in 3D	25
		Merge 3D	26
		Renderer 3D	27
		Software Renderer	27
		OpenGL Renderer	28
		3D Viewer	29
		Navigating the 3D View	30
		Transforming with the Views	30
		Transparency Sorting	30
		Material Viewer	31
		Transformations	31
		Parenting	33
		Cameras	33
		Importing Cameras	34
		Applications Supported for Camera Import	34
		Parented Cameras	35
		Lighting	35
		Light Hierarchies	36
		Lighting Options	36
		Shadows	37
		Shadow Maps	37
		Multiplicative and Additive Bias	38
		Shadow Softness	38
		Transmittance	39
		Materials	40
		Illumination Models	40
		Material Components	41
		Transmittance	42
		Textures	43
		Composite Materials	44
		Reflections and Refractions	45
		The Export Bumpmap Tool	49
		Geometry	52
		Primitive Shapes	52
		Visibility	52
		FBX Format	53
		FBX Export	54
		Fog 3D and Softclipping	55
		Matte Objects	55
		Material and Object IDs	56
		Point Clouds	57
		Stereoscopic	59
		What is a WPP?	61
		Different Coordinate Spaces	61
		The Scene Input	61
		Position Pass Concept	61
		The “Invisible Sphere”	62

Polyline Types	63	Follow Published Points	75
Bezier Polylines	63	Double Polylines	75
Polylines, Effect Masks and Motion Paths	63	Converting a Single Polyline to a Double Polyline	76
B-Spline Polylines	64	Adding Softness to a Segment	76
Polyline Toolbars	64	Adding Additional Points to the Curves	76
Selecting a Polyline	65	Locking/Unlocking Point Pairs	76
Polyline Creation Modes	65	Converting Spline Types	77
Protection Modes	66	Types of Motion Paths	78
Closing and Opening Polylines	66	Path Modifier	78
Selecting and Adjusting the Polyline	67	Displacement Paths	78
Moving Polyline Points	67	Displacement Splines and Path Timing	78
Edit Points	67	Motion Paths	78
Smoothing a Polyline Segment	67	Types of Points	79
Linearizing a Polyline Segment	68	Locked Points	79
Deleting Selected Points	68	Unlocked Points	79
Bezier Handles	68	Look at the Spline Editor	80
Edit Points	68	Locking and Unlocking Points	80
Reduce Points	69	XY Paths	80
Shape Box	69	Switching Default Paths	81
Other Polyline Options	70	General Path Techniques	81
Stop Rendering	70	Path Centers	81
Showing and Hiding Points and Handles	70	Copying and Pasting Motion Paths	81
Importing and Exporting Polylines	70	Removing Motion Paths	81
ASCII Format	70	Recording	81
Roto Assist	71		
Roto Assist	71		
Native Format	71		
Mask Types	72		
Effect Masks	72	2 Interface Overview	82
Garbage Mattes	72	Toolbars	83
Pre-Masking (DVE, Glow and Highlight Tools)	72	Main Interface	83
Effects Mask	72	Display Views	83
Mask Tools	73	Control Panels	84
Common Mask Features and Controls	73	Time Ruler	84
Bitmap Masks	73		
B-Spline Masks	73		
Ellipse, Rectangle and Triangle Masks	74	3 Navigating the Interface	85
Mask Paint	74	Panning	86
Polygon Mask	74	Scaling	86
Wand Mask	74	Layout	86
Ranges Mask	74		
Polygon and B-Spline Animation	74		
Polygon Change Spline	74	4 Comments Tab	88
Removing Animation from a Polyline Mask	74		
Adding and Removing Points from an			
Animated Mask	75	5 Console Introduction	90
Published Points	75	Console Toolbar	91
Publishing a Point	75	Console Toolbar Buttons	91
Publish Point to Path	75	Console Entry Field	92

6

Time Ruler	93
Current Time	94
Keyframe Ticks	94
Changing the Display Format	94
Ranges	95
Global Start and End Range	95
Render Start and End Range	95
Visible Range Slider	95
Render Button	95
Playback Buttons	95
Audio Mute	96
Composition Quality Options	96
High Quality	96
Motion Blur	96
Proxy and Auto Proxy	97
Selective Updates (Some, None and All)	97

7

Bins	98
Overview	99
Bins Layout	99
Bins Hierarchy	100
Connecting to Bins on Another System	100
Adding an Entry to the Bin Server List	100
Accessing Remote Bins	101
Permissions	101
Managing the Bins	101
Items That Can Be Stored in the Bins	101
Opening and Closing the Bins	101
Interface Overview	101
Bins Context Menu	102
Working with Bins	102
Adding Items to the Bins	102
Compositions and Footage	102
Tools	102
Tool Settings	102
Footage Stills vs. Sequences	103
Creating a New Bin Folder	103
Renaming a Bin Folder	103
Deleting Items from a Bin	103
Deleting a Bin Folder	103
Adding an Item to the Current Composition	104
Footage	104
Compositions	104
Tools and Macros	104
Tool Settings	104
Working with Clip Items	104
Stamp Files	104

8

Stamp Playback	105
Scrubbing Footage	105
Jog and Shuttle	105
Backing Up the Bins	106
Display Views	107
Types of Display Views	108
Left/Right Display Views	108
Floating Display Views	109
Hardware Frame Buffers	109
Display View Context Menu	110
Displaying a Tool's Output in the Views	110
Viewing Images with Drag and Drop	111
Viewing Images with the Context Menu	111
Viewing Images with Keyboard Shortcuts	112
Removing Images from the View	112
Viewing Images from Disk	112
Viewing Images Using Other Methods	112
Position and Layout	113
Layout Toolbar	113
Splitter Bars	114
Fullscreen Floating Views	114
Fullscreen Fixed Views	114
Panning and Scaling the Image	114
Panning	114
Scaling	115
Using the Mouse	115
Using the Keyboard	115
Using the Context Menu	115
Rotating	116
Using the Mouse	116
Previews	116
Settings	117
Size	117
Network	118
Shoot On	118
Frame Range	118
Configurations	118
Playing Back the Preview	119
Mouse Shortcuts	119
Keyboard Shortcuts	120
File Sequence Previews	120
Onscreen Controls	120
Enabling/Disabling Onscreen Controls	121
Making Fine Adjustments to Controls	121
Toolbars	121
Display View Toolbar	121
Tool - Specific Toolbars	122

A/B Buffers	122
Switching Between Buffers	123
Subviews	124
Displaying/Hiding The Subview	124
Changing the Subview Type	124
Swapping the Subview with the Main View	124
View and Subview Types	124
Changing the View Type	125
2D Viewer	125
3D Image Viewer	125
Histogram	125
3D Histogram	126
Vectorscope	126
Waveform	126
Navigator	127
Magnifier	127
Image Info	127
Color Inspector	127
Metadata	128
Choosing Color Channels	128
Viewing Non-Color Channels	128
The 3D View	129
Wireframe	129
Changing 3D Viewpoints	129
Lighting	130
Shadows	130
Transparency	131
Grid	131
Vertex Normals	132
Quad View	132
Quad View Layouts	133
Effect Masks	133
Guides	134
Enabling and Disabling the Guides	134
Selecting a Guide	134
Guide Frame Aspect	135
Look Up Tables (LUT)	135
Opening the LUT Editor	136
Managing LUTs	137
Display Settings	137
LUT Curves	137
LUT Files	138
Settings and Macros	138
LUT Preferences	138
How Lookup Tables Work in Fusion	138
LUT Processing Order	138
Enabling and Editing Image LUTs	139
Enabling and Editing Buffer LUTs	139
Stacking LUTs Together	139

Types of LUTs	139
Save and Load LUT Settings	140
Setting a Default LUT	140
View Preferences and Settings	140
View Preferences	140
View Settings	140
Saving View Settings	141
Loading View Settings	141
Loading and Saving Defaults for a View	141
Status Bar Information	141
General Display Options	141
Display Options Sub Context Menu	141
Display View Label	143
Controls List	144

Flow Editor	145
Overview	146
Viewing, Scaling and Panning the Flow	147
Displaying the Flow	147
Resizing the Flow	147
Panning the Flow	148
Context Menus	148
Flow Context Menu	148
Tool Context Menu	149
Adding Tools	149
Using the Menus	149
Using the Toolbars	150
Using the Context Menu	151
Using the Bins	151
Using the Add Tool Script	152
Using the File Browser	152
Connecting and Disconnecting Tools	153
Connecting Tools	153
Disconnecting Tools	154
Replacing Connections	154
Tracing Pipes Through the Flow	154
Branching	155
Inserting Tools	155
Drag and Drop	155
Shift Insert	155
From Toolbars and Menus	156
Replacing Tools	156
Deleting Tools	156
Connection Types	156
Moving Tools and Organizing the Flow	157
Moving Tools with the Mouse	157
Lining Up Tools	157
Align to Grid	157

1010

Removing a Filter	192	Viewing, Scaling and Panning the Timeline	205
Time and Value Editors	192	Viewing the Timeline	205
Time Buttons	192	Panning/Scrolling the Timeline	205
Time	192	Scaling the Timeline	205
Time Offset	192	Using the Toolbars	205
Time Scale	192	Using the Axis Labels	206
Value Buttons	192	Using the Mouse	206
Value	193	Using the Keyboard	206
Value Offset	193	Working with Tools in the Timeline	206
Value Scale	193	Selecting Tools	206
Looping Splines	193	Re-Ordering Tools in the Tree	206
Set Loop	193	Moving Tools	206
Loop	193	Trimming Tools	207
Relative Loop	194	Extending the First/Last Frame	207
Ping-pong	194	Showing Keyframe Values	207
Set Pre-loop	194	Working with Splines in the Timeline	207
Duplicate Repeating Splines	194	Changing Keyframe Time	207
Reversing Splines	195	Duplicating Spline Keyframes	208
Time Stretching	195	Time Stretching	208
Shape Box	196	Showing Keyframe Values	209
Ease In/Out	197	Timeline Filters	209
Reducing Points	197	Applying Filters	209
Guides	198	Creating Filters	210
Working With Guides	198	Configuring Filters	210
Creating Guides	198	Removing Filters	210
Deleting Guides	198	Timeline.def	211
Moving Guides	198	Timeline Options	211
Jumping to Guides	199	Collapse/Open All	211
Sharing Guides with the Timeline	199	Current Time Mode	211
Autosnap	199	Display Point Values	212
Autosnap Context Menu	199	Line Size	212
Autosnapping and Guides	200	All Lines Size	212
Importing and Exporting Splines	200	Sort Order	212
Exporting a Spline	200	All Tools	212
Importing a Spline	200	Hierarchy	213
Additional Spline Options	200	Reverse	213
Spline Context Menu	201	Animated	213
Scale	201	Guides	213
Options	201	Creating Guides	213
Show	202	Working With Guides	213
Tool Splines	202	Deleting Guides	213
		Moving Guides	213
		Show Guide List (Shift-G)	214
		Renaming Guides	214
		Sharing Guides with the Spline Editor	214
		Autosnap	214
		Autosnapping and Guides	215
		The Spreadsheet Editor	216
		Selecting a Tool to Edit	216

11

Timeline Editor	203
Timeline View	204
Tool Menu	204
Current Time Indicator	204
Timeline Toolbar	205
Spreadsheet	205

	Selecting Multiple Tools to Edit	216		Comments Tab	234
	Selecting Animation to Edit	216		Scripting Tab	235
	Inserting Keyframes	216		Scaling and Panning the Control Window	235
	Moving Between Cells of the Spreadsheet	216		Control Animation	236
	Timeline Tool Color Coding	217		Animation Curves	236
	Tool Context Menu	217		To Animate a Control Using the Default Animation Curve	236
				To Animate a Control Using a Non-Default Animation Curve	237
				To Set a Keyframe on An Animated Control	238
				Removing Animation	238
				Connecting a Control to An Existing Animation Curve	239
				Connecting One Control to Another Without Animating Them	239
				Using Simple Expressions and Pickwhipping	240
12	Script Debug	218	14	Preferences	241
	Breakpoints	220		Preference Settings	242
	Watch Area	220		3D View Preferences	244
	Debugging Tutorial	220		Grid	244
	Watching Variables	221		PerspectiveViews	245
	Stepping Over and Breakpoints	221		Fit to View	245
	Fixing the Bug	221		Default Lights	245
				Cluster Preferences	246
				Slave Groups to Render With	246
				Frames to Render	247
				Default Preferences	247
				Default Animate	247
				Auto Tools	248
				Global Range	248
				Time Code	248
				Flow Preferences	248
				Force	249
				Options	249
				Frame Format Preferences	250
				Default Format	251
				Settings	251
				Color Depth	252
				General Preferences	252
				Usability	252
				Controls	253
				Auto Save	254
				Proxy	254
				Layout Preferences	255
				Program Layout	255
				Document Layout	255
				Window Settings	256
13	Tool Controls	223			
	The Controls Area Interface	224			
	Maximizing the Control Area	224			
	Viewing the Control Window for a Tool	225			
	Hiding the Control Window for a Tool	225			
	Control Area Display Preferences	225			
	Selecting a Tool with the Control Header	225			
	Viewing and Previewing a Tool with the Control Header	226			
	Viewing a Tool's Splines with the Control Header	226			
	Types of Controls	226			
	Slider Controls	226			
	Thumbwheel	227			
	Range Controls	227			
	Checkboxes	227			
	Drop-Down Lists	228			
	Button Arrays	228			
	Color Wheel and Picker	228			
	The Windows Color Picker	229			
	The Color Wheel	229			
	Picking Colors from an Image	229			
	Gradients	230			
	Gradient Colors	231			
	Gradient Interpolation Method	231			
	Context Menus	231			
	Tool Context Menu	231			
	Control Context Menu	231			
	Common Controls Tab	232			
	Masking	233			
	Object and Material ID Masking	233			
	Motion Blur	234			

Loader Preferences	256
Defaults	257
Cache	257
Memory Preferences	258
Caching Limits	258
Interactive Render	259
Network Preferences	259
General	260
Email Notification	260
OpenCL Preferences	260
Options	260
Debugging	261
Preview Preferences	262
Options	262
Type	263
Path Maps Preferences	263
Script Preferences	264
Login	265
Options	265
Spline Editor Preferences	265
Spline Editor Options	265
Autosnap Points	266
Autosnap Guides	266
Autoscale	266
LUT View Options	266
Splines Preferences	267
Autosmooth	267
B-Spline Modifier Degree	268
B-Spline Polyline Degree	268
Tracker Path Points Visibility	268
Tracker Path	268
PolyLine Edit Mode on Done	268
Allow Selection of Points From Multiple Polylines	269
Timeline Preferences	269
Filter	269
Settings for Filters	269
Tool List	270
Timeline Options	270
Tweaks Preferences	270
Network	270
File I/O	271
Area Sampling	272
Open GL	272
View Preferences	273
Bins/Security Preferences	274
Users	274
Bins/Server Preferences	276
Bins/Settings Preferences	277

I/O AVI Preferences	278
AVI Video Compression Options	278
I/O QuickTime Preferences	279
QuickTime Video Compression Options	280
EDL Import Preferences	280
Import Properties	280
Managing Fusion Preferences in Larger Environments	281
Preferences for Each Artist	281
Multiple Preferences	283
About Comp and Global Preferences	283
So What About Frame Format Then?	283
Conclusion	284

15

Tracking Explained	285
Steady Position	286
Original Motion	286
Steady Position 1	287
Steady Position 2	287
Unsteady Position	288
Original Motion	288
Unsteady Position 1	289
Unsteady Position 2	289
Steady Angle	290
Original Motion	290
Steady Angle	290
Unsteady Angle	291
Original Motion	291
Unsteady Angle	292
Steady Size	293
Original Motion	293
Steady Size	293
Unsteady Size	294
Original Motion	294
Unsteady Size	295
Offset Position	296
Original Motion	296
Offset Position 1	296

A1

Frame Formats	298
Setting the Frame Format	299
Setting Up the Default Frame Format for New Compositions	299
Changing the Frame Format in Existing Compositions	300
Fields	300
Combining Fielded and Full Frame Images	301

Setting Field Processing for the Flow	302	Submitting a Composition	322
Pixel Aspect	303	Flipbook Previews	322
How Pixel Aspect Affects the Display Views	303	Disk Cache	322
Setting Pixel Aspect for Creator Tools	304	Clustering	322
Setting Pixel Aspect for Loader Tools	305	Configuring Slaves for Clustering	323
Combining Images with Different Pixel Aspects	305	Enabling Clustering on a Tool	323
Color Depth and Processing	306	Removing Clustering from a Tool	323
Supported Color Depths	306	Remote Administration	323
Setting Color Depth	307	Installing Updates Remotely	324
Combining Images with Different Color Depths	308	Setting the Render Master Remotely	325
Advantages of Floating Point Processing	308	When Renders Fail	325
Detecting Out of Range Values	308	Automatic Rejoining of the Queue	325
When to Use Floating Point	309	eyeonServer	326
Clipping Out of Range Values	310	Frame Timeout	326
		Heartbeats	326
		Managing Memory Use	327
		Override Composition Settings	327
		Render Several Frames at Once	327
		Simultaneous Branching	327
		Things You Need To Know	327
		Time Stretching	327
		Linear Tools	328
		Saving to Multiframe Formats	328
		Video Hardware	328
		Troubleshooting	328
		Check the Render Log	328
		Check the Composition	328
		No Slaves Could Be Found	328
		The Composition Could Not Be Loaded	329
		The Slave(s) Stop Responding	329
		The Slave(s) Failed to Render a Frame	329
		Check the Slaves	329
		Check the Network	329
A2 Network Rendering	311		
TCP/IP and Fusion	312		
Ensuring That the Network Functions Correctly	312		
Using a Third-Party Render Manager	312		
Setting Up Network Rendering	313		
What is the Render Master?	313		
Rendermanager	313		
Setting Up the Render Master	313		
Preparing Render Slaves	314		
Choosing the Render Master Used By Slaves	315		
The Render Manager	315		
Opening the Render Manager Dialog	316		
Adding Computers to The Slave List	316		
Scanning for Slaves	316		
Manually Adding Slaves	316		
Removing Computers from the Slave List	317		
Loading and Saving Slave Lists	317		
The Render Queue	317		
Adding Compositions to the Queue	317		
Removing Compositions from the Queue	318		
Saving and Loading Queues	318		
Re-Ordering the Queue	319		
Groups	319		
Assigning Slaves to Groups	320		
Assigning Compositions to Groups	320		
Using Multiple Groups	320		
The Render Log	320		
Preparing a Composition for Network Rendering	321		
Paths	321		
Using UNC Names	321		
Using Mapped Drives	321		
Fonts	321		
Plugins	322		



Getting Started

1

Getting Started

What is Fusion?

Fusion is a compositing program, or more accurately, software used to composite images. Look up the word composition in any dictionary and the most relevant definition would likely be: “combining of distinct parts or elements to form a whole.”

Fusion is designed to combine images from many sources and compose them into a final whole. Modern media assembles a multitude of raw materials into the final product that we see in film, television or via the Internet. Virtual sets are blended with live action material. Computer generated actors fill in where live stuntmen can't. With each passing year the real and virtual parts of production become more tightly intertwined.

So the job of the compositing artist is to take all of the discreet elements and layers that make up a shot, and turn them into a single image that respects the goals and artistic vision of the production. To achieve this, the artist draws on techniques as diverse as color correction, greenscreen replacement, tracking, particles, blur and glow effects.

Fusion draws on more than 20 years in visual effects post production, providing a complete toolkit for addressing any compositing challenge.

Some Key Features

In addition to providing in excess of 200 tools for image processing, Fusion also provides a large number of features to improve an artist's effectiveness and efficiency. The following is a summary of some of the key features provided by Fusion:

- > Region of Interest and Domain of Definition renderer with the interactive responsiveness of a full frame based renderer
- > Concatenated transformations preserve image sharpness
- > Offscreen pixels are preserved through infinite canvas provided by Domain of Definition
- > Complete toolset for compositing in 2D and 3D simultaneously
- > Reads and writes more than 100 production file formats
- > Resolution Independent
- > Comprehensive HDRI and floating point color support
- > Automation through either Python or eyeonScript (Lua) programming languages
- > ASCII based project format for easy integration with asset management
- > RAM and disk-based caching systems
- > Strongly multithreaded native 64-bit application takes full advantage of modern hardware
- > GPU-accelerated 3D rendering and display
- > Real time color calibration through all common 2D and 3D look up table formats
- > Network rendering not just of final projects but also for previews and disk caches
- > Background and clustered network rendering directly to RAM and disk-based caches
- > Bins server provides facility-wide access to frequently used tools, macros and settings
- > Macro system for repackaging many tools into one simplified but powerful tool
- > Easy to use and well-documented SDK

System Requirements

Before installing Fusion, review the following system requirements.

Hardware Requirements and Recommendations

Processor

Fusion will take advantage of Intel and AMD processors equally well. It is recommended to have at least a Pentium 4 or Athlon 64 processor to operate Fusion. The 64-bit version of Fusion requires dual core and hyperthreaded processors, and thrives when multiple processors are available.

Graphics

Fusion requires OpenGL capable ATI or NVIDIA graphics cards. Available graphics memory should be 2GB or greater. The Fusion interface requires a screen resolution of 1280 x 1024 pixels or greater at 32-bit per channel (true color). Higher resolutions are suggested.

Fusion's graphics card requirements will be higher for 3D intensive composites. For example, 3D scenes employing large numbers of textures, or very large textures will require significantly more memory for the graphics card. Additionally Fusion makes more extensive use of the GPU for many 3D operations. A modern graphics card will have a big impact on interactivity and performance in Fusion.

High-end workstation class graphics cards should not be required; a decent mid-class gaming class card with large amounts of memory is generally sufficient.

Fusion Render Slaves will be unable to render OpenGL 3D scenes unless they have an OpenGL-capable graphics card. Software rendering of 3D scenes will not require a graphics card.

Memory

Fusion is designed to make use of every kilobyte of memory available. A minimum of 8GB of memory is required, and 16GB or greater is strongly recommended.

USB Port

Fusion uses a hardware lock as part of its licensing system. This is a small device that connects to the USB port.

Hard Drives

Fusion requires a minimum of 1GB of hard disk space to install. When working with images stored locally to the workstation, the speed of the drives can be a major factor in the performance of Fusion. If the majority of the images used in a composite come from a remote fileserver or SAN, the network will be a much bigger factor in rendering speed than the drive performance.

Software Requirements

Operating Systems

Fusion 7.0 supports the following operating systems:

- > Microsoft Windows 7 64-bit
- > Microsoft Windows 8 64-bit

QuickTime

Fusion requires Apple QuickTime to be installed on the workstation and render nodes in order to read and write to the MOV and PNG file formats.

Note: Due to the poor Open GL support in most virtualization platforms Fusion does not currently support operation under virtualization.

The Fusion Render Slave should operate under virtualization, although it is more efficient to run directly on the hardware.

Compositions

The term ‘Composition’ is used to refer to the Fusion project file. This chapter describes how to create, load, and save compositions. A default installation of Fusion will create a new empty composition when it is launched.

Composition is frequently abbreviated to Comp by Fusion artists.

Creating a New Composition

Any of the following methods can be used to create a new composition:


- > Select File > New from the menus at the top of the Fusion window
- > Click on the ‘New’ toolbar button
- > Use the CTRL + N keyboard shortcut while the Fusion window is active

Saving a Composition

Any of the following methods can be used to save the current composition to the file system

- > Select File > Save from the menus at the top of the Fusion window
- > Select File > Save As from the menus at the top of the Fusion window to save under a new name
- > Click on the ‘Save’ toolbar button
- > Use the CTRL + S keyboard shortcut while the Fusion window is active
- > If the composition has not previously been saved, the above methods will ask the artist to provide a file name using the standard File Dialog. If the composition has been saved, the above methods will save over the existing file. The file extension .comp will automatically be added to the end of the file name chosen by the artist.

Compositions that have unsaved changes will display an asterisk (*) marker next to the composition’s name in the Fusion windows title bar, as pictured below:



To save a composition with a new name select File > Save As from the menus at the top of the Fusion window.

Opening a Composition

Any of the following methods can be used to open a composition from the file system:

- > Select File > Open from the menus at the top of the Fusion window
- > Click on the ‘Open’ toolbar button
- > Use the CTRL + O keyboard shortcut while the Fusion window is active
- > Drag the composition file from the file browser into an empty Fusion window, or to the Time Ruler of the current composition

- > Drag the composition file from the bins into an empty Fusion window, or to the Time Ruler of the current composition
- > Double-click on a composition file in the operating systems file browser

The first three options will all display the standard File Dialog.

Closing a Composition




Any of the following methods can be used to close the current composition

- > Select File > Close from the menus at the top of the Fusion window
- > Click on the 'Close' icon in the top right of the Composition's window

If the composition has unsaved changes a dialog box will appear asking if the artist wishes to save before closing.

Auto Save



Auto saving is a feature that allows Fusion to save the composition automatically to a temporary file at pre-set intervals. This is to protect the artist from loss of work due to power loss, bugs in the software or accidental closure. A default installation of Fusion has the auto-save feature disabled.

To enable auto save for new compositions open File > Preferences from the menu, then locate Global > General > Auto Save in the preferences dialog.

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

If an autosave file is present when Fusion loads a composition, a dialog will appear asking if the artist wishes to load the autosaved or original version of the composition.

The Composition File Format

Composition files are saved to the file system as human-readable plain text files. Using plain text files to describe the project makes it easier for studios to integrate Fusion into existing pipeline and asset management solutions.

These files can then be opened and edited using any standard text editing program. Fusion is installed with the SciTE text editor, which recognizes the format of Fusion compositions and provides additional features such as folding and syntax highlighting. Never use a formatting word processor, such as Microsoft Word, to modify a composition file, as these will generally save additional formatting information, making the composition unreadable to Fusion.

```

1  - Composition {
2      CurrentTime = 1,
3      RenderRange = { 1, 100, },
4      GlobalRange = { 1, 100, },
5      HiQ = true,
6      Version = "Fusion x64 6.0 build 365",
7      SavedOutputs = 0,
8      HeldTools = 0,
9      DisabledTools = 0,
10     LockedTools = 0,
11     AudioOffset = 0,
12     Resumable = true,
13     OutputClips = {
14         "comp:\\out.0001.tif",
15     },
16     Tools = {
17         Bump1 = Bump {
18             Inputs = {
19                 Input = Input {
20                     SourceOp = "Transform1",
21                     Source = "Output",
22                 },
23                 HeightScale = Input {
24                     Value = 2.0001,
25                 },
26                 MaterialID = Input {

```

Generally speaking, the advantages of saving a composition in ASCII 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 as plain text. Specifically, a composition that contains a very large number of animated polygon style masks may become abnormally large, with sizes in excess of 100MB. In these cases it can be desirable to enable saving of compressed compositions in the GlobalGeneral preferences.

Multiple Document Interface

Fusion uses a Multiple Document (MDI) interface that allows it to have multiple composition projects open simultaneously. Opening or creating a composition does not close any other compositions that may be open at the time. Instead, the compositions are moved to the background.

A list of currently open compositions can be seen at the bottom of the Window menu. The artist can switch between open compositions using one of the following methods:

- > Selecting Window > composition name from the menus at the top of the Fusion window

- > Using the CTRL + TAB keyboard shortcut to toggle between available compositions
- > Using the Window > Cascade or Window > Tile Horizontally or Window > Tile Vertically options to automatically arrange the available Compositions so that they are all visible

Having multiple compositions open at once makes it simple for the artist to switch between different compositions and makes it possible to copy/paste between compositions or to compare parts of one composition with another.


Tabbed Layout




Views in Fusion can be arranged to allow for an optimum efficiency based on your personal Workflow.

They can be arranged both in the main GUI Window, as well as within floating views. Right-click in the upper right corner of any view or tab to bring up the Add View menu. Select the desired view from the list.

They can be arranged both in the main GUI Window, as well as within floating views. Right-Click in the upper right corner of any view or tab to bring up the Add View Menu. Select the desired view from the list.



You can change the order of tabs by dragging them with the mouse to the desired position. When you drag a Tab out of its original position, the View-Arrange menu will appear.




It consists of a cross shaped area with 5 droppers, which appear in the middle of the view the mouse is hovering on, and 4 individual droppers, which appear at the outer perimeter of the Fusion interface.

The droppers in the center cross-shaped dropper allow you to add that Tab to the view the mouse is hovering over.

Dropping it on any of the 4 Top/Left/Right/Bottom droppers splits the current view and adds the Tab at the respective position. Dropping it onto the center dropper opens a new tab in the current view rather than splitting it. You can also drop the Tab onto an existing set of tabs to add it to the set.


Dropping it on any of the 4 droppers at the perimeter will split the entire Interface at the respective position to add the Tab. Right-click in the background area of the upper right corner of any view to bring up the View Layout menu to Remove or Undock a tab.

You can also select to show or hide tabs here.



Tools

In Fusion a tool is an operation that affects one or more images. Tools are responsible for loading images from disk, saving the results of the composition, merging multiple images together, and applying effects like blurs, glows, color corrections and more. This chapter discusses how to add tools to a composition, connect them together, and other basic operations.



Tools are often referred to by artists as nodes, tiles or operators. These terms are synonymous.

When a tool is added to a composition, a tool tile is added to the Flow editor, and the tools name appears in the Timeline editors as well. The flow is the only interface where an artist can easily manipulate the connections between tools, so typically this is where the majority of time is spent when working with tools.


Tool Basics

Tool Tiles

Almost every tool in a composition adds a tool tile to the flow. These tiles display the inputs and output for the tool as small colored icons around the edges of the tile. The arrow-shaped nodes represent the image inputs, and the grey or red square nodes represent the tool's processed output.

The tools name will appear in the center of the tile. Tool tiles can be displayed as a small rectangle or as a larger square. The rectangular form displays the tools name in the center of the tile, while the square form shows either the tools icon or a thumbnail of the image it is outputting. More detail on setting the display options for the tool tiles can be found in the Flow Editor chapter.

Hovering the mouse over a tool will display a pop up with more information about the tool and the image or data it is currently processing.




Inputs and Outputs

All tools have one output. They generally also have one or more inputs. Tools are connected together by drawing connections between the output of one tool and the input of one or more tools. Inputs and Outputs can transmit multi-channel and single channel image data, 3D scene and material data, or particle information. The exact type of data required by each input and produced by each output is described in the tool reference documentation for the specific tool.

Tools cannot have more than one output, however, that one output can be connected to an unlimited number of inputs on other tools.

In the example image below a Loader tool's output is connected to a Blur tool's input. The Blur tool's output is then connected to a Color Corrector.



Holding the mouse pointer over a tool's input will cause a small pop-up containing the input's name to appear.

Tool Controls

When a tool is selected its options and settings will appear in the Control panel on the right side of the Fusion interface. See the Control panel documentation for more information.




Adding Tools

Tools are added to the composition's Flow Editor using one of the methods described below.

Using the Menus


The Tools menu at the top left of the Fusion window displays every tool known to Fusion. These tools are organized into general categories. To add a tool to the composition from this menu, navigate down the category list and select the desired tool.



If another tool is currently active on the flow, the new tool is automatically inserted downstream of the existing tool. The new tool is inserted into the existing connection of the active tool.

If there is no tool currently active tool in the flow, the new tool will be placed in the first empty space in the flow.

Using the toolbars



The Toolbars can also be used to add tools to the flow. These small bars display an abbreviated form of each tool's name (i.e., BC for the Brightness/Contrast tool). Hold the mouse over any toolbar button to display a pop-up with the unabbreviated name of that tool.


Clicking on a tool's button in a toolbar adds it to the composition exactly as if the tool had been selected from the Tool menu (as described above). The artist 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.

Tools can also be dragged from the toolbars into the left and right display views. This allows tools to be added to the project even if the flow is not currently visible. Tools that are dragged onto a view from the toolbar are added to the flow immediately after the currently viewed tool, regardless of the currently active tool.

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

Using the context Menu

The Flow context menu (displayed when right-clicking on an empty part of the flow) and the Tool context menu (displayed when right-clicking on a tool tile) contain one or more submenus that can be used to add tools to the composition.



Add or insert tool

Selecting a tool from the Add Tool or Insert Tool submenus will insert a tool into the flow at the position where the mouse pointer was when the context menu was opened.

If the mouse pointer is positioned over a connection between two tools, the new tool will be inserted into the flow already connected to that pipe (between the existing tools).

Replace Tool


The Replace Tool submenu only appears when the context menu is opened by right-clicking directly on an existing tool's tile in the flow. Selecting a tool from the Replace menu will replace the current tool with a new one.

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

Using the Bins

The bins can display thumbnail icons for each of the tools and third party plugins that are known to Fusion. To toggle the bins on and off, select Ctrl + B from the keyboard, or select File > Bins from the menu.

To add a tool from the bins to the project, drag and drop the tool's icon from the bin into the flow.




The bins for a new install of Fusion must be configured to show the tools. They are not automatically placed in the hierarchy. To add the tools to the bins, right-click in the bins and select **Create Tool Folders** from the context menu that appears.

Bins can contain much more than just tools. They provide a useful place for storing shortcuts to saved tool settings, groups, macros and media, making the bins more flexible than even toolbars. For more information, refer to the bins chapter of this manual.

Using the File Browser

Dragging and dropping files from the operating system's file browser is an efficient way to quickly add footage or saved tool settings to a flow.



Footage

Dragging a file from the file browser to the flow will automatically add a loader to the flow, directly loading the file or file sequence. If multiple files are dragged into the flow, a separate loader is added for each file.


Holding **Doc_key-Shift** while dragging a file from a file sequence will create a Loader tool that is already trimmed to use that file as a still. This can be useful when loading an image taken by a digital still camera.

Saved Tool Settings


Dragging a tool's settings (.settings) file from the file browser to the flow will add the tool (or tools) saved in that file to the flow, configured with the settings that were saved into the file.

Selecting Tools

In order to move, copy, view or adjust a tool's controls it is first necessary to select the tool. This can be done by left-clicking directly on a tool's tile, or 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. A tool selected this way is known as an active tool. The tool tile for the active tool will turn bright yellow. Only one tool can be actively selected at a time.




One or more tools can also be selected by left-clicking and dragging the rectangle to surround the tool or tools to be selected. The tools within the rectangle will be added to the selection group, and the color of the tool tiles will change to blue.



Tools can be added or removed from a Selection group by holding down the CTRL key while making selections. A Selection group may or may not contain an active Selection tool.

In the image below there are four tools selected, one of which is active. Two other tools are not selected.



Deleting Tools


To delete a tool or tools from the composition, first select the tool or tools and then press the Delete key on the keyboard. Alternately, right-click on one of the selected tools and then select Delete from the context menu that appears.

Moving Tools

To move tools around in the flow, select one or more tools and then left-click and drag to move them to a new location. All of the currently selected tools will be moved. Larger flows can become disorganized and difficult to understand, so several techniques for organizing the flow exist.

Connecting Tools


Build a flow by connecting tool tiles together using their inputs and outputs. For example, to connect the output of a Background tool to a Color Corrector tool, the artist would left-click on the output node on the Background tool and drag a connection over the input node of the Color Corrector. A thin line will be drawn between these two nodes to represent that connection.



The connection between tools does not always have to be drawn in output to input order. As long as the target input is not already connected to something else it is entirely valid to drag a connection from the input to the output of another tool instead.

Disconnecting Tools

The pipe connecting two tools is divided into two halves the output half and the input half. To disconnect a tool's inputs or output, click on the input half of the connection between the tools (or click on the input arrow itself) and drag to pull the connection away from the tool.




Holding down the Shift key while moving the tool can disconnect an entire tool. If the Shift key is still held when releasing the tool then all existing connections to that tool will be removed.

Replacing an Input

Dragging a connection over to an input that already has a connection will create a new connection that replaces the old one.

Branching

The output of any tool can be connected to multiple inputs on other tools, which is a technique called branching. Branching permits a single image to be simultaneously processed by two or more streams of tools.



Only one connection per input is allowed but a tool output can have any number of branches and outbound connections.

Copy, Cut and Paste

Copying Tools

There are four ways to copy the currently selected tools to the system clipboard.

- > Select Copy from the Tool context menu
- > Choose Edit > Copy from the Fusion menus
- > Click on the Copy icon in the toolbar
- > Use the Ctrl-C keyboard shortcut

Cutting Tools


Cutting removes the selected tools from the composition and places them in the system clipboard. There are four ways to cut the currently selected tools:

- > Select Cut from the Tool context menu
- > Choose Edit > Cut from the Fusion menus
- > Click on the Cut icon in the toolbar
- > Use the Ctrl-X keyboard shortcut

Pasting Tools

Pasting will make a copy of the tool(s) currently in the clipboard in the composition..

- > Select Paste from the Flow context menu
- > Select Paste from the Tool context menu
- > Choose Edit > Paste from the Fusion menus
- > Click on the Paste icon in the toolbar
- > Use the Ctrl + V keyboard shortcut



The location of the pasted tools in the Flow 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 left-click in the flow and then use the CTRL + V keyboard shortcut, or to use the context menus. In either case, the tools will be placed at the location of the pointer when the mouse was clicked.

To replace a tool already in the flow with the contents of the clipboard, right-click directly on the tool to be replaced and then select Paste from the tool's context menu. 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 the pasted tool is a creator (one that does not have any inputs) the pasted tool will automatically be combined with the active tool using a Merge tool. This auto-merge behavior can be changed with the Default_Preferences panel in the Global Preferences.

Pasting Tool Settings

Fusion also offers the ability to paste the settings of the tool from the clipboard into another tool. This function has no keyboard shortcut but can be accessed by right-clicking directly on the tool into which the copied settings will be placed.

When copying the settings from one tool to another, only the settings that match between two tools will be copied. For example, when pasting settings from a merge to a transform, the center control values, the blend, and the size and angle controls will be copied from one to the other.


Using a Text Editor

As of Fusion 5.1 and later, tools can be copied from the flow and pasted into a Text Editor. This will display the tools' data as it would be saved to disk. The text can then be edited and pasted back into the flow, either as a new tool, or to replace the settings of an existing tool.

Renaming Tools

When a tool is added to the composition it will automatically be assigned a name that reflects the name of the tool followed by a number. For example, the first Blur tool added to a composition will be called "Blur1."

A tool can be renamed by selecting Rename from the tool's context menu or by using the F2 keyboard shortcut. If multiple tools are selected a dialog will appear asking for a name for each tool.




Valid tool names use alphanumeric characters only and cannot contain any spaces. A tool name may not start with a number. This ensures that the tools name remains a valid identifier in the eyeonscript (Lua) and Python languages. Fusion will automatically replace spaces with _ and strip away invalid characters.

Tool Defaults and Settings

Saving New Defaults

When a tool is added to the composition its controls are set to the default values. The artist can override the default settings with their own custom settings by right-clicking on the tool and selecting Settings > Save Default.



When Fusion is constructing the tool it first looks for a .settings file with the same name of as the tool in the folder defined by the Globals > Path Maps > Defaults preference. If one is found then the settings stored in that file will be used instead.

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

Removing a Saved Default

To remove a saved default for a tool either delete the .settings file from the Defaults folder or right-click on the tool and select Settings > Reset Default.

Saving Tool Settings

- > The current settings for one or more tools can be saved to disk for later retrieval. These settings can then be loaded from disk, added to the toolbars or bins and retrieved from the Tool context menu.
- > The following methods can be used to save settings for one or more tools.
- > Select and then right-click on the tool(s) and then choose Settings > Save As from the context menu. A file browser dialog will appear pointing at the folder specified by the Globals > Path Maps > Settings preference.

-
- > Select and then drag the tool(s) into the bins. This automatically displays the same dialog described above, and also adds a shortcut for the saved settings into the bin.
 - > 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.

Loading Saved Settings

To load a setting to an existing tool select Settings > Load from the tool's context menu and locate the .settings file using the file browser. The Settings submenu also displays a list of the last few .settings created with that tool for easy access.

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

To create a custom icon for a saved setting create a 32-bit bitmap (.bmp) with exactly the same name as the .settings file and save it to the same folder. The bitmap will be used as the icon for the file when it is added to the bins. The bitmap can be any size, but it is suggested to keep it small for best performance.


RoI and DoD

Fusion uses advanced techniques called Region of Interest (RoI) and Domain of Definition (DoD) to dramatically improve rendering.

Domain of 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 resolution 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 tool's effect to, pixels that lie outside the visible portion of the image. This concept is often referred to as an 'infinite workspace' or 'infinite canvas.'



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 Fusion would be required to process every pixel in the image. With a DoD Fusion is 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 in green.




For the most part, the DoD is calculated automatically and without the need for intervention from the artist. 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 will often be much smaller.

Media loaded from disk will typically also have the DoD default to the full image width. The exception is media stored in the OpenEXR or Fusion RAW file formats. These formats are capable of storing the DoD of the image, and Fusion will apply this when loading the image.


The DoD is established as soon as the image is created or loaded into the composition. From there it passes downstream, where tools combine it with their 'Region of Interest' in order to determine exactly what pixels should be affected by the tool. The Region of Interest concept is described in more detail later in this chapter.

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 mouse hovers over the tool tile for a moment in the flow, as pictured below.



The DoD is shown as two xy coordinates indicating the corners of a rectangular bounding box (in pixels). In the screenshot above, the DoD rectangle is from 116,94 to 204,14.

In addition to the DoD used by each tool, the tooltip also shows the current image dimensions. It may also display a field named 'data,' which describes the portion of the image that was last rendered by the tool after combining the DoD with the Region of Interest. Note that this information will only appear in the tooltip if the tool has rendered at least once.




It is also possible to set the DoD for an image manually using the Set Domain tool (found in the Miscellaneous category). 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 Set Domain tool can be used to animate a DoD that covers the character and ignores the rest of the image.

Sampling Outside the DoD

While the DoD's impact on Fusion composites is mostly positive, there are places where it significantly changes the way artists approach a composite. Most frequently this occurs in places where the effect must use or sample pixels that lie outside the current DoD.

The following images best demonstrate the issue. The first image shows an ellipse mask, the second shows what happens when using a Channel Boolean tool to invert the image without DoD, and the third shows the same result, but with a DoD.




As you can see, the Invert operation is applied only within the area of the DOD. This is usually not the desired result.

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 renderer what pixels are actually present in the image, the RoI tells the renderer which pixels actually need to be rendered. When a tool renders, it combines the current RoI with the current DoD to determine what pixels should be affected.

Artists often use the RoI to restrict rendering to a small region of the image. This can significantly improve performance when working with very large or complex compositions.



The RoI button in the 2D display view toolbar can be used to turn the RoI on or off for that display view. The Region > Enable Region option in the display views context menu will also toggle the RoI. Clicking on the small arrow next to the RoI toolbar button will display the same list of options shown in this context menu.

When RoI is enabled a rectangular RoI control appears in the display view. 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.



Click and drag on the edges of the rectangle to adjust the size of the RoI. Click and drag on the corners of the rectangle to adjust both dimensions at once. To move the RoI without adjusting its dimensions click and drag on the small circle found at the top left corner of the RoI rectangle.


A faster way to set the RoI is to select Region > Set and then click drag in the view to set a RoI rectangle directly.

While the RoI is active Fusion will only request rendering of the pixels inside the region when it displays an image in that view. Flipbook previews created in that view 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, OpenEXR, TIFF and many uncompressed file formats are some of the formats that support this feature.

Changes to the viewed image size will cause the pixels outside the RoI to be reset to the image's background 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.

The following images show a color correction with an RoI applied when Auto Proxy mode is on, and when it 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.

Working in 3D

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 in that all images are either captured using a camera with freedom in all three dimensions, in a scene that has real depth, or has been created in a true 3D modeling and rendering application.

Fusion's advanced 3D environment provides the artist with a complete OpenGL-accelerated 3D environment, including support for geometry, text and particles. This makes tasks such as the ones listed below much easier to perform:

- > Convert 2D images into image planes
- > Match cameras from 3D matchmoving software, such as Boujou, Syntheyes, and PF Track.
- > Import cameras from popular 3D software packages, such as Maya, Softimage XSI, LightWave and 3DS Max.
- > Create rough primitive geometry or import mesh geometry from FBX scenes.
- > Create realistic surfaces using illumination models and other real time-capable tools.
- > Rendering with realistic depth of field, motion blur and supersampling.
- > Use OpenGL-accelerated 3D particle systems.
- > Extrude and bevel text.
- > Shadow casting across geometry.



Fusion's 3D toolset includes over 40 different tools. This chapter examines the tasks they perform and discusses how these tools can be combined to produce effective 3D composites within Fusion.

The functionality of OpenGL rendering depends on the hardware. To provide a maximum range of compatibility there are some environment variables that can force Fusion to use specific Cg profiles. If you have some trouble rendering a scene using OpenGL or displaying it in the viewport, try to set the environment variables `CGGL_LATEST_FRAGMENT_PROFILE = arbf1` and/or `CGGL_LATEST_VERTEX_PROFILE = arbvp1`

The 3D Scene

Tools that output a 3D scene description instead of an image create a 3D scene. Examples are tools that generate primitive geometry, import geometry, and create lights and cameras. Tools that modify a 3D scene to warp, deform or duplicate geometry usually also output 3D scenes. Almost all of these tools are collected together into the 3D category of tools.

There is no need to specify whether a composition will be 2D or 3D when it is created, as a composition can seamlessly combine any number of 2D and 3D scenes together. 3D geometry is added to the scene using Image Plane 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 environmental maps.

When these tools are viewed in one of Fusion's displays, a 3D viewer is used to allow the artist to navigate and manipulate the scene in 3D. The interactive 3D viewer is highly dependent on the workstation's graphics card, relying on support for OpenGL and GPU rendering. The amount of on-board 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 tool cannot be connected directly to the input of a Blur. A Render 3D tool must first be used to create a 2D image from the scene for the Blur tool to operate.

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 tool and an Image Plane tool are separate until they are somehow combined into the same scene.

The Merge 3D tool 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 two inputs for background and foreground. 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 tool presents very few options in the control panel. The main control panel shows just a single checkbox that the artist can use to determine whether light present in the scene is passed downstream to the output. This mechanism is used to restrict light to a portion of the scene. By default this checkbox is disabled.


This allows the artist to light elements in the scene without worrying about how the light might affect geometry introduced further downstream. For example, the artist may wish 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 tool control panel shows the common 3D Transformation tab. These transformation controls 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.

Additionally, the Merge 3D tool control panel shows the common 3D transformation tab. These transformation controls 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.

As a result of the two features described above, the Merge 3D tool forms the basis of parenting for the Fusion 3D environment.

Most 3D primitives also provide a 3D scene input that can be used as an alternative to the Merge 3D tool.



Renderer 3D

In order to generate a 2D image from a 3D scene, the scene must first be rendered using a Renderer 3D tool.


The renderer uses one of the cameras in the scene to produce a 3D image. If no camera is found, a default perspective-style view is used. This default view rarely provides a useful angle and it cannot be adjusted, so most 3D scenes include at least one camera.

The image produced by the renderer can be any resolution, up to Fusion's 32,000 x 32,000 pixel limit. The Renderer 3D tool is equivalent to other creator tools, such as the Background, Fast Noise or Text tools. As such, it provides the usual options for fields processing, color depth, and pixel aspect.

Fusion provides two different types of renderer by default: a GPU-accelerated OpenGL renderer which produces results almost identical to those shown in the interactive 3D views, and a software renderer which is slower but supports more advanced rendering techniques such as soft shadows or extremely large textures.

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 card, ensuring that the results of the render are always the same whatever the configuration of the machine used to render.



The software render is required to produce soft shadows. This feature is not supported by the OpenGL renderer.


The software renderer can easily handle textures much larger than the graphics card's maximum texture size. It is recommended for maximum quality when working with large images as textures. As the computer systems used as render slaves in network render farms frequently have either weak graphics cards, or no graphics cards at all, the software renderer is generally preferred when rendering a composite through the render farm.

OpenGL Renderer

The OpenGL Renderer relies on the GPU and memory on the graphics card of the workstation to produce the image. The textures and geometry are uploaded through the systems bus to the graphics card, and OpenGL shaders are used to produce the image. This can be orders of magnitude faster than the software renderer, but limits the OpenGL renders to the capabilities of the graphics card.

With the introduction of very fast GPUs and performance-optimized shader languages like NVIDIA's CG language, the OpenGL renderer has become much more flexible. This speed makes techniques like super sampling and depth of field possible, when they would have previously been much too slow for production use. As a result, the OpenGL renderer in Fusion is very appropriate for final renderings.

Note that the OpenGL Renderer in Fusion respects the color depth option in the radioactive tab of the Renderer 3D tool. This can cause slowdowns on certain graphics cards when rendering to int16 or float32.




Displaying and rendering textures larger than the graphics card's maximum texture size is very difficult for OpenGL, and it either fails entirely, or slows the rendering immensely. Most OpenGL-accelerated cards available at the time this was written allowed a maximum texture size of 2048x2048 or 4096x4096, depending on how much memory was installed on the graphics card.

Textures larger than this maximum cannot be displayed directly by OpenGL. Fusion will resize images larger than the maximum to half the maximum texture size of the graphics card to ensure the image can be displayed in the views and rendered by the OpenGL renderer.


Contact the graphic cards vendor to determine the maximum supported texture, or use freely available utilities such as the OpenGL extension viewer.

3D Viewer


Viewing a tool with a 3D output in any display view will switch the display type to a 3D view. Initially, the contents of the scene will be displayed through a default perspective view. To change the viewpoint right-click in the view and select 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 left or right corner of the display view.



In addition to the usual Perspective, Front, Top, Left and Right viewpoints, the Camera submenu will also display all of the cameras and lights present in the scene as potential viewpoints. It is even possible to display the scene from the viewpoint of the geometry by selecting it from the context menu's Camera > Other submenu.



The view can't be switched to a viewpoint that is not present in the displayed tool's 3D output. Consider the flow in the image below. Two image planes and a spotlight tool are combined in a single Merge tool, which then connects to a renderer. Viewing either image plane in a view will show only the image plane generated by that tool. Viewing the merge will show both image planes.



For a detailed description of the options available in the 3D viewer, please see the 3D View section of the Display Views chapter.

Navigating the 3D View

For the most part panning and scaling of the 3D display view uses the same controls as the 2D view. To pan, hold the middle mouse button and drag the mouse. To dolly (zoom) the view, hold the `Doc_mouse_ImImiddle` and left mouse buttons and drag the pointer to the left or right. Holding down `CTRL` while spinning the scroll wheel will also dolly the camera.

It is also possible to rotate around the view by holding down the `Alt` key and the middle mouse button and dragging the mouse pointer to the left and right.

There are also some useful hotkeys for navigating:

A - Fit all objects in the view

F - Fit to selection (or Fit All if nothing is selected)

C - Rotates the view to look at the center of the currently selected object but doesn't move the view's position.

Furthermore, selecting a 3D tool in the flow area also selects the associated object in the 3D views. You can select a 3D tool in the flow, activate the 3D view (using the middle mouse button) and hit F to fit the view on the selected object.

Transforming with the Views

When the display view is set to look through an object in the scene, like a camera or spotlight, the usual controls for panning and rotating the view will now directly affect the position of the object.

For example, to adjust the camera position when looking through the camera, placing the mouse in the view with a `Doc_mouse_ImImiddle` and left mouse buttons drag will dolly the camera.

It is even possible to view the scene from the perspective of geometry like image planes and shapes by selecting the object from the `Camera > Others` menu. The same transformation techniques will then move the position of the object.

Transparency Sorting

While generally the order of geometry in a 3D scene is determined by their Z-position, sorting every face of every object in a large scene can take an enormous amount of time. To provide the best possible performance, Fusion defaults to using a Fast Sorting mode in the OpenGL and interactive displays. The Fast Sorting mode uses the order of the inputs connected to the Merge 3D to determine the sort order for all objects except the one closest to the viewer. 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 toggled using the Fast button in the display toolbar, or from the Transparency menu of the views context menu. The Renderer 3D tool also presents a Transparency menu when the Renderer Type is set to OpenGL.

The Software renderer type will always use the slower and more accurate method.



The basic rule is when a scene contains transparency, use the Full/Quick Sort mode, and otherwise use the Z Buffer. If the Full Sort method is too slow, try switching back to Z buffer.

Material Viewer

Viewing a tool that outputs a 3D Material in a display view will switch the display type to a Material viewer. The default state for this view shows the material applied to a lit 3D sphere rendered with OpenGL. The type of geometry, the renderer, and the state of the lighting can all be adjusted from either the display views toolbar, or the context menu. Each view presents 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 view and scaled to fit. It is not possible to pan or scale the Material viewer. The view can be rotated on a single axis to provide a different angle on the material, however, by holding down the Alt key and the middle mouse button and dragging the mouse pointer to the left and right.


Clicking and dragging in the view with the middle mouse button held down can adjust the position of the light. Light can also be toggled on and off using the display views toolbar button.

For more information on the Material Viewer refer to the Display Views chapter of this manual.



Transformations

Most of the tools found in the 3D category have Transformation controls that are collected together into a single tab. The controls found in this tab affect how the object is positioned, rotated and scaled within the scene. The Transformation tab is labeled with a typical 3D coordinate tripod as its icon, as pictured below.



The Offset controls are used to position the object in the scene. The Rotation controls affect the object's local rotation (around its own center). The Scale sliders affect its size.


Onscreen Transformation Widget

When an object is selected, it will display an Onscreen Transformation Widget in the display views. The widget has a position, rotation and scale mode. Use the Transform toolbar buttons to switch the mode of the selected control or press q for Position, w for Rotation and e for Scaling.

In all three modes, the red portion of the widget indicates the object's local X-axis, while green and blue are used for the Y- and Z-axis respectively. Click-drag in the center of the widget to apply the transformation without constraint, or click-drag directly on an axis to constrain to that axis. Holding the ALT key down while click dragging in the view allows a selected control to be adjusted without requiring the artist to click directly on the control.



If the Scale's Lock XYZ checkbox is still selected, only the overall scale of the object can be adjusted (using the red axis). Click-drag on the center of the widget or on the red axis. Adjustments will affect all three axes equally. Unlock the controls to scale the object along individual axis separately.




Pivot

In 3D, objects are rotated around a vector (the so-called rotation axis). A vector gives just the direction. By default, this vector goes through the object's center. Force the rotation axis to go through any 3D point using the Pivot control.

Target

Enabling the Use Target checkbox will reveal another set of XYZ controls and a new position control will appear in the display views. A dashed line is drawn between the target and object center in the display views to show the relationship between these controls.




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 to face its target.

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 to Fusion's node-based approach to 3D is that Parenting between objects becomes implicit in the structure of the 3D flow. The basis for all Parenting is the Merge tool.

For example, picture a scene with two cameras that are both pointing at exactly the same location, except one is offset from the other by a small amount. This type of rig was frequently used to simulate left-eye/right-eye stereoscopic cameras before this feature was built into cameras.




Cameras

One way to connect cameras together would be to publish the XYZ Offset controls and the XYZ Animation controls on one of the cameras, and then connect the matching controls on the other camera. A much faster way would be to use a Merge 3D to create a parent transformation for the two cameras, as depicted in the composition below.



In the illustration above, the Merge 3D tool has been renamed to Camera Rig to make its purpose clearer. Anytime the camera rig is adjusted, apply a transformation to the merge's center. To change the relationship between the cameras, for example to offset them further from each other, adjust the individual camera's transformation.

Any transformation applied to a merge is applied to all objects present in the scene at that merge, including cameras, lights and geometry.




A frequent task for compositors involves matching the cameras from footage or from a 3D scene in the composite. To simplify this task, Fusion provides a flexible Camera tool. Each camera can be used as either a free camera or a target camera. Select and adjust Angle of View, Focal Length, Aperture and Clipping planes to match the behavior of real scenes or 3D cameras.

To allow for exact matching between applications, Fusion can import camera paths and positions from a variety of popular 3D applications.

Importing Cameras

To import a camera from another application, locate and click the Import Camera button. Using the file browser that appears, locate and select the scene in one of the supported 3D scene files. A dialog box with several options will appear.



If multiple cameras are present in the imported scene, it is possible to choose the one to import from the list on the left. Checkboxes are provided to select the channels of motion that will be imported.

Fusion is able to import animation splines from Maya and dotXSI directly to their own native spline formats. Animation applied to cameras from 3Ds Max and LightWave must be sampled on each frame. When the Force Sampling checkbox is enabled, Fusion will sample each frame of the motion, regardless of the format.

Applications Supported for camera Import

3Ds Max (.ASE)

Autodesk's 3Ds Max uses a binary file format so Fusion is unable to read the camera directly from the saved scene. Instead, export the camera to an ASCII ASE file. Consult the application's documentation for details. Fusion must sample the camera motion, generating a keyframe for each frame.

LightWave (.LWS)

NewTek's LightWave scene file is text based so Fusion can easily extract information about cameras from the scene directly. Fusion must sample the camera motion, generating a keyframe for each frame.

Maya (.MA)

Alias's Maya has a native ASCII scene format, so Fusion can extract information about cameras, lights and locators directly from the .ma file. Fusion can import Maya motion as native splines, only importing as many keyframes as seen in the original spline.

Softimage XSI (.dotXSI)

Softimage uses a binary format for its native scene so Fusion can import the scene to a text-based format using the dotXSI format. Fusion's support for dotXSI allows it to import native splines, only importing as many keyframes as are present in the original spline.

Filmbox (.FBX)


Fusion can import cameras from FBX scenes using the Import FBX option in the File menu. This creates a camera with the correct settings. The Import Camera button in the Camera3D tool cannot be used to import from FBX at this time. See below for details.

Parented Cameras

At this time, Fusion is unable to import many Parented or rigged Cameras. If you have trouble importing a camera from any scene, try baking the camera animation in the 3D application before importing it into Fusion.

Lighting

Fusion provides four different types of lights for use in 3D scenes: ambient, directional, point and spotlights. These lights can be combined to form very detailed lighting environments. A scene without lights will use a default directional light when the scene is displayed in the views. The default light will disappear once a light is added to the composition by an artist. To ignore lighting in the display views, toggle the Light button in the display view toolbar.



Lighting must be enabled in the Renderer tool. Tools that create or load geometry also present lighting options which can be used to choose whether an object is affected by light or not. When lighting is disabled in either the display views or final renders the image will appear to be lit by a 100% ambient light.

Ambient Light

Ambient light 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. Use ambient light to set a base light level for the 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 seemingly directional light source.

Point Light

A point light is a well-defined light that has a clear source. A light bulb is an excellent example of a point light.

Spotlight


A spotlight is an advanced point light that produces a well-defined cone of light with falloff, and generally decays over distance. This is the only light type in Fusion capable of producing shadows.



All of the light tools display on-screen widgets in the views when selected. In the case of the ambient light, the position, rotation, and scaling of this widget has no affect on the results. The directional light can be rotated, but position and scale will be ignored. The point light ignores rotation. All three transformations apply to the Spotlight tool.

Light Hierarchies


Lights normally do not pass through Merges. This provides a mechanism for controlling which objects are lit by which lights. For example, consider 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.



To have lights pass through merges to affect downstream objects, click on the Merge tool's Pass Through Lights option.


Per-Vertex Lighting vs. Per-Pixel Lighting

When the OpenGL renderer calculates the effect of lighting on geometry it defaults to the per-pixel model, which calculates the lighting for each individual pixel. This is required for any advanced shading or projection. An older option called per-vertex is also available, which calculates the effect only at the vertices of the geometry. This mode can be enabled in the display view context menu's Lighting Calculation submenu, or in the Render tool's control panel. The per-vertex lighting mode is provided for compatibility with older compositions only, however, and should be avoided in new compositions.



Lighting Options

Most tools that generate geometry have an extra panel of options under the roll out labeled Lighting. These Lighting Options are used to determine how the object reacts to lights and shadows in the scene.



Affected By Lights

If this checkbox is enabled, lights in the scene will affect the geometry.

Shadow Caster

If this checkbox is 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 until shadows are enabled using the display view toolbar button. Shadows will not appear in the output of the Renderer tool unless the Shadows option is enabled for that renderer. To prevent a spotlight from casting shadows, disable the Enable Shadows checkbox in the tool's control panel.

See the Spotlight tool reference for a more detailed description of the shadow controls.



Shadow Maps

The shadow map is one or more internal depth map layers. Each pixel in a depth map specifies the pixel's depth in the scene. This is used to assemble the shadow layer that is inserted into the scene.

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 shadow, or the more falloff in the cone, the larger the shadow map will need to be to produce useful quality results. The shadow map is square. Set the value of the Shadow Map Size control to set the size of the depth map in pixels.

Generally, there is a point of diminishing returns where increasing the size of the shadow map no longer improves the quality of the shadow. This point can usually be found by trial and error. 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 when Fusion is doing fast interactive previews, such as autoprox and LoQ renders. A value of .4, for example, represents a 40% proxy.

Multiplicative and Additive Bias

Shadows are essentially textures applied to objects in the scene, occasionally resulting in Z-fighting where the portions of the object that should be receiving the shadows render over top of the shadow instead. The top picture below shows typical shadow map Z-fighting, with the corrected shadow shown below.




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.

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. Animating the bias may be required to get a proper result for the frames.

Shadow Softness

By default, the spotlight generates shadows without soft edges. Additional options produce shadows with constant and variable 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. Many artists use softness to cheat the aliasing rather than increasing the shadow map to preserve memory and avoid exceeding the graphics cards capabilities.



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 shadows 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.



Selecting the Variable option causes the Spread, Min Softness and Filter Size sliders appear. 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 shadows distance from the geometry increases. These controls are used to limit the shadowmap 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 the Spotlight tool reference for additional details on these controls.

Transmittance

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. Selecting the Force All Materials Non-Transmissive overrides this, causing the shadow map produced by this tool to ignore transmittance entirely.

See the Materials section for more details.

Materials

In order to render a 3D scene, the renderer must take into account the shape of the object as well as the appearance. The geometry of an object could be said to describe the shape of the object, while the material applied to the object describes its appearance. Fusion comes with a range of options for applying materials and textures to geometry.

Tools that describe the geometry's response to light are called Illumination models. Typically, these tools are named for the authors of the mathematical technique used by the tool. The Blinn, Cook Torrance, Ward and Phong tools are the most frequently-used illumination models. These tools are found in the 3D > Material category of tools.

Most materials also accept textures, which can be used to refine the look of an object further, giving it a rough organic appearance or a smooth reflective sheen. Textures can be simple 2D images or they may be more complex bumpmaps, heightmaps and reflection maps. These tools are typically found in the 3D Texture category.


Materials can also be combined together to produce elaborate and highly-detailed composite materials. The term shader is often used when describing materials that are combined in this way. Materials are calculated on the graphics cards GPU whenever possible, and to do this Fusion compiles all the Material tools into a program called a shader, which is executed on the graphics cards GPU.

Each tool that creates or loads geometry into the 3D scene also assigns a default basic material, with controls built directly into the tool. The default basic material is a simplified version of the Blinn illumination model. It is possible to override this basic material using one of several tools that output a 3D Material. These materials provide a greater degree of control over how the geometry reacts to light, providing inputs for diffuse and specular texture maps, bumpmapping and environmental maps, which mimic reflection and refraction.



Illumination Models

Illumination models are advanced materials for creating realistic surfaces like carbon fiber, 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. In Fusion there are four different models: Blinn, Phong, Cook Torrance and Ward. The image below shows their characteristics, in the order they were named.



Basic

The Basic material provides basic control over the diffuse, specular and transmittance components of the material. It only accepts a single texture map for the diffuse component. The Basic Material controls are found in the Material tab of all tools that load or create geometry. Connecting any tool that outputs a 3D material to that tool's Material input will override the Basic material and the controls in the Material tab will be hidden. See the Common 3D Controls section of the Tool reference for a complete description of this material's options.

Blinn

The Blinn material is a general-purpose material that is flexible enough to represent both metallic and organic surfaces. At first glance it appears identical to the Basic material, but the Blinn tool allows for a greater degree of control by providing additional texture inputs for the specular color, intensity and exponent (falloff), as well as for reflection and refraction environment maps and bumpmap textures. Blinn is generally more efficient than the Phong illumination model.

Cook Torrance

The Cook Torrance material uses Fresnel equations to simulate the behavior of light when moving between media of differing refractive indices. Otherwise, it is similar to the Blinn illumination model. This illumination model is primarily used for shading metal or other shiny and highly-reflective surfaces.

Phong

The Phong material produces a result that is commonly thought to model the behavior of light against polished or plastic surfaces.

Ward

The Ward illumination model is ideal for simulating brushed metal surfaces, as the highlight can be elongated in along the U or V directions of the mapping coordinates. This is known as an anisotropic highlight.

Material Components

Regardless of whether the object uses the Basic, Blinn, Cook Torrance, Phong, Ward materials, or a shader that combines several materials, all illumination models share certain characteristics that must be understood.


Diffuse

The diffuse component of a material describes the appearance of an object where light is reflected from the object at an angle that leads away from the current viewpoint. This diffuse color and texture could be considered to be the base appearance of an object, before taking into account highlights, reflections, and bumpmaps. The opacity of an object is generally set in the diffuse component of the material.


Specular

The specular component of a material describes the appearance of an object where the light is reflected directly back at the current viewpoint. This causes a highlight that is generally brighter than 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 usually described as color, intensity and exponent. The specular color determines the color of light that reflects from a shiny surface. Specular intensity describes how large the highlight will be. The following images shows a split view of a sphere with 50% white as the diffuse color, and 100% white as the specular color. The right side of the sphere has the specular intensity set to 0, while the left side shows the image with specular highlights.



The specular exponent controls the falloff of the specular highlight. The smaller the value, the sharper the falloff and the larger the specular component will be. The following image shows the sphere from the above example with a specular exponent of 9 and again with an exponent of 200.




Transmittance

The transmittance component describes 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 component is essential to creating the appearance of stained glass. This component is described in greater detail below.

Note that Fusion allows for adjusting the opacity and transmittance of a material separately. The opacity determines how transparent the actual surface is when it is rendered. This might be a bit counter-intuitive to artists unfamiliar with 3D software at first. 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.


The illustrations below cast a shadow through the following image. The image contains a green-red gradient from left to right. The outside edges are transparent, and the inside is a small semi-transparent circle.



Transmissive surfaces can be further limited using the Alpha and Color Detail control.


Alpha Detail

When this slider is set to 0, the alpha channel of the object is ignored and the entire object casts a shadow. If it is set to 1, the alpha channel determines what portions of the object cast a shadow. Note that the OpenGL view will always cast a shadow from the entire object, ignoring the alpha. Only the software renderer supports alpha in the shadow maps. Below is an example image with Alpha Detail set to 1. 0. The shadow is only cast where the alpha channel is greater than 0.



Color Detail

Color Detail modulates light passing through the surface by the diffuse color and texture colors. Use this to throw a shadow that contains color details of the texture applied to the object. Increasing the Color Detail slider from 0 to 1 brings in more of diffuse color and texture color into the shadow. Note that the alpha and opacity of the object is ignored when transmitting color, allowing an object with a solid alpha to still transmit its color to the shadow. Below is an example image with Color Detail set to 1.0. The gradient is clearly visible in the shadow.



Transmission Color


The Transmission Color determines how much color is passed through the object. For an object to have transmissive shadows, set the Transmittance color to (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.

Textures


Texture maps are typically either 2D images or 3D materials that are used to modify the appearance of a specific component in an Illumination model. This is done by connecting an image (or another material) to the inputs on the Material Tools tile in the Flow Editor. When a 2D image is used, the UV mapping coordinates of the geometry are used to fit the texture over the image, and when each pixel is rendered the material will multiply the value of the component against the value of the corresponding pixel in the map.

Texture maps are used to modify the diffuse color, specular color, specular exponent, specular intensity, reflection, refraction and bumpmap components of a material. The most commonly used texture maps are applied to the diffuse color component.

The following example shows how to control the roughness of a Cook Torrance material using the image produced by a Fast Noise tool.




A tool that outputs a material is frequently used instead to provide a finer degree of control over the texture. In the following example a Texture 2D tool is used to translate the texture in the UV space of the object.




Composite Materials

Building complex materials or shaders is as easy as connecting the output of a Material tool to one of the Texture inputs of another material or Texture tool. When a Texture input is supplied with a material instead of a 2D image, it will combine the results of that material with its own calculation. The following example shows how to combine an anisotropic highlight with a Blinn material.



The output of the Blinn, including its specular, is used to determine the diffuse color of the Ward material. There are times where this is not desirable because the output of the Blinn is relit by the Ward material. The Channel Boolean Material can be used to add the Ward material's specular component to the Blinn material with a greater degree of control.



Reflections and Refractions

The illumination models in Fusion support the use of environmental reflection maps created by the Reflect tool in the 3D Material tool category. This tool creates an environmental reflection map, which can be used to simulate reflections and refractions on an object. Reflections are direct reflections of the light that hit an object, while refractions simulate the distortion of light seen through semi-translucent surfaces.

The reflections and refractions produced in Fusion are not accurate, ray traced solutions, as these would be much too slow in the context of a compositor.

Instead, Fusion uses a technique common in games called environment mapping to produce an approximation that is good enough for many scenarios. Environment maps assume an object's environment is infinitely distant from the object and renders it into a cubic or spherical texture surrounding the object.

Fusion provides the Cube Map and Sphere Map tools to help create environment texture maps, which will produce more realistic results than a simple flat texture mapped directly to the object.




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 using 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 tool will not respond to lights in the scene. For this reason, the Reflect tool is usually combined with the Blinn, Cook Torrance, Phong or Ward tools.

Additionally, reflections and lighting may not interact in a realistic fashion. Environment mapped reflections are a texturing effect that is blended over the background color. That makes it possible to apply the reflection either before or after the lighting model. The images below show different approaches to adding reflections to a material.

In this example, the Blinn material is connected to the Background Material input of the reflection. This causes the Blinn to light the reflection texture.



In this example, the Reflection texture is connected to the Diffuse Color component of the Blinn, causing the texture to be multiplied by the diffuse color.



The Channel Boolean Material can be used to provide more control over how these are combined, as discussed in the Sphere Map or Cube Map is connected then the reflection/refraction vectors created by the tool are used to map the texture to the object. Connecting a 2D image or a texture created by the Create Texture tool causes the Reflect tool to use the standard UVW texture coordinates instead.

The Reflect tool provides the following texture inputs:

- > **Reflect. Reflection. Color.** Material typically receives the environment reflection image
- > **Reflect. Reflection. Intensity.** Material typically receives a texture for scaling the constant/by angle strength of the reflection
- > **Reflect. Refraction. Tint.** Material typically receives the refraction environment image.


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 the any material or texture used for the Background Texture input.

Working with reflection and refraction can be tricky sometimes. Here are some hints to make work easier.

- > Keep it subtle. Typically, environment mapped reflections use just a hint of reflection, between 0.1 and 0.3 strength. Use higher values only for glossy surfaces like chrome.
- > Bumpmaps can add detail to the reflections/refractions. Note that a bumpmap also lacks an Illumination model, so connect it to the Illumination model used by the Reflect tool.
- > When detailed reflections are not required, use a relatively small cube map, such as 128 x 128 or 256 x 256 pixels, and blur out the reflections.
- > In photorealistic renderers the reflection is added to the lighting information of the surface. In Fusion it will blend between the reflection color and the color of the surface. For behavior closer to the photorealistic renderers, create the reflection separately and add it to the Illumination model using a Channel Boolean Material tool.
- > The refraction is evaluated only when the ray hits the object. There is no additional computation when the ray would leave the object.
- > In practice it can be hard for the human brain to distinguish between a physically correct (ray traced) refraction or an environment refraction. As long as everything looks warped, the brain doesn't notice the precise optics.
- > The alpha of refracted pixels is set to 1 even though the pixels are technically transparent.
- > If the refraction is not visible even when a texture is connected to Reflect. Refraction. Tint. Material, double-check the alpha/opacity values of the background material.

Bumpmaps


Bumpmapping helps add details and small structures to the object using textures. Bumpmapping is a lighting effect only; it does not modify the geometry of the object or change its silhouette. Bumpmapping requires that lighting be on for it to work (as it is a lighting effect), and there must be a lighting model. Bumpmapped reflections/refractions do not require lighting to be on.



Bumpmaps can be applied to materials and objects using the Bumpmap tool. Typically an image containing the bump information is connected to the Bumpmap tool, and this one is connected to the bump input of an illumination model or another material node. There are two ways how to create a bumpmap for a 3D (material) tool.

Using a heightmap

A heightmap is an image where the color value of a pixel represents the height value. 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 luminosity of a pixel in the heightmap that determines the bumpiness, but rather how the luminosity changes in the neighborhood of a pixel. This is why you need to set the scale or detail high on a Fast Noise tool, for example, in order for it to produce bumpiness, since there is not enough variation from pixel to pixel.




Using a Bumpmap

A bumpmap is an image containing normals stored in the RGB channels used for modifying the normals of the geometry.

Typically bumpmaps you author from other 3D applications or image processing software are int8. Normals are 3 float values (nx, ny, nz) whose components are in the range [-1, +1]. Because you can store only 0, .255 or [0, 0, 1, 0] in an int8 image, the normals are compressed/packed from the range [-1, +1] to the range [0, 1] by multiplying by 0.5 and adding 0.5. Compressed/packed and uncompressed/unpacked are terminology used when dealing with bumpmaps. You can use a Custom Tool to do the decompression/compression. When moving from color to normal you want to do $(nx, ny, nz) = (2*r-1, 2*g-1, 2*b-1)$. When moving from normal to color you want to do $(r, g, b) = (0.5*nx+0.5, 0.5*ny+0.5, 0.5*nz+0.5)$.

The Create Bumpmap tool expects a compressed bumpmap.

Connecting a bumpmap directly to the Bump input of a material will result in incorrect lighting. Fusion uses a different coordinate system for doing the lighting calculation. The Create Bumpmap tool will do the conversion of the bumpmap to work correctly.



The Export Bumpmap Tool

The Create Bumpmap tool helps to convert a heightmap into a bumpmap. This is a 2D tool, and does this conversion only. The output will be an image; BumpMap produces a 3D material. This allows you to manipulate the bumpmap with Fusion's image processing tools. When you are done, use the Create Bumpmap tool in Bumpmap mode to connect it to Material nodes.

Help, my bumpmapping isn't appearing or looks bad!

- > Make sure you have the tools connected correctly. The heightmap/bumpmap should feed into a Create Bumpmap tool and then, in turn, should feed into the Bumpmap input on a lighting model tool.
- > Check that the original image has enough detail/contrast. Smoothly varying heightmaps, like a default Fast Noise, don't have enough variation from pixel to pixel to visibly affect the per-pixel normals.
- > Change the color depth of the heightmap to get better normals. You can also play with the Bumpmap Texture Depth options on the Create Bumpmap tool itself.
- > Play with the height scale on the Bumpmap tool. If you are feeding it a black and white image like Text+, then a heightscales of 1 works and a heightscales of 50. 0 makes it look horrible. If you are feeding it a Fast Noise, then you may want to jack up the heightscales to 50 to get the bumpiness.
- > Make sure if you are feeding the Bump tool a heightmap, it's set to heightmap, and if you're feeding it a bumpmap, it's set to bumpmap.
- > Check if HiQ is on/off. Some Fusion tools like Text+ produce an antialiased version in HiQ that will substantially improve bumpmap quality.

Projection Mapping


Projection is a technique for texturing objects where images are mapped onto objects using a Camera or Projector tool. This can be useful in many ways: texturing objects with multiple layers, applying a texture across multiple separate objects, projecting background shots from the camera's viewpoint, image-based rendering techniques and much more.

Projected textures can be allowed to slide across the object. If the object moves relative to the Projector 3D, or alternatively, by grouping the two together with a Merge 3D they can be moved as one, and the texture will remain locked to the object.

Basically, there are three different ways to do projection mapping.


Using the Projector/Camera Toolset to Project Light

In this mode the Projector or the Camera acts as a light with all lighting features. This mode requires per pixel lighting and lighting must be enabled in order to see some results. It is possible to choose if the projector behaves like a spotlight or an ambient light. Alpha channels cannot be projected. Overlapping projections add together like any other light tool. An internal clipping plane (at around 1.0 distance from camera) limits how close the Projector/Camera can get to the projection receivers. Have a look at the Projector tool documentation for more details.




Using the Projector/Camera Toolset to Project a Texture onto a Catcher Material

In this mode the projecting tools act as real texture projectors; there is no light contribution anymore. As a consequence, a Catcher tool 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. Overlapping projections can be combined in various ways (mean, median, blend) via the Catcher tool. An internal clipping plane (at around 1.0 distance from camera) limits how close the Projector/Camera can get to the projection receivers. This mode needs per pixel lighting.




Using the UVMaP Tool with its MapMode Set to Camera/Projector

This mode requires a camera and an UVMaP tool downstream of the objects to which the texture is being applied. The UV Map tool gathers the information from the camera and creates a new UV-map for the upstream objects; these are used for texturing. In order to work properly, the object must be tessellated sufficiently.



Textures are assigned to the object or its materials like any other texturing technique. It supports alpha channels as well as bumpmapping. The UVs can be locked to the vertices at a chosen frame. This locking only works as long as the mesh is not changing over time, since Fusion must be able to match up the mesh at the reference time and the current time. Its transform can be changing, but the mesh itself cannot. To be more specific, vertices may not be created or destroyed or reordered (e.g., projection locking will not work on particles because they get created/destroyed, a cube with its tessellation slider animated, a duplicate tool with its time offset set to non zero so addition meshes get created over time)..



Geometry

Primitive Shapes

Fusion supports several basic primitive shapes for assembling a 3D scene.

Image Plane

The Image Plane is the basic tool used to place a 2D image into a 3D scene.

Shape

Shape can create planes, cubes, spheres, cylinders, cones and toruses.

Cube

The Cube tool creates a cube with six inputs that allow mapping of different textures to each of the cube's faces.

FBX Mesh

The FBX Mesh tool is used to import mesh geometry from an FBX file. At this time, only static geometry is supported. This tool can only import triangles. The model may need to be converted to triangles when exporting it to the FBX format. The first texture applied to the mesh will also be imported, if available.

Text 3D


The Text 3D tool is a 3D version of the 2D text tool. This version supports beveling and extrusion but does not have support for the multi-layered shading model available from the 2D text tool.

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 3D renderer instead of the particle renderer. See the Particle Suite chapter for more detail.

Visibility

The Visibility options are found on all 3D source tools. They control object visibility under various conditions.



Visible

If the Visibility checkbox is not selected, the object will not be visible in the display views, nor will it be rendered into the output image by the Renderer 3D tool. A non-visible object does not cast shadows.

Unseen by Cameras

If the Unseen by Cameras checkbox is selected, the object will be visible in the display views (unless the Visible checkbox is turned off), except when viewed through a camera. The object will not be rendered into the output image by the Renderer 3D tool. Shadows cast by an Unseen object will still be visible when rendered by the software renderer, though not by the OpenGL renderer.

Cull Front Face/Back Face

Use these options to cull (eliminate) rendering and display of certain polygons in the geometry. If Cull Back Face is selected, all polygons facing away from the camera will not be rendered, and will not cast shadows. If Cull Front Face is selected, all polygons facing towards the camera will likewise be dropped. Selecting both checkboxes has the same effect as deselecting the Visible checkbox.


Ignore Transparent Pixels in Aux Channels

In previous versions of Fusion, the SW/GL renderers rejected transparent pixels. To be more specific, the SW renderer rejected pixels with $R=G=B=A=0$, and the GL renderer rejected pixels with $A=0$. This is now optional. The reason you might want to do this is to get aux channels (e.g., normals, z, UVs) for the transparent bits. For example, suppose in post you want to replace the texture on a 3D element that is transparent in certain areas with a texture that is transparent in different areas; it would be useful to have transparent bits set aux channels (in particular UVs). As another example, suppose you are doing post DoF. You will probably not want the Z channel to be set on transparent areas, as this will give you a false depth. Also keep in mind that this rejection is based on the final pixel color including lighting, if it is on. So if you have a specular highlight on a clear glass material, this checkbox will not affect it.

FBX Format


The Filmbox FBX format is a scene interchange format first developed by Kaydara to facilitate the transfer of information from motion capture systems to popular 3D systems. Since Kaydara was acquired by Alias, this format has been made available for more general use and is employed extensively as a way of moving scene information from one scene to another.

For exact details on importing FBX scenes, see the Importing Appendix later in this manual.



Importing An FBX Scene

Use the Import FBX utility found under the File>Import menu to create a composition with the same lights, cameras and basic geometry found in an FBX file.



Importing FBX Geometry

The FBX Mesh tool is used to import mesh geometry from an FBX file. At this time, only static geometry is supported. The first texture applied to the mesh will also be imported, if available.

The imported geometry may be enormous compared to the rest of the scene. Different 3D applications use different units to measure their 3D scenes, and Fusion will treat the scales as equal to its own system. For example, one popular 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.

Use the scale slider in the FBX Mesh tool to reduce the scale of such files to something more reasonable for the scene.

FBX Export

You can also export a Fusion 3D scene to other 3D packages using the FBX Export tool. The exporter allows writing into different file formats like *. dae or *. obj. The tool also allows export of Geometry, Cameras and Lights, and export of animation. The animation data is baked frame by frame. Textures and materials cannot be exported. The following example shows a simple scene in Fusion and the exported .fbx in Softimage. Each sphere is a single object after the export.

Fog 3D and Softclipping


The Fog 3D tool helps to create realistic fog in the 3D scene. It is based on the depth channel, therefore this tool does not create a volumetric fog.




The Fog 3D tool works well with depth of field and antialiasing supported by Fusion's OpenGL renderer. Since it is not a post processing tool (like the Fog tool found in the Tools/Deep Pixel menu), it does not need additional channels like Coverage or Background color. Furthermore, it supports transparent objects.

Matte Objects

Matte Objects are used for masking out parts of the scene using objects. 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. Matte objects are not boolean objects; they do not change or create any geometry. The effect of a matte object is global to the scene. To change the properties for downstream tools, use the Override tool.





Is Matte

When activated, objects whose pixels fall behind the matte object's pixels in Z do not get rendered.


Opaque Alpha

Sets the alpha value of the matte object to 1. This checkbox is only visible when the IsMatte option is enabled.

Infinite Z


Sets the value in the Z channel to infinite. This checkbox is only visible when the IsMatte option is enabled.

Matte objects in the view cannot be selected in the display view unless the Show Matte Objects view option is selected. However, it is always possible to select the matte object by selecting its tool in the flow.




Material and Object IDs

Any tool in Fusion can use the Object ID and Material ID auxiliary channels to generate a matte or mask. The settings used to accomplish this are found in the Common Controls tab of each tool.




Both IDs are assigned automatically upon tool creation 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. The Override 3D option offers a comfortable way to change the IDs for several objects. The Renderer will write the assigned values into the framebuffers during rendering, when the output channel options for these buffers are enabled. Empty pixels have an ID of 0. It is possible to assign a value of 0 manually to an object or material, but this is not advisable because a value of 0 tells Fusion to set an unused ID when it renders. Fusion allocates a 16-bit image buffer for rendering the ID channels. The maximum value of 65535 is reserved internally, so it is possible to use a value range from 0 to 65534.




Point Clouds

The Point Cloud tool was designed specifically to make it easier to work with locator clouds generated from 3D tracking software. Applications, such as Syntheyes, Boujou and Maya, will often generate hundreds or even thousands of tracking points. Seeing these points in the scene is important to assist with lining up the image, but bringing each point in as an individual locator or geometry can impact the application's performance dramatically.



The Point Cloud tool can import point clouds written into scene files from matchmoving software. Importing a point cloud works like importing a camera browse to the scene file and selecting the cloud to import from the scene. The entire point cloud is imported as one piece of geometry, which is a significantly faster approach.

Many 3D trackers allow for the naming of individual tracking points, as well as setting tracking points on points of interest. The Point Cloud tool will quickly find these points and publish them. A published point in the cloud can be used to drive the animation of other parameters.




The point cloud is represented as onscreen controls, not geometry. To see it in the views, first select the Point Cloud tool in the flow.

Finding a Point in the Cloud

To find a point in the point cloud, select the Point Cloud tool and right-click in a 3D View. Select Find from the Point Cloud's submenu at the bottom of the context menu. Type the name of the point and click OK.

If a point that matches the name entered is found, it will be selected in the point cloud and colored yellow.



This is a case-sensitive search. A point named tracker15 will not be found if the search is Tracker15.

Renaming a Point in the Cloud

Use the Point Cloud context menu to rename a selected point. This only works for a single point. A group of points cannot be renamed.

Publishing a Point


Use the context menu to publish a selected point or points. A published point exposes its XYZ positions for connections to other controls in the scene. This is useful for connecting objects to the motion of an individual tracker.

Stereoscopic

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 tool for the (virtual) left and right eye. Fusion supports several methods for displaying the 3D-scene in stereo. Have a look at the Anaglyph tool documentation for more details. Furthermore, it is possible to combine two different cameras for a stereo camera rig, or assign different materials for the object when viewed through the left or right camera.




The stereo functionality is also available when working interactively with the scene. It is possible to examine objects, materials and lights in stereo without having a camera or renderer.




Stereo Camera

There are two ways to setup a stereoscopic camera. The common way is to simply add a Camera tool and adjust the eye separation and convergence distance parameters.




The other way is to connect another camera to the RightStereoCamera input port of the Camera tool. 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 materials per eye.



The Position tools in Fusion offer a completely new way of working with Masks and Volumetrics for footage containing XYZ Position channels.

ZtoWorld offers the option to create those channels out of a Z Channel and 3D Camera information.

For the sake of this overview, we will refer to the World Position Pass as WPP.


Position Pass Concept

What is a WPP?

The WPP reflects each pixel's XYZ position in the original scene as RGB color values. So if the face from which the pixel was derived from in the original scene sits at 0/0/0, the resulting pixel will have an RGB value of 0/0/0 and thus will be black. If said face sits at 1/0/0 in the original scene, the resulting pixel will be fully red.


Of course if the coordinates of our face are something like -60/75/123, the pixel will hold those values as well. As you can see here, due to the huge extent 3D scenes can have, the WPP channel should always be rendered in 32-bit floating point to provide the accuracy needed.

This image shows a 3D Rendering of a scene with its center sitting at 0/0/0 in 3D Space and the according WPP. For better visibility the WPP has been normalized.




Different Coordinate Spaces

WPPs can be rendered in different Coordinate Spaces. These include World Space, Eye Space and Object Space. The image below depicts how those different spaces look. However, the tools in Fusion require the WPP to be rendered in World Space.



The Scene Input

The tools offer a Scene input, which can either be a 3D Camera or a 3D scene containing a camera. While the camera is vital for the ZtoWorld tool, Volume Mask and Volume Fog will still generate their output without any camera attached or with the camera position set to 0/0/0. However, connecting such a camera, which lines up with the original camera the WPP has been rendered from, or setting the camera's position manually greatly improves accuracy and look of the resulting Fog.



The “Invisible Sphere”


The example scene shown so far has an empty background, meaning there is nothing in the scene apart from the ground plane and the cubes.

If applying fog to a scene like that which is larger than the ground plane, the result will look similar to the “w/o Sphere” example shown below, because with no WPP information outside the ground plane the resulting value will be 0/0/0, so the fog fills that area as well.

To get around that you can add an invisible bounding sphere to your scene to create “dummy” WPP values to help the fog tool create the correct volume as shown in “with Sphere” below.



Polylines, Effect Masks and Motion Paths




Polylines appear whenever a control is animated with a path or when a tool's effect is masked with a polygon or baseline mask. They also appear in the Paint and Grid Warp tools. In a more basic form, polylines are used to control the animation of every single animation in Fusion, through the Spline Editor.

These simple shapes are used by Fusion for just about everything. As a result, they are extremely flexible, with a considerable amount of controls, modes and options.

This chapter offers an overview of polylines and their operation. It will not discuss how to create paths or effect masks, as these are covered in depth in the Effect Masks chapter and the Motion Paths chapter later in this manual.

Polyline Types




Bezier Polylines

Bezier Polylines are shapes that are composed of key points and Bezier handles. Several points together are used to describe the overall shape of a polyline.

Each key point has a pair of Bezier handles associated with it that is used to define the exact shape of the polyline segments passing through each key. Adjusting the angle or length of these Bezier handles will affect whether that segment of the polyline is smooth or linear.

3D software packages, such as Maya, 3Ds Max or LightWave, 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 keypoint and two handles for smoothness, the B-spline polyline has just one point, defining 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.


Polyline Toolbars



Whenever a tool that contains one or more polylines is selected, the polyline is shown on all display views and the Polyline Toolbar will appear. The toolbar contains several buttons that make switching polyline modes and options easy to access.

To identify each of the buttons, place the mouse pointer over the button until a tooltip that describes the button's function appears. Clicking on a button will affect the currently active polyline (see Polyline Focus below) or the selected polyline points, depending on the button.

To change the size of the toolbar icons, add labels to the buttons, or make other adjustments to the toolbar's appearance, right-click on the toolbar and select from the options displayed in the toolbar's context menu.




Selecting a Polyline

Keyboard shortcuts and toolbar selections will only affect the currently selected polyline. It is possible to have several polylines in the display view at once, so it is important to be able to switch between polylines effectively.

In addition to selecting a polyline by clicking on one of its keyframes or segments, use the following methods to make a polyline active.

- > **TAB and SHIFT-TAB Cycling.** To cycle between available polylines, use the TAB and SHIFT-TAB keyboard shortcuts.
- > **Context Menu.** Right-click in the view and choose the desired polyline by name from the Controls > Select menu.
- > **View Toolbar.** To the right of the display view toolbar, click on the small arrow icon. This will display a list of all of the controls that are currently displayed in the view. Choosing a polyline from this menu will select the polyline.

Polyline Creation Modes



Several different modes are available for creating and modifying polylines. The specific mode used when a polyline is first added will depend on whether the polyline is part of a path or a mask. To change modes, select from the Polyline toolbar and use the keyboard shortcut associated with that mode or select the mode from the spline's context menu.

Each of the modes is described in more detail below, together with the keyboard shortcut that is used to set the active polyline into that mode. Take time to explore each of the Polyline Creation modes offered by Fusion to become comfortable with their operations.

Click Append (SHIFT-C)

This mode is the default mode for mask creation. Simply click the mouse pointer in the position where the new point is to be added. Each new point is automatically connected to the last created point. Use Click Append mode to quickly define the rough shape of the mask, then switch to Insert And Modify mode to refine the mask further. It is also useful when the overall shape of a motion path is known, but the timing is not yet available.

To close a polyline created in this mode, place the mouse pointer over the first point created and click when the cursor changes shape. A closed polyline is automatically switched to Insert and Modify mode. Click Append mode is rarely used with paths.

Insert and Modify (SHIFT-I)

New points are added to the polyline between existing points on the polyline. Dragging the keys or handles modifies existing points on the polyline. Points can only be added along the existing polyline.

Insert and Modify mode is the default creation mode for motion paths. Any time a control that is animated with a motion path is moved, a new key is created and added to the end of the polyline, extending or refining the path. Masks, which are created in Click Append mode, automatically switch to Insert and Modify mode when the mask shape is closed. This makes it easier to add additional points and refine the shape of the mask.

Draw Append (SHIFT-D)

The Draw Append mode creates a polyline shape directly on the view, like drawing with a pencil or a paintbrush. Click-drag the left mouse button to draw the desired shape for the path, or use a pen and tablet. This mode is ideal for use in conjunction with a tablet, and for the creation of garbage mattes and effect masks around complex shapes.

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 (SHIFT-M)

Modify Only mode allows existing points on the polyline to be modified, but new points may not be added to the polyline. Be aware that it is still possible to delete points from a polyline in this mode.


Done (SHIFT-N)

The Done mode prohibits the creation of any new points, as well as further modification of any existing points on the polyline.

Closing and Opening Polylines


There are several ways to close a polyline, which will connect the last point to the first. When creating the polyline, place the mouse pointer over the first point created and the pointer will change to indicate that clicking on the points will close the polyline.

Closed polyline masks will normally generate a solid mask, although the Mask controls can also be used to create outline shapes. Used more often for masks, closed motion paths are also used to create seamlessly looping motions.



You can also close a polyline using the keyboard shortcut Shift-O, the Close button on the polyline toolbar, or by selecting the Closed option from the polyline's context menu. All of these options are toggles that can also be used to open a closed polygon.

Selecting and Adjusting the Polyline



Polyline Points Selection

To select one or more points on a polyline, click directly on the point or box-select the points by click-dragging a rectangle around them.

To add or remove points from the current selection, hold the Shift or Ctrl keys down while making the selection.

Use CTRL-A to select all of the points on the active polyline.

Moving Polyline Points

The selected polyline points can be moved using either the keyboard or the mouse.

Using the Mouse

Once the points are selected, simply click and drag on the selected points with the mouse. Holding down SHIFT or ALT before click-dragging will restrict movement to a single axis.

Using the Cursor Keys

The cursor (arrow) keys on the keyboard can also be used to move selected points. For smaller movements hold down CTRL while using the cursor keys, and for larger movement hold down SHIFT.


Edit Points

The E keyboard shortcut will display an Edit Points dialog that can be used to manually enter positions for one or more selected points. This dialog is described in detail further below.

Smoothing a Polyline Segment

Smoothing a spline segment reveals the Bezier control handles, making it possible to apply shaping to a polyline curve. There are three ways to smooth the selected points on an active polyline.

Use the keyboard shortcut Shift-S while the segment is selected. Use the Smooth button on the Polyline toolbar, or select Smooth from the polyline's context menu.




Linearizing a Polyline Segment

To make certain that a polyline segment is perfectly straight, that segment needs to be linearized. A Linear Segment has no Bezier handles. The segment is always drawn in a straight line between two points on the polyline.

There are three ways to make the selected points on an active polyline linear: use the keyboard shortcut Shift-L while the segment is selected; use the Linear button on the polyline's toolbar, or select Linear from the polyline's context menu.

Deleting Selected Points

Delete a selected point or group of points by pressing the Delete key, choosing Delete from the context menu or by clicking the Del button in the toolbar. The shape of the polyline will change to reflect the removal of these points.



Deleting all of the points in a polyline does not necessarily delete the polyline itself. The polyline controls still exist and are simply waiting for points to be created. To delete a polyline, delete the tool or modifier that created the polyline.

Bezier Handles

Each control point has two Bezier handles that describe tension through the point, helping to determine the shape of the curve. These Bezier handles only appear when the point is selected.

Click-drag on a handle to make adjustments to the curve of the segment that emerges from the point. The handle on the opposing side of the point will also move to maintain the relationship between these two handles.

To break the relationship between handles and adjust one handle independently, hold down CTRL while dragging the point. Subsequent changes will maintain the relationship, unless the CTRL key is held during each adjustment.

To constrain the motion of the handle to a single axis, hold ALT down while dragging the handle.

Hold SHIFT down while moving a handle point to adjust the length of both handles without changing the angles.

Edit Points

To display the Edit Points dialog, press E on the keyboard or select Edit Points from the polyline's context menu. The Edit Points dialog provides numeric controls for precise control over the position of polyline points in Fusion. When a single point in the polyline is selected, it will bring up a dialog box that contains the X- and Y-axis values for that point. This dialog box can then be used to edit the location of that point manually by entering the desired value.

It is possible to perform minor mathematical equations in this dialog box. For example, typing 1.0-5 will move the point to 0.5 along the given axis.

When multiple points are selected, all points will be moved to the same position. This is useful for aligning points to the same plane on a specific axis. Changes made to the dialog will affect all selected points. This is useful for lining up several points along a given axis.

In addition to absolute values for the X- and Y-axis of each point, it is possible to adjust points using an offset from their current position. To change the mode of the Edit Points dialog, click once on the label for the axis. The label will change from X to X-Offset or from Y to Y-Offset.

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 point to be affected is drawn on the display view surrounded by a larger box. Change which point in the selection is affected by using the next and previous buttons that appear at the bottom of the dialog.

Reduce Points

Selecting Reduce Points from the polyline's context menu or toolbar will bring up a dialog box to decrease the number of points used to describe the polyline. Fusion will attempt to maintain the overall shape of the spline while eliminating redundant points from the path. To reduce the amount of points in the path, move the slider to the left.

When the value is 100, no points will be removed from the spline. Use smaller values to eliminate more points. This mode is extremely useful when used in conjunction with the Draw Points Spline Creation mode, or to simplify an editable paint stroke.

Transforming Individual Points

Select the points to be transformed and:

- > Hold down the T key and click-drag to twist.
- > Hold down the S key and click-drag to scale.
- > Hold down the X key and click-drag to scale horizontally only.
- > Hold down the Y key and click-drag to scale vertically only.
- > Hold down the O key and click-drag to offset the points perpendicularly from the polyline.

The position of the mouse pointer when the transformation begins becomes the center used for the transformation.

Shape Box

The Shape Box permits free scaling and transformations of the polyline shapes. Toggle the shape box on or off using the Shape Box toolbar button, the Shape Box option in the context menu, or the CTRL-B keyboard shortcut.

If there are selected points on the polyline when the Shape Box mode is enabled, the shape box is drawn around those points. Otherwise, draw the shape box in the position desired by click-dragging a rectangle.

A four-sided box will be drawn. Select a handle in the middle of a side to resize the polyline along a specific axis, or select a corner handle to shape the polyline along both axes.

Holding down CTRL while dragging a shape box point will apply adjustments from the center of the shape box, constraining the transformation to the existing proportions of the shape box. Holding down SHIFT while dragging a corner handle affects only that handle, allowing skewed and non-uniform transformations.

Other Polyline Options

Stop Rendering

Fusion will attempt to render the tool's result while points are moved on a polyline, providing interactive feedback of the effect the changes will have.

This can be distracting in some situations and can slow down performance on a complex effect. To restrict this behavior so that Fusion only renders when the points stop moving, toggle the Stop Rendering button in the toolbar or select this option from the polyline context menu.

Showing and Hiding Points and Handles

It is often difficult to identify individual points when they are placed closely together. Fusion offers the ability to display both points and their Bezier handles, just points, or just handles. Select which display mode is used with the Show Key Points and Show Handles toolbar buttons, or from the polyline's context menu.

Use these options to simplify the screen display when adjusting points placed closely together and to avoid accidentally modifying controls and handles that are adjacent to the intended target.

Importing and Exporting Polylines

Fusion provides methods for importing and exporting polyline shapes in to 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.

ASCII Format

To save a polyline to an external ASCII file, right-click on the view and select Export from the context menu. A submenu will appear that contains several options for exactly how the shape will be described in the text file.

Point

Selecting Point as an option will cause Fusion to divide the mask into equally spaced points while attempting to preserve the exact shape.

Bezier

Select Bezier to cause the output file to contain the precise coordinates of each point in the shape, as well as the position of its associated Bezier handles.

Key

Select Key as an option to cause only the position of the existing key points to be exported, ignoring the position of Bezier handles. It is not possible to select more than one of the above three options. Additionally, the coordinate system to be used by Fusion must be selected when exporting the points.

Relative

The Relative option will describe the X and Y coordinates as offsets that are relative to the center of the mask.

Absolute

The Absolute option will describe the coordinates of the points and/or handles as values between 0 and 1, using Fusion's native coordinate system. Once the desired options are selected, choose Export Shape from the context menu and provide a name and path for the output file. To save the timing of a motion path with the file, select Export Path. To import a saved ASCII file, create a polyline and select Import from the polyline's context menu.

A full description of the file format used is available in the Appendix of this manual.

Roto Assist



Roto Assist is an option in the polyline's context menu and toolbar. If it is On, on click-add or freehand add points, the point will be placed on the closest edge if it is within the defined distance. Snapped points are indicated by a cyan outline.

Roto Assist

- > **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.

All Roto Assist options can be saved to and read from the comp.

Native Format

To save a polyline shape in Fusion's native format, right-click on the header of the Mask tool or the path modifier and select Settings > Save As from the context menu. Provide a name and path for the saved file and select OK to write a file with Fusion's *.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, create a new polyline of the same type, then select Settings > Load from the mask's context menu.

You can also copy the tool to the clipboard, then paste it into any text editor.

Effects Mask

Under normal circumstances, 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 only.

Any number of Effect Masks of any time 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.

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.

Mask Types

Effect Masks

Effect Masks are applied post-effect; the tool's result is calculated over the entire image first and then the mask is used to filter the result. Complex mask shapes are created and animated using a combination of primitive shape masks or more complex polyline shapes.

Effect masks are by far the most common type of mask applied to tools in Fusion. Although just about every tool in Fusion supports masking, there are a few where effect masks make little to no sense, notably Savers, Time tools, and Resize, Scale and Crop tools.

Garbage Mattes

Garbage Mattes are a special type of effect mask that is 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. This is accomplished by setting the alpha channel within the region of the mask to either black or white.

Use garbage mattes to remove rigging from blue and green screen shots and to isolate a section from an image not suited to chroma or color keying.

The option of creating a garbage matte is only available from the context menu when one of the keying tools is selected. The various garbage mask types and shapes are identical to effect masks described in this chapter.

Any number of garbage mattes can be applied on a single tool, however, they must all perform the same function. For example, if a garbage matte is added to an ultra keyer to exclude a rigging element from the key by setting the alpha channel to black, all other garbage mattes added to that ultra keyer must also be set to make the alpha transparent.

To choose whether garbage mattes applied to a keying tool will set the alpha channel to opaque or transparent, locate the garbage matte's drop-down menu in the tool's controls and select **Make Transparent** or **Make Solid** from the menu.

To mix garbage mattes set to different modes, use a separate matte tool immediately after the keyer in the flow.

Pre-Masking (DVE, Glow and Highlight Tools)

An additional type of masking, known as Pre-Masking, is available to the DVE, Glow and Highlight tools. Unlike effect masks, a pre-mask is applied to the tool before the effect. 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.

The following mask tools are described in greater detail in the Tool Reference.

Common Mask Features and Controls

Most Mask tools share several controls in common. These are documented in the Tool Reference Mask category. See Common Mask Controls.

Bitmap Masks

Bitmap Masks use the image from another tool to create the mask channel. 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.

A Bitmap mask is not required to generate a single channel mask from a multi-channel image – it is possible to connect an image directly to the Mask input of any tool. Using the Bitmap mask tool provides more detailed control over the resulting mask. This makes it possible to soften the mask, invert it, or clip the ranges used.

Direct connection of an image to the mask input is useful when the image already contains the exact mask desired. When direct connection is used, the common controls for that tool will display a drop-down menu, which can be used to select the channel that will generate the mask.

A Bitmap Mask tool must be used if the mask is going to be combined with other mask tools. See Bitmap Mask in the Tool Reference.

B-Spline Masks

B-spline Masks are user-created shapes made with polylines that are drawn using the B-spline method of calculating smoothness. They behave identically to polyline shapes when linear, but when smoothed, the key points influence the shape through tension and weight. This generally produces smoother shapes while requiring less control points.

B-splines are often used to rotoscope organic shapes, like people and animals.

The smoothness of a B-spline is determined by the tension of the control points. To adjust the tension of a B-spline's control points, select the point, hold down the W key and click-drag the mouse pointer to the left and right to increase or decrease the tension of the curve through that point.

See B-spline Mask in the Tool Reference.

Ellipse, Rectangle and Triangle Masks

These are primitive shape masks. See Rectangle, Ellipse, and Triangle in the Tool Reference.

Mask Paint

Mask Paint allows a mask to be painted using Fusion's built-in vector paint tools.

See Mask Paint in the Tool Reference.

Polygon Mask

Polygon masks are user-created shapes made with polylines that are drawn using the Bezier method of calculating smoothness. This is the most common type of polyline and the basic workhorse of rotoscoping.

See Polygon in the Tool Reference.

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 pixel of a similar color is also included in the mask. This type of mask is ideal for secondary color correction.

See Wand in the Tool Reference.

Ranges Mask

Similar to Bitmap mask, the Range 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, Range allows spline-based selection of low, mid and high ranges, akin to Color Corrector.

See Ranges in the Tool Reference.

Polygon and B-Spline Animation

When a Polygon Mask is added to a tool in the flow, it is automatically animated using a Polygon Change spline.

Polygon Change Spline

To morph the shape of a mask over time, simply change the current frame to a new time and make the required changes to the mask. A new point will be added to the Polygon Change spline. This one keyframe controls the position of all of the points that comprise the mask for that frame. Fusion will automatically morph the shape and position of the mask between the two keyframes. Additional points can be added to the mask at any point to refine the shape as areas of the image become more detailed.

The Polygon Change spline controls the timing of the mask's animation. To adjust the overall timing of the mask animation, edit the Polygon Change spline using the Spline Editor.

Removing Animation from a Polyline Mask

If it is preferred that the Polygon mask remain static, open the mask's controls and right-click where it says Right Click Here For Shape Animation and select Remove Bezier Spline from the context menu. To animate the Polygon mask again, select Animate from the context menu of this label.

Adding and Removing Points from an Animated Mask

When adding points to an animated mask, the new point exists at all keyframes and Fusion will do its best to fit the new point into the shape at all keyframes. It is recommended to check the composite at each keyframe to ensure that the result is as desired. Deleting a point removes that point from all keyframes in the animated mask.

Published Points

Although the polychange spline offers rapid animation of the entire shape, it does not allow a specific point's animation to be controlled explicitly. A complex mask requires more precise timing for key points on the polyline. It is possible to isolate individual points on a polyline by publishing them.

Publishing a point on a polyline removes that point from the polychange animation spline. Once removed, this point can be connected to paths, modifiers, expressions, and trackers.

Publishing a Point

To publish a selected point or points, click on the Publish Points button in the polyline toolbar. Alternately, select Publish Points from the polyline's context menu. A new coordinate control will be added to the polyline mask controls for each published point, named Point 0000, Point 0001, and so on.

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 context menu.

Publish Point to Path

When a point ID is published, any animation already applied to that point is destroyed. To keep the animation, Fusion also offers the Publish Point to Path option. To publish the selected points and convert their existing animation to a path, right-click on the polyline and select Publish Points to Path from the polyline's context menu. Alternately, click the associated button in the polyline toolbar.

Follow Published Points

It is often necessary to have points in a polyline mask that are between two published points follow the motion of the published points and still maintain their relative offset and shape. For this reason, key points in a polyline mask can be set to Follow Published Points using the polyline's context menu.

When a key point of an effect mask is set to follow points, the point will be drawn as a diamond shape rather than a small box.

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 softness control available in all mask tools provides a mechanism for softening the entire mask equally. Softening part of the curve while keeping other portions of the curve sharp is often required when rotoscoping in a frame with a lot of motion blur.

This form of softness is called Non-Uniform Softness, which is accomplished in Fusion 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 curve. The inner curve is the original shape from the single polyline, whereas the outer curve is used to determine the curve's softness. The further the outer curve gets from the inner curve, the softer that segment of the shape becomes.

Converting a Single Polyline to a Double Polyline

To convert a mask into a double polyline, right-click in the view and select **Make Double Polyline** from the mask's context 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.

The control points on the outer curve are automatically parented to their matching points on the inner curve. This means that any changes made to the inner curve will also be made to the outer curve. The relationship is one way. Adjustments to the outer curve can be made without affecting the inner curve.

A dashed line drawn between the points indicates the relationship between the points on the inner and outer curves.

Adding Softness to a Segment

The shape of the outer curve is drawn using a dashed line instead of a solid line to help distinguish it from the inner curve. To select the outer curve, use the Tab key to cycle between controls until the dashed outline is visible, or select the outer polyline using the context menu's **Controls > Select** menu.

Once the outer polyline is selected, click-drag any of the points away from the inner polyline to add some softness to the mask.

The further the outer curve segment is from the inner curve, the softer the mask will be in that area.

Adding Additional Points to the Curves

It is not necessary for every point on the inner curve to have a match on the outer curve, or vice versa. Add additional control points to refine the shape of either curve. There is no need for a point on either curve to be parented to any other.

Each polyline has its own polychange spline and stores its animation separately. If a point is adjusted on the inner curve that is parented to a point on the outer curve, a keyframe will be set on both polychange splines. Adjusting a parented point on the outer curve only sets a keyframe for the outer curve's polychange spline. If a point that is not parented is adjusted, it will only set a keyframe on the relevant polychange spline.

Locking/Unlocking Point Pairs

Additional points that are added to the inner and outer curves may need to be parented together. To parent the selected points, right-click in the view and choose **Lock Point Pairs** from the context menu for either spline. This will cause the selected point on the outer curve to become parented to the selected point on the inner curve.

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 view and deselect **Lock Point Pairs** from the context menu.

Converting Spline Types

Just because a shape using a B-spline or Polyline mask has been created does not mean that the mask type has to be used. Convert from B-spline to Bezier, or Bezier to B-spline, as needed. Right-click in the display view and select Convert Bezier Spline To B-Spline or Convert B-Spline To Bezier from the spline's context 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 curve 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 post-conversion.

Motion Paths

Motion Paths are polylines that apply movement to the positional controls of layers, effects and masks over time. Points can be inserted into any part of the motion path to create any shape imaginable.

Motion paths can only be added 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.

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 two types of Motion Paths: Displacement Paths and XY Paths. A displacement path uses two splines to record the path, one for shape of the path and a displacement spline that represents 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 displacement path is the default type for motion paths, and most documentation in this chapter assumes that this type is used.

Path Modifier

A path is a Modifier tool. When the path is added to a control, controls for the path will be visible in the Modifiers tab of the tool control window.

Displacement Paths

Creating a Displacement Path

To add a motion path to a control, change the current time to the frame where the motion should begin and position the control in its starting position.

Right-click on the control in the display view and select Set Key Creating Path from the context menu for that control.

The control's text boxes will turn white to indicate that the parameter is now animated and to note that any changes made to the control will cause a new keyframe to be created.

The first point of the path will be created and displayed in the view.

To continue adding points, change time and move the control to the location desired for this frame. Fusion will automatically create a keyframe on the motion path at the current time.

A polyline will be drawn from the original keyframe to the new one. The position of the control at frames between the two keyframes is automatically calculated.

Continue adding points by changing time and moving the control until the entire motion path is described.

Upon completion, set the polyline to insert and modify mode by pressing Ctrl-I 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.

Displacement Splines and Path Timing

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

Displacement splines are used to control the speed of a control's movement along its path. To slow down, speed up, stop or even reverse the motion of the control along the path, adjust the values of the points for the path's displacement in the Spline Editor.

A displacement value of 0.0 in the Spline Editor indicates that the control is at the very beginning of a path. A value of 1.0 indicates that the control is positioned at the end of the path.

Each locked point on the motion path will have an associated point on the displacement spline. Unlocked points will not have a corresponding point on the displacement spline. Locked and unlocked points are described in more detail below.

Types of Points

Displacement paths in Fusion are composed of Locked and Unlocked Points. Whether a point is locked or not is determined by how it was added to the polyline. Each has a distinct behavior, as described below.

Locked Points

Locked Points are created by changing the current time and moving the animated control when the keypoint is created (as described earlier in Creating a Motion Path). 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.

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 views as smaller, solid square points.

To demonstrate the difference between locked and unlocked points, add a Text tool to the flow. Type a word in the Text tool (or anything else), then position the text's center in the upper left corner of the image.

Make certain that the current time is frame 0, right-click on the text center and select Animate from the context menu.

This will create the first locked point of the path. Advance to frame 45 and move the text center to the center of the screen. This sets the second locked point. Display the Spline Editor to look at the Path 1: Displacement Spline that was created.

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.

Select the keyframe at frame 45 in the displacement spline and move it to frame 50. The motion path is now 50 frames long, without making any changes to the motion path's shape.

Now 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.

Advance to frame 100 and move the text center to the bottom right corner of the screen. This will create an additional locked point and set a new ending for the path.

Currently, the center is the selected control. Select the motion path spline by using the tab keyboard shortcut to cycle controls until the path is selected, or choose the path using the control's context menu. At the moment, the path is in click append mode. To insert some points, switch the polyline by right-clicking on the path and selecting Insert And Modify from the context menu.

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. Adjust these points so that they are positioned like below.

Look at the Spline Editor

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 (because the current time was not changed before the points were added). Take advantage of this behavior by making the text stop briefly in the center.

Select the point at frame 50 in the displacement spline and hold down the Ctrl key while dragging it to frame 65. The point will be copied at frame 65.

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. Render a preview of frames 0 to 100 to see the results so far. Remove the preview when done.

Switch back to working in the view. Delete the unlocked point added between the first point and the middle point. The spline segment of the motion path should 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. Step through a few of the frames in the scene to confirm this.

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 allows the full power of Fusion's motion paths to be harnessed.

Locking and Unlocking Points

It is possible to change an unlocked point into a locked point, and vice versa, by selecting the point(s) and choosing the Lock Point option from the context 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, right-click on the control and select Modify With > XY Path from the context menu.

At first glance, XY paths work like displacement paths. To describe the path, change frames and position the control where it should be on that frame, then change frames again and move the control to its new position. Fusion automatically interpolates between the points. The difference is that no keyframes are created on the onscreen path.

Look in the Spline Editor to find the X and Y channel splines. Changes to the controls position will be keyframed on these splines. The advantage to the XY path is that it becomes very easy to work with motion along an individual axis.

Switching Default Paths

Change the Default Path type used when animating a coordinate control to XY path (if this is the preferred type of animation). Open the Global preferences and look under the default's tab for the drop-down list in the default animate section labeled Point With. Change this from the current value to XY path. The next time Animate is selected from a coordinate control's context menu, an XY path will be used instead of a displacement path.

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 manor.

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.

- > To copy a motion path, right-click on the path's control header and select Copy from the context menu.
- > To cut a motion path out of a tool, right-click on the path's control header and select Cut from the context menu.
- > To paste the copied path over another path, right-click on the path's control header and select Paste from the context 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 From Path X: Value.

Recording

Using this mode allows both of the control's spatial and temporal information to be animated at the same time. 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 context 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 suite the needs of a project precisely. Adjust the path's displacement spline in the Spline Editor to more correctly match the required motion.



Interface Overview

2

Interface Overview

This chapter should familiarize the reader with the names and locations of the main elements of the Fusion interface. A detailed description of each interface element can be found later in this documentation.


Toolbars

The top of the screen contains the standard application menus, which are always located at the top of the application window. In a default Fusion layout there will also be three standard toolbars. These toolbars can be detached and moved to other parts of the interface. The Tool, Macro and Script toolbars can be customized to the artist's preferences. Additional toolbars can also be created as needed.



See the Menus chapter and Toolbars chapter for more information.

Main Interface



Display Views

The Display Views are used to display 2D images and 3D scenes created within the composition. Display views may also be used to show custom interfaces for analyzing images, such as histograms, waveform monitors, color inspectors and metadata.

See the Display Views chapter for more information.

Work Panel

The work panel is a tabbed region below the display views that hosts 6 interfaces for interacting with Fusion. These are the:

- > Flow Editor
- > Timeline Editor
- > Spline Editor
- > Scripting Console
- > Comments


Control Panels

The control panel is a tabbed region on the right side of the window that is used to display the sliders, menus and edit boxes that make up a tool's controls. There are two tabs:

- > Tool Controls
- > Modifier Controls

Time Ruler

The Time Ruler is the name given to the controls found beneath the work area. The Time Ruler contains controls for initiating playback as well as setting composition options like current frame, quality, proxy level and motion blur.



See the Time Ruler chapter for more information.



Navigating the Interface

3

Navigating the Interface

This chapter explains how to 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. These will be described in more detail in the chapter dedicated to that panel. Examples of this include the Flow's Navigator window and the Spline Editor's axis labels.

Panning

To pan or reposition the display of any part of the Fusion interface hold the middle mouse button down and drag in the desired direction. The artist can also click-drag on the scroll bars that appear on the right and left sides of each interface element when off screen elements are present.

If the mouse does not have a middle mouse button, hold the Ctrl-Shift keys down while clicking the left mouse button.

Scaling

To smoothly adjust the scale of a panel place the mouse pointer over that interface panel and hold the left and middle mouse buttons down while dragging the pointer. Moving the pointer to the left will zoom out and moving to the right will zoom in.


To scale a panel in steps place the mouse pointer over the interface and hold down the middle mouse button. While holding down the middle mouse button, click on the left mouse button to increase the scale, or the right mouse button to decrease the scale.

The + and - keys on the numeric keypad can also be used to scale the interface panels.


Layout

The size of each interface panel can be set using preset configurations or they can be adjusted manually with the mouse. The display views and work panels are complementary the more space used to display the work panel, the less space available to show the display views.


To resize a panel manually left click and drag anywhere along the raised border surrounding the edges of the panel.



The Layout toolbar provides four preset panel layouts. In a default installation the toolbar is found beneath the menus at the top left of the window.



The top of the tool controls area can be aligned to the bottom of the display views or to the top of the screen. 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 F4 keyboard shortcut will collapse all interface panels except the one under the mouse pointer, allowing it to take over almost the entire screen. The F4 key is a toggle, selecting it again restores the previous layout.




Comments Tab

4

Comments Tab

The Comments Tab (F9) provides a rich text edit box that can be used to store comments and notes about the composition.



This can be useful for maintaining a todo list for personal purposes or for collaboration between several artists working on the same composition.

Right-clicking in the text box will display a context menu with options for cutting/copying/pasting text to/from the clipboard. Alternately, use the Windows standard keyboard shortcuts Ctrl-X, Ctrl-C and Ctrl-V.






Console Introduction

5

Console Introduction

The Console Tab contains a text area used by Fusion to display messages produced while working with the composition. This includes error messages, basic information messages and output from running scripts.



Additionally, the console provides a scripting interpreter that can be used for immediate mode scripting. If the console is not the active view, any new messages sent to the console will cause a small icon to appear in the tab, indicating that a new message is available.

Console Toolbar

The Console Toolbar contains seven buttons. One button is used to clear the console, four are filters for the messages displayed in the console and two switch between Lua and Python code.



Note: The filters do not destroy messages sent to the console. They merely choose whether the messages are hidden or displayed.

Console Toolbar Buttons

Hold the mouse pointer over each button for a moment to display a popup dialog that describes the button's function.

	Cls	Click on this button to clear all text displayed in the console.
	Show Error Messages	Click on this button and the console will display error messages in the console.
	Show Log Messages	When this button is selected, the console will display general information messages in the console. For example, after a render is complete, the console will show information about when the render started, how long each frame took, and the average render time per frame for the composition.
	Show Scripting Messages	When this button is selected, the console will display the output of flow and tool scripts run by the artist.
	Show Input Messages	When this button is selected, the console will echo any text entered in the text box below the console.
	Lua Mode	The console will now interpret any code entered in the entry field as Lua.
	Python Mode	The console will now interpret any code entered in the entry field as Python.

Console Entry Field



The text box immediately beneath the console can be used to enter scripting commands directly.


Any Fusion script commands that are entered into this dialog will be executed the moment enter is selected. For details on using the Console Entry Field and scripting, please refer to the online scripting manual that was installed with Fusion.



Time Ruler

6

Time Ruler



Current Time

The Current Time field (a.k.a. Current Frame) indicates which frame of the project is currently displayed. When tools are displayed in the views it is the current time that is rendered and displayed. Adjustments made to animated controls add keyframes at the current time. The text field at the top-right side of the ruler indicates the current time, and can be used to enter a new time. The Ruler also shows a marker to indicate where the current time is within the currently visible frame range. Left-clicking in the ruler will set the current time to the time clicked on.

The following keyboard shortcuts are useful for manipulating current time:

[and]

Advances the current frame forward or backward in time by 1 frame.

Shift - [and]

Sets current time to the project's Global Start or Global End Frame.

Alt - [and]

If a tool is currently active then the Alt-[and] keys can be used to set the current time to the next keyframe for any animation curves attached to that tool's controls. 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, vertical ticks on the Time Ruler will indicate the associated keyframe positions.

Changing the Display Format

The current time is typically shown as just the current frame number, however this can optionally be configured to show current time in either SMPTE timecode or the film format Feet+Frames instead. To change the current format open the View menu at the top of the interface window and enable the Show Timecode option. Use the Time Code Format menu to select which display format is used.

The options in the compositions frame format preferences panel determine the framerate used for Timecode, or the number of frames found in a foot of film.



Note Many Fusion edit boxes will evaluate mathematical expressions simple maths - for example, typing $2.0 + 4.0$ in most edit boxes will result in the value 6.0. Unfortunately, SMPTE timecode uses the + symbol as a separator symbol, so this edit box will not evaluate mathematical expressions which use the + symbol correctly, even when the display format is set to Frames mode.

Ranges

Global Start and End Range

The Global Start and End controls tell Fusion what the composition's valid frames are in a project. Video editors should think of global range in Fusion as the equivalent to the length of the Timeline in an editing program. The current time cannot be changed to a frame outside the Global Range.

Shift-dragging a tool tile from the Flow Editor to the Time Ruler will automatically set the Global and Render Range to the valid extent of the tool.

Render Start and End Range

The Render Start and End Range controls set the range of frames that will be used for interactive playback, disk caches and previews or final renders. The range is normally visible in the time slider as a grey highlight over the global time. Frames outside the Render Range will not be rendered or played in the flow, although the current time can be changed to these frames to see what the image looks like.

Dragging a tool tile from the Flow Editor to the Time Ruler will automatically set the Render Start and Render End Range to the valid extent of the tool.

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 most artists will leave this at the default value (showing the entire global range), there are times when 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.

Render Button


Clicking the Render Button will cause Fusion to display the composition's Render Settings dialog, where the artist can confirm the render options and initiate rendering of the composition to disk (assuming the composition contains Saver tools). SHIFT clicking on the button will skip the dialog, using default values for the rendering (full resolution, high quality, motion blur enabled).

Playback Buttons



The playback buttons in the Time Ruler can be used to initiate playback of the composition forward or 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 that Fusion will render the composition to display currently viewed tools at the current frame, then advance current time to the next frame (or previous frame) and repeat. There is no assurance of real time playback at the project's full frame rate, unless the frames are already cached in memory.

Right-clicking on either the frame advance or previous frame button will display a context menu that can be used to set a frame increment. This increment will be respected by playback as well as the frame advance buttons and keyboard shortcuts [and]. This is often useful for rotoscoping on fixed increments, or inspecting a frame one field at a time.



The loop button is a toggle that determines whether playback will stop when the current time reaches the render start or end, or whether playback will loop.

Audio Mute

The Audio button is a toggle that can be used to determine whether the composition attempts to play back any audio associated with the clip. Additionally, right-clicking on this button will display a context menu that can be used to select a .wav file that can be played along with the composition, and to assign an offset to the audio playback.

Composition Quality Options



High Quality

When working with Fusion 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, Fusion does not normally render images interactively at the same quality that is required for final renders. More complex and time-consuming operations such as area sampling, anti-aliasing and interpolation are skipped to produce results more rapidly. Enabling High Quality (HiQ) mode will force Fusion to produce a full quality render instead.

To enable HiQ mode, click on the HiQ button at the bottom of the screen. Interactive processing will take longer, but the image seen onscreen will be identical to that produced during a final render.

Motion Blur

When the Motion Blur button in the toolbar is toggled off, Fusion will ignore the motion blur options enabled on any of the tools in the composition. 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 interactively. Enable proxy mode to force Fusion to process one out of every x pixels interactively, instead of rendering every pixel. The value of x is decided by right-clicking on the Proxy button and selecting a proxy ratio from the button's context menu.

If a lower resolution version of footage 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 footage.

Auto Proxy

The Auto Proxy button enables Fusion to interactively degrade the image while adjustments are made to controls, snapping back to normal resolution when the control is released.

For example, while dragging a center control, Fusion may only process the image at a third 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.

Set the auto proxy ratio by right-clicking on the Auto Proxy button and selecting it from the context menu.

Selective Updates (Some, None and All)

Interactively, 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.

Some

During final renders, Fusion only renders the tools that directly contribute to the output saved to disk. This is known as the Some mode of selective update because only some of the tools in the flow are rendered. Some is the default mode.

None

None 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 None mode is enabled the display views will be highlighted with a red border to indicate that the display views are not being updated to reflect changes taking place in the composition.

All

All forces all of the tools in the flow to render. This can be useful when interactively adjusting small, fast compositions.



Bins


7

Bins

Overview

Fusion's bins provide an easy means of accessing commonly-used tools, settings, macros, compositions and footage. Use the bins to keep all media resources close at hand or to improve the collaborative workflow.

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




Bins Layout

The bin's detailed window can be set to one of several layouts. The default layout is a straight icon view, using large icons.


The contents of bins can be viewed as a table of icons, a simple list or as a detailed list. The icons can be adjusted to small, medium, large or huge.

To change the layout for a bin, right-click in the bins and select the desired options from the Icon Size and View submenus.



Bins Hierarchy

Bins are organized in a hierarchy tree of folders similar to the Windows file system.




Projects and materials can be organized in whatever form best suits your workflow and temperament.

Best of all, bins can be shared with others running Fusion on the network. Run one or more standalone bins servers to be shared by everyone in a facility.

Connecting to Bins on Another System

To connect to a remote system and display its bins, open the preferences and locate the Global > Bins > Servers panel in the list.




This panel shows a list of the systems, with buttons below to allow new entries to be added to or deleted from the list.

Adding an Entry to the Bin Server List

Click the Add button to add an entry to the list. The text controls below the button will become enabled for editing. Enter the system name or IP address of the system hosting the bins to be connected.

Fusion does not attempt to test or connect to the system until the Bins dialog is opened. If it cannot connect, the system will still be listed, with access denied or unavailable marked next to the name of the system in the tree.



There is no practical limit to the number of bins that can be accessed.

Accessing Remote Bins

Remote Bins show in the Bins list as another top-level item in the tree on the left side. The system is shown by name. If the system is unavailable, or requires a password, the label will also read unavailable or access denied. Open and navigate remote bins just like local bins.

Permissions

Remote users access to the bins can be restricted, requiring a login to view the contents, make changes or create folders.

Managing the Bins

Items That Can Be Stored in the Bins


Bins support any footage recognized by Fusion, as well as compositions, tools, saved tool settings, macros and tool groups.

Opening and Closing the Bins

The bins are a separate window that float over the main Fusion interface. To open and close the bins window, select File>Bins from the menu, click on the Bins button in the main toolbar, or use the keyboard shortcut Ctrl-B.

Interface Overview

The bins window is separated into two panels. The bins tree on the left side shows the relationships between the bins, and details of the selected bin's content are shown the right.




Bins Context Menu

Most operations in the bins are accessed via the Bins context menu. Open the context menu by right-clicking on an empty area of the bin. Right-clicking on media or any other item in the bins will show the same menu with additional options for manipulating that item. There is no context menu available from the tree panel.

Working with Bins

Adding Items to the Bins

There are three kinds of items that can be placed in a bins folder: footage, tools and tool settings, and saved flows. To add items in to the flows, right-click inside the bin and select Add Item from the context menu.



A file browser will appear. Select the footage, files or flows to be added to the bins.

Compositions and Footage

Footage and compositions can be added to the bins by drag-dropping them into a folder from the Windows file browser, or by right-clicking in the required folder and selecting Add Item from the context menu.

Tools

To add a tool with default settings to the Bins, drag and drop the tool's tile from the flow into the bins folder.



To add some Fusion tools and all third-party plug-ins, drag and drop or load the .plugin file for that tool from the Plugins directory. For example, to add the particle tools to a bin, locate the file particles.plugin in the Fusion > Plugins directory, then drag and drop the file into a bin. Use the Create Tool Folders option in the context menu to quickly create folders containing all of the tools currently known to Fusion, organized by tool category.

Tool Settings

To save the settings of a tool or tools and add them into the bins, select the tool or tools in the flow and drag it into the desired bin. A dialog will appear to choose where on disk the new settings file will be saved. That setting will be stored in the bin. In addition, load the .setting file directly from disk using the Add Item option in the context menu.


Footage Stills vs. Sequences

Fusion automatically identifies file sequences on disk and loads them as sequences rather than stills, so it is not necessary to select more than one file from a file sequence. Fusion will automatically identify the sequence as a single clip when it loads it into the bins.

To ignore the sequence, just add the clip into the bins as a single file or still. Hold the Shift key down while adding the clips to the bins and each file will be seen as a still. This technique applies to dragging and dropping as well as using the file browser.

Creating a New Bin Folder

To create a new Bin Folder, right-click in the Bins and select New Folder from the context menu.



Renaming a Bin Folder


To rename a bin folder, right-click on the folder icon from its parent bin and select Rename from the context menu or press the F2 keyboard shortcut.

Deleting Items from a Bin

To remove an item from the bin, select the item and press Delete, or right-click on the item and select Delete from the context menu.


Deleting a Bin Folder

To remove a bin folder and all its contents, right-click on the folder and select Delete from the context menu, or press the Delete key while the bin folder icon is selected. You cannot undo removing a bin folder.



Adding an Item to the Current Composition

To add the items in a bin to the composition, either drag the item from the bins to the composition or double-click on the item. Each type of bin item behaves a little differently when added to a flow.



Footage

Dragging a footage item into the flow editor creates a new loader that points to the footage on disk. Still files or single images will be set to loop automatically. Footage that is added to the bins is referred to as a clip.

Compositions

Double-click a composition to open it. Dragging a comp item onto an open composition will have no effect.

Tools and Macros

Dragging a tool from the bins to the flow will add that tool into the flow with the default settings. Dragging a tool from the bins into a display view will insert that tool after the currently viewed tool.

Tool Settings

Dragging tool settings from the bin to the flow will add the tool(s) into the flow with the saved settings applied. Dragging from the bin into a display view will insert the tool after the currently viewed tool.


Working with Clip Items

Stamp Files

Stamp Files are low resolution, local proxies of footage (referred to as a clip) that Fusion uses to quickly play back the individual frames in an image sequence or movie.

Although there have been vast improvements in the interactive playback from disk, stamps remain extremely useful for quickly previewing clips stored on a network server, or for very large Cineon and DPX files.

To create a stamp file for a clip, right-click on the item in the bins and select Create Stamp from the context menu.



The status bar at the top of the Bins window shows Fusion's progress as it creates the stamp in the background. Other stamps can be queued as well.

It is possible to continue working with the composition while the stamps are created.

Stamp Playback

To play back a clip, press the right mouse button and select Play from the context menu.



Fusion plays the clip back in its image icon.


Scrubbing Footage

In addition to being able to play back the stamp, Fusion also offers the ability to scrub through a clip's frames in the bins. To scrub, right-click on the footage icon in the bin and drag the mouse left and right. The footage will appear in the left display view of the currently open composition. Dragging to the left will move back through the frames and dragging right scrubs forward.

If a composition is not currently open, the frames will not be able to display in the left view and nothing will happen.

Jog and Shuttle

Fusion can scrub in one of two modes. Jog works the same as the jog shuttle on VTRs, where the clip icon is the pause position. Right-click and drag to the right for forward jog and to the left for reverse jog.




Shuttle works the same way as a dial shuttle. Move the mouse to the right to scrub forward or move to the left to scrub back frame by frame.

To switch between Jog and Shuttle mode, select the desired option from the Scrub Mode submenu in the context menu.

Backing Up the Bins

The configuration and items saved in the bins are all stored as text files in a folder called Libraries, which is automatically created beneath the directory containing eyeonServer.exe.



To back up the bins, or to copy them to another installation of Fusion, copy the files in this folder. Usually there is only one file here, called Library.library



Display Views




Display Views

Without Display Views, it would be difficult to determine exactly what is happening to the images as the composition is assembled. The Fusion display views show not only 2D images, they also show any 3D environment that might exist in the composition.

Display views also offer the power to analyze images using waveforms, histograms and vectorscopes.


By default, Fusion opens with two display views visible side by side across the top of the program interface. Additional floating display views can be created, or an external video monitor can be used if supported hardware is present on the workstation, such as one of the boards manufactured by Leitch or other manufacturers.



The display views are OpenGL-accelerated and rely upon the presence of a good OpenGL-accelerated display device.


Types of Display Views

Left/Right Display Views




The primary display regions in Fusion are the Left and Right Display Views. Located immediately above the Flow, Timeline and Spline regions, these display views are always available for use.

Floating Display Views



In addition to the large display views above the flow, an unlimited number of Floating Display Views can be created. These floating views are excellent for taking full advantage of a dual monitor configuration. Floating views can also be set to full screen mode to make the best use of screen real estate.

To create a new floating display view, select Window > New View from the menu bar at the top of the screen. The position and configuration of the new view can be saved in the layout preferences, if required.




Hardware Frame Buffers

When supported, or if video hardware is present in the computer, Fusion can take advantage of Hardware Frame Buffers to display images on an external NTSC, PAL or HD monitor.

The hardware display cannot be used to manipulate onscreen controls like center crosshairs or masks. They are extremely valuable for evaluating the images in the same format to which they will be output, and for determining color fidelity and accuracy with a properly-calibrated output monitor.

See the appendices at the end of this manual for a complete list of supported video I/O hardware.

Display View Context Menu



The display view provides a wide range of configuration options via a dynamic context menu that can be accessed by right-clicking in the view. Many of these options also have keyboard shortcuts or can be found in the View toolbar along the bottom of the view.

All of the options found in the Display View Context Menu are explained in depth throughout this chapter.


Displaying a Tool's Output in the Views



There are several methods to display the output of a tool in the display views. When a tool is displayed in one of the views, a status light appears on the left side of the tool's tile in the flow. To determine in what view the tool is displayed, hold the mouse pointer directly over the status light until a tooltip appears.

Generally, the leftmost indicator is for the left display view. The next one to the right is for the right display view. Additional indicators are used for floating views and video output devices.

The name of the currently-viewed tool is displayed on the horizontal toolbar at the bottom of the display view.



Viewing Images with Drag and Drop

From the Flow View

Click-drag any tool in the flow into a display view, then release the mouse button to display that tool's output.

From the Tool Controls

Click-drag the tool's control header from the control view on the left into a display view to display that tool's output.

From the File Browser


Drag and drop an image from the operating system's file browser into a display view to view that image. This will not add a tool into the flow. It is used only to preview an image rapidly before actually adding it to the flow.

Holding down the Alt key while dragging an image tool into a display view will create a flipbook preview in that view instead. 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 view. This is similar to selecting Create Preview On from the tool's context menu. Hold down Shift+Alt instead to skip the dialog and immediately start rendering the preview with the last-used settings.

Viewing Images with the Context Menu


From the Flow View

Right-click on a tool's tile in the flow to display its context menu. Select View On, then choose the display view to use.




From the Tool Controls

Right-click on the tool's control header on the left to display the context menu for that tool. Select View On, then choose the display view to use.



From the Timeline

Right-click on the tool's name in the tree menu to the left of the Timeline to display the context menu for that tool. Select View On, then choose the display view to use.



Viewing Images with Keyboard Shortcuts

Each display view in Fusion is assigned a numeric shortcut between 1-9, which can be used to quickly display the result of the currently selected tool. Pressing '1' accesses the left display view. The right display view is accessed with '2,' and successive floating views or hardware frame buffers are assigned numbers in the order in which they are created.

To determine the shortcut assigned to a display view, bring up the View On menu in any tool's context menu. The shortcut assigned is displayed to the left of the tool's name in the menu. These shortcuts are available for use when the Flow Editor or the tool controls are active.


Note: As of Fusion 5.1 and later, the numeric keypad may also be used for display view shortcuts, in addition to the numeric keys arranged horizontally above the keyboard.

Removing Images from the View

To remove all images from all display views in one action, press the tilde (~) key. This key is usually found to the left of the 1 key on American keyboards.


Viewing Images from Disk

Use the display views to show stills, videos or sequences, even if they have not been loaded into the composition using a loader. From the desired view, click the right mouse button and select Load Image from the context menu. This will open a file browser. Navigate to the folder with the footage or image to be viewed and click OK. Still images will be displayed on the view statically. For image sequences and multi-frame formats, the transport controls will appear in the bottom of the control area to play, review, and loop the sequence.



Viewing Images Using Other Methods

Each tool tile has a small strip along the bottom that shows a small white indicator for each view. The strip is generally not visible until the tool is displayed in a view, selected, or the mouse pointer is positioned over the tile.



By hovering the mouse pointer over an indicator, a small popup will identify the view associated with the indicator. Generally, the leftmost indicator is for the left display view. The next one to the right is for the right display view. Additional indicators are used for floating views and video output devices.

When the indicators are visible, click directly on them to display or remove the tool's output image from a display. A white indicator shows that the view is displaying the image, otherwise, the indicator is black.

Position and Layout

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

Once the desired layout has been achieved, use the Grab Document Layout button in the Layout Preferences to remember the layout for newly-created documents. The Grab Program Layout button will remember the size and position of any floating views, and enabling Create Floating Views checkbox will automatically create them again when Fusion starts up.

The current view layout is always saved with the comp. Turn off the Recall Layout switch in the Layout preferences to ignore the saved layout when opening a comp.

Layout Toolbar

The Layout Toolbar at the top of the interface (shown below) can be used to configure the display views to several preset layouts.



First Button

This option maximizes the left display view, aligning the tool controls with the top of the screen.

Second Button

This splits the available display region equally between the left and right display views, aligning the tool controls with the top of the screen.

Third Button


This fully maximizes the left display view, aligning the tool controls with the bottom of the display area.

Fourth Button

This fully maximizes the display region, splitting the left and right views equally. The tool controls stay aligned to the bottom of the display region.

Splitter Bars

The relative sizes of the left and right display views can be adjusted using the horizontal splitter bar between them.



Click-drag the splitter to increase or decrease the amount of space used by one display view. The adjacent display view will adjust to suit the new layout.

The amount of vertical space available for both display views can be adjusted by click-dragging on the horizontal scroll bar between the display region and the work region.


Fullscreen Floating Views

Floating views are not affected by the buttons in the Layout toolbar. Adjust the size of a floating view using the usual method of click-dragging on the outside of the window border, or by double-clicking on the window's titlebar. Additionally, right-click in the display and set the floating view to fullscreen from the context menu. The fullscreen option only appears once an image is present in the view.

Fullscreen Fixed Views

To maximize any fixed display view so it fills the screen, place the mouse pointer over the view and press the F4 keyboard shortcut. To restore a maximized view to its original settings press F4 again. This shortcut can also be used to maximize the workspace below the fixed views.

Panning and Scaling the Image



Both Panning and Scaling work the same way as in all other parts of Fusion. Also see the Scaling the Flow chapter in this manual.

Panning

To pan in any display view, place the mouse pointer in the view and hold the middle mouse button down. Drag the mouse in the desired direction. If the mouse does not possess a middle mouse button, hold Ctrl+Shift and click-drag the view (with the left mouse button) for the same effect.

When viewing a 3D scene, panning will change 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.

Note that 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 Ctrl+1.

Scaling

The following methods can be used to scale an image or 3D scene once it has been displayed in a view.

Note that 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

Mouse Drag

Place the mouse pointer over the view, then hold the middle and left mouse buttons down while dragging the mouse to the left and right to scale smoothly in and out. The scaling is centered on the starting position of the pointer.

Mouse Click

Place the mouse pointer over the view, then hold down the middle mouse button. While holding down the middle mouse button, click once on the left mouse button to increase the scale or the right button to decrease the scale by a fixed amount. The scaling is centered on the position of the pointer.

Mouse Wheel


Place the mouse pointer over the active view, hold down the Ctrl key, 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 pointer.

Using the Keyboard

To change the scale by a fixed value, press the + or - keys on the numeric keypad while the display view is active. Ctrl+1 will set the image or scene back to its default of 100% scale, and will recenter it as well. Ctrl+F will fit the image or scene into the visible region.


Using the Context Menu

Press the right mouse button while the pointer is positioned over the display view and select Scale from the display view’s context menu.



Using the View Toolbar

The scale button in the Display View Toolbar can also be used to adjust the scale of the view. Click anywhere on the Scale button to bring up a menu of several preset scales to choose from, or click on the Fit button beside it to fit the image or scene to the view window.



Rotating

When viewing a 3D scene or 3D histogram, the display view can also be rotated. This can be done with the mouse.

Normally, view rotation “orbits” the point of view, to allow easy viewing from any angle. When viewing through a camera in the scene however, this changes to rotate (pitch and yaw) the camera instead.

For more information about working in 3D you might want to have a look at the 3D Chapter of this manual.

Using the Mouse

Mouse Drag


Place the mouse pointer over the view, then hold the middle and right mouse buttons down (alternately, Alt key and middle button) while dragging the mouse to the left and right to rotate the view. The rotation is centered on the middle of the view.

Previews

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

RAM Flipbook Previews

RAM Flipbook Previews are preview renders that exist entirely within RAM. They are created by holding down the Alt key while dragging a tool into the display views, or by selecting Create > Play Preview from the tool’s context menu. A Preview Render dialog will pop up.




You can change the settings or leave them at their default.

If you Shift-Alt-Drag the tool into the view either the default preview settings or the settings you have chosen on the last preview render will automatically be applied.

The settings dialog will not show up and rendering the preview will start right away.

Settings



HiQ

Will render in full image quality like a final render would do.


MB

Will render with motion blur applied if any tool is set to produce motion blur.

Some


Will render only those tools specifically needed to produce the image of the tool you're previewing.

Size




Depending on the amount of available RAM, not all frames of the desired render range will fit. Fusion dynamically calculates the currently available memory and displays how many frames in any of the given resolutions will fit into RAM. Choose the setting that delivers the best Quality/Length ratio for your preview.

Network




Available Render Slaves or copies of Fusion running on the network, can be used to generate the preview. See the Network Rendering chapter for more information.

Shoot On




Sometimes its desirable to not render every single frame but every second, third or fourth to save render time and get faster feedback. Use this setting to adapt to your needs.

Frame Range




This field will by default take your current Render In/Out Range, but can be modified if needed to set the render range.

Configurations




Once you're done setting your personal preview configuration, this can be saved permanently for later use.



Click on Add, then type in a new name for your personal settings.


Playing Back the Preview




Once the preview is done rendering, the flipbook will be loaded into the view, temporarily overriding any tool currently displayed in the view.

Playback controls appear in the bottom right corner of the interface beneath the tool controls. A frame counter appears in the corner of the view, and a play/pause button is added to the right of the view's bottom toolbar.

To remove the preview and destroy the flipbook, right-click on the view and select Remove Preview from the context menu.



Other options on the view's context menu enable looped or ping-pong playback, or Fullscreen mode, which hides the rest of Fusion's interface. View controls are not displayed in Fullscreen mode. To exit Fullscreen mode, deselect it from the context menu, or press the Esc key.



It is also possible to “trim” the flipbook 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. Ctrl+drag the scrub bar to set or change the frames to be played, or right-click on it to set or reset the In and Out points via a menu.

Mouse Shortcuts

While the flipbook preview is showing, it can be played or scrubbed using only the mouse. These shortcuts work in Fullscreen mode as well.

Playback

Double-click the left mouse button anywhere in the view to start playback.

Scrubbing

Hold down the right mouse button and drag the mouse left or right to scrub through frames in the flipbook.

Keyboard Shortcuts

While the flipbook preview is active, playback and frame advance can also be done with the keyboard. These shortcuts work in Fullscreen mode as well.

Playback

The spacebar can be used to start or stop playback. Hold down Shift while pressing the spacebar to play in reverse.

Forward/Reverse

The left and right arrow keys will move the flipbook to the previous or next frame. Hold down Shift as well to jump back or forward 10 frames at a time. Ctrl+left will return to the first frame of the flipbook, while Ctrl+right jumps to the last frame.

File Sequence Previews

When a loader or saver is selected, the file sequence or footage loaded by that tool can be played back in the display view by pressing play on the transport controls that appear in the bottom of the tool controls region.


When the play button is selected, the file preview becomes active and Fusion will load the images as quickly as possible and play them to the display view. By default, the left display view is used for this purpose. That default can be changed using the Preview preferences. By changing these preferences, it is also possible to simultaneously play the file sequence to installed video hardware (this may decrease playback performance).

As with flipbooks, trimming In and Out points can also be set for file sequence previews, using the scrub bar in the transport controls that appear at the bottom right of Fusion's window. Ctrl+drag the scrub bar to set or change the frames to be played, or right-click on it to set or reset the In and Out points via a menu. By default, trimming done here 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 context menu.

While the file sequence preview is enabled, normal viewing of tools to that display view is disabled. To return to normal operations, select another tool or blank area of the flow.

Onscreen Controls


The display views show Onscreen Controls, or widgets, for manipulating the controls of the currently selected tool. Common widgets include crosshairs, angle indicators, polylines and paint strokes. Each of these controls can be manipulated directly in the display view using the mouse or keyboard.



The display view shows the controls associated with the tools that are selected in the Flow, Timeline or Spline Editor. If multiple tools are selected, the controls for all of those tools are shown.

The controls shown are determined only by the tools that are selected, not by the tool displayed in the views. For example, a downstream blur is easily viewed while manipulating the controls for a selected polygon mask or merge.


Enabling/Disabling Onscreen Controls



The onscreen controls for a view can be disabled so that they do not interfere with viewing the image.

To disable the onscreen controls, click once on the Show Controls button in the View toolbar or select Show Controls in the Options submenu of the display view context menu. Alternately, use the Ctrl-K keyboard shortcut.

Note that some tools like Masks allow disabling of their onscreen controls on a per-tool basis to reduce visual cluttering.



Making Fine Adjustments to Controls

Adjusting a control with the mouse often lacks the precision required for fine work. In cases such as this, the keyboard can often be used to make small finite adjustments to the selected control.

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

Holding the Ctrl key down while using the arrow keys reduces the scale of each step by a factor of ten. Holding Shift down increases the scale of each step by a factor of ten.

Toolbars

There are two toolbars in the display view, a horizontal view 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.


Display View Toolbar



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

This toolbar is dynamic. The buttons displayed may change depending on the type of viewer. A 3D display view shows a set of buttons for controlling the display of lights and shadows, for example, but these buttons are missing from a 2D view since they would have no meaning in that context. To learn more about a button's use, hold the mouse pointer over the button. A small popup will describe the button and its function.


Tool - Specific Toolbars



In addition to the toolbar displayed at the bottom of each view, Fusion also has a space reserved to the side of the view for Tool-Specific Toolbars. Text, masks, paths, paint strokes and the 3D environment are some of the tools that make use of this toolbar.

Customize the size and type of buttons displayed in this toolbar using its context menu, accessed by right-clicking on the toolbar with the mouse.

Customizing the View Toolbars



Right-click on the toolbar to display its context menu. This menu is used to change the size of the icons used for the buttons, and also to determine whether text labels are used to represent the button's functions.

A/B Buffers

Each display view has two buffers for viewing images. Each buffer can be considered a complete and separate view. Generally, all that is seen is the default A buffer.



Note: Buffers can be set to different display view types. It is entirely possible to have an A buffer showing the image with a B buffer showing a histogram.

Switching Between Buffers

Toolbar

Click on the View toolbar button labeled A or B to switch to the associated buffer.

Keyboard

Use the comma (,) key to set the active view to the A buffer, the period (.) key to set the view to the B buffer and the (/) key to enable the split wipe mode (described below).

Split Wipe Images

The AIB buffers can be used to enable a split display, where the separate buffers are displayed onscreen simultaneously. The displays are separated by a red splitter bar, which can be adjusted in any direction or angle by click-dragging on the splitter.



To enable the split wipe, click on the AIB button.



Turn off the split wipe mode by switching to either the A or B buffer directly using the toolbars or keyboard shortcuts.

While split wipe is enabled, new tools can still be dragged and dropped into either buffer to change the image displayed. Panning or zooming either buffer will keep the other buffer 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 be no longer visible in the view. Holding down Ctrl+Alt and clicking anywhere in the image will cause the splitter to jump to the current mouse position.

Holding down Shift while dragging the splitter bar will “snap” it to the nearest 45-degree angle.

Subviews



A Subview is a smaller display view that docks within the main view. It is connected to the same tool as the main view and is often used to show somewhat different information about the image.

For example, the RGB channels may be viewed in the main view, with the alpha channels in a subview. Any of the view types mentioned above can be used in a subview. In fact, several view types are designed for use only in the subview.

For the most part, the subview is nothing more than a miniature view, with its own context menu and options. It responds to the same keyboard shortcuts and navigation controls as any other display view, which means that most of the subviews can be panned or zoomed in.

Displaying/Hiding The Subview



Use the V key to toggle the subview on and off in the active view. Alternately, click on the subview button in the view toolbar, or select Views > Subview > Enabled from the context menu.

Changing the Subview Type

To change the subview type, click on the small arrow next to the Subview button in the display view toolbar. This will show a popup menu of all available subview types. Select one from the list or right-click in the subview and select the desired view type.


Swapping the Subview with the Main View

Use the Shift-V keyboard shortcut to swap the main view and the subview, or select Views>SubView>Swap from the context menu. This is not always valid, as certain view types can only be used in the subview. In such cases, the swap will not take place. Also see the View Type chapter for more information.

View and Subview Types


As well as showing images, views can be changed to show a variety of different bits of information about the image. Not all views are valid at all times. For examples, the 3D viewer is not available for a 2D tool and many of the viewers are only valid for subviews.

Changing the View Type




To change the View Type, open the display view's context menu and select the desired type from the view's submenu.

2D Viewer




The 2D Viewer is the basic 2D display view for showing images. When used as a subview, a different tool than the one used in the main view can be displayed by dragging and dropping the tool into the subview. This is the only subview type that is not slaved, showing the same tool that is displayed in the main view.

3D Image Viewer



The 3D Image Viewer is the basic 3D display view, available when viewing a tool from the 3D category.


Histogram



The Histogram is an analysis tool that can be used to identify problems with contrast and the dynamic range in an image. This is essentially a graph that shows the frequency distribution of colors in the image. The horizontal axis shows the colors from black to white. The vertical axis shows the number of pixels in the image that occur at each level. Histograms are excellent tools for examining the overall contrast and distribution of colors in an image. In Fusion 5.2 and later, the Histogram viewer shows out-of-range colors found in floating point images.

The Histogram viewer will also display Gradient information. Currently, the only tools in Fusion that can output gradients are the From Image and Perturb modifiers. Drag the modifier's titlebar into the view to see the gradient represented as a histogram.

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. Navigate the 3D cube like any other 3D view (Alt-MMB to rotate).

Vectorscope




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

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.

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 display view. This is useful when zooming in on an image in the main view.

Magnifier



The Magnifier can only be used in a subview. It shows a zoomed-in version of the pixels under the mouse pointer in the main view.

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.

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 mouse pointer. It will show all channels present, even the auxiliary channels such as Z buffer, XYZ normals and UV mapping channels.

The color is not sampled from a specific pixel, but rather it is sampled from several pixels. The default setting will sample from a 3 x 3 array of pixels around the mouse pointer. This can be changed by opening the View panel in the preferences and changing the value of the Color Picking Area Size.

Metadata


```
brightness = 0
camera_firmware_version = 20.1.3#20
camera_id = A
clip_aspect_ratio = 1.78
clip_aspect_ratio_denominator = 9
clip_aspect_ratio_numerator = 16
clip_id = 004
user_timecode_preference = 1
white_balance_kelvin = 5600
white_balance_tint = 0
TimeCode = 14:43:28:14
exposure_compensation = 0
```

Depending on the amount of metadata in your image this might present a wealth of information.

Choosing Color Channels

Fusion's display views and subviews are not limited to displaying only the full color version of the image. The separate Color and Depth Channels that comprise the image can be viewed separately.

Use the Channel toolbar button (shown below), the view's context menu, or keyboard shortcuts to change the channel displayed in a view.



Click the triangle on the side of the toolbar button to show the list of channels. Clicking elsewhere on the button will switch between Full RGB and Alpha, or the currently viewed channel.

The following hotkeys will also change 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 Fusion view has only support for the RGBA and Z channels. To view the auxiliary image data stored in RLA or OpenEXR files, it may first be necessary to copy the data from the auxiliary channel to one of the color channels using a Channel Boolean tool. Many auxiliary channels employ a 32-bit data type, so ensure the image is processing at 32-bit per channel before applying the Channel Boolean, or else the data may end up being incorrectly displayed.

The 3D View

The 3D View is only available when viewing a tool from Fusion's 3D category of tools or viewing some particle systems. It shows a representation of a composite in a true OpenGL-accelerated 3D environment.


This chapter is closely related to the Working in 3D chapter of this manual.

Rotating, Panning and Scaling the View

For the most part, navigation in the 3D environment is identical to the navigation of 2D images in the view. Panning and zooming the view are identical. To rotate in the view, hold the Alt key down and click-drag left or right with the middle mouse button. Alternatively, hold down the middle and right mouse buttons together while dragging horizontally.

Wireframe


The Wireframe view shows not only 2D Images on image planes, it also shows true geometry, such as that generated by the particle system, text tool, FBX mesh import and basic primitives from the 3D toolset. To display these elements in wireframe rather than shaded mode, click the Wire button in the View toolbar or select 3D Options > Wireframe from the view's context menu.



Changing 3D Viewpoints

By default, the 3D view shows a free perspective camera to get a look at the 3D scene from any angle. Change the view to look through a camera in the scene, a spotlight, or any of the usual orthographic views.

To change the 3D Viewpoints, right-click in the view and select the Camera submenu from the context menu. The choices include Perspective, Front, Top, Left and Right.




If one or more camera objects are present in the 3D environment, these will also be listed as options in this context menu. Spotlights and other lights or objects in the scene will be listed in the Other submenu. Selecting one of these will show the point of view of the camera, light or object chosen. While looking “through” an object in this manner, rotating, panning or zooming the view will instead change the rotation, position and scale of the camera, light or other object.

The “Copy PoV to” submenu is very useful for instantly positioning and orienting a camera, light or other object to match the current point of view shown in the display view. Choosing an object from this submenu will copy the view’s position and angle to the object’s Translation and Rotation controls. Note that this is done in the object’s own coordinate space; any transformations downstream of the object are not taken into account.

A small text label is usually drawn in the bottom corner of the view to indicate which of the available viewpoints is currently displayed. Right-clicking on this label or the axis control above it acts as a shortcut to the Camera submenu, allowing easy selection of cameras and other viewpoints.

Lighting

A default 3D view has Lighting turned off. Enable the display of lighting by clicking on the Light button in the View toolbar, or by right-clicking and selecting Lighting from the 3D Options sub context menu.




If there are no lights in the scene, default lighting will be used. Adding a light to the scene will disable the default lighting, as will unticking the Default Lights option.

To see the effects of a light upon the scene, view the Merge3D tool that combines the light and the scene, or a tool downstream from this. Viewing just the Light tool by itself will show only an empty scene, with nothing to be illuminated.

Shadows

A default 3D view has Shadows turned off. To see shadows in the interactive OpenGL view, enable them by clicking on the toolbar button labeled Shad, or by right-clicking and selecting Shadows from the 3D Options sub context menu.




Enabling shadows will automatically turn on lighting, if this feature is not already turned on. The Shadows shown in the display view will generally be hard edged, as the Open GL renderer cannot display soft edged shadows.

Transparency

A default 3D view uses a method of hiding obscured pixels known as Z-buffering. This method is extremely fast, however it may in some cases produce inaccurate results when displaying transparent pixels, particularly when there is more than one layer of transparency in the scene.


An alternative is the Sorted method. This method can be significantly slower in some scenes, but provides accurate results no matter how many layers of transparency are in the scene.



Change between the two transparency sorting methods by clicking the button labeled Fast in the View toolbar on or off, or selecting one of the options from the Transparency sub context menu.

Grid

By default, the 3D view shows a Grid that 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. The grid is used to help provide a point of reference in the 3D scene. To disable the display of the grid, right-click and select Grid from the 3D Options sub context menu.



Vertex Normals

To display the normals for geometry in the scene, right-click in the view and select 3D Options > Vertex Normals from the context menu that appears. Normals indicate what direction each face of the geometry is facing, and they are used when calculating lighting and texturing on an object.


Quad View

The display view generally only has one surface on which to show an image. It is possible to split the view into four parts. This is called a Quad View. This is most commonly used in conjunction with the 3D view to allow a view of the 3D environment from several different angles at one time.

Although there are four parts to the view, all four will be viewing the same image. When using a Quad view to display a 3D environment, there is a choice 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.


When used with a 2D environment, each part can be used to show a different view type. For example, the image in one part can be shown while the other parts show the alpha channel, vectorscope and histogram.

Enabling and Disabling the Quad View



To toggle the display of the Quad view, select Views > Quad View from the display view's context menu or use the Ctrl-Q keyboard shortcut. A Quad button on the View toolbar is also seen when viewing a 3D environment.


Quad View Layouts



There are a variety of Quad View Layouts, ranging from four equally-sized panels to having three small panels across the bottom of a larger single panel.

To try the various layouts, toggle Quad View on, then select one of the choices found in the Views > Quad Layout's display view context menu.

Effect Masks



Fusion's Effect Masks are used to restrict the effect of a tool to a specified portion of an image. Masks can be polylines, basic primitive shapes, paint strokes or bitmaps from other tools.


To add an effect mask to the currently displayed tool, click on the button for the desired type of mask in the View toolbar or right-click in the display view and select Effect Mask from the context menu.

The toolbar buttons can only be used to add effect masks. Garbage and pre-masks are added via the view's context menu. For more information, see the Effect Mask chapter.

Guides

A number of standard Guides can be displayed in the views. The four guides are: Monitor Safety, Safe Title, Center and Film.

Enabling and Disabling the Guides



Context Menu

Select Guides > Show Guides from the display view's context menu.

Keyboard

Use Ctrl-G to enable or disable the display of guides in the active view.

Selecting a Guide

Monitor Safety

This indicates the safe area viewable on most monitors and TV screens.

Safe Title

This indicates the safe area for titles viewable on all TV and monitor screens.

Center

This shows the Center point and X- and Y-axis of the view.


Film

This displays the Film guides as specified in the Frame Format > Preferences screen. Some frame format defaults have a corresponding default for the guides, whereas some will require customization.

Film Guide 1 contains four text boxes. These boxes specify the offset from the edges of the image for the Left Guide, Top Guide, Right Guide and Bottom Guide, 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.

Film Guide 2's text box is used to set the aspect ratio of the projection area.


Guide Frame Aspect



The Guide Frame Aspect can be used to help show how the current image relates to a specified frame format. The frame aspect is usually set to Default. This forces the frame aspect to the same resolution as the image displayed in the view.


When the frame aspect is set using the Guides > Frame Aspect context menu, the guides will represent the exact boundaries of the specified format. The guides will display the correct frame format and any image area outside of that will be dark grey.

Look Up Tables (LUT)




Look Up Tables (LUT) are used to help match the appearance of a view to its eventual output medium. It is essentially a color modification tool that affects only the view, not the image data used for image processing.

In Fusion, different display views can have different LUT values, ensuring that what is seen is what will be output.




Opening the LUT Editor

To access the Look Up Table editor, select LUT > Edit... from the view's context menu.



Or, select Edit... directly from the LUT flyout menu under the view.



The views are the primary area where color correction is assessed. It is, therefore, crucial that the view is an accurate representation of what the footage will look like when it is back on film.

As with the Color Curves tool, the LUT editor uses spline-based color correction. The splines in this editor can be manipulated and shaped like other splines in Fusion.

Since the purpose of the display LUT is to provide an unchanging correction for the monitor or the file's color space, however, these splines cannot be animated.

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.

Managing LUTs



There are a variety of ways to create and use different display LUTs in Fusion.

Display Settings

Choosing Settings > Save New from the display view's context menu will save all the view's settings, including all LUT curves and gain/gamma values. You can save these under different names, and each set of settings can be reloaded at any time by choosing Settings > name. Save Default will make these settings the standard for all new comps.

LUT Curves

The display LUT edit window can be used to import and export LUT curves for reuse. Right-click on the LUT curve editor and choose Export LUT to get a file browser. Currently it is possible to export the LUT curves as either ASCII or Saved format. The ASCII (.lut) 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.

The Import LUT option will load these files (and other formats) back into the curve editor, or alternatively, if the file has been saved in Fusion's LUTs folder (as specified in Prefs > Global > Paths > LUTS:), it will appear in the LUT file menu list, as described below. LUTs can be exported and imported in this way between views, 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 display view toolbar's LUT popup, or the view's LUT context submenu. 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

Fusion can also use any tool, macro or group of tools as a view LUT, by selecting the tool(s), right-clicking on them and choosing Save Settings. A file browser will appear, and a .setting file describing the tools can be saved in the LUTs folder. This is the most flexible approach, as it allows almost any combination of tools to be used as a display LUT, but is also the slowest. This is because the LUT tools must be rendered solely on the CPU, whereas the other methods are GPU-accelerated when using modern graphics cards.

LUT Preferences

The View Preferences can be used to manage saved display view settings, as well as specifying color and depth defaults for each display view. See the Preferences appendix later in this manual for details.

How Lookup Tables Work in Fusion

A Look Up Table is a table of values used to transform an image. This is generally used to correct for variances in the monitor or the source color space of the image. It is possible to apply a LUT to all displays (regardless of type), or individual LUTs for each 2D display view. Multiple LUTs can be chained together, making it possible to apply a standard Log/Lin color space transformation followed by a 2D LUT curve.

A wide range of LUTs are supported, including industry standard 2D and 3D LUTs, shader fragments embedded in Fuse tools, and macros assembled from standard Fusion tools. Generally LUT processing is performed on the graphics card's GPU in real time. Macro based LUTs will not process in real time, taking as long to render as the tools they contain.

Image LUTs


A different image LUT can be applied to each view, including separate LUTs for the A and B buffers in the view. These LUTs can only be applied to 2D images. Multiple image LUTs can be chained together.

Buffer LUTs

The buffer LUT is applied to all of the view, regardless of contents, including 3D scenes, 3D materials, and sub-view 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.

LUT Processing Order

When a tool's output is displayed in the views the data produced is processed before it is shown in the displays. The processing order is applied as follows:



2D images first have the image LUT applied (if any), and the result composited over the checker underlay. 3D scenes are instead rendered with OpenGL.

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 view, and any controls are drawn on top.

Enabling and Editing Image LUTs

The LUT button in the display view toolbar can be used to toggle the image LUT on and off. The default LUT type is a standard Fusion View LUT, which is initially flat, so there will be no apparent change to the image. To select a different LUT type click on the small arrow to the right of the button and a LUT menu will appear with a list of all known LUT types, including applicable macros and fuses.

The LUT menu can also be found as a submenu in the display view's context menu. Selecting the Edit option from the LUT menu will display a dialog showing a control panel to set the options for each LUT applied.

Enabling and Editing Buffer LUTs

To enable a buffer LUT for all the display views right click in the display view and select the `doc_menuGlobal Options` -> Buffer LUT -> Enable option. The `doc_menuGlobal Options` -> Buffer LUT submenu will contain a list of all known LUT types, including applicable macros and fuses.

Buffer LUTs are often useful for applying monitor corrections, which do not usually change between projects.

Stacking LUTs Together

The display view context menu can be used to stack image LUTs into a processing chain. For example, when loading a Cineon in Log color space, the artist may wish to first apply a Log->Lin conversion, followed by a Fusion View LUT to apply a color calibration.

To insert a LUT after the current one in the processing chain select the type from the `doc_menuLUT` -> Add new submenu in the display view's context menu. To remove a LUT from the stack select the LUT from the `doc_menuLUT` -> Delete submenu.

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

Types of LUTs

Fusion View LUT

The Fusion View LUT is the default and most frequently used LUT type provided by Fusion. A filename selector can be used to load any one of the 2D and 3D LUT formats supported by Fusion. See the File LUT tool section for details.

It also 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 display LUT is to provide an unchanging correction for the monitor or the file's color space, however, these splines cannot be animated.

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.

Log-Lin View LUT

Similar to the Cineon Log tool, this LUT allows linear image data to be converted to logarithmic data, or vice versa. This can be particularly useful when used in conjunction with supplied LUT files that expect logarithmic data.

Macro LUTS

Any Macro tool can also be used as a LUT, simply by saving the macros .setting file to one of the directories specified by the LUTs pathmap in the global pathmap preferences. On Microsoft's Windows Vista this is usually the C:\Users\Public\Documents\Blackmagic Design\Fusion\LUTs\ folder, or the equivalent. Fusion may need to be restarted before macros added to the folder will appear in the menu.

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.

Fuse LUTS

A Fuse of type CT_ViewLUTPlugin can be applied as a LUT to a display view. Shaders can be written and embedded into the LUT for real time processing. See the Fuses chapter for more information.

Fuse LUTs require shader-capable graphics hardware, and cannot be applied in software.

Save and Load LUT Settings

The LUT -> Save and LUT -> Load display view context menu options can be used to save the current LUT settings to disk or to load previously saved LUTs from disk. 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.

Setting a Default LUT

The default LUT applied when a new composition is created can be assigned in the Global -> View Preferences panel.

View Preferences and Settings

View Preferences

The default settings for each display view can be changed using the View panel in the Preferences. The position and size of each floating view can also be saved using the Layout menu in the Preferences.

For details on using the layout and view preferences, consult the Preferences appendix in this manual.


View Settings

It is often preferable to switch between entirely different view configurations while working. For example, while keying, the image is in the main view and the alpha channel is in a subview. View 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 view settings to help reduce the amount of effort required to change from one configuration to another.

Saving View Settings

Once a view is set up as required, select Setting > Save New from the view's context menu to create a new setting. Enter a name for the settings and the overall configuration of the view will be stored in the preferences file.



Loading View Settings

To load previously saved settings, select the settings by name from the Setting menu in the view's context menu.

Loading and Saving Defaults for a View

The view can also be returned to its defaults using the Load Defaults option in the Settings portion of the view context menu. To save a new set of defaults for the view, select the Save Defaults option.

Status Bar Information

Position: X: 0.61642	829	Y: 0.45215	342	Color: R: 0.01961	G: 0.02353	B: 0	A: 1.0
----------------------	-----	------------	-----	-------------------	------------	------	--------


The status bar at the bottom of the Fusion interface will provide the exact RGBA and Z values for the pixel beneath the mouse pointer if the mouse is over one of the display views. Additional information about the X and Y coordinates of the mouse and the exact pixel position are also displayed in the status bar.

The Color Inspector subview can also be used to show more information about the pixels in the view. See Subviews earlier in this chapter for more details.

General Display Options

Display Options Sub Context Menu

The Display Options Sub Context Menu contains several options to display the views. All of the options listed here can be found in the context menu. Several of them have a button in the View toolbar as well.



Locking the Display

Default: Off

A display can be prevented from updating when it is locked. The display will not update the image from the viewed tool until the view is unlocked. The tool still processes and the new image is queued for display in the view. Toggle this option using the Ctrl-L keyboard shortcut, or by clicking on the Lock button in the View toolbar when viewing a 2D image.

Alpha Overlay

Default: Off

When this option is enabled, the view will show the Alpha channel overlaid on top of the color channels.

Overlay Color

Default: White

Change the color used for the alpha Overlay from the default white using the options in this menu.

Checker Underlay

Default: On

The Checker Underlay shows a checkerboard beneath transparent pixels instead of black. This is the default option for 2D views. Toggle this option by clicking on the Checker button in the 2D View toolbar.

Smooth Resize

Default: On

The Smooth Resize option causes the display view to use a smoother interpolated resizing method when zooming into an image in the views. When smooth resize is disabled, scaling uses the faster nearest neighbor method. This option is enabled by default and can also be toggled by clicking on the SmR button in the View toolbar.

Follow Active

Default: Off

Enabling the Follow Active option will cause the view to always display the currently active tool in the flow.

Show Controls

Default: On

The Show Controls option shows or hides onscreen tool controls in the view. This option is handy when the controls are not necessary or are getting in the way of evaluating the image. This option is toggled using the Ctrl-K keyboard shortcut or by clicking on the Show Controls button in the View toolbar.

Show Square Pixels

Default: Off

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.

Fusion will automatically perform aspect correction when viewing non-square images so that the image seen in the views does not appear squashed or stretched. Although this aspect correction helps to correct the shape of the objects in the image, it can also cause aliasing in the image. Fusion uses a relatively fast resize algorithm to make the correction in the views. This aliasing does not affect the final result. It is an artifact of the display view only. Regardless, its presence can obscure real problems in the image.

Enabling the Show Square Pixels option will override the aspect correction in the view to confirm that this ‘tearing’ is an artifact of aspect correction and not a problem with the composite.

Show square pixels can also be toggled on and off using the 1:1 button in the View toolbar.

For more information on the effect of aspect on the composite and how Fusion handles aspect ratio, see the Frame Format chapter in this manual.

Pixel Grid

Default: Off

Enabling this option will show a light black grid that outlines the exact boundaries of pixels in the image when the image is scaled in past a certain threshold. Toggle this option with the Grid toolbar button.

Show Full Color Range

Default: Off

Enabling the Show Full Color Range option causes Fusion to automatically normalize any image displayed in the view. 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.

This allows for the visualization of values outside of the normal viewing range, particularly when working with floating point images. 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.


Toggle this option with the Related View toolbar button.

Display View Label



Each display view has a text label at the bottom right of the View toolbar. This Display View Label indicates the name of the tool currently viewed in the display. Clicking on the label will select the viewed tool and, if the tool is not currently visible in the flow area, it will highlight and pan the tool in the flow editor.

Controls List



Click on the small arrow button to the far right of the label for a popup list of the onscreen controls for the selected tools. The currently active control is marked on the menu. Choosing a different control from the menu will activate it for use. This can be handy when several controls are in the exact same position.

If multiple tools are selected, the menu will provide a complete list of all of their onscreen controls. Selecting a control from the list will highlight and activate this control on the view.

Pressing Tab or Shift+Tab will cycle through the onscreen controls, activating each in turn. The view's context menu has further options for control management in the Controls submenu.




Flow Editor

9

Flow Editor

Overview




The Flow Editor is displayed below the large image display and is Fusion's primary workspace. The flow provides a schematic-type overview of the structure of the project. Each tool is represented by a tile in the workspace, with connections to and from the tile to indicate how each tool relates to the overall processing of the image result.


This workspace is known as the flow because images move through the connections like water through pipes. By watching the flow of images through the pipes, an intuitive overview of the current image processing operations can be formed.

This view offers tremendous flexibility for connecting tools, including the ability to branch a tool's output into separate streams and collapse entire groups of tiles to view a thumbnail image of a tool's, or several tools', outputs.

The following illustration shows a simple flow. This flow contains a loader that brings footage in from disk and that loader is connected to a resize tool. The resize tool is, in turn, connected to a saver for rendering the result to disk.



The next illustration shows a more complex flow. It combines multiple layers of footage and demonstrates the idea of branching.




Note: Previous versions of Fusion would save projects to disk using the file extension .flw. This was commonly referred to as a flow. The use of the term flow to refer to a part of the interface as well as the saved project caused some confusion. Fusion renamed the saved project to .comp (for composition). As of version 5, the term flow refers to this portion of the interface only.

Viewing, Scaling and Panning the Flow

Displaying the Flow


The Flow Editor is displayed by default whenever a new flow is created or when an existing flow is opened. Since the Flow Editor shares the same space in the interface used by the Spline Editor, Timeline Editor and Console, switch between these views as needed.



To display the flow when it is not currently visible, either select the Flow tab from the tabs at the top of the work area or press the F5 key on the keyboard. When the flow is the active area of the interface, it will be surrounded by a yellow border.

Resizing the Flow


To resize the flow, place the mouse over the edge of the raised border between the display view and the flow. Click and drag on this boundary to adjust the size allocated to the display view and the Flow Editor respectively.



To maximize or minimize the flow, place the mouse pointer over any part of the flow and hit the F4 key once. This will expand the flow to cover the display views or return a maximized flow editor to its original state. A maximized flow takes over all of the space used by the display views.

Panning the Flow

To pan the view of the flow, hold the middle mouse button down and drag in the desired pan direction, or click-drag on the standard scrollbars at the right and left sides of the flow.



If the mouse does not have a middle mouse button, hold the Ctrl-Shift keys down while clicking the left mouse button. This will have the same effect as holding down the middle mouse button.

If a tool that is selected is not visible in the flow or in any other part of the Fusion interface, the display will automatically pan the view to show that tool's tile in one of the four corners of the visible area.


Click-drag in the displayed flow navigator to change what portion of the flow is visible. See below for more information on the navigator.

Context Menus

There are two context menus in the flow editor where most functions in the flow are accessed: the Tool and the Flow context menus.


Flow Context Menu

Right-clicking on an empty region of the flow will display a context menu with options and functions for manipulating the flow display in general.



Tool Context Menu

Right-clicking directly over a tool in the flow displays the tool's context menu with options and functions that are specific to that tool. More details on the various options of the context menus can be found at the end of this chapter.




Adding Tools

The flow is always empty when a new project is started, since no tools or loaders have been added. Tools are added to the project in the flow using one of the methods described below.

Using the Menus

The tool's menu at the top of the screen can display every single tool available to Fusion. They are organized into general categories. To add a tool to the flow from this menu, navigate down the category list and select the desired tool from the menu's categories.




If another tool is currently active on the flow, the new tool is automatically inserted downstream of the existing tool. If there is no current tool, the new tool stays disconnected from the rest of the flow and it will be placed in the first available space found by Fusion in the flow.

Using the Toolbars

The Toolbars can also be used to add tools to the flow. These small bars display an abbreviated form of each tool's name (i.e., BC for the brightness/contrast tool). Hold the mouse over any toolbar button for a moment to display a pop-up with the unabbreviated name of that tool.

To add a tool to the flow using the toolbar, click on the button once with the mouse, or drag the tool into the flow from the toolbar and drop it on the flow.


A simple click on the toolbar will add the new tool to the flow or it will be inserted depending on the same conditions described in Adding Tools from the Menus (as described above).



Choose where the tool is added to the flow by dragging the toolbar button into the flow. This offers the ability to replace one tool with another or to insert the new tool between two others. These operations are described in greater detail further below.

Tools can also be dragged from the toolbars into the large and small views. This allows tools to be added to the project even if the flow display is not currently visible. Tools that are dragged onto a view from the toolbar are added to the flow immediately after the currently viewed tool, regardless of the currently selected tool.


The toolbars can be customized to display different tools. They can be docked at the top of the screen to either side, or they can float. To customize toolbars, right-click on the them and use the dialog box that appears.



For details on customizing toolbars, refer to Toolbars in this manual.

Using the Context Menu

Right-clicking in the flow will display the flow's context menu, which may contain one or more tool submenus. These are copies of the tool menu described above and can be used as a shortcut for adding tools to the flow.




Add or Insert Tool

Selecting a tool from the Add Tool or Insert Tool submenus will insert a tool into the flow at the position where the mouse pointer was when the context menu was opened.

If the mouse pointer is positioned over a connection between two tools, the new tool will be inserted into the flow already connected to that pipe (between the existing tools).

Replace Tool

The Replace Tool option only appears when the context menu is opened by right-clicking directly on an existing tool's tile in the flow.



Selecting a tool from this menu will replace the current tool with a new one.

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

Using the Bins

The Bins can display thumbnail icons for each of the tools and third party plug-ins that are known to Fusion. To add a tool from the bins to the project, drag and drop the tool's icon into the flow. The bins usually start empty. To add the standard Fusion tools to the bins, right-click in the bins and select Create Tool Folders.

Bins can contain more than tools. They contain saved tool settings, tool groups, macros and footage, making bins much more flexible than toolbars. For more information on the bins, refer to the Bins chapter in this manual. To toggle the bins on and off, select Ctrl-B from the keyboard, or select File > Bins from the menu.

Footage

Dragging a file from the file browser to the flow will automatically add a loader to the flow, directly loading the file or file sequence. If multiple files are dragged into the flow, a separate loader is added for each file.


Saved Tool Settings

Dragging a tool's settings (.settings) file from the file browser to the flow will add the tool (or tools) saved in that file to the flow, configured with the settings that were saved into the file.

Connecting and Disconnecting Tools

Each tool in a project has its own tile in the flow. These tiles show the inputs and output for the tool as small colored nodes attached to the tile. The arrow-shaped nodes represent the image inputs, and the grey or red square nodes represent the tool's processed output.


A tool must have at least one output node. It generally also requires one input node, unless it is a creator tool. Almost all tools also have at least one mask input, which is usually colored blue. A tool cannot have more than one output. A classic example of a multi-input tool is Merge, which has a background input, a foreground input, and a mask input.



Hold the mouse pointer over a tool's input for a moment and a small pop-up containing the input's name will appear.


Connecting Tools

Build a flow by connecting tools together using the inputs and outputs. For example, to connect the output of a Background tool to a Color Corrector tool, click on the output node of the Background tool and drag the connection over the input node of the Color Corrector. A small line, or pipe, will be drawn between these two nodes to represent the connection.



Disconnecting Tools

The pipe between two tools is divided into two halves the output half and the input half. To disconnect a tool's inputs or output, click on the input half of the connection pipe or click on the input arrow itself and drag to pull the connection away from the tool.




An entire tool can be disconnected by holding down the Shift key while moving the tool. If the Shift key is depressed when letting go of the tool, all existing connections to that tool will be removed.

Replacing Connections

To replace an existing connection on a tool with another one, simply drag the new connection over the existing one.

Tracing Pipes Through the Flow

Several visual cues exist to help you follow the pipes through the flow. As an option, pipes can always been drawn on top of any other tools.



Placing the mouse over a specific tool will cause the pipes connected to that tool to become highlighted. Placing the mouse directly over a pipe will cause a tooltip to appear with details on its source output and destination input.

Branching

The output of any given tool can be connected to multiple inputs on other tools, which is a feature in Fusion called branching. Branching permits a single image to be separately processed by two or more streams of tools.




Only one connection per input is allowed but a tool output can have any number of branches and outbound connections.

Inserting Tools

Use one of the following methods to insert a tool into an existing connection between tools.

Drag and Drop

Click on the tool to insert from the bins or toolbars and drag it into the flow. While still dragging the tool, position the mouse over the pipe of the tool to be connected so that the pipe is highlighted in blue, then let go. This will cause the new tool's first input to be connected to the upstream tool and the output to be connected to the downstream tool.



Shift Insert

If the tool to be inserted is already present on the flow, hold the Shift key down and drag the tool over the pipe where it should be inserted. When the mouse is positioned correctly, the pipe becomes highlighted. Release the mouse button to insert the tool.


From Toolbars and Menus

Select the tool in the flow to which the new tool will be connected. Click once on the toolbar button of the tool to be added, or select the desired tool from one of the tool menus. The new tool will be added immediately after the selected tool.

Replacing Tools

To replace a tool, drag the new tool onto the flow from the bins or toolbar and place the mouse pointer directly over the tool to be replaced.

When the mouse button is let go, a confirmation dialog will appear asking if the existing tool is to be replaced with the new tool.




An extremely powerful feature is the fact that if the new tool has settings that are identical to those in the original tool, these settings will be copied to the new tool. For example, replacing a Transform tool with a Merge will copy the existing center, angle values from the transform to the Merge.

Deleting Tools


To delete a tool or tools from the flow, select the tools in the flow and press the Delete key on the keyboard. Multiple tools can also be selected and deleted. Alternately, right-click on the tool(s) and select Delete from the context menu that appears.

Connection Types

By default, Fusion uses direct or linear pipes to connect tools together.



Switch to the orthogonal connection style used in previous versions of Fusion by selecting Options > Orthogonal Pipes from the flow context menu.



Moving Tools and Organizing the Flow

Larger flows can become disorganized and difficult to understand so several techniques for organizing the flow exist to assist.


Moving Tools with the Mouse

To move a tool, simply left click on a tool and drag it to its new position in the flow. All currently selected tools will be moved.

Lining Up Tools

Right-click on an empty section of the flow to display the flow's context menu. Choose Line Up from this menu to cause all of the tools to line up vertically and horizontally.

To line up only the currently selected tools, right-click on one of the tools and select Line Up from the menu. Tools that are not currently selected will be left as they are.



Align to Grid

Tools can snap to an invisible grid as they are positioned on the flow. Enabling the Arrange Tools > To Grid option in the flow's context menu will enable this feature. As tools are positioned in the flow, a faint rectangle will indicate to where that tool will snap when the tool is released.


To set the Arrange Tool > To Grid option as the default for new compositions, simply set the Preferences > General > Arrange To Grid checkbox to On.

Groups

Groups are containers on the flow that can hold multiple tools in a small container. They are excellent for keeping the visual complexity of a flow to a minimum.


Groups can contain other groups as well as tools. There is no limit to the number of tools that can be contained within a group.

A group can be collapsed or expanded. A single tile on the flow represents a collapsed group, where a group that has been expanded is represented by a subflow showing the tools contained within the group.



Creating Groups


To create a group, select the tools to be added to the group. Right-click on one of the selected tools and choose Group from the context menu that appears. Alternately, you can use the Ctrl-G keyboard shortcut.



The tools contained within the group disappear from the flow and are replaced by a single tile that represents the collapsed group. If any of the tools in the group have any unconnected, or free, inputs when the group is created, those inputs will not be visible in the group. When a group is created to be saved to disk for later re-use, connect the input to another tool before creating the group if that input needs to be visible.


Expanding and Collapsing Groups

A collapsed group is represented by one tile on the flow. To open a collapsed group, click on the small icon to the left of the group's name in the tile.



This will reveal a subflow list of the tools within the group. The subflow is a separate area that can be scaled, resized and panned independently of the main flow. To collapse the group again, click on the minimize icon at the top right corner of the subflow.


The controls for the various tools in a group can only be accessed if the group is expanded. Double-click on a tool tile within a group to display that tool's controls. The controls will disappear again when the group is collapsed. Tools can only be added, inserted or deleted from a group while that group is expanded.



Use the Ctrl-E keyboard shortcut when a group is selected to switch between expanded and collapsed modes for that group.

Ungrouping Tools

To remove the group without deleting the tools within it, right-click on the group and select Ungroup from the context menu. The tools inside the group will be placed back in the main flow.




Deleting Groups

Deleting a group is no different from deleting a single tool. The group is removed from the flow, including any tools contained within that group.

Renaming Groups

To rename a group, select the group and hit the F2 key, or right-click on the group and select Rename from the context menu.



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. Use one of the following methods to save a group's settings.

Drag And Drop

Drag a group from the flow into the bins. A dialog will appear that requests the name of the new settings file and the path where it should be saved on disk. The .settings file will be saved in the specified location and placed in the bins for easy access in the future.

Context Menu

A group can also be saved by right-clicking on the group and selecting Settings > Save As from the context menu.

Locking, Panning and Scaling the Group's Subflow

Click on the small hand icon in the title bar of an open group to enable/disable mouse panning and scaling for the group's subflow. This defaults to enabled for new groups.

Displaying Tile Pictures


It is possible to expand a tool's tile so that it displays an informative thumbnail image of either the tool's icon or the rendered tool's output. These small thumbnail images can be an enormous help in visualizing the operation of a flow.

Turning On The Tile Pictures

To turn on tile pictures for one or more tools, first select the tool(s) that are to have their tiles enabled. Right-click on one of the selected tools and choose Show > Show Tile Pictures from the menu. Additionally, enable tile picture using one of the following four options.


Force Active Tile Pictures

This choice makes the active tool display a tile picture.



Force Source Tile Pictures

Display tile pictures for all source tools in the composition. For example, Loaders, Text, and BG tools will all show full tile pictures.



Force Mask Tile Pictures

This choice will cause all mask tools in the flow to show a tile picture.

Force All Tile Pictures

This option will cause all tools in the flow to show a tile picture.

Show Thumbnails

To switch between showing the thumbnail image rendered by the tool and the generic icon that represents the tool, use the Show Thumbnails option in the flow context menu.

Thumbnails That Do Not Display

A tool may not display its tile picture under certain circumstances, as described below. In these cases, the default icon for that tool is displayed instead of an image.

Concatenated Tools

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. Affected tools include transforms and merges.

Particle Tools

The only particle tool that generates an image is the pRender tool. All other particle tools display their default icon instead of an image.

3D Tools

The only 3D tool that generates an image is the Render 3D tool. All other 3D tools display their default icon instead of an image.

Failed Tools

If a tool fails to process for some reason, it will not produce a thumbnail image. This usually happens when an upstream loader has no footage, is out of range or if the tool's valid region has been trimmed in the timeline.

Pass Through


Tools that have been set to Pass Through mode are effectively disabled and do not generate an image to display in the Thumbnail view.

Viewing Tools

Any tool in the flow can be viewed in a display view or through the frame buffer of supported video hardware. There are several methods for viewing tools.


For more information about working with the image views, refer to the Display Views chapter in this manual.

Using Drag And Drop




Drag a tool into the views by selecting it in the flow and dragging the mouse over the view to see its image. The tool will appear to move at first but will snap back into its original location once the mouse leaves the area of the flow editor.

You can also drag the tool's header from the control area into the view.



Using Tool Tile View Indicators

A tool displayed in one of the views will have a series of indicators displayed under its tile in the flow. One indicator is present for each available view. The indicator will turn white for the view in which that tool is displayed.




To see what view is related to which indicator, hold the mouse pointer over the indicator until an informative tooltip appears. Click on one of these indicators to toggle the display of that tool's image in the related view.

If a tool is not currently showing in any of the views, holding the mouse over the tile will make the indicators visible and available for clicking.

Using the Context Menus

View the output of a tool on one of the image views by right-clicking on the tool and selecting View On from the context menu for that tool.



Using Keyboard Shortcuts

Each image view 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. The left display view is 1, the right display view is 2, and additional floating displays and hardware frame buffers are assigned numbers in sequence.

Tooltips and the Status Bar

Tooltips


Hold the mouse pointer over any tool to display a pop-up Tooltip containing that tool's name, aspect, resolution and color depth.

Any notes that have been added to the tool's comments tab will also be displayed in the tooltip.

Status Bar


The Status Bar is at the bottom of the screen, to the left of the interface. It also displays the name, aspect and resolution of the tool beneath the mouse pointer.

Copying and Pasting Tools



Copying Tools

To copy a tool on the flow to the clipboard, right-click on the tool and select Copy from the context menu or select the tool and choose Copy from the edit menu at the top of the screen. The keyboard shortcut Ctrl-C will also work.




Cutting Tools

To cut a tool on the flow and place it in the clipboard, right-click on the tool and select Cut from the context menu or select the tool and choose Cut from the edit menu at the top of the screen.

If the Flow view is active (with a highlighted boundary), the keyboard shortcut Ctrl-X will also perform a cut, as will the View menu's Cut option and the Cut toolbar icon.

Note: This will remove the selected tool from the flow.

Pasting Tools



To paste a tool or tools back into the flow from the clipboard, right-click in the flow and select Paste from the context menu. The tools will be inserted at the location clicked to bring up the menu.

To replace a tool already in the flow with the contents of the clipboard, right-click on the tool to be replaced and select Paste from the tool's context menu instead.

To insert a tool from the clipboard between two other tools, right-click directly on a pipe between two tools and select Paste from the context menu.


The standard keyboard shortcut Ctrl-V can be used to perform a paste, as can the View menu's Paste option, and the Paste toolbar icon. In these cases, the tool will be placed immediately after the currently active tool, and will be automatically connected to it. If the active tool is in the middle of a chain, the new tool will be inserted, with connections being made and other tools being moved as necessary.

If the pasted tool is a source tool with no primary image input, a Merge tool will be inserted instead, the new source tool will be placed at the bottom left corner of your tools and connected to the Merge's foreground input. This auto-merge behavior can be changed with the Defaults page in Preferences.

If no tool is currently active, or if multiple tools are on the clipboard, they will be placed at the bottom right of your tools. If the tool on the clipboard is a source tool, it will be placed at the bottom left instead.

Pasting Tool Settings

Fusion also offers the ability to paste the settings of the tool from the clipboard into another tool. This function has no keyboard shortcut but can be accessed by right-clicking on the tool into which the copied settings will be copied.



When copying the settings from one tool to another, only the settings that match between two tools will be copied. For example, when pasting settings from a Merge to a Transform, the center control values, the blend and the size and angle controls will be copied from one to the other.

Using a Text Editor

As of Fusion 5.1 and later, tools can be copied from the flow and pasted into a text editor. This will display the tools' data as it would be saved to disk. The text can then be edited and pasted back into the flow, either as a new tool, or to replace the settings of an existing tool.

Instancing Tools

Normally, when a tool is pasted into the flow, a copy of that tool is created. Changes made to the copied tool are not made to the pasted tool.


There are times when two tools need to have the exact same settings at all times. This can be done by publishing each control on the source and connecting the matching controls on the target tool. This becomes prohibitively complex and time consuming on tools with several hundred controls.

The solution is to link the tools as an Instance. Instanced tools share the same controls a change to one tool affects all of the tools in the instance.

A control on a tool is visibly instanced when a yellow border surrounds the text edit box to the control's right.

Creating an Instance


To create an instance, copy the tool to be used in the clipboard using Ctrl-C or by selecting Copy from the tool's context menu. Right-click in the flow and select Paste Instance from the flow's context menu or use the keyboard shortcut Ctrl-Shift-V.



De-Instancing an Instanced Tool


To remove the instance link from a tool, right-click on the tool and select De-Instance from the context menu. A tool cannot be re-instanced once this has been done (although undo still works).

An instance can also be recreated by copying the source and pasting the instance again.




De-Instancing a Control

Occasionally, only one control in an instanced tool will need to be adjusted without having to de-instance the entire tool. In this case, right-click directly on the control and select De-Instance from the control's context menu. The yellow border around the control's edit box should disappear.




Re-Instancing A Control On An Instanced Tool

Re-instance a de-instanced control by right-clicking on the control and selecting Re-Instance.



Showing Instance Links Between Tools


The composition shows the links between instanced tools as green lines.



This can be disabled until needed, or to reduce clutter in the composition. To disable the instance links, right-click in the flow and select Options > Show Instance Links from the context menu.

Finding Tools

Modern effects shots and compositions can be quite detailed. Compositions of several hundred tools are not rare, and locating tools can be challenging. To quickly locate tools in the flow, select the Find option from the flow's context menu. This will display a dialog to enter a search term. The dialog also offers several checkbox options that can be used to filter the search further.



Select Find Next or Find Previous to find and select one tool in the flow that has a name matching the search term, or select Find All to select all of the tools in the flow that match the selected name.

This is particularly useful if you want, for example, to bypass all resize tools. Find all will select all of them and you can pass them through by pressing Ctrl-p.


Macro Tools

The creation, editing and use of macros are covered in depth in the Macro Tools chapter later in this manual.

Flow Options

The Flow Options are contained in the flow context menu. This submenu contains several choices that can be used to customize how the flow looks and behaves.

Flow Options' Context Menu



Pipes Always Visible

This option will cause pipes to be rendered on top of the tools instead of beneath them.

Aspect Correct Tile Pictures

This forces the display of tile pictures to be aspect corrected, which is slower but more accurate.

Full Tile Render Indicators

When a tool is rendering, a small green light usually appears in the bottom right corner of the tile. Enabling this option causes the entire tool to flash green while it is rendering, which makes it easier to identify what tool is processing in a large, complex flow.

Show Grid

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

Show Instance Links

When this option is enabled, the flow draws a green line between an instanced tool and its parent.

Auto Remove Routers

Orphaned Elbows or Router nodes, that means Routers with no input or output, are deleted automatically from the flow.

Show Navigator

Enabling this option displays a small overview window of the entire flow in the Flow Editor's top right corner. The keyboard shortcut V can be used to toggle this option's state on an active flow.

Auto Navigator

The Navigator will automatically come on when any tool is outside the visible flow area and go off when all tools are within the flow area.

Build Flow Vertically/Horizontally

This option is used to determine whether new tools will be added beneath the current tool or to the right of the current tool.

Orthogonal/Direct Pipes

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

Tool Context Menu


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

View On

Select an option from this submenu to display the tool's output on one of the available image displays or hardware frame buffers.

Create/Play Preview On

Create a preview of the tool's output on one of the available image displays or hardware frame buffers. Also create a preview by holding down Alt while dragging the tool into the display view. Hold down the Shift key when selecting the preview from the menu, or hold Shift-Alt to skip the render dialog and to start creating the preview immediately with the default settings or the last settings used to create a preview.




Settings

The options in this submenu are used to save a tool's settings to a .setting file, to load .settings files from disk, as well as to create or remove saved defaults for that tool.

Cache To Disk

Select this option to display a dialog used to create and manage the Disk Cache. The disk cache is covered in more detail in its own chapter later in this manual.



Modes

Update

This option is selected by default. If it is turned off, the tool will still allow changes to its settings and parameters but it will not update or process an image unless the Update option is turned back on.

Locked

Locking a tool will prevent changes from being made to the tool's settings.

Show Controls

This option is on by default. If Show Controls is turned off, the tool will no longer appear in the tool controls area. Its onscreen controls will no longer appear in the display views.

Pass Through

A tool that has been passed through will not process its images or masks when rendered. The tool will simply Pass Through any images it receives in its inputs.

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.

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.

Set Render Range

This menu command will Set the Render Range of the flow to the valid range of frames available to the tool. Especially useful when you want to limit the render range, for example, to the valid frames of a Loader.




To set the render range you could also drag the tool directly onto the Time Ruler.

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 [Displaying Tile Pictures](#) for more information.

Rename

Rename changes the name of the selected tool(s).



Swap Inputs

This menu command will swap the Primary and Secondary inputs on a multiple input tool, such as the foreground and background of a Merge tool. The keyboard shortcut Ctrl-W will also swap the inputs of the selected tool.

Cut and Copy

Cut and Copy the tool to the clipboard. See [Copying and Pasting Tools](#) 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

Collapse the selected tools into a group or uncollapse 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.


Edit Splines

Switch the view to the Spline Editor. The splines for the selected tool(s) are displayed and enabled for editing.

Tool Tile Color Coding and Icons

By default, the tile display color for all tools in the flow is light grey. Certain tiles will change color to indicate special properties or selection states.

Tile Properties/Selection States



Green

Green tools are source tools, either loading footage from disk or creating images from scratch.

Yellow

A Yellow tool tile is the actively selected tool. This overrides all other tool coloring.

Dark Blue

A Blue tool tile is used to indicate a selected tool.

Light Blue

A Light Blue tool tile indicates a 3D tool.

Purple

Particle tools are all colored Purple.

Pink

Savers are the only tools that are displayed in Pink.

Red

A Red tool tile indicates that the tool has failed to render for some reason.


Passed Through Tool Colors

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


Mode Icons

The following icons are used in the tile pictures to indicate that a tool has a special mode or option applied to it.

Cache to Disk



Locked






The Spline Editor

10

Spline Editor

The Spline Editor is where the animation of various controls is manipulated in Fusion. All animated controls are connected to a spline curve, as described below. One or more splines can be displayed in the editor for adjustment.



What Are Splines?

All animation in Fusion is controlled by splines, which are mathematical curves that describe the value of a control at any given point in time. Splines are represented on a graph, with time as one axis and the value of the control as the other.


The advantage of using splines to represent animation is that splines do not require a value to be set for each and every frame of a project. For example, if a value of 1.0 is set on a control at frame 1, followed by a value of 10.0 at frame 10, the animation spline would automatically apply a value of 2.0 at frame 2, a value of 3.0 at frame 3, and so on.

The values explicitly specified are known as keyframes. The values between keyframes are smoothly interpolated or calculated based on the shape of the spline.

Creating Animation Splines

From the Context Menu

To create a spline, in the tool controls right-click on the control to be animated and select one of the following options.



Animate

Selecting Animate from the context menu will connect to the control to the default spline type. This is usually a Bezier Spline, unless the default spline has been adjusted in the defaults panel of the preferences.

Other spline types can be used via the Modify With context menu's submenu.

Modify with Bezier Spline

Bezier Splines are the default curve type. Three points for each keyframe on the spline determine the smoothness of the curve. The first point is the actual keyframe, representing the value at a given time. The other two points represent handles that determine how smoothly the curve for the segments leading in and out of the keyframe are drawn.

Modify with >B-Spline

B-splines use a single point to determine smoothness of the curve. Instead of using a keyframe and three handles, a single keypoint determines the value as well as the smoothness of the curve.

Modify with >Natural Cubic Spline

Natural Cubic Splines are similar to Bezier splines, except that the Bezier handles for this spline type are locked and invisible, and always set to make the smoothest possible curve. This spline type is almost never used.


Modify With Wqq>NURB Spline

A NURB spline uses NURB-type spline interpolation between points on the spline.

See the Modifier Chapter for more details on the individual Spline types.

Removing and Deleting Animation Splines


To remove an animation spline from a control, right-click on the control and select Remove [tool control's name] from the context menu. The spline tool will be disconnected from the tool control. 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.



Spline Editor Interface


The Spline Editor Interface has three main working regions: the Graph view, Tree view and Toolbars.

Graph View




The Graph View is the largest region of the interface. It is here that the actual animation spline is displayed for editing. There are two axes in the Graph view. The horizontal axis represents time and the vertical axis represents the spline's value.

Tree View



The Tree View provides a mechanism for determining what splines are visible in the Graph view. The name of each spline in the project is shown beneath the tool that contains the animated control. The checkboxes beside each name are used to show whether that spline is currently shown in the Graph view and whether the spline is enabled for editing.

Toolbar



The Toolbar across the top of the spline editor represents the most common operations applied to a spline, or splines, in the Graph view.

The various operations represented by the toolbar are all accessible from the Graph view's context menu, but the following buttons provide a faster shortcut.

The operation of each button is described later in this chapter.

Current Time Indicator

The thin yellow bar that runs vertically through the Spline Editor represents the current time of the flow. Click-drag on the Current Time Indicator to change the current time.

Status Bar

The Status Bar at the bottom left of the main interface constantly displays information about the position of the mouse pointer along the time and value axis whenever the pointer is over the Graph view.


Time: 81.2, BrightnessContrast3Gamma: 1.58897, Slope: 0.00725

Context Menus

There are two context menus accessible from the Spline Editor, as described below. The options found in these menus are described in detail through the rest of this chapter.


Spline Context Menu

Right-click anywhere in the Graph view to display the Spline context menu.



Guide Context Menu

To see the Guide Context Menu, right-click on the grey region of the horizontal (time) axis at the top of the Graph view.




Viewing, Scaling and Panning the Spline Editor

Viewing the Spline Editor

Using the Tabs

Click the Spline Editor Tab located at the top of the flow area to switch the view to the Spline Editor.



Using the Keyboard

Press F8 to switch to the Spline Editor.

Using the Context Menus

Right-click within the Flow or Timeline views and select Edit Splines from the context menu.

Scaling the Spline Editor

It is often necessary to adjust and scale the display area within the Spline Editor to ensure that the splines are visible. There are a number of different ways this can be accomplished.


Using the Toolbars

- > The Scale buttons are located in the top left corner of the Graph view.
- > The Fit button attempts to rescale the view so that all currently active splines fit within the Graph view.
- > The buttons with the magnifying glass icons will scale the Graph view horizontally or vertically.



Using the Axis Labels

- > Place the mouse pointer over the darker labels for the horizontal or vertical axis and click-drag to resize the Graph view on that axis only.
- > The view is scaled with the position of the mouse pointer as center.



Using the Mouse

- > Place the mouse pointer over the Graph view, then click and hold the middle mouse button. With the middle mouse button still held down, click once on the left mouse button to zoom in and once on the right button to zoom out.

Using the Keyboard

- > Keyboard shortcuts can be used to scale the display of the Graph view while the Spline Editor is active.
- > The + and - keys on the numeric keypad will scale in and out of the view.
- > Ctrl-R will enable scale to rectangle mode, which allows a click-drag of a rectangle around the area of the Graph view. That area will be scaled to fit inside the view.

Panning the Spline Editor

To pan the Spline Editor, click the middle mouse button while dragging the mouse to the left or right or use the scrollbar at the bottom of the Graph view.


In general scaling and panning the Spline Editor works the same as in all navigatable parts of the Fusion interface. See the Flow Navigation Chapter for more details on this topic. Displaying Splines For Editing

The menu on the left of the Spline Editor is a tree menu. The top level of the menu shows the name of any tool in the flow that has an animated control.

The next level, which is revealed by clicking on the + symbol next to the tool's name, shows all of the names of the animated controls in that tool.


Tool Status Checkbox

Beside the name of each spline is a checkbox that indicates the spline's status. Changing the selection mode of the checkbox will change the status of that spline. There are three selection modes for each checkbox: Active, Viewed and Disabled.




Active

When the checkbox is enabled, that spline is displayed in the Graph view and is Active. This spline can then be edited.



Viewed

When the checkbox is enabled but the check is grey, the spline is visible in the Graph view and is read-only. It cannot be edited.



Disabled


When the checkbox is cleared, the spline is not visible in the Graph view and cannot be edited.



Clicking directly on the checkbox will toggle it between these three states. Changing the state of the checkbox at this level will set the mode for all splines attached to that tool.

Selection States

There are three Selection States that determine how the tree behaves when a checkbox or label is selected to activate a spline. These buttons are located in the spline toolbar, just above the tree. They are labeled All, None and One.



All

Clicking on this button activates all splines for editing.

None

All spline checkboxes will be set to grey, or Viewed mode, on the first click of this button. On the second click, all checkboxes will be cleared.

One

This option is a toggle. When the One button is selected, only one spline in the tree is active and visible at a time. Clicking on any spline's checkbox will set it to active and all other splines will be cleared.

Sel

When active only the Spline of the currently selected tool(s) will be displayed.

Show all control

By default only animated controls are displayed. Activating this switch will display all controls of all currently displayed tools. 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 controls, this might lead to cluttering and slow downs of the interface. We recommend using this option in conjunction with the Sel option to limit the number of tools and controls displayed and yield optimum performance.

Labels

Clicking directly on the name label for a spline in the spline tree will activate that spline for display and editing, while disabling any other splines that are currently active.


Working With Splines

Renaming Splines




The default name for a spline is the name of the control it animates. It may be useful to assign more descriptive names to a spline, especially if a composition is shared between several artists. Right-click on the name of the spline in the tree menu and select **Rename Spline** from the context menu.

Show All Inputs



Selecting this option will show all inputs on all tools, regardless of whether they are animated or not. This mode can get very slow with large compositions. It is handy when animating a bunch of controls quickly, since adding a point to the flat curve for any control will automatically animate the control with a Bezier spline.


Changing Spline Colors



Each spline in the Graph view 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 Tree view will also be drawn in the same color as the spline.

- > To change the color of the spline, right-click on the name of the spline in the Tree menu and select **Change Color** from the context menu.
- > Select the new color from the dialog box that appears and click **OK**.

Selection Groups



It is possible to save the current selection state of the splines in the tree, making selection groups that can easily be reapplied when needed. To create a selection group, open the spline's context menu and choose **Save Current Selection**. A dialog will appear to name the new selection.


To reapply the selection group, choose the selection group by name from the **Set Selection** menu in the context menu. Other context menu options will allow selection groups to be renamed or deleted.

Working with Keyframes


Once a spline is displayed in the Graph view and enabled for editing, the spline's keyframes (and thus the animation) can be manipulated in a variety of ways.

Selecting Keyframes

- > To select keyframes on a visible spline, click directly on the keyframe, or click-drag a small rectangle around the keyframe.
- > To select multiple keyframes, click-draw a rectangle that encloses all of the keyframes to be selected.
- > To add or remove a keyframe from the current selection, hold **Ctrl** down while selecting the points. This will remove the currently selected keyframe and add unselected points to the selection.
- > Use the **Ctrl-A** shortcut or choose **Select All** from the Spline Editor's context menu to select all keyframes from the active splines.



Moving Keyframes



Keyframes can be moved with the mouse, keyboard or the edit point controls. A keyframe cannot be made to overlap an existing point. If a key exists at frame 2 and frame 5, the key at frame 5 cannot be moved to frame 1.

With the Mouse

- > Click-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 the Alt key down before dragging the point to constrain its motion to a single axis.

With the Keyboard


- > Use the cursor keys to move selected keyframes by small fixed amounts. The left and right cursor keys will move the keyframes by a single frame. The up and down cursors will adjust the value of the keyframes by a small amount.
- > The current scale of the Graph view determines the degree of vertical movement applied to the value with each keypress. The closer the zoom in the spline, the finer the adjustment.
- > Hold the Shift key down while making keyboard adjustments to a keyframe to increase the amount of adjustment.

With the Edit Points Controls

- > The Edit Points Controls are found in the toolbars as wide buttons with text labels that usually say Value and Time. These allow explicit values to be entered for selected keyframes. These controls are explained in more depth later in this chapter.



Adding New Keyframes




New keyframes are automatically added to a spline whenever the value of the control that the spline is animating is changed. 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.

Points can also be added to a spline by clicking directly on the spline in the position where the new keyframe is to be added. A new keyframe can be added at the current time by pressing Shift-K on the keyboard.

Selecting Set Key from the context menu adds a new keyframe along the active spline wherever the position of the mouse pointer was when the context menu was called.

Adding Equal Keyframes



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 the Delete key on the keyboard. This will only remove the keyframes. It will not remove the spline, even if all of the keyframes are deleted.

Copying and Pasting Keyframes

Fusion allows the values from one spline key point to be copied to another. These values can even be copied from one spline to another.

Copying Keyframes

To copy a value or values, select the desired keyframes and press Ctrl-C or select Copy from the context menu. To quickly copy points from one segment of a spline to another, hold Ctrl down and drag the points to where they will be pasted.

Pasting Keyframes

Place the mouse cursor at the time where the points are to be pasted and select Paste from the context menu. The points will be pasted into the current spline.


Linear and Smooth Splines

Ease In / Out...	T
Flat	F
Smooth	Shift+S
Linear	Shift+L
Smooth Points	Shift+A
Step In	I
Step Out	O
Reverse	V
Isolate	
Duplicate	▶
Set Loop	▶
Set Preloop	▶
Gradient Extrapolation	

Although a keyframe defines the value of an animated control at a specified time exactly, there is some flexibility in the way the segment between two keyframes is treated.


Linear

A Linear segment will effectively take the shortest route, which is a straight line between two keyframes. To make the selected keyframe(s) linear, press the L key or click the Linear button in the toolbar.




Smooth

A smoothed segment will provide a gentle keyframe transition. To smooth the selected keyframe(s), press the S key or click the toolbar's Smooth button.




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 cause the value of the previous keyframe to be held, then jump straight to the value of the next keyframe. Step Out will cause the value of the selected keyframe to be held right up to the next keyframe.




Step In and Step Out modes can be set for selected keyframes by clicking on the toolbar buttons for each mode, or by right-clicking and choosing the appropriate option from the context menu. The keyboard shortcuts I and O can also be used to enable Step In and Step Out on selected keyframes.

Reverse



Both time and value of the keyframes will be reversed.

Bezier Handles



All Bezier spline keyframes have a pair of control handles to shape the spline as it passes through the key point. These handles can only be seen when the keyframe is selected. Click-dragging on a keyframe's handles will adjust the slope of the segments passing through the spline.

By default, the two control handles on a key point are locked together so that if one moves, the one on the other side moves with it. This maintains a constant tension through the keyframe. There are situations, however, when it is desirable to modify these control handles separately for a more pronounced curve or effect.

To temporarily break the lock between the control handles, select the key point to be modified. While holding down the Ctrl key, click on the control handle. They will now move independently of each other, as long as the Ctrl key is held down.

To make the Spline Editor treat all Bezier handles as independent, right-click in the Graph view and select Independent Handles from the Options context menu.

Showing Key Markers



A region immediately below the labels of the horizontal time axis has been set aside to show markers that indicate the position of each keyframe. The display of these markers can be enabled by right-clicking on the Graph view and selecting **Show > Key Markers** from the context menu, or by clicking on the button labeled **Show Key Markers** in the toolbar.

When the key markers are visible, they can be used to adjust the position of the keyframes in time, by click-dragging on the markers.


Filtering the Spline Editor

A complex composition can easily contain dozens, if not hundreds, of animation curves. As a composition grows, locating a specific spline can become more difficult. There are two ways to filter the splines shown in the spline editor: display selected tools only or create a filter to show only certain tools.

A filter allows the tools shown in the Spline Editor and Timeline to be limited to specific types and classes of tools. For example, a filter that shows only Creator tools and Loaders can be created, or one that only shows Color Corrector and Color Curves tools can be added.


Filters are applied simultaneously to the Spline Editor and Timeline.

Show Selected Only



Selecting the **Sel** button above the tree in the Spline Editor will enable the **Selected Tools Only** mode. When this mode is active, only splines for the tools that are currently selected in the flow will be visible.

Creating a Filter



To create a filter, open the preferences and select the **Global Timeline** panel. Click on the **New** button to create a new filter and name the new filter in the dialog box. In the tree menu, select by category or by individual tools using the checkboxes.

Close the preferences to save the new filter. Access this dialog quickly by choosing the Create/Edit Filters option from the Filter button in the spline toolbar. Also see the Timeline Filtering chapter for more information on this.

Applying a Filter

To apply a filter to the Spline Editor, click on the Filter button to the far right of the spline toolbar and choose the desired filter by name. The filter will be applied to both the Spline Editor and the Timeline.


Removing a Filter

To remove a filter and show all tools in the Spline Editor again, click on the Filter button again and choose Show All.

Time and Value Editors

The Time and Value buttons are each accompanied by two edit boxes that can be used to enter precise values for the selected keyframes. Clicking on the buttons will switch between three modes, each described below.

Time Buttons



Time

The text box shows the position of the current frame in time. Entering a new frame number into the edit box will move that key 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

By clicking once on the button to switch to Time Offset mode, the text box to the right can now be used to offset the frame where the selected keyframes exist. 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 on the button a second time will switch to Scale mode. The selected keyframes will be scaled by whatever value is entered here. For example, entering a scale of 2 will move a selected keyframe from frame 10 to 20.

Value Buttons



Value

The text box shows the value of the currently selected keyframes. Entering a new number into the edit box will change the value of the keyframe. If more than one keyframe is selected, the values of the keyframes are averaged. Entering a new value will cause all keyframes to adopt that value.

Value Offset

Clicking once on the button will enable the Offset mode. Entering a new number 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 a second time on the button will enable the 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.


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. Fusion offers a variety of ways to repeat a selected segment.


Set Loop



To repeat or loop a selected spline segment, select the keyframes to be looped. Select Set Loop from the context menu or click on the Set Loop button in the toolbar. The selected section of the spline will be repeated forward in time until the end of the global range, or until another keyframe ends the repeating segment.




Loop



Change the looped segment by modifying one or any of the key 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.

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.

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. Relative looping can be enabled on the selected segments from the context menu or the toolbar.

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 can only be enabled from the context menu. There is no toolbar button for this mode.


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 can only be accessed from the Repeat context menu, which reveals a submenu with all of the looping modes described above. Selecting any of these modes will open a dialog where the number of times the segment is to be copied can be entered.


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 context menu or use the V hotkey. The group of points is immediately mirrored in the Graph view. Points surrounding the reversed selection may also be affected.



Time Stretching




Time Stretching allows for proportionally stretched or squashed selected splines, or segments on that spline. To enable spline stretching, select this option from the mode submenu in the context menu or click on the Time Stretching button in the toolbar.


If more than one keyframe is selected on the spline, the outer keyframes will be surrounded by two vertical purple bars. If no keyframes are selected, click-drag a rectangle to set the boundaries of the bars.

To disable the Time Stretching mode, click on the toolbar button again or reselect Modes > Time Stretching from the context menu.

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 in time.




Shape Box




The Shape Box spline transformation mode is similar to the time stretcher, however, it adjusts the scaling of keyframe values as well as time.

To enable the Shape Box, select Mode > Shape Box from the context menu or click on the Shape Box button in the toolbar. A purple rectangle will be drawn around any points selected when the mode is enabled. If no points are selected, or to redraw the rectangle, click-drag a new region.




Scale the segments enclosed within the box by click-dragging on the edges of the box. Hold the Ctrl key down while moving a point to scale its complementary (opposite) point equally. Hold the Shift key down while dragging a corner of the shape box to adjust the corner without any constraints.

The Shift-B keyboard shortcut will enable and disable the Shape Box mode.



Ease In/Out



The Spline Ease dialog provides a more precise way to adjust the length of Bezier shape handles attached to selected keyframes. To access the spline ease dialog, select Edit > Ease In/Out from the spline context menu or press T on the keyboard.




The appearance of your toolbar changes to present you the Ease In / Out control.

Use the sliders to adjust the handles to your liking.



Activating the Ease Both switch will apply the same amount of ease to the In and Out handles.


Reducing Points



The Reduce Points dialog can be used to decrease the number of points in a spline. Fusion will attempt to maintain the overall shape of the spline, while eliminating redundant points from the path.

Select Reduce Points from the context menu to display the Reduce Points dialog. The dialog that appears contains a slider with a range from 0 to 100 that represents the accuracy of the spline.

When the value is 100, no points will be removed from the spline. Use smaller values to eliminate more points.



Guides


Guides are markers created to help identify important frames in a project. They may indicate a frame where a footfall occurs in the scene, the moment that someone passes through a portal in the image, or any other important event in the composite.

Guides are created in the small area between the time axis labels and the Graph view. They have the appearance of a small blue triangle and a blue line that extends from them vertically through the Graph view.

Working With Guides

Holding the mouse pointer over a guide will cause a tooltip to appear with the frame to which that guide is set, together with the name assigned to the guide (if any).

Creating Guides



To create a guide, click in the space immediately above the Graph view, below the time axis labels.


Deleting Guides

To delete a guide, click-drag it up beyond the axis labels and release the mouse. Alternately, right-click directly on the guide and select Delete Guide from the context menu.

Moving Guides


Click-drag a guide to reposition it along the time axis.

Jumping to Guides



Clicking on a guide in the guide list will change the current frame to the guide's frame.

Sharing Guides with the Timeline




Right-click in the guides area and select Edit Guides or press Ctrl-G to open a dialog box with the names of the guides in the composition. Clicking on the name of a guide will cause the current time to change to that guide's frame. Beside the names of the guides, there is a pair of checkboxes, one for the Spline Editor and one for the Timeline. By default, guides are created active in both portions of the Fusion interface. To prevent a guide from appearing in either view, de-select the appropriate checkbox.

Autosnap

To assist in precisely positioning spline keyframes and control points, the Spline Editor has an Autosnap function. Autosnap controls can be accessed through the Spline Editor's context menu. Autosnap only affects horizontal (time) snap. Control points are always free to move in vertical (value) range.

Right-click in the Graph view and select Options > Autosnap, then select one of the following options from the context menu.

Autosnap Context Menu



None

This allows free positioning of the keyframes.

Frame

Keyframes will snap to the nearest Frame.

Fields

Keyframes will snap to the nearest Field.

Autosnapping and Guides

By default, a newly-created guide will snap to the closest frame. Guides that are moved with the mouse will also snap to the current frame. This behavior can be changed by selecting Options > Autosnap Guides > None or by selecting Options > Autosnap Guides > Field from the context menu.

Keyframes that are moved close to a guide will snap to the guide, allowing for precise timing. This behavior is on by default and it can be disabled by selecting Options > Autosnap Points > Guides in the spline's context menu.

Importing and Exporting Splines


Spline shapes can now be imported and exported from ASCII text. 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 easier to identification.

A detailed description of the format used when exporting a spline can be found at the end of this manual in the ASCII Appendix.

Exporting a Spline

To export a spline, make sure it is active in the Spline Editor. Right-click directly on the spline to display the spline's context menu. Select Export from the context menu and then select from the three format options in the submenu. A file browser dialog will appear, allowing for the selection of a name and location for the output file.



Importing a Spline


It is also possible to import ASCII splines from disk to an animated control. To import a spline, first animate the control. Then, in the Spline Editor, right-click on the animation spline and select Import Spline from the context menu. Select the .spl file that contains the new spline curve from the file browser that appears. This will replace any existing animation on the spline.

Additional Spline Options

There are a few options in the context menus that have not been covered above. They are mostly options governing the operation or display of the Spline Editor.

Spline Context Menu

Scale



Scale to Fit

Scales the visible splines to fit them into the view.

Scale to Rectangle

Allows for a rectangle to be drawn in the Spline view to zoom the view to the extents of that rectangle.

Default Scale

Resets the scaling of the Spline view to the default values.

Zoom In / Zoom Out

Performing the same operation as the [+] and [-] keys on the numpad.


Auto Fit

Like the Scale to Fit option, but attempts to scale the view to fit all splines, even when new splines get selected or added.

Auto Scroll

Only attempts to fit the view in X direction.

Options



Fit Times

When enabled, the Graph view 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 view 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.

Independent Handles

Enabling this option causes all of the Bezier handles in a keyframe to be independent from their twin on the same keyframe, which is exactly the same as when the Ctrl key is held down while moving the handle.

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 sort of protection against accidentally moving the guides.

Show

Tips

This option enables and disables the display of Tooltips that can appear whenever the mouse hovers over a keyframe in the Graph view.

Tool Splines



The Spline view can also be used to display and modify any tool's splines, like in this case 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 not sufficient for some artists.

The Spline view offers an excellent working environment in this case.




Timeline Editor

Timeline Editor

The Timeline Editor is a powerful tool for adjusting the timing of elements in a project. This interface makes it easy to trim and extend clips, to adjust the timing of one spline, or to trim the valid region of a tool.

Timeline View


The Timeline View is the dominant part of the Timeline Interface. Bars represent the time range of each tool. The horizontal axis labels in this view represent time and are usually labeled with frame numbers. The timeline bars themselves are color-coded with the same color system as the tool tiles in the flow.



Tool Menu

The Tool Tree is a hierarchical view of the flow. It is displayed on the left side of the screen. Each tool is displayed by name in this Explorer-like interface tree.

Tools with effect masks or animation splines can be opened further, revealing a submenu for each spline animation, modifier and mask attached to that tool.



Current Time Indicator

The Current Time Indicator is a vertical yellow bar that runs through the Timeline view to indicate the position of the current frame or time.

Click-dragging this bar left or right can change the current frame of the project. As current time is changed, the frame number is displayed to the right of the toolbars to assist with precise positioning.

Hold down Ctrl+Alt and click in the Timeline view to jump the current time indicator to the mouse pointer position.

Timeline Toolbar

The Timeline Toolbar sits above the Timeline view, providing access to various timeline filters, scaling modes and the spreadsheet.



Spreadsheet

The Spreadsheet view offers an alternate means of visualizing time in a project. The spreadsheet display is much like a financial spreadsheet, showing the numeric values of the selected tools, including trimming and keyframes.


Key Frame :	25.0	30.0	40.0	50.0	60.0	70.0	106.0
Polygon4_Polyline	0	1	2	3	4	5	6

See The Spreadsheet Editor for details on displaying and using the Spreadsheet Editor.

Viewing, Scaling and Panning the Timeline

Viewing the Timeline

Click the Timeline Tab located at the top of the flow area to switch the work area view to the Timeline Editor. The F7 keyboard shortcut will also open the Timeline tab. Pressing the F4 key with the cursor over the Timeline view maximizes the view on the screen. Press again to return it to the previous size.



Panning/Scrolling the Timeline

To pan the Timeline view left or right, click on the middle mouse button and drag to the left or the right. Alternately, use the standard scrollbars at the bottom of the view.

Scaling the Timeline


To change the scale of the Timeline view, use one of the following methods:

Using the Toolbars

The scale buttons are located in the top left toolbar. The Fit button attempts to re-scale the view so that all currently-visible tools fit within the Timeline view. The buttons with the magnifying glass icons will scale the graph view in or out. To zoom into the selected region, use the Zoom To Rectangle button by click-dragging the rectangle.

Using the Axis Labels

Place the mouse pointer over the darker labels for the horizontal axis and click-drag to resize the Timeline view. The view is scaled with the position of the mouse pointer as center.



Using the Mouse

Place the mouse pointer over the Timeline view. Click and hold the middle mouse button. While the middle mouse button is held down, click once on the left mouse button to zoom in and once on the right button to zoom out.

Using the Keyboard

Several keyboard shortcuts can be used to scale the display of the Graph view while the Spline Editor is active. The + and - keys on the numeric keypad will scale in and out of the Timeline. Use Ctrl-R to enable scale to rectangle mode.

Working with Tools in the Timeline


Selecting Tools

To select a single tool in the Timeline, either click on the tool's name or on its bar in the Timeline. To add another tool to the selection, hold down the Ctrl key while selecting a tool. Holding Ctrl while selecting an already selected tool will remove it from the selection.

Note: Selecting a tool from the Timeline tree also selects the tool's tile in the flow, with its controls in the control window to the right.


Re-Ordering Tools in the Tree

To re-order the tools listed in the Timeline tree, click on the name of the tool and drag it to a new position in the list and release the mouse button.



Moving Tools

To change the position of a tool in time, click-drag on the middle of the tool's bar in the Timeline. The mouse pointer will resemble a bar with two arrows pointing in either direction, as pictured below.




Trimming Tools

To trim a tool in the Timeline, click-drag on either end of its bar in the Timeline. The mouse cursor will change to a vertical bar with a single arrow when the mouse pointer is in the right location to trim.

Trimming tools in the Timeline has a different effect depending on whether the tool is a source tool or an effect. Trimming a Loader tool will modify the clip time start and clip time end of the loaded sequence or file.

When any other tool in Fusion is trimmed, 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 tool to start at frame 80 in the Timeline will effectively prevent it from rendering and consuming memory or processor time until needed.

Extending the First/Last Frame

To extend the First or Last Frame of a Loader tool's clip, hold down Ctrl while dragging the ends of the bar in the Timeline. See the Loader tool reference for details on extending frames.

Showing Keyframe Values

Clicking on the Val button in the timeline ruler will reveal text boxes beneath each keypoint on a spline bar, showing the spline's current value. Click in the text box to edit the value.



Working with Splines in the Timeline

The Timeline also shows any splines that are attached to a tool. When the splines are collapsed, they are drawn over the tool's timeline bar. Clicking on the plus sign beside the tool's name in the menu will expand the display so that each spline has its own bar in the Timeline.

Keyframes for the splines are normally represented as vertical bars rather than points to make them easier to manipulate.

Changing Keyframe Time

To change the timing of a spline's keyframe, use one of the following methods:


With the Mouse

Click-drag on the keyframe to position it to a new time. Multiple keyframes can also be moved in this fashion by box-selecting or Ctrl-clicking on the desired keyframes. There is no restriction to moving keyframes from a single spline. Keyframes from several splines can be selected and moved together.

The exact position of the dragged keyframe in time is displayed to the right of the Spline toolbar when dragging, as well as in the status bar.

With the Toolbar

The Time Edit box in the Toolbar allows a new position in time to be entered for the selected keyframe or keyframes. Clicking on the Time button will switch it to Time Offset mode. Clicking again will set it to Time Scale mode. Both modes can be used to adjust the timing of multiple keyframes at once.



With the Spreadsheet

When the name of the tool is selected in the Tool menu, the keyframes will be shown in the Spreadsheet view. The timing of each keyframe can then be edited directly in the spreadsheet.

See the documentation for the spreadsheet later in this chapter for additional details.


Duplicating Spline Keyframes

Keyframes and multiple keyframe selections can be duplicated. Select the keyframe(s) to be duplicated and, while holding down the Ctrl key, drag the spline segment to a position without keyframes. The cursor will change from the forbidden symbol to a sub selection cursor. The keyframes can be dropped without over-writing existing keyframes.

Time Stretching

Time Stretching allows a selected spline, or segment on that spline, to be proportionally stretched or squashed. To enable spline stretching, select this option from the context menu or click on the Time Stretching button in the toolbar.

If more than one keyframe is selected on the spline, the outer keyframes will be surrounded by two vertical purple bars. If no keyframes are selected, click-drag a rectangle to set the boundaries of the bars. The time stretcher can affect more than one spline at a time.




Click-drag on the purple vertical bars to stretch or shrink the timing of the spline segments within the bars. Move these bars back and forth to stretch or squash the spline segment in time.

To disable the Time Stretching mode, click on the toolbar button again or re-select Time Stretching from the context menu.

Showing Keyframe Values

Clicking on the Val button in the timeline ruler will reveal text boxes beneath each keypoint on a spline bar, showing the spline's current value. Click in the text box to edit the value.




Timeline Filters

Flows with many tools can quickly become difficult to visualize in the Timeline. Timeline filters can now be applied to filter out the display of tools that are not necessary to the current operation. Use pre-made timeline filters, or create new ones as needed, in the Timeline preferences.

Applying Filters


To apply a Timeline filter, select the Filters button above the tree in the Timeline view, as shown below.



All of the available filters will be shown at the bottom of the menu that appears. Select the desired filter from the list to apply it to the Timeline.

Creating Filters

Open the preferences dialog (File > Preferences) and select Timeline in the options tree on the left. This will display the dialog shown below.



Click on the New button to create a new filter. Provide a name for the filter in the dialog box that appears. New filters will default to having all tools enabled for display in the Timeline.

Configuring Filters


Use the checkboxes in the tool tree to enable or disable the display of specific tool types and tool categories in a filter.

The Invert Categories and All Categories On/Off buttons can apply global changes to all of the checkboxes, toggling the selected states as described.

When the filter is configured, select the Save button to close the preferences dialog.

Removing Filters

To disable the currently applied filter, select the Show All filter from the Tools button in the Timeline. To permanently remove a filter, use the Delete button in the Timeline Preferences dialog when the filter is selected.



Timeline.def


Timeline filters that are created are stored on disk in the file `timeline.def`, found in the directory where `Fusion.exe` is located. This file can be copied from one system to another to maintain settings on a new installation.

Timeline Options

Collapse/Open All

The Collapse/Open All button is located on the far left of the Timeline toolbar, in the second row. It will be labeled with a + or a - symbol, depending on the current state of the Tool menu.


Clicking this button when the label is a minus (-) sign will collapse the Tool menu, closing all open branches. When the label is a plus (+) symbol, all branches of the tree menu will be opened.



Current Time Mode

The Current Time button, labeled CT, is located at the far left of the Timeline toolbar in the top row. When this button is enabled, it filters the timeline display so that only the tools that are valid at the current time are displayed.

This mode is extremely useful when working on a complex composite with many elements that appear and disappear.



Display Point Values

The Display Point Values option causes all keyframes in the flow to be displayed as points instead of vertical bars. Select Options > Display Point Values from the context menu to enable or disable this option.


Line Size

The Line Size option controls the height of each timeline bar 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 bar, right-click on the bar in the Timeline view and select one of the five options for line size in the submenu that appears. The default line size is small.


All Lines Size

This submenu provides the same choices as the Line Size menu but it will affect all bars currently visible in the Timeline.



Sort Order

The Tool tree has several sort modes to determine the order in which the tools are displayed in the Tool menu. These sort modes can be toggled by clicking on the button between the CT (current time) and the Fit button at the top left of the tool menu.



All Tools

The default sort order is All Tools. This will display the tools in linear order, scanning the flow from left to right and top to bottom.

Hierarchy

Hierarchy sorting will display tools by starting with the most background element through to the most foreground, 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.

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.


Guides

Guides are markers to help identify important frames in a project. They may indicate the frame where a footfall occurs in the scene, the moment that someone passes through a portal in the image, or any other important event in a composite.

Guides are created in the small area between the time axis labels and the Timeline view. They have the appearance of a small blue triangle and a blue line that extends from them vertically through the graph view.

Creating Guides

To create a guide, click in the space immediately above the Timeline view and below the time axis labels.



Working With Guides

Holding the mouse pointer still over a guide will bring up a tooltip with the frame to which that guide is set, together with the name assigned to the guide (if any).

Deleting Guides

To delete a guide, click-drag it up beyond the axis labels and release the mouse. Alternately, right-click directly on the guide and select Delete Guide from the context menu.

Moving Guides

Click-drag a guide to reposition it along the time axis.

Show Guide List (Shift-G)

Right-clicking in the guides area and selecting Show Guide List will display a small dialog box that lists the current guides in the composition. Double-click on a guide's name to set the current time to the guide's time, and enable/disable the display of the guide in the Spline and Timeline Editors.

Renaming Guides


By default, a guide is named after the time (frame) where it is set. When using guides to navigate, it is often useful to give the guides more descriptive names. To rename a guide, right-click on the guide and select Rename Guide from the context menu.

Sharing Guides with the Spline Editor

Right-click in the guides area and select Show Guide List or press Ctrl-G. A dialog box will appear with the names of the guides in the composition. Clicking on the name of a guide will cause the current time to change to that guide's frame.

There is 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 created as visible in both portions of the Fusion interface. To prevent a guide from appearing in either view, deselect the appropriate checkbox.



Autosnap

To assist in precisely positioning spline keyframes and control points, options to Autosnap to a field, frame or guides can be accessed through the Timeline's context menu. Autosnap only affects horizontal (time) positioning.

Configuring Autosnap

To configure autosnapping, select one of the following options from the Options>Autosnap Points menu in the context menu.

None

None allows free positioning of keyframes and tool bar edges.

Frame

Frame forces keyframes and tool edges to snap to the nearest frame.

Field

Field forces keyframes and tool edges to snap to the nearest field.


Guides

When enabled, the keyframes and tool bar edges snap to guides in the timeline.


Autosnapping and Guides

By default, a newly-created guide will snap to the closest frame. Guides that are moved with the mouse will also snap to the current frame.

This behavior can be changed by selecting Options > Autosnap Guides > None or Options > Autosnap Guides > Field from the context menu.



One of the many advantages to guides is keyframes that are moved close to a guide will snap to the guide, allowing for precise timing. This behavior is on by default and can be disabled by selecting Options > Autosnap Points > Guides in the spline context menu.



None

Guides can be placed anywhere.

Frame


Frame forces all guides to snap to the beginning of the nearest frame.

Field

Field forces all guides to snap to the beginning of the nearest field.

The Spreadsheet Editor

The Spreadsheet Editor is located beneath the Timeline Editor. To reveal it, click on the Show > Hide Spreadsheet button in the toolbar. The spreadsheet will appear at the bottom of the Timeline interface.



To resize the display of the spreadsheet, place the mouse pointer over the border between the spreadsheet and the rest of the Timeline. The pointer icon will change to a two-headed arrow. Click-drag the border to resize the available space reserved for the spreadsheet.

Selecting a Tool to Edit

To display a tool's timing in the spreadsheet, select the tool in the Timeline Editor. The in and out points of the selected tool will appear in the keyframe's line of the spreadsheet. Clicking on the cell that contains the value will make it possible to edit that value. This has the same effect as trimming the tool in the Timeline.

Selecting Multiple Tools to Edit

Multiple splines and tools can be edited together in the spreadsheet. Select each tool in the Timeline Editor by Ctrl-clicking on the tool's name in the tree view.

Selecting Animation to Edit

To edit an animation applied to a tool in the Spreadsheet Editor, select the animation spline in the Timeline. This will display each of the keyframes on that animation spline in the spreadsheet. The values can now be edited by clicking in the cell and typing a new value for the keyframe.

Inserting Keyframes

New keyframes can be added by clicking in an empty keyframe cell and entering the desired time for the new keyframe. The new keyframe will be inserted into the spreadsheet. A new value for the animation can now be entered.

Moving Between Cells of the Spreadsheet

Select the Tab or Shift Tab keys to move between cells in the Spreadsheet Editor.

Timeline Tool Color Coding

The same color model used by tiles in the flow is also used in the Timeline Editor. A single exception is that a saver will show red instead of pink if it fails to save frames for any reason.


See the Flow Editor's tool tile color coding and icon section for details on what colors represent what categories.

Tool Context Menu

Most of the options found in the Tool Context Menu can also be found in the Timeline.

To open the Tool Context menu, right-click on the Name label for the tool in the tree menu or right-click on the tool's vertical bar in the Timeline view.

For complete documentation of the options in the tool context menu, see the Flow Editor's tool context menu section.







Script Debug

12

Script Debug

The Script Debugger can be added to any view by right-clicking on the tabs and selecting 'Add View.'



The debugging features includes common functions like breakpoint, stepping through code line by line as well as a Watch area and CatchScript functionality.

The Buttons and Shortcuts for debugging are:

-  Ctrl - O Open Script
-  Ctrl - S Save Script
-  Catch Scripts

Catch Scripts automatically loads, and runs under debug, scripts that are executed from the UI by means of the Scripts menus. While normally only scripts executed from within the debug window will be run under debug, that doesn't always provide the needed context for how the script may be run.

Turn on catch scripts, then load, for example, AddTool.eyonscript into the debugger. This could be done manually, or by running the script by pressing ctrl-space. In both cases it should auto-load into the debug window.

Set a breakpoint somewhere in the script. The next time you run that script, it should break into debug.

Comp: Lua Language State Selector

This allows switching between the Global and Comp-specific environments of Lua and Python. Observe that the Global environments will not generate any output in the Console.

-  F5 Execute Script. Use Ctrl-F5 to run it without Debug Mode.
-  Ctrl-Break Break
-  F11 Step Into (single step, stepping into function calls)
-  Shift-F11 Step Out (run until return from current function)
-  F10 Step Over (single step, stepping over function calls)
-  F9 Toggle breakpoint on current line

Breakpoints

Breakpoints pause the execution of the code in the buffer before executing the code.

You can toggle the breakpoint on the line of the text cursor by pressing F9.


With F10 only one line of the buffer is executed and then the execution is paused until you either run the rest of the script or make another step.

F10 steps over function calls. This means that if a function is triggered it will be completed and the execution is stepped to the next line.

F11 Steps-Into function calls. The stepping is resumed inside the function.

You can leave the function call into the outer function by pressing Shift-F11 to Step-Out.

Watch Area



This view allows you to monitor variables of the buffer by showing at any point what values they store in memory.

To watch a variable simply write the variable name in the Watch window. On the next debug execution the variable value is displayed beneath the variable in the Watch window. This is only true as long as the variable is valid. After the script has finished execution you won't see the variable value anymore.

To watch multiple variables simply use one line per variable.

Debugging Tutorial

In this tutorial the debugging functions will be explained in practice.

Copy and Paste the following Lua code in the Script Area

```

function SafeSum(a, b)
    return a + b
end
j = 10
for i = 1, 20, 2 do
    print ( SafeSum(i, j) )
end

```

Now simply run the code by hitting F5.

The Log window will show a similar output to this:

```

-----run script-----
[string "Main"]:2: attempt to perform arithmetic on local 'b' (a nil value)

```

This tells us that there is a problem on line 2 of the code.

Watching Variables

Open the ScriptWatch by pressing the Watch button. In the Watch window write the following lines:

```
a  
b
```

Now go to line number two of the ScriptEditor and press F9 to create a breakpoint. A red dot next to the line number marks the breakpoint. Press F5 to run the script.

A yellow arrow indicates the position where the execution stopped. This should be exactly on the red breakpoint on line 2. Note that the ScriptEditor is greyed out. You cannot alter code while it is being executed.

In the ScriptWatch the values of a and b should be visible.

```
a  
  1  
b  
nil
```

Stepping Over and Breakpoints

Note that the execution on a breakpoint is paused before the line was executed. So the next action on line 2 will be:

```
return a+b
```

Obviously Lua does not know how to compute the result of 1 and nil. Press F10 to step over that line and see that this line generates the error message that was displayed earlier in the log window. Since an exception like this stops the execution of the script, the yellow arrow disappears while the code becomes editable again.

That was pretty fast. Let's go over the whole script line by line and see where the problem starts. First, disable the breakpoint by moving the text cursor to line 2 and hitting F9 again. The red dot disappears.

Now press F10 to start the execution of the whole script step by step. You will notice by the arrow that the function SafeSum is evaluated first, which means that it is stored in memory but not yet executed. Press F10 again until the arrow is at line 5.

```
k = nil
```

To watch the variable k simply add k as a new line to the script watch. Now press F10 to step over to line 6, the for loop. Also add the i variable as a new line into the ScriptWatch.

Hit F10 again. The arrow pauses at line 7, the print statement. In the ScriptWatch k is nil while i is evaluated as 1. In order to Step into the SafeSum function press F11 this time. The arrow jumps to line 2 inside the SafeSub function.

As we can see in the ScriptWatch a is 1 and all the other variables are nil. Why is i also nil? The ScriptWatch only shows the variables in the current scope. This means that it only shows the values of the variables in the current part of the execution.

Inside the SafeSum function there is no access to the i or k variables, just as inside the for loop there is no access to the a and b variables of the SafeSum function's scope.

Fixing the Bug

As we know the next execution will trigger the exception. But now we know where to look for the error. The k variable is nil and is passed to the SafeSum.

Let's make sure the Sum function checks for this first; after all, it is called Safe Sum. Change the function accordingly:

```
function SafeSum(a, b)
  if a == nil and b == nil then return nil end
  if a == nil then return b end
  if b == nil then return a end
  return a + b
end
```

This will check first:

- > If either of the variables is nil.
- > If both are it will return nil.
- > If a is nil it will return b.
- > If b is nil it will return a.
- > If neither a or b is nil it will return the sum.

Let's go over the script once again. Go to the print statement at line 10. Hit F9 to mark it as breakpoint. Hit F5 to execute the buffer.

It will pause at the breakpoint. Hit F11 to Step into the SafeSum function.

Hit F10 a few times and you will see that the last if statement triggers the return of a, since b is nil.

The execution will jump to the for loop. Use F10 to run through the loop by stepping over the SafeSum function.


Now that we are happy we can either toggle the breakpoint of or use Shift-F5 to resume the execution without the debugging. This means that the script will be executed and breakpoint will be ignored.



Tool Controls

13

Tool Controls



Every tool added to a flow in Fusion has a set of parameters, or controls, which determines how that tool behaves. When that tool is selected, the controls for that tool appear in a window on the right side of the screen in the Controls Area. More than one tool can be selected at a time.

Each control window has a header. This is a small title bar for the window that displays the tool's name and a button to collapse or open the window. The control header is also a good place to access the tool's context menu.

Modifiers can also be attached to a tool's controls. Modifiers are tools like expressions, calculations, trackers and paths. When a modifier is connected to a control on a selected tool, its controls will also appear. These controls appear separately in their own modifiers tab for the controls window.

The Controls Area Interface

Maximizing the Control Area



A small button labeled with a pair of arrows can be found at the top of the controls area to the left of the tabs. Click on this button to toggle the Control Area between maximized and regular modes.

In maximized mode, the tool controls are aligned with the top of the window, taking up the maximum available height for the display of controls.


In regular mode, the tool controls are aligned with the bottom of the display views, making the best possible use of horizontal space for the display views.

Viewing the Control Window for a Tool

► general

To open a tool's controls, select the tool from the Flow, Timeline tree or Spline Editor tree. If multiple tools are selected, the control windows for all selected tools will appear.


Hiding the Control Window for a Tool



The control window for a tool can be hidden from display so that the controls never appear, even if the tool is selected.

- > To hide a control window for a tool, right-click on that tool in the flow, or on the control header, and de-select Modes>Show Controls.
- > To re-enable the display of the control window for a tool, right-click on the tool's tile in the flow and re-select Modes>Show Controls.

Control Area Display Preferences



There are two different ways to display a tool's controls. The default option shows only the control windows for the selected tools. The alternate option shows the controls for all tools in the control window at once.

To switch to this mode, select File>Preferences from the menu and open the General Preferences. The following options are available.


Auto Control Close

When the Auto Control Close option is enabled, the control window for the selected tool is the only open window. Every other control window is automatically closed whenever a new tool is selected.

Auto Control Hide

When the Auto Control Hide option is enabled, the only control window visible is the one for the currently selected tools. All other control windows are hidden to reduce clutter. This is enabled by default. Turn it off to see all tool controls at once.

Selecting a Tool with the Control Header




A tool becomes the active tool when the tool's control header is clicked. Any onscreen controls for that tool will be displayed in the views that contain an image. This is exactly the same as clicking on a tool in the flow to select it. When a tool is the actively selected tool, its control header changes color.

Viewing and Previewing a Tool with the Control Header

► general

To view a tool in one of the display views, click-drag the tool's control header into one of the display views. To make a preview in one of the display views, hold the Alt key while dragging the tool header into the view.



Viewing a Tool's Splines with the Control Header

► general


To view all of the animated parameters associated with a tool, drag and drop the tool into the Spline Editor. The splines attached to the tool (if any) will automatically be selected and displayed for editing.

Types of Controls

Although a few tools use highly-customized interface elements, the vast majority of controls are modified using a mix of sliders, angle wheels, checkboxes and so on. The following is a description of the most common control types found in Fusion.

Slider Controls

Slider Controls are used to select a single value from a range of values. The artist can click and drag on the sliders handle to adjust the value, or enter a new value manually into the edit box to the sliders left.



Clicking on the gutter to the left or right of the handle will increase or decrease the value by a set amount. Holding down the Ctrl or Shift key while clicking on the gutter will adjust the values by smaller or larger increments respectively.

Clicking on the sliders handle directly also makes it possible to adjust the values using the left and right cursor keys. The Ctrl and Shift modifiers can again be used to modify the value change to larger or smaller increments.

While slider controls will usually present with a minimum and maximum value, entering a value directly in the edit box outside that range will often expand the range of the slider to accommodate the 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 indicator will appear above the bar. Clicking on this circle will reset the slider to its default.

Thumbwheel

A Thumbwheel control is identical to a slider except that this control does not have a maximum or minimum value. Thumbwheel controls are typically used on angle parameters, although they do have other uses. To use a Thumbwheel control, click on the center portion and drag left or right.



Fine-tune the parameter by clicking the arrowheads at either end of the control. The artist can also click directly on the control and then use the left and right cursor arrows to adjust the controls values. As with the slider control, the Ctrl and Shift modifiers can be used to increase or decrease the size of the control change

Values can also be entered directly into the edit box to the controls right.

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

Range Controls

Range Controls are used to specify a low value and a high value. To adjust the values, click and drag the handles on either end of the Range bar. To adjust the high and low values of the range simultaneously, click and drag from the center of the Range bar.


Expand or contract the range symmetrically by holding down the Ctrl key and dragging either end of the Range bar.



The Range control is actually two separate controls, a Low Range and a High Range. These can be animated and modified individually.

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. When this control is animated a value of 0 means the control is off while a value of 1.0 or greater will turn it on.



Drop-Down Lists

Drop-Down lists are used to select one option from a list of options. Clicking on the control will reveal the full list of available options for selection. Once the list is open, clicking on any of the entries in the list will cause that entry to be selected.

Only one option can be selected at a time, and the currently selected option will be displayed in the control when the menu is closed.



When a drop-down box control is animated the integer value 0 represents the first item in the selection, integer value 1 represents the second entry in the list, and so forth.

Button Arrays

Button Arrays are groups of buttons that allow the artist 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.



Color Wheel and Picker


The Color Picker is displayed wherever a tool parameter requires a color as its value. The selected color is shown in the 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 to visualize the transparency.



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

The Windows Color Picker

Clicking on the Pick button will display the operating systems standard Color Selector dialog. This dialog can then be used to select a color, and clicking on OK will apply that color to the control.



As the behavior of this control may change with the operating system used, please consult your operating systems documentation for additional details on the use of this dialog.


The Color Wheel

Clicking on the small arrow button above the Pick button will reveal a Color Wheel with a luminance bar, which can be used to select a color.



Picking Colors from an Image

To select a color directly from an image displayed in one of the Fusion display views, click and hold the Pick button. The cursor will change to the Eyedropper tool. While still holding down the mouse button, move the eyedropper over the view. A popup swatch will appear above the mouse pointer showing the selected color and its values.




When the color displayed is the one desired, release the mouse button to set that color.

The color picker normally selects from a single pixel in the image. Adjust the size of the selection by holding the Ctrl key and dragging while picking. The size change applies to all color pickers until the size is changed again.

Gradients

The Gradient control bar is used to create the colors that a gradient follows from start to finish, using keyframe colors set along the length of the bar. The gradient control is actually nine or ten controls in total, which are described below.



Gradients have their own context menu, accessed by right-clicking directly on the gradient bar. This provides options for animating, publishing and connecting one gradient to another. It also reveals one gradient-specific modifier that builds a custom gradient by sampling colors from an image produced by a tool in the flow.

Gradient Type

Use this button array to select the method used to draw the gradient. Select from linear, reflect, square, cross, radial and angle gradient types.

Start and End Position

The Start and End Position controls are represented by two crosshairs in the display views, as well as a set of X and Y controls. They determine where the gradient physically begins and ends.

Gradient Colors

This bar type gradient control is used to select the colors through which the gradient will transition when moving from the start position to the end position. Only two color positions are available by default, the start and end colors. To modify one of the colors, select the triangle below the color on the bar. Use the Color Selector controls below to set the color.

New keyframes are added to the gradient control by clicking just beneath the Gradient bar. The position of the keyframes can be adjusted by click-dragging on the keyframe. To copy an existing keyframe, hold the Ctrl key down before dragging the keyframe.

To remove a keyframe, click-drag the keyframe straight up past the gradient bar or select the keyframe and click on the red X button.

Gradient Interpolation Method

This button array determines what color space is used to calculate the colors between keyframes.

Offset

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

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.

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.


Context Menus

Tool Context Menu

To display the Tool Context Menu from the control window, right-click on the tool header across the top of the controls. This displays the same tool menu options that would be the result of right-clicking on a tool in the flow. See the chapter on the Flow Editor for more information on these options.

Control Context Menu

The context menu for individual controls can also be accessed. For example, right-clicking on a slider will provide the Control Context Menu, with options to animate the control or add additional modifiers. These options are described in more detail through the rest of this chapter.




Common Controls Tab

The Common Controls Tab is labeled with a radioactive symbol and can be found on virtually every tool found in Fusion. The common controls are generally found even on third-party plugin tools. This tab collects together controls that would normally have appeared on every tool.

The following controls are generally found in this tab. Some controls may not appear on all tools. For example, motion blur settings have no meaning on a Cineon Log tool.

Blend

The Blend control is found in all tools (except Loader and Creator tools). Use the Blend control 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.



Process When Blend is 0.0

The tool is processed even when the input value is zero.

Red/Green/Blue/Alpha Channel Selector

These four checkboxes can be used to limit the effect of the tool to specified color channels. This filter is often applied after the tool has been processed.

For example, if the red checkbox on a Transform tool is deselected, the transformation will first be applied to the image, then the red channel from the original input will be copied back over the red channel of the result.

There are some exceptions, such as tools where deselecting these channel selectors causes the tool to skip processing that channel entirely. Tools that do this will generally possess a set of identical checkboxes on another tab in the tool. In this case, the state of the common control channel selector checkboxes are instanced to the channel selectors found elsewhere in the tool.

Masking**Apply Mask Inverted**

Enabling the Apply Mask Inverted option inverts the complete Mask channel for the tool. The Mask channel is the combined result of all masks applied to a tool.

Multiply By Mask

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

Object and Material ID Masking**Use Object/Use Material (Checkboxes)**

Some 3D software can render to file formats that support additional channels. Notably, the RLA and RPF file formats support 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 present. The specific Material ID or Object ID affected is chosen using the next set of controls.

Pick (For Use Object/Use Material)

Use these controls to select which ID will be used to create a mask from the Object or Material channels of an image. Use the Pick button in the same way as the color picker to grab IDs from the image displayed in the view. The image or sequence must have been rendered from a 3D software package with those channels included.

Correct Edges

This checkbox only appears when the Use Object ID or Use Material checkboxes are selected. It toggles the method used to deal with overlapping edges of objects in a multi-object image. When this control is toggled on, the Coverage and Background Color channels are used to separate and improve the effect around the edge of the object. If this option is toggled off (or no Coverage or Background Color channels are available), aliasing may occur on the edge of the mask.

For more information on the Coverage and Background Color channels, consult the Auxiliary Channels chapter.

Motion Blur

Motion Blur

This toggles the rendering of Motion Blur on the tool. When this control is toggled on, the tool's predicted motion is used to produce the motion blur caused by the virtual camera's shutter. When the control is toggled off, no motion blur is created.

Quality

Quality determines the number of samples used to create the blur. A quality setting of 2 will cause Fusion to create two samples to either side of an object's actual motion. Larger values produce smoother results but increase the render time.

Shutter Angle

Shutter Angle controls the angle of the virtual shutter used to produce the Motion Blur effect. Larger angles create more blur but increase the render times. A value of 360 is the equivalent of having the shutter open for one whole frame exposure. Higher values are possible and can be used to create interesting effects.


Center Bias

Center Bias modifies the position of the center of the motion blur. This allows for the creation of Motion Trail effects.

Sample Spread

Adjusting this control modifies the weighting given to each sample. This affects the brightness of the samples

Comments Tab




The Comment Tab is marked by a small blue icon that contains the letter I in white. This tab is found on every tool. It contains a single text control that is used to add comments and notes to the tool. Click in the entry area and type the text.

When a note is added to a tool, a small blue icon with the letter I appears on the control header. To see the note in the flow, hold the cursor over the tool tile for a moment. 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 for details.

Scripting Tab

The Scripting Tab is present on every tool in Fusion. It contains several edit boxes used to add scripts that process when the tool is rendering. For more details on the contents of this tab, please consult the Scripting documentation.



Scaling and Panning the Control Window

Resizing the Control Window



To increase the overall width of the Control Window, click-drag on the vertical splitter that separates the controls from the work area and displays. Most controls will resize together with the window. This allows for greater precision with sliders, angles and other tools.

The controls have a minimum width of approximately 280 pixels. They will begin to be cropped rather than be resized past this point.

Panning the Control Window

The control window can be panned up or down to reveal portions that have scrolled out of the visible portion of the screen. To pan, either click drag on the scrollbar that appears to the right of the control window, or hold down the middle mouse button and drag while the mouse pointer is positioned over the control window.

Control Animation

Virtually any control found in Fusion can be animated over time. Controls default to static, with no animation. Animate a control by attaching it to an animation curve. Multiple controls can all be connected to the same animation curve.

Animation Curves

When animating a control, the control is connected to an Animation Curve. These curves are made up of splines from one of four spline types. By default, Fusion uses a Bezier spline when a control is animated. This is the type of spline found in most software packages. Change the default spline type from Bezier to one of the other types by changing the Default Animate option in the Globals>Defaults panel of the preferences.

Bezier Spline

The shape of a Bezier Spline is determined by keyframes that represent the value of a control at a frame. The values between keyframes are determined by interpolation. With a smooth the curve, two handles attached to each keyframe determine the weighting of the interpolation between values.

Natural Cubic Spline


A Cubic Spline also uses keyframes to represent the value of the control in time, and appears identical to a Bezier curve when the interpolation between points is linear. When the interpolation is smooth, a Cubic spline uses a different method of calculating interpolation that eliminates the need for handles.

NURB Spline


NURB splines also use keyframes, but the value of the keyframe only directly maps to the value of the control when the spline is linear. When the spline is smooth, the keyframes define the shape of the curve with no direct relationship to the value. Although it can be used to create very naturalistic motions, this spline type is rarely used because it is so difficult to control.

Also see the Spline Editor chapter of this manual.

To Animate a Control Using the Default Animation Curve




Right-click on the control and select **Animate** from the context menu. This will attach the control to an animation curve and set a keyframe for the current value at the current frame.




The Value fields get colored to indicate that they are animated. In addition to that, animation indicator bars appear to the left of each animated control. Blue indicates that the value is generally animated. Green means that a Keyframe exists at the current position in time.

To Animate a Control Using a Non-Default Animation Curve



Controls can also be connected to a plethora of modifiers or to other animated or published controls on other tools. Right-click on the control and select the desired curve or modifier type from the context menu's **Modify With** submenu.


To Set a Keyframe on an Animated Control



Keyframes are automatically updated whenever the value of a control attached to an animation curve is changed. If there is a keyframe already at the current time, the keyframe is updated, otherwise, a new keyframe is created.


To set a new keyframe without changing the value of the control, right-click on the control and select **Set Key On Tool Name Spline Name** from the context menu. This is handy when the animation is to hold a specific value for a length of several frames.

Removing Animation




To remove an animation spline from a control, right-click on the control and select **Remove Bezier Spline** or **Remove Cubic Spline** from the context menu.

Connecting a Control to an Existing Animation Curve



Right-click on the control and select the name of the animated control to be connected to from the context menu's Connect To submenu.

Connecting One Control to Another without Animating Them




It is often useful to connect two controls together without either control having to be animated. In this case, one control must be published first. This is done by right-clicking on the control and selecting Publish, then connecting the second control as described above.

For example, to connect a control to a Shake modifier, place the pointer over the control. Press the right mouse button and select Modify With from the context menu that appears. Then select the modifier that is to act on that control. The context menu will list only modifiers that are valid for that control.

For more information on connections and modifiers, see the Connections and Modifiers chapter later in this manual.

Using Simple Expressions and Pickwhipping



You could also use simple expressions by entering an equal sign into the value field, followed by Return. When using simple Expressions and Pickwhipping, values can be connected and combined visually without the need to publish a value first.




Preferences

Preferences

Most artists will want to customize Fusion to suit their own methodology. A wide variety of optional settings are available to configure Fusion's behavior. These settings are accessed via the Preferences dialog. The Preferences dialog can be opened from the File menu at the top of the interface. Select File > Preferences.

The dialog box is divided into a tree 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. Each new composition inherits a copy of the preferences. The Composition Preferences can be further modified without affecting the global prefs, or any other flow that is currently open.

Preference Settings



The Preferences tree's first entry is assigned to the global preferences. Opening the global preferences branch reveals the following sections.

3D View

Offers control over various parameters of the 3D Views.

Appearance

Controls the Appearance of the interface and assigns lookpacks.

Cluster

Used to configure the classes and options to Cluster render interactively with a group of machines.

Default

Used to select Default behavior for a variety of Fusion options, such as global range, timecode display, and automatic tool merging.

Flow

Used to set default options for display and handling of the Flow Editor in new compositions.

Frame Format

Used to select the default Frame Format when adding new creator tools, the frame rate for playback and the default pixel aspect ratio applied to footage.

General

Contains options for the General interface, such as auto save, default proxy settings and autosmoothing of splines.

Layout

Used to save the exact physical Layout of Fusion's interface elements.

Loader

Allows the artist to set options for depth and aspect ratio as well as define the local cache settings.

Memory

Used to control Memory management both interactively and at render time.

Network

Used to configure Network rendering options, such as selecting a render master, email notification and whether the machine can be used as a render slave.

OpenCL

Containing options for the use of OpenCL devices and tools.

Path Map

Used to configure virtual path names used by loaders and savers for footage and path preferences.

It also sets the default folders used by Fusion to locate comps, macros, scripts, tool settings, disk caches and more.

Preview

Set up for Preview creation and playback options.

Script

Configures passwords used to execute Scripts from the command line, and programs for use when editing scripts.

Spline Editor

Allows the artist to set various spline options for Autosnap behavior, handles, markers and more.

Splines

Contains options for handling and smoothing of animation splines, Tracker path defaults and more.

Timeline

Used to create and edit Timeline/Spline filters, and set default options for the Timeline Editor.

Tweaks

Miscellaneous settings for tweaking Fusion's behavior when loading frames over the network and queue/network rendering.

View

Used to manage display view settings and set default control colors and color picking ranges.

Bins

Contains two panels: a servers panel used to select which remote Bins servers to which Fusion will attempt to connect, and a settings panel for default settings related to the bins.


I/O

The I/O branch contains sub-entries for configuring display playback and default settings of AVI and QuickTime (QT 3 and up) format files. It is also used to configure video hardware like the dpsReality, etcz.

Import

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

3D View Preferences



Grid

Controls drawing of the Grid in 3D views with various parameter options.

Grid Antialiasing

Depending on Graphics drivers, the antialiased bits of the lines might sort incorrectly in zbuffer. To turn the grid off completely, use the 3D View > Right-click > Grid option, rather than preferences.

Size

Increasing this will increase the number of grid lines drawn

Scale

Overall scaling factor for the grid. This is useful, for example, if your geometry is so big that the grid is tiny or not visible.

Units

Since units are not defined in Fusion, currently a 'unit' is whatever the user wants it to be.

PerspectiveViews**NearPlane**

Default near clipping plane. Due to perspective matrix precision, this can't be set lower than 0.05.

Camera3D

Setting Near Plane too low and Far Plane too far results in loss of depth precision in the view. Setting Near Plane less than 1.0 results in an especially large loss of precision.

So if you have a really large scene and want to set the far plane over 1000, we recommend that you set the Near Plane = 1.0.

Eye Separation / Convergence / Stereo Mode

Apply when stereo is turned on in the 3D view

Fit to View**Fit Selection**

Controls what happens when you press F in the 3D view with one or more objects selected.

Fit All

Controls what happens when you press F in the 3D view with no objects selected.


Default Lights

Controls the default light in the 3D views.

The default light is used when lighting is turned on and there is no light in the viewed 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.

Cluster Preferences



Clustering allows the artist to use all available workstations, render nodes or parts thereof to render frames and send them to your local cache while you're still working interactively on your comp.

Based on the complexity of your comp this can provide major speed ups and faster playback. This can be particularly useful in client-attended sessions where every minute counts.

The Cluster Preferences configure settings and defaults for tool clustering. For details on Clustering, see the Network Rendering chapter of this manual.

Slave Groups to Render With

Available Groups

This list control shows all of the groups available for network rendered clustering. These classes define which machines are used for when a cluster render is sent to the render manager.

See the Network Rendering chapter for details.

Selected Groups

This edit box displays a list of the Selected Groups. Type a list of groups manually into this control, as well as selecting them from the list above.

Frames to Render

Entire Render Range/Side Frames


Select one of these two radio buttons to make clustering render either the entire Render Range from start to finish, or frames to either side of the current time.

Render X Frames Around the Current Time

If Side Frames is enabled, this control can be edited to select the number of frames around the Current Time that will be cluster-rendered.

Default Preferences

The choices made here are used to determine Fusion's behavior when new tools are added to the flow, and when controls are animated.



Default Animate

The Default Animate preferences can be used to change the type of modifier attached to a control when the Animate option is selected from its context menu. The default option is Nothing, which tells Fusion to animate controls using the internal defaults.

Fusion's internal default 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 to use as the default. For example, change the default type used to animate positions by setting the Point with control to XY Path.

Choices shown in this menu come from installed modifiers, which are valid for that type of control. These include third-party plug-in modifiers like those in the KroKroDove plug-ins, as well as native modifiers installed with Fusion.

Auto Tools

Defines which tools Fusion will add automatically for the most common operations, like Background tools, to provide a valid canvas for other tools and Merge operations to combine imagery.

Background

When set to none Fusion will use its internal defaults when adding Backgrounds, which is a standard Background tool. You can choose from a wide variety of available tools including 2D and 3D tools to tailor your pipeline to your needs.

Merge

When set to None, Fusion will use its internal defaults when connecting the outputs of two tools or dragging multiple Loaders on the flow area, which is 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, Fusion will not create any Merges when you drag multiple clips from explorer onto the flow area.


Global Range

These controls define the default Global Start and End frame used when creating new compositions.

Time Code

Use this option to determine if new compositions will default to showing SMPTE Time Code or frames to represent current time (Feet + Frames).

Flow Preferences



Force

Active/All/Source/Mask Tile Pictures

When these options are selected, Fusion will display pictures in certain tools tiles on the flow, rather than showing just a small and featureless grey tile. Active sets pictures for the actively selected tool, All enables pictures for all tiles, Source and Mask enable tile pictures for just Source and Mask tools.

The picture shown will either be a thumbnail of the image rendered by the tool, or the tool's default icon, if the Show Thumbnails option is off (below) or when the tool is a concatenated transform.

All Tool Names

Selecting this option (default) will make all of the tools in the flow display their names in the tool tiles.

Show Modes/Options

This 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

Arrange to Grid

This enables the flow's Snap to Grid option to force the tool layout to align with invisible markers.

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 in new compositions.

Auto Remove Routers

Routers or "elbow nodes" in the flow are considered to be 'orphaned' if the tools connected to either the input or output are deleted. Orphaned Routers are automatically deleted on loading the comp when this option is active.

Pipes Always Visible

When enabled, the connecting pipes between tools will be drawn over the top of the tool tiles.

Keep Tile Picture Aspect

Enabling this option forces tool tile thumbnail pictures to preserve the aspect of the original image in the thumbnail.

Full Tile Render Indicators

Enabling this checkbox 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.

Show Instance Links

This option can be used to select whether instance tools will show links between parent and instance as green lines between the tools.

Navigator

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 the method employed in previous versions of Fusion.

Build Direction

When auto-building or laying out a flow, Build Direction controls whether layers are organized horizontally or vertically.

Scale

This determines the default scale of the Flow Editor when a new flow is created.


Pipe Grab Distance

This determines how close the mouse cursor must be (in pixels) to the pipes in the flow when selecting them.

Group Opacity

This slider controls the opacity of expanded groups in the flow.

Frame Format Preferences



Frame Format Preferences select the format settings throughout Fusion's interface. The settings in this preferences dialog are described in greater detail in the Frame Formats chapter of the manual.

Default Format

Use this drop-down list 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 any resolution.

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

Width/Height

This specifies the Width or Height in pixels of the current format.

Frame Rate

Enter or view the frames per second played by this format. This sets the default Frame Rate for previews and the saver tool.

Has Fields?

Creator/Loader tools added to the flow will be in Fields process mode.

Film Size

This control 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

This controls the pixel Aspect Ratio of the chosen frame format.

Guide 1

Guide 1 displays custom Guides in the views and defines the visible portion of film frames. The four values display the current left, top, right and bottom guide positions respectively. 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 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

This button creates a new default setting in the drop-down menu. Type a name for the setting in the dialog box that appears.

Copy Button

This copies the current setting to create a new one for customization.

Delete Button

This button will remove the current setting from the default drop-down list.


Color Depth

Use these menus to select the Fusion color mode for processing images. 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- and 32-bit float per channel uses even more system resources and is best applied to film and HDR rendered images.

Independent settings exist for preview renders, interactive renders and full (final) renders. Generally, these options are ignored by the composition unless a Loader or Creator tool's Color Depth control is set to Default.

For a detailed explanation of how Color Depth is handled by Fusion, please consult the Frame Formats chapter in this manual.

General Preferences



These presets affect the behavior of Fusion's interface to suit personal preferences.

Usability

Auto Clip Browse

When this control is selected, Fusion automatically displays the file browser whenever a new loader or saver is added to the flow.

New Comp on Startup

When checked, Fusion will create a new, blank flow each time the program is started.

Summarize Load Errors

When loading comps that contain unknown tools (e.g. comps that have been created on other computers with plugins not installed on the current machine) Fusion will summarize missing tools in the console rather than presenting a dialogue for every missing tool.

Save Compressed Comps

This option enables the saving of compressed flows. 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. Note, though, that compressed comps are not ASCII based and thus cannot be edited with a text editor, unless saved again uncompressed.

Show Video I/O Splash

This toggles whether the Fusion Splash image will be displayed onto the video DDR board in your system.

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 down 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 zooms in the currently active window. With this option active it will zoom in the window underneath the pointer, 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

Determines the numeric scale Fusion uses 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.

Controls**Auto Control Close**

When selected, multiple tool control headers will be visible but only one tool control header will be opened at any given time. Otherwise, any number of tool controls may be opened at the same time. This setting has no affect unless the Auto Control Hide checkbox is cleared.

Auto Control Hide

When selected, only one of the tool controls for the currently selected tool will be made visible. Otherwise, all tool control headers will be visible.

Auto Control Advance

When selected, the Tab key and Enter key will cause the keyboard focus to advance to the next edit box within a tool's control. When cleared, 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, Fusion will only show tool controls for the active tool. By default, it is enabled showing controls for all selected tools.

Combined Color Wheel

When this option is enabled, Fusion's Color Corrector shows one Color Wheel for the master, shadow, midtones and highlight channels.

Grab Distance

This slider ranges from 1 to 10 and defaults to 5. It designates the active area around the pointer and should be modified only if there are difficulties in selecting points for modification in paths and spline curves. Smaller values represent more accurate selection.

Auto Save**Auto Save**

When enabled, Fusion will automatically save a comp to a backup file at regular intervals. If a backup file is detected when attempting to open a file, Fusion will offer the choice of loading either the backup or the original.

After a crash or power failure, select the backup. If the backup file is saved, it will overwrite the original file. If the backup file is closed without saving, Fusion will delete it without touching the original file.

When checked, an Auto Save is performed at intervals defined by the user.

Save Before Render

When selected, Fusion will automatically save the comp before it performs a preview or final render.

Delay

Use this control to set the interval between Auto Saves. Take care to set the interval in full time notation(i.e., 1000). Just typing 10 causes an Auto Save to occur every 10 seconds.

Proxy**Update All, Selective, No Update**

This sets the Update mode used by new compositions. The default mode is Selective, which 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.

When the images are being processed at a fraction of their original size, the display view areas will use the original image size for display purposes. In other words, even though smaller images are being used for processing, the image viewing scales 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 context menu.

Layout Preferences



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.

There are a lot of options but, in practice, simply organize the interface the way it is preferred on startup and when a new composition is created, then open this preferences panel and click on the three buttons.

Program Layout

Grab Program Layout

Pressing this button causes Fusion to store the program's overall current position and size.

Run Mode

This menu is used to select the program's default mode at startup.

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 views will be recorded when the Grab Program Layout button is used.

Document Layout

Recall Layout Saved In Composition

When checked, all layout settings in the controls below will be recalled when a saved composition is loaded.

Grab Document Layout

Pressing this button causes Fusion to store 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 causes Fusion to store the size and position of the selected Window.


Run Mode

Select the default run mode for the selected Window.

Use Grabbed Position and Size

When checked, the selected Window will be created using the stored position and size.

Loader Preferences



Defaults

Loader Depth

Defines how files are handled when adding a loader.

“Format” means that Fusion will automatically set the correct bit depth for the loader, depending on the file format and the information in the file’s header.

“Default” sets the bit depth to the value specified in the Frame Format preferences.

Cache

Disable All Local Caching

No files at all will be cached from the network. This overrides all other local cache options.

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 (and in the Path Maps Preferences) depending on the settings below.

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, to reduce network cluttering.

Cache Path Separator Character.

Fusion will re-build the path of the original files in LoaderCache. 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

Determines Fusion’s behavior if the original files can’t be found:

“Fail” behaves exactly as the default Loader in Fusion. The Loader will not process, which may cause the render to halt.

“Load Cache” loads the cache although no original file is present. Note: This may lead to out-of-date files. “Delete Cache” clears the cache, so that the original files will be loaded..

Cache Location

A copy of the LoaderCache path set in the Path Maps Preferences.


Explore

Opens the LoaderCache path in Explorer.

Clear All Cache Files

Deletes *all* cached files present in the LoaderCache path.

Memory Preferences



Occasionally, it will be necessary to adjust the Memory Preferences in order to make the best use of available memory on the computer. For example many artists prefer a higher cache memory for interactive work, because it allows for more frames to be held in memory for faster playback. For final renders the cache memory is often reduced, so that there's more memory available for simultaneous processing of tools or multiple frames at once being rendered.

Caching Limits

Limit Caching To

This slider is used to set the percentage of available memory Fusion will use for the interactive tool cache. Available memory refers to the amount of memory installed on the system, not the amount currently available.

When Fusion's interactive cache reaches the limit defined in this setting, it starts to destroy 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. Fusion will notify on the console for 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, Fusion will always make certain this amount of physical memory is 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 memory is running out frequently.

Final Render

These settings apply to Fusion's memory usage during a rendering session, either preview or final, with no effect during an interactive session.


Render Slider

Adjust this control to set the number of frames that Fusion will render at the same time.

Simultaneous Branching

When checked, more than one branch of a flow will be rendered at the same time. When running low on memory, turn this off to increase rendering performance.

Network Preferences



These Preferences are used to set up and control Fusion's network rendering. The majority of settings are found in the Render Manager dialog.

Submit Network Render Compositions

Enter the Master Name and IP address of the computer that will manage all network renders sent from this machine into the text control in this section.

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 artist 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 selected, this copy of Fusion will accept network render compositions from other machines and manage the render. It does not necessarily mean that this copy will be directly involved in the render but it will submit the job to the slaves that are listed in the render manager dialog.

Allows This Machine to Be Used as a Network Slave

When selected, this computer 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 force Fusion to ignore groups and priorities configured in the Render Manager. Compositions submitted from this system for network rendering will always be assigned to every available slave.

Email Notification

Notify Options

These checkboxes cause Fusion to send an email 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. 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.

OpenCL Preferences

Sets how Fusion handles OpenCL acceleration

Options

OpenCL tools

“Disable” - No OpenCL tools or plug-ins will be used. Fuses may still require OpenCL.

“Enable” - OpenCL will be available for tools and plug-ins, if appropriate drivers are installed.

“Auto” - OpenCL will be available in full Fusion, but will be disabled in Render Nodes.

Device

“GPU” - All OpenCL code will be rendered on the graphics card, if drivers are available.


“CPU” - OpenCL code will be rendered on the CPU, if drivers are available.

“NVIDIA” NVIDIA OpenCL drivers will be preferred.

“AMD” AMD/ATi OpenCL drivers will be preferred.

“Auto” - GPU rendering is preferred. If 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 AMD and NVIDIA options are intended to force the use of one vendor over the other.



Debugging

Verbose Console Messages

More information will be shown in the Console. For example, Startup Logs, Compiler Warnings and Messages.

Performance Stats in Metadata

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

Allows specific command-line arguments to be given to the OpenCL compiler.

Clear Cache Files

Deletes the contents of the AllDataOCLCache folder, where compiled OpenCL binaries are stored and cached.

Preview Preferences



Use these controls to set the default size of Previews and the type of preview created by Fusion.

Options

Render Previews Using Proxy Scaling

When checked, Fusion will scale the images down to the preview size for the Loader and Creator tools. This causes much faster rendering. If this control is de-selected, frames will be rendered at full size, then scaled down.

Skip Frames To Maintain Apparent Framerate

When checked, Fusion will skip frames during playback of flipbooks and file sequences to maintain the framerate setting.

Show Previews for Active Loaders

Determines if the preview playback controls on the bottom right of the screen are shown when a Loader with a valid file is activated.

Show Previews for Active Savers

Determines if the preview playback controls on the bottom right of the screen are shown when a Saver with a valid file is activated.

Display Files Sequences On

This designates which view is used for the interactive and file sequence playbacks as well as for the scrubbing function in the bins.

Type

AVI/Flipbook

Use this setting to select whether previews are created using RAM-based Flipbook, or using the Microsoft Video for Windows codec selected in the AVI Preferences. When the Flipbook option is selected, the rendered frames are spooled into RAM. No image compression is used. The color depth of the Flipbook images is defined by selection of the provided options.

In almost all cases, use Match Screen Depth to use the configured color depth of the video card. Select from the other options if there are problems displaying or playing the Flipbook preview. When the Flipbook option is selected, a conservative estimate of the number of frames that can comfortably fit within available memory is displayed beside each of the Preview Size options.

Use Image Cache

This forces all previews to be played from the interactive Image Cache in memory; custom Flipbooks are never created that displace existing caches.


Note: For the best Flipbook display speed, the video card should be capable of double buffering in its current resolution and color depth mode.

Path Maps Preferences

Path Maps are virtual paths used to replace segments of filepaths with variables. For example, define the path 'movie_x' as actually being in X\Shows\Movie_X. Using this example, Fusion would understand the path 'movie_x\scene_5\scan.000.cin' as actually being X\Shows\Movie_X\scene_5\scan.000.cin.

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

Also, path maps are used in Windows render farms to bypass the different filename conventions. For example, the path to the footage may be x\shots\footage on Window.



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 systems temporary folder. Fusion refers to the folder where Fusion is installed

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

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 mapped 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 above 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 considered when applying mapping to existing filenames.

Script Preferences



Login

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 login regardless of this setting. See the Online 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). SciTe is recommended for its ease of use and Fusion-specific coloring of the code.

It can be found in your [Fusion]\Utilities\SciTE folder.

Spline Editor Preferences



Spline Editor Options

Independent Handles

Handles and newly created keyframes will be 'broken' rather than 'connected' so the In and Out handle can be moved without affecting each other. You can revert this by pressing "S" on selected keyframes to smooth them out and reconnect the handles.

Follow Active

The splines are only displayed for the currently active tool(s), similar to clicking on the “sel” switch in the Spline Editor itself.

Show Key Markers

Small colored triangles will be displayed at the top of the Spline Editor timeruler 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**None / Field /Frame**

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**None / Field /Frame**

When moving or creating Guides these will snap to the Fields or Frames or can be moved freely.

Autoscale**None**

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**

Handles and newly created keyframes will be ‘broken’ rather than ‘connected’ so the In and Out handle can be moved without affecting each other. You can revert this by pressing “S” on selected keyframes to smooth them out and reconnect the handles.


Show key Markers

Small colored triangles will be displayed at the top of the Spline Editor timeruler to indicate points on LUT splines. The colors of the triangles match the colors of the splines.

Show Tips

Toggles if tooltips are displayed or not.

Splines Preferences



Autosmooth

Automatically smooths out any newly created points or keyframes on the splines selected in this dialogue. This is valid for both animation splines as well as masks.

Spline

Creates smooth points on Bezier spline modifiers.

B-Spline

Creates smooth points on B-spline modifiers.

Polylines

Creates smooth points on Bezier spline polylines.

B-Spline Polylines

Creates smooth points on B-spline polylines.

LUTs

Creates smooth points on LUT splines.

Paths

Creates smooth points on Motion Path.

Mesh

Creates smooth points on Meshes when adding a Grid tool.

B-Spline Modifier Degree**Quadratic**

Creates a B-spline modifier of second degree

Cubic

Creates a B-spline modifier of third degree

B-Spline Polyline Degree**Quadratic**

Creates a B-spline polyline of second degree

Cubic

Creates a B-spline polyline of third degree

Tracker Path Points Visibility**Hide**

After creating a Tracker, Tracker path points will be hidden in the view.

Show always

After creating a Tracker, Tracker path points will be displayed in the view.

Show on hover

After creating a Tracker, Tracker path points will only be displayed if the mouse cursor is hovering over a point.

Tracker Path**Bezier Polyline**

After a Tracker is created, the resulting Tracker path will be a Bezier polyline.

B-Spline Polyline

After a Tracker is created, the resulting Tracker path will be a B-spline polyline.

XY Splines

After a Tracker is created, the resulting Tracker path will be an XY path.

PolyLine Edit Mode on Done**Insert and modify**

After closing a polyline, the Polyline edit mode will be switched to 'Insert and Modify.' Existing points can be modified by clicking on them and new points can be created by clicking anywhere on the polyline where there are no points.

Modify only

After closing a polyline, the Polyline edit mode will be switched to 'Modify Only.' Only existing points can be modified by clicking on them. No new points will be created in this mode.

Allow Selection of Points From Multiple Polylines

Enables selecting and editing points from multiple polylines at the same time. When deactivated, only the currently active polyline can be modified.

Timeline Preferences

This preferences dialog is used to create and configure the filters used in the Timeline and Spline Editors.


Filter

Filter

This drop-down list shows all filters currently available. Select a Filter to enable it for editing in the Settings dialog below. The Show All filter is read only and cannot be edited.

Filter to Use

Use this menu to automatically apply a filter to new compositions.



Settings for Filters

Set/Reset All

Click this button to toggle on or off all checkboxes in the tool tree below. This button has no affect when the selected filter is Show All.

Invert All

This button Inverts the current checkbox selections in the tool tree menu. Selected checkboxes will be cleared and cleared checkboxes selected.

Tool List

The Tool List is a tree menu that displays every tool Fusion has loaded, organized according to category. A checkbox beside each tool indicates if that tool is visible in the Timeline when the current filter is applied.

New Button

Click on this button to create a New timeline filter.

Copy Button

Click on this button to make a Copy of the current timeline filter.

Delete Button

Click on this button to permanently Delete the current timeline filter.

Timeline Options**Current Time**

When selected new comps will default to showing only tools that are valid for the current time in the Timeline.

Selected Tools

When this checkbox is selected only the tools that are selected in the flow will appear in the Timeline.

Display Point Values

Enable this checkbox to cause the Timeline to show points instead of bars to represent keyframes on splines.

Animated Tools

When this checkbox is selected only the tools that are animated in the flow will appear in the Timeline.

Show Spreadsheet

Enable this checkbox to make the Spreadsheet display.

Autosnap Points/Autosnap Guides

These options control the precision of the Autosnap for guides and points.

Tools Display Mode

Use this menu to select a sorting order for tools listed in the Timeline tree.

Tweaks Preferences

This section contains more advanced settings that should be left alone unless specific problems are encountered.

Network**Maximum Missed Heartbeats**

Heartbeats are the signals or pings that the Render Master sends to the Render Slaves to check that they are still functioning. If no response is received, the Render Slave is either not available (power off, crashed, Fusion closed), or too busy to respond (usually because of excessive swap file usage).

This control determines how many heartbeats in a row a slave can miss before the Render Master assumes that the slave will no longer be available for additional rendering. If a slave becomes unavailable, it will be removed from the list for the current render and any frames assigned to it will be reassigned to the other machines in the Slave List.

Heartbeat Interval


The Heartbeat Interval refers to the amount of time in milliseconds that passes between each heartbeat sent by the Render Master.

Load Composition Timeout

The value of this control determines how many seconds the Render Master should wait before assuming that a Remote Slave failed to successfully load a flow. If the render log indicates that slaves are being dropped from a render because of a packet 3 timeout, try increasing this value on the Render Slave to give the flow more time to load.

Last Slave Restart Timeout

This control determines the amount of time a Render Master will wait after all of the slaves listed have become unavailable before aborting the render. This gives the slaves time to restart and rejoin the render in the case of a localized power outage or temporary network connection loss.



File I/O

Enable I/O C Cancelling

This checkbox enables I/O Cancelling, an experimental feature in Fusion. When a file-read operation is aborted, such as will occur when changing frames before an interactive render is complete, there may still be data queued up for transfer from disk. This occurs because Fusion reads ahead on disk, anticipating the data load for the next operation.

I/O Cancelling enables a feature of the operating system that allows these queued operations to be cancelled when the function that requested them is aborted. This can improve the responsiveness of the software, particularly when loading large images over a network.

Enabling this option will specifically affect the performance of Fusion while loading and accessing TIFF, VideoPump and other formats that perform a large amount of seeking.

This option should only be enabled with the understanding that it has not been completely tested with all hardware and OS configurations. Never enable it without thorough testing of drive loads from both local disks and network shares. If the system behaves normally, by all means enable the checkbox.

Enable Direct Reads

This checkbox enables Direct Reads of data from the drives into memory, using a method that is more efficient when loading a large chunk of contiguous data. The advantage is that the CPU load is reduced somewhat for I/O operations. The disadvantage is that this feature employs an undocumented ability of the operating system. As such, it may have other effects that are unknown.

Read Ahead Buffers

This slider determines the number of 64K buffers Fusion will use to read ahead in a file I/O operation. The more buffers, the more efficient Fusion is likely to be at loading frames from disk, 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 very large images, 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 with 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 Display View as float16 instead, which may improve playback performance.

Texture Depth

Defines in what depth images are uploaded to the Display View.

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 display view control used with Merge and Transform tools, and displays 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.

None

This 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 tweak can automatically disable the display view's Smooth Resize setting 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)

Running Fusion 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 loss or even crashes.

Use 10-10-10-2 framebuffer

If your graphics card and monitor support 30-bit color, this setting will render display views with 10 bits per primary accuracy, instead of 8 bits. Banding is greatly reduced when displaying 3D renders or images deeper than int8.

View Preferences**Saved View Settings**

This edit box lists the current saved settings for the views. Create a saved setting from the display view's context menu. For more details read the Display View chapters documentation.

Use the Rename and Delete buttons to manage the selected entries in the list.

Settings for View

Each display view has its own preferences. This drop-down list is used to select the display view affected by changes to all other settings in this tab.

Control Colors

Active/Inactive These determine the color of the Active/Inactive onscreen control.

Color Picking Area Size


Use these width/height controls to set the number of pixels sampled when using the color picker in the views. This setting also affects the Color Inspector subview.

Displayed Depth range

Sets the default for how Fusion expects Z-buffers created from other software.

Display LUT plugins

Lists the available Display LUTs and activates the selected one as default.



Bins/Security Preferences

Users

Users List

This dialog lists the servers that are currently in the connection list. Select one of the entries to edit its settings using the User and Password edit boxes below.

Add

Use this button to add a new user to the list. Clicking on this button adds an entry to the list. Set up the user name and password for that user.

User

This editable dialog 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.

Remove

Click on this button to remove the selected entry.

Permissions

The administrator can set up different permission types for users.

Read

This will allow the user to have read-only permissions from 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 Preferences



This preference is used to add remote Bin servers to the list of bins Fusion will display in the Bins dialog. Connect to servers by hostname or IP address.

Servers

This dialog lists the servers that are currently in the connection list. Select one of the entries to edit its settings using the User and Password edit boxes below.

Add

Use this button to add a new user to the list. Clicking on this button adds an entry to the list. Set up the user name and password for that user.

Remove

Click on this button to remove the selected entry.

Server

This editable dialog shows the name or IP address of the remote Bin server for the selected entry in the list.

User

This editable dialog shows the username for the selected Bin server item. If the username is unknown, try Guest with no password.

Password

Use this field to enter the password for the Bin user entered in the Users list.


Library

Specify another library but the default “Library” on the server.

Application

Defaults to “eyeonServer” but can be replaced to hook up to other compatible programs like customized inhouse server engines.

Bins/Settings Preferences



Stamp Quality

This percentage slider determines the compression ratio used for Stamp file creation. Higher values offer better quality but take up more space. The default is 75%.


Stamp Format

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

Options

When Open bins on startup is checked, the Bins will open automatically when Fusion is launched.

I/O AVI Preferences



The controls in this tab configure the default settings of Previews and Savers using Microsoft Video For Windows.

Settings For

Select either Saver or Preview from this drop-down list to display and modify AVI defaults. When Saver is selected, the settings are used as the defaults for the Format tab of the Saver tool when the Saver is set to save out AVIs. When Preview is selected, the settings are used to determine what AVI software codec is used to compress a preview AVI.

AVI Video Compression Options

Compressor

This menu displays the AVI codecs available from the 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 control 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 AVI.

Limit Data Rate To X KB/Sec

When checked, the Data Rates of the resulting AVI file will be limited to the amount specified. Not all codecs support this option.

Enter the Data Rate used to limit the AVI in kilobytes (KB) per second, if applicable. This control will have no affect if the Limit Data Rate To option is not selected.

Sample Image

Fusion will send a Sample Image to the compressor using the settings specified. The image displayed is a sample of the result.

The file size of the compressed image will be displayed as a percentage of the original in text above the logo.


Configure

Many AVI codecs offer built-in options that can be configured, in addition to the standard Video for Windows settings.

About

The About button displays a dialog box usually containing information about the selected AVI codec.

I/O QuickTime Preferences



The controls here are used to configure the default settings of Savers using Apple's QuickTime multimedia interface.

QuickTime Video Compression Options

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.

Sample Image

Fusion will send a Sample Image to the compressor using the settings specified. The image displayed is a sample of the result.

The file size of the compressed image will be displayed as a percentage of the original in text above the logo.

Configure Button

If the selected Codec has configurable options, clicking on this button will display those options for editing in their own dialog box.

If there are no options available in the compressor, this button is greyed out.

EDL Import Preferences

Use these options to determine how Fusion creates composition from imported EDL files.

Import Properties

Flow Format

This drop-down list selects the process mode 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, Fusion will use shot names stored in DPS EDL files to locate the footage.



Managing Fusion Preferences in Larger Environments

The point of this document is to outline a technique by which you could allow artists to have independent Fusion preferences for elements such as layout, flow building order, etc., while still maintaining a degree of control over critical elements, such as frame format, render master, script paths, and memory usage.

On Windows XP systems Fusion 7.0 will create and store its preferences in a sub-folder of the main installation called profiles\default.

On Vista and Windows 7 systems this may appear in C:\ProgramData\eyeon\Fusion\Profiles\Default, due to changes in the basic permissions model. Regardless of who is logged in, all users on the computer use the preferences stored in this folder.

Preferences for Each Artist

We will want to move these preferences to a network location that is accessible to all users instead.

For the purposes of this document, let's assume we have created a folder on V to store the preferences. For example:

v\eyeon\fusion\preference_library

To instruct Fusion to use this folder for preferences, we must create an environment variable with the name FUSION_PROFILE_DIR and set the value to the path to this folder.

```
FUSION_PROFILE_DIR=v\eyeon\fusion\preference_library
```

Note If there are spaces in the path name, do NOT wrap the path in quotes.

Now, whenever Fusion starts, this environment variable will instruct it to look for preferences in the v\eyeon\fusion\preference_library\default folder. If none exist, it will create a sub-folder called Default, and put a default copy of the preferences within. We want to take it further, and have it store a different set of preferences for each user. So we set another environment variable with the name FUSION_PROFILE

```
FUSION_PROFILE=jane
```

Now, Fusion will look for the preferences in the folder v\eyeon\fusion\preference_library\jane instead. The trick used frequently here is to set FUSION_PROFILE to the name of the logged in artist with a login script (or by launching Fusion with a batch file).

So now we have a unique set of preferences for each artist. How can we maintain central control over critical settings, such as Frame Format, Pixel Aspect or memory usage? For this we will need one more environment variable, called FUSION_MasterPrefs.

```
FUSION_MasterPrefs=v\eyeon\fusion\preference_library\master_locked.prefs
```

Example for a Preference File

In this example, we have set the value of FUSION_MasterPrefs as a fully qualified path to a file called master_locked.prefs. This file contains a partial preferences file, assembled by copying portions of a Fusion preference file. The following example shows the format for one possible scenario:

```
{
  Locked = true,
  Comp = {
    AutoSave = {
      Enabled = true,
      Delay = 300,
      OnRender = false,
    },
  },
  Global = {
    Tweaks = {
      DisableCancelIO = true,
      DisableDirectReads = true,
      ReadAheadBuffers = 4,
    },
  },
  Paths = {
    Map = {
      ["Scripts"] = "v:\eyeon\fusion\scriptsFusion:Scripts",
      ["LUTs"] = "v:\eyeon\fusion\LUTsFusion:LUTs",
      ["Defaults"] = "v:\eyeon\fusion\defaults",
      ["Fuses"] = "v:\eyeon\fusion\fusesFusion:Fuses",
      ["Macros"] = "v:\eyeon\fusion\macrosFusion:Macros",
      ["Settings"] = "v:\eyeon\fusion\settingsFusion:Settings",
    },
  },
}
```

This example is in eyeonscript's table format. It turns on the AutoSave preference, and sets the AutoSave interval to 300 seconds. It also disables certain IO tweak settings, and assigns a few pathmaps that allow Fusion to load Scripts, LUTs, Tool Defaults, Fuses, Macros and settings from a network location as well as the usual local folders.

The first entry in the table (on line 2) is "Locked=true." This tells Fusion the preferences in this file should be considered locked - unchangeable by the artist once Fusion has been opened.

Multiple Preferences

What if we have preferences we want to suggest, but which we want to allow the artist to change later if they so desire? We simply add a second path to the FUSION_MasterPrefs variable containing those suggestions. We separate the paths with a semi-colon (no spaces), as shown in the following example:

```
FUSION_MasterPrefs=v\eyeon\fusion\preference_library\master_locked.prefsv:\eyeon\fusion\preference_ library\
master_unlocked.prefs
```

The master_unlocked preferences are identical to the master_locked preferences in format, we just remove the "Locked=true" line. Note that these unlocked master preferences will only really apply when a new set of default preferences are created, say when a new artist starts Fusion for the first time. Once the artist-specific preference file is created, these will be overridden by the artist's personal choices. That lets an artist make their own choices about whether the flow snaps to grid, builds tools horizontally or vertically, etc.

About Comp and Global Preferences

Internally Fusion has two classes of preferences, called Comp and Global. Global preferences apply to Fusion generally, and would affect all compositions equally.

An example Global preference is the Render Manager setting in the Network panel. Regardless of what composition is currently open, the Render Manager is not going to change. Comp preferences are preferences that will be copied to new compositions when they are created. Once a composition is created, it carries its own copy of these preferences from then on.

In the locked example fragment above, the AutoSave preferences are found in the Comp table, and the Tweak and Path preferences are Global preferences.

To provide a clearer example, imagine we have used the locked Master Preference file to set the default width, height, pixel aspect and frame rate for new compositions to values appropriate for a 2K Anamorphic project. These Frame Format settings are all considered to be Comp preferences. The artist would be unable to change the values that appear in the preferences dialog's Global/Frame Format panel. When an artist creates a new composition, its default frame format would be set to 2K Anamorphic. Since each composition carries its own copy of the Frame Format preferences, it would thereafter be possible for the artist to change the frame format of that specific composition to whatever they desired.

This reflects the fact that many facilities will work on more than one project at a time, and they may not necessarily all share the same frame format requirements. Generally, you should only use the locked Master Preferences to set the Frame Format to a fixed value if all your compositions will always use the same settings.

So What about Frame Format Then?

What if you do want to make sure that all artists use the 'approved' frame format for a project? A common source of errors in many facilities occurs when a project requires a non-square pixel aspect in the renders (i.e., for an anamorphic render), but the artist accidentally sets the pixel aspect to square pixels instead. The best way to deal with this is to check the composition when it is submitted for rendering - if the pixel aspect does not match the intended value for the project, then make the artist confirm that the composition is correct before allowing the render to take place.

You could also have a script attached to the OnSave and OnSaveAs that checks the compositions pixel aspect against the projects expected setting, and warns or corrects the artist when they Save.

Since this approach involves a fair bit of scripting, we'll handle that in a future document.

Conclusion

The approach described above outlines a reasonable way to get some centralized management over the artists' preferences while still allowing each artist a measure of personal choice over factors such as layout. If you have questions or require additional assistance, please feel free to contact us using the contact information below.

Software Tech Support for Fusion 7 Studio:

US: support-usa@blackmagicdesign.com

+1 408 954 0500

Software Tech Support for Fusion 7:

<http://forum.blackmagicdesign.com/>




Tracking Explained

15


Tracking Explained

Steady Position




Original Motion

The Steady Position Mode can be used to stabilize footage in both X and/or Y to remove camera shake and other unwanted movement. This picture shows the original motion of the tracked pattern.




Steady Position 1

When you connect a center control 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. 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 Position mode will run in the opposite direction of the original motion.




Steady Position 2

To explain this further one could say that the resulting point is mirrored around the center between the first frame of the original motion and 0.5/0.5.




Unsteady Position




Original Motion

The Unsteady Position can be used to de-stabilize a foreground to match the motion of the tracked background. This picture shows the original motion of the tracked pattern.




Unsteady Position 1

When you connect a center control to the Unsteady Position of the Tracker, it will be placed at 0.5/0.5 (the center of the screen) 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 is basically an offset in the same direction as the original motion.




Unsteady Position 2

To explain this further, one could say that the resulting motion, once the offset to 0.5/0.5 is acquired, follows exactly the original motion.




Steady Angle



Original Motion

The Steady Angle Mode can be used to stabilize footage in both X and/or Y to remove camera shake and other unwanted movement. This picture shows the original motion of the tracked pattern.


To use Steady Angle and Steady Size you need at least two tracked patterns in your Tracker. With just one point you can only apply (Un)Steady Position.




Steady Angle

When you connect a control, like 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.




Unsteady Angle



Original Motion

The Unsteady Angle Mode can be used to destabilize footage in both X and/or Y to match the rotation of a foreground plate to the tracked background. This picture shows the original motion of the tracked pattern.


To use UnsteadyAngle and Unsteady Size you need at least two tracked patterns in your Tracker. With just one point you can only apply (Un)Steady Position.




Unsteady Angle

When you connect a control, like for example the Angle of a Transform, to the Unsteady 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 match the rotation of the original motion.

So if the actual angle at frame 10 is 15 degrees, the result of the Unsteady Angle will be 15 degrees as well.




Steady Size



Original Motion

The Steady Size Mode can be used to stabilize the size of the footage to remove camera shake and other unwanted movement. This picture shows the original motion of the tracked pattern.


To use Unsteady Angle and Unsteady Size you need at least two tracked patterns in your Tracker. With just one point you can only apply (Un)Steady Position.




Steady Size

When you connect a control, like 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. From there on the resulting Size of the Steady Size mode will 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 $1 - (1.15 - 1) = 0.85$.




Unsteady Size



Original Motion

The Unsteady Size Mode can be used to destabilize the footage to match the size of a foreground plate to the tracked background. This picture shows the original motion of the tracked pattern.


To use Unsteady Angle and Unsteady Size you need at least two tracked patterns in your Tracker. With just one point you can only apply (Un)Steady Position.




Unsteady Size

When you connect a control, like for example the Size of a Transform, to the Unsteady 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. From there on the resulting Size of the Steady Size mode will match the size changes of the original motion.

So if the actual size at frame 10 is 1.15, the result of the Unsteady Size will be 1.15 as well.




Offset Position




Original Motion

Other than the modes discussed so far which refer to the median value of all trackers in one Tracker node, the Offset Position is available for each single tracker in the Tracker node and refers to that single tracker only. This picture shows the original motion of the tracked pattern.




Offset Position 1

When you connect a center control to the Offset Position of the tracker, it will follow exactly the path of that tracker.



Offset Position 2

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.





Frame Formats

A1

Frame Formats

One of the first tasks that will most likely be done with Fusion is loading footage from the hard drives into Fusion. Depending on the type of work, this footage may be NTSC, PAL, HD, film or even multimedia.

Each of these formats has different properties, from pixel aspect ratios to field order, and resolution to color depth. Since Fusion can mix and match material of different formats together seamlessly in a single composite, it is important to note how Fusion configures and combines materials of different formats when loading and merging them together.

Setting the Frame Format


The first thing to do when starting on a new project is to set the preferences to match the intended format of the final output.

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

If the same frame format is used day after day, the global frame format preferences should match the most commonly used footage. For example, with a HD Altitude system where the majority of the footage will be high definition, it will make sense to set up the preferences to match the HD frame format usually used.

The preferences are described in greater detail in their own chapter in this manual. This chapter focuses solely on the behavior of the frame format preferences.

Setting Up the Default Frame Format for New Compositions




Open the Preferences dialog from the File menu. The preferences are organized into separate trees one for the globals and new comps, and one for showing the preferences for each composition that is loaded into Fusion at that moment.

Open the tree for globals by clicking on the + sign next to its label. Click on the Frame Format entry in the list to open its controls with a dialog like the one pictured below.

The options selected here determine the default frame format for each new composition. Changing frame format options in the Globals section does not affect currently open compositions. It affects only the default used when creating new compositions.

Changing the Frame Format in Existing Compositions




Each composition also presents its own copy of the preferences settings in the tree menu to the left. Because each comp maintains a copy of its own preferences, there can easily be one comp with a default PAL format open while also working on a second comp set to Film defaults.

To change the settings for the current comp, access the Frame Format controls listed under that comp's name in the tree menu.

For additional details on the settings in this preferences panel, see the Preferences documentation.

Fields

Although it is standard to say that the frame rate for NTSC is 30fps (or 29.97fps) and PAL is 25fps, this is actually misleading. Televisions display twice this number of frames per second. This is accomplished by storing two separate images in each frame of video. Each image contained in the frame is called a Field.



The two fields are combined into a single frame by alternating scanlines (rows of pixels).

When a television receives the frame and displays it, the television first shows every second scanline, leaving a blank line between each to avoid squashing the image. It then fills in the blank lines with the scanlines it missed in the first pass. In this manner, two images are stored in a single frame by interleaving the images into alternate scanlines. The television displays a new field every 1/60th of a second (NTSC) or every 1/50th of a second (PAL).

This technique dramatically increases the smoothness of motion but at the expense of image quality. Both NTSC and PAL employ fields in this manner, as do several of the HD standards. There are no commonly used film formats that use fields to increase apparent frame rate.

Although each frame of NTSC video is said to be 720 x 486 pixels, each frame actually contains two half-height fields, for an actual resolution of 720 x 243. PAL video frames are 720 x 576 pixels, with half-height fields at a resolution of 720 x 288. The ordering of these fields is different between NTSC and PAL.

When processing fielded video for effects and image processing, it is important that each field is treated as a separate frame. Because of the field order difference between NTSC and PAL, to separate each frame of video into its component fields, Fusion requires the field order for that footage. To do this, select the Field Order from the Process Mode control in the Loader.

If the first scanline at the topmost row of pixels in an image is the upper field and the second row of pixels is the lower field, it is easy to see how each video format alternates fields.

NTSC Fields

With NTSC Fields, the first field is stored in the lower, or second, scanline. This is occasionally referred to as Odd Field Order because the first field is stored on the odd numbered scanlines (rows 1, 3, 5, 7...).

PAL Fields

With PAL Fields, the first field is stored in the upper, or first, scanline. This is occasionally referred to as Even Field Order because the first field is stored on the even numbered scanlines (rows 0, 2, 4, 6...).

HD Fields

With HD Fields, the first field is stored in the upper, or first, scanline, as with PAL.

Fusion will automatically process a frame as two separate fields if the Loader's process mode is set to anything other than Full Frames. The second field is always shown as x.5 (as in frame 1, 1.5, 2, 2.5...).


Combining Fielded and Full Frame Images

Although the combining of fields and frames is handled by Fusion, it takes some experience to understand exactly what happens.

Accidentally processing fielded footage as if it were frames could seriously affect the quality of the results, particularly if warps or transformations are applied to that layer. Fields could be incorrectly blended together, causing ghosting and stuttering on playback.

Conversely, processing full frame footage as if it were fields does no serious harm but may lower the general quality of the results, mostly because there is less information for Fusion to work with when calculating blurs and other filter type effects.

Hold the mouse pointer over a tool's tile in the flow view to see if it is outputting fields or frames. The tooltip that appears will read Field x Frame Size if the tool is processing fields, otherwise it will only read Frame Size. This information will also appear immediately on the status bar.



When combining field-processed footage with full frame-processed footage in a merge, the merge will output the same type of images, frames or fields as the tool connected to the background input. In other words, when merging frames over fields, the result will be fielded, and when merging fields over frames, the result will be full frames.

For example, picture a flow with two loaders, connected to a merge. The loader that is connected to the background input of the Merge tool is set to process as NTSC fields. The foreground loader is set to process as full frames. The merge will output fields, and any tools connected to the output of the merge will process the footage as separate fields.

The Merge tool is particularly good at handling foreground and background images with varying sizes, aspects and field orders. Most other multi-input tools (such as Dissolve or Channel Boolean) should be given the same size and type of images.

Note: For this reason, the comp's overall Process Mode, whether the saved results contain fielded motion or not, is determined by default by the Process Mode setting of the bottom layer in the composition, the first background or loader that is merged over all other layers. This can also be overridden by the Saver's process mode setting, as described below.

The basic rule of thumb is simple. Always set loaders to the type of footage they are loading. It is also recommended to set the generated source tools (such as background or fast noise) to Full Frames to maximize the amount of information available to Fusion.

Setting Field Processing for the Flow

Since the mode of the tool in the background-most position determines the processing and field ordering when merging images together, one way to get the exact field order is to start every flow with a black Background tool that is set to the desired Process Mode, and with the color as R0, G0, B0, A0 (black with no alpha channel).

Another way is to set the Process Mode in the Saver tool. This defaults to Auto, which will use the bottom layer's mode. It can, however, be forced to Full Frames or NTSC or PAL Fields. This effectively overrides the default process mode and renders the entire flow with fielded motion or unfielded frames, as desired.

Note: This does not override how the images are loaded or created by other source tools, only how those images are processed. It will change how the animation is completed with field or frame accuracy, but not the size of the images that are processed. In all cases, footage will be re-interlaced, if necessary, and saved as frames on disk.

For example, a still image can be loaded as Full Frames, with title text (created with a full frame Text+ tool) merged over the top and set to scroll upward. Saving this using a Saver with Process Mode set to Auto will result in frame-accurate animation the frames on disk will not contain fielded motion. Setting the Saver's Process Mode to NTSC Fields will result in field-accurate animation; the text will be merged over the background twice as often.

The resulting frames will have the appropriate scanlines extracted and interlaced together and the frames saved to disk will contain smoother, fielded motion.

As another example, consider a non-animated title (created as a full frame) merged over fielded video footage (loaded as fields). By default, this would be saved with field motion, as that is the mode of the bottom layer (the footage). By forcing the Saver's Process Mode to Full Frames, only every second field is rendered. Every other field of the video footage would be dropped to avoid getting fielded motion in the results. The remaining field is interpolated up to full frame size, then saved.

Pixel Aspect

Pixel Aspect refers to the dimensions of an individual pixel. While computer screens and most film formats use square pixels, many video formats use rectangular pixels. When a computer is used to produce images for display on a device that does not use square pixels, it is important to account for the pixel aspect of the source material as well as the eventual display device.

Pixel aspect affects the way that images are displayed in the views, in addition to the way they are merged and transformed (especially rotation). If the pixel aspect of the footage is set to match its format, the images will be distorted.

Fusion represents pixel aspect as a ratio between width and height. A pixel with an aspect of 11 is perfectly square. A pixel with an aspect of 1:2 is twice as high as it is wide. Listed below are the pixel aspects for some of the more common formats.

- > NTSC (D1) - 0.9 1
- > NTSC (widescreen) - 1.2 1
- > PAL (D1) - 1 0.9375
- > PALPlus (widescreen) - 1 0.703125

How Pixel Aspect Affects the Display Views

When an image is displayed with a non-square pixel aspect in the display views, Fusion must transform the image slightly to compensate for the fact that the image is viewed on a computer monitor with square pixels. This ensures that the correct geometry is always viewed -- circular shapes are not squashed into ellipses, squares do not become rectangles.


A very fast method is used to apply this transformation, which prevents it from slowing down the display and playback. The side effect to using a faster transformation is that the image quality may suffer slightly and occasional steps, or aliasing, may appear on diagonal edges.

If there is this type of aliasing in the image, try overriding pixel aspect compensation in the view by clicking on the display view toolbar button labeled 11.



This will show the true square pixels without compensation. If the aliasing disappears, there will be nothing to worry about when the image is displayed in its final output format.


Alternatively, enable Smooth Resize (from the view context menu's Options submenu). This uses a filtered aspect transformation that eliminates the aliasing, but slows image display, sometimes quite noticeably.



Another way to detect apparent aliasing problems arising from aspect compensation is to view the image in a display that matches the image's pixel aspect. For example, if text for NTSC broadcast is created, viewing the image on the framebuffer of a supported device, such as Decklink cards, will help identify if there is a real problem.

Setting Pixel Aspect for Creator Tools

When a new Creator tool is added to the flow in the composition, the pixel aspect for the image created is determined by the default in the comp's frame format preferences. Set a new aspect for the Creator tool with the Pixel Aspect control in the tool's Image tab. Right-clicking on the label of this control will display a menu of default values for each format, or a value can be typed in manually.



Setting Pixel Aspect for Loader Tools

The handling of Pixel Aspect for loaded footage is decided by the Pixel Aspect button array under the Import tab of the Loader's tool controls. The following options are available.



From File

This is the default mode for new loaders. When the From File button is selected, Fusion first looks to see if the pixel aspect for the image is stored in the header of the image file on disk. Certain formats (notably TIFF) can contain this information. If the aspect is present, that value will be used. Otherwise, the default value specified by the frame format preferences for the flow is used.

Default

The pixel aspect defined in the flow's Frame Format preferences is used, even if pixel aspect information for the image is stored in the file header.

Custom

Clicking on this button will expose a pixel aspect control that can be used to manually set the aspect for this footage. Try right-clicking on this control to display a menu of default values for each format, or type a value in manually.


If the pixel aspect of a tool is uncertain, hold the mouse over the tool's tile in the flow and a tooltip will appear with details on the resolution, color depth and pixel aspect of the image last processed by that tool, as well as on the status bar.

Combining Images with Different Pixel Aspects

When combining footage with different pixel aspects with a Merge or Dissolve tool, the output of the tool will have the same pixel aspect as the background input. The foreground input will be correctly transformed to produce a correct image. Most other multi-input tools (such as Dissolve and Channel Boolean) ignore pixel aspect differences.

Color Depth and Processing

In Fusion, the Color Depth of an image refers to the number of bits used to represent each color channel. An 8-bit color channel will therefore use 8 bits of data for each pixel in the Red, Green, Blue, and Alpha channels, for a total of 32 bits per pixel. The higher the color depth, the more colors are available to represent an image and the more accurate the manipulation of those colors will be.



Another factor to consider when discussing color depth is whether the color is stored and processed as integer values or floating point values. Integer processing limits the color so that it can never be brighter than white or darker than black. Float processing allows out of range data to exist. File formats such as OpenEXR, Cineon, HDRI and some variations of TIFF, FLX and RPF can all store out of range color values.

Supported Color Depths

8-Bit (Integer)

Generally, 8-bit color processing is used for all video formats (although some equipment capable of producing video at 10-bit depths now exists). 8-bit is the lowest color depth and gradients produced at this depth may have visible steps that prevent them from appearing smooth. Also, significant gamma or color corrections may produce more visible banding.

16-Bit (Integer)

16-bit color depth doubles the amount of precision used for images, eliminating problems with stepping in gradients and improving the resolution of the color. 16-bit color comfortably includes all of the colors that can be represented in standard digital film processing. Processing 16-bit per channel color takes longer and uses more memory than 8-bit processing.

Using 16-bit processing with footage that is stored in formats that only contain 8-bit data (for example, from tga files, bmp files, AVI or QuickTime movies, etc.) will not reduce banding in the original data but will provide greater precision when processing the tool with additional effects. Additionally, many 10-bit and even 8-bit YUV video formats will see some benefit when loaded with 16-bits of RGBA rather than 8.

16-Bit (Float)

This color depth sacrifices a bit of the dynamic range from 16-bit integer to allow storage of color values greater than white and less than black. This mode contains more than enough dynamic range for most film purposes, yet requires significantly less memory and processing time than is required for 32-bit float images.

32-Bit (Float)

32-bit float is like 16-bit float, except with a much greater range of values for even greater precision.

Setting Color Depth

As with image resolution, field order and pixel aspect, each Loader or Creator tool can define the color depth used to process the image it produces. 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.


Setting the Interactive and Preview depth to 8 bits per channel can improve system performance while assembling a flow, allowing the work to be faster, while final renders still get full 16-bit color resolution.

We do not recommend that integer color depths are mixed for interactive vs. final renders, since the way that float handles out of range values may produce significantly different results to color depths using integer values.

When adding a creator tool such as background or text, the composition's frame format preferences are used to determine its color depth. Later, if the color depth in the composition's preferences is changed, the color depth of that tool will also change.

Override the default color depth used for source tools by setting the Color Depth control located under the tool's Common Controls tab. When any option other than Default is specified, that image will always be processed at the specified depth, regardless of the setting in the preferences.


Loaders handle color depth somewhat differently than Creator tools. The Color Depth control for Loaders is found under the Import tab, not under the Common Controls tab. There is an additional option in the button array, labeled Format. This is the default setting for new Loaders.



Fusion uses the color depth that makes the most sense for the file format stored on disk. For example, load a TGA file from disk and the color depth for the tool 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.

As with source tools, setting the Color Depth control to Default will force the Loader to process at whatever color depth is configured in the flow's preferences. The other color depths in the array can be selected to force that image to process at a specified color depth, regardless of the preference or format.

If the color depth at which a tool processes is uncertain, place the mouse over the tool's tile in the flow and a tooltip will appear, listing the color depth for that tool as well as on the status bar.



Combining Images with Different Color Depths

Fusion can combine images with different color depths together 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

Floating point images use significantly more RAM and take much longer to render than any other color depth. So what advantage does Floating Point Processing offer to offset the disadvantages?

There are two major advantages to floating point processing. The first is that floating point values do not lose accuracy through integer rounding. The second is the preservation of out of range values.

Detecting Out of Range Values

Although floating point processing preserves out of range values, they are still shown as black or white on the display views. This can make it difficult to determine the overall dynamic range of an image. To discover if there are out of range values in a viewed image, click on the Show Full Color Range button in the display view toolbar.




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 the Display Views chapter for details.

When to Use Floating Point

Out of range values are frequently used in film work, as film has a very high dynamic range and maintains color detail even in heavily over- or under-exposed frames. While it is possible to work with film frames using integer data, one must either clip all out of range values, losing all detail in the image highlights and shadows and reducing the range of possible color corrections, or apply some tricks to fit more of the color data within the limited integer range while still processing and viewing the image as it is intended to appear. Float processing makes working with film considerably easier by preserving highlight detail without the need for tricks.



Images with brighter than white color detail can also give very nice, natural results when blurred, glowed, color corrected, or even just when faded or dissolved.

There is also some value to using float color depths with NTSC, PAL or 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).

Accuracy

To demonstrate 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 255 (white), so our pixel is a bit dark.

Imagine that the gain of that pixel is halved with a Color Gain tool. Now the pixel's red value is half of 75, or 37.5. We cannot, however, store decimal or fractional values in 8 bits, so we must round that value to 37. Now we double the brightness of the pixel with another Color Gain tool. Now the pixel's R value is 37 multiplied by 2, or 74.

We lost a full value of precision due to integer rounding in this example, a problem that can result in visible banding over several color corrections. Similar problems arise when merging images together, or transforming them. The more operations that are applied to an image, the more color precision is lost to rounding. The problem is much less noticeable at 16 bits per channel, although only because the differences are much smaller.

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.

Out of Range Values

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

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

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.

Alternatively, clip the range by adding a Change Depth tool and switching to one of the integer color depths.



Network Rendering

A2

Network Rendering


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

Fusion can submit compositions to be rendered by other copies of Fusion, as well as to the the Fusion Render Node, which is a non-interactive license of Fusion available at a reduced cost.

TCP/IP and Fusion

Fusion relies on the presence of the TCP/IP network protocol to control network rendering across multiple machines. Refer to the Windows reference documentation or the network administrator for help to install TCP/IP on the network.

Ensuring that the Network Functions Correctly



Launch a command prompt for the render Master station and use the ping command to test the connection to each slave.

Let's assume the address you want to test is 192.168.0.10

At the command prompt type: ping 192.168.0.10

The response should look like this:

```
Pinging 192.168.0.10 with 32 bytes of data:
Reply from 192.168.0.10: bytes=32 time<10ms TTL=128
Reply from 192.168.0.10: bytes=32 time<10ms TTL=128
Reply from 192.168.0.10: bytes=32 time<10ms TTL=128
Reply from 192.168.0.10: bytes=32 time<10ms TTL=128
```

When there is a response from a slave, the network link to that machine is functioning correctly and Fusion should be able to use it to network render.

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. eyeon provides a script called RenderTool.eyeonscript 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.exe program 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 rendermanager.exe and eyeonScript.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 File>Preferences 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 dialog is described in detail later in this chapter.

Preparing Render Slaves


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

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 context 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 Ctrl-M 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. Alternately, right-click in the Slave List and select this option from the context menu.



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 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, or from the Slave List context menu.




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 Ctrl-Shift 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 drag-dropping the .comp file into the Render Manager's queuelist from the File Explorer.

From the Command Line

Also add compositions to the queue using the Renderslave.exe utility, which can be found in Fusion > Utils. 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 context 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 Queue list be rendered sooner. Click-drag an entry to a new position in the list to rearrange the order of the compositions within the queue. If a composition that is already complete (status is done) is 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 context 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 context 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 **Ctrl-Shift** 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 Queue list and select **Assign Group** from the context 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\netest1.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

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

Disk Cache

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

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 context 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 context 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 context 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 Overseer.exe, 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 context 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.exe. Nodes will not allow themselves to be updated with Fusion.exe.

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 context 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.

eyeonServer

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


eyeonServer 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 eyeonServer will relaunch the slave after a short delay, allowing the slave to rejoin the render.

eyeonServer 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 eyeonServer cannot detect the problem. Hung processes like this are detected and handled by frame timeouts, as described below.

Frame Timeout

Frame timeouts are a failsafe 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 context 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 Speedsix's Smoke and Rain, and eyeon's own Trails tool are the only ones that cannot render properly over the network, although there may also be AE Plugin adapter tools with this limitation. 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 Targa, Cineon, 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

Like AVI and QuickTime video, hardware boards, such as the Leitch Altitude, Matrox DigiSuite and NewTek Video Toaster, do not normally have the ability to directly accept rendered frames from a rendering network. As with AVI and QuickTime movies, these boards usually require frames to arrive in order. Always render to separate sequential file formats like Targa, Cineon, JPEG, and so on, and spool the images to the desired hardware format afterwards. Currently, the only widely used DDRs that can be used as a destination for network rendering is Leitch's line of products. Specifically, the dpsReality, Quattrus and Altitude boards can all accept network rendered files using the Virtual Tape File System. Consult the board's documentation for details on how this is done.

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 became full or because of any other reason for which Fusion might normally be unable to render a frame. In this case, the Render Manager will attempt to reassign that failed frame to a different slave. If no slave can render the frame, the render will fail. Try manually rendering that frame on a single machine and observe what happens.

Check the Slaves

Fusion's Render Manager incorporates a number of methods to ensure the reliability of network renders. Heartbeats are used to detect network or machine failures. In this event, failed slave's outstanding frames are reassigned to other slaves where possible. In rare cases, a slave may fail in a way that the heartbeat continues even though the slave is no longer processing. If a slave may have failed (although the Render Master may not have detected it) and you do not want to wait for the Frame Timeout, simply restart the Fusion workstation or Fusion Render Node that has hung. This triggers the heartbeat check, reassigns the frames on which that slave was working, and the render should continue. Heartbeats may fail if the system that is performing the render is making extremely heavy use of the Swap file or is spending an extraordinary amount of time waiting for images to become available over a badly lagged network. The solution is to provide the node with more RAM, adjust memory settings for that node or upgrade the network bandwidth.

Check the Network

At the Render Master, bring up the Network tab of the Preferences dialog box and click Scan. If a slave is not listed as running, the Render Master will not be able to contact it for network rendering. Alternatively, bring up a command prompt and ping the slaves manually. If the remote systems do not respond when they are up and running, the network is not functioning and should be examined further.