



VUE₂₀₁₆

VUE 2016 Infinite & xStream — Reference Manual

e-on software, SARL
✉ 68 avenue Parmentier
75011 PARIS – FRANCE
☎ Phone +33 1 4314 2815
Fax +33 1 4355 3671

Web : www.e-onsoftware.com
Info : www.e-onsoftware.com/info
Support : www.e-onsoftware.com/support
Feedback : www.e-onsoftware.com/feedback

Trademarks

VUE, VUE Pioneer, VUE Frontier, VUE Esprit, VUE Studio, VUE Complete, VUE Infinite, VUE xStream, VUE GeoDesign, PlantFactory, PlantFactory Exporter, PlantFactory Artist, PlantFactory Designer, PlantFactory Studio, PlantFactory Producer, The Plant Factory, LumenRT, LumenRT VIZ, LumenRT Architect, LumenRT Studio, LumenRT GeoDesign, Carbon Scatter, Ozone, CloudFactory, CloudFactory Ozone, Cornucopia3D, Mover, SolidGrowth, Dynamic Motion Reaction, Solid3D, RenderCow, SmartCow, NewCow, Hybrid 2.5D, AccuShadows, EasyGI, SmartGraph, EcoSystem, Cornucopia, Cornucopia3D, “The Natural 3D Studio”, “Natural 3D”, “The Art of Natural 3D”, “Solutions for Natural 3D Environments”, “Digital Nature”, “Immersive Nature” are trademarks or registered trademarks of e-on software, SARL.

Windows, Windows 95, Windows 98, Windows NT, Windows 2000, Windows XP, Windows Vista, Windows 7, Windows 8 and Windows 10 are registered trademarks of Microsoft, Inc.

QuickTime, Mac and OS X are registered trademarks of Apple, inc.

Flash and Photoshop are registered trademarks of Adobe, Inc. Maya.

Softimage and 3ds Max are registered trademarks of Autodesk, Inc.

LightWave is a registered trademark of NewTek, Inc.

Cinema 4D is a registered trademark of Maxon, Inc.

Poser is a registered trademark of Smith Micro Software.

RealMovie is a registered trademark of RealNetworks, Inc.

ArchiCAD is a registered trademark of Graphisoft, Inc.

MicroStation is a registered trademark of Bentley, Inc.

CityEngine is a registered trademark of ESRI, Inc.

SketchUp is a registered trademark of Trimble, Inc.

Revit is a registered trademark of Autodesk, Inc.

All other product and brand names mentioned in this manual are used for identification purposes only. They may be trademarks or registered trademarks, and, as such, remain the exclusive property of their respective holders.

Copyright

VUE Program, PlantFactory Program, LumenRT Program, Carbon Scatter Program, CloudFactory Program ©1997-2018 e-on software, SARL. All rights reserved.

VUE Documentation, PlantFactory Documentation, LumenRT Documentation, Carbon Scatter Documentation, CloudFactory Documentation ©1997-2018 e-on software, SARL. All rights reserved.

This manual, as well as the software described in it is furnished under a license agreement and may only be used or copied in accordance with the terms of such license agreement. Information in this document is subject to change without notice and does not represent product specification or commitment on the part of e-on software, SARL. Algorithms used inside this software were developed for visual performance only, and may not constitute a precise simulation of real phenomena.

Warning: This software and the accompanying documentation are protected by U.S. copyright law as well as by international intellectual property conventions. Any reproduction, sale, transfer, transcription, storage in a retrieval system, translation into any language in any form or distribution by any means whatsoever of this software or accompanying documentation, in part or in full, without the prior written permission from e-on software, SARL is strictly forbidden. Any such act shall constitute a copyright violation and shall be prosecuted to the fullest extent of the law.





Contents

1 Getting Started

Introduction	51
System Requirements	51
Installation	51
<i>Default Folders for All Versions</i>	52
<i>Installing Vue Infinite and xStream</i>	53
Installing VUE Standalone	53
For Floating Licenses	53
Installing VUE as a Plugin	53
Technical Support	54
<i>Included Support</i>	54
<i>Complimentary 30 Day Maintenance</i>	54
<i>Standard Maintenance</i>	55
<i>When You Can Get Under Maintenance?</i>	55
Activating Your Product	55
<i>Activating Your Product</i>	55
Welcome Dialog	56
<i>Feeling “At Home”</i>	59
Checking Video Board Compatibility	59
Updating VUE	60
<i>Automatic Updates</i>	60
<i>Canceling Updates</i>	61
<i>NewCow™ Network Updates</i>	61
Importing Files from Previous Versions	61
<i>Importing Locked Content</i>	62
Memory Management and Fault Protection	62
<i>Memory Management</i>	62
<i>Fault Protection</i>	63
<i>OpenGL Crash Interceptor</i>	63
<i>Embedded Error Reporting</i>	64
<i>Compatibility Mode</i>	64
Downloading the Extra Content	65



Interface Overview	66
Dialog Bar	67
Customizing the Interface	68
Selecting an Interface “Model”	68
Further Customization	68
<i>Docking</i>	68
<i>Undocking</i>	69
<i>Resizing the ViewPorts</i>	70
3D Views	70
Active View	70
Maximizing / Resizing Views	71
Full Screen Mode	71
ToolTips	71
View Display Options	72
Saving Default Viewport Configuration	74
Quick Render	74
Saving Pictures	75
Channels	75
Multi-Pass, Mask and G-Buffer	76
Additional ViewPorts	76
Adding a viewport	77
To display an additional viewport	77
Displaying a camera that is not active	77
Double Action Icons	78
Unfoldable Icons	78
Top Toolbar	79
Left Toolbar	82
Scene Information Bar	84
Object Properties Panel	85
Aspect Tab	86
Preview Color and Options	88
Numerics Tab	92
Animation Tab	94
Animation	94
Forward Dynamics (Linking and Tracking)	95



Camera Control Center	97
<i>Render Preview</i>	97
<i>Camera Controls</i>	98
Adjusting a Spotlight	100
Synchronized Cameras and Spotlights	100
World Browser.	100
<i>Status Bar</i>	101
<i>Objects Tab</i>	102
<i>Materials Tab</i>	113
Increasing Responsiveness in Large Scenes.	114
Material Types	115
Editing Multiple Materials	116
Material Preview	116
The Bottom Toolbar.	116
<i>Library Tab</i>	118
<i>Links Tab</i>	121
Animation Timeline	126
<i>The Main Timeline</i>	127
Auto-Keyframing	128
Moving around in the Timeline	129
Constraining the Current Time Slider	129
Rendering the Animation	129
<i>The Properties Timeline</i>	130
<i>The Animation Graph</i>	132
<i>The Animation Preview</i>	133
Visual Browsers	134
<i>Collections</i>	135
Loading other Files.	137
Search	137
Browser Options	138
Virtual Collections	140
<i>Downloading Collections</i>	141
<i>Scene Versions</i>	141



Grouped Dialogs	143
HiDPI	146
Saving Scene	146
Thumbnail Previews in Windows	147
<i>How to use it</i>	148
VUE's Installed Content	148
VUE's Sample Scenes	149
Understanding Vue	150
Operations	150
<i>Clipping Plane</i>	150
<i>Drag and Drop</i>	151
<i>Popup Menus</i>	151
3D Coordinates	152
<i>World Space</i>	152
<i>Object Space</i>	153
<i>Material Mapping Coordinates</i>	154
Ray Tracing	156
<i>Render Engine</i>	156
<i>Soft Shadows</i>	157
Ray-Traced Soft Shadows	158
Shadow Map	158
Area Lights	159
<i>Blurred Transparencies and Reflections</i>	159
<i>Reflection Maps</i>	161
<i>Depth of Field</i>	162
<i>Motion Blur</i>	163
<i>Hybrid 2.5D Blurring</i>	165
Distributed Ray-Tracing	165
Hybrid 2.5D	165
Limitation of Hybrid 2.5D	166
Enabling Hybrid 2.5D Blurring	167
Fast Hybrid 2.5D	167
Light Related Features	168
<i>Lens Flares</i>	168
<i>Glow</i>	169



<i>Volumetric Effects</i>	170
Volumetric and Spectral Atmospheres	170
Volumetric Lights	170
Godrays	171
Volumetric Materials.	171
Volumetric Clouds	172
Spectral Clouds.	172
<i>Hypertextures</i>	172
<i>Caustics</i>	172
Physically Accurate Caustics.	173
<i>Dispersion</i>	174
<i>Ambient Occlusion, Global Illumination and Radiosity</i>	175
<i>HDRI and Image Based Lighting</i>	177
<i>Illumination Baking</i>	178
<i>Sub-Surface Scattering</i>	179
Absorption.	180
Multiple Scattering.	180
<i>Displacement Mapping</i>	181
Advanced Effects Quality	182
EcoSystem Technology	183
Framing and Picture Composition	185
Making Better Pictures	185
Choosing the Viewpoint	185
Moving the Perspective Camera	186
Picture Format.	187
Center of Interest	187
Balance.	188
Light and Shape.	189
Patterns, Colors and Textures	189
The Power of Lines.	190
Conclusion	190
Options and Preferences	191
General Preferences	191
<i>Load/Save Options</i>	192
<i>Preview Options</i>	193



<i>Connections to E-on Software Website</i>	194
<i>Object Options</i>	195
<i>Poser Import Options</i>	196
<i>Render Options</i>	196
<i>Undo/Redo Options</i>	197
<i>Configuration</i>	198
<i>EcoSystem Options</i>	199
<i>User interface options</i>	199
<i>Gamma Options Dialog</i>	200
Display Options	202
<i>3D View Display Quality</i>	203
<i>OpenGL Texturing Options</i>	205
<i>OpenGL Lighting</i>	205
<i>OpenGL Atmosphere Preview</i>	206
<i>Preview Dynamic EcoSystems</i>	206
<i>View Options</i>	206
<i>Maintain Instant Draw Frame Rate</i>	209
<i>Degraded Modes</i>	210
<i>View Background Options</i>	211
<i>OpenGL Clipping</i>	211
Units & Coordinates	211
<i>Length Units</i>	212
<i>Snapping Grid Resolution</i>	213
<i>Order of Rotations</i>	214
<i>World Coordinate System</i>	214
<i>Spherical scene</i>	214
<i>Sea Level</i>	215
Operations	216
<i>Customizing Keyboard Shortcuts</i>	217
<i>Triggers and Trigger Modifiers</i>	218
<i>Shortcuts Files</i>	219
<i>Additional Texture Map Folders</i>	219
<i>VUE Content Folder</i>	219
<i>Image Viewer</i>	220
<i>Preferred Documentation Media</i>	220
<i>Triggers and Modifiers</i>	221



Cornucopia3D™	225
<i>Integrated Cornucopia3D Store</i>	225
<i>Copy-Protected Items</i>	226
What's New in VUE 2016 Release 2.	227
<i>Additional ViewPorts</i>	227
<i>Panoramic Rendering</i>	227
<i>Multi-Materials in the Object Properties Box</i>	227
What's New in VUE 2016 Release 3.	228
<i>Optimized scene saving process</i>	228
<i>Path Tracer renderer improvements</i>	228
<i>Cloud Modulation</i>	229
What's New in VUE 2016 Release 4.	229
<i>Stereoscopic render</i>	229
<i>180 Degree VR Panoramas</i>	230
<i>Path Tracer Renderer Improvements</i>	230

2 Building Scenes

Creating Objects	233
Primitive Objects	233
Infinite Planes	234
Alpha Planes	235
Terrains	235
Plants	237
<i>Plants From PlantFactory</i>	238
Polygon Meshes	238
<i>Rigged Meshes</i>	239
3D Text	240
Rocks	241
A Rock as an EcoSystem Specimen	241
Planets	242
MetaClouds	242
Lights	243
<i>Simple Light Sources</i>	243
<i>Area Lights</i>	245



Group Objects	246
<i>Boolean Operations</i>	246
<i>Metablobs</i>	246
<i>Hyperblobs</i>	246
<i>Working with Group Objects</i>	247
Ventilators	248
Cameras	248
Camera Target	248
Advanced Options	249
Editing Objects	250
Selecting Objects	250
<i>Inside the 3D Views</i>	251
<i>Inside the World Browser</i>	251
<i>By Category</i>	251
<i>Walking through a Selection of Objects</i>	252
<i>Deselecting Everything</i>	252
Moving Objects	252
Rotating Objects	253
<i>Inside the 3D Views</i>	253
Resizing Objects	253
<i>Inside the 3D Views</i>	253
Twisting Objects	254
<i>Inside the 3D Views</i>	254
<i>Using the Numerics Tab</i>	255
Gizmo Manipulators	255
<i>Selecting the Appropriate Gizmo</i>	255
<i>The Position Gizmo</i>	256
<i>The Rotation Gizmo</i>	257
<i>The Size Gizmo</i>	257
<i>Gizmo Coordinates</i>	258
<i>Customizing Gizmo Behavior</i>	259
Changing Object Material	260
Changing Object Preview Color	261
Editing Lights	261
Point Light and Quadratic Point Light	262



Directional Light	264
Spot Light and Quadratic Spot Light.	265
<i>Lens Flare</i>	266
<i>Light Gel</i>	267
<i>Volumetric Light</i>	268
<i>Shadow and Lighting</i>	269
Light Panel	269
Editing Terrains	270
Editing Bodies of Water	271
Geometry	272
Use Global Wave Control	272
Waves.	273
Foam Along Coasts	273
Foam Over Waves	273
Underwater Caustics.	274
<i>Changing the MetaWater Material</i>	274
Editing Plants	275
<i>Applying Wind</i>	275
Ventilators.	276
<i>Editing Plant Geometry</i>	276
Editing Polygon Meshes	277
<i>Polygon Mesh Options Dialog</i>	277
Mesh Information	278
Modifiers.	278
Turbo Smooth	279
Illumination Baking	279
<i>Published Parameters</i>	281
<i>Decimating Imported Objects</i>	281
<i>Baking Objects to Polygons</i>	282
Direct Re-Posing of Rigged Meshes.	284
Animation Creation	286
Editing the Torus	287
<i>Torus Options Dialog</i>	287



Editing 3D Text	287
Editing Alpha Planes	288
<i>Alpha Plane Editor</i>	288
Billboards	289
Editing Planets	290
Boolean Objects.	291
<i>Boolean Union</i>	292
<i>Boolean Intersection</i>	293
<i>Boolean Difference</i>	294
Metablobs	296
<i>Metablob Options</i>	297
Hyperblobs.	298
<i>Hyperblob Options</i>	299
Using Hyperblobs in EcoSystems	300
Ventilators	300
Ventilators.	301
Directional Ventilators.	302
Replacing Objects	302
Saving Objects	304
Working with Pixologic ZBrush	304
<i>Installing GoZBrush for Vue</i>	304
<i>From VUE to ZBrush</i>	305
Scatter -- Replicate Objects.	305
<i>Scatter Objects Tab</i>	306
<i>Replicate Objects Tab</i>	306
Camera Options and Framing	308
Framing.	308
<i>Frame Guides</i>	309
Display Framing Strips	310
Safe Frames	310
Field Grids	310
Using the Camera.	311
<i>Managing Cameras</i>	311
<i>Camera Target</i>	314
Selection and Visibility	315



Focusing on Objects	315
<i>Camera Backdrop Options</i>	316
<i>Advanced Camera Options</i>	317
Advanced Camera Options	318
<i>Camera Manager</i>	326
Rendering	328
Bucket Rendering	328
Render Options	329
<i>Preset Render Settings</i>	330
<i>Renderer</i>	331
<i>Render Destination</i>	333
<i>Render What?</i>	334
<i>G-Buffer / Multi-Pass Options</i>	334
<i>Enable Diagnosis Buffer</i>	335
<i>Enable Relighting</i>	336
<i>Render Quality</i>	337
<i>Advanced Effects Quality</i>	339
<i>Anti-Aliasing</i>	339
<i>Indirect Lighting Solution</i>	340
<i>Picture Size and Resolution</i>	341
<i>Locking User Defined Aspect Ratio</i>	343
<i>Add Information Strip</i>	343
<i>Panoramic View</i>	344
<i>Stereoscopic rendering</i>	345
Convergence modes	346
<i>Render Area</i>	347
<i>Rendering the Selected Render Area</i>	347
<i>Memory Optimization</i>	348
<i>Closing the Dialog</i>	348
G-Buffer Multi-Pass Options	349
<i>G-Buffer</i>	349
Rendering	349
Channels	350
Saving Pictures as RLA or RPF Files	351
Saving Animations as RLA or RPF Files	351



<i>Multi-Pass</i>	352
Rendering Components	352
Extra Components	353
Layer Masks	354
Object Masks	355
Material Masks	355
EcoSystem Material Masks	355
Cloud Masks	356
Saving as Multi-Layer Files	356
Sub-Ray Options	358
Blur Rendering Options	360
<i>Number of Passes</i>	361
Advanced Effects Options	362
<i>General Tab</i>	363
<i>Photon Maps Tab</i>	364
Anti-Aliasing Options	365
<i>Anti-Aliasing Options</i>	365
Object Anti-Aliasing	367
Anti-Aliasing Strategy	367
Subrays	368
Texture Filtering	369
Texture Anti-Aliasing	369
Batch Rendering	371
Rendering To Screen -- The Render Display Window	373
<i>Current Render Display</i>	374
<i>The Render Stack</i>	376
<i>Comparison</i>	378
<i>Reapplying Camera Settings to Current Camera</i>	378
<i>Hiding the Stack</i>	378
Post Render Options	379
Non Photo-Realistic rendering	380
Film Settings	381
Lens Glare	382
Post Processing	383



Non-Photorealistic Rendering.	384
<i>Activating NPR and selecting style presets</i>	385
From the main toolbar	385
From the Advanced Camera Options	386
From the Post Render Options in the render stack.	386
<i>Outline Tab</i>	386
<i>Line Style Tab</i>	396
<i>Line Distortion Tab.</i>	398
<i>Line Cloning Tab.</i>	399
<i>Shader Tab</i>	402
Saving styles.	404
Path Tracer.	405
How to Use It	406
Tips.	408
Saving (Exporting) Images	409
<i>Saving Animations.</i>	410
Printing Pictures	411
<i>Print Preview</i>	411

3 Importing and Exporting

Importing Objects	415
Import Options.	417
Poser Import Options.	419
<i>Animation Import.</i>	420
<i>Re-Posing.</i>	421
<i>Rendering Using the Poser Shader Tree</i>	423
Exporting Content	424
Export Options.	425
General options	426
Output format	428
Export geometry	428
Export 2D heightmap	429
Export levels of detail	429
Export billboard	429
Export materials	430



Map Output Options	431
Rigging and animation	432
Export sky.	433
Exporting Objects	437
Exporting Entire Scenes	438
EcoSystem Export	439
Limitations	440
Export Preview	440
<i>Export Objects on a Per-Object Basis</i>	441
PTex Export Preview	442

4 The Editors

Terrain Editor	445
Terrain Types	446
<i>Heightfield Terrains</i>	446
<i>Procedural Terrains</i>	447
Procedural Terrain Presets	447
Touching Up Procedural Terrains	448
Mapping Modes.	448
Changing Mapping Mode	449
Spherical Terrains.	450
<i>Creating and Manipulating a Spherical Terrain</i>	451
<i>Basic Spherical Scene</i>	451
<i>Planet Spherical Scenes</i>	452
<i>Planetary Mapping Mode</i>	453
<i>Objects in Spherical Scenes</i>	453
<i>Atmospheres in Spherical Scenes</i>	454
Editing Terrains	455
<i>Handling the 3D Terrain View</i>	456
<i>Terrain Editor Top Toolbar</i>	457
<i>Terrain Editor -- Predefined Terrain Styles</i>	461
<i>Terrain Editor -- Sculpting Terrains</i>	468
<i>Terrain Editor -- Brush Editor</i>	471
<i>Terrain Editor -- Painting Materials</i>	474
<i>Terrain Editor -- Procedural Altitudes</i>	477



<i>Terrain Editor -- Zones</i>	479
<i>Terrain Editor -- Effects</i>	480
<i>Terrain Editor -- Exporting Terrain Geometry</i>	485
<i>USGS Digital Elevation Model Data</i>	487
Hint	488
User Touch-up Graph Node -- Heightfield Terrains	488
Text Editor	490
<i>Interface Elements</i>	491
Toolbar	491
Text Preview	493
Text Tab.	494
<i>Text</i>	494
<i>Text Style</i>	495
<i>Using Vector Graphics</i>	495
Bevel Tab.	496
Extrusion Tab	497
Materials Tab	499
Text Styles	500
Plant Editor	501
Plant Preview	502
Plant Preview for TPF Plants	503
<i>Mesh Resolution</i>	503
<i>Published Parameters</i>	504
Toolbar	504
Editing Plants	507
<i>Trunk and Branches</i>	507
Description	507
All Subsets	508
Empty Subsets	508
Trunk and Branch Materials	508
Trunk and Branch Settings.	510
<i>Leaves and Petals</i>	511
Description	511
All Subsets	511
Empty Subsets	511



Leaf and Petal Materials	511
Leaf and Petal Settings	512
<i>Creating Variations of the Same Plant</i>	513
<i>Creating New Species</i>	513
<i>Exporting Plants</i>	513
Leaf Editor	514
Response To Wind Options	515
<i>For VUE .veg files</i>	515
<i>For PlantFactory .tpf files</i>	516
Using PlantFactory Plants With Presets	516
Atmospheres	519
Loading an Atmosphere	519
Sun Tab	522
Light Tab	525
Global Lighting	529
Light Color	530
Auto Decay Sunlight Color	530
Changes for Photometric Atmospheres	531
Sky Tab	532
Clouds Tab	533
<i>Simple Cloud Animation</i>	536
<i>Volumetric Clouds</i>	536
<i>Spectral Clouds</i>	537
<i>Cloud Layers as Objects</i>	537
<i>Cloud Layer Zone</i>	538
<i>Advanced Cloud Material Editor</i>	539
Color & Density Tab	539
Large scale density Tab	543
Lighting & Effects Tab	544
Fog and Haze Tab	547
<i>Fog</i>	549
Standard Atmosphere Model Only	549
<i>Haze</i>	550
Volumetric Atmosphere Model Only	550



Sky, Fog & Haze Tab	551
Sky	551
Fog and Haze	552
Global Settings	553
Wind Tab	554
<i>Adjusting Breeze</i>	555
Breeze Settings	555
Gusts of Wind	556
Influence of Wind Intensity.	556
Fluttering of Leaves	556
Breeze Preview	557
Effects Tab	558
Stars	559
Rainbow	560
Ice Rings.	560
Default Lens Flares	561
Environment Map	562
Separate Illumination Map	564
Default Reflection Map	565
Use Environment Map Beyond Atmosphere	565
Rain & Snow Tab	566
Saving an Atmosphere.	568
Photometric atmospheres in xStream	568
<i>3dsMax</i>	569
V-Ray	569
Mental Ray	570
<i>Maya</i>	570
<i>Cinema 4D</i>	571
<i>Lightwave</i>	571
<i>SoftImage</i>	572
<i>Natural Film Response</i>	572
Automatic Sun Softness	573
How to use it	574



Light Editor	575
Lens Flare Tab.	575
Ring	576
Color Shift.	577
Random Streaks	577
Star Filter	577
Reflections.	578
Fading	578
Lens Flare Reflections Editor	579
Polygonal Reflections	579
Adding Interpolated Reflections	580
List of Reflections	580
New, Load, Save	583
Gel Tab	583
Volumetric Light Tab	584
<i>Volumetric Light Controls.</i>	585
Smoke/Dust in Volumetric Light Beams	585
Shadow Tab	586
Enabling Shadows	586
Using Shadow Maps	587
Projected Hard Shadows	587
Quality	588
Softness	589
Additional Information	590
Lighting Tab	591
Light Attenuation	591
Variable Color.	592
Photometric Settings	593
Influence Tab	594
Specular and Diffuse Components	594
Objects Influenced by Light	594
Material Editor	596
<i>Types of Materials</i>	598
<i>Multi-Materials</i>	598



Common Material Controls	599
New, Load, Save	599
Defining the Material	599
Type	601
Effects	602
Material Hierarchy	603
Layering Materials	603
Material Previews	604
Store	605
Tabs	605
Basic Material Editor	606
<i>Simple Materials in Basic Editor.</i>	607
Color Frame	607
Bump Frame	607
Transparency Frame	608
Other Settings	609
<i>Mixed Materials in Basic Editor</i>	609
<i>Layered Material in Basic Editor</i>	611
Advanced Material Editor	612
<i>Driving Material Settings with Functions.</i>	613
<i>Published Parameters</i>	613
Simple Materials	614
<i>Material Layer</i>	614
<i>Color & Alpha Tab</i>	614
Mapped Picture	615
Procedural Colors	620
Natural Grain	623
<i>Bumps Tab</i>	624
<i>Highlights Tab</i>	627
Highlights	629
<i>Transparency Tab</i>	631
<i>Reflections Tab.</i>	635
<i>Translucency Tab</i>	638
<i>Effects Tab</i>	640
<i>Presence Tab.</i>	644



Mixed Materials	646
<i>Materials to Mix</i>	647
Distribution of Materials	647
Smooth Blending Strip	648
Blending Method	648
<i>Alpha</i>	649
<i>Influence of Environment</i>	649
Influence of Altitude	650
Influence of Slope	650
Influence of Orientation	650
Altitude Range	650
Coordinate System	651
Volumetric Materials	651
<i>Color and Density</i>	652
Density Production	652
Volumetric Settings	652
<i>Lighting and Effects</i>	653
Lighting	653
Flare	655
Volume Shaded Material	655
Hypertextures	655
Origin of Material	655
Velocity of Material Origin	656
Global Transformation	656
Dissolve Near Objects	656
Layered Materials	657
<i>Multi-Layer Materials</i>	658
Adding a Layer	658
Changing the Order of Evaluation	658
Influence of Layers	659
Material Snapshots	659
<i>The Shared Material Layer</i>	659
Grouped Materials	660
<i>Alpha Tab</i>	661
<i>Presence Tab</i>	661
Slope Constraint	662



Influence of Orientation	662
Two-Sided Material	663
<i>Materials to Mix Tab</i>	663
<i>Alpha Tab</i>	664
Function Graph	665
Description	665
The Basics	666
<i>What Is a Graph?</i>	666
<i>Input and Output Nodes</i>	666
<i>Output Data</i>	667
<i>Multiple and Master Outputs</i>	667
<i>Nodes</i>	668
<i>Types of Data</i>	669
<i>Links</i>	669
SmartGraph™	670
Presentation	670
<i>Top Toolbar</i>	671
Node and Function Previews	671
Toolbar Icons	672
<i>Nodes Toolbar</i>	675
<i>Function Graph</i>	682
Node Selection	682
Adding/Replacing Nodes	683
Connecting Nodes	683
Published Parameters	684
<i>Node/Link Details</i>	684
Node Details	684
Extracting Constant Values	685
Link Details	685
Multi-edition of nodes	686
MetaNodes	686
Creating a MetaNode	686
Editing a MetaNode	686
Building a MetaNode Interface	686
Saving and Re-using MetaNodes	687
Example Use	687



Locking MetaNode Content	688
Scene Graph Approach	688
<i>Object Graphs</i>	689
Connecting Graphs.	689
External Dependencies	689
Recall Dependency	690
Exporting Values	691
Function Node Preview	692
Curves and Extension	693
Function Output Observer	693
Cyclic Nodes.	695
<i>How to Use Them</i>	695
Noise Nodes	696
<i>Common Parameters</i>	696
Scale	696
Wavelength	696
Origin	696
<i>Cellular Patterns</i>	696
Chipped, Crystals, Pebbles	697
Drought	698
Voronoi.	698
Voronoi (Altitude)	703
Voronoi (Generalized)	705
<i>Distributed Patterns</i>	708
Round Samples and Round Samples (2D)	708
Square Samples and Square Samples (2D)	722
<i>Flat Patterns</i>	728
Varying Blocks, Clumps, Water Cress	728
<i>Line Patterns</i>	729
Lines, Fabric	729
Cracks	730
Sparse Cracks	731
<i>Math Patterns</i>	733
Onion, Wavelet, Step (Vertical), Step (Gradual), Tooth (Rectangular), Tooth (Triangular), Tooth (Gaussian), Radial Sine, Sine Wave, Triangular Wave, Leopard, Saw Teeth, Water Wave	733



Spiral	740
Rectangular Wave	742
Other Patterns	746
Dots, Water (Calm), Water (Rough), Granite.	746
Perlin Noises	747
Common Parameter	747
Linear, Value, Gradient	747
Value-Gradient (Variable), Linear-Value-Gradient (Variable).	749
Value-Gradient, Linear-Value-Gradient.	750
Square Patterns	753
Random Altitudes, Squares, Squares (Pairs), Stones, Square Blobs, Square Stones	753
Fractal Nodes	756
Common Parameters.	756
Basic Repeater	767
Simple Fractal.	775
Grainy Fractal	776
Terrain Fractal	777
Terrain Fractal 2	781
Rocky Mountains Fractal and Eroded Rocky Mountains Fractal.	782
Fast Perlin Fractal	785
Variable Roughness Fractal.	786
Variable Noise Fractal	792
Three Noise Fractal	796
Open Ocean	802
Color Nodes	807
Color Creation Nodes	807
Color Map	807
2 Color Output	808
Linear Interpolation 2	808
Spline Interpolation 2	808
3 Color Output	808
Linear Interpolation 3	809
Spline Interpolation 3	809
Color Variation	809
Color Brightness Variation	809



Two Color Variation	810
Natural Color Blend 2.	810
Color Variation Map.	811
Terrain Color Patterns	811
Color Correction Nodes.	812
Common Settings.	812
Gamma.	812
Gain	812
Brightness	812
Contrast	813
HLS Shift	813
HLS Color Shift.	813
Filter	813
Perspective	814
Color Blender	814
Texture Map Nodes.	814
Mapping Nodes	814
Texture Map	814
Projected Texture Map	815
Animation Map.	816
Projected Animation Map Node.	817
Blended Image Node.	818
Blended Grayscale Image Node	818
Image Sample and Multi-Image Sample Nodes	819
Ptex node	820
UV Coordinates Node.	820
Filter Nodes	821
Environment Sensitive Filters	821
Altitude	821
Slope	821
Altitude and Slope	822
Orientation	822
Environment	822
Patches.	822



<i>Recursive Nodes</i>	824
Strata Processing Data Parameters	825
Strata Positioning Features	825
2nd Output: Detect Rough Areas	826
Confined Strata	826
3D Stratification	826
<i>Input Filters</i>	827
Filter	827
Partial Filter	827
Offset ($X + a$)	827
Opposite ($-X$)	827
Multiply (aX)	827
Divide (a/X)	828
Brightness-Contrast ($aX + b$)	828
Parabolic ($aX^2 + bX + c$)	828
Absolute	828
Gamma	828
Bias	828
Gain	829
Power	829
Gaussian	829
Floor	829
Ceiling	829
Clamp	829
Clip	830
Smooth Clip	830
Map	830
Smooth Map	830
Quantize	831
Saw Wave	831
Absolute Wave	831
Sine Wave	831
Threshold	831
Smooth Threshold	832



Constant Nodes	832
Constant Number	832
Constant Color	832
Constant Coordinates	832
Constant Vector	832
Random Constant Number	832
Connectable Constant	832
Turbulence Nodes	833
Simple Turbulence	833
Misc Style Turbulence	842
Advanced Turbulence	848
Combiner Nodes	851
Blender	851
Combiner	851
Color Combiner	852
Add	853
Subtract	853
Multiply	853
Divide	853
<i>Image Combiner Node</i>	854
Math Nodes	854
Conversions	854
Vector To RGB	854
RGB To Vector	854
RGB To HLS	854
HLS To RGB	855
Color To Brightness	855
Vector Operations	855
Offset	855
Rotation And Twist	855
Projection	855
Matrix Transformation	855
Decomposer	856
Composer	856
Length	856
Normalize	856



Dot Product	856
Vector Product	856
Vector Quantization	856
Other Math Nodes	858
Sine	858
Arc Cosine	858
Floor	858
Fractional Part	858
Invert	858
Power	858
Square Root	858
Multiply	858
Dynamics Nodes	859
Link Relationship	859
Track Relationship	859
Derivative	859
Integral	859
Delay	860
PID Controller	860
Distance Constraint	861
Axis Constraint	861
Grid Constraint	861
Acceleration Limiter	861
Speed Limiter	861
Low Pass Filter	861
Examples	862
<i>Creating Turbulence</i>	862
<i>Slope Dependent Scale</i>	862
<i>Variable Color-Texture Mapping</i>	863
Heightfield Nodes	864
Erosion nodes	864
Available outputs	865
Presets	866
Global Parameters	872
Terrain Properties	872
Flow Parameters	873



Weathering Parameters	873
Flow Surface Output.	874
<i>Slope node</i>	874
<i>Convexity node</i>	874
<i>Blur node</i>	875
<i>Terraces node</i>	875
<i>Auto-mapping node</i>	879
Layout Nodes	880
<i>Area Demarcation Node</i>	880
Available parameters.	881
<i>Spline Proximity node</i>	881
How to use it	881
Node outputs	884
<i>Scene Objects Mask Node</i>	885
Spline Editor	886
<i>Material Effect</i>	888
<i>Geometry Effect</i>	890
<i>Terrain Effect</i>	892
Editing Filters	893
Description	893
<i>The Curve</i>	894
<i>Smooth Filters</i>	894
<i>Toolbar</i>	895
<i>New, Load, Save</i>	896
Profile Tab	897
<i>Adding Key Points</i>	897
<i>Modifying Key Points</i>	898
<i>Deleting Key Points</i>	898
<i>Clamping modes</i>	898
Influence of Environment Tab	899
<i>Dependent of Altitude</i>	900
<i>Dependent of Slope</i>	900
<i>Dependent of Orientation</i>	900



Editing Color Maps	901
Description	901
<i>Adding Key Colors.</i>	902
<i>Adding Key Opacities.</i>	902
<i>Modifying Key Colors & Opacities</i>	903
<i>Manipulating Multiple Key Colors</i>	903
<i>Deleting Keys</i>	904
<i>Advanced Opacity Control</i>	904
<i>New, Load, Save.</i>	904
Selecting Colors.	905
<i>Quick Color Selection.</i>	905
<i>Color Selection Dialog</i>	906
Favorite Colors	907
Summary of Materials	908
<i>Loading, Editing and Scaling Materials</i>	909
Interface Colors	910
<i>Flat Interface Colors.</i>	910
<i>New, Load, Save.</i>	913
Macros	914
Recording Macros	914
Playing Macros	914
Macro QuickLaunch	915

5 EcoSystems

Painting EcoSystems	920
EcoSystem Painter	920
<i>EcoSystem Painter Tools.</i>	922
<i>Painting With Brushes</i>	924
Building the EcoSystem Population.	924
Influencing the EcoSystem Population	924
Global Settings	926
<i>EcoSystem Population</i>	927
Paint What?.	927
<i>Global EcoSystem Object</i>	928
Multiple Global EcoSystem Objects.	929



<i>Using the Paint Function to Modify EcoSystem Materials</i>	929
<i>EcoSystem Display Options</i>	930
<i>Brush Editor</i>	930
General Tab	931
Environment Tab	934
<i>Creating an EcoSystem Content Brush</i>	936
<i>Selecting EcoSystem Instances</i>	937
Selecting EcoSystem Instances	938
Manipulating Selected Instances	939
The Vue EcoParticle System	941
Setting up an EcoParticle System	942
<i>General Tab of the Material Editor</i>	942
<i>Particle Characteristics Dialog</i>	943
Properties Tab	943
Evolution Tab	944
Periodic Emission Tab	944
Collision Tab	945
Death Tab	945
<i>Bake Particle Motion Dialog</i>	945
<i>Global Particle Configuration Dialog</i>	946
General	946
Collision	947
Forces	947
Rendering	948
<i>Density Tab of the Material Editor</i>	948
<i>Scaling & Orientation Tab of the Material Editor</i>	949
<i>The Particles Effector</i>	949
Effector Editor	950
Animating an EcoParticle System	950
Limitations of the EcoParticle System	950
EcoSystems in the Material Editor	952
Temporary Global Settings for Quality Display of EcoSystems	954
General Tab	955
EcoSystem Population	955
Underlying Material	957



Fast Population Mode	958
Display Options.	958
Distribution	959
Density Tab	961
Overall Density	961
Placement	961
Offset from Surface.	962
Slope Influence	963
Clumping	963
Variable Density	963
Decay near Foreign Objects.	963
Scaling & Orientation Tab	965
Overall Scaling	965
Maximum Size Variation	965
Direction from Surface.	966
Rotation	966
Variable Scaling.	967
Shrink at Low Densities.	967
Lean Out at Low Density	968
Color Tab	968
Color Correction	968
Color at Low Densities	969
Variable Color.	969
Presence Tab	970
Altitude Constraint	971
Slope Constraint	971
Influence of Orientation	971
Animation Tab	972
Converting Objects To EcoSystem Instances	974
EcoSystem Export	975

6 Tutorials

Introduction	979
Quick Reference How-To's	979
<i>Rendering and Saving a Picture</i>	<i>979</i>



<i>Rendering a Full Screen Picture</i>	980
<i>Rendering a High Resolution Picture</i>	980
<i>Creating a Rock Material</i>	980
<i>Mapping a Material Using a Picture</i>	981
<i>Creating Fog at Low Altitudes</i>	981
<i>Making a Custom Canyon Terrain</i>	982
<i>Creating Terrains with Snow Covered Tops</i>	982
<i>Adding a Sense of Depth to Scenes</i>	983
<i>Mixing Materials with Tilted Stratum</i>	984
<i>Creating Undulated Furrows</i>	984
<i>Adding Color Variations to Materials</i>	985
<i>Modifying a Terrain</i>	986
<i>Modifying a Terrain Externally</i>	986
<i>Importing Multi-Part DEMs</i>	987
Elaborate Feature Tutorials	988
Boolean Objects.	988
<i>Building a Simple House</i>	988
Clipping Terrains	990
<i>Making a Stone Arch</i>	990
Fuzzy Materials	992
<i>Modeling Clouds</i>	992
Additive Materials.	995
<i>Faking Volumetric Lights</i>	995
Using Pictures inside Scenes.	997
<i>Making a Road Sign</i>	998
<i>Glowing Neon Lights</i>	999
Distant Planets	1000
Underwater Scenery	1001
Convincing Image Based Lighting	1003
Varying Materials on EcoSystem Populations	1005
Creating a Piston Rig.	1008
Interactive Procedural Terrain Modeling.	1013
Painting Materials on Terrains	1017
Camera Mapping	1019
Animated Dynamic EcoSystems	1024



Clouds Modulating with Landscapes	1029
Creating a Planet	1032
Creating a Road Using a Spline	1036
Building a Complete Scene	1040
Shaping out the Forest	1040
Adding a Distant Background	1041
Tuning the Atmosphere	1043
Adding Vegetation	1044
Importing the Subject	1045
Tuning the Light	1046
Animation Tutorials	1048
A Pursuit	1048
Getting a Spin	1051
Opening a Window	1053
Angular Paths	1054
Drop and Bounce	1057
Flickering Lights	1059
Animating a Fish	1061
A Steam Power Train	1065
Making Waves	1068
Animating Clouds	1071
Moving Caustics	1073
Stoned Frog	1074
Animating Groups and Boolean Objects	1077
Using Spin and Vibrate Effects	1079
Dying Plants	1081
Animating an Ocean	1084
Using Animated Billboards to Create a Fire	1085
Dying Flesh	1088
Loose Dynamics	1090
Crashing Meteor	1091
EcoSystem Phasing	1095



7 Animating Scenes

Animating from Scratch	1101
Animation Properties Tab	1101
<i>Types of Motion</i>	1101
Animation Wizard	1103
<i>Step 1: Introduction</i>	1104
<i>Step 2: Selecting a Motion</i>	1105
<i>Step 3: Global Animation Settings</i>	1106
Repeat Mode	1106
Main Axis	1107
Speed Mode	1107
<i>Step 4: Advanced Effects</i>	1108
Spin	1108
Note on Spin and Pivots	1109
Vibrate	1109
<i>Step 5: Object Path</i>	1111
Adding Way Points	1111
Editing Way Points	1112
Inserting Way Points	1112
Deleting Way Points	1112
Scrolling/Zooming the View	1112
Restrictions	1112
The Tunnel Case	1113
<i>Step 6: Animation Setup</i>	1113
<i>Step 7: Animation Preview</i>	1114
<i>Step 8: Conclusion</i>	1114
Animating with the Timeline	1115
Navigating Inside the Animation	1116
Keyframes	1117
Animation Properties	1118
<i>Standard Primitive, Polygon Mesh, Group, Boolean Object, and Metablob and Plant</i>	1118
<i>Torus</i>	1119
<i>Terrain</i>	1119
<i>Plant</i>	1119



<i>Directional Light</i>	1119
<i>Point Light and Quadratic Point Light</i>	1120
<i>Spot Light and Quadratic Spot Light</i>	1120
<i>Ventilators.</i>	1120
<i>Camera</i>	1120
Animating Objects	1121
<i>Creating the Animation</i>	1122
Published Parameters for Animated Objects	1122
Published Parameters for Animated Textures	1123
<i>Working with Keyframes</i>	1123
Keyframe Types	1123
Selecting Keyframes	1123
Moving Keyframes	1123
Adding Keyframes	1124
Copy-Pasting Keyframes	1124
Deleting Keyframes	1124
Modifying the Value of a Keyframe	1124
Keyframe Tangents.	1125
Keyframe Values	1127
Quaternion vs. Euler Rotations	1127
<i>Editing Paths in 3D Views</i>	1128
Selecting Way Points.	1128
Gray Way Points	1129
Moving Way Points	1129
Editing Way Points	1129
<i>Changing Rigged Mesh Motion</i>	1129
<i>Multi-Spins</i>	1129
<i>Animating Plants.</i>	1130
Breeze	1130
Animating Wind	1130
Animating Plant Geometry	1131
<i>Animating Torus Thickness</i>	1131
<i>Camera Switching</i>	1132
<i>Animated Post Processing and Motion-Blur Length</i>	1133
Animating Post Processing	1133
Animating Motion Blur Length	1133



Animating Materials.	1134
<i>Material Surface Animation</i>	1134
<i>Material Velocity Animation</i>	1135
<i>Complete Material Animation</i>	1136
<i>SmartGraph Material Animation</i>	1136
<i>Published Parameters for Animated Materials</i>	1137
Animated EcoSystem Population	1137
Animating the Atmosphere	1138
<i>Atmosphere Keyframes</i>	1138
<i>Animating Clouds</i>	1139
<i>Animating the Sun</i>	1139
Working with Animations.	1140
<i>Pasting Animation</i>	1140
<i>Destroying Item Animation</i>	1140
Destroying Object Animation	1140
Destroying Material Animation	1140
Destroying Atmosphere Animation	1141
<i>Shifting the Start of an Animation</i>	1141
<i>Changing the Duration of an Animation</i>	1141
Using Time Splines	1142
<i>Editing Time Splines</i>	1143
The Curve	1143
Smooth Time Splines	1143
Toolbar.	1144
New, Load, Save	1145
Adding Key Points	1146
Modifying Key Points	1146
Deleting Key Points	1146
Animation Toolbox	1147
Look Ahead	1148
Smoothed Velocity	1148
Changing the Duration of an Animation	1149
Spin and Vibrate	1149
<i>Motion Options Dialog</i>	1149
Influence of Acceleration on Pitch	1150



Linking and Tracking	1151
Linked Hierarchies	1151
Tracking Objects	1152
<i>Loose Linking and Tracking</i>	1152
<i>Forward Dynamics Options</i>	1153
Partial Links	1153
Loose Dynamics	1154
Object Graph	1154
<i>Camera Mapping</i>	1155
Rendering the Animation	1157
<i>Animation Render Options</i>	1158
Channel Files	1158
Animation File Formats	1159
Frame Rate	1160
Frame Resolution	1160
Timecode	1160
Renderer	1161
Miscellaneous	1162
Closing the Dialog	1163
<i>Advanced Animation Options</i>	1163
Flicker Reduction	1163
Field Interlacing	1165
Pixel Aspect Ratio	1165
Automatic Illumination Baking	1166
<i>Animation Preview Options</i>	1167
<i>Recovering TMP Files from an Aborted Render</i>	1168
Technical Notes	1169
Rotating Look Ahead and Track Objects	1169
Synchronizing Cameras and Lights	1170
Import Synchronization Data	1170
<i>VueSynch</i>	1171
Installing Plug-Ins	1171
Generating Synchronization Data	1171
Synchronizing Objects	1172



Import Motion Tracking Information	1173
Importing Vertex Keyframe Animation	1174
Exporting Animation	1175

8 Appendices

Mouse and Keyboard Operations	1179
Mouse Operations	1179
<i>Inside any dialog and editor</i>	1179
<i>Inside the 2D/3D Views</i>	1180
Selecting and Dragging Objects	1180
Choosing transformation axes	1180
Using Bounding Box Handles	1182
Industry-standard navigation triggers.	1183
Legacy navigation triggers	1183
<i>When Editing Paths in 2D/3D Views</i>	1184
<i>Inside the World Browser</i>	1184
<i>Inside the Timeline</i>	1185
<i>Inside the Animation Wizard Path Editor</i>	1186
Mouse Wheel Operations	1186
Keyboard Operations.	1187
<i>Interface Shortcuts</i>	1187
<i>Object Creation</i>	1189
<i>Object Edition</i>	1190
VUE xStream	1192
Standalone and Integrated Modes	1192
VUE Licenses	1192
VUE Installation	1193
<i>32bit versus 64bit Versions</i>	1193
<i>Mental Ray Configuration Files</i>	1194
<i>V-Ray Renderer</i>	1194
Supported Versions of the Host Applications.	1194
Adding the VUE Menu and Toolbars.	1195
<i>3DS Max</i>	1195
<i>Cinema4D</i>	1196
<i>LightWave</i>	1197



<i>Problem with Menu Display -- Incomplete Menus</i>	1197
<i>Maya</i>	1199
<i>Softimage</i>	1200
xStream Options Dialogs	1200
<i>xStream Options Dialog</i>	1200
<i>General Tab</i>	1201
Scene File	1201
Time	1202
Scale	1202
Light Options	1203
VUE Proxies	1203
Misc.	1204
<i>Render Options Tab</i>	1205
Render VUE Scene	1205
LightWave renderer (LightWave only)	1206
Mental Ray & Vray (only applications supported for Vray and Mental Ray)	1206
Post Processing	1206
Render Quality	1207
Final Gather Matching	1207
<i>VUE Light Options Dialog</i>	1208
Native Lights	1209
VUE Lights	1211
<i>VUE Scene Options Dialog</i>	1212
Spherical scene	1212
Sea Level	1212
<i>VUE Render Options Dialog</i>	1213
xStream Integration	1214
<i>VUE Commands</i>	1214
Scene Display in Host Application Viewports	1217
<i>VUE Dialogs</i>	1217
<i>Proxy Objects</i>	1218
<i>VUE Proxy Objects</i>	1218
Scene Loading Example	1219
VUE Splines in a Host Application	1220



<i>Editing the VUE Elements</i>	1220
<i>Special Case for Lights and Cameras</i>	1221
<i>Animating Your VUE Elements</i>	1222
<i>Further Edition of the VUE Scene</i>	1222
<i>Display of the Proxy Objects.</i>	1223
3DS Max.	1223
Cinema4D	1224
LightWave	1224
Maya	1224
Softimage	1224
<i>Editing Proxy Objects.</i>	1224
<i>Scene Updates.</i>	1225
xStream EcoSystem Painter	1227
Painting EcoSystems in 3DS Max	1228
Painting EcoSystems in Cinema4D	1228
Painting EcoSystems in Softimage.	1228
Notes about Painting	1228
Rendering with VUE xStream.	1229
<i>Render Quality.</i>	1230
<i>G-Buffer and Multi-Pass</i>	1231
<i>Post-Processing</i>	1231
<i>Advanced Rendering Effects</i>	1232
Global Illumination	1232
Saving Global Illumination Prepass for Reuse.	1233
Motion Blur	1234
Motion Blur with 3DS Max.	1234
Motion Blur with Cinema 4D	1234
Camera Motion Blur with Mental Ray	1235
Motion Blur Settings in Cinema 4D.	1235
Depth of Field	1235
Saving Your Work with VUE xStream	1236
Automatic Backup of 3DS Max	1236
Incorporating VUE Scenes inside Native Scenes	1236
Additional Content.	1236
Sharing Scenes with Other Users	1237



Export Animation to VUE	1237
Converting VUE Objects into Native Objects	1238
FBX Import	1244
<i>Objects</i>	1244
<i>Camera</i>	1244
Camera Synchronization	1244
Supported features from Maya	1245
Nuke EXR Exporting	1245
<i>Compositing</i>	1245
<i>Relighting</i>	1246
<i>Tutorial</i>	1246
E-on License Server	1247
Node-Locked vs. Floating Licenses	1247
About the License Server	1247
<i>Purchasing a License Server</i>	1248
Using a Floating License from VUE	1248
Network Rendering	1249
RenderCows vs. RenderNodes	1249
Description	1250
<i>HyperVue™ Network Rendering Manager</i>	1250
<i>RenderCow™</i>	1251
<i>Additional RenderCow Licenses</i>	1251
<i>Installing RenderCow Packs</i>	1252
Setting Up RenderCows	1253
<i>Installing a RenderCow</i>	1253
<i>Launching RenderCows at Boot</i>	1253
<i>RenderCow Port Numbers</i>	1254
<i>CPU Affinity</i>	1254
<i>Installing a RenderCow on the Computer Running VUE</i>	1255
Rendering With HyperVue™	1255
<i>Configuring HyperVue™</i>	1255
<i>Putting RenderCows to Work</i>	1257
<i>Starting a Render</i>	1257
<i>Monitoring the Render Farm</i>	1257
<i>Aborting a Render</i>	1258



Managing RenderCows™	1258
<i>Showing RenderCow Status</i>	1259
<i>Adding RenderCows During Render</i>	1259
<i>Pausing a RenderCow</i>	1259
<i>Shutting Down a RenderCow</i>	1259
<i>Updating RenderCows</i>	1260
RenderNodes	1260
RenderNode Licenses	1261
<i>RenderNode Network Options</i>	1261
Temporary Folder	1261
Command to Execute	1262
Arguments in the Command	1262
<i>Sample Setup for BackBurner</i>	1263
Hints on Setting Up Your Network Rendering Manager.	1264
<i>Setting Up RenderNodes</i>	1264
Updating RenderNodes	1264
<i>Controlling RenderNodes via Command Line</i>	1265
<i>Saving Global Illumination Prepass for Reuse</i>	1266
Reusing a GI Prepass File	1267
xStream Network Rendering	1267
<i>Saving Native Scenes for Network Rendering</i>	1268
<i>VUE Licenses for Network Rendering</i>	1268
<i>Setting-up Your Workstation for a Network Render</i>	1269
<i>Behavior without License</i>	1269
<i>Rendering with Mental Ray Satellites</i>	1270
Python Scripting	1272
Python Scripts	1272
Running Python Scripts	1273
Startup Scripts.	1273
Running Python from Command Line	1274
Creating Python Scripts	1275
Python Console	1275
Python Documentation.	1275



Hot Tips	1276
Resuming Render.	1276
Using Layers for Faster Display	1276
Resuming Animation Render	1277
Using Motion Blur in Still Pictures	1277
Avoiding Object Detection in Wizard.	1278
Rendering Time Dependent Materials.	1278
Reducing Animation Render Times	1279
Compressing Video.	1279
Troubleshooting	1280
Scenes Take Ages to Render.	1280
Camera Moves by Itself	1281
Bright Fringes Appear at Wall Base in Radiosity.	1281
Render Time Estimation Is Pessimistic	1282
Long Preparation Time for Small Image.	1282
Program Crashes Randomly	1282
Noise Appears in Volumetric Effects.	1283
Atmosphere Is Different in Preview and Final	1283
Undesired Lens Flares Appear on Lights	1284
Missing Details	1284
Vector Graphics Don't Load in Text Editor	1284
Invisible Objects.	1285
Unable to Select Objects	1285
Objects Don't Render	1286
Close-up Materials Look Like Tiles.	1286
Look Ahead Objects	1286
Objects Overreact to Motion	1287
Materials Don't Move With Objects	1287
Objects Keep Getting Animated	1288
Animations Flicker	1288
<i>Texture Filtering</i>	1289
<i>How It Works</i>	1289
<i>In Practice</i>	1289
<i>Texture Filtering Will Influence 2 Components at Render.</i>	1290



VUE 2016 Infinite & xStream — Reference Manual

Animations Pulsate	1290
Dark Triangles Appear on Terrains.	1291
Wrong Material Scale	1291
Soft Shadows Look Noisy	1292
VUE Objects Lose Relative Positions	1292
RenderCow Not Responding	1293
Unable to Export Object	1293
Maya Mental Ray Renders Black.	1294
Native Objects Don't Reflect VUE Objects	1294
Volumetric Plugin Conflict	1295
Vue Artist Modules	1296
License Agreement	1297
License	1298
Copyright	1299
Restrictions	1300
Content and Content Limitations	1300
Term.	1301
Limited Warranty	1302
Remedies	1303
US Government Restricted Rights	1303
Opening Objects From The Command-Line	1304
Third Party Licenses	1305
Alembic File Format	1305
Apache License 2.0.	1308
Boost	1312
CURL License	1312
FBX SDK	1313
Google-Breakpad	1314
MIT License	1315
Ptex Licensing	1316
zlib License	1317



Section 1

Getting Started





Introduction

System Requirements

VUE is a 64 bit application, designed for the Windows® 64 bits XP, Vista, Windows 7, Windows 8 and Intel Mac OS X platforms.

Like all 3D packages, it is highly demanding in terms of computer power. Although the application is totally multi-threaded to ensure the smoothest possible response, you have to realize that there is a lot going on when you work in VUE. This is why we feel that running it on reasonably recent and reasonably fast computer is best suited. We recommend a minimum of 4GB of RAM.

If you find that the program is not responding as quickly as you would like it to, there are a certain number of actions that you can take that will help speed it up. Please turn to [Options](#) and [Preferences](#) for a complete description of these actions.

The minimum resolution for operating VUE is 1200×768 , in hi-color or true-color modes (at least 16 bits per pixel). For better comfort, we recommend a resolution of at least 1600×1200 in true color.

Installation

The VUE installation files are downloaded in .zip file format. Unzipping this file into a work directory on your hard drive will give you all the files you need to install the software on your computer.



Default Folders for All Versions

By default, VUE installs in the following folders:

Windows Vista, Windows 7, Windows 8 or Windows 10:

- Program files: `c:\Program\ Files\e-on\ software\`
- Configuration files: `c:\Users\[User\ Name]\AppData\Roaming\e-on\ software\`

Prior to the 2016 version, the VUE content files were installed in the following directory:

- `c:\\ [User\ Name]\My\ Documents\e-on\ software\`

With 2016, these content files are now installed in this directory:

- `c:\\ programdata\e-on\ software\`

This is a hidden directory. You will need to have the capability of seeing this directory and its folders turned on the the Windows Folder option of the Control Panel.

Macintosh:

- Program files: `Applications/`
- Configuration files: `Users/[User\ Name]/Library/Application\ Support/e-on\ software/`

Prior to the 2016 version, the VUE content files were installed in the following directory:

- `Users/[User\ Name]/Documents/e-on\ software/`

With 2016, these content files are now installed in this directory:

- `/Users/Shared/`

VUE will also create an empty content directory structure for your own personal conyent in a location that you can designate during installation. By default, this is created on the PC in the `/Users/[Your User Name]/Documents/e-on software/` directory. On the Mac, this directory with its set of folders is created in the `/Users/[Your User Name]/Documents/e-on software/` directory.

You can change the destination folders for the Program files and Content files at installation time.

You can also change the [Vue Content Folder](#) after installation.

[Installing Vue Infinite and xStream](#)

[Installing the Artist products](#)



Installing Vue Infinite and xStream

Installing VUE Standalone

During installation, you will have to decide whether you want to install VUE as a node-locked, or as a floating license. Nodelocked is used for single computer installations, like personal computers and laptops. Installing a floating license requires the [e-on License Server](#) to be installed. If you decide to install a nodelocked license, you will have to enter the product's serial number during the installation. This is the number, in the form of "VUE99INF-aaaaa-aaaaa-aaaaa-aaaaa-aaa-xxxxx", that is written on your registration card or your notification email (where "a" represents a letter, and "x" represents a digit).

Note:

This number is confidential, and should not be communicated to third parties.

You have the option of overriding the location of the installation of the program files and the content files. VUE will run from any location you specify on your hard drive(s). If you enter a different location, and change your mind, there is a Reset button at the bottom of this dialog that will reset the location fields back to the default installation locations.

Also, during the installation process, you will have the choice of installing all of the software, or only parts of it. Since only what is necessary is actually installed on your hard-drive, we recommend you choose the Typical installation mode.

For Floating Licenses

Be sure you have chosen **Floating** during the VUE installation. If you have installed VUE as nodelocked and later decide to run it as a floating license, you will have to do a complete uninstall of the current installation and a reinstallation of the program, being sure you have chosen **Floating** as the option.

1. Run the update on your VUE.
2. Add your VUE serial number to the LicenseServer and activate it from the license Server as with previous versions.
3. Restart VUE. You should receive your VUE license from the License Server.

Installing VUE as a Plugin

Installing VUE xStream can involve both installing VUE as a standalone product or as a plugin in a host application.

For more information, refer to the general introduction to [VUE Pro as a Plugin](#).

xStream



Technical Support

Included Support

If you experience difficulties installing or using the software, the first thing we recommend is that you visit our website and read through the frequently asked questions to see if there is already an answer to your problem in there.

If not, please visit www.cornucopia3d.com/forum.

As a registered client of e-on software professional products, you benefit from the following standard support services:

- Complimentary 30 Day Maintenance,
- Phone-based installation troubleshooting for 90 days following your purchase,
- **Knowledge Base**
- Web-based **Technical Support**,
- **Registered User Forums**,
- Free **Software Updates**.

Also, please visit the forums at www.cornucopia3d.com/forum.

Complimentary 30 Day Maintenance

All purchases of VUE include 30 days of complimentary maintenance.

With the complimentary 30 Day Maintenance Plan, you receive:

- Unlimited, priority web-based technical support,
- Access to EEF releases (Expedited Engineering Fix),
- Free upgrades during the maintenance period (access to pre-release versions not included).

The 30 Day Maintenance Plan is automatically added to your account after the first activation of your product.

If you decide to extend this 30 Day Maintenance by subscribing to a Standard Maintenance Plan, the yearly maintenance period will extend the 30 days, providing a total of 13 months coverage.



Standard Maintenance

This is an annual maintenance contract that includes a number of benefits:

- Free upgrades during the subscription period, including free upgrades to all .5 and full versions,
- Unlimited, priority web-based technical support,
- Access to EEF releases (**Expedited Engineering Fix**),
- Access to “Maintenance Only” forums.

When You Can Get Under Maintenance?

Since the annual subscription period starts from the date of activation of your copy of VUE, we recommend that you purchase maintenance together with your product, or that you add maintenance within **30 days** of purchasing. This will offer you 13 months of Maintenance coverage.

If you don't purchase maintenance during this time, you will not be able to get your product under maintenance until you purchase an update to VUE.

Activating Your Product

Activating Your Product

You must activate your copy of VUE Infinite/xStream before you can operate the product to its full extent. It is possible to use the product without activating it, but it will then operate in “Personal Learning Edition” (PLE) mode (renders will be watermarked and files created with the product will only be useable on the computer that was used to create them). Of course, these restrictions will disappear as soon as you activate the product (you will have to resave the files that were created in PLE mode in order to use them on other systems).

If you are running VUE as a node-locked product, when you open VUE for the first time after installation, the *User Registration* box displays. Enter your name, Company name (optional) and your Serial number.



A screen displays with three options:

- You can opt to automatically activate online. This option attempts to connect to your e-on software account and activate the product there.
- You can opt to activate your product manually. Use this option if the machine you are running VUE on is not connected to the Internet or if you have already received your Activation key file.

A screen displays with your INST- code. If you need to activate on another machine you will need this code to do it. If you have received an activation key file, you can point to it now to activate. Or activate online.

- The third option is to run this program unactivated but it will then operate in “Personal Learning Edition” (PLE) mode (renders will be watermarked and files created with the product will only be useable on the computer that was used to create them). Of course, these restrictions will disappear as soon as you activate the product (you will have to resave the files that were created in PLE mode in order to use them on other systems).

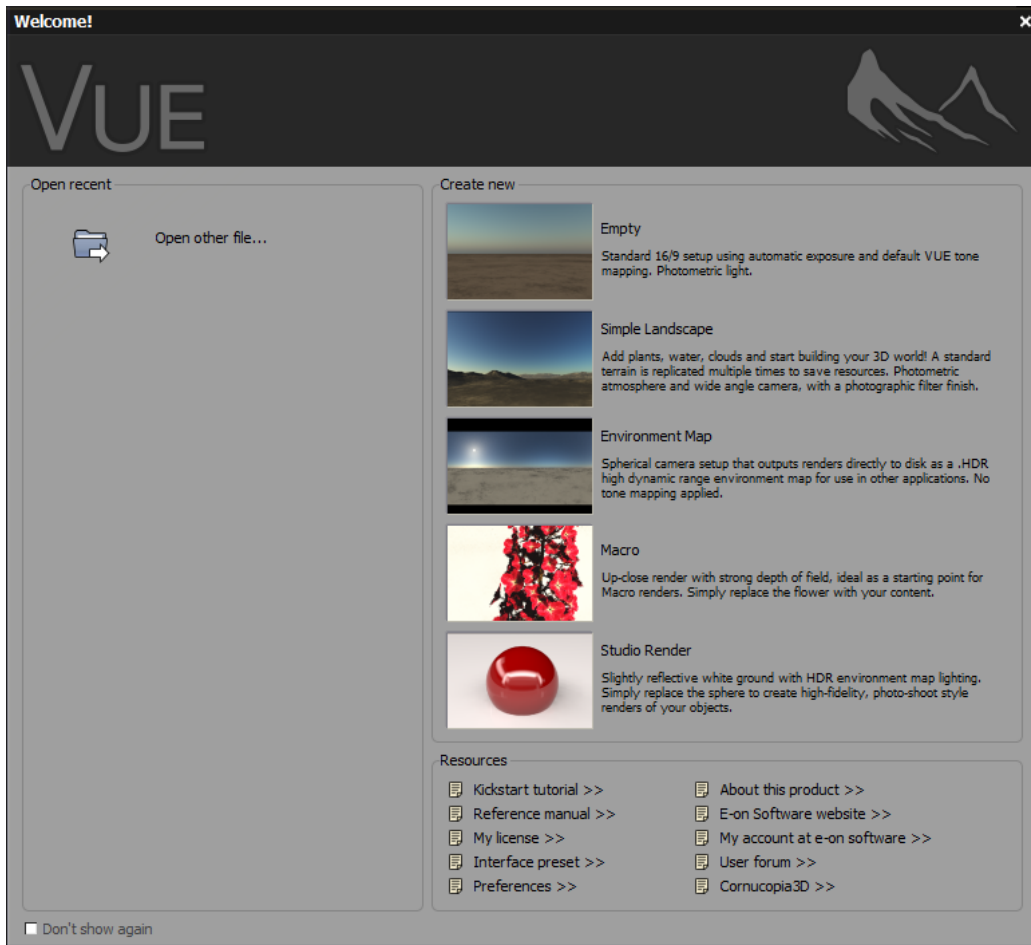
If you go to your account page to activate, you will receive an email with an activation key file attached. Place this key file in a directory that gets backed up frequently and is safe from accidental erasure. The next time you open VUE, click to locate the activation keyfile, locate it on your hard drive and your product will be activated.

If you are running VUE as a floating license, instead of seeing the *User Registration* box when you open VUE, you will be asked to locate the *License Server*. Your VUE license will be activated from there. For more information about floating licenses and the VUE License Server, refer to [E-on_License_Server](#).

Welcome Dialog

The *Welcome Dialog* appears when you open VUE.





Startup Screen

It displays the most recent files you've worked on for selection (on the left side) or you can browse to open another file.

On the right is a list of preset files to choose from for starting a new image file.

There are links to various resources:

- **Kickstart Tutorial:** provides access to the beginning tutorial for VUE.
- **Reference Manual:** accesses the VUE wiki
- **My license:** displays the installation (INST-) code for this product.
- **Interface preset:** you can select a keyboard interface for VUE that matches other



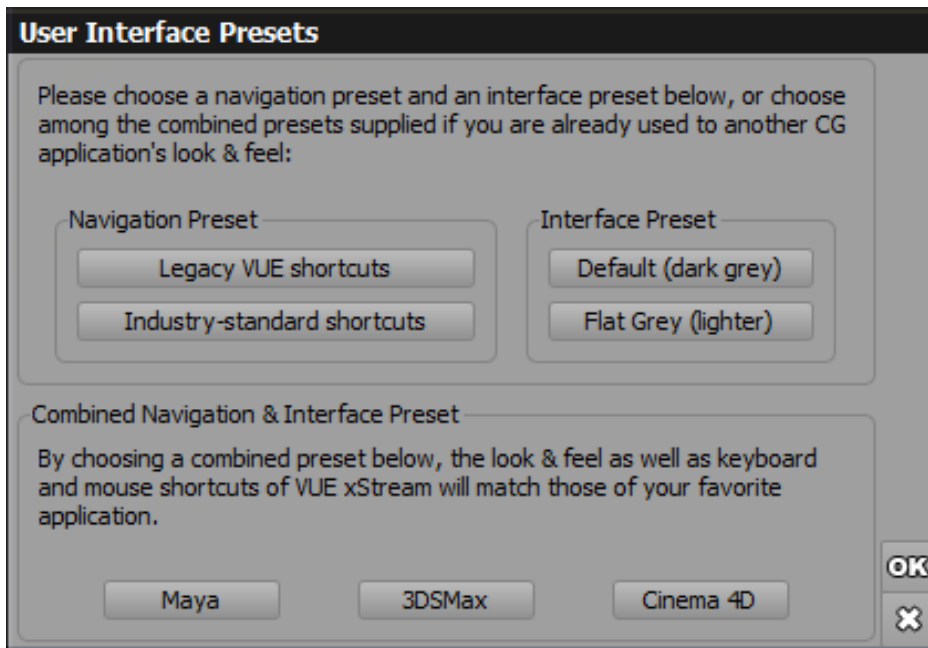
VUE 2016 Infinite & xStream — Reference Manual

software products like Max, Maya and others.

(see *Feeling at Home* section below)

- **Preferences:** opens the VUE Options page where you can set your preferences.
- **About this product:** displays product name, registered users name and the build number of the most current update applied.
- **e-on software Website:** displays the front page of the e-on website.
- **My account at e-on software:** opens your account page at e-on.
- **User Forum:** takes you to the Registered Users Forum at the e-on website.
- **Cornucopia 3D:** takes you to the first page of the www.cornucopia3d.com website.

If you don't want to see this dialog next time you launch VUE, check the **Don't show this dialog again** option at the bottom of the dialog. If you decide you want to see it again, you can access the screen from the **Help>Show** startup panel option on the VUE menu.



User Interface Presets Dialog



Feeling “At Home”

If you are used to working with other 3D software packages, you have probably grown very accustomed to the particular keyboard shortcuts of that package. Learning a whole new set of keyboard shortcuts is probably not something you are looking forward to! This is why we have implemented the ability to set the interface to match as closely as possible, the interface of other popular 3D packages.

The first time you launch VUE, a dialog will popup letting you select the type of interface you are familiar with so that VUE’s interface can be adjusted to make you feel “at home” as much as possible. Simply select the type of interface of your choice and press OK. The keyboard shortcuts and color scheme of the VUE interface will be adjusted to match your selection.

Checking Video Board Compatibility

Because VUE makes extensive use of advanced OpenGL features, low quality video boards with obsolete drivers will not perform appropriately.

To ensure that the operation of your software is as smooth as possible, a video board checking mechanism has been implemented in VUE. What this does is check that your video board and driver are on our Qualified Hardware list.

If a potential problem is detected with your video board, either because your video board doesn’t appear on our list of Qualified Hardware, or because you are using an obsolete driver, a warning will appear explaining what the problem is. If the problem is severe, you will be given the option to disable OpenGL hardware acceleration. It is highly recommended that you either ensure that your video board meets our standards, or that you disable OpenGL hardware acceleration. Failing to do so may result in highly unreliable performance.

We update our Qualified Hardware list regularly. Please help us by taking the time to report any issues or driver fixes to our [feedback center](#).



Updating VUE

No software is ever perfect. This is why e-on software regularly releases software updates through its website. These updates can provide new features as well as bug fixes. Keeping your software up-to-date by regularly downloading and installing these updates is recommended for optimal performance.

Automatic Updates

Because keeping track of these updates can be a time consuming process, VUE features an optional automatic updating technology. This technology requires that VUE have access to the internet in order to connect to the e-on software website to periodically check for new updates. The very first time you launch VUE, a popup will appear asking you whether or not you wish to authorize such connections. If you refuse, or if your computer is not connected to the internet, the automatic update feature will be disabled.

You can check for updates anytime on demand by selecting the menu command **File | Check For Update**. However, this again requires that VUE have access to the internet. If this is not the case, or you do not wish to authorize such connections for any reason, you will have to visit the Software Updates page of our website (**Support | Software Updates** from our site menu) and manually retrieve the updates there.

If automatic updating is enabled, the application will periodically attempt to connect to the e-on software website to see if any software updates are available. If a new update is found on the e-on website, a short description of the update will appear so that you can decide whether or not you wish to install the update. If you decide to install the update, VUE will download and install the update automatically. A backup of your software will be made so that you can remove the update if you decide that you don't want to keep it. Please wait until the process completes before continuing.

If you don't want VUE to check for new updates automatically, uncheck the corresponding option in the new update prompt.



Canceling Updates

If for any reason, you decide that you do not want to keep the last update, you can uninstall it and restore the previous version by selecting the menu command **File | Cancel Last Update**.

After a few minutes of processing, the update will be removed and the previous version restored.

NewCow™ Network Updates

VUE's updating technology features an advanced network updating component called *NewCow™*. This unique technology automatically takes care of updating the network rendering nodes (*RenderCows*) installed on your network. If a software update is available for the *RenderCows*, it will be downloaded together with the main application update, and installed on the render nodes on demand. Please click [Updating_RenderCows](#) for further details on the *NewCow™* technology.

The latest rendernode update can be installed in the same way as the main application update.

Importing Files from Previous Versions

VUE features an automatic import filter that converts files created with Vue d'Esprit 4 and later, so importing these files is possible. However, please understand that due to the major changes in the render engine, the Terrain Editor and spectral lighting, files from previous versions may not produce identical results as in the version where they were created.

Note:

Plants created with Vue d'Esprit 4 or older versions of VUE cannot be edited and do not react to wind in later versions.



Importing Locked Content

If your copy of VUE is an upgrade from a previous version of VUE and you acquired locked content for that previous version (they are identified as being 'Locked to your license(s)'), you can use that content in VUE natively. If this is not the case (e.g. you did not purchase an upgrade), you should re-download your locked content from Cornucopia3D after activating your copy of VUE.

Memory Management and Fault Protection

Memory Management

VUE features advanced memory management technologies such as texture and geometry virtualization. When the amount of RAM required to process a specific scene exceeds the amount of physical RAM available on your system, your system will automatically store and retrieve data to disk (virtualization).

Virtualization happens automatically, without any user action being required. The only side effect of virtualization, is that the application will become more and more sluggish as it becomes increasingly large.

If you find that application response is becoming extremely slow, your memory may need reorganizing. Select the menu command **File | Purge Memory** to automatically reorganize the system's memory and ensure memory defragmentation and cleaning up of any data that is not immediately required (for instance, if you delete a very large object, this object stays in memory in case you decide to undo this operation – by purging the memory, the object will be removed from RAM and stored on disk, until it is completely removed when the delete operation goes out of the undo list).



Fault Protection

Fault protection is the generic term that covers the different technologies that have been implemented in VUE in order to avoid as much as possible application crashes and loss of data, as well as improve the behavior of the application on particular setups.

The most noticeable effects of these fault protection technologies will be some warnings when your system is getting low on RAM, as well as automatic scene saving when the system returns a memory allocation failure. The system may return such an error although system monitoring tools indicate that there is still a lot of free memory available in the system – this type of error is caused by what is known as memory fragmentation, and is, generally speaking, a result of the fact that, unlike most other applications, 3D applications often require massive chunks of memory to operate. The risk of a memory allocation errors occurring increases if your total memory consumption exceeds half of the total memory available.

Whenever VUE's fault protection technology intercepts such a memory allocation error, it will attempt to save the current scene. Usually, the application will crash very shortly after saving the scene (if not during saving...). Next time you restart the application, VUE will automatically detect the backup scene and offer to reload it. However, because the system was in a very unstable state at the time of saving this scene, you should be advised that it may not be valid and could lead to another application crash.

OpenGL Crash Interceptor

There is a system in VUE that will try to intercept OpenGL crashes and make a backup of the scene prior to signaling the user of the fault.


When that happens, the user has two choices: either to restart Vue or to try to continue working with VUE. The second choice exists to allow for possible editing of the suspected problem area before saving the scene and eventually restarting Vue.

During this time, all OpenGL zones will be grayed out.

When VUE is restarted, the user is given the option of loading the backup scene that was created just after the crash. Refer back to the Memory Management section for more information.



Embedded Error Reporting



E-on Software Error Report

A software problem has caused Vue 7 Infinite to close unexpectedly. We apologize for the inconvenience this has caused you.

An error report was generated. By sending us this report, you will help us analyze the cause of the problem. A support ticket will be automatically created for this report on the e-on software support center.

E-on software technicians will send notifications about this ticket to the e-mail address listed in your account (the account that was used to register/activate this product).

If this product is not registered/activated, no notifications will be sent. Please review the support ticket online (the ticket ID will be provided after the error report is successfully sent).

[Privacy Policy](#) [View Report Details](#)

Error Description

Title:

Step-by-step description:

Error Report

If Vue should encounter an error, it will bring up the E-on Software Error Report dialog. By filling out and sending the report, a Crash Report is automatically created for this error on the e-on software support center. These reports provide vital feedback to our maintenance effort. If this product is not activated, no notification will be sent. While all of these reports are read and routed to the development group when necessary, you probably will not receive any reply from support technicians to this Crash Report.

Compatibility Mode

If the application crashes for any reason other than running out of memory, the next time you restart it a message will appear offering you to enable the Compatibility Mode. Compatibility mode has been designed to minimize the risks of incompatibilities between the application and the particularities of each user's setup (video board driver incompatibility, conflicting applications, etc.). What compatibility mode does is disable the features in Vue that may cause the most problems (typically, advanced multi-threading and previewing options).

Note:

Don't enable compatibility mode if you think you know why the application crashed. Enable it only if you find that the application crashes randomly without any apparent



reason.

Following is the exact list of features disabled by the compatibility mode:

- Hardware acceleration of OpenGL in 3D views: VUE will use the slower, but 100% compliant OpenGL mode instead,
- Background draw thread: some video boards have issues with this; also, this feature requires a lot more OpenGL resources,
- Decimated mesh previews: decimating meshes requires a lot of system resources,
- Boolean and Metablob previews: like the decimated mesh previews, this requires a lot of system resources,
- Automatic rendering of material previews: having too many threads running simultaneously can sometimes cause issues with certain setups,
- Automatic scene preview: this requires a lot of system resources.
- Multiple undo-redo operations: this also requires a lot of system resources.

Once you have enabled the compatibility mode, you can re-enable these feature one by one using the [Options dialog](#) until you find the feature(s) that cause(s) problems with your particular setup.

Downloading the Extra Content


Prior to the 2016 version, the VUE program was distributed with the installation files and a .zip file of Extra Content. The Extra Content is no longer distributed, but the individual files are made available by simply selecting them in the browsers.

For example, if you wish to open the scene *Room Radiosity* in the /Samples 3 folder, select File>Open to open the scene browser and select *Room Radiosity*. The file will be downloaded and open in VUE.

If you wish to keep a copy of this file on your hard drive, you would have to save it.



Interface Overview



A snapshot of the main software interface

- | | |
|--------------------------------------|--------------------------------------|
| 1. 3D Views | 8. Object properties panel |
| 2. Selected objects are drawn in red | 9. Material of selected objects |
| 3. Resize handles | 10. Camera Control Center |
| 4. Rotate handles | 11. World Browser |
| 5. Triangle means unfolded icon | 12. Selected objects are highlighted |
| 6. Dull icons are disabled | 13. Camera view |
| 7. Square dot for double action icon | 14. Status bar |

The VUE interface was designed with three goals in mind: ease-of-use, workflow and clarity.

This is why the interface is designed “in layers”. What this means is that every user works









with the program at his/her own level of proficiency. Although the initial impression may be that of excessive simplicity, as the user works along and delves in deeper, he will realize – de facto – how many options and possibilities are truly available but couldn't be seen at first glance. Facing all these possibilities right from the start would have been overwhelming.

The final layout dedicates as much possible space to the *3D Views*, because this is where the user spends most of his time. Three large panels are available on the right, displaying all useful scene and object information within mouse reach, thus allowing for quick navigation inside the scene.

Dialog Bar

On the right edge of most dialogs you will notice a bar that contains a number of icons. This bar is called the dialog bar. The icons in this bar depend on the dialog:

-  **OK:** click this icon to close the dialog and accept the modifications.
-  **Cancel:** click this icon to close the dialog and cancel all modifications.
-  **Help:** click this icon to display contextual help information.
-  **New:** click this icon to reset the settings in the dialog.
-  **Load:** click this icon to load settings from the disk.
-  **Save:** click this icon to save settings to the disk.



Customizing the Interface

Selecting an Interface “Model”

The first time you launch Vue, a dialog will pop up asking you how you want to set up the program’s interface. If you are used to working with another 3D application, you will have the opportunity to automatically customize the Vue interface so that your favorite keyboard shortcuts work the same, the colors of the interface are familiar, and the overall feel of the application is as close as possible to your favorite 3D application.

You can change the interface “model” anytime using the [Load interface preset](#) button in the *Options* dialog.

Further Customization

You can modify the colors of the interface using the [Interface Color Editor](#).


You can further customize the behavior of the interface and change [keyboard shortcuts](#) using the *Options* dialog.

Docking

[Toolbars](#) and dockable dialogs may be rearranged at will in the [main window](#). You can tell if a dialog is dockable by just moving it in the interface. If it is dockable, available docking areas will show in the interface. Freely floating windows can be moved by left-clicking on their **upper bar** and dragging. While dragging, trapezoids show up in the other windows. You can dock anywhere within the trapezoid. Moving the window and dropping it on such a trapezoid will dock it there.

Trapezoids sometimes have a gradient of lighter borders that allow you to dock along side a group of areas. The nearer the border the bigger the group.





Docking Node Options Toolbar to New/Load/Save Toolbar

Undocking

Windows that are already docked have a **grip bar** on their left or top edge. Left-click on the grip and drag it to undock the window.

If you are satisfied with your current window organization, you can lock it using the **Lock Workspace** command of the **Display | Workspace** option of the Vue menu. The lock workspace will prevents you from docking or undocking inadvertently, but you will still be able to resize any areas using the moveable splitters. In the **Display | Workspace** menu you can load default and saved docking presets as well. Docking configuration is preserved from one Vue session to another.



Resizing the ViewPorts

All of the viewports in the workspace can be resized and moved around in the workspace. For example, you might decide to move the Top View to the lower left. You can rename the viewports by using the **View Display Option** icon in the title bar of each viewport. You can resize the Object properties, Camera Control and World Browser areas as well.

Once you have your workspace defined the way you like, use the **Display | Workspace** command from the menu to save the layout and lock the workspace. You can save several different workspace setups and change between them if you wish.

3D Views

These are the large windows sitting in the middle of the interface. This is where you build your scenes. By default, these windows display four different views of your scene: the *Top view* displays your scene as seen from above, the *Front view* displays the scene as if looking at it from straight ahead, and the *Side view* as if looking at it from the right. Since all of these 3 views are orthogonal projections, they are also known as the orthogonal views. They are ideally suited for moving, rotating and sizing objects. The last view, which is the bottom right one, is the *Main camera view*. It displays a preview of your scene, as seen from the camera. If you move the camera around, you will notice that this view changes interactively. This view enables you best to handle picture framing and composition.


When a *Spot Light* or *Quadratic Spot Light* is selected, the *Main camera view* can also be used to look at your scene from this light source origin and adjust the light.

Active View

At any given time, there is one of these views that has a highlighted title bar. This is the **Active** view. Keyboard operations (e.g. nudging with the arrows) will always be directed to the active view. Simply click in a view to activate it.



Maximizing / Resizing Views

You can maximize a 3D View anytime by double-clicking on the title bar of the view. If you don't like the 4-view layout, you can maximize the main view, and do everything there. It's a matter of personal preference. You can also maximize a view by clicking the **Toggle Current View / Four Views** icon () in the top toolbar, or by selecting the corresponding menu items **Display | Toggle Current View / Four Views** and **Display | Maximize | xx View**. To toggle back to the 4 views, simply do this again.

Views can be moved around by clicking and dragging them with the right/Ctrl mouse button. You can also zoom into or out of them by pressing Control while you drag. Furthermore, the Main camera view can be panned by Shift dragging it.

When the mouse cursor is on top of the window separators, the cursor changes indicating that you can click and drag to modify the proportions of the views. The ratio between the various views is indicated in the Status Bar as you drag the mouse. Drag the mouse outside the view ports to restore previous settings.

Full Screen Mode


By hitting **Alt-Enter**, you can maximize the view ports so that they fill-up the screen, providing all available space to the view ports. In this mode, the menu bar and the other tool-bars are hidden, but you can revert to the standard layout by hitting **Alt-Enter** again. When Full Screen mode is activated, all menu commands can be accessed via the popup menu.

ToolTips

If you let the mouse cursor stand still over the interface, a ToolTip will pop up, telling you what is under the cursor. This works for icons, controls in the *Object Properties* and *Timeline* panels, as well as for objects in the views. The latter is particularly useful to find out which object will be selected when you click.




View Display Options

The  button appears in the title bars of all the views. Clicking on this button displays a popup-menu that lets you customize the behavior of that view.

The options in this menu are:

- **Maximize/Restore:** select this command to maximize the view when all views are displayed, or to switch back to 4 views when a view is currently maximized. This is the same as double-clicking on the title bar.

Following this option is a list of view types that you can use to modify the type of the view. A check appears in front of the current view type:

- **Top View:** select this option to make the view display your scene from above,
- **Front View:** select this option to make the view display your scene from the front,
- **Side View:** select this option to make the view display your scene from the side,
- **Main Camera View:** select this option to make the view display your scene as seen through the currently active camera,
- **Perspective View:** select this option to make the view display the scene in perspective, from a viewpoint that is not related to the active camera. If you select this view, several different icons display on the title bar. () The icon to the far right toggles between perspective and main camera view. The icon second from right has two functions: **Copy current view to activate camera view** or **Store current view in new camera**.

Options related to the display quality of the objects in the view:

- **Wireframe Box:** the least detailed, but also the fastest,
- **Filled Box:** same as wireframe box, except the box is solid,
- **Wireframe:** useful when you want to see through objects,
- **Flat Shaded:** almost as good as smooth shaded, only a little bit faster, and
- **Smooth Shaded:** the best quality, and also the slowest. This is the default.
- **Interactive Ray-Tracing:** uses the Ray Tracer renderer.

This setting acts as a maximum display quality for the objects in the view. For instance, if a sphere has a “Wireframe box” display quality setting, and the view has a “Wireframe” display quality setting, the sphere will be displayed as a wireframe box. Of course, since the [display quality options](#) are defined independently for each view, you can have a view in Wireframe and another one Shaded.

The second set of options are relative to the view’s fog settings:

- **Show Fog in View:** this option turns OpenGL fog on for the view. Fog is useful



to give an additional information on the distance to objects. When this option is selected, the fog density options become available.

- **Density From Atmosphere:** when this option is selected, the density of the OpenGL fog is automatically matched as closely as possible to that of the fog in the scene.
- **Adjust Fog Density:** this option is only available when the Density From Atmosphere options isn't selected. It lets you adjust the density of the fog manually. If you select this option, the mouse cursor will turn to an Up-Down arrow. Click in the view and drag the mouse up or down to increase/reduce fog density.
- **Refresh sky:** this option refreshes the sky if you've made any changes.

The next group contains one or two options depending on whether the current view is the main view or one of the orthogonal views:

- **Show Only Objects From Active Layer:** this option will hide all objects from the view that are not in the currently *active layer*. Although the objects don't appear in the view, they are still in the scene and will be rendered (unless they are disabled from rendering) as any other object. This option is handy when your scene becomes complex because it enables you to only display in the views those objects that are part of the layer you're working on at a given time. Using different *preview options* helps keep your scene sorted by hiding or displaying items differently.
- **Frame Guides...:** this option is only available in the *Main camera view*. If you select this option, the *Frame Guides* dialog will popup letting you configure safe areas and visual guides.
- **Zoom Extents Selected:** this option is for use with the Space Navigator.

After these are the lighting options:

- **Light From Scene:** if this option is selected, the first 8 lights in the scene are used to light-up the *3D view*. If it isn't selected, the *3D view* gets its light from a source placed on top of the viewer's left shoulder.
- **Shadows:** this option is only available in the *Main camera view*. If this option is selected, objects placed above the ground cast a vertical shadow onto it. This is useful because it gives a better idea of the altitude of the objects above the ground.

The Display options menu of the main view offers a few additional options:

- **Show Last Render in Back:** this is a truly fantastic option when it comes to finding/placing an object precisely. What it does is draw the objects in the view on top of the last rendered picture. So you get the OpenGL preview of your objects on top of the real render! This option can significantly slow down the 3D display on computers that are not equipped with an OpenGL optimized video board...
- **Show Only Selected Objects:** this option is only available when the above option is enabled. Because objects are usually displayed in OpenGL as smooth shaded, they will cover-up the picture that is in the background. In order to see the background,



you'd have to turn all objects to wireframe, or hide them manually. What this option does is hide all objects that are not selected from the *Main View* so you can see the render in the background without having to do the above.


- **Display 3D View:** select this option to display the OpenGL 3D view.
- **Display Color Channel:** select this option to display the color channel of the last rendered picture.
- **Display Alpha Channel:** select this option to display the alpha information in the last rendered picture.
- **Display Depth Channel:** select this option to display the depth information in the last rendered picture.
- **Non Photo-Realistic Render:** select this option to enable this rendering feature.
- **Display Multi-Pass, Masks and G-Buffer:** select this option to display the different Multi-Pass render components, masks, or the contents of the G-Buffer (this option is only available if you have enabled rendering of [Multi-pass](#), [mask](#) or [G-Buffer](#) information).

This also allows you to display any of the [diagnosis buffer information](#).

Saving Default Viewport Configuration

Once you have customized your 3D views (by choosing display options explained above, or activating Interactive Path Tracer in a given view, for instance), you can save your configuration as default, using the menu entry **Save Current Viewport Configuration As Default**, in *Display* menu. The viewport configuration is also saved within your scenes. If you load a scene which uses a different viewport configuration, VUE will offer you to choose between the scene's configuration and your default viewport configuration.


Quick Render

Click the **Quick Render** icon () in the title bar of a view to do a quick render of the view. The rendering will always take place inside the view, and will always be performed using the [internal renderer](#). If the view is an orthogonal view, an orthogonal camera will be used and all atmosphere effects will be disabled.

By default, the quick render is done using the **Preview** [preset render quality](#). However, if you long-click on the icon, a menu will appear letting you select the preset render quality to be used for the quick renders.




Saving Pictures

When you have finished rendering a picture you can save it to disk by clicking the small **Save Displayed Picture** icon () in the view's title bar, or use any of the other methods to [save renders](#).


Channels

When it renders a picture, Vue generates three different types of information for each pixel in the picture: the color of the pixel (**Color channel**), whether an object is visible in that pixel (**Alpha channel**), and the distance to the closest object at that pixel (**Depth channel**, also known as Z-Buffer).

This information is known as channels. All three channels of information are available when the rendering of the picture is complete. This extra information is extremely useful when you want to do some post-processing on your pictures (e.g. compositing, or applying Photoshop filters that are distance sensitive, such as blur).

To display each channel after rendering a picture, use the group of buttons that are at the right hand end of the *Main view's* title bar () , or select the corresponding menu option from the **View Display Options** menu:



Color channel: click this button to display the color information in the picture. Long-clicking changes the button to  and opens the Render Display window.



Alpha channel: click this button to display the alpha information in the picture.




Depth channel: click this button to display the depth information in the picture.




Non Photo-Realistic Render: click this button to enable this rendering feature. These buttons are not available until you have rendered a picture. To switch back to *3D view*, click once on the title bar.



Multi-Pass, Mask and G-Buffer

The Multi-Pass, Mask and G-Buffer icon () expands into a menu that lets you display the different render passes, object or material masks as well as G-Buffer channels.

These Multi-Pass, Mask and G-Buffer viewing options are only available if you have rendered a picture with the [Multi-Pass](#) or [G-Buffer data](#) collection options enabled.

If you select a G-Buffer channel, you can change the G-Buffer layer using the  arrow icons. The number in between the arrows indicates the rank of the currently displayed layer. Because there is only one layer of information defined in Multi-Pass and Mask renders, these options are not available when displaying multi-pass render components of masks.

This icon is also used to display the Diagnosis Render Buffer information. To render with this option, you need to set this on the *Render Options* panel (within the **Render What** section). Unlike the other buffers, this information is available for all renders, even **Preview** quality.

To the top and to the left of these views are two toolbars. Before taking a closer look at these, we would like to introduce you to two special types of icons available in VUE.

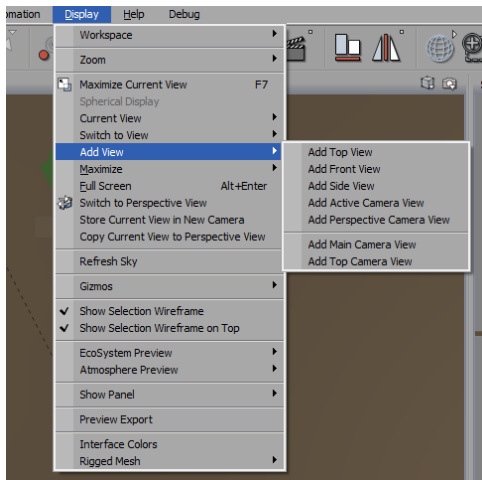
Additional ViewPorts

Previously, the number of OpenGL views inside VUE was limited to four. It is now possible to have up to 16 views at the same time.

There are two major changes concerning camera handling. Until now, the perspective view was unique to the scene. Now it's a camera specific to a viewport. This means you can have up to 16 perspective cameras in a scene, not only one. Before, you could only take the viewpoint of the active camera. This meant that when you had several cameras in your scene, you could only ever view one at the same time. You are now able to display the point of view of whatever camera you choose, additionally to the active camera.




Adding a viewport



Additional ViewPorts – Selecting a ViewPort

To display an additional viewport



Additional ViewPorts – Displaying an Additional ViewPort

Open the Display menu and choose Add View. Then, in the submenu, choose the type of viewport you want to add.

Displaying a camera that is not active

To make a viewport display a camera from the scene that is not the active one, just open the options of the viewport and choose a camera.

In the image above, choosing one of the menu entries between “Main Camera View” and “Camera 5 View” will make the viewport display the chosen camera, regardless of the active camera. Having duplicate Perspective Views Like always, if you want to switch to




the Perspective View, you have lots of possibilities. One of them is to open the viewport options and choose “Perspective View”. To profit from the 16 possible perspective cameras, you just have to switch several viewports to their “Perspective View” mode. Then, naturally, their position and orientation remain independent.

Double Action Icons

Some icons in the toolbars can perform differently, depending on the way you click on them. If you click on them, they will perform the default action, as expected. However if you do any of the following, they will perform an alternate action:

- Click with the right mouse button (Ctrl click on Mac),
- Shift (or Control) click,
- Click and maintain the button depressed until the icon changes to the alternate display.


Double action icons are identified by a small white square dot on the right border (e.g. the Render icon .

Unfoldable Icons

Other icons in the toolbars can have multiple actions, also depending on the way you click on them. If you click on them in the normal way, they will perform the default action, as displayed in the icon. However, as with the Double action icons, if you do any of the following, they will unfold to display a set of alternate actions:

- Click with the secondary mouse button,
- **Or** click and maintain the button depressed until the icon changes to the alternate display.

To select an alternate action, drag the mouse through the unfolded icon, onto the requested action, and then release the button. The default action of the unfoldable icon becomes the last selected action.










Unfoldable icons are identified by a small white triangle on the right border (e.g. the Sphere icon unfolds to reveal other Primitives .



Top Toolbar



This toolbar is placed at the top of the interface, immediately below the menu. It provides shortcuts for most common operations like file manipulation or undo. It contains the following icons:

-  **New...:** click to open a new file.
-  **Open...:** click to select a file to open from the Scene Browser.
-  **Save / Save As...:** click to save the current scene. Right-click to open the Save As dialog. [More information about save scene modes](#)
-  **Cut:** click to cut the selected object/plant/terrain.
-  **Copy:** click to copy the selected object/plant/terrain.
-  **Paste:** click to paste what was previously cut or copied.
-  **Duplicate / Scatter/Replicate Objects:** click to duplicate the selected object/plant/terrain. Right-click opens the Scatter/Replicate Objects screen for making multiple copies and defining placement.
-  **Undo... / Undolist:** click to undo the previous action. Right-click to display previous actions that can be undone.
-  **Redo... / Redolist:** click to redo the previous action. Right-click to display previous actions that can be redone.





Record macro: click to record a macro.



Play a macro or tutorial: click (or right-click) to display a dialog with available macros or tutorials listed.



Atmosphere Editor / Load Atmosphere...: click to open the Atmosphere Editor for the atmosphere currently loaded. Right-click to open the Atmosphere Browser to select a new atmosphere.



Edit Object: click to display the Polygon Mesh Options dialog for the currently selected object.



Paint EcoSystem: click to display the EcoSystem Painter dialog to set up and create an EcoSystem by painting.



Show Material Summary: click to display all the materials used in the current scene.



Show Timeline / Animation Wizard: Click to display the Timeline. Right-click to open the Animation Wizard.



Alignment Tool: only available if two or more items are selected in the scene. This opens the Align Tool dialog for you to define alignment.



Flip Vertical Axis / Flip Horizontal Axis: click to flip the vertical axis of the selected item. Right-click to flip the horizontal axis of the selected item. Re-click to undo.

- **Select by ...:** right-click to unfold this icon giving you three options when selecting items in a scene



Select by Wireframe Color:





Select by Object Material:



Select by Object Type:

Click the icon again to cancel the selection.



Frame Selected Area / Frame All/Selected Objects: Click, then select an area to frame. Right-click to frame the selected item.



Zoom Into View: Click to zoom in closer in the view port.



Zoom Out of View: Click to zoom out in the view port.



Toggle Current View / Four Views: Click to toggle between one viewport and four. The current (selected) viewport becomes the enlarged one.



Browse Previous Render (R) / Save Color Picture: Click to browse previous render in the Render Display. Right-click to save the last render.



Select Render Area: Click to select a render area in the Main Camera View. Click again to remove the selected render area.



NPR Options: This turns on the Non-Photo Realistic options and displays the windows for setting up this option



Render / Render Options: Click to render the scene using the settings as they currently exist in the Render Options dialog. To change the render settings, right-click to display the Render Options dialog.

Clicking on empty parts of the toolbar will deselect all selected objects.























Left Toolbar



This toolbar is also known as the object bar. It provides shortcuts for creating, grouping, and selecting objects.


- **Water / Ground / Cloud / Enable Rain / Enable Snow** ().



- Sphere / Cylinder / Cube / Pyramid / Cone / Torus / Plane / Alpha plane ()
 - Text Editor ()
 - Heightfield Terrain/ Procedural Terrain / Load Procedural Terrain Preset ()
 - Create plant icon () The alternate action is **Load plant species** ()
 - Rock / Load rock template ()
 - Add cloud layer / Load cloud layer / Add metacloud / Metacloud from preset / Metacloud primitive ()
 - Planet ()
 - Spline / Road ()
 - Particle System / Particle System From Preset ()
 - Particles Effector / Directional Ventilator ()
 - Load object icon () The alternate action is **Save object** ()
 - Point light / Quadratic point light / Spot light / Quadratic spot light / Directional light / Light panel / Daylight Portal ()
 - Group Objects / Ungroup Objects ()
 - Boolean difference / Boolean union / Boolean intersection ()
 - Metablob object icon () The alternate action is **Hyperblob** ()
 - Drop object icon () The alternate action is **Smart Drop** ()
- Clicking on empty parts of the toolbar will deselect all selected objects.



Scene Information Bar




The scene information bar is the bar that sits on the right hand side of the interface. It is comprised of 3 sections:

- the [Object Properties](#) panel that displays information on the currently selected object(s),
- the [Camera Control Center](#) that displays a real-time preview render of the scene, and
- the [World Browser](#) that shows a hierarchy of all objects in the scene.

You can resize the Scene Information bar by dragging its left edge with the mouse. All the controls in the bar will be resized, including the [Preview Render](#). You can make the Preview Render larger that way.



Object Properties Panel



Object Properties panel – Nothing selected

This is a contextual panel that displays information relative to the selected object(s).

If no objects are selected, the panel is empty (see opposite).

The Object Properties panel is made of 3 tabs, identified by the following pictograms:



Aspect: this tab displays information on the visual aspect of the selected objects.




Numerics: this tab gives numerical control over the position, rotation, size, twist and pivot position of the selected objects.




Animation: this tab gives control over the animation characteristics (motion, linking, etc.) of the selected objects.

The title of the panel is the name of the selected object (or “Mixed objects”) followed by the number of selected objects inside square brackets, if several objects are selected. If



you don't need this panel, click the () button to fold it up. Click again to unfold it. Double-clicking the title bar does the same.

Aspect Tab

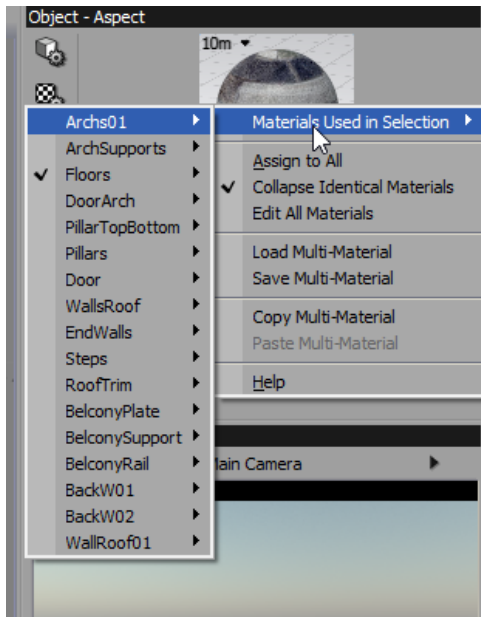


Object Properties panel – Aspect tab

The **Aspect Tab** of the *Object Properties* panel displays the material of the item currently selected in the material preview. This small box also contains a lot of information about that material and presents some material editing options.

Preview of material assigned to the object. Materials define the aspect of the object's surface. You can edit the material by double-clicking this picture.





MultiMaterials Selection

If several objects are selected, and they use different materials, or if the selected object is made of several materials (e.g. plants or imported objects), if you right-click on the displayed material, menu entries corresponding to the different materials will be grouped in a submenu.

Also, a set of arrows will appear on the bottom of the preview, together with the number of materials in the selected object. You can browse through the different materials using the left and right arrows, or you can use the bottom arrow to display a list of the different materials.

This material popup menu displayed by selecting a material will also exhibit the following commands:

- **Assign To All:** this option will replace all materials of the object with the currently displayed material.
- **Collapse Identical Materials:** when this option is active, identical materials that are assigned to different material zones in the object will only appear once in the list of materials. If you uncheck this option, all material zones will be displayed, letting you access individual material zones that have the same material assigned to them.
- **Edit All Materials:** select this option to edit all materials simultaneously. The [Material Editor](#) will appear, letting you edit the currently displayed material. All changes made to this material will be replicated in the other materials (provided the



materials are of the same type – e.g. all materials are “simple materials”, or all materials are “mixed materials”). You can also edit multiple materials simultaneously using the [Materials](#) tab of the *World Browser*.

- **Load Multi-Material:** this loads a material that has been saved as a multi-material. The preview of these materials in the Visual Material Browser displays as a mosaic of all of the materials within the one *.mat* file.
- **Save Multi-Material:** this saves all of the materials in the object or plant as a multi-material, or one *.mat* file containing several materials.
- **Copy Multi-Material:** this is used to copy a multi-material from this object/plant.
- **Paste Multi-Material:** this is used to paste a multi-material copied from another object/plant.

There is a row of icons running down the left side of the box. Depending on the item selected these icons may or may not be available.

- **Edit selected object:** this opens the editor for the type of object selected. If a tree is selected, the *Plant Editor* will open; if an object is selected, the *Polygon Mesh Options* panel displays.
- **Edit material:** this opens the *Material Editor* for the material currently displayed.
- **Edit all materials:** this opens the *Material Editor* for you to make a change to all materials for that particular object currently selected. For example, you can reset the highlighting for all materials using this option.
- **Load material:** this opens the Materials Browser for you to select a replacement material for the material currently selected.

Preview Color and Options

The **Preview options** icon displays a menu when clicked. This menu lets you define the display options of the selected object in the *3D views*. These settings are global to all the views. The first set of options relate to the quality of the preview. These options are the same as those of the views:

- **Wireframe Box:** the least detailed, but also the fastest,
- **Filled Box:** same as wireframe box, except the box is solid,
- **Wireframe:** useful when you want to see through the object,
- **Flat Shaded:** almost as good as smooth shaded, only a little bit faster, and
- **Smooth Shaded:** the best quality, and also the slowest. This is the default.

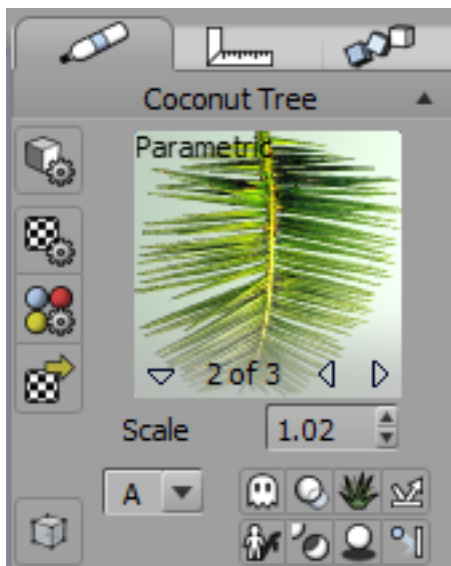
Keep in mind that these settings can be overridden by the [view settings](#). If the object has a display quality that exceeds that of the view, it will be displayed in that view using the view's quality setting.



The second set of options relate to the visibility of the object in the views. These options are:

- **Locked:** the object will be displayed gray and transparent. It isn't possible to select locked objects when clicking in the *3D views*. This is useful when you don't want to select an object, but still need to see it in the views for reference. Locked objects are displayed in gray in the *World Browser*. You can still select them there. Infinite planes are created Locked.
- **Hidden:** the object won't be displayed in the *3D views*. It won't be possible to select it either. Hidden objects are displayed in pale gray and italic inside the *World Browser*. You can still select them there.
- **Main View Only:** the object will only appear in the main view, but won't be displayed in the other (orthogonal) views. Planets are created as "Main view only". Selected objects are visible in all views, whatever their visibility setting. Removing one of the above options won't necessarily make the object appear (e.g. if it is inside a locked or hidden [layer](#)).

Selected objects are visible in all views, whatever their visibility setting.



Object Properties panel – Multiple materials

If the object currently selected has multiple materials, you'll find a downward arrow in the left corner of the picture. This will list all of the materials for this object, allowing you to select a different material, or will give you the option of collapsing identical materials or editing all materials.



Directly under the picture, the current **Scale** of the selected material is displayed. This only affects the size of the material when it is rendered in your scene. If several objects are selected that use materials at different scales, this item will be blank.

The **Preview color** is a drop-down list that lets you select the color of the object when displayed in the *3D views*. When you select a new color, the selected objects will be flashed inside the views to show the new color. Assigning a different color to your objects is good practice; since it lets you identify them more easily. By default, objects are created a dark gray, or the same color as you last selected in the list. Plants are always created green, lights yellow and planets pale blue.

Neither the **Preview Color** nor any of the **Preview Options** will affect how an object renders.

Next to the **Preview Color**, you will find a set of small icons.



Hide from render: this option lets you hide objects from rendering. When this option is set, the object will not appear in rendered pictures or animations. It will however still appear in the 3D Views (unless you explicitly hide the object from the views as well). This is ideal for creating objects that are used as scene helpers (e.g. objects that are tracked but mustn't be rendered). When an object is hidden from render, its icon is crossed out in the World Browser (you can also hide an object from render by clicking this icon in the [World Browser](#)).



Render area occluded by object in G-buffer: this option is relevant only when generating G-Buffer information. If you select this option, rendering will not stop when it encounters this object; instead, it will continue gathering information about what is behind it. Using this information, it subsequently becomes possible to remove objects from rendering during a post-processing phase or, for instance, to perform accurate motion blur effects without any missing information issues. You can also force the rendering of all occluded regions using the option in the [G-Buffer and Multi-Pass Options](#) dialog. This is the same as enabling the Render occluded option on all objects in the scene. Make sure this option is not selected if you only want to render occluded information for certain objects.



Ignore object(s) when populating EcoSystems: when this option is selected, the object will have no influence on the population of EcoSystems that are sensitive to the [presence of foreign objects](#), even if they are placed right in the middle of an EcoSystem.



Ignore indirect lighting: this option is enabled only when either the Global Il-



lumination or Global Radiosity [lighting models](#). In some cases, you may decide that the benefit of computing indirect lighting on certain objects may not be worth the investment in rendering time. This option is designed for such cases: indirect lighting will not be evaluated on objects that have this option set. The objects will however still participate in the illumination of other surrounding objects (e.g. by generating a dark halo around their bases).



Don't cast shadows: when this option is selected, the object will cast no shadow on other objects.



Don't receive shadows: when this option is selected, the object will receive no shadows from other objects in the scene.



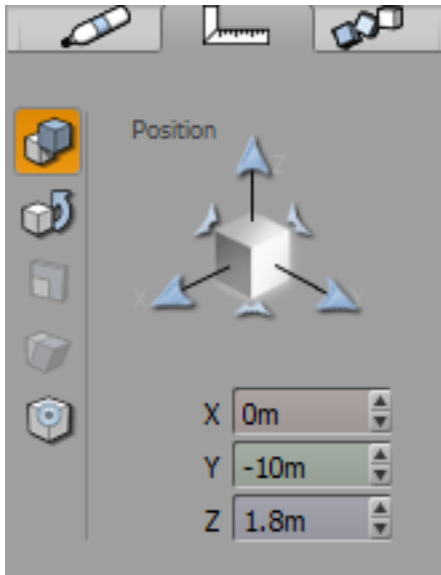
Only shadows: when this option is selected, the object will not render. Only its shadow will render.



Enable collisions: select this option to enable EcoParticle collision with the selected object.



Numerics Tab



Object Properties panel – Numerics tab

This tab lets you enter precise numerical values for position, rotation, size, twist and pivot position along each *axis*. The fields are color-coded to correspond to the 3D axes in the viewports.

This tab is divided into 5 sub-tabs: **Position**, **Orientation**, **Size**, **Twist** and **Pivot position** each represented by an icon in the panel's left icon bar. Press the required icon to display the corresponding numerical values. Drag the mouse over the axes to get visual feedback on the type of modification the object will undertake. Click-drag the controls to adjust the values, or type values in the corresponding fields.






Position: numerical values given for position are relative to the World origin, which is the center point of the orthogonal *3D Views* when you create a new scene. Clicking in between axes lets you move the object inside a plane.



Orientation: numerical rotation angles relate to the object for Pitch and Roll, but relate to the World for Yaw.







Size: numerical values for size relate to the object. They are independent of the object rotation. Consequently, using Numerics to size an object that has already been rotated is a quick and precise method. Clicking in between axes lets you resize the object along both axes simultaneously. Clicking in the middle of the control lets you resize the object globally. The **Lock sizing proportions** () is a toggle button that locks the proportions of the object; meaning that, when you change the size along one axis, the sizes of the other two axes are adjusted in such a way that the proportions of the object are maintained. The **Show actual object sizes** () is a toggle button that will display the real size of the object when selected (otherwise, internal dimensions will be displayed instead – usually not very useful, but provided for compatibility with previous versions). The **Resize around opposite corner** button () is also a toggle that indicates whether the object is resized around its center, or around the opposite corner (as when resizing objects inside the views).




Twist: Numerical values for twisting are less straightforward. Basically, they will twist one axis of the object towards another axis. This is rather difficult to visualize, so the best is to try it out. However, please understand that, due to complex matrix operations, twisting and untwisting objects in several directions may not restore the initial object conformation.



Pivot position: use the Pivot position sub-tab to display the location of the object's pivot point. The pivot point is the point around which the object will be rotated or scaled (e.g. the pivot point of a window would be at the windows hinge). Press the **Show pivot** toggle button () to show/hide the pivot in the *3D views*. The pivot point is displayed by a green handle that you can grab and drag to a new location using the mouse. Alternatively, you can enter numerical values for the location of the pivot point. You can elect to display the position of the pivot in world coordinates, or in object coordinates by using the **Relative coordinates** toggle button ()

Note:

If expressed in world coordinates, the position of the pivot point is modified when you move the object. However, the relative position to the center of the object remains the same. Press the **Reset pivot** button () to reset the pivot: a left click will reset it to the origin of the object, a right click to the center of the object. The origin is the point of the object with local coordinates (0, 0, 0), whereas the center is the apparent “middle” point of the object. Origin and center are usually

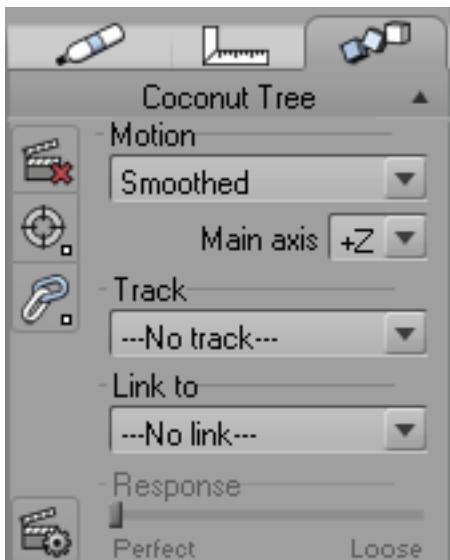


the same point except for some objects (usually imported from an outside source) where the origin can be some distance off the apparent center.

Animation Tab

This tab deals with object animation and forward dynamic hierarchy (linking and tracking).

Animation



Object Properties panel – Animation tab

To animate an object, select a type of **Motion** other than **Not animated**. Basically, [types of motion](#) let you specify how the object will react to its motion (e.g. airplanes bank as they turn, automobiles follow the surface of the ground...). If you don't know which to pick, select **Standard**, which is the standard type of animation found in usual 3D packages.



When you select a type of motion other than **Not animated**, the Animation Wizard pops-up to help you setup your animation easily. The instructions that appear on screen are relatively straightforward, but it is recommended that you read the section on the [Animation Wizard](#) to fully understand each setting.


To destroy an objects animation, select **Not animated** from the drop-down list or click



the **Forbid Animation** icon (see below).

Whenever an animated object is selected (or the selected object tracks another one – see below), the **Main Axis** drop down list becomes active. This list lets you select which axis of the object should be pointing in the direction of travel (useful only for objects that have the **Look Ahead** property, or which axis of the object should be pointing at the tracked object).



- **Forbid animation:** the forbid animation icon () is the topmost icon in the icon bar. As long as this icon is selected, the object will never become animated. If you click this icon and the object is currently animated, a warning will appear informing you that all animation data will be lost.
- **Non switchable camera:** there is an additional icon () that appears just below the **Pick link object** icon when the selected object is a camera. When selected, this icon forbids animation of the camera (as above) and also prevents camera switcher keyframes being created in an animation when you [switch to this camera](#). If the camera is already being used in the camera switcher, selecting this option will remove all switches to this camera.

Click the  icon in the icon bar to display the *Animation Toolbox* for the selected object. This toolbox lets you adjust the global behavior of animated objects.

Forward Dynamics (Linking and Tracking)

Forward dynamics is a hierarchical animation feature that greatly simplifies the animation of complex structures. Basically, it lets you link some objects to others (the link parents), thus building a hierarchy of objects. When an object is linked, modifying the link parent will automatically affect the linked object.

To link an object to another one, you may either:



- select the parent object from the **Link** drop-down list,
- press the  icon or select the menu command **Object | Pick link parent**, then click on the parent object in the *3D views* (objects will be highlighted as the mouse passes over them), or
- press the  icon or select the menu command **Object | Pick link parent**, then click on the parent object inside the *World Browser*. This last method is the only one that lets you link to objects that are inside groups or Boolean objects.

You can decide how the linked object will be affected by modifications made to the link parent by checking or unchecking the link options boxes (**Position**, **Rotation**, **Size**, **Join**).

Object tracking lets you decide that the selected object will always point in the direction of the tracked object. To select which object is tracked by the selected object, use any of



the following methods:

- select the parent object from the **Track** drop-down list,
- press the  icon or select the menu command **Object | Pick tracked parent**, then click on the parent object in the *3D views* (objects will be highlighted as the mouse passes over them),
- press the  icon or select the menu command **Object | Pick tracked parent**, then click on the parent object inside the *World Browser*. This last method is the only one that lets you track to objects that belong to groups or Boolean objects.

When a track is defined for the selected object, the Main axis drop-down list becomes active. This lets you define which axis of the object will be pointing in the direction of the tracked object (e.g. a camera that tracks an object should have a **+Z** main axis so that it looks straight at the object it tracks; selecting **+X** will make the camera look straight up (90° upwards from the tracked object direction).

While standard forward-dynamic linking or tracking produces exact motion, this motion usually looks unnatural and jolty. This is because computers perform linking and tracking in a much too perfect manner, whereas a real operator would have a hard time following a rapidly moving target (he would always be catching up or compensating overshoots).

VUE is now able to simulate this “human behavior” by introducing loose [tracking and linking](#) algorithms with the **Response** slider. You can very easily customize the reactivity of the forward dynamics engine for each object, from standard (i.e. perfect reactivity) to slack responsiveness.

Camera Control Center





Camera Control Center

The *Camera Control Center* is made of two parts: the render preview, and the camera controls.

The title bar of the panel displays the name of the currently active camera. If you don't need this panel, click the (🏠) button to fold it up. Click again to unfold it. Double-clicking on the title bar does the same.

The camera control center is now split into two parts: the scene preview and the camera controls. Therefore, they are now independently dockable.

Render Preview

The top half of the panel displays a real-time thumbnail render of your scene. Modifications you make to your scene are immediately reflected in the preview. This can prove extremely useful, especially when adjusting subtle parameters, such as atmosphere and lighting. You can increase the size of the preview by adjusting the width of the [Scene Information](#) bar.

There is a number of ways you can act on the speed and reactivity of the preview; if you right-click on the preview (Ctrl + click on Mac), a menu will appear that lets you customize its behavior.

- **Auto-Update:** the first option in this menu instructs VUE to automatically refresh the preview each time a modification takes place. This is checked by default, but if you find that the program is not reacting as fast as you would like, disabling this feature may be a good way of speeding things up (especially when dealing with



scenes that contain lots of advanced rendering features, such as volumetric lights, depth of field, etc.). Click on the preview each time you want it to be updated.

- **Show Framing Strips:** this option (selected by default) tells VUE to only render the part of the scene that is framed by the camera. Black frames will be displayed around the picture, as required by the aspect ratio of the picture. If this option isn't selected, the entire preview area will be rendered, regardless of the picture's [aspect ratio](#). This option is only available when the picture's aspect ratio has been set to something other than 4/3 (the aspect ratio of the preview).
- **Preview Quality:** this option lets you select the overall precision of the render preview. If you find the preview is too slow, select **Fast** quality. Using **Best** quality is generally not recommended, except on really fast machines.
- **High Priority:** by default, VUE renders the preview at the same time as it completes other pending tasks (e.g. refreshing all material/function previews, refreshing all dialogs, drawing the detailed version of the 3D views). This obviously slows down the preview significantly, and, when a timely response from the render preview is what you need, you may want to select the High Priority option. When selected, this option will postpone all background tasks until rendering the preview is complete – with the consequence that the other tasks will be delayed accordingly. Also, clicking on the preview when the High Priority option is not selected will temporarily grant maximum priority to the preview.

Note:

in order to increase responsiveness, breeze is not applied to plants in the render preview.

Camera Controls

The lower half of the panel features a group of controls that let you manage cameras and easily move around in your scene.



Pan: click and drag on this control to move the camera left-right and/or top-bottom. The movement is done in the camera plane, which means that if the camera is pointing down and you pan towards the top, you are actually moving the camera forwards. Depending on the part of the control you click on, either one or both directions of movement are possible. Movement is blocked by the [clipping plane](#).



Move camera back/forth: click and drag on this control to move the camera forwards or backwards in the direction it is pointing. If the camera is pointing down



and you drag forwards, you are actually moving the camera down. Also, make sure you understand the difference between moving back/forth and modifying the focal process (movement of the camera is blocked by the clipping plane).



Rotate camera: this icon rotates the camera 360 degrees.



Rotate camera around selection: this icon rotates the camera around a selected object. If no object is selected, the camera selects the first object in front of the camera.

Note:

you can slow down the camera controls by holding down the **Ctrl** key as you move this can be customized using the [Operations](#) tab of the *Options* panel.



Focal: click and drag on this control to adjust the focal length of the camera. Dragging the mouse up will zoom into the scene, dragging down will zoom out. The camera doesn't move in the process.



Copy to Camera: The first icon is the one that will be shown when you are in perspective view. Click on this icon to store the perspective view to the active camera. The second icon is the one that will be shown when you are in main view. Click on that icon to store the main view camera settings to the perspective view. Right click on either of this two icons will create a new camera based on the current camera settings. You won't be prompted to name the camera. Instead, VUE will automatically name it as 'Camera ##', where ## is an automatically incremented number.



Secure Active Camera: enable this option to change the main view in perspective view whenever you modify the camera settings with the mouse. This will permit you to play with the perspective camera, and once you are satisfied with the settings you can then store them into the main camera.



Switch to Perspective View / Previous Camera / Next Camera: the first



of these three icons right above the *Render preview* lets you switch the current camera to perspective view. The other two icons let you circulate through the list of stored cameras in the World Browser. If these cameras are animated and you switch cameras at a non zero time, a camera switch keyframe will be created in the *Timeline*. Please turn to [Camera Switching](#) for details on switching cameras in an animation.

In between the **Switch to Perspective View/Previous/Next** camera icons, you will see the name of the active camera. Please turn to [Camera Manager](#) for more details on managing cameras.

Adjusting a Spotlight

If the selected object is a [spotlight or quadratic light](#) and the **View through** option is selected, the controls will act upon the position, orientation and spread angle of the spotlight, rather than the camera.

Note:


The *Render preview* still shows the scene as viewed from the active camera.

Synchronized Cameras and Spotlights

If the current camera or the selected spotlight is synchronized, the camera controls will be disabled. This is because you cannot modify the animation of a synchronized object (to avoid losing synchronization) with [cameras and lights](#).

World Browser

The *World Browser* is at the bottom of the [Scene Information](#) bar, underneath the *Camera Control Center*.

The *World Browser* is used for fast navigation inside your scene. Click on the **Expansion** icon () to expand the *World Browser* so that it fills up the entire right column.

The *World Browser* displays a list of all the objects in the scene.

The World Browser displays four tabs. Each tab is designed to let you access a different type of information regarding your scene.



Objects: this tab lets you view the objects in the scene; you can list the objects



using different sorting methods.



Materials: this tab displays all the materials used in the scene, together with their hierarchy (for mixed materials).



Library: this tab displays all the objects that are used several times in the scene (including all EcoSystem populations). Using this tab, you can modify all copies of the same object simultaneously.



Links: this tab displays all linked items in the scene (texture maps and imported objects).

At the bottom of each tab is a small toolbar that display tools frequently used in each specific tab.

Status Bar

This is the general information bar that sits right at the bottom of the interface. The left side of it displays current render status, and menu command help messages.

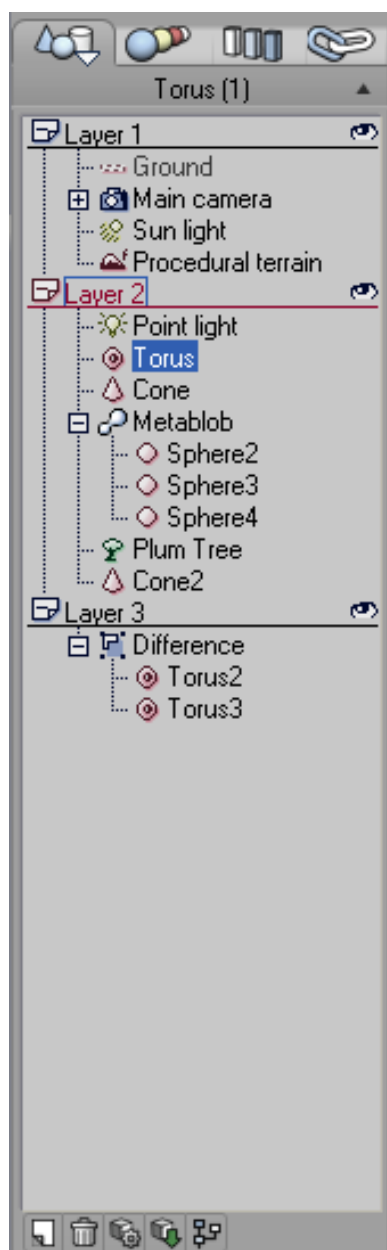
To the right, you will find an indication on the number of processors on your computer as well as the number of objects and the number of lights in your scene. While the only limit to these figures is the power of your machine, we feel that keeping them under 200 objects and 20 lights would be reasonable.

The rightmost figure is an estimate of the equivalent number of polygons that your scene would require if it were modeled inside a “standard” 3D package. A drop-down list lets you switch to displaying the free resources or used GPU resources instead.



Objects Tab

Organization of the List



World Browser – Objects Tab

The first tab in the *World Browser* displays a list of all the objects in the scene. This list of objects can be organized in several ways; when the **Objects** tab is the active tab, clicking on the tab will display a popup menu that lets you change the way objects are displayed:

- **Organize In Layers:** in this type of organization, a list of the different layers in the scene is displayed, with all the objects that belong to this layer appearing below it. This is the default type of organization. Read [here](#) for a more detailed description of Vue layers.
- **Sort By Names:** in this type of organization, the objects are displayed in a list that is sorted by name in alphabetical order. This is useful when you want to locate an object in a complex scene and you only remember its name.
- **Sort By Size:** in this type of organization, the objects are displayed in a list that is sorted by order of increasing size. Tiny objects will appear at the top of the list, whereas very large objects will be displayed at the bottom. This is useful for locating very small objects.
- **Sort By Types:** in this type of organization, the objects are displayed in a hierarchical list sorted by object types (e.g. Spheres, Terrains, Boolean operations). This is useful to pinpoint objects of similar nature.

Using the List of Objects

You can select objects by clicking on their name in the list. This also works for objects that are placed inside groups, which means that you don't have to

Ungroup and **re-Group** groups or Boolean operations to modify one of their members! Click again, and the name of the object changes into a label where you can edit the object name (press **Enter** to confirm the new name).

You can move up and down in the list using the Up/Down arrows.

And you can drag and drop objects from one position to another. You can also drag objects into or out of groups or Boolean objects!

Pressing **Control** while dragging objects will result in a copy of the dragged objects being generated, and placed at the drop point.

Multiple objects can be selected by **Control** clicking on them. **Shift** clicking will select all objects in the marked range.



Pop-Up Menu

If you right-click on any object (plant, object, terrain) in the list of objects in the World Browser, a menu displays. The options available depend on the type of object you've selected.




- **Frame Selected Objects:** selecting several objects and clicking this option puts a frame around them for easy moving as a group.
This applies for the perspective camera as well as the other cameras.
- **Show All Layers:** if some layers have been hidden, this restores all to show.
- **Hide All But This Layer:** this option hides all but the selected layer from view. This does not hide from render.
- **Show Current Layer:** this unhides a selected layer.
- **Lock Current Layer:** if a layer is locked, its objects are displayed in gray and cannot be selected in the viewports.
- **Hide Current Layer:** this hides the contents of the layer in the viewports. This does not hide them from render.
- **Collapse All But Current Layer:** this collapses all open layers except the selected one.
- **Expand All Layers:** this opens all layers to show contents. *World Browser Objects Tab*
- **Cut:** Removes selected object from scene. The object can be restored with a *Paste* operation.
- **Copy:** click to copy the selected object. Use a *Paste* operation to add the copy of the selected object to the scene. This works with cameras as well as other objects in the *World Browser*.
- **Paste:** adds an object to the scene from a previous *Cut* or *Copy* operation. This works with cameras as well as other objects in the *World Browser*.
- **Delete:** removes selected object from scene. Unlike *Cut*, it cannot be restored with a *Paste* operation. You would need to use *Undo*.
- **Select All:** selects all objects in the scene.
- **Deselect All:** deselects all objects in the scene.
- **Hide From Render:** hides the selected object from a render.
- **Enable Collision:** when toggled on, EcoParticles will collide with a bounce off this object.
- **Group Objects:** places the selected objects into a bounding box.
- **Ungroup Objects:** removes the bounding box from selected objects.



- **Replace By (Keep Proportions):** use this command to replace an object with the object of your choice, keeping the proportions of the original object. All copies of the object will be replaced by the new object. If you select multiple objects to replace, all selected objects will be replaced by the new object.
- **Replace By (Fit Object):** use this command to replace an object with the object of your choice, fitting the new object into the bounding box of the original object. All copies of the object will be replaced by the new object. If you select multiple objects to replace, all selected objects will be replaced by the new object.
- **Bake to Polygons:** converts any selected object in a Vue scene into a polygon mesh approximation. The word “approximation” is important here, as some objects used in Vue simply don’t have a polygonal equivalent (for instance, spheres or plants).
- **Convert to Area Light:** converts the selected object to an area light. This action cannot be undone.
- **Add Daylight Portal to Object:** adds a daylight portal to the object selected.
- **Change Material:** opens the Visual Material Browser for new material selection.
- **Edit Material:** opens the *Material Editor* for editing of material.
- **Edit Object:** If the item selected is an object, the *Edit Object* dialog displays. If the item selected is a plant, the *Plant Editor* opens. If the item selected is a cloud layer, the *Atmosphere Editor* displays.
- **Save Object:** saves the selected item as a Vue object (.vob). This includes terrains, plants and cloud layers.
- **Rename:** allows you to rename an item.
- **Help:** displays the help manual.

The Bottom Toolbar

At the bottom of the *World Browser* is a small toolbar. In the **Objects** tab, the effects of the buttons in this toolbar are:

-  **New layer:** click this button to add a new [layer](#) to your scene.
-  **Delete selected object(s):** click this button to delete the selected object(s) and/or layers. This button is only available when one or several objects or layers are selected.
-  **Edit selected object:** click this button to open the Object Editor for this object



(e.g. the Plant Editor for a plant). This button is only active when the selected object is editable.



Export selected object: click this button to [export the selected object](#). This button is only available when a single exportable object is selected.



Edit Objects Graph: Click this button to open the [object graph](#) for this object.

Object Identification

[Objects](#) are depicted by small pictures placed in front of them. These pictures are designed to facilitate object identification. The signification of the pictures is:



Infinite Plane: this object is typically a water, or ground plane.



Sphere: this object is a Sphere,



Cylinder: this object is a Cylinder,



Cube: this object is a Cube,



Pyramid: this object is a Pyramid,



Cone: this object is a Cone,



Torus: this object is a Torus,



Plane: this object is a Plane,

















Alpha Plane: this object is an Alpha plane,



3D Text: this object is 3D Text object,



-  **Terrain:** this object is a Terrain,
-  **Symmetrical Terrain:** this object is a Symmetrical terrain,
-  **Skin Only Terrain:** this object is a Skin Only terrain,
-  **Symmetrical Skin Only Terrain:** this object is a Symmetrical Skin Only terrain,
-  **Procedural Terrain:** this object is a Procedural terrain,
-  **Procedural Symmetrical Terrain:** this object is a Procedural Symmetrical terrain,
-  **Procedural Skin Only Terrain:** this object is a Procedural Skin Only terrain,
-  **Procedural Symmetrical Skin Only Terrain:** this object is a Procedural Symmetrical Skin Only terrain,
-  **Plant:** this object is a Plant,
-  **Global EcoSystem:** this object is a Global EcoSystem,
-  **Rock:** this object is a Rock (the icon appears yellow if the rock's illumination is baked),
-  **Polygon Mesh:** this object is a Polygon Mesh object (i.e. an imported object. In VUE Pro versions, the icon appears yellow if the object's illumination is baked),
-  **Animated Mesh:** this object is either a Rigged Mesh or an Animated Poser object (In VUE Pro versions, the icon appears yellow if the object's illumination is baked),
- 



Bone: this object is a Bone,



Point Light: this object is a Point light,



Quadratic Point Light: this object is a Quadratic point light,



Spot Light: this object is a Spot light,



Quadratic Spot Light: this object is a Quadratic spot light,



Directional Light: this object is a Directional light (e.g. the sun),



Light Panel: this object is a light panel,



Daylight Portal: this object is a daylight portal,



Light Emitting Object: this object is a light emitting object,



Group: this object is a Group,



Boolean Union: this object is a Boolean union,



Boolean Intersection: this object is a Boolean intersection,



Boolean Difference: this object is a Boolean difference,



Metablob: this object is a Metablob object,



Hyperblob: this object is a Hyperblob object,



MetaCloud: this object is a MetaCloud object,

- 

Cloud: this object is a Cloud layer,

- 

Ventilator: this object is a Ventilator,

- 

Python Object: this object is a Python object created by a python script,

- 

Camera: this object is the Camera,

- 

Spline: this object is a Spline group,

- 

Particles Effector: this object affects EcoParticles. It will usually be hidden from render.

Object Renaming





You can rename objects by clicking twice on their name. This will open a little field letting you enter the new name for the object. You can rename several objects simultaneously using the **Rename...** command from the popup menu. The *Object Renaming* dialog will appear, letting you enter a new name for all the selected objects. If the **Keep object numbering** option is selected, any digits that appear at the end of the object names will be appended to the new name.

Hidden From Render / Switch Off


You can also show/hide an object from render using the [Aspect](#) tab of the *Object Properties* panel.


When an object is hidden from render, or when a light is turned off, a small cross will appear on top of the object's identification picture (as described above) indicating that the object is now hidden from render (or that the light is turned off). Objects that are hidden from render (or lights that are turned off) still appear in the list of objects and in the *3D Views*, but they don't appear when you render the picture (or animation). Hiding objects from render is particularly useful when creating helper objects.

Groups, Boolean Objects and Metablobs -- MetaClouds

Group objects also have a symbol in front of them (PC:  or , Mac:  or ).









The  symbol means that the object is **folded up**; members of the group are not displayed inside the browser. Clicking on the picture will unfold the contents of the group.

The  symbol means that the object is **unfolded**; members of the group are displayed underneath the group, and are linked to it in the list using a dotted line. You can access any of the members directly by clicking on them. Clicking on the picture will fold the group back up.

The ability to select objects that are inside groups makes for tremendously convenient use. Objects can be modified even once they are inside a group. Better still, objects may be added or removed from groups simply by dragging them into, or out of the group (in order to do so, the group has to be unfolded).

MetaClouds are a slightly specific type of group in that they can only contain MetaCloud primitives. You cannot add any other type of object to a MetaCloud.

Camera Group

The camera in the list of objects also has the group symbol (PC:  or , Mac:  or ) in front of it. If you unfold the camera group, you will see a list of all the different cameras in the scene. The current camera is identified by the active camera symbol () , whereas other cameras are identified by the inactive camera symbol (). Click on an inactive camera to select it, and double-click to make it the active camera. Notice that the name of the camera group changes to reflect the currently active camera.


Note:

You can also select inactive cameras by clicking on them in the *3D Views*, and you can switch active cameras by double-clicking on the camera to activate.


Layers





All the objects that you create are automatically placed inside layers. In order to see the layers, you need to organize the list in layers.


When the list is organized in layers, the objects will be displayed under the layer they belong to. Layers act as organizers for the objects inside your scene. Objects will render the same, whatever layer they are placed in (unless you tell the renderer to ignore specific layers).

To create a new layer, simply press the **New layer** button () at the bottom of the *World Browser*. A new layer will automatically be added at the bottom of the list, and new objects will subsequently be placed in this layer. You can add as many layers as you want.



To delete a layer, click on the layer to delete, and then click on the **Delete selected object(s)** button (). All objects in the selected layer will be deleted, and the selected layer will be removed. If you delete all of the objects within a layer, the layer will not be deleted automatically. The empty layer may remain or be deleted manually.

The contents of the layers may be visible or hidden, depending on whether they are unfolded or not. Unfolded layers are depicted by a  ( on Mac) on the left of the browser. To fold up and hide what the layer contains, click on that picture. The layer folds up, and the button changes to  ( on Mac). Please understand that objects inside the folded up layers still exist, although they are not listed in the *World Browser*. Folding layers up is a good way of limiting the number of objects displayed by the browser.

Layers can be hidden from render globally. Just check the **Hide from render** icon () next to the layer name to hide the entire contents of the layer from the render. If you choose not to hide the entire layer, you can still hide any item in that layer by checking them individually.

Objects may be moved from one layer to another by dragging them out of the old layer and dropping them into the new one. Layers can be moved freely up or down in the hierarchy, whether they are hidden, locked or active.

Clicking on the name of a layer will select all objects inside the layer. Clicking again will let you rename it (press **Enter** to validate the new name).

On the right side of the layer is a little picture indicating the state of that layer. This can be any of the following:



Active layer: objects inside active layers are visible and may be selected in the *3D Views*.



Locked layer: objects inside locked layers are displayed in gray and cannot be selected in the *3D Views*.



Hidden layer: objects inside hidden layers are not visible and cannot be selected in the *3D Views*. All layers can be hidden.

To change the state of a layer, click on the state picture. The state of the layer cycles through Active / Locked / Hidden. Double-clicking on the state picture always activates the layer.

The ability to change the state of a layer is the key to their organizational power. The fact that you can lock or hide layers means that you can decide to temporarily hide away parts of your scene by putting them into a hidden layer. In doing so, objects will still be there



for rendering, but won't clutter up your *3D Views* as you concentrate on another part of the scene. When you need them back, just click on the layer state, and here they come! Alternately, putting objects in a locked layer will keep them visible in the *3D Views* (e.g. for reference), while not hindering selection of other objects you are working on. Objects can still be copy/pasted to locked or hidden layers.

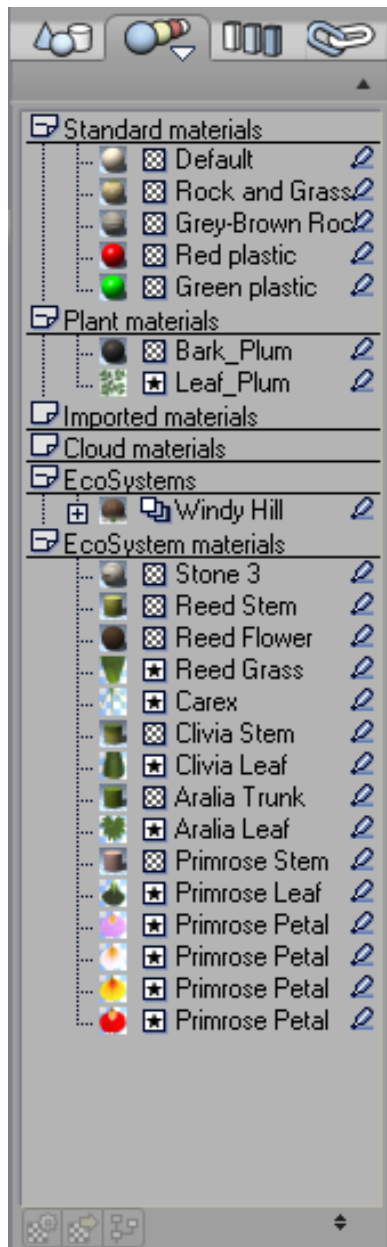
All layers can be freely moved up and down in the World Browser. This doesn't affect the image in any way.

The state of a layer has absolutely no influence on the way objects are rendered (unless expressly asked for – see *Render Options*). Not convinced? Take a look at a sample scene. Objects in the scenes are usually organized in layers. Try activating or hiding the different layers, and you will see how cluttered the scene can become, making selection and progress tedious.

When you start a new scene, all objects are, by default, positioned in the first layer. Successive objects that you create will be placed in the layer that has the focus. The name of the layer that has the focus appears in red. It is the last active layer you used. Selecting an object from an active layer will give focus to that layer.



Materials Tab








World Browser Materials Tab

The second tab in the *World Browser* provides a handy list of all materials used in the scene.

The materials are gathered into 6 categories:

- **Standard materials:** this category holds all standard materials (i.e. materials that do not belong to another category),
- **Plant materials:** this category holds the materials used for plants,
- **Imported materials:** this category holds the materials that were created when an object was imported from another application,
- **Cloud materials:** this category holds the cloud materials used in the [Clouds](#) tab of the *Atmosphere Editor*,
- **EcoSystems:** this category holds all the materials that define [EcoSystems](#) (that is EcoSystem materials, mixed or layered materials with one of the sub-materials being an EcoSystem),
- **EcoSystem materials:** this category contains all the materials used on the population of the various EcoSystems in the scene.

In a similar way to how layers operate, the contents of each category can be folded up. Unfolded categories are depicted by a  ( on Mac) on the left of the browser. To fold up and hide the contents of the category, click on that picture. The category folds up, and the button changes to  ( on Mac). Although the materials are no longer displayed in the list, they still exist, and unfolding the category will show them back. Empty categories are depicted by a . Clicking on this has no effect.

Clicking on a material name will select all the objects that use this material. This is a handy way of checking which objects use a given material. Click again and a small label will appear, letting you modify the name of that material. Material names that are framed by a gray rectangle indicate that the material is used by a selected object, but other objects using the material are not selected.

Double-click on the material name to open the *Material Editor* and edit this material. Please look [here](#) for details on the *Material Editor*. You can also edit cloud materials directly. This is much faster than opening the [Atmosphere Editor](#) and browsing to the desired layer.

Increasing Responsiveness in Large Scenes

In large scenes, with many materials, the list of materials in the *World Browser* can become hard to work with. You can select a group of materials to view, limiting the materials to a certain criteria.




Right-click on the **Materials** tab to display the choices. Your options are:





- **Show Only Materials of Selected Objects:** select an object(s) and choose this option to display only those materials for the selected objects. If no objects are selected, the *World Browser* will display empty.
- **Show Only Materials Names:** this displays names only, without the small icons, simplifying the displayed list.
- **Hide Materials of EcoSystem Specimens:** this hides materials used in objects of EcoSystems.


These options are saved in the User preferences so it is applied to all scenes in each session of Vue. These options can be toggled on/off.

Material Types


The different types of materials are identified by a symbol in front of the material name:

-  **Simple material:** this material is a simple material (see [here](#) and [here](#)). It can be used to define the entire material, or sub-materials of mixed materials, or layers of layered materials.
-  **Bitmap material:** this material is a simple material with its colors based on a bitmap image.
-  **Mixed material:** this material is a mixed material (see [here](#) and [here](#)).
-  **Layered material:** this material is a layered material (see [here](#) and [here](#)).
-  **Volumetric material:** this material is a volumetric material (see [here](#)).
-  **EcoSystem material:** this material is an EcoSystem layer (see [here](#)).

Mixed and layered materials appear in the *World Browser* as a hierarchy, letting you access sub-materials or layers directly. They are identified by a small symbol in front of them (PC:  or , Mac:  or ). EcoSystem materials will also appear that way, letting you access the underlying material directly.

The  symbol means that the material is **folded up**; sub-materials or layers of this material are not displayed inside the browser. Clicking on the picture will unfold the contents of the material.



The  symbol means that the material is **unfolded**; sub-materials or layers of this material are displayed underneath the material name. You can edit any of the sub-materials directly by double-clicking on it. Clicking on the small symbol will fold the material back up.

This feature is very handy when material hierarchies become complex. Because you can access remote sub-materials directly, you don't have to open all the intermediate material levels first (this can also be done using the material popup menu). The material hierarchy of individual materials can also be accessed using the [Material Hierarchy](#) of the *Material Editor*.

To the far right of each material and material group listed is a small icon. This icon has three states:


- The default setting for this icon is to render normally.
- Click once and the material is hidden from render.
- Click again to render the material as a flat color. Right-click to open the *Color Selector* to change the color.

This can be useful if you want to see the exact impact of the environment settings on the edited layer, or if you want to concentrate on a given layer without losing the other layers (in this case, just hide all other layers temporarily).

Editing Multiple Materials

You can edit all the materials that are assigned to a given object or group of objects by selecting the **Edit All Materials** option from the material preview popup menu in the *Object Properties* panel (see [Changing Object Material](#)).

Material Preview

In front of each material name is a tiny preview of the material used to facilitate identification of the material. You can adjust the size of the preview using the **Preview size** () control in the toolbar at the bottom of the *World Browser* (see below). Click on this control and drag the mouse up to increase the size of the preview. Drag down to reduce the size.

The Bottom Toolbar

At the bottom of the *World Browser* is a small toolbar. In the **Materials** tab, the effects of the buttons in this toolbar are:



- **Edit material:** click this button to open the [Material Editor](#) and edit this material. This button is only available when a single material is selected.





-

Replace material: click this button to open the Material Browser and replace the selected material with another one. This button is only available when a single material is selected.



-

Edit material graph: click this button to directly access the graph of the selected material (see [here](#) for details on the Function Graph).



-

Preview size: this button is used to resize the material previews (see above for details).



Library Tab



World Browser Library Tab

The **Library** tab displays a list of all the objects in the scene that are used several times.

If the multiple copies of the object were created by copy-pasting, duplicating or replicating, the objects will be identified as **Master objects**. If the multiple copies were created using an EcoSystem, or the *EcoSystem Painter*, they will appear under the **EcoSystem population** category.

If you select a master object on this list, all the copies of this object will be selected.

If you right-click on an object/plant in an EcoSystem, a menu displays. One option is to go into **Master Object Edition Mode**, the other option is **Materials**, which displays a submenu of materials for that item with options to edit, copy or save the material for that item. This provides quick access for changing a material for all EcoSystem items.

Popup Menu

The options of the popup menu in the **Library** tab are:

- **Replace By (Keep Proportions):** use this command to replace a master object with the object of your choice, keeping the proportions of the original master object (see [Replacing Objects](#)). All copies of the master object will be replaced by the new object. If you select multiple objects to replace, all selected objects will be replaced by the new object.
- **Replace By (Fit Object):** use this command to replace a master object with the object of your choice, fitting the new object into the bounding box of the original master object (see [Replacing Objects](#)). All copies of the master object will be replaced by the new object. If you select multiple objects to replace, all selected objects will be replaced by the new object.
- **Edit Object:** if the selected item is editable, use this command to edit it. This will open the corresponding type of editor. Any modifications made to the master object will be applied to all the copies of the object.
- **Master Object Edition Mode:** if you select this option, a veil will appear, cloaking all the objects in the *3D Views* except for a single copy of the master object. If the object is a group, a Boolean or a Metablob object, this mode lets you edit the different components of the master object directly in the views. To exit the Master Object Edition Mode, simply deselect the master object or reselect this menu command. Any modifications made to the master object will be applied to all the copies simultaneously.
- **Convert to individual objects:** this removes the duplicated object from the Master objects list so edition of each object can be made separately.

Note:



you can edit an individual copy of a master object like any other object. If you modify a copy of a master object, the “connection” with the master object will be automatically broken.

The Bottom Toolbar

At the bottom of the World Browser is a small toolbar. In the **Library** tab, the effects of the buttons in this toolbar are:



Edit selected master objects: if the selected items are editable, use this button to edit them. This will open the corresponding type of editor. Any modifications made to the master object will be applied to all the copies of the object.



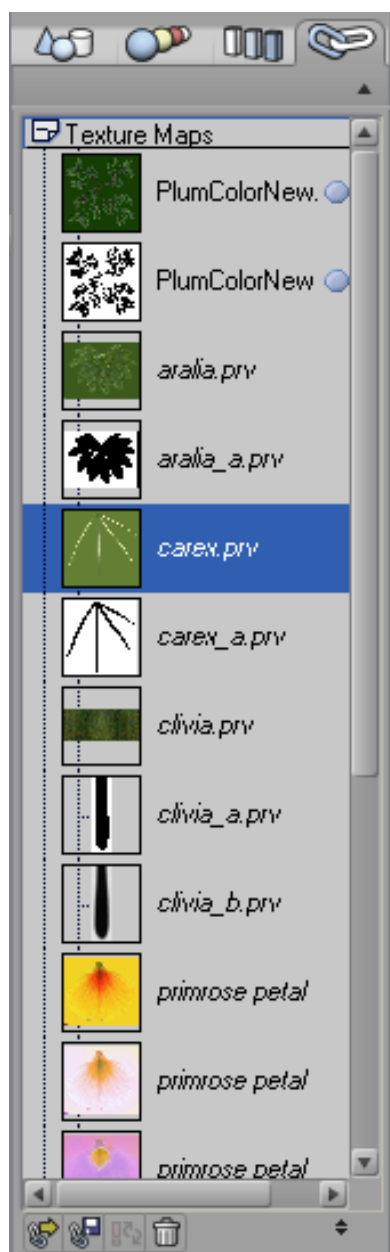
Master object edition mode: click this button to activate the master object edition mode; A veil will appear, cloaking all the objects in the 3D Views except for a single copy of the master object. Using this mode, you can resize or rotate the master object. If the object is a group, a Boolean or a Metablob object, you can edit the different components of the master object directly in the views. To exit the Master Object Edition Mode, simply deselect the master object or reselect this menu command. Any modifications made to the master object will be applied to all the copies simultaneously.



Delete selected master objects: click this button to delete the selected master object(s) and all their copies. This button is not available for EcoSystem populations.



Links Tab








World Browser Links Tab

The last tab in the *World Browser* maintains a list of all the external items that were loaded into VUE, and the way they are linked.

The list is organized in two categories:

- **Texture maps:** this category displays all the texture maps (imported pictures) used in the scene together with a small icon indicating the way the texture map is linked. Please check the details on the *Material Editor – Mapped Picture* for an example of loading a texture map.
- **Objects:** this category holds a list of all the objects that were created in another 3D application and imported into Vue (see [here](#) for details on importing objects from other applications).

In a similar way to the categories in the **Materials** tab, the contents of each category can be folded up. Unfolded categories are depicted by a  ( on Mac) on the left of the browser. To fold up and hide the contents of the category, click on that picture. The category folds up, and the button changes to  ( on Mac). Although the items are no longer displayed in the list, they still exist, and unfolding the category will show them back. Empty categories are depicted by a . Clicking on this has no effect.

Texture Maps

The first category displays a list of all the texture maps used in the scene. Clicking on a texture map name will select all the materials that use this texture map, and all the objects that use these materials. This is a handy way of checking which objects use a given texture map. Missing textures will be identified as broken links.

Double-click a texture to view the texture map at full resolution (using external viewer).

Alongside the name and preview of the texture map, you will notice a small symbol that identifies the way the texture map is linked. Possible linking options are:

•

No symbol: with this linking option, only the name of the file is saved in the scene. When you reload the scene, the texture map will have to be at the same location in order to be successfully loaded. If you modify the texture map in an external application, the version used by VUE will not be updated until you reload the scene.

• 

Synchronized: this is similar to the previous option, except that the texture map is automatically reloaded if it has been modified in an external application (a prompt will appear offering to reload modified texture maps).

• 

Incorporated: with this option, the texture map is saved together with the scene.




You don't have to worry about modifying or deleting the original file, because a copy of this file will be stored inside the scene. Of course, this results in much larger scene files, but is very useful when you need to transfer items to another party or publish them; you don't need to worry about including appropriate texture map files. If the texture map is modified in an external application, it won't be modified inside VUE until you reload it.

- 

Incorporated and synchronized: this is similar to the above option, except the texture map will be automatically reloaded if it is modified in an external application.

You can toggle the linking options of a texture map using the popup menu (see below).


Texture Map Preview

In front of each texture map filename is a tiny preview of the texture map used to facilitate identification. You can adjust the size of the preview using the **Preview size** () control in the toolbar at the bottom of the *World Browser* (see below). Click on this control and drag the mouse up to increase the size of the preview. Drag down to reduce the size.

Imported Objects

Whenever you import an object created with another 3D application, this object will be listed in the *Objects* category of this list.

Clicking on the name of an object in this list will select the corresponding object. Double-clicking on the name of the object will open the [Polygon Mesh Options](#) dialog.

When you import an object from another 3D application, you have the possibility of decimating the object (see [here](#) for details on polygon decimation). Objects that have been decimated are identified by a small pictogram () to the right of the object's name. You can remove decimation by re-importing the object without decimation (see below).

Popup Menu

The popup menu of this tab offers the commands below. Menu commands apply to the selected items or the item under the mouse cursor at the time of displaying the menu if no item is selected.

- **Incorporated:** select this menu option to toggle the incorporated status of the texture map under the mouse cursor. This option is only available if the item under the mouse cursor is a texture map.
- **Synchronized:** select this menu option to toggle the synchronized status of the texture map under the mouse cursor. This option is only available if the item under the mouse cursor is a texture map.



- **Locked:** This option locks the material image and the material to this object only. This operation cannot be undone.
- **Downsample:** this option allows you to non-destructively change the current resolution of the material. Your options are:
 - To halve resolution
 - To divide it by four
 - To divide it by eight
 - To select a custom downsampling coefficient

This can be reversed by selecting the **Original size** option.

- **Replace Link:** selecting this menu command will display a standard Picture or Object File Browser letting you select a picture or an object file that should replace the selected item.
- **Export Link:** selecting this menu command will display a Picture or Object File Browser letting you select the name of the file under which you would like to save an incorporated texture map or an imported object. This command is only available if the item under the mouse cursor is an incorporated texture map or an imported object (see [here](#) for details about exporting objects).
- **Reload Link:** select this menu command to reload a texture map or an imported object that has been modified in an external application. For VUE Pro users, if the reloaded item is an imported object, and this object has been decimated, it will be reloaded with the same level of [decimation](#) as previously imported.
- **Reload Without Decimation:** this menu command is only available if the item under the mouse cursor is an imported object, and this object has been decimated. When you select this command, the object will be re-imported without being decimated. Using this feature, you can import a large object, decimate it to facilitate placement and test rendering, and then re-import it without decimation when you are ready for the final rendering.
- **Reload All Without Decimation:** this command will reload the full geometry of all imported objects that have been decimated.
- **Incorporate All Texture Maps:** this command will simply incorporate all the texture maps that are not yet incorporated (see above for an explanation of incorporated texture maps). This is very handy if you want to make sure that all the texture maps are included with the scene before you transfer it to another party.
- **Copy all Texture Maps in a Folder:** this option copies all textures from the imported objects into a folder for future reference.



The Bottom Toolbar

At the bottom of the World Browser is a small toolbar. In the **Links** tab, the effects of the buttons in this toolbar are:



Replace link: clicking this button will display a Standard File Browser letting you select a picture or an object that should replace the selected item.



Export link: clicking this button will display a standard Picture or Object File Browser letting you select the name of the file under which you would like to save an incorporated texture map or an imported object.



Reload link: click this button to reload a texture map or an [imported object](#) that has been modified in an external application. If the reloaded item is an imported object, and this object has been decimated, it will be reloaded with the same level of decimation.




Delete object: click this button to delete the selected object(s). This button is only available when the selected item is an imported object.



Preview size: this control is used to resize the texture map previews.



Animation Timeline



The complete Timeline (Properties, Graph and Preview Timeline unfolded)




1. Navigation & keyframing
2. Animated items list
3. Property animations
4. Graph legend
5. Property graph
6. Preview/Render animation
7. Animation Toolbox
8. Path display options
9. Show in Graph & Property TimeSpline
10. Animation preview

The animation *Timeline* is Vue's animation control center.

To display the animation Timeline, select the menu command **Display | Display Timeline** or click the **Display Timeline** icon (🎬). This displays the *Animation Wizard* together with the *Timeline* (you can disable this feature using the checkbox on the first page of the Wizard). The *Animation Wizard* helps you easily setup an animation of your scene. If you don't want to use the Wizard, just press **Escape** or click **Close**.


The *Timeline* panel is subdivided into 4 sections: the *Main Timeline*, the *Properties Timeline*, the *Animation Graph* and the *Animation Preview*. Initially, only the *Main Timeline* is visible.

You can display the other sections of the *Timeline* by pressing the corresponding unfold icons:


-  displays the *Properties Timeline*
-  displays the *Animation Graph*
-  displays the *Animation Preview*.

Click those icons again to fold-up the corresponding part of the *Timeline*.

The Main Timeline



The Main Timeline

The *Main Timeline* summarizes the animation information contained in your scene. It comprises a set of controls () that are used to navigate and preview the animation. Under this set of controls you will find the **Current time edit box**. This box displays the current time inside the animation. You can modify it by entering a new value, or by dragging the Current time slider to a new position. The current time can be displayed as:





- Seconds
- SMPTE
- Frames

To the right of the animation controls, you will find a ruler. It is this ruler that displays summarized information about your animation. It allows you to position events in time, and is graduated in seconds.

The solid bar in the middle of the ruler is the **Duration** bar. It depicts the total duration of the animation. It is not possible to adjust the length of this bar manually, because it



automatically adapts to the duration of the longest property animation in the scene.

This **Duration** bar is struck through by a yellow line. This yellow line indicates the active part of the animation. This is the part that will be played if you press the **Play** icon () or rendered if you press the **Generate preview** icon (). You can modify the active part of the animation by dragging the handles situated at both ends of the line () and (). The line will shrink or expand accordingly. By default, the end of the active part of the animation “clings” to the end of the **Duration** bar. If you modify the end of the active part of the animation, this is no longer the case. You can restore it by dragging the end slider to a negative time value.


The **Current time slider** sits above the **Duration** bar. Drag it to the left or to the right to change the current time. The views will be automatically updated to show the scene at the new time.

If the **Current Time | Constrain Current Time** option is selected from the *Timeline* menu, the current time slider will not be able to move beyond the beginning and end of the active part of the animation.

Below the **Duration** bar are drawn the **Keyframe handles**. Only the keyframes belonging to the currently selected objects are displayed in the Timeline. This reduces clutter when many keyframes have been defined in a scene. If no objects are selected, the keyframes of all objects, materials and atmosphere in the scene appear.

Auto-Keyframing


By default, keyframes appear each time you define a new property for an object, at a different time. This is called auto-keyframing. For instance, if you select an object, modify the current time and then drag the object to a new position, a position keyframe will be created. This newly created keyframe will appear under the Duration bar. You can modify the time at which a keyframe occurs by dragging it to the left or to the right.

The **Auto-Keyframing** icon () lets you control the way keyframes are created.

If you click on this icon, a menu will appear with the following options:

- **Enable Auto-Keyframing:** select this option to enable/disable auto-keyframing. When auto-keyframing is disabled, keyframes will not be created automatically; the Auto-Keyframing icon will be changed to the **Manual Keyframing** icon (see below).
- **Add Keyframe to all Properties:** select this menu command to add a keyframe to all the animation properties of the selected objects at the current time.
- **Add “xxx” Keyframe,** where “xxx” is the name of an animation property: selecting this menu command will add a keyframe to the “xxx” property at the current time.



When auto-keyframing is disabled, the Auto-Keyframing icon changes to the **Manual Keyframing** icon (). Click this icon to create a new keyframe (when auto-keyframing is disabled, keyframes are not created automatically as soon as you change an object's property. If you check an object's property and click the icon, appropriate keyframes will be created. If you don't, the animated object will revert to its original state as soon as you change current time). If you long-click the Manual Keyframing icon, the keyframe menu will appear.

Moving around in the Timeline


To **Scroll** the *Timeline*, press the right mouse button, and drag the mouse to the right or to the left (Ctrl drag on Mac).



To **Zoom** in or out of the *Timeline*, press Control as you drag the mouse down (zoom out) or up (zoom in) with the right mouse button down.

Constraining the Current Time Slider

You can constrain the Current Time slider to the length of the animation, so that the slider cannot go before the beginning of the animation, or move further than the end of the animation. This feature is controlled using the menu command **Constrain Current Time** on the popup menu.

Rendering the Animation

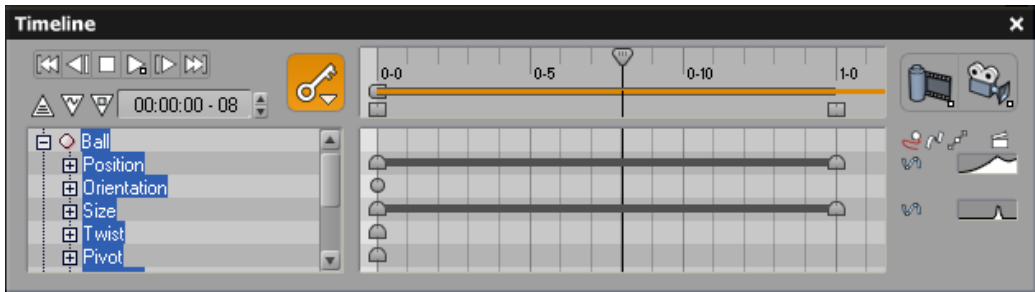
Use the  icon to refresh the [animation preview](#). Activate the alternate action () of this icon to access the [Animation Preview Options](#) dialog.

Use the  icon to render the animation (read more [here](#)). Activate the alternate action () of this icon to directly access the [Advanced Animation Options](#) dialog.


You can resize the *Timeline* horizontally. You can also resize it vertically. If the *Animation Preview Timeline* is displayed, resizing will add new lines of previews. If the *Properties Timeline* is displayed, resizing will add new lines to the Properties Timeline. If the *Animation Graph* is displayed, resizing will increase the size of the graph. If both the *Animation Graph* and the *Properties Timeline* are displayed, you can adjust the respective size of these two elements using the resizer handle in between them.







The Properties Timeline



The Properties Timeline unfolded

When you press the  button, the *Timeline* expands to include the *Properties Timeline*.

The *Properties Timeline* gives you a detailed view of all the animated items of your scene. To the left, you will find a list of all these items (animated objects, animated materials, animated atmosphere...). This list operates much like the [World Browser](#) except that all objects can be expanded (they are preceded by a  / , like groups in the *World Browser*). Clicking on the  /  will display all of the animatable properties of the object (e.g. position, orientation, size...). However, unlike in the *World Browser*, pressing **Delete** will not destroy the object, but only destroy its animation.

This list of animated items also introduces 3 new identification pictograms:



Animated material: this element is an animated material (either *Surface*, *Velocity* or *Complete animation*; read everything about Animating Materials [here](#)).



Animated cloud: this element is an animated cloud material (either *Surface*, *Velocity* or *Complete animation*; read everything about Animating Clouds [here](#)).



Animated post processing: indicates that some global post-processing applies to all cameras, and that this post-processing is animated (see [here](#) for details on post-processing).

The ruler to the right displays a list of all keyframes defined for the properties, together with a solid bar that depicts the total duration of the properties animation. Keyframes are automatically added each time you modify a property of an object (e.g. a new Position keyframe will be added to the object animation each time you move it at a different time).



You can modify the time at which the keyframe occurs by dragging it with the mouse. To delete a property keyframe, just click on it and press **Delete** (or select the **Delete keyframe** from the *Timeline* popup menu).

To the right of the properties ruler, you will find a set of controls that let you further customize the behavior of your objects animations. The first three relate to the way the path followed by the object is displayed in the *3D Views* (they are only useful if your object actually moves during the animation):



Persistent path: select this option to display the path even when the object is not selected.




Show path as ribbon: this affects the way the path is displayed in the 3D Views. Instead of being displayed as a thin line, it will be depicted by a ribbon, giving an idea of the orientation of the object as well as its position.



Show tangents: this will display the tangents of the objects path at each keyframe.
Note:

It is not possible to modify these tangents directly (you have to modify the path to modify the tangents), because it is spline driven. It is possible to edit these tangents directly in the timeline's animation graph, however.


The last button () opens the [Animation Toolbox](#) for the property.

Opposite each animation property, you will find the **Show in curve** () toggle button. When this button is pressed, the corresponding animation property will always appear in the animation graph (see below), even when the property is not selected. This is useful e.g. if you want to adjust the curves of one property according to that of another property.

Alongside the Show in curve button you can see a curve. This is known as the **Time spline**. This [Time spline](#) lets you fine tune the flow of the animation.



The Animation Graph







The Animation Graph unfolded

When you press the  button, the *Timeline* expands to show the *Animation Graph*.

The *Animation Graph* provides advanced control over the way each animation property evolves with time. The list to the left of the graph shows all the animation properties of the currently selected objects. If you select one of these properties, one or several curves will appear on the graph to the right of the list, showing the evolution of the property over time.

The number of curves that are displayed for each property depends on the type of the property:


- If the selected property is a color or a position, three curves will appear in the graph; one for each component. The property will be preceded by a  /  symbol in the property list, indicating that it can be expanded, letting you select individual components of the animation properties (XYZ for a position and RGB for a color).
- The **Orientation** property is also preceded by the  /  symbol. By default, orientation is processed using quaternion arithmetic. If you want to access the individual angles of rotation, you need to convert the animation to the [Euler](#) model. This is done automatically when you expand the animation property (after displaying a warning). You can then access the individual rotations in the graph.
- All other properties are displayed as a single curve in the graph.

By default, the colors used to display the curves in the graph are Red, Green and Blue for X, Y and Z. You can change these colors by clicking on the color swatch alongside each animation property/component.


You can select single or multiple [keyframes](#) in the graph using standard selection tools. You can drag keyframes around in the graph to modify keyframe times and values (the keyframes of some properties do not correspond to “values” per se – for instance, a terrain geometry keyframe is not a “value”. For such properties, you can only modify the keyframe time).

You can add keyframes to a specific component of an animation property (e.g. add a keyframe to the X coordinate of the Position property) by selecting the appropriate component in the list, and selecting the **Add Keyframe** command from the popup menu in the graph.



The Animation Preview




The Animation Preview unfolded

The last part of the *Timeline* features a real preview render of your animation. Click on the  button to display it. Vue renders a small preview of the animation, and then plays it repeatedly inside a tiny window. Press Escape or click to stop playing the animation.

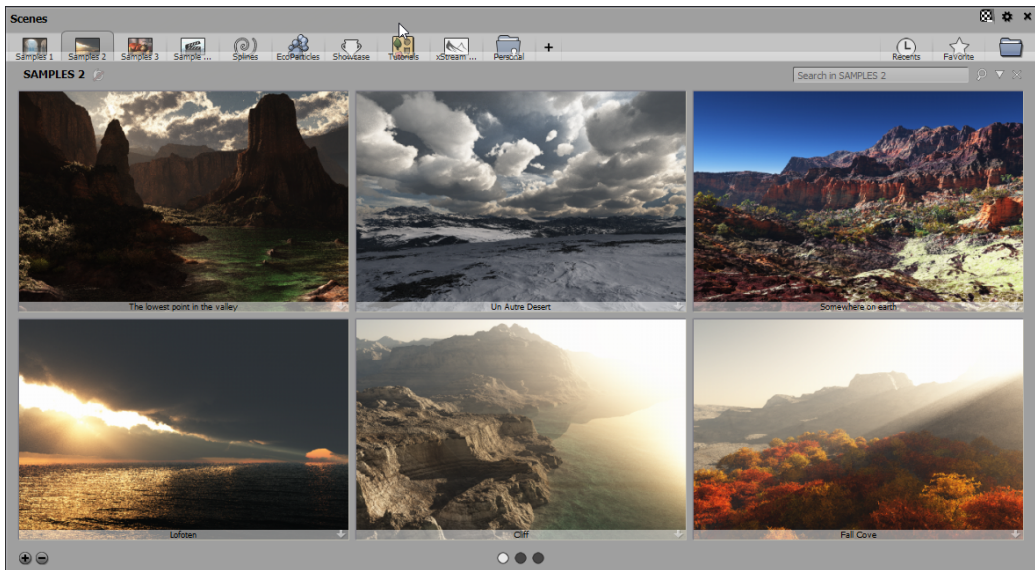
A set of tiny windows appears, displaying successive frames of the animation. The window that is framed by a solid line displays the frame at the current time. If you drag the Current time slider, the frames will scroll accordingly. Clicking on a frame sets the current time to that of the frame.

Press the **Render animation preview** icon () to refresh the preview. Activate the alternate action () of this icon to access the [Animation Preview Options](#) dialog.

You can play-back the animation preview alone (i.e. without animating the *3D Views* for smoother playback) at any time by activating the alternate **Play** icon ()




Visual Browsers



Selecting a plant using the plant browser

Whenever you need to select a file, Vue will display a Visual Browser to help you make your selection.

The Visual Browser lets you select files using preview pictures instead of filenames. Moving the mouse over one of the preview pictures (or clicking on the  icon) will display it a larger size next to the selected item in a popup, along with the title of the file and a short text description. The name of each item is displayed by passing the cursor over the item.

Underneath the text description, you will also notice the name of the file. This is useful when you want to delete a file, or move it to another folder.

To select an item, simply click on it.

If the items in this collection display on more than one page, use the dots at the bottom of the browser to display other items currently not visible.


Right-clicking an item will display a menu. You can Show a ToolTip, Add to Favorites list, Remove the image from the browser and Load the file. You can also delete items in



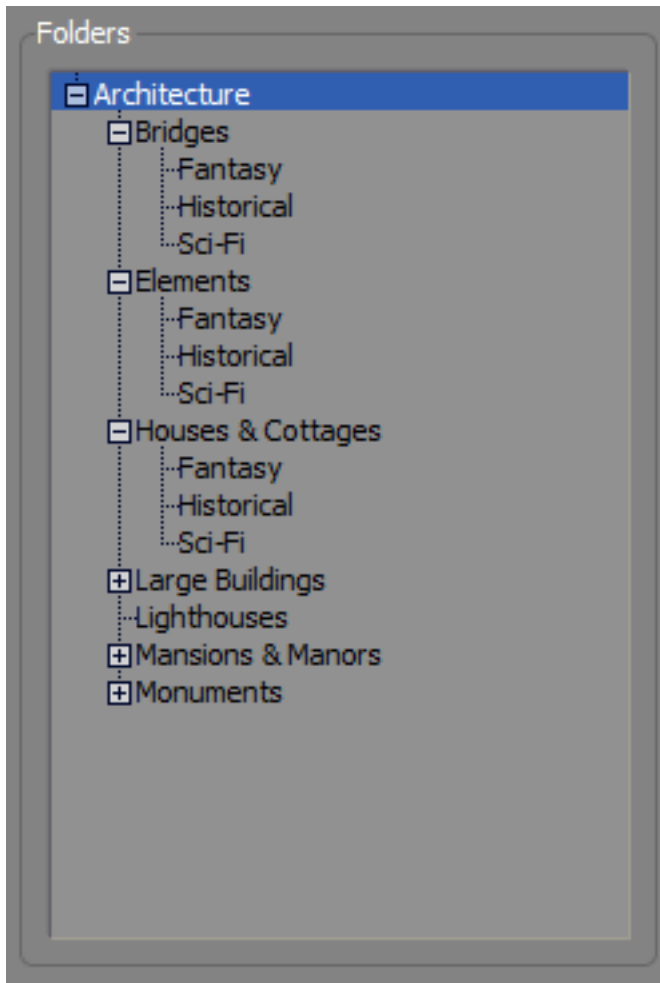
the list by highlighting them and hitting the **Delete** key. A prompt will appear asking if you want to delete the item completely (i.e. delete the file from your hard drive) or if you just want to hide the item from the collection.

Some visual Browsers are non-modal, so that you can leave them open and drag content from the browser into the scene anytime.

Collections

Inside the Visual Browser, items are organized in collections. The collections are displayed as icons () across the top of the browser. Running the mouse over the icons displays the name of the collection. The items that are currently displayed belong to the highlighted collection. If the collection has subfolders, these will be listed on the left side of the browser. Click on any folder in the list to display contents.





Selecting a subfolder from a category

In the lower left of the browser are plus (+) and minus(-) signs. Click on the plus sign to enlarge the pictures of the items in the browser; click on the minus sign to display those images smaller.

Collections are in fact shortcuts to directories containing items of the requested type. You may add as many collections to the list as you like by creating new folders in the root folder of the type of item displayed by the Visual Browser (e.g. **Materials** folder for Material Browser) or by clicking the plus sign icon at the end of the current list. Browse to the folder location on Standard File Browser that appears. Browse to the new collection's folder, and select the desired folder. You will be prompted for a name. This is the name



the collection will have inside the Visual Browser. Click **OK**, and wait for a couple of seconds while Vue builds the item previews and displays them).

If you add files to a folder that is listed by a Visual Browser as being a collection, the corresponding new items will automatically be added to the previews.


Collections may be removed from the list by right-clicking the icon for that collection and select **Hide Collection**. A prompt asks you to confirm the removal of the currently highlighted collection.

Note:

Removing a collection does not delete any actual files from your hard disk.

The contextual menu of the tabs can also change the icon of a collection or move the tabs right or left.


Loading other Files

If the file you wish to load is not displayed in any of the available collections, you may access it directly by clicking the folder icon at the top right of the dialog(). A Standard File Browser will appear, letting you browse and select your file.

You can bypass the Visual Browser completely when loading items by checking the *Use System Browser* option in the [General Preferences](#) tab of the Options panel or in the browser options.

Search

You can search/filter items by using the search toolbar:

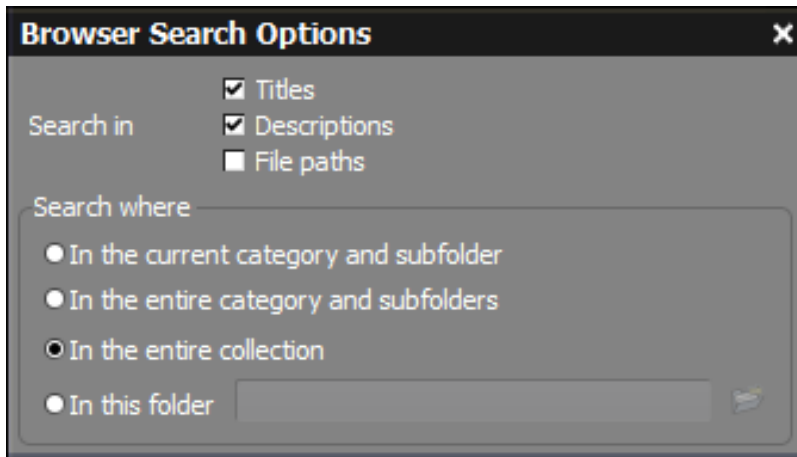


Typing on the text box will trigger a research on the whole collection on icon names. While there is a custom input on the text box, when you navigate inside the categories, the visible previews will be filtered by this text. You can clear the filter by clicking on the icon



The icon  opens the search option panel.






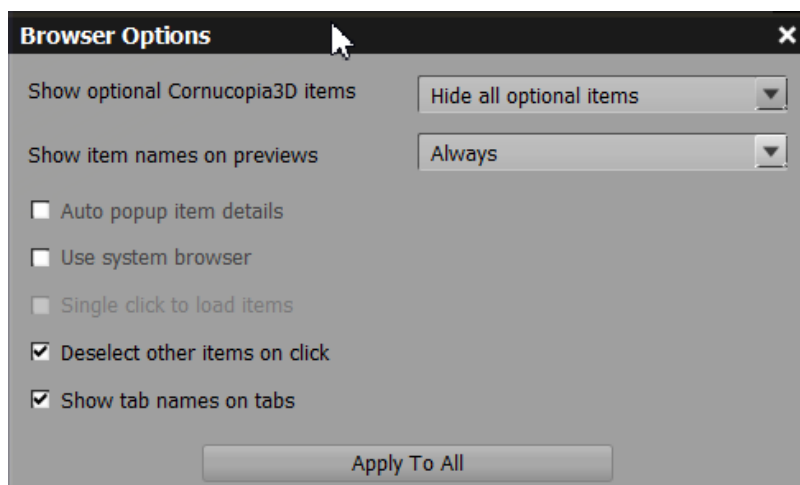
Search Option Panel

- **Titles:** Search on the titles of items
- **Descriptions:** Search on the descriptions of items
- **File Paths:** Search on the name of the items file on the hard disk.
- **In the current category and subfolder:** When you validate a research, you will search only on the current subfolder.
- **In the entire category and subfolders:** When you validate a research, you will search only on the whole selected category.
- **In the entire collection:** When you validate a research, you will search on all items in the collection, on any of the categories.
- **In this folder:** Select a folder to search on the hard drive.

Browser Options


There are some general options that can be selected by clicking the () icon. This displays the






Browser Options

- **Show Optional Cornucopia3D items:** You have three options for this option
 - Show all items : This shows all optional items available from Cornucopia 3D as well as those items that are on your hard drive.
 - Show most popular icons: This shows only a subset of items available from Cornucopia 3D.
 - Hide all optionals items: This hides all optional items available from Cornucopia 3D.

You can click on icon  just under the tabs to switch this setting temporarily without opening the options dialog.


- **Show items names on preview:** Filter if you want to see the name of files on the previews according to the size of the previews. You will still have the files name by moving the mouse over a preview.
- **Auto popup item details:** This will deactivate the auto tooltip on previews and show the  icon)
- **Use system browser:** This allows to bypass the visual browser when selecting an item to load. Your operating system browser displays instead. Since the visual browser will no longer show, the only way to revert this option is to uncheck “Use system browser” in the Options dialog, [General Preferences](#) tab.
- **Single click to load items:** You only have to click once to choose item.
- **Deselect other items on click:** When you click on an item in the browser,




any other items that have been checked are unchecked, or deselected. Having this unchecked may cause problems when selecting items for an EcoSystem. If you select one item, you will get all checked items in the browser into your EcoSystem.

- **Show tab names on tabs:** This displays the name of the categories on the tabs.
- **Apply to All:** All these options are defined for the current browser. This button apply the current options to all the browsers.

Directly below, is the Favorites icon (). Click to open your list of Favorites.


Next to that icon is the icon () that will display your recent choices.

In the lower right corner are the icons to increase/decrease the items' size in the browser ()


Virtual Collections

Some collections have a slightly different behavior, because files from these collections are not necessarily available on your Hard Disk. These collections are known as Virtual.

The whole purpose of Virtual Collections is to offer you a large selection of files, without requiring massive amounts of Hard Disk space.

When you select a file from a Virtual Collection, the following pictogram appears in the preview: . This means that the file you have selected was not physically copied to your Hard Disk at the time of installation. Instead, a reference was created to this file on one of the product's Extra CD.

If you load the file, a dialog box will appear instructing you to insert a given CD. Simply insert the CD in your CD ROM drive and press **OK**.

If you don't want to have to fetch the CD next time you use that file, select the option **Copy the file to my Hard Disk**. When you press **OK**, the file will be first copied to your Hard Disk, and then loaded into the program. Next time you select this file in the collection, you will notice that the  pictogram has disappeared from this file's preview.

You can also decide to copy all the files from the collection to your Hard Disk by checking the option **Copy all files in this collection to my Hard Disk**. When you press **OK**, all the files in the collection will be physically copied to your Hard Disk, and the collection will no longer be virtual.

If you don't want to use files in this collection any more, press the **Remove collection** button.

If you have to locate the file manually, press the **Locate file manually** button. A



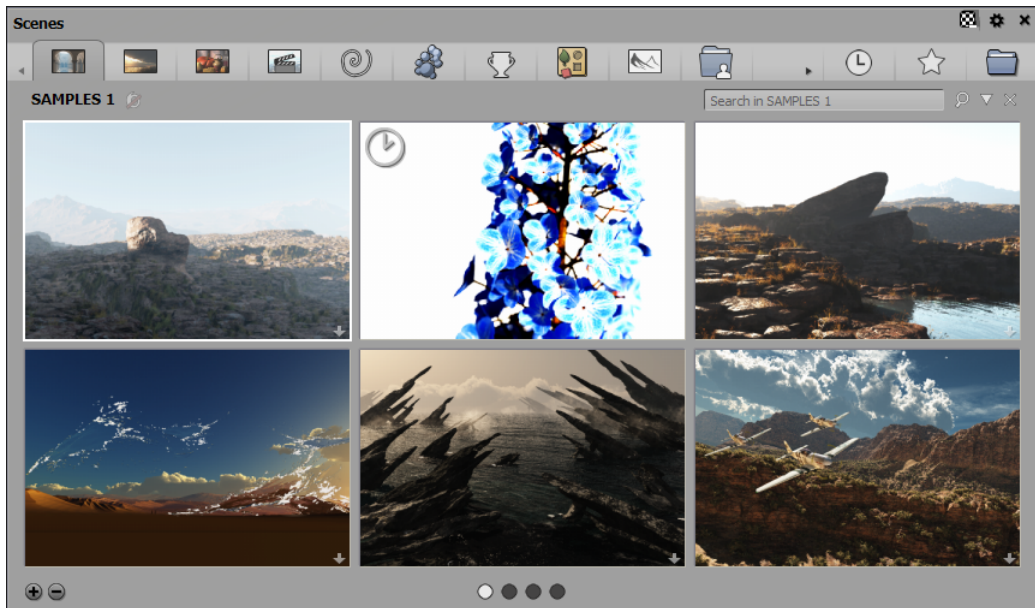
Standard File Browser will appear letting you browse to the location of the selected file.

Downloading Collections

If you wish to download all of the scenes (or objects), select the icon in the upper right of the browser. This will download all of the scenes to your hard drive so you won't need a continuous Internet connection to get the scenes, downloading one scene at a time.

Scene Versions

This section is specific to the *Scene* browser. When several backup versions of a scene exist, a pictogram representing a clock appears on its preview, to let you choose between the available backup versions.

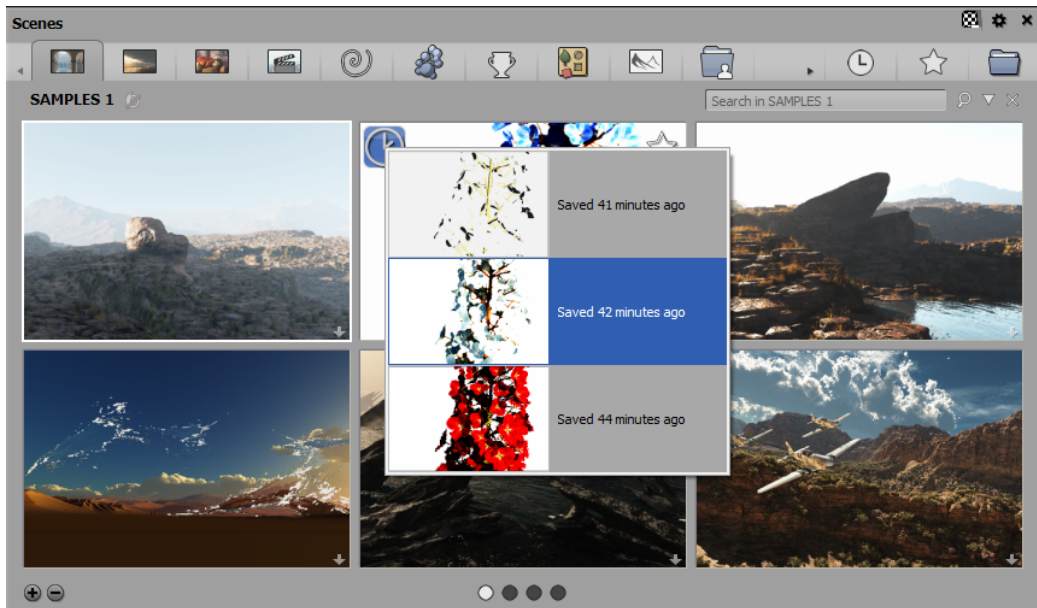


Easily identify scenes with backup versions

When you click on the pictogram, you can see the list of available versions, and select any of the latter.



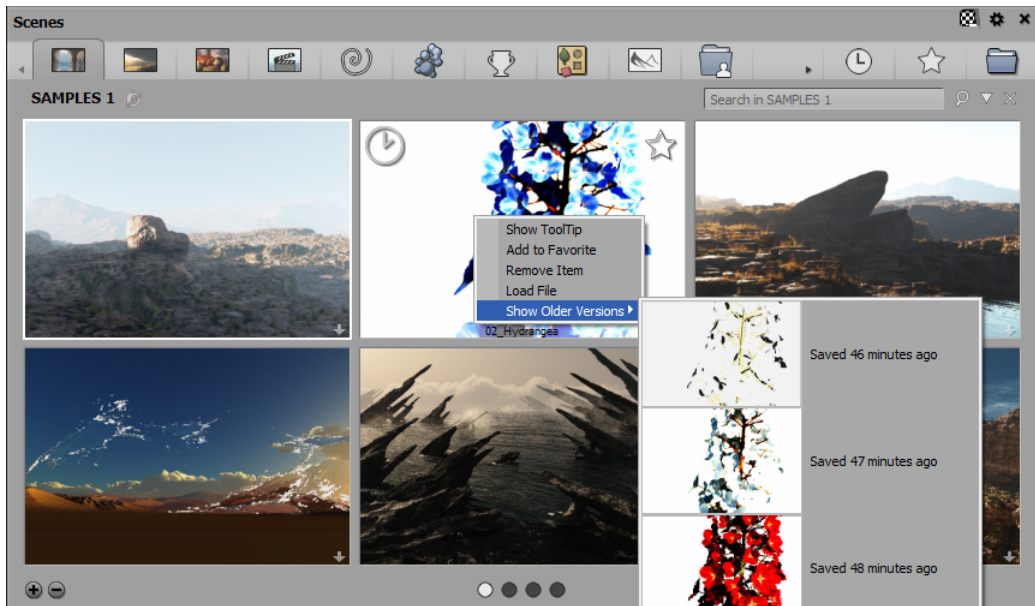
VUE 2016 Infinite & xStream — Reference Manual



Choose the backup version you want to load

You can also display this list from the contextual menu which pops up when you right-click over a scene preview.





Contextual menu for scene browser

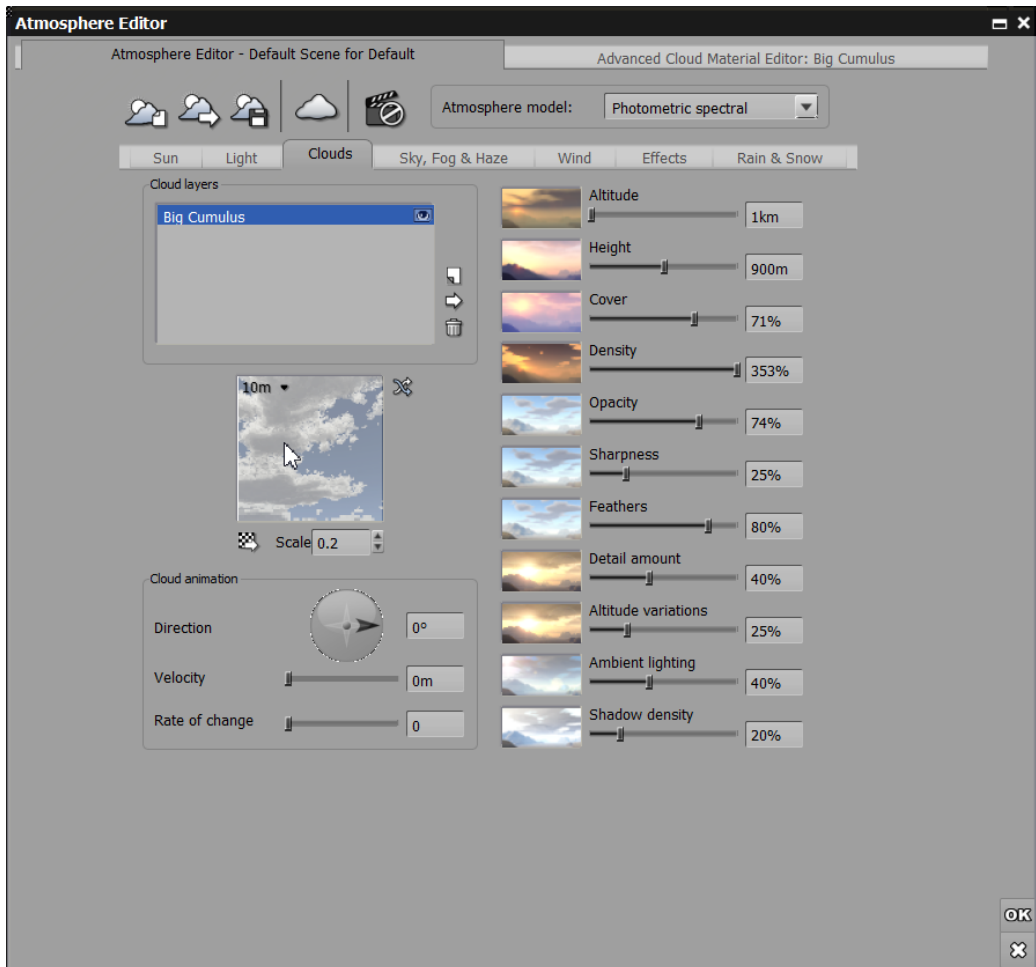
Note:

When you remove a scene from the browser using the *Remove Item* menu command, and choose to remove the file from disk too, all backup versions of the file will be removed from disk too.

Grouped Dialogs

Related dialogs or panels can now be grouped for easier use and the saving of space in the user interface.

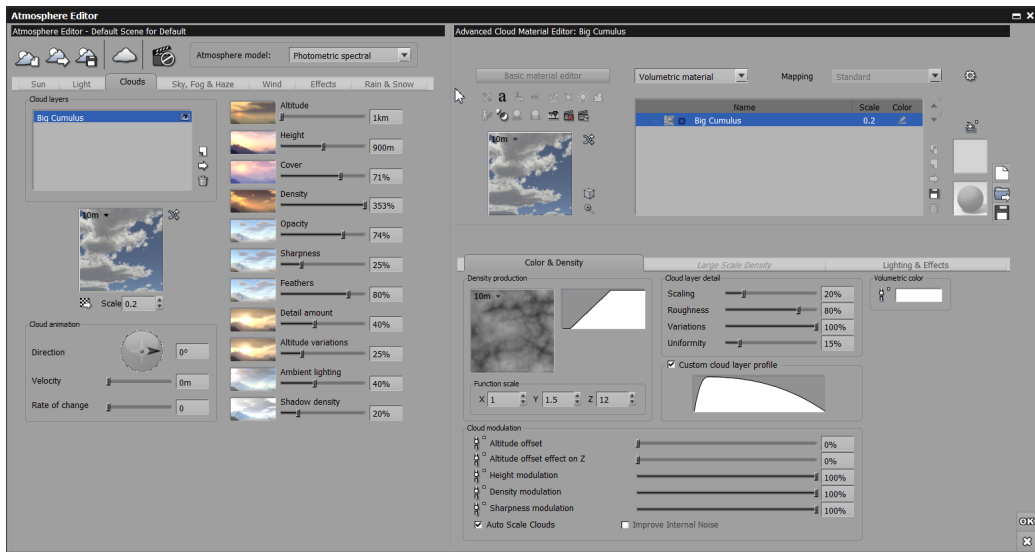




Panels grouped

When you are on a dialog (for example, the *Atmosphere Editor*), important sub-dialogs like the *Advanced Cloud Material Editor* can optionally open in the same window. By default, the sub-dialog will be tabbed with the first one. A dialog can have many sub-dialogs (*Atmosphere editor* can have *Advanced Cloud Material Editor* or *Function Graph* as a sub-dialog, or both).






Sub-Dialog attached at right

Another option is to attach the sub-dialog to the right of the first dialog.

Each dialog has **OK/X** buttons. If there is more than one dialog grouped together, every dialog will be affected by the **OK/X** on the main dialog. In other words, if you select **X** on the main dialog, all dialogs will close.



A menu can be accessed by right-clicking on the tab/caption of the dialog. You can choose to ungroup sub-dialogs by right-clicking on a tab and selecting to use the dialogs as floating (separate dialogs) or to have the second dialog attached to the right of the first dialog. If, after you've selected to ungroup the dialogs and use them as floating, you wish to group them again, this is set in the main or first dialog.



HiDPI

HiDPI, or High Dots Per Inch, is a high res display available on PC's with Windows 7 and higher. It's the PC equivalent of the Mac Retina screens.


This means VUE scales the user interface with Windows scaling.

To activate, open the Windows Control panel and select Appearance and Personalization>Adjust screen resolution. Changing a value on this screen will make the VUE user interface larger by the same ratio. (With Windows 7, 125% and 150% are proposed).

Windows 10 latest updates allows to have a setup with screens with different DPI (example: a normal screen and a 4K screen). In that case, VUE will rescale automatically the dialog that changes screen.

Saving Scene

There are 3 modes to save scenes:



- **Save as...** is the default mode (Fast-save mode, it creates a .fsvue file): This mode is the fastest ; it should be used during the creation of your project, as long as you edit it alone, on your machine. Note that this mode replaces user snapshots which used to exist before VUE 2016 Release 3. Unlike scene snapshots, which were lost when VUE was closed, .fsvue scenes are persistent. They are based on the same incremental system as snapshots (heavy data such as EcoSystem populations are not written each time you save the scene, except if they were actually modified). Those .fsvue files do not incorporate texture maps or EcoSystem variations. Therefore, you should use another mode to share a projects with other artists (see below)

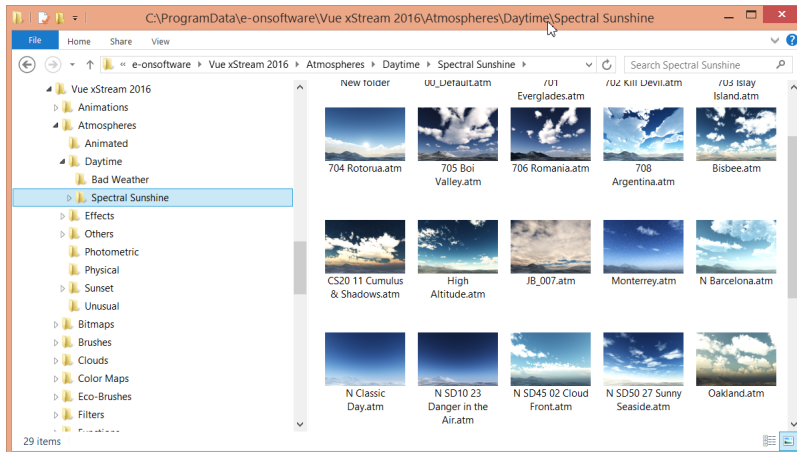
Warning:

.fsvue files are not compatible with previous versions of VUE.



- **Consolidated Archive:** incorporates all dependencies so that the scene can be easily shared with other artists. This mode will produce .vue files, which are compatible with previous versions of VUE.
- **Network Rendering:** should be used for scenes which are about to be rendered on a RenderNode network: in addition of all dependencies, it also incorporates data which will speed up the rendering process (.tpf plants species render cache, typically)

Thumbnail Previews in Windows



Folder View

VUE now displays its content directories as preview thumbnails. Even without VUE being launched, you are able to see the preview of VUE files. Also, when opening a native browser window inside VUE, you can still see the preview, but it won't be in a special control at the bottom of the window, it will be directly integrated inside the window.

You must have the option to **Use System Browser** selected on the *Options* panel, *General Preferences Tab*.



How to use it

Open Windows Explorer and navigate to a folder containing a format supported by VUE. The example shown here is a folder of atmospheres. When using the Windows Explorer launched by VUE, do not forget to change the display mode to '*Small/Medium/Large/Extra Large*' icons to see the previews of the files.

At this time, this feature is only available for Windows. It will be available later for the Mac.

VUE's Installed Content

VUE installs with materials, atmospheres, terrains and other content. The default location for this content is now installed in these locations.

On the Mac, the default location is `/Shared/Vue xxx/`

On the PC, the default location is `C:/ProgramData/e-onsoftware/Vue xxx/`

The content is installed with subdirectories for `/clouds`, `/materials`, `/objects`, and so on.

The location of these directories can be changed at installation.

If you should use the default location for the installed content, on the PC, this location is a hidden directory/file. So to see the directory of installed items while browsing in the Windows File Explorer, you must have the option to see these hidden directories/files turned on in Windows. To turn this on, on the Windows File Explorer window, from the menu, select **Tools>Folder Options**. On the *Folder Options* dialog, select the **View** tab. On the list that displays, under **Files and Folders**, select **Show hidden files, folders, and drives**. Click on **Apply**, then **OK**. Then you will be able to see this contents directory.

During installation, you also have the option of creating a new user content directory. This is in addition to the VUE installed content directory. Or you can point to an existing directory that you use for your content.



VUE's Sample Scenes

Pre-2016 version, you received a .zip file with all of the sample scenes included. These scenes are now available from the Scene Browser. You just have to click on the scene you wish to use and that individual scene is downloaded to your hard drive. You can then save it to your hard drive if you wish.

If you do wish to download all of the sample scenes at once, there is a “checkerboard” icon in the upper right corner of the browser. Click that to download and install all samples scenes to your hard drive.



Understanding Vue

Operations

Clipping Plane

You may have noticed that objects below the ground are not visible in the main 3D view. This is because they are automatically clipped below the ground, to improve the visual comprehension of the view.

This feature can be turned off by using the [Options](#) dialog in the View Options section of the Display Options tab. It may also be activated in the orthogonal views using that same dialog.

To decide which plane is going to be the clipping plane, Vue looks through all of the layers of the scene, to find the first infinite plane that is perfectly horizontal, and facing upwards (so if you rotate the ground plane so that it is no longer horizontal, it won't be a valid clipping plane any more, and will not be retained). If you add water to the scene, another infinite plane will be created above the ground. It will be placed in the *World Browser* just before the ground, thus becoming the first available clipping plane. This is why everything under the surface of the water is now invisible (which is usually what is wanted). If you drag the ground above the water in the *World Browser*, the ground will become the new clipping plane again (e.g. when you want to see underwater).

It is important that you understand how Vue decides which is the clipping plane, to avoid getting confused after having inadvertently manipulated the infinite plane that happens to be the current clipping plane.

Note:

Do not confuse this with the OpenGL clipping planes that define the range of minimum and maximum distances from the camera of objects displayed in the OpenGL *3D Views* (this can be adjusted using the [Display](#) tab of the *Options* dialog).



Drag and Drop

Drag and drop is everywhere in Vue:

- you can drag a material onto objects inside the *3D Views* or the *World Browser*, or onto other materials,
- you can drag materials onto several objects using the Shift key to extend selection,
- you can drag materials from the *Summary of Materials* onto other objects or materials,
- you can drag functions, filters and color maps from one material to another,
- you can drag items from the Visual Browsers into your scene, for instance, drag materials from the Material Visual Browser onto objects, or drag objects from the Object Visual Browser into the scene, and
- you can drag objects around inside the *World Browser*, into groups, out of groups, duplicate them, move them from one layer to another.


Popup Menus

Popup menus (also known as contextual menus) are available almost everywhere in Vue. A popup menu is a menu that lists operations relative to the item under the mouse cursor. It is called by clicking on the desired item with the right mouse button (Ctrl + click on Mac). In the 3D views, you have to be careful not to move the mouse in between the moment you press the button and the moment you release it; if not the views will be moved and no popup menu will appear.



3D Coordinates

World Space



The three axes of world space

To understand how objects are positioned and oriented relative to each other, we have to define a coordinate system. In Vue, this is (very classically) constructed from 3 axes, all at right angles from one another.

The center of the world, also known as the origin, is the point located in the middle of the orthogonal views, when you create a new scene. All positions are indicated relative to this point. The ground is also initially positioned at altitude 0.

The vertical axis is known as the Z axis, with positive numbers representing points above the ground, and negative ones representing points under the ground. Although this may seem unusual, it is the correct International Unit System.

Looking at *Top view*, the X axis is the horizontal one, with positive values representing points to the right of the origin, and negative values representing points to the left of the origin.


Also in *Top view*, the Y axis is the vertical one, with positive values representing points above the origin in this view, and negative values representing points under the origin.

Note:




If you are familiar with other coordinate systems, you can configure Vue to use another ‘up axis’. This is available on the Options panel, Units & Coordinates Tab.

Object Space



World space



Object space

Same material, different space


Object space is linked to an object, and is independent from the orientation of the object inside the World. You may rotate, stretch and twist an object in any way you like, object space will still indicate the same axes for the object, because it is relative to that object.

It is important to understand the difference between the two coordinate systems, and when each one is used. The *3D views* operate only in world space, that is, object independent coordinates. If you resize an object using the resize handles inside the 3D Views, the resulting object depends on the objects orientation: try creating a cube, resizing it vertically, and then rotating it 45°. Now create another cube, rotate it 45°, then resize it vertically. The result is not the same.




The **Numerics** tab of the *Object Properties* panel operates in object space. Try repeating the above operations, this time using the Numerics control. As you will notice, the results are identical in both cases.

Material Mapping Coordinates



Standard




Cylindrical

Same material, different mapping


Materials can be expressed in either coordinate system. Imagine you have a cube, made out of a black and white checker pattern. If you rotate the cube, the checker pattern won't fit the object any more. The solution consists in defining the material as being in Object space, so that the axes of the checkerboard match those of the object. Obviously, this is not what you would want all the time.

Both of the coordinate systems can be represented in either one of 4 modes:





Standard



Spherical

Same material, different mapping

- **Standard:** this is the standard (Cartesian) coordinate system, where X and Y represent the coordinates of the point in the horizontal plane, and Z represents the vertical elevation.
- **Cylindrical:** X represents the distance to the vertical axis, Y represents the angle (in the horizontal plane) of the line that joins the point to the origin, and Z represents the vertical elevation. Cylindrical mapping is best suited for cylindrical objects.
- **Spherical:** X represents the distance to the vertical axis, Y represents the angle (in the horizontal plane) of the line that joins the point to the origin (the heading), and Z represents the pitch of that same line. Spherical mapping is best suited for spherical objects.
- **Parametric:** in this mode, the mapping coordinates are automatically adjusted in such a way that they are independent on the size of the object. This mode is particularly useful when mapping e.g. a picture on a cube, because resizing the cube will not affect the number of times the picture is mapped on the cube.

Since cylindrical, spherical and parametric mapping modes are computed relative to the



origin, they give best results when expressed in Object coordinates (because in these coordinates, the origin is the center of the object).

Ray Tracing

Ray tracing is probably the most powerful 3D imaging technique for rendering photo-realistic pictures. It's incredibly complex too! The major drawback of this technique is long render times. It's the usual tradeoff: quality or speed.

Ray tracing was not invented for computers, as most people think, but for capturing the physics of rainbows. And this was back in 1637! The idea is that a picture is the result of the interaction of light with the objects that build up a scene. The optical principle behind ray tracing is that light follows exactly the same path, whether it travels forwards, or... backwards. So, what ray-tracing does, is follow all those light rays reaching the camera back through the scene, up to the light source. For each and every pixel of a picture, the render engine will trace the ray of light back into the scene, decide what objects were hit by the ray, bounce that ray back if the surface is reflective, transmit it through the object's surface if it is transparent, and eventually head for light sources, making sure nothing is getting in the way, that would shadow the object... A simple scene made out of a water plane and a ground plane requires tracing 7 rays per pixel ! (And this is if you are not using any elaborate features, like soft shadows, blurred reflections or transparencies, depth of field, super-sampling; you could go up to hundreds...).

This complexity is the reason why ray-tracing yields such beautiful results. And also why it is so slow.

For Vue, we wanted the best. So we chose ray-tracing... However, numerous optimizations and special ways of displaying pictures as they render have enabled us to produce what we think is a render engine fast enough to be enjoyable.

Render Engine

This is the process that converts the mathematical description of 3D scenes into 2D pictures. It scans each line of the picture, deciding what should be the color of every pixel of the line. When **Tile rendering** is selected it proceeds in several passes, doubling the resolution of the picture at each pass. The first pass renders groups of 16×16 pixels, the second groups of 8×8 , the third 4×4 , the fourth 2×2 , and finally, the last pass renders the picture at full blown resolution. In doing so, the user rapidly sees his picture take shape and can decide early if the result is going to be as expected, or decide to stop



the process by pressing **Escape**. The render engine may be customized to a large extent through the *Render Options* dialog. This is explained in detail in the section on Framing and rendering.

Soft Shadows



Standard ray traced shadows



Soft, natural shadows

Soft shadows

Real shadows seldom look as sharp as their computer generated counterpart. You always find smooth transitions between light and shadow; the farther the object casting its shadow, the smoother the transition. This is because light emanates from a surface rather than from a single point.

Such subtle effects can be captured by Vue.



Ray-Traced Soft Shadows

When the render engine must decide if a point is standing in the light or in shadow, it casts a ray at the light source. If the ray reaches the source without hitting anything opaque, the point stands in the light, otherwise, it stands in the shadow. To generate soft shadows, the render engine sends a bunch of rays, aimed at the whole surface of the light. It then compares the results to decide if more rays are needed. When enough rays have been cast, it computes an average luminosity for the point. The quality and efficiency of the result depends on that of the algorithm that decides how rays are cast. This technique is so powerful that it captures soft shadows cast by one complex object (e.g. a tree) onto itself.

Soft-shadows are turned on, on a “per light” basis, using the Object Properties panel. Use them with care, due to the extra computation they require. 5° is a good **Softness** value.

Shadow Map



Shadow Mapped shadows

While ray-traced soft shadows have the distinct advantage of being technically accurate, they have a severe drawback: they are very time consuming – especially if you want to get rid of noise artifacts, because you will have to increase the number of samples taken for each light. This is why Vue offers the shadow map alternative.

Shadow maps are a good way of rendering smooth soft shadows with very little noise and, most importantly, they are significantly faster to compute. To be more precise, the efficiency of shadow maps when compared to ray-traced shadows increases with the scene's complexity. This is an important rule to keep in mind. Of course, shadow maps have drawbacks too: performance is obtained at the cost of a loss of accuracy.

There are several ways of finding a good compromise between performance and accuracy when using shadow maps.



Here is a basic description of how shadow maps work: the shadow map is an approximation of the scene representation from the light's point of view. The light's field of view is divided into cells, each cell containing information about encountered objects. The limited number of cells reveals the main shadow map weakness: accuracy. Too few cells will result in a coarse approximation of the scene (and shadows will therefore suffer from a lack of precision), whereas too many cells will result in a very accurate representation of the obstacles encountered by rays leaving the light source, but will result in a huge memory consumption as well as a drastic loss of performance.

Vue uses an advanced proprietary shadow mapping technology known as *AccuShadows™*. This technology is similar in some ways to the deep shadow mapping technology. It allows for the efficient processing of shadows created by transparent objects as well as directional light sources (e.g. sunlight).

The behavior of shadow maps can be customized using the [Shadows tab](#) of the Light Editor.

Area Lights


Unlike simple light sources, area lights emit light from all points of their surface. Because objects placed in front of area lights will only block part of the light emitted by the area light source, they will create nice soft shadows. Area lights are very useful to recreate smooth, real-life-like lights. They are unfortunately a lot slower to compute than the other types of light.

The [Lights](#) section has more details on area lights.

Blurred Transparencies and Reflections



Standard ray-traced reflections



Blurred reflections. Note how the farther tropic is more blurred

2 tropics reflecting in a cylinder

Once again, the standard ray-traced, mirror-surfaced sphere (on a checkerboard?) is a poor approximation of what would really happen in nature. Perfect reflections or transparencies rarely occur. Instead, as reflected (or transmitted) light travels further, it spreads out, due to imperfections on the surface of the objects. Therefore, the reflected (or transmitted) picture eventually gets blurred. Although some natural surfaces are really smooth, this effect always ends up happening, as light travels further.

Vue can also capture this extremely subtle effect, in much the same way as that used for generating soft shadows. Instead of casting one reflected (or transmitted) ray, it casts a whole bunch, spread out at an angle dependent on the quality of the surface. It then decides if enough rays were cast, and, if so, computes the average color. As for soft shadows, the quality and efficiency of the result depends on that of the algorithm that decides how rays are cast.

Blurred reflections and transparencies are turned on, on a “per material” basis, using the **Reflections** and **Transparency** tabs of the *Material Editor*. 5° is a good value. You can even decide to vary the amount of blurring over the surface of the material! Use blurring with care, however, due to the extra computation it requires.



Reflection Maps



Spring with Reflection Map for faster reflections

Computing true reflections can sometimes be very costly, especially for blurred reflections that require distributed ray-tracing (see above), which can suffer from noise artifacts in addition to a strong increase in rendering time. In order to speed up rendering of reflections while getting rid of noise artifacts, Vue offers the possibility to use reflection maps.

A reflection map is a bitmap that is used as if it were mapped onto a static sphere enclosing the scene. Thus, for reflective objects using a reflection map, this bitmap will be “reflected” at their surface instead of the real surrounding environment. This feature can speed-up rendering time amazingly, but should only be used when reflections don’t need to display the true surrounding environment. For instance, this technique is particularly useful for chrome-like objects which must exhibit a reflective behavior without necessarily having to produce truly accurate reflections in order to look realistic.

Like [Shadow maps](#), reflection maps trade-off picture quality for improved render times. Reflection maps usually produce satisfactory results when used in animations.


Reflection maps can be assigned globally, or can be used on a per-material basis:

- To find out how the default reflection map is assigned, turn to section [Default Reflection Map](#).
- To assign a reflection map only to a given material, edit the material and use the [Reflection Map](#) settings in the *Material Editor*.


You can also force the render engine to use [reflections maps](#) instead of ray-traced reflections for all objects using the **Force use of reflection map** option in the *Render Options* dialog.



Depth of Field



clear version



blurred version

Realistic depth of field

Depth of field is what makes some parts of a picture blurred, while other parts, in focus, are clear. It defines the depth of the band inside which objects are in focus. This band is centered on the focus distance. All real-life lenses have depth of field. Even the human eye. This is because the depth of field of a lens is related to the proportion of light it lets through. This proportion of light coming through the lens is called the aperture. The higher the aperture, the shorter the depth of field.

Note:



In the pictures, look how the black sphere appears in the blurred picture. No amount of post-processing applied to the clear picture will ever capture this.


With computer generated pictures, luminosity is not really a problem, since the lens is defined mathematically. Hence, it becomes possible to have a very large depth of field, so large, in fact, that everything is always in focus. Although it is straightforward, it is not natural... And don't think realistic depth of field can be achieved by blurring parts of a clear picture in a paint program: realistic depth of field can only be generated at render time, because of the way light is spread by the lens (see sample renders on the right).

Vue can also capture realistic depth of field. This is done by taking into account a real lens aperture, rather than considering it to be a pinpoint. Instead of casting one single ray through the center of the pinpoint aperture, a bunch of rays are cast from all over the aperture, and traced into the scene. Results are then compared to decide if more rays are required. If not, an average value is computed and displayed.

Depth of field is turned on by increasing the **Blur** value from the *Object Properties* panel, when the camera is selected. The focus distance is adjusted using the **Focus** control. Beware, depth of field will make render times several times longer...

In Vue, rendering of depth of field can be done using either a [standard distributed ray-tracing approach](#), or a [Hybrid 2.5D approach](#).

Motion Blur



An example of motion blur.

Motion blur is a natural phenomenon that makes rapidly moving objects look blurred when they are photographed. The reason why this happens lies in the principles of traditional photography: light enters the camera through the lens, passes a shutter and exposes the film. The longer the shutter remains open, the more light reaches the film. So that it is exposed properly, the film requires a given amount of light that depends on its sensitivity. Which means that the shutter must remain open some amount of time, usually no longer



than a small fraction of a second.



Giving a sense of speed to a still picture

However, this fraction of a second is long enough for a rapidly moving object to change its position between the start and the end of the exposure. This is why it appears blurred. You can see on the previous picture an example of motion blur on a shiny black sphere moving from left to right. Note the blurred highlight and shadow.

Although this may sound like a defect, it is actually very close to what happens inside the human eye: since the human brain can only process something like 24 pictures per second, rapidly moving objects are seen blurred. If you present the human eye with a set of perfectly sharp pictures of an object that moves rapidly, the resulting animation will be perceived as flickering. Because the eye is used to seeing rapidly moving objects blurred.

Vue handles motion blurring as efficiently as possible. It does however adversely affect render times. To simulate motion blurring, Vue samples the scene at many different times inside the frame, (obviously concentrating its efforts on moving objects). It then averages the results of these samples to produce the desired blurring. Vue takes the simulation of motion blur very far, since every aspect of an animation supports motion blurring (including animated materials), resulting in perfectly smooth animations.

Note:

Motion blurring is only available from **Broadcast** preset render setting upwards (refer to the [Render Options](#) section for further details on preset render settings, and on how to turn Motion Blur on).



Simulating motion blur is the only way to provide smooth animations. But it is not only useful for animations. It can also dramatically increase the realism of still pictures, giving an impression of speed to moving objects (check the *Animation Wizard* introduction picture opposite for an example).

In Vue, motion blurring can be achieved using a standard distributed ray-tracing approach, or a Hybrid 2.5D approach. Please read below for a discussion on both approaches.

Hybrid 2.5D Blurring

Distributed Ray-Tracing

Motion blur and depth of field are features that are usually rendered with distributed ray-tracing. This method is a physically accurate approach that converges towards the exact solution as the number of samples per pixel increases. Unfortunately, if the number of rays is not sufficient, distributed ray-tracing suffers from noise artifacts, especially in areas that display a large amount of blur. To reduce this noise, a solution is to boost quality by increasing the number of samples per rendered pixel. This method works but can lead to dramatically slow renderings. To avoid such a loss of performance, Vue offers the possibility to compute an approximation of the motion blur and depth of field effects, known as Hybrid 2.5D blurring.


Hybrid 2.5D

Hybrid 2.5D is a method that is totally noise-free and much faster to compute. Also, computation time is not very dependent on scene complexity. The one drawback is that this method is not as physically accurate as its distributed ray-tracing counterpart.

In most cases, for complex scenes, using hybrid motion blur or depth of field rather than distributed ray-tracing will speed-up your renderings to an amazing extent. Of course, since it is built on approximations, this method has a few limitations that will be detailed below.



Limitation of Hybrid 2.5D



Same example with Hybrid 2.5D Motion Blur

To render motion blur, Vue uses a Hybrid 2.5D technology that produces smooth results very quickly. Basically, Vue looks at the speed at which all the pixels in the image are moving (or how blurry they are), and spreads them out according to their velocity (blur level).

This method of rendering motion blur is very efficient, but it has some limitations. The main limitation of Hybrid 2.5D effects is that information regarding what is lying behind blurred objects is missing. This is why Hybrid 2.5D blurring will always appear stronger than ray-traced blurring.

You can reduce this defect by increasing the number of passes. This will reconstruct a reasonable amount of information about what is taking place in the background – especially for depth of field, as blurred objects in the foreground will cover significant areas of the background. This will however slow down rendering time accordingly.

Generally speaking, rendering strong depth of field will require a large number of passes to produce accurate results. On the other hand, motion blur, which is usually subtle, doesn't require as many passes (actually, one pass is usually sufficient). Keep in mind that rendering with 10 passes will take exactly five times longer than rendering with 2 passes.

This is an important point: avoid using a too large number of passes.

Another limitation is that Hybrid 2.5D motion blur and depth of field don't work well together. If you decide to use both effects together, you might get totally unexpected results. This is due to the fact that for performance reasons, each blurring effect has its own approximation method which is not compatible with the other.

When your scene exhibits both motion blur and depth of field, it is recommended that you use distributed ray-tracing rather than Hybrid 2.5D.

Finally, concerning motion blur, you might notice that for moving objects that cast shad-



ows on the scene (on the ground, for instance), these shadows will not be blurred. This is unavoidable as a projected shadow not only depends on the object that casts it, but also on the receiving object that is not necessarily animated. Moreover, it is not always possible to tell which object is responsible for a specific shadow.

This is not true for camera animation. When the camera is moving, the whole scene looks like it is moving. As a result, shadows will be blurred as expected. More generally, Hybrid 2.5D motion blur shows better results when associated to a camera movement, especially camera rotation. Scenes with animated camera orientation are the best candidates for Hybrid 2.5D motion blur.

Enabling Hybrid 2.5D Blurring

Hybrid 2.5D is automatically selected when you enable Broadcast or Superior [rendering presets](#). Ultra uses Distributed ray-tracing instead.

Hybrid 2.5D can be enabled and adjusted using the [Blur Rendering Options](#) dialog.

Fast Hybrid 2.5D


The **Fast Hybrid 2.5D** option uses a new algorithm for depth of field generation. It is based on image blur like Hybrid 2.5D but uses a faster color spreading algorithm and works in conjunction with distributed ray tracing. Usually several passes are required to get all of the distributed ray-tracing noise smoothed out.

Systematic object anti-aliasing is incorporated inside Fast Hybrid 2.5D. Therefore anti-aliasing settings become linked to the depth of field settings. This means that only systematic anti-aliasing becomes available, and the minimum number of rays per pixel becomes equal to the number of depth of field passes (changing either of them changes both values).



Light Related Features

Lens Flares



Lens flare caused by looking at a bright light

Lens flares appear when the camera is pointed towards a bright source of light. They are caused by light rays being reflected and refracted inside the lens – or the eye – and are generally considered by photographers and directors as an unwanted effect. It is precisely in order to avoid – or minimize – this defect that camera lenses are often equipped with sun visors.

However, this defect is also a great addition to Computer Graphics, because it adds a touch of imperfection to an otherwise too perfect look, thus largely enhancing realism.


Because lens flares happen inside the camera, they are not affected by the rest of the scene. They won't appear in the reflections of other objects either.

In Vue, lens flares are generated at the time of rendering, which means that you get a very quick impression of what the result will be, without having to wait for the render to complete. Please read [here](#) for details on how to setup lens flares in your scenes.



Glow

Glow is a haze of light that appears around certain objects. It is an interesting feature that can help you achieve amazing effects. It is generally useful when representing very hot materials that will illuminate the air around them. Although glowing materials seem to illuminate their whereabouts, they do not cast any actual light. Glow works best with self-illuminating materials (i.e. luminous materials).



A glowing sphere placed behind a checkerboard sphere

Glow is a post-processed effect meaning that it is added after the rendering of the 3D scene is complete.

Because of this, glow has a number of limitations:

- The most important one is that you won't be able to judge the effect, until rendering is complete.
- The second limitation is that there will be no glow displayed if the glowing object is completely hidden behind another object. Glow will begin to show as soon as a tiny part of the glowing object becomes visible. As a result, a glowing object that disappears behind another one will see its glow vanish suddenly, whereas you would expect the glow to disappear more gradually.

If these limitations are not acceptable, you can reproduce the effects of glow by using volumetric materials/lights. However, you must realize that setting up the effect is going to be a lot more complex, and render times will be notably increased.



Volumetric Effects

Volumetric effects, common in everyday life, add a degree of realism to your scenes. Instead of computing the interaction of light and materials at the surface of objects, volumetric effects will compute these interactions inside the entire volumes.


Because of this added dimension, and the resulting increase in complexity, volumetric effects can play a significant role in creating ever more realistic environments. The counterpart is a dramatic increase in render times...

Volumetric and Spectral Atmospheres

Vue's [volumetric](#) and [spectral atmosphere](#) models are incredibly elaborate models that compute the interaction of light with the various types of particles in suspension in the Earth's atmosphere all along the path followed by the light.

For each ray of light that it processes, Vue evaluates the density of all the components of the atmosphere (humidity, gases) along the ray, and calculates the corresponding scattering of light. These results are then integrated to produce such realistic effects as the reddening of sunlight close to the horizon.

Volumetric Lights



Beams of light created by a group of spheres in a volumetric light

Volumetric lights will give physical volume to the beams of light. If an object is placed in front a volumetric light, the shadow of the object will be visible inside the beam of light.

Volumetric lights are particularly well suited when trying to achieve a dramatic environment.

In real life, beams of light are created by light bouncing off the surface of numerous tiny particles in suspension in the air (e.g. dust, smoke, ...).

When computing a volumetric light, Vue determines whether each point inside the volume



is actually exposed to the light, or is being shadowed by an object placed between it and the source. Vue evaluates the brightness of the air at any given point by combining the exposure of all the points along a ray.

These lights can be customized using the [Volumetric Light Options](#).

Godrays

Godrays appear when the sunlight is partially hidden by thick, obscure clouds. This effect is particularly noticeable when facing the sun, and the cloudscape is dense, but with holes in it.

Just like in the real world, godrays only happen under very specific conditions – they are not an easy effect to achieve.

Although they are similar, godrays and volumetric lights are two different types of effect in Vue. Godrays take into account the shadowing produced by clouds – and only by clouds – whereas volumetric lights take into account the shadowing produced by other objects in the scene – and not the clouds.

Godrays are only possible when using the Spectral atmosphere model. They are enabled using the [Godrays](#) option in the *Atmosphere Editor*.

Volumetric Materials



Sample use of volumetric effects

Volumetric materials are extremely useful when recreating objects that don't have a well established frontier, such as clouds of dust, gas, smoke, etc.

When rendering [volumetric materials](#), Vue evaluates the density of the material in many



points along each ray, and computes the resulting global density and lighting.

Volumetric Clouds

Volumetric clouds are very similar to volumetric materials in the way they are processed. For increased realism, they use a special algorithm that takes into account the internal lighting of the clouds.

Spectral Clouds


Spectral Clouds are a particularly advanced type of volumetric effect that is specifically optimized for the rendering of extremely realistic clouds. They take into account the subtle interaction of light with the water particles that form the clouds. Spectral clouds are used in spectral cloud layers as well as the standalone MetaClouds.

You can edit [Spectral Clouds](#) through a specific subset of volumetric materials in the *Material Editor*.

Hypertextures

Hypertextures are an interesting hybrid between solid materials and volumetric materials that are very well suited for rendering porous materials (such as sponge or corroded metals). They are created using a specific subset of volumetric materials in the *Material Editor*.

Caustics



Caustics for different IORs – left to right: Water, Glass, Crystal

If a transparent material has a greater density than air (i.e. an Index of Refraction greater than 1), it bends the rays of light that cross its surface. This bending not only occurs on what you see through a magnifying glass, it also occurs on rays of light coming from



a light source and crossing the magnifying glass; it focuses the rays of light onto a given point (this is how you burn a piece of paper using a magnifying glass, by concentrating all the light onto a small area). Because all the rays of light are focusing onto one point, all the other points behind the magnifying glass get darker. The total amount of light behind the glass is still the same, but its distribution changes.


The fact that light behind a transparent object is being concentrated onto some areas rather than being distributed equally over the surface of the shadow is called a caustic.

The higher the index of refraction of the material, the more concentrated the light will be; so the brighter the central spot, and the darker the rest of the shadow. This behavior is correctly simulated by Vue.

Physically Accurate Caustics

When rendering caustics, you can instruct Vue to generate physically accurate caustics. To generate the physically accurate caustics, Vue follows the path that light would follow inside the material, in order to determine the exact location where light is focused.


The processing of physically accurate caustics is a lot more complex than the default caustic effect, but can produce interesting results when rendering detailed refractive materials. Physically accurate caustics will also capture caustics created by reflective materials (such as in the metal ring opposite).



Caustic refractions in a glass prism



Reflective caustic



Caustics with dispersion

Physically accurate caustics

Reflective caustics happen when light gets bounced off the surface of reflective materials. Reflective materials reflect the light in a preferred direction. As a result, depending on the geometry of the object, light will either be “focused” by concave areas of the object, creating caustic “hot spots”, or “scattered” by convex areas of the object.

In the ring image opposite, you can see a caustic “hot spot” at the center of the ring, where light is concentrated by the concave interior of the ring – and you can also see a slightly brighter area around the outside of the ring, where light, reflected by the convex exterior of the ring, is scattered.

Note:


Because by construction the surface of infinite planes is infinite, infinite planes cannot generate physically accurate caustics. When such an effect is required, you should use standard planes instead. While infinite planes cannot generate accurate caustics, they can still receive them from other objects.

Dispersion

Dispersion takes place when the different wavelengths in the light are not refracted the same. This results in the well known spectrum effect after light goes through a crystal prism. This effect is simulated by Vue.



Ambient Occlusion, Global Illumination and Radiosity




Standard render with 15% sunlight

In the standard ray-tracing lighting model, objects receive light from the different light sources placed in the scene. If an object is not directly lit by one of the sources of light in the scene, it is considered uniformly dark. Obviously, this is not the way light behaves in the real world, because objects that are not directly lit still receive light from the other objects around them. To compensate for this fact, the standard ray-tracing model adds a uniform light known as ambient light. Unfortunately, this only constitutes a very crude approximation of the way light is scattered around in the real world.

A number of techniques have been developed recently to improve this basic lighting model. These techniques are usually referred to under the generic term of “Global Illumination” techniques. Unfortunately, these techniques work best when the scene is confined in a relatively small area – typically not the case when rendering outdoor scenery. Our engineers have put a significant amount of effort into developing lighting models that are capable of simulating the subtlety of natural light throughout extensive outdoor scenery. The result of this research has been included in Vue as 3 different models of increasing accuracy. Unfortunately, the more refined the model, the slower the rendering speed. This is why, despite the use of highly optimized algorithms, when using the Global radiosity model (which is the most accurate model) you should expect render times that are at least an order of magnitude slower than with the ray-tracing model.






Global Illumination render. Notice the shadows around the object

The simplest of these models is actually only a refined version of the ray-tracing model, where the intensity and color of the ambient light is defined by the color of the sky around the objects. As a result, objects that are in the shadow and looking away from the sun may take a different color tone than objects pointing towards the sun. This model is known as **Global Ambience** and is the fastest to compute among the advanced lighting models.

Vue's **Ambient Occlusion** and **Global Illumination** models are more accurate than global ambience in that they take into account the light cast by the sky onto the scene, as if the sky were constituted of an infinity of small colored lights. The shadows cast by each one of these little lights is taken into account, resulting, among other realistic effects, in very soft shadows appearing around the base of objects.

The main difference between Ambient Occlusion and Global Illumination is that, in the Ambient Occlusion model, only objects that are close enough together will cast soft shadows on each other. As objects get further apart, the shadowing effect diminishes until it totally disappears at the Ambient Occlusion range. Because Vue only searches for occluding objects within a given range, rendering with Ambient Occlusion is usually faster than with Global Illumination. Ambient Occlusion also leads to lighter shadows.

Despite the use of very advanced optimizations, the rendering of all these subtle shadows results in much longer render times than with the standard ray-tracing model. The effects of global illumination are particularly noticeable when there is a lot of ambient light in the scene.



Global Radiosity render. Notice how the picture is brighter due to color bleeding.

The main limitation of the Global Illumination model is that it really only scatters shadows around objects. To model the full subtlety of natural light, you need to scatter light around objects. This is what the **Global radiosity** model does by taking into account the light cast by each object onto all the other objects in the scene. With this model, a bright red object will cast red highlights onto other objects around it. The radiosity model is obviously much slower to compute, but will yield incredibly pleasant and realistic results, and is essential to achieve convincing interior renders.

The other drawback of global illumination techniques is that they are often complex to setup, with dozens of exotic parameters that you can act upon to improve quality or speed. Because we think that artists should not be asked to cope with such complex parameters, our engineers have developed a unique technology known as *EasyGI™*. What this does basically is allow you to adjust the quality of the global illumination renderings with one single slider!

These lighting modes are activated on the [Light Tab](#) of the Atmosphere Editor.

HDRI and Image Based Lighting



Typical Image Based Lighting render with a reflective sphere.

HDRI stands for High Dynamic Range Image. It is a special picture file format (**.HDR**) that is capable of describing pictures with very important variations in light intensity. In other words, this file format is capable of storing pictures with details in the very high lights as well as in the very dark areas. Imagine a photo that would be so bright in places that it would actually light your face as you look at it.

HDRI pictures are created by combining several identical pictures taken at different exposures. Although HDRI images could be any shape and could be used to map any object, they are usually designed to be mapped on a sphere and used as an environment map.

One of the typical applications of HDRI images when setup as an environment map is



to use the lighting information in the picture to illuminate the scene. Each pixel in the HDRI picture is then considered as a source of light and traced into the scene to determine illumination. This is known as **Image Based Lighting**, and is a technique commonly used by the motion picture industry to ensure that the lighting of a CG scene matches exactly that of the real environment (captured as a HDRI image).

Setting up a scene for Image Based Lighting (IBL) can be a little tricky, but Vue actually lets you do this within a few clicks! Also, because Vue lets you combine IBL with its own atmosphere engine, you can create stunningly realistic pictures easily (by matching the atmosphere of the scene with that of the background image you will avoid the typical IBL image discrepancies between the CG rendering and the atmosphere seen in the background image).

Setting up a scene for Image Based Lighting is done from the [Effects Tab](#) in the Atmosphere Editor.

Illumination Baking



Sample baked illumination rendering

[Illumination baking](#) is a very useful technique when rendering static objects in global illumination or global radiosity. What illumination baking does is create a separate texture channel for the baked object, evaluate the amount of indirect lighting that reaches the object at each point of its surface, and store this illumination information in the texture map.

That way, after the illumination has been computed and stored once, it can be retrieved as many times as necessary, very quickly. If you are rendering an animation, this will result in a dramatic increase in rendering speed. This is typically the case when rendering architectural projects. If some objects in the scene are animated, the baked illumination may not be correct throughout an entire animation (because indirect lighting is influenced by surrounding objects). However, depending on the level of lighting accuracy required by your projects, this may not be an issue.



Illumination baking can also be a fantastic time saver when working on stills. Because creating a satisfactory still usually involves numerous renders, you end up evaluating indirect lighting numerous times. By baking the illumination once and for all, you can save a lot of time in subsequent renders.



Part of illumination map for above sample

Because only indirect lighting is taken into account when baking illumination, the position and intensity of lights can be changed without affecting the validity of the baked illumination (when using the global radiosity lighting model, this is only true to the extent that illumination of envrioning objects is not affected too severely by the changes in lighting conditions). Direct lighting (including shadows) will be recomputed as expected.

Illumination baking only works with polygon meshes. If you wish to bake the illumination of a non-polygonal object, you will first have to convert the object to polygons (see [here](#)). If you want to bake the illumination on the ground plane around an object, you should create a plane below the object and bake the illumination on that object. Baking illumination on the entire ground plane will lead to poor results (due to the size of the ground plane).

You can export baked illumination information together with objects.

Sub-Surface Scattering



Rendered without





*Rendered with
subsurface scattering*

Translucent materials react to light in a very different way than “regular” materials. With a regular material, incident light is either diffused, reflected, or refracted. With translucent materials, the light is also absorbed by the surface of the material and re-emitted at a point that is not the same as the point where it arrived.

This results in very subtle light showing up in parts of translucent objects that would otherwise be in the shadow. Typical translucent materials are marble, jade, human skin, fruit flesh, milk, orange juice, etc.

The technology used to render translucent materials is known as sub-surface scattering. In order to properly render translucency, it is important to take into account both aspects of sub-surface scattering: absorption and multiple scattering.

Absorption

Also known as “single-scattering”, absorption is what happens when light travels through a translucent material and bounces as it hits imperfections in the material. Depending on the properties of the material, this light will bounce either backwards, forwards, or in all directions. The light picks up the color of the material as it goes. In the case where light is bounced “mostly forwards”, this effect will be visible when the source of light is placed behind the translucent object (typically what you see when you place your finger on a bright light: a red glow appears on the thinner parts of the finger). It’s also absorption that makes a wax candle glow as it’s illuminated by the flame above.

Multiple Scattering

With multiple scattering (sometimes also referred to as “diffusion”), light that penetrates inside the material bounces off all the little defects inside the material, in random directions. Eventually, part of that light makes it back to the surface again, at a point that is different from where the light penetrated the material.

While absorption is relatively easy to implement, multiple scattering is a lot trickier. Don’t be fooled by the simple interface: there is a lot going on under the hood when rendering multiple scattering, and memory requirements to properly handle this effect



can be staggering.

Translucency can be activated for your materials using [Material Editor](#) settings.

Displacement Mapping



Bump (left) vs. Displacement (right) Mapping. Notice how the displaced sphere appears larger.

Displacement mapping is the ultimate way of adding detail to the surface of your objects. Instead of only fooling the eye into seeing detail (by modifying the normal vector to the surface), displacement mapping really adds those details to the geometry of your objects. Because they are true geometry, those details will cast shadows and create very pleasing soft shadows when rendering with global illumination.



with Bump





with Displacement

Same brick material

The typical example of the use of displacement mapping is in the case of rendering a brick wall. If you use a bump map to simulate the fact that the bricks stick out of the mortar, looking at the wall from a low angle of incidence will show that the wall is flat. On the other hand, if you use displacement, the edge will no longer be flat, and you will see the bricks protruding from the wall.

[Displacement mapping](#) dynamically adds detail to the geometry of your objects. The higher the render resolution, or the closer you look at the objects, the more detail gets automatically added, so that your objects will always look fully detailed.

Advanced Effects Quality

The quality of the visual effects generated by Vue depends on the time the render engine spends rendering them. For instance, when rendering volumetric effects, the quality of the result depends on the number of samples evaluated for each ray in the image. If the number of samples isn't sufficient, noise will appear in the volumetric effects. If there are too many samples, rendering times will be horrible.

This is true for all types of volumetric effects, and, generally speaking, for all the advanced visual effects available in Vue.

In order to provide an idea of the results of advanced effects during the development process without slowing render too much, Vue automatically adapts the render quality of the advanced effects (and henceforth the time they take to render) to the overall Render quality setting (see the [Render Options](#) dialog). Usually, when you design a scene, you will be working in “Preview” render quality (the default). At this setting, Vue generates



a very rough approximation of the advanced effects that will look very noisy or crude, but will render quickly. As you switch to higher render quality settings, the rendering of the advanced effects will automatically improve and the “noisiness” will disappear.

The quality of individual advanced effects can be adjusted by a **Quality boost** setting. If you increase the quality boost, more time will be spent generating that particular effect, so it will look better, but the render time will increase...

Note:

it is generally bad practice to increase the quality boost setting of advanced effects to achieve perfect results in “Preview” mode, because as you switch to higher quality settings to produce the final image, the advanced effects quality will automatically increase too, resulting in an excessively high amount of time being spent rendering those effects. It is better to adjust the quality boost settings of the advanced effects once you switch to the final render quality you are going to use to output your work.

EcoSystem Technology




Typical EcoSystem forest

EcoSystem™ is the name of e-on software’s revolutionary set of patented technologies to distribute, manage and render millions of instances of plants or objects in your scenes. With this technology, you are able to recreate the millions of plants, trees and rocks that are required to create convincing environments.

EcoSystems behave like materials in the sense that they define the aspect of an object. What this means is that, whether you want grass or sand at the surface of an object, the interface is the same (or at least very similar). Like other materials, *EcoSystem* materials are assigned to objects. When you assign an *EcoSystem* material to an object, Vue will populate the surface of the object with instances of the EcoSystem population. The



EcoSystem population is the most important aspect of an *EcoSystem* material; it lists all the elements that will be distributed at the surface of your objects. These elements can be plant species, imported objects, rocks... They can even be animated elements!



EcoSystem based on imported objects

The *EcoSystem*TM technology uses advanced instantiation techniques to distribute the elements of the *EcoSystem* population over the surface of your objects. When creating *EcoSystems* based on plants, the technology is combined with *SolidGrowth*TM to create a wide variety of plants of the selected species.

The fact that *EcoSystem*TM works with instances means that you can get tremendous visual complexity at a relatively low cost in terms of system resources: the actual geometry of the objects is not duplicated. Instead, “virtual” copies are used (this is not the case with rocks and *SolidGrowth* plants where variations are automatically created).

*EcoSystem*TM uses a patented rendering technology that dramatically speeds up the rendering of the millions of instances required to create convincing environments. Without this technology, it would not be possible to render such scenes in a reasonable amount of time.

Like other types of materials, *EcoSystem* materials are based on functions and parameters that control the aspect of the *EcoSystem*. The most important parameter (aside from the actual type of elements to be placed at the surface of the object) is density. The density parameter controls the number of instances of the *EcoSystem* population list that will be placed at the surface of the object. Items are automatically distributed at the surface of objects in a realistic way (no items are placed where they could fall down!).

New to the third generation of the *EcoSystem* technology is the ability to dynamically populate infinite expanses. And the latest addition to *EcoSystems* is faster processing times for smaller instances, for example, gravel.

[EcoSystem materials](#) are created and modified in the Advanced Material Editor.



Framing and Picture Composition

In this section, you can read about the basics of framing and picture composition. If you are well aware of these topics, you won't find anything there that you don't already know, so we suggest you simply [skip the section](#).

Making Better Pictures

This section doesn't aim to give you a list of rules that must absolutely be followed to guarantee successful pictures. And the reason for this is simple: such rules do not exist! Framing and composing pictures is an art form. And, as such, anything can be imagined; no rules will ever replace artistic talent.

However, what we intend to do is attract your attention to the importance of careful framing and composition, by detailing a series of concepts that should be understood and taken into account. Just like a photographer wanting to make a picture of a landscape would spend hours deciding on the best viewpoint, the best picture format, and the best time of the day, you, the virtual photographer, should think about this when you make pictures. With a tremendous advantage over a real photographer: you can act upon every aspect of the picture... Plus you can fly!

Choosing the Viewpoint


The viewpoint affects to a large extent the final result of a picture. In some cases, moving the camera slightly can completely modify the composition of a picture. A good way of improving the quality of your pictures is to take the habit of moving the camera around the scene, in order to find the best possible viewpoint. Storing several viewpoints (using **Display | Store Camera**) is also good practice, because it gives you the opportunity to compare several framing attempts.

Most scenes are composed of a foreground, a background, and a subject. Relative positions (and importance) of these elements can be modified by moving the camera around. You may choose to give more importance to the foreground, for instance by moving the camera



down. Or you could choose to display less sky by moving the camera upwards, while still aiming at the same subject... Of course, some subjects may offer several interesting aspects. In that case, it all depends on which aspect you think should be emphasized.

Moving the Perspective Camera



In the OpenGL view, when using *Main camera view* or *Perspective view*, it is possible to move the camera along the horizontal or vertical axis by holding Ctrl+Shift+Right mouse button (this shortcut can be changed in the *Options* dialog, **Operations** tab).

Which axis is used depends on the position of the mouse when the buttons are pressed, if the mouse is close to the window's horizontal central axis, movement will be constrained to this horizontal axis. Else it will be constrained to the vertical axis.

The following schema shows a window. Black lines are the axes of the window. If the buttons are pressed when the mouse is in the green area, movement will be constrained to the horizontal axis. If the buttons are pressed when the mouse is in the blue area, movement will be constrained to the vertical axis.



Picture Format

Unlike real photographers, users of Vue may choose the format of their pictures at any time. You should use this advantage to give more impact to your pictures. Picture formats are selected from the *Render Options* dialog (menu **Render** | **Render options**).

Usually, horizontal formats make for calm, peaceful pictures. On the contrary, vertical formats make for more dynamic pictures. If you are making a picture of a sunset on the horizon, a long stretched horizontal format will probably be the one you want.

But, of course, this is not the rule. Sometimes, best results are achieved by breaking preconceptions; framing a landscape vertically will display interesting details in the foreground and the background.

Center of Interest

Any picture whatsoever will always have more impact if the composition attracts attention to a particular element. This means that, before deciding on the way you will handle a subject, you should decide on the center of interest of the picture.


Sometimes, it can be quite straightforward: it could be a tree, a hill, or a lake. If you find no distinctive feature, keep looking, or make one. Once this research is complete, figure out how you could best exploit it, and avoid any other details that could compete with it. There are two ways of doing so; they can be used separately or combined.

- Frame out any details that could distract attention,
- Build up contrasts in color and shades between the main subject and the environment.


You could also use depth of field to blur out the background of the main subject (beware of render times though!).



Balance



Black dots mark the strong points



The principles of composition should never be followed to the letter. However, understanding them will help you make more pleasing and balanced pictures.

Following a basic principle, the main subject of a picture should be placed on a strong point. If you divide your picture into three vertical and three horizontal boxes, the strong points of the picture are where the lines (known as strong lines) intersect. De-centering the main subject onto one of these strong points usually yields good results, provided some elements are there to counterbalance (for the sake of clarity, this is not the case in the opposite illustration). However, if your picture only has one point of interest, it is better to nearly center it, because this balances a background devoid of any particular interest. Obviously, if the main interest of a subject is that it is symmetrical, intensifying that symmetry by precisely centering it yields best results.

To give impact to a picture, you can also use lines or colors to guide the eye towards the main subject. Converging lines will guide the eye to the point where they meet. You can achieve the same result by using a gradation in the colors: dark in the foreground, with colors becoming brighter as they get closer to the subject.

In the opposite picture, the castle in the foreground (i.e. the main subject), although de-centered, is counterbalanced by the opposing castle in the distance. The fact that the two castles seem to be trapped each in an opposite corner of the picture emphasizes the



rivalry between them.

Light and Shape

Lighting conditions affect the way we see shapes. Light helps the eye to perceive bumps, dips, curves and perspective. If you light up your subject from the back, you'll get a dark picture of it, with no relief. If you move the light so it illuminates your subject from the side, the gradation of shades produced will let your eye understand the surface of the subject.

Optimal intensity and direction of light depend on the subject, and what you want it to express. Early morning or late afternoon lights produce long, stretched shadows. They are best suited when you want to attract attention to details on the surface of objects. They also usually give warmer, more pleasing colors to the eye.

If you want to reveal soft, round shapes, such lights would not be suitable, because they increase contrast between bright and dark areas. You'd have to select a soft, diffused light, with a lot of ambient light, and probably soft shadows.

Patterns, Colors and Textures

Patterns are made of repeating shapes, colors or lines. They have fantastic visual power.

You can make patterns even more pronounced by using low-angled lights that will add a succession of bright and dark areas, and make the relief more visible.

With a well chosen light, you can render the roughness or the softness of any given surface, where other lighting conditions would only show a flat surface. If you want to insist on the texture (e.g. surface roughness) of an object, you should use low-angled lights. This makes bumps more pronounced, and will darken any recessed areas.

Powerful lights will hide small details but are particularly well suited for shiny surfaces.



The Power of Lines

Lines are often the basis of the composition of a picture. If you take a careful look at your scene, you will realize that some lines or strong contours give it strength.

Lines can be used to balance a picture by attracting the eye to the main subject, and by creating links between other objects in the scene, or they can create conflicts. With low-angled lights, strong shadows can give incredible power to a picture. Also, lines can give the idea of depth, due to perspective. The lines could converge outside the picture, thus giving the impression that they are escaping from the screen.

The shape of dominant lines often affects the atmosphere of a picture. Sharp edges and angles express energy and aggressiveness. Round curves express calm and peace...

Conclusion

Understanding, and taking in, these basic concepts will help you get more visual power out of your pictures. Experimenting is the key to success. Isn't it worth trying?



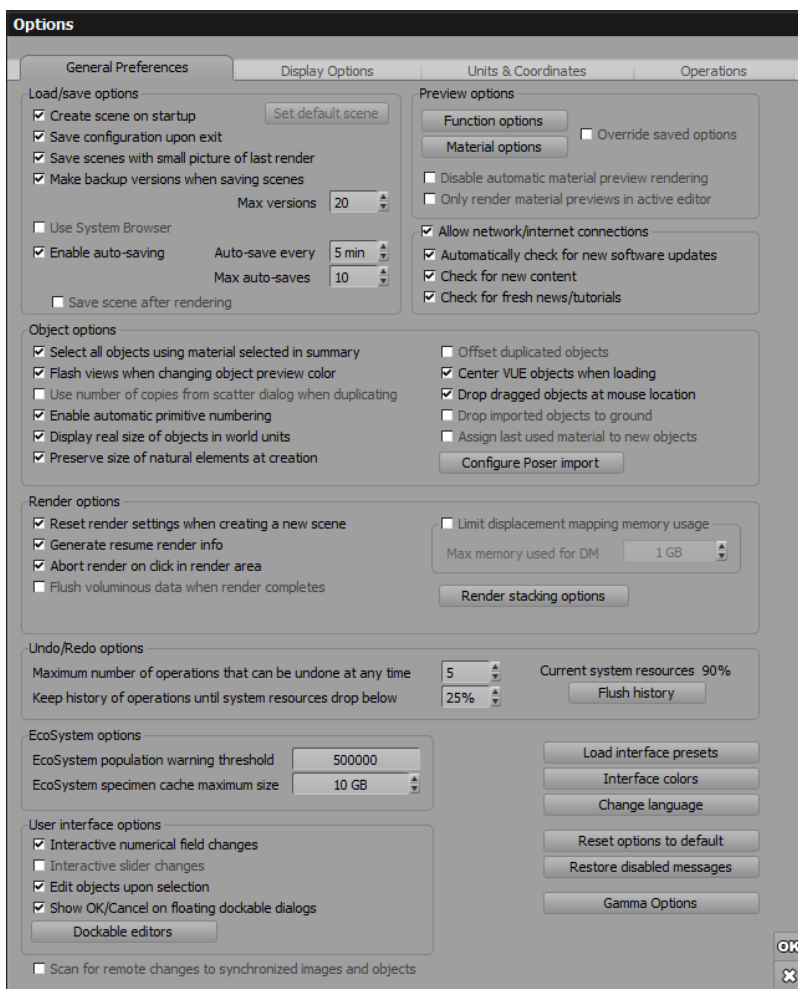
Options and Preferences

To display the *Options* dialog, select the menu command **File | Options**. This dialog lets you customize the way Vue behaves.

General Preferences



Load/Save Options



The Options dialog – General Preferences tab

- **Create empty scene on startup:** when this options is enabled, a default scene is automatically created upon startup based on the default atmosphere.

You can make the current scene become the default scene by pressing the Set default scene button. If you want to revert to the factory default scene, you will have to erase the file *Environment/Default Startup Scene.vue* .

- **Save configuration upon exit:** unchecking this will prevent VUE from saving



your configuration when you exit the software. It is not recommended.

- **Save scenes with small picture...:** when you save a scene, VUE stores a thumbnail preview of it inside the file. This is then used in the Scene Browser. Disabling this option yields black previews for all your scenes.
- **Make backup versions...:** unchecking this box prevents VUE from making a backup copy of your scene when you save it. Although making a backup is good habit, since it can avoid losing data, you may want to turn this feature off because of the extra disk space used. Backup copies have the **.bak** / **.fsbak** extension instead of **.vue** / **.fsvue**. Backup scenes can be restored from the menu **File > Previous Versions**, or directly from the [Visual Browser](#).
- **Max versions:** You can adjust the total number of backups of a given file using the **Max versions** setting.
- **Use System Browser:** when this option is selected, Vue will bypass the Visual Browsers completely whenever you want to load a scene or a picture. You can locate content using your operating system's browser. Each Visual Browser can also be individually configured to use the System Browser or not (see [Browser Options](#)). When the General Preferences' checkbox for "Use System Browser" is neither checked nor unchecked (it shows a light grey square instead), it means some but not all of the Visual Browsers have the option checked.
- **Enable auto-saving:** when this option is selected, a backup copy of your scene is made automatically, at intervals you can set. Intervals are from every minute to once an hour. You can also stipulate how many autosaves that will be made. When the maximum number of autosave files is reached, the oldest is deleted to create the new one.
- **Save scene after rendering:** this feature saves the scene when rendering is finished.

Autosave files are available in the menu **File | Previous Versions**. Autosave files are deleted when the current scene is closed. If you wish to keep an autosave, you have to revert your scene to that version and save the scene.

Preview Options

The controls in this group let you customize the previews of materials and functions, using the [Preview Options](#) dialog. If you select **Override saved options**, when you load a material or a function, these options will override those that were saved together with the material or function.

If you check the **Disable automatic material preview updating** option, material preview will no longer be generated automatically. Whenever a material preview becomes obsolete, a small triangle will appear on top of it. Simply click the preview to update it.



Depending on your work habits, this option can be useful if you find that your computer is having difficulties keeping all materials up to date as you work.

Checking the **Only render material previews in active editor** option is useful if your system become very slow when editing complex hierarchies of materials. When this option is selected, only the material or function previews that are in the topmost editor are refreshed.

Connections to E-on Software Website

The options in this group control the way Vue connects to your local network or the e-on software website. Vue regularly connects to the e-on software website in order to check for software updates, new content and fresh news/tutorials. Vue only connects to the e-on software website, and any exchange of information is both secure and in keeping with e-on software's privacy policy (you can review this policy at <http://www.e-onsoftware.com/privacy>).

- **Allow network/internet connections:** if you disable this option, Vue will never attempt to connect to your local network or the e-on software website. Network rendering will not be possible if this option is disabled. We recommend that you do not forbid network/internet connections, as this would adversely impact your Vue experience.
- **Automatically check for new software updates:** if you disable this option, you won't be notified when new updates are available for your software (new automatic updates are never installed without your prior approval). You should visit the e-on website regularly in order to perform software updates manually.
- **Check for new content:** if you disable this option, you won't be informed when new content is available for your software. You should visit the e-on website regularly in order to check for new content manually.
- **Check for fresh news/tutorials:** if you disable this option, you won't be informed of the latest news and tutorials for your software. You should visit the e-on website regularly in order to check for the latest news manually.



Object Options

- **Select all objects using material...:** when you click on a preview of a material inside the Summary of materials dialog, all the objects that use the material in the scene are selected. Unchecking this box will disable this feature.
- **Flash views when changing...:** when you select a new Preview color from the drop-down list of the *Object Properties* panel, the selected objects are temporarily deselected to show that they have taken on the new color. Uncheck the box if you don't want this to happen.
- **Use number of copies...:** when duplicating this option instructs Vue to use the number of copies indicated in the Scatter/Replicate Objects dialog when duplicating objects (i.e. if 10 copies is indicated and this option is selected, duplicating an object will actually generate 10 copies of it). Please read [here](#) for details on the *Scatter/Replicate Objects* dialog. This option is not set by default, and is provided for compatibility with previous versions.
- **Enable automatic primitive numbering:** when you create a new object, the name of the object automatically reflects the number of objects of the same type that have already been created in this scene (e.g. “Sphere 2” means this is the second sphere you create inside this scene. The first sphere may have been deleted. Copied / duplicated objects are not counted). Unchecking this box will disable this feature.
- **Display real size of objects...:** when this option is selected, the size of objects as displayed in the **Numerics** tab of the *Object Properties* panel indicate the actual size of the object in real world units. If you deselect this option, the size will indicate an internal value (this option is only provided for compatibility with previous versions).
- **Preserve size of natural elements at creation:** When this option is checked, all “natural” Vue primitives (rocks, plants, terrains, water surfaces) will be rescaled according to current internal unit settings upon creation, thus keeping a standard physical size.
- **Offset duplicated objects:** when you duplicate objects, the copy of the objects are offset by one nudge unit in each direction. You can disable this feature by unchecking the box. As a result, duplicated objects will be the exact copies (including position) of the initial objects. This can be useful when you are doing precise alignment.
- **Center Vue objects when loading:** this option will move Vue objects that you load, so that they always appear at the center of the views. If you uncheck this option, objects will retain their position as defined in the *.vob* file. It is useful to uncheck this option when loading several files that each represents a different part of a single object.
- **Drop dragged objects at mouse location:** If dragging an object from the Object



Browser, it will drop at the current mouse location.

- **Drop imported objects to ground:** when this option is checked, objects that you import will be automatically dropped to the ground plane.
- **Assign last used material to new objects:** If checked, any objects added to the scene will be assigned the same material as the first object added to that scene.

Poser Import Options

The **Configure Poser import** button opens the *Poser Import Setup* dialog. There are several Poser SDK's that you can choose from:

- **Poser SDK from September 2012:** This one adds compatibility with the SR 3 of Poser 9 and PoserPro 2012.
- **Poser SDK from October 2013:** This one adds compatibility with Poser Pro 2014.

You might need to experiment with these to find out which one works best for you.

You also need to browse to your Poser application that you will be importing from. This setting is required to import Poser content.

Render Options

- **Render small picture of objects...:** when you save an object, Vue renders a thumbnail preview of it, for use in the Objects Browser. Unchecking this will yield black previews for all objects.
- **Reset render settings...:** when you create a new landscape, render settings are automatically reset to default (e.g. Preview render quality). Unchecking this box will disable this feature.
- **Generate resume render info:** if this option is selected, Vue will automatically generate the information required to let you resume later a rendering session that you decided to interrupt. Because the generation of resume information can take some time, this information is not generated when rendering in **OpenGL** or **Preview** modes. Disabling this option will allow faster interruption of renderings by skipping the generation of the resume information in the more advanced rendering modes – but you won't be able to resume rendering. Please also note that any modification of the scene will make it impossible for you to resume rendering later. You can also control the creation of resume render information using the **Render | Generate Resume Render Info** menu command. This command is only available if an advanced render mode is selected (**Final** and up). The effect of the menu command



is identical to that of the checkbox. This option only pertains to renders to screen; this **Resume Render** is not available for network rendering.

- **Abort render on click in render area:** when this option is selected, clicking in the render area will abort the rendering process. If this option is deselected, you will have to press Escape to stop rendering.
- **Flush voluminous data when render completes:** during the rendering process, Vue generates a lot of data that can be preserved in order to speed up subsequent renders of the same scene. When this option is selected, this data will be automatically discarded when the render completes, thus freeing up memory.
- **Limit displacement mapping memory usage:** This allows you to specify the lower limit for memory cache used for displacement mapping. A lower limit means that swapping to disk will occur more frequently. So it sets the balance between in-memory and on-disk displacement data storage.
- **Render stacking options:** This button opens the *Render Stacking Options* dialog where you can specify what renders you wish to keep in the renderstack, sets the size and location of the render stack. These settings are also available on the [The Render Display Window](#)

Undo/Redo Options

The options in this group are used to control the number of operations that you can undo or redo.

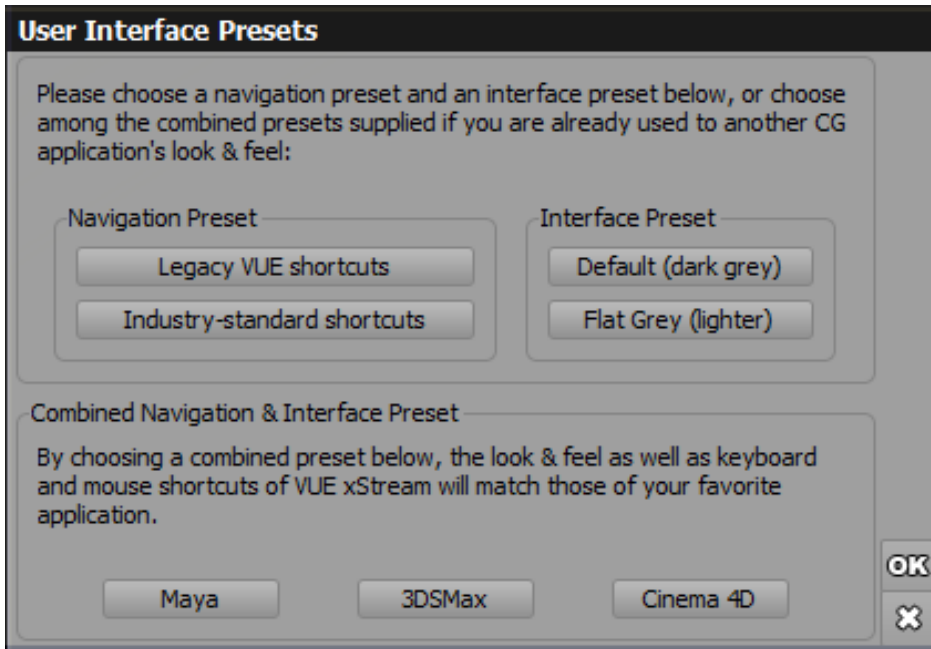
- **Maximum number of operations that can be undone...:** this setting controls the number of operations that you can undo by hitting the Undo button in the [Top Toolbar](#). Provided system resources are sufficient (see below), you will be able to undo/redo that number of operations at any time.
- **Keep history of operations until system resources drop below:** this option lets you configure the maximum system resources that can be used before Vue stops storing undo/redo operations. When system resources drop low, the number of operations that can be undone may be reduced. The default value is 25%; you shouldn't set this to 0%, as it will eventually lead to a total system lockup.

The **Current system resources** field displays the percentage of system resources that are available. If you have lots of undo/redo operations and system resources drop low, you can press the **Flush history** button. This will discard all undo/redo operations and free-up some system resources.



Configuration

- Load interface presets:



User Interface Presets Dialog

VUE's interface defaults to a dark grey theme, and can be changed to a lighter theme named "Flat Grey". It can also be automatically configured to emulate that of other popular 3D applications.

Select this option to display a dialog that lets you select the software interface you would like to emulate.

- **Interface colors:** click this button to display the [Interface Colors Editor](#) and customize the colors of the interface.
- **Change language:** click this button to change the interface's language. You will have to restart the application for the change to take effect.
- **Reset options to default:** click this button to reset all options to the factory settings. All modifications made to the *Options* dialog will be lost.
- **Restore disabled messages:** click this button to restore all messages that have been hidden by selecting the "Don't show this message again" option in the message box. This is particularly useful if you have created a default behavior for one of the messages by checking the option and want to change this default behavior.



- **Gamma Options...:** click this button to display the [Gamma Options](#) dialog.

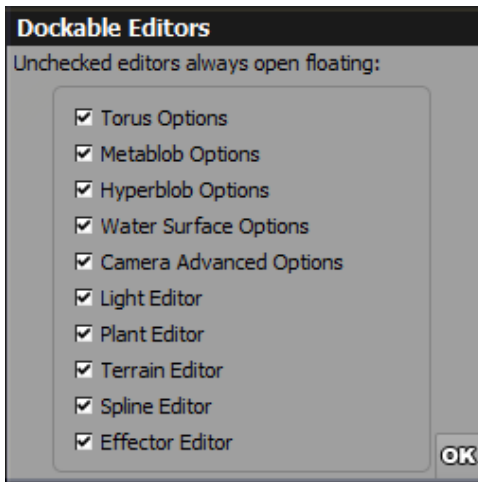
EcoSystem Options

- **EcoSystem population:** this allows you to set a threshold limit on an EcoSystem population. If an EcoSystem population is going to be greater than this number, a warning message is displayed.

User interface options

- **Interactive numerical field changes:** If this option is selected (the default), the interface is automatically updated as you enter numerical values in the input fields. For instance, if you enter a new position using the Numerics field of the *Object Properties* panel, the position of the selected objects will be updated as you type the new position. If you prefer that the interface is updated only when you have finished entering the value and press Enter or switch to another control, deselect this option.
- **Interactive slider changes:** If this option is selected, changes are made in synch with the slider movement.
- **Edit objects upon selection:** If checked, and you select an object is selected, immediately you go into edit.
- **Show OK and Cancel buttons:** in undocked dialogs. These buttons will display in the lower right corner if this option is selected.
- **Scan for remote changes images and objects:** When you enable synchronization of images and imported objects when you import either an object or a texture map, and check the **Synchronized** option in the 4th tab of the World Browser (*Links* tab), VUE periodically checks if the corresponding files have been modified by an external program, and proposes to reload the file in this case. This process can be really time consuming if you have hundreds of maps or objects in your scene, so this checkbox can be used to disable this process globally without having to edit each map or object previously selected in the *Link Tab* of the World Browser.

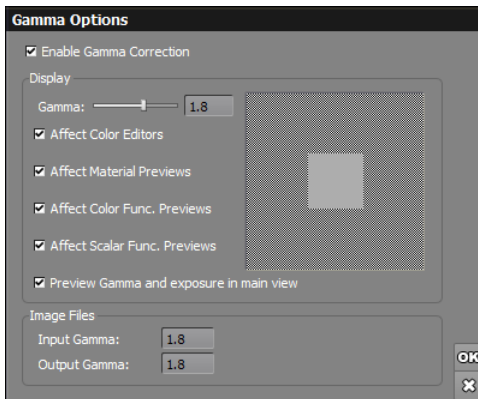




'Dockable Dialogs'

You can select the dialogs to dock or not dock. These dialogs can be selected on the **Dockable Dialogs** dialog.

Gamma Options Dialog



Gamma Options dialog

The Gamma correction features lets you specify the input and output gamma that Vue will use internally. The gamma correction is an operation applied to intensities, to transform them from a non-linear space to a linear one (and inversely).



Monitors, cameras, and all of the other devices that display or capture images are non-linear by nature. When displaying a mid-gray value on a monitor, for instance, the intensity of light that is actually displayed on the screen is not half of the one of a pure white. To compensate for this, a gamma correction is introduced so that the above becomes true.

In Computer Graphics, this has an important effect on the rendering engines, as all of the computations should really be computed in linear space, which can provide a more accurate realism when contributions from several lights are added, for instance.

You can choose to save the images with this gamma correction applied. But in some cases, you may want to keep the image in linear space so that you can open the image in another application and disable the input gamma correction in the input of the other application. Only at the end of the pipeline, you then add a gamma correction corresponding to the device on which you want to display your image.

Using gamma correction, you can be assured that all people working on a scene will see the same result, even when monitors have different behaviors.

When accessed from the *Options* panel, the gamma settings become global settings which can be changed for specific images using the *Advanced Camera Options* panel or the *Post Render Options* panel.

- **Enable gamma correction:** this is a global setting that enables/disables the entire gamma correction option.
- **Display:** using these controls you can specify the gamma correction that Vue needs to apply to the displayed images.
- **Gamma:** use the slider to setup the gamma correction by referring to the preview on the right. When the square area within the preview matches the intensity of the surrounding area, the gamma is correctly set. By default, the gamma correction is only applied to the render. Using the next options you can extend the gamma correction to:
 - The Color Editors
 - The Material Previews
 - The Color Function Previews
 - The Scalar Function Previews


Note:

The gamma correction is not applied to the OpenGL preview.



Image Files

- **Input Gamma:** the gamma correction that is applied to all of the bitmaps that will be loaded in the render engine. For any texture map node, even one eventually connected to the color output node, evaluating through either its grayscale or alpha output won't apply gamma correction. Only its color output will be gamma corrected.
- **Output Gamma:** gamma correction that is applied to the render output.

Once you have set up these options, you can still choose to change the gamma on a per-image basis in the case of **Input** textures. If you want to override the gamma for a specific texture map, you can switch from the system gamma (set up here) to a specific setting by choosing the **Override Gamma** option () located to the right of the image preview in the *Material Editor* or any texture map nodes in the *Function Graph*.

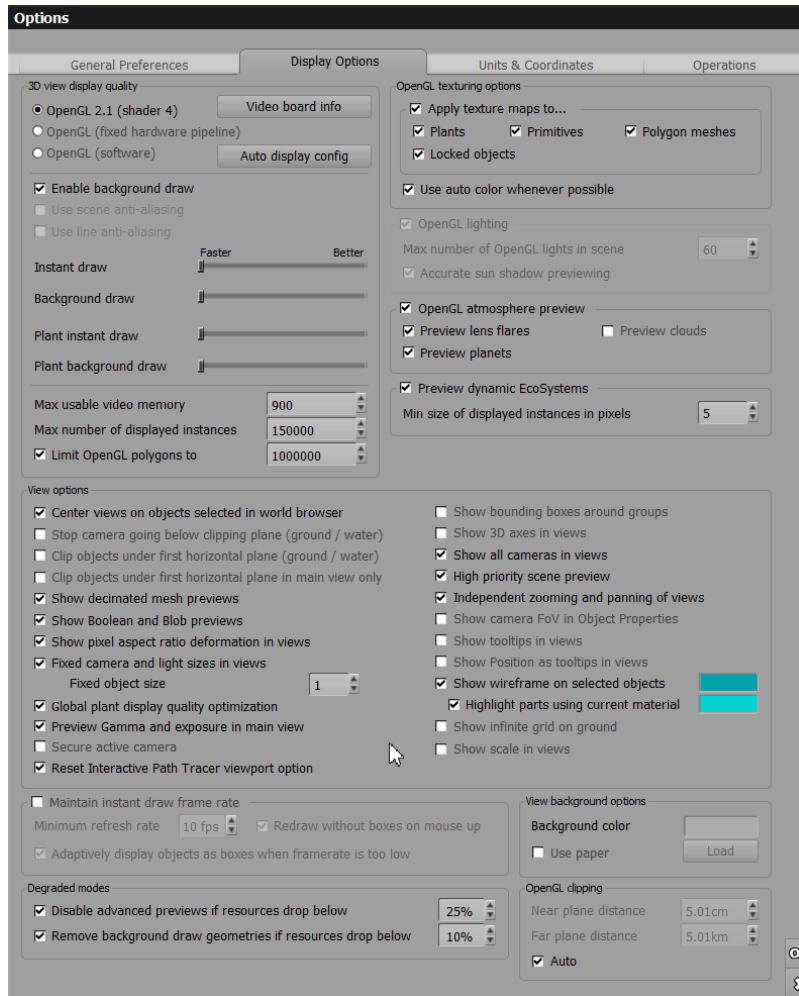
Rendering

Gamma settings can also be applied to your finished render by adjusting the settings using the **Edit** button by the slider showing the **Current Display Gamma**.

Display Options



3D View Display Quality



The Options dialog – Display tab

- **OpenGL 2.1 (shader 4):** select this option to enable OpenGL previewing of the scene using this display engine, based on the programmable shader supported by recent graphics cards. When you first start Vue, if your card supports this feature, it will be automatically selected. This setting provides per pixel lighting for diffuse and highlights. It will also enable the shaded billboard feature for EcoSystem previews.
- **OpenGL (fixed hardware pipeline):** select this option to enable OpenGL pre-



viewing of the scene using the fixed hardware pipeline in your video board. An OpenGL optimized video board is recommended (although not necessary) to get best performance from this option. Unfortunately, it has been frequently observed that some video board drivers don't correctly implement all the required OpenGL features used by Vue, which may lead to system or software instability. You should test the compatibility of your video board (see below) before using hardware acceleration.

- **OpenGL (software):** if this option is selected, Vue will not use the hardware acceleration in your video board, but will emulate the OpenGL features, using a default – fully compliant – software implementation of OpenGL instead.
- **Video board info:** this button displays your video card information and the display quality you are currently using.
- **Auto display config:** this button resets your display mode to the previous setting.
- **Enable background draw thread:** in order to provide speedy feedback while still offering detailed previews, Vue uses two different levels of 3D view quality. The first is the instant feedback, the second is a multi-threaded background update (background draw thread) using a higher level of detail. If you don't want to use the background draw thread, uncheck this option (e.g. to avoid the views switching quality all the time). If you experience frequent random crashes on your system, this is probably due to your video board driver not supporting the multi-threaded drawing. If you don't want to disable OpenGL altogether, turning this option off may help improve stability.
- **Use line anti-aliasing** and **Use scene anti-aliasing:** select these options to enable OpenGL anti-aliasing of lines and anti-aliasing of the entire scene. Depending on your video board, anti-aliasing can affect performance significantly.

The **Instant draw** and **Background draw** sliders let you customize the level of detail used for both types of previews. We generally don't recommend increasing the quality of Instant draw, since this has to be quick to remain responsive. You can however turn it down, if you feel that program response is on the slow side.

Since **Background draw** takes place in the background, the associated 3D view quality can be very high. Obviously, background draw quality should be higher than Instant draw.

Because plant preview is particularly demanding in terms of video processing, there is a separate set of sliders to adjust the quality of the preview of plants in both previewing modes. **Plant instant draw** and **Plant background draw** sliders let you customize the preview quality of plants independently from the preview quality of other objects.

- **Max usable video memory:** This setting only comes into play when using full quality near the camera. To a lesser degree, it affects the changing to another geometry representation, such as box, wire box, flat shaded or smooth shaded.



- **Limit OpenGL polygons to:** Because some video boards may have issues displaying the millions of polygons of Vue scenes, this option will automatically limit the number of polygons that the video board has to draw in order to avoid OpenGL driver issues. When the polygon limit is reached, additional objects will be displayed as boxes instead of their full polygon representation. You can try increasing this value if you feel that it is not appropriate for your work, however, if the application starts crashing because of video board driver errors, you should revert to the factory setting.

OpenGL Texturing Options

This group of controls let you customize the way object texturing is previewed in the 3D Views. Normal and UV mapping is now displayed correctly in OpenGL previews.

- **Apply texture maps to...:** if this option is selected, the objects will be previewed with their texture maps. This is particularly nice for plants, but also applies to all objects that are mapped using images. You can enable or disable texture maps independently on **Plants**, **Primitives**, **Polygon meshes** and **Locked objects**. When this option is selected for locked objects, these will be shown with textures applied. Otherwise, they will be displayed gray.
- **Use auto-color whenever possible:** Vue will always attempt to determine a color that is representative of each material in the scene. If this option is selected, this color will always be used to preview the objects, except when specifically instructed to use another color.

OpenGL Lighting

The unique control in this group lets you adjust the number of light sources used for previewing. OpenGL supports a maximum of 8 light sources. Each time you create a light source in Vue, an equivalent light will be created in OpenGL, until the maximum number of lights that can be used is reached. You can accelerate the previewing of the scene by reducing the number of lights used by OpenGL.

If you are using the OpenGL 2.1 (shader 4) display option this options will be unavailable and greyed out. If you wish to change these options and use them, you need to use the OpenGL Hardware display option. If you change these settings and switch back to OpenGL 2.1 (shader4), these settings will revert back to the default.

Accurate sun shadow previewing: Use this setting to enable real shadows projected from the sun. The sun becomes the first directional light in the object list. This option disables **Clip objects under the first horizontal plane in main view only**.



OpenGL Atmosphere Preview

The **OpenGL atmosphere preview** option is only available when OpenGL is enabled. When you enable previewing of the atmosphere, Vue colors the sky according to the atmospheric settings in your scene. Clouds are not taken into account to generate this preview.

When the Atmosphere preview option is selected, you can also indicate whether you want to preview the lens flares and planets in the scene.

If you enable **Preview lens flares**, an OpenGL version of the lens flare will be created that looks like the actual lens flare that will be rendered.

Enable **Preview planets** to represent a texture mapped preview of the planets in the 3D Views.

Enable **Preview clouds** to view cloud layers in the OpenGL preview.

These options can be turned on/off easily on a per scene basis from the **Display | Atmosphere Preview** item on the Vue menu.

Preview Dynamic EcoSystems

- **Preview Dynamic EcoSystems:** Check this option to view the dynamically populated instances that are closest to the camera.
- **Min size of displayed instances in pixels:** Enter a value to limit the display of EcoSystem instances in preview mode.
- **Max number of displayed instances:** Enter a value to define the upper limit of EcoSystem instances you want displayed in preview mode. This is valid for both dynamic and static EcoSystems.


View Options

- **Center views on objects selected...:** turning off this feature will stop the views from moving each time you select an object using the *World Browser*.
- **Stop camera going below clipping plane...:** this is useful to avoid getting black renders when the camera goes inside the ground. It can be turned off for special cases when the camera should go below the clipping plane (e.g. underwater scenery).
- **Clip objects under first horizontal plane:** checking this will make all objects



invisible under the clipping plane, in all views (not only in *Main view*).

- **Clip objects... in main view only:** checking this will make all objects invisible under the clipping plane in the main view. This is useful for the comprehension of the *3D View*.



Preview of meshes (from front to rear): original frog ~20000 polygons, slight decimation ~10000 polygons, standard ~5000 polygons, strong decimation ~2000 polygons

- **Show decimated mesh previews:** in order to provide instant feedback, polygon meshes that comprise many polygons are only partially drawn. This means that only certain polygons of the object are drawn, resulting in a ghost-like preview with lots of holes. When this option is enabled, Vue will attempt to display a simplified version of the object (a decimated version) that retains the overall outline of the initial object. Although the resulting geometry may be seriously distorted, it usually retains enough to be identifiable. The amount of decimation depends on the preview quality settings. This decimated version of the object is created in a background task (to avoid slowing you down in your creative process). It may take a couple of seconds to compute and display the decimated version. Of course, the full object geometry will be used for rendering.
- **Show Boolean and Blob previews:** if this option is selected, Vue automatically builds a polygon approximation of the result of any Boolean or Metablobs/Hyperblobs operations that you create in your scene. This polygon approximation is created as a background task (to avoid slowing you down in your creative process). It may take a couple of seconds to compute and display the polygon approximation. Of course, the exact object geometry will be used for rendering. This option is not available if the **Background draw thread** is disabled.
- **Show pixel aspect ratio deformation in views:** when this option is selected, the effect of non-square pixel aspect ratios (see [here](#)) appears in the OpenGL previews.
- **Fixed camera and light sizes in views:** if this option is selected, the camera and light icons in the *3D views* will remain the same size when zooming.
- **Global plant display quality optimization:** Checking this option will boost display quality of plants depending on the overall polygon count in the scene. This



gives the highest display quality whenever possible.

- **Preview Gamma and exposure in main view:** Check this option to have your gamma and exposure settings reflected in the OpenGL preview of the *Main camera view*.
- **Show bounding boxes around groups:** If this option is selected, a dotted bounding box will be drawn around groups of objects.
- **Show 3D axes in views:** when this option is enabled, a small graph displays the axes of the views in the lower-left corner of the *3D Views*.
- **Secure active camera:** enable this option to change the main view in perspective view whenever you modify the camera settings with the mouse. This will permit you to play with the perspective camera, and once you are satisfied with the settings you can then store them into the main camera.
- **Show all cameras in views:** when this option is selected, all the cameras in the scene will be displayed in the views. The active camera is the only one displayed with a viewing frustum. When all cameras are shown in the views, you can activate another camera by double-clicking on it.
- **High priority scene preview:** if this option is selected, the Thumbnail Scene Preview in the [Camera Control Center](#) will be drawn first (before all other background tasks). This option is equivalent to selecting the **High Priority** option in the scene preview. This option is not set by default.
- **Independent zooming and panning of views:** by default, all orthogonal 3D Views are zoomed and panned simultaneously. If you would rather have a different zoom and pan setting for each view, select this option.
- **Show camera FoV in Object Properties:** when this option is selected, the size of the camera lens will be displayed as the horizontal Field of View rather than focal length.
- **Show tooltips in views:** If this option is selected the tooltips indicating “Ground”, “Sea” and so on will display in the viewports.
- **Show wireframe on selected objects:** when this option is selected, a wireframe will appear on top of the selected objects, letting you locate the selected objects more easily. The default color for this wireframe is red.

The way this wireframe is displayed is also affected by the menu option **Display | Show Selection Wireframe On Top**: if the menu option is selected, the wireframe will always appear on screen, regardless of whether the object is hidden by other objects or not. This option can be toggled on or off, because there are cases when it can get in the way (e.g. when placing an object on top of a terrain, the wireframe can get in the way of seeing how the object is positioned relative to the surface).

- **Highlight parts using current material:** this option works in conjunction with the above. When it is selected, the objects (or parts of objects) that use the current



material (as displayed in the *Object Properties* panel) will appear with a brighter wireframe. This way, you can easily see what parts of your objects use a specific material.

- **Show infinite grid on ground:** when checked, a grid displays on the ground plane in the *Main camera view* and in the *Top view*. It shows at any level of zoom or camera altitude. There are three scales in the grid, multiples of 1, 5 and 10 internal units. 1's will show lightly, whereas 10's will show brightly. During transitions between order of magnitude, the grid lines representing those multiples will smoothly fade into their new scale multiple.
- **Show scale in view:** when checked, a small scale bar displays in the lower left corner of all viewports. It indicates the current magnification level. It will match the grid cells, and allows you to know what distance represents one grid cell at all times.
- **Secure active camera:** enable this option to change the Main View in Perspective View whenever you modify the camera settings with the mouse. This will permit you to play with the perspective camera, and once you are satisfied with the settings you can then store them into the main camera.
- **Reset Interactive Path Tracer viewport option:** The **Interactive Path Tracer** is activated by default in the main view. Furthermore, the activation of the **Path Tracer** is remembered directly in the scene file, so that loading a scene that was used with the **Path Tracer** automatically activates it. If you don't want this behavior, you can prevent it by setting this option to on.

Maintain Instant Draw Frame Rate

If you select this option, Vue will attempt to maintain a minimum frame rate when refreshing the OpenGL views (when dragging objects or moving inside the views). Whenever the scene becomes too complex to display at the indicated frame rate, Vue will display objects as boxes. The number of objects displayed as boxes is adapted dynamically in order to achieve the desired frame rate. Objects far from the camera are displayed as boxes in priority.

Note:

This setting only affects the instant draw. The background draw is always displayed according to the quality setting indicated.

- **Minimum refresh rate:** use this setting to indicate the minimum frame rate that you want the OpenGL views to be displayed at. Warning: do not set this parameter too high, or your views will always be displayed as boxes. The default 5fps (frames per second) is usually a good compromise.



- **Redraw without boxes on mouse up:** when this option is selected, the OpenGL views are redrawn completely as soon as you release the mouse button after dragging the views/objects. This avoids having to wait for the background draw thread to complete drawing the views to get an idea of object placement when some objects have been drawn as boxes.
- **Adaptively display objects as boxes when framerate is too low:** when this option is selected, Vue will display distant objects as boxes when resources become depleted. This enables the application to discard display data required for the detailed OpenGL views, thus freeing up some memory for use in other, more vital tasks. The number of objects displayed as boxes is automatically adapted according to system resources.

Degraded Modes

When your system resources become dramatically depleted, Vue may no longer have sufficient resources to complete its tasks. When this situation occurs, and in order to maintain core functionality for as long as possible, degraded modes will kick into action.

Degraded modes are a special mode of operation whereby the application gives up some of its “peripheral” processing in order to focus on “vital” tasks. This includes suspending some background tasks and simplifying the *3D Views*.

The following options let you fine tune the actual level of system resources that trigger the various degraded modes:

- **Disable advanced previews if resources drop below:** select this option to automatically suspend the advanced preview threads (mesh decimation, Boolean operations and Metablob previews) when the system resources drop below the indicated threshold. These threads require a comfortable amount of resources to perform their task. By suspending them, some resources are freed for use by more vital tasks. They will be automatically restarted as soon as resources become plentiful again. Set the resources threshold for this event.
- **Remove background draw thread geometries if resources drop below:** when this option is selected, Vue will display distant objects as boxes when resources become depleted. This enables the application to discard display data required for the detailed OpenGL views, thus freeing up some memory for use in other, more vital tasks. The number of objects displayed as boxes is automatically adapted according to system resources. Set the resources threshold for this event.
- **Disable background draw thread...:** select this option to automatically switch off the background draw thread when the system resources drop below the indicated threshold. By suspending this thread, more resources are freed for use by vital tasks. This thread is automatically restarted when resources become plentiful again.



View Background Options

- **Background color:** this control lets you select the color that will be used for the background of the views.
- **Use paper:** select this option to add grain/color to the background of the views (as if they were drawn on paper). The first time you select this option, a Standard File Browser will appear, prompting you to select the paper to use in the views. Vue ships with a set of predefined papers placed in the *Papers* sub- folder of the *Environment* folder. Press the **Load** button to change the background paper used in the views. You can create your own custom papers; a paper is a seamless black and white picture with a size of 64×64 pixels.

OpenGL Clipping

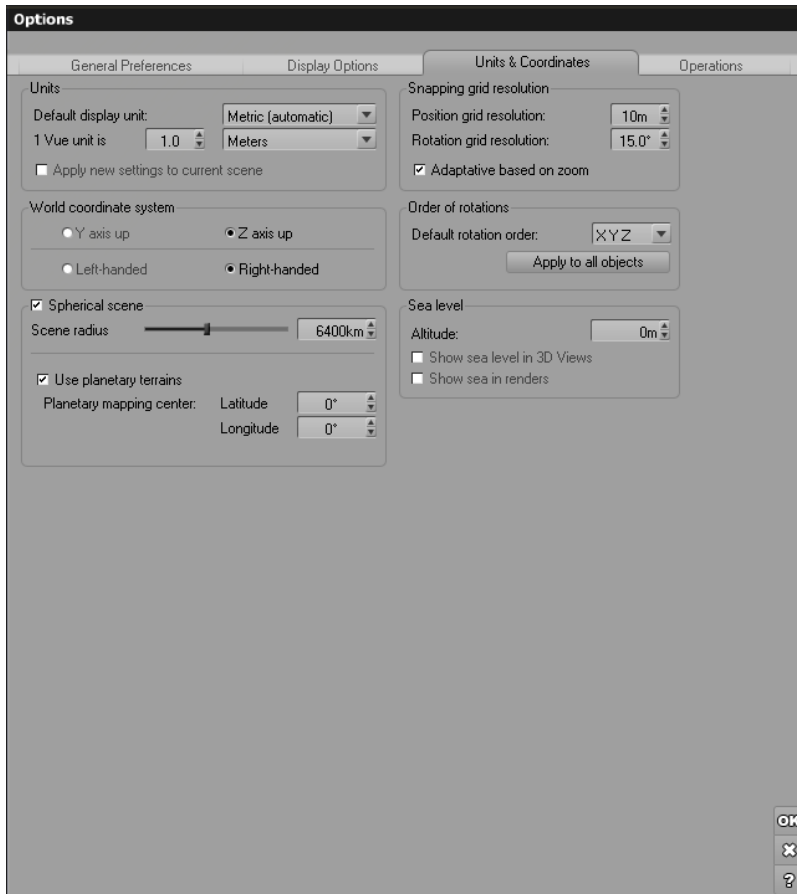
This option lets you define the OpenGL view clipping planes:

- **Near plane distance:** this defines the distance to the clipping plane near the camera. Any object that is closer to the camera than this distance will not appear in the OpenGL views.
- **Far plane distance:** this defines the distance to the clipping plane far from the camera. Any object that is further from the camera than this distance will not appear in the OpenGL views.
- **Auto:** select this option to let Vue automatically take care of these settings.

Units & Coordinates



Length Units



The Options dialog – Units and Coordinates tab

- **Default display unit:** This drop-down list box lets you define which measurement unit will be used to display lengths in Vue:
 - **Metric (automatic):** all measures will appear in metric units, automatically selecting the unit that is most appropriate for each particular measure (i.e. millimeters for very small objects, and kilometers for very large objects).
 - **Imperial (automatic):** all measures will appear in imperial units, automatically selecting the unit that is most appropriate for each particular measure (i.e. inches for very small objects, and miles for very large objects).
 - **Meters:** all measures will appear expressed in meters.



- **Centimeters:** all measures will appear expressed in centimeters.
- **Millimeters:** all measures will appear expressed in millimeters.
- **Kilometers:** all measures will appear expressed in kilometers.
- **Inches:** all measures will appear expressed in inches.
- **Feet:** all measures will appear expressed in feet.
- **Yards:** all measures will appear expressed in yards.
- **Miles:** all measures will appear expressed in miles.
- **Vue Unit:** all measures will be displayed in Vue units.
- **1 Vue unit is:** use the Vue unit setting to precisely specify how long a Vue unit will be. When changing this value, Vue automatically adjusts the atmosphere's aerial perspective scaling factor so that the atmosphere renders identically independent from the current Vue unit value.
- **Apply new settings to current scene:** check to apply these settings to current scene only.

Snapping Grid Resolution

The two settings of this group let you control the resolution of the grid when moving and rotating objects.

Snapping to the grid is enabled when you press Shift at the same time as you move or rotate an object. When snapping is enabled, the object will jump between positions or angles on the grid.

Use **Position grid resolution** to set the resolution of the grid when moving objects.

Use **Rotation grid resolution** to set the resolution of the grid when rotating objects.

Adaptive based on zoom: If this option is checked, this grid will show at any level of zoom or camera altitude, and at any panning position. There are three scales in the grid, multiples of 1, 5 and 10 internal units. 1's will show lightly, whereas 10's will show brightly. During transitions between order of magnitude, the grid lines representing those multiples will smoothly fade into their new scale multiple. This happens each time a viewport's zoom approaches a round decimal power value (1, 10, 100...).



Order of Rotations

This drop-down list box lets you define the order in which the rotations will be applied for all the new objects you create.

If you want to change the order of rotations for all existing objects, click the **Apply to all objects** button.

World Coordinate System

This group lets you configure the axes of the world coordinate system. By default, Vue uses the Z axis as the vertical axis, but if you are more familiar with the Y axis being vertical, this is the place to change: click **Y axis up** to have the Y axis vertical, or click **Z axis up** to use the default conformation.

- **Left-handed** and **Right-handed**: this lets you define whether your coordinate system is left handed or right handed (the default).

Spherical scene

This section is for setting up scenes that make use of spherical terrains, either complete planets or partial curved terrains. These properties should probably not be checked as a scene default.

- **Spherical scene**: Check this option to enable spherical scene properties in the current scene.
- **Scene radius**: this sets the size of the planet you are creating.
- **Use planetary terrains**: this will reform all of the infinite planes currently in your scene (and any you might add) into a spherical shape.
- **Planetary mapping center: Latitude/Longitude**: This mapping allows you to control the latitude/longitude settings when using a large scale map of cloud layers. It is also used to control the latitude/longitude parameter of the new planetary image mapping node which is used to map a world map for a planetary terrain.



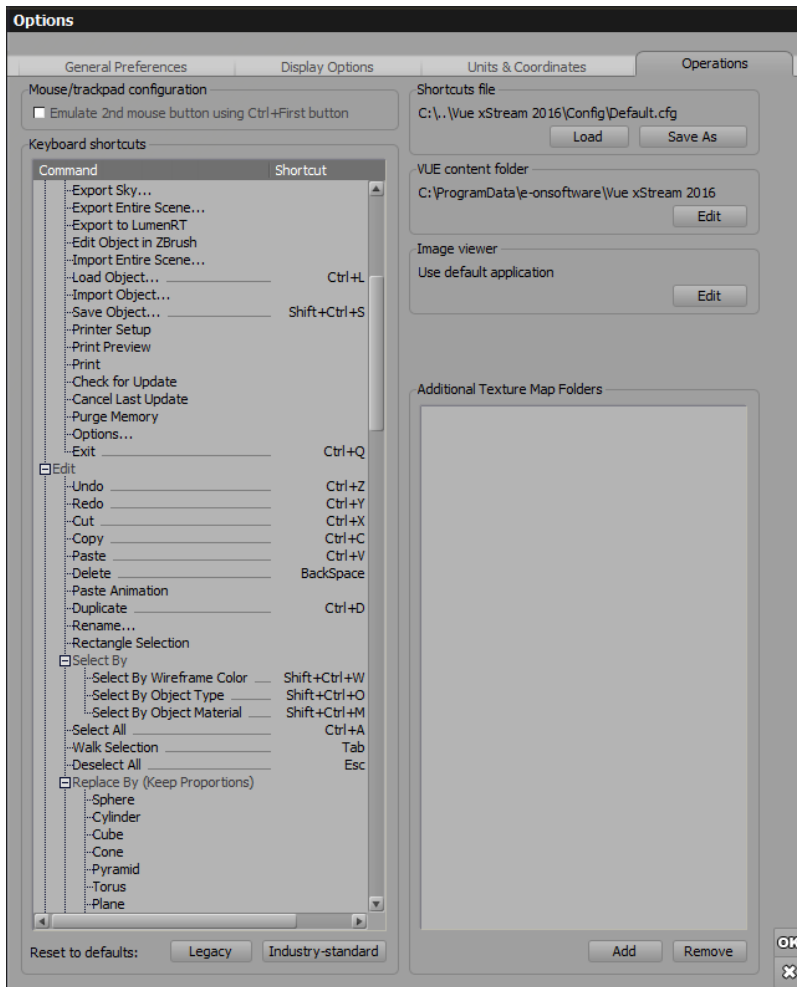
Sea Level

- **Altitude:** this sets the default for the water plane in your scene. This default affects all of the scenes you create using a water plane.
- **Show sea level in 3D views:** this gives you a visible plane in your views as a reference. A sea level plane will show in the World Browser, but be invisible in renders.
- **Show sea in renders:** this will give you a visible water plane and it will show in the World Browser as **Sea**.

If you don't check either option, sea level is still present and its value is define by default as $z=0$, or whatever value you give in on this screen.

Operations





The Options dialog – Operations tab

This tab lets you redefine keyboard shortcuts for all keyboard operations, store these settings in user configuration files, and define additional search paths for texture maps.

If you wish to emulate the Right Mouse Button (RMB) using *Ctrl+Left/Default button* for mouse and trackpad, check the box at the top of the dialog. This is available for both Windows and Mac computers. On the Mac, this should be very useful when using a trackpad.

On all platforms, when two different physical keys exist for a single character or action (eg. the two *Shift* keys on either side of the keyboard), both keys are treated as the same



key.

For Windows machines, the *Windows* key can now be used and assigned a shortcut. A keystroke assigned to this key would be listed as *Win+T* for example. The Windows key must have another key assigned to it as a combination. Otherwise, its function reverts back to displaying the Windows start-up menu.

For Macs, the *Command* and *Control* keys can be mapped independently to shortcuts.

Customizing Keyboard Shortcuts

This section lets you create new or alternate shortcuts for various commands. That way you can reassign the shortcuts you are most used to in your other applications.

In the **Keyboard shortcuts** list appear all the menu commands with their existing shortcuts alongside them.

To create a new or alternate keyboard shortcut, simply click on the line and type your new shortcut. If the shortcut is already assigned to another command, a prompt will appear, asking what you would like to do.

Click outside the list of commands to close the *Type new shortcut* invitation.

If you want to remove an existing shortcut, right click on the shortcut to be removed.

Note that when the combination you are trying to use is already assigned to another keyboard shortcut or mouse trigger, a message will tell you so and you will be able to cancel the new assignment or remove the conflicting shortcut from its original assignment to make it available.

By pressing one of the buttons at the right of the **Reset to default** label below the list of shortcuts, you can reset all the shortcuts to the factory settings:

- The *Legacy* defaults refer to the navigation defaults in VUE up to and including version 2016 Release 4.
- The *Industry standard* defaults refer to a navigation preset new to VUE version 2016 Release 5, which conforms to more widely accepted industry standards for 3D navigation.

You can also use the **Load interface presets** command in the **General Preferences** tab to load typical shortcuts from other popular 3D applications.

Customized mouse and keyboard shortcut presets can be saved in individual files (see below).



Triggers and Trigger Modifiers

Command	Shortcut
Trigger Modifiers	
-Cancel Current Operation	Esc
-Snapping	Shift
-Constrain To Axis	X
-Slow Camera Controls	W
-Duplicate Object	Alt
-Ignore Preemptive Gizmos	Space
Trigger Modifiers (Picking)	
-Toggle Selection State	Shift
-Select All Under Mouse	A
-Pick Next Object Behind Selection	B
-Select Group Components	Ctrl
Trigger Modifiers (Scaling)	
-Change Scaling Pivot	Alt
-Resize Globally	Shift
Triggers	
-Panning	Shift+Right MB+Drag
-Zoom	Ctrl+Right MB+Drag
-Default Trigger	Left MB+Drag
-Viewport Contextual Menu	Right MB
-Select Object Under Mouse	Left MB
-Draw Selection Frame	Left MB+Drag
Triggers in 3D Views	
-Orbit Around Selection	Middle MB+Drag
-Rotate Viewpoint/Camera	Right MB+Drag
-Move Camera Back/Forth	MWheel
-Frame Selected Area	Shift+Ctrl+Right MB+Drag
Triggers in Editors & Previews	
-Panning	
-Zoom	
-Orbit Around Selection	Right MB+Drag
-Rotate Viewpoint/Camera	
-Move Camera Back/Forth	
Triggers in Camera Views	
-Roll Camera	Shift+Right MB+Drag
-Translate Camera constrained to axis	Alt+Ctrl+Right MB+Drag
-Panning	
-Zoom	
-Orbit Around Selection	
-Rotate Viewpoint/Camera	
-Move Camera Back/Forth	
Triggers in 1D/2D Views	
-Panning	Right MB+Drag
-Zoom	
Triggers in Terrain Editor	
-Invert brush switch	V
-Change brush size	Shift+Left MB+Drag
-Change brush flow	Ctrl+Left MB+Drag
-Change brush falloff	Alt+Left MB+Drag
Triggers in Spline Editor	
-Toggle tangent angle constraint	A
-Toggle tangent length constraint	L
-Align tangent	Shift+Ctrl+A
-Modify Control Points	Ctrl

Triggers and Trigger Modifiers in the Shortcuts list

Triggers are shortcuts involving mouse buttons, typically used in the 3D viewports to manipulate the viewpoint, draw a selection rectangle or a render area, *etc.*

Trigger modifiers are additional keys that can be pressed during a triggered operation to modify its effect. For example, pressing *W* during camera rotation or panning will slow down the movement, allowing more precise operations.

Please refer to the detailed [Triggers and Modifiers](#) article for more information on categories and specific triggers.



Shortcuts Files

You can save and restore all your mouse and keyboard settings in a configuration file using the extension *.cfg* . That way, several people working on the same computer can have different shortcut mappings. You can also take the configuration file with you to another computer and restore it there to keep your preferred mapping.

Click the **Load** button to display a *Standard File Browser* and load the desired configuration file (the current settings are saved to the current file location before loading the new one). Likewise, the **Save As** button will let you select a new file to store the current user configuration.

Additional Texture Map Folders

You can define additional search folders for your texture maps (pictures) using the controls in this group. When VUE cannot find a texture map at the expected location, it will automatically check the *Bitmaps* folder to see if the texture map can be found there. If the texture map is not in the *Bitmaps* folder, VUE checks the folder where the file that is being loaded is located (this search is not recursive).

If you have defined additional texture map folders, VUE will also search in these folders to see if the requested texture map can be located there. That way, if you have gathered all your texture maps in folders, you can define them as additional folders, and VUE will automatically check there if it cannot find a given texture map elsewhere. This search can be made recursive.

To add a new folder, simply click the **Add** button. A *Standard Folder Browser* will appear, letting you select the new folder to be added. Click the **Recursive** checkbox to let VUE search in the sub-folders of this folder.

To remove a folder from the search path, highlight a folder and click **Remove**.

VUE Content Folder

You can easily change the location of your VUE content folder. To change to a new folder, simply click the **Edit** button. A *Standard Folder Browser* will appear, letting you select the content folder.



Image Viewer

Use this **Edit** button to browse to the program you wish to use to view your bitmaps that you are using in the *Material Editor* or *Plant Editor*.

Preferred Documentation Media

Check to indicate whether you want to view the documentation .pdf file or the online wiki when you select **Overview** on the *Help* menu.



Triggers and Modifiers

Command	Shortcut
Trigger Modifiers	
-Cancel Current Operation	Esc
-Snapping	Shift
-Constrain To Axis	X
-Slow Camera Controls	W
-Duplicate Object	Alt
-Ignore Preemptive Gizmos	Space
Trigger Modifiers (Picking)	
-Toggle Selection State	Shift
-Select All Under Mouse	A
-Pick Next Object Behind Selection	B
-Select Group Components	Ctrl
Trigger Modifiers (Scaling)	
-Change Scaling Pivot	Alt
-Resize Globally	Shift
Triggers	
-Panning	Shift+Right MB+Drag
-Zoom	Ctrl+Right MB+Drag
-Default Trigger	Left MB+Drag
-Viewport Contextual Menu	Right MB
-Select Object Under Mouse	Left MB
-Draw Selection Frame	Left MB+Drag
Triggers in 3D Views	
-Orbit Around Selection	Middle MB+Drag
-Rotate Viewpoint/Camera	Right MB+Drag
-Move Camera Back/Forth	MWheel
-Frame Selected Area	Shift+Ctrl+Right MB+Drag
Triggers in Editors & Previews	
-Panning	
-Zoom	
-Orbit Around Selection	Right MB+Drag
-Rotate Viewpoint/Camera	
-Move Camera Back/Forth	
Triggers in Camera Views	
-Roll Camera	Shift+Right MB+Drag
-Translate Camera constrained to axis	Alt+Ctrl+Right MB+Drag
-Panning	
-Zoom	
-Orbit Around Selection	
-Rotate Viewpoint/Camera	
-Move Camera Back/Forth	
Triggers in 1D/2D Views	
-Panning	Right MB+Drag
-Zoom	
Triggers in Terrain Editor	
-Invert brush switch	V
-Change brush size	Shift+Left MB+Drag
-Change brush flow	Ctrl+Left MB+Drag
-Change brush falloff	Alt+Left MB+Drag
Triggers in Spline Editor	
-Toggle tangent angle constraint	A
-Toggle tangent length constraint	L
-Align tangent	Shift+Ctrl+A
-Modify Control Points	Ctrl



Triggers and Trigger Modifiers in the Shortcuts list

Triggers and their *Modifiers* are sorted into several categories. To allow more freedom of customization, some *Triggers* appear in several categories. For example, the *Orbiting* trigger can be set to *Alt+Left MB+Drag* in the scene viewports but overridden by the simpler trigger *Left MB+Drag* in the object editors' 3D previews.

Trigger Modifiers

Under the *Trigger Modifiers* category, you will find triggers which apply to camera and object interactive transformations, as well as two more generic triggers:

- The *Cancel Current Operation* shortcut (usually *Escape*) can apply to a wide range of operations,
- The *Ignore Preemptive Gizmos* modifier can be used to prevent inlaid or overlaid interface elements like 3D gizmos in the scene viewports from activating when hovered.
 - This is because those interface elements have priority over other operations, which in some circumstances make it difficult to effect other operations like picking an object when looking at a dense area of your scene where many objects overlap.
- Modifiers applying to transformations are pretty standard:
 - *Snapping* snaps movement to the intersection of an implicit grid, or rounds up the angle of a rotation (it has no effect on *Scaling* operations),
 - *Constrain To Axis* limits the degrees of freedom of an action,
 - *Slow Camera Controls* will slow down the effect of camera, viewpoint or object transformations,
 - *Duplicate Object* allows to duplicate an object at the beginning of a transformation in order to work on a copy rather than on the originally selected object.

Under the *Trigger Modifiers (Picking)* category, you will find triggers modifiers applying to the corresponding *Select Object Under Mouse* operation.

Under the *Trigger Modifiers (Scaling)* category, you will find triggers modifiers applying to the scaling operations available using the *gizmo* handles in the 3D views.

Triggers

Under the top-most *Triggers* category, you will find triggers which operation applies to many situations. Typically, *Panning* and *Zooming* can refer to translating the camera as well as dragging a node in the *Function Graph* or moving the *Time Line* left or right when working on an animation.



The *Default Trigger* deserves a special note: it simply refers to the default mouse button used to start some operations that can be activated from an interface button or with a keyboard shortcut. For example, if you activate the *Select Render Area* feature using the button in the *Top Toolbar*, or using the *Ctrl+Alt+S* keyboard shortcut, you will have to use the *Default Trigger* (usually *Left MB+Drag*) to actually select the area's rectangle in a 3D view.

Triggers in 3D views

The nested *Triggers in 3D views* category adds triggers meaningful only in 3D context (some of the triggers apply in orthographic viewpoints as well).

- *Orbit Around Selection* rotates the viewpoint around the selected objects, or around the scene's center point. When no object is selected, be careful when orbiting as the scene center can be far off the current frustum of your viewpoint.
- *Rotate Camera/Viewpoint* rotates the camera or the viewpoint's origin around itself.
- *Move Camera Back/Forth* (a.k.a. *Dolly*ing) moves the camera "inside" the scene or "away" from it along a linear path. Combined with the camera's *Altitude lock* feature, this operation can be made similar to a "walking" mode, typically when operated over a terrain. As a convenience, this trigger will also zoom in or out of an orthographic viewpoint, even if in that case it only comes down to shrinking or enlarging the view extent.
- *Frame Selected Area* will let you draw a rectangle on a view and try to adjust the viewpoint settings to focus on this area.

Editor & Previews vs. Camera Views

The nested *Triggers in Editor & Previews* category repeats triggers already mentioned to allow to override them in the 3D preview of an object editor like the *Terrain Editor*, *Plant Editor*, *Text Editor*, etc.

The nested *Triggers in Camera Views* category repeats triggers already mentioned to allow to override them in a perspective view of the scene, and adds two triggers specific to this context:

- *Roll Camera* rotates the camera around its forward (*look at*) axis, changing the angle also known as the camera *attitude*.
- *Translate Camera constrained to axis* (kept for compatibility with previous versions – see also modifier mentioned earlier in the *Trigger Modifiers* category) allows to translate the camera along the X or Y axis as described in [this article](#).



Triggers in 1D/2D Views

The nested *Triggers in 1D/2D Views* category repeats the generic *Panning* and *Zooming* triggers to allow to override them in the specific context of non-3D views like the *Function Graph*, the *Time Line*, the *Function Node Preview*, etc.

Triggers in Editors

The *Triggers in the Terrain Editor* and *Triggers in the Spline Editor* categories hold mouse and keyboard shortcuts meaningful in the context of these editors only.

This allows to define shortcuts for actions like raising or shrinking the *Terrain Editor Brush* radius, flow or fall-off distance between two brushing actions, without having to move the mouse to the *Paint* panel to operate the interface. See the section on *Sculpting Terrains* in the *Terrain Editor* documentation.



Cornucopia3D™

Cornucopia3D™ is an online community of Vue users. On the *Cornucopia3D.com* website, you will find:


- User forums where Vue users share their experience and knowledge of Vue,
- Galleries where you can display your work,
- Classrooms and tutorials to help you rapidly master the product,
- Regular contests where you can pit your talent against other members, and
- An online store where you can purchase additional content for use inside your Vue scenes.

So don't hesitate to visit the friendly crowd at www.Cornucopia3D.com!


Integrated Cornucopia3D Store

Cornucopia3D also hosts a service designed to help you to easily and rapidly find the quality content you need to complete your projects. To this effect, the *Cornucopia3D* technology places “ghost” items into your collections. These items don't “physically” exist on your hard disk. This grants you the ability to instantly browse through a massive amount of high quality content from within Vue without having to exit your Vue session – content among which you will probably find what you need for your projects.

New *Cornucopia3D* items are regularly downloaded from the [Cornucopia3D](http://www.Cornucopia3D.com) website and shown in your collections, so that you always have access to the most up-to-date content.

You can customize how *Cornucopia3D* items appear in your collection anytime by using the Cornucopia3D item display options () in the Browser Option in the upper right of each browser.

- **Hide all items:** select this option to remove all Cornucopia3D items from the list.
- **Show best items:** use this option to only show a short selection of the best Cornucopia3D items available. While this option does not let you see all the content available, you will still see a few items that we believe may be of the most interest to you.
- **Hide all items:** select this option to view all Cornucopia3D items of the collection.

Cornucopia3D items are identified by a little pictogram () that is visible in the bottom-right corner of the item's preview. These items need to be downloaded before you can use them. If you load a *Cornucopia3D* item, the *Cornucopia3D Item* dialog will popup to provide additional information about the item, including the file size and cost of the



item (most of the *Cornucopia3D* items are available for a small fee, but some of them are available for free!).

When the *Cornucopia3D Item* dialog appears, a connection is established with the *Cornucopia3D* website in order to determine the cost of the item as well as other information (including the amount left on your *Cornucopia3D* account; if the item is available in different formats, a dialog will also appear to present the different versions and let you decide the one you want).

To acquire the *Cornucopia3D* item, simply press the **Get item** button. You will be asked for your account id and password (this is defined at the time of registering your product) to authenticate the charge on your *Cornucopia3D* account, after which the item will be automatically downloaded and opened in Vue. The item will be automatically saved on your hard drive so that you can load it directly next time you want to use it.

If for any reason Vue cannot connect to the *Cornucopia3D* website, a message will appear inviting you to connect to an URL manually, using your standard web browser.

You may also shop for *Cornucopia3D* content anytime by going to the site using your standard web browser (www.cornucopia3d.com) and logging in with your account ID and password.

Copy-Protected Items

Some items on *Cornucopia3D.com* are sold copy-protected. These items are identified as being 'Locked to your license(s)'. Copy-protection was implemented in Vue as a response to the request of artists who broker their content on *Cornucopia3D*, so that copy-protection translates into lower prices for you, the customer.

At e-on software, we felt that copy-protection would be an acceptable limitation, provided that:

1. Copy-protection meant cheaper prices for you (in exchange for the protection of their content, we require from our brokers that they sell copy-protected content cheaper), and
2. Copy-protection is totally transparent and hassle-free.

This is why e-on software has developed a state-of-the-art copy-protection technology (e-on software holds a number of patents on that technology) that works totally “behind the scenes”:

- When you download a copy-protected item, it is automatically prepared so that it can be read on all your registered products without even requiring you to enter a key. To you, the item will behave exactly like a non-copy-protected item.
- When you later upgrade your copy of Vue, your copy-protected content will work in



the new version.

- If an item is sold as copy-protected and non-copy-protected versions, you can “upgrade” your version of the item to non-copy-protected for the difference in cost between the two versions.
- You can exchange a scene that contains copy-protected items with a friend. The copy-protected items will appear as boxes, but your friend will still be able to modify the scene. When he sends the scene back to you, you will see the complete scene with his modifications. If your friend has also purchased the copy-protected content, he will see the scene exactly like you do.

Note:

Although we like that our users have the ability to choose between copy-protection and no copy-protection, Cornucopia3D does not “impose” copy-protection. It is the brokered artist’s choice to sell content with, without or with and without copy-protection.

What's New in VUE 2016 Release 2

Additional ViewPorts

Previously, the number of OpenGL views inside VUE was limited to four. It is now possible to have up to 16 views at the same time, including multiple perspective views.

Panoramic Rendering

VUE now provides a direct way to render VR panoramas.

Multi-Materials in the Object Properties Box

For objects with multi-materials, if you right-click on the material preview, the menu entries corresponding to the different materials will be grouped in a submenu.



What's New in VUE 2016 Release 3

UX/UI improvements

- Major performance/responsiveness/stability improvements
- Higher quality display in Render Scene preview for soft shadows
- Better handling of editor window manipulations
- Support for HiDPI screens on MacOS

Optimized scene saving process

- VUE now has 3 modes to save scene : Default, Consolidated Archive, Network Rendering. The default mode now creates .fsvue file type which is smaller (up to 2 times smaller on disk) and faster (up to 4 times faster to write) compared to previous .vue files.

HDR Rendering

- Major improvements in dynamic range depth when using Spectral Atmospheres
- New optional Sun Visibility parameter for Spectral Atmospheres
- Automatic de-activation of Tone-Mapping when saving HDR pictures as .hdr
- Automatic de-activation of Tone-Mapping when saving HDR pictures as .exr

Path Tracer renderer improvements

- Adds support for Spectral Cloud layers
- Support for “Don’t cast shadows” and “Disable Indirect Lighting” material properties
- New Translucent material type
- Now supports Soft Shadows and Blurry Reflections/Refractions



Cloud Modulation

- New Cloud Opacity modulation
- Access local coordinates for color settings

xStream Integration

- Adds support for 3ds Max 2018
- Optimized native spline handling

xStream

Other

- Faster computing of Spline operations for HeightField Terrains
- Network Rendering: New command line options to set color/alpha/depth passes for both stills and animations rendering. See [the RenderNode documentation](#)
- New [Python callbacks](#) for setting and getting Color, Alpha and Depth channels saving locations

What's New in VUE 2016 Release 4

Stereoscopic render

VUE allows you to render your scenes in Stereo for Stills and Animations. Stereoscopic rendering is available for both panoramic (360° and 180°) and non-panoramic renders. Stereo renders are compatible with most of VUE's Render Options, such as multi-pass or High Dynamic Range (HDR) for instance.

Interpupillary Distance is user defined (default: 6.4 cm) as well as final image layout (top/bottom or left/right). For maximum customization, VUE offers 3 convergence modes for handling parallax: Parallel, Converged and Off-axis.

Stereo images and videos are automatically recognized by Stereo-compliant platforms (such as Youtube for instance).



180 Degree VR Panoramas

Rendering in 360° can sometimes take a lot of computing resources. VUE xStream 2016 now offers VR180 Panoramic still and animation rendering, which in essence will render a 180° vertical and horizontal FOV of your scene. VR180 will be ideal to save some rendering time, for those instances where you don't need to render what's behind the camera!

Multi-Pass Rendering Improvements

Ideal for compositors, VUE offers optimized handling of semi-transparent materials to prevent background color to bleed within object mask color data. For instance, Cloud passes or tree leaves masks won't show the sky in the background.

Additionally, you can now embed the alpha channel for each mask pass (when the output format supports it).

Path Tracer Renderer Improvements

- The VUE Path Tracer now supports normal mapping.
- You can specify the OpenCL devices that should be used for the Path Tracer: All Available Devices, GPU Devices Only, or CPU Devices Only.

Misc. UX/UI Improvements

- VUE xStream 2016 now supports CPUs with more than 64 cores on Windows.
- Improved 4K/HiDPI display compatibility.
- Adds support for Maya 2018 (MentalRay and V-Ray – incl. viewport 2.0) and Cinema 4D R19.
- Better responsiveness and faster computation of Procedural Terrain within the Terrain Editor.
- Alembic compatibility updated to version 1.7.3.

xStream



Section 2

Building Scenes






Creating Objects

All scenes are created from a set of predefined objects. These objects fall under 15 categories:


-
- | | |
|----------------------|----------------------|
| 1. Primitive Objects | 8. Rocks |
| 2. Infinite planes | 9. MetaClouds |
| 3. Alpha planes | 10. Planets |
| 4. Terrain objects | 11. Lights |
| 5. Plants | 12. Group objects |
| 6. Polygon meshes | 13. Ventilators |
| 7. 3D Text | 14. Cameras |
| | 15. Imported Objects |
-

Primitive Objects

Primitive objects are “mathematically pure” objects. They are defined by a primitive mathematical equation that the render engine solves each time it has to consider the object. This equation defines the position and shape of the object. Although this may sound pretty complex, primitive objects are among the easiest to use, and their mathematical complexity is hidden away by user-friendly tools. There are 7 primitive objects available in Vue: **Sphere**, **Cylinder**, **Cube**, **Cone**, **Pyramid**, **Torus**, **Plane** and **Alpha Plane**.

These objects are created either by clicking on the second icon from the left toolbar (if the requested primitives icon is not available, you will need to unfold it ) or by using the items from the **Object | Create** sub-menu. Keyboard shortcuts are also available for advanced users (they are indicated in the menu).

Since all of these objects can be moved, resized in any direction, rotated, and twisted, they can yield an incredible variety of shapes.

Arbitrary materials may be assigned to primitive objects using the **Change object material** item from the **Objects** menu, or by clicking the **Load material** button () on




the *Object Properties* panel.

Infinite Planes

Infinite planes are not so different from primitive objects, in the sense that they are also defined by a mathematical equation. However, unlike primitive objects, infinite planes are unbounded objects. That means they extend infinitely in every direction, separating the world in two halves. One half will be outside the object, the other inside.

There are 3 different types of infinite planes available, although they differ only by the material assigned to them, and their initial orientation. These are: **Water** (or Sea), **Ground**, and **Cloud** planes to add rain or snow.

Like primitive objects, infinite planes can be moved, rotated, twisted (except for cloud planes), and have materials assigned to them. Resizing an infinite plane will yield no result, since it is, after all, infinite.

These objects are found on the left toolbar, topmost icon (if the requested infinite plane icon is not available, you will need to unfold it ). They can also be created using the menu items of the **Object | Create** sub-menu.

All infinite planes are created horizontally. Water (or Sea) and Ground planes are created with their “inside” underneath them.

Default landscapes are created with a single Ground plane, positioned at altitude 0. While there can be multiple ground and cloud planes, there can only be one water plane, which represents sea level.


Infinite planes are displayed by a surface and a normal vector. In order to find out which side of the plane is “outside”, you need to look at this vector. It is situated on the outer side of the infinite plane.




Alpha Planes

The Alpha plane is a variation of the Plane primitive that is designed to facilitate the setup of transparency mapped images (images with alpha information).

Like other objects, they may be moved, resized, rotated, twisted. Although it is perfectly possible to assign a material to an Alpha plane, Alpha planes are specifically designed to avoid having to do so. It would seem more logical to create a simple Plane primitive instead.

To create an Alpha plane, either unfold the second icon from the left toolbar and select the corresponding icon () , or select the menu command **Object | Create | Alpha Plane**.

When you create a new Alpha plane, the *Alpha Plane Options* dialog appears. This dialog lets you define the aspect of the Alpha plane. Please click [here](#) for details on editing Alpha planes.

The *Alpha Plane Options* dialog is simply a shortcut designed to help you create the correct material for the plane. This material can be modified using the standard [Material Editor](#), but you can also re-open the *Alpha Plane Options* dialog anytime by double-clicking on the Alpha plane object in the *3D views*, or by clicking the **Edit** button () on the top toolbar when the Alpha plane is selected, or by selecting **Edit Object** from the **Objects** menu.


Terrains


Terrains are constructed using complex fractal algorithms to recreate mountainous structures. They are a special type of Polygon mesh designed to efficiently handle massive amounts of polygons. Terrains are the basic construction block used for building landscapes. Like other objects, they may be moved, resized, rotated, twisted, and have materials assigned to them.

Terrains come in two different types: standard and procedural. Standard terrains use a fixed size grid to represent the terrain altitudes. They are also known as “heightfields”, and are the most straightforward type of terrain.



Procedural terrains use a significantly more complex technology to build and refine the terrain according to the distance at which you are observing it. This technology is able to dynamically adjust the level of detail of the terrain, so that it appears to be infinitely detailed. The altitudes of this type of terrain are generated using a complex mathematical procedure. Luckily, VUE is able to hide away all this complexity by providing a selection of presets from which you can pick. You can always customize the look of these terrains later.

To create a terrain, right-click on the terrain icon on the left of the screen and select which type of terrain you want – standard heightfield, procedural or infinite from the expanded icon (). The shape of the terrain is created randomly and will be unique. No two terrains will ever yield the same shape.

Once a terrain is generated, it can be modified by accessing the *Terrain Editor*. This is done by double-clicking on the terrain in the *3D Views* or *World Browser*, or by clicking the **Edit** icon () on the top toolbar when the terrain is selected, or by selecting **Edit Object** from the **Objects** menu.


Inside the *Terrain Editor*, you will be offered a large variety of tools (such as erosion, manual editing, special effects...) that will let you shape the terrain as you like. There are also controls for painting materials directly on terrains. For a complete description of the *Terrain Editor*, please see [here](#).

Alternately, you can choose to create standard terrains directly inside the *Terrain Editor*, by long-clicking on the terrain icon in the left toolbar, or by selecting the menu command **Object | Create | Heightfield Terrain | Heightfield Terrain in Editor**.

There is also a spherical terrain in VUE, which basically creates a planet or a curved terrain (part of a planet). This is presented in more detail [here](#).



Plants



The Visual Plant Browser

Plants are probably the most unique feature of VUE. They are generated using *Solid-Growth™*, a set of incredibly complex algorithms designed to grow a whole variety of plant forms. Plants are very complex objects using multiple materials. They are often made of an equivalent of tens of thousands of polygons.

Like other objects, they may be moved, resized, rotated, and twisted.

They can also have [wind](#) applied to them, and they will react to any [global breeze settings](#) in the *Atmosphere Editor* section.


To create a plant, either click on the sixth icon from the top in the left toolbar (🌳), or select the item **Plant** from the **Object | Create** sub-menu. If this is the first time you create a plant in the current session, a random species is selected from the available ones. Successive plants will all be of the same species. When a plant is created, it grows randomly inside the computer, following rules characteristic of the species. Since the resulting plant was grown at random, no two plants from the same species will ever look identical. If you look for a particular shape inside a given species, try growing several plants, and retain the one you like best.

To change the active plant species, right-click (Ctrl + click on Mac) on the vegetation button, or select **Load Plant Species** from the **Object | Create** sub-menu. You will be presented with a browser of available species. When you select one of them, a plant from the corresponding species will be grown.

Once a plant is created, it can be modified by accessing the *Plant Editor*.

This is done by double-clicking on the plant in the *3D Views* or the *World Browser*, or



by clicking the **Edit** button () on the top toolbar when the plant is selected, or by selecting **Edit Object** from the **Objects** menu.

Inside the *Plant Editor*, you will be offered a large selection of tools that will let you modify the shape of the plant as well as create entirely new plant species.

Note:

More plant species are available from www.cornucopia3d.com.

Plants From PlantFactory

Plants that are created in e-on's PlantFactory can also be used in VUE. While the file extension for these plants is different, .tpf, VUE still recognizes them and makes use of them in the same way as SolidGrowth .veg files.

These plants can grow, be acted upon by wind, and be modified in the VUE Plant Editor. They replicate like VUE SolidGrowth plants with variations.

Since these plants tend to be more complex, a quality control has been added when using these plants in an EcoSystem to lighten the plant for processing times.

These plants are available for purchase from www.cornucopia3d.com or you can use the PlantFactory product to create your own plants.

Polygon Meshes


Polygon meshes (sometimes also called “3D models”) are objects designed by assembling numerous flat 3D faces, called polygons (polygons are sometimes referred to as triangles, quads or n-gons depending on the number of edges in the polygon). These objects are created in other 3D applications, and may be imported into VUE using the supported file format converters using various [supported file formats](#).


The strong point of polygon meshes is their ability to model any shape. However, when creating a “standard” shape, such as a sphere, using primitive objects is far more efficient, and yields better visual results (because you don't see any broken edges on the sides of objects, caused by the limited number of polygons in the object).

Like other objects, polygon meshes may be moved, resized, rotated, twisted, and have materials assigned to them. Some polygon meshes may actually have different materials




assigned to their different parts.

Polygon mesh objects are loaded into VUE as any other type of object: select the menu command **File | Load Object**, or click the **Load Object** icon () and select the object of your choice. Some polygon meshes are **rigged meshes**, meaning that their geometry will change over time.

To import a polygon mesh from another application, select the menu command **File | Import Object**, or click the **Load Object** icon () in the left toolbar, and then click **File** inside the browser. You will be presented with a Standard File Browser, letting you select the file to import.

Note:

Importing external objects is a long and complex process, so it may be time consuming.

Although no control over the shape of the polygon mesh is available in VUE, the aspect of the surface can be modified using the [Polygon Mesh Options](#) dialog. This is accessed by double-clicking on the polygon mesh in the *3D views*, or by clicking the **Edit** button () on the *Top toolbar* or in the *World Browser* toolbar when the polygon mesh is selected, or by selecting **Edit object** from the **Objects** menu.

Inside the *Polygon Mesh Options* dialog, you can indicate if the surface of the object should be smoothed, as well as the amount of smoothing required.

Rigged Meshes

VUE supports a special type of polygon mesh that is known as rigged meshes. These meshes are special in the sense that they are imported together with animation information. When you import a rigged mesh and you open the [Timeline](#), you will notice that these objects are listed as being animated objects. If you drag the current time slider, you will see that the geometry of the objects is modified as you change the current time. For instance, if you load a rigged model of a man with walk animation included, when you drag the current time slider, you will see the man walk.

Rigged meshes are particularly useful to instantly add life to your scenes. For instance, if you are creating an architectural rendering of a building, adding a few characters walking around the building will make the project look much more convincing and lifelike.

With rigged meshes, you can change their animation by using pre-defined motion files. Motion files contain a set of parameters that describe the way rigged meshes should be modified over time. By assigning a different motion file to a rigged mesh, you could, for instance, change the animation of a walking man into that of a running man.



The nice thing about motion files is that they apply indifferently to all sorts of rigged meshes. For instance, you could assign the aforementioned running man motion to a model of a walking woman; this would turn the model into a running woman.

Note:

However, because some motion files are designed for very different types of models, assigning new motions from one type of character to another may sometimes lead to strange results.


Not all motion files define an animation. There are some motion files that simply define a static pose (e.g. that of a sitting character). When you assign such a “pose” to a rigged character, the character does not appear in the list of animated objects in the [the Properties Timeline](#).

The *Skeleton Editor* can be used when adding or changing the motion of your rigged meshes.


3D Text

3D Text is also a special type of polygon mesh that is created by extruding and shaping 2D text into the third dimension.

Like other objects, text objects may be moved, resized, rotated, twisted, and have materials assigned to them.

To create a 3D Text object, click the **Text** icon () on the left toolbar. The *Text Editor* appears, letting you enter the characters in your text, and define their shape. When you are happy with the text, click **OK** and the text will appear in the *3D Views*.

For greater flexibility, all the letters in a 3D text are created as separate objects (polygon meshes), and they are grouped into a single 3D Text object.


You can access the *Text Editor* again after creating the text object by double-clicking on the text object in the *3D Views*, or by clicking the **Edit** button () on the *Top toolbar* or in the *World Browser* toolbar when the text object is selected, or by selecting **Edit Object** from the **Objects** menu.



Rocks

Rocks are a special type of polygon mesh created randomly using complex algorithms. In much the same way as vegetation creation algorithms, these algorithms are designed to produce rocks that are always different.

Like other objects, rocks may be moved, resized, rotated, twisted, and have materials assigned to them.

To create a rock, either click on the **Rock** icon in the left toolbar () or right-click to display the Rock Template Browser. Use the Rock Template Browser as you would with the Plant Browser: select a rock to add it to the scene. Now, if you click on the **Rock** icon again in the left toolbar, you will create a rock that is the same rock type you selected in the browser, but the rock shape will be slightly different, just like variations with plants. Creating a rock may take a few seconds as all rocks are made from various fractals and noises which allows for their difference in shapes and sizes.

If you double-click on a rock, the *Polygon Mesh Options* dialog displays. You can achieve some effects by modifying the smoothness of the mesh – but the default rock settings usually work best.

A Rock as an EcoSystem Specimen

The Rock Template Browser is also enabled in the *EcoSystem Population* dialog where you can use it to drag/drop directly from the Rocks Browser to the population list. When you select to add a rock, the Rock Template Browser displays for you to make your choice. When you add a rock, VUE generates about 20 different variations (instances) from the template chosen, and uses them randomly when populating.


Rock generation in EcoSystems may take some time. Once this is done, however, these same variations will be available for future EcoSystems.

Rocks can be exported.




Planets

Planets are the only objects in VUE to be placed beyond the atmosphere's cloud layers. Like other objects, they can be moved, resized and rotated. They cannot have materials assigned to them. By default, planets have the **Main View Only** viewing option.

To create a Planet, either click on the **Planet** icon in the left toolbar () , or select the menu command **Object | Create | Planet**. Use the *Object Properties* panel for [planet editing](#).

MetaClouds


MetaClouds are stand-alone clouds that can be moved around, rotated and resized just like any other object. They are a good complement to cloud layers and are particularly useful when you need large, prominent clouds such as cumulonimbuses, or if you want to place a cloud at a precise location without having to fiddle with the intricacies of procedural cloud layers.

To create a MetaCloud, either click on the **MetaCloud** icon in the left toolbar () , or select the menu command **Object | Create | MetaCloud**.

Note:

MetaClouds can only be created when using the [Spectral Atmosphere model](#).

Like plants, MetaClouds are constructed randomly from a set of rules that define their overall aspect. So each time you create a new MetaCloud, it will be different from the previous MetaCloud.


The rules that define the overall aspect of MetaClouds are gathered in MetaCloud model files. You can change the MetaCloud model by right-clicking (Ctrl + click on Mac) on the **MetaCloud** icon () to display the MetaCloud model Visual Browser. This browser displays all the available MetaCloud models. Simply select the model of your choice and press **OK** to create a new MetaCloud of the desired model.

MetaClouds are made out of the association of many spherical cloud primitives. These cloud primitives can be modified individually: they can be moved or resized to a certain extent (only proportional resizing is possible, and they cannot be rotated), letting you



customize the shape of your clouds to fit your needs.

A special type of volumetric material is assigned to the MetaCloud in order to capture the way real clouds would interact with light and the atmosphere. The same material is assigned to all the parts of the MetaCloud.

You can add new primitives to a cloud by expanding the MetaCloud in the [list of objects](#) and selecting one of the parts of the MetaCloud. The MetaCloud icon on the left side of the user interface will then change to the **Add MetaCloud Primitive** icon (). Clicking this icon will add a new primitive to the MetaCloud, that you can place or resize as needed. You can delete MetaCloud primitives by selecting the primitives to be deleted and pressing **Delete**.

MetaClouds can be animated either globally, or by animating individual MetaCloud primitives. You can achieve incredible “cloud morphing” effects this way.

Lights

Seven types of light sources are available in VUE: five “basic” types of light (Point lights, Quadratic point lights, Spot lights, Quadratic spot lights and Directional lights) as well as 3 types of “area” lights (Light Panels, Daylight Portals and Light Emitting Objects). Each type of light source has a different way of casting light.

Simple Light Sources

Simple light sources emit light from a single “mathematical” point. This is an approximation of reality, where light is in fact emitted by the entire surface of the light source (usually a very hot surface too). Simple lights are easier to compute than the more realistic “advanced” light sources, but result in unnaturally sharp transitions between light and shadow. To circumvent this problem, you can simulate the behavior of realistic lights by assigning a fake surface to the light source, which will result in smoother transitions between light and shadow, known as “soft shadows”.

Soft shadows are turned on by setting the **Softness** of the light to a non zero value. This control is available in the *Object Properties* panel, when the light is selected. The greater the value, the larger the “surface” of the light, and the more gradual the transition from light to shadow. However, since soft lights are much more computationally demanding than standard lights, it is recommended that you use such effects judiciously.




Following is a list of the different types of simple lights and the way they cast light:

- **Point lights** and **Quadratic point lights**: emit light in all directions. Light is cast from the center, with an intensity that reduces proportionately to the distance from the center. They function like a typical light bulb. Quadratic point lights are identical to standard point lights, except that light intensity decays more rapidly.
- **Spot lights**: emit a cone of light around one direction. Light is cast from the center, with an intensity that reduces proportionately to the square of the distance from the center. Two settings let you adjust the angle of the cone (spread) and the speed at which light falls off on the side of the cone.
- **Directional lights**: emit light in one direction. All light rays are parallel. Directional lights are also known as Infinite lights, since they are best used to capture infinite (or near infinite) light sources, such as the sun. Since the light source is far away, the intensity of the light does not vary inside the scene. Although directional lights are displayed in the *3D Views* as little suns, their position inside the scene is not relevant. Only the direction at which they point is important. Selecting the option that makes directional lights always point at the camera avoids misunderstanding this, since it links the direction of the light to the position of the light source (it is usually good practice).

Note:

The light is not necessarily in front of the camera.

Light sources can be moved, resized, and rotated, but they cannot be twisted. Resizing a light source only affects the power of limited range lights (point and spot lights), as opposed to global lights such as the sun. Since directional lights are infinite, they are not affected by resizing. You cannot assign materials to light sources, but you can define the color of the light they emit.

To create a simple light, click on the **Light** icon in the left toolbar. If the requested type of light is not directly available, you will have to unfold the icon (). Alternatively, you can create lights by selecting the requested item from the **Object | Add Light** sub-menu.


Default scenes are created with a single white directional light, usually named “Sunlight”.

The [Editing Lights](#) section gives more details on the different light settings.



Area Lights

VUE features several types of area lights:

- **Light Panel:** this is a rectangular panel of light. Light is emitted from the entire surface of the rectangle. You can change the aspect-ratio of the panel by resizing it. To create a light panel, right-click (Ctrl + click on Mac) on the **Light** icon in the left toolbar and select the last icon to the right () , or select **Light Panel** from the **Object | Add Light** sub-menu.
- **Daylight Portal:** a daylight portal is used in interior scenes to give you more accurate interior lighting in relation to the exterior light of the scene. The portal is most often applied as a rectangular area light placed in the window of a room scene. The intensity and color emitted by the daylight portal in the interior scene is obtained from the sky outside the window. To create a daylight portal, select the object that you want a daylight portal applied to, and select **Add Daylight Portal to Object** from the menu in the World Browser. Or, you can just select the Daylight Portal from the Light icons on the left of the user interface and then place the panel manually in the scene.
- **Light Emitting Object:** light emitting objects are objects that emit light from all points at their surface. Any object can be converted into a light emitting object by selecting the menu command **Convert to Area Light** from the **Object** menu, or from the popup menu in the *3D Views*. When you convert an object into a light source, the materials of the object will be automatically converted into light gels. So if your object is red in some parts and green in others, the light emitted by the object will automatically be red in some parts and green in others.

Warning:

once an object has been converted to an area light, it cannot be converted back to a normal object.

The *intensity of the light* emitted by area lights is proportional to the surface of the light. So the larger the light, the more powerful it will be. You can also adjust the power of the light source (as well as other parameters of the light) using the controls in the *Object Properties* panel.



Group Objects

Group objects can be of two types:

- the first type behaves as simple “bags” into which member objects are placed to organize the scene; such objects are known as **Groups**, the **3D Text** group being a special kind of group that enables editing of the text using the *Text Editor* (read more about 3D Text [here](#) and see [here](#) for details on the *Text Editor*),
- the second type of group objects operates on the member objects; depending on the type of the actual operation, these objects are classified as **Boolean objects** or as **Metablobs/Hyperblobs**.

Boolean Operations

Boolean operations come in three flavors: **Union**, **Intersection** and **Difference**. **Boolean objects** can be made easily by combining member objects together to yield an incredible variety of new shapes.

Metablobs

Metablob objects “blend” their different member primitives together as if they were melted together. Metablobs are great for modeling organic shapes. You can only use basic primitives to create [Metablob objects](#).

Hyperblobs





Hyperblobs are Metablobs that make use of HyperTextures and can be used to create highly detailed and realistic rock shapes and rock formations. Jagged grottos, broken scree, and an infinite variety of rock and stone shapes can be generated using Hyperblobs. Hyperblobs can be baked to remove any parts of the HyperTexture that are disconnected from the main object (an artifact of standard HyperTextures).




Working with Group Objects

Like other objects, group objects may be moved, resized, rotated, twisted, and have materials assigned to them. All of their member objects will be moved, resized, rotated, and twisted accordingly. If you assign a new material to a group object, all member objects will take on that material. If all member objects of a group object don't use the same material, the picture of the material displayed in the *Object Properties* exhibits a pair of arrows to let you browse the different materials on the [Aspect Tab](#).

Member objects may be added, removed or modified inside a group object using the *World Browser*.

To create a group object, first select all the objects you would like to be grouped together, then click the requested icon: for simple groups, click the group icon on the left toolbar (). For Boolean objects, select the requested operation from the unfoldable Boolean operation icon (). If the operation you want to use is not directly available, you will have to unfold the icon (). To create a Metablob object, select the Metablob icon (). Alternately, you can use the menu commands from the **Object | Group objects, Object | Make Boolean Object** or **Object | Make Metablob Object** menus.

You can group together any number of objects, of any type (except the camera). Boolean objects cannot include light sources. Metablobs can only include basic primitives.

To ungroup objects, select the group that you wish to destroy, and click the **Ungroup** icon (). Alternately, you can use the menu item **Object | Ungroup**.

To “un-Metablob” a group of primitives, simply ungroup it as above.

If a group is made up of the same type of objects (e.g. a group of polygon meshes, toruses...) it can be edited like if it were a single polygon mesh or torus, etc.



Ventilators


In VUE, ventilators are objects that are used to control the wind applied to plants in a very local manner. They have no influence on the scene, other than affecting trees and plants.

In Vue, ventilators come in two flavors:

- **Directional ventilators** that blow wind in a preferred direction, and
- **Omni ventilators** that blow wind in all directions equally.

Like other objects, ventilators can be moved around, rotated and resized. You cannot assign materials to ventilators. Resizing a ventilator will increase the intensity of the wind generated by the ventilator. You can also control the intensity of the wind generated by the ventilator using the *Object Properties* panel when a [ventilator](#) is selected. Ventilators can be turned into “attractors” by giving them a negative wind intensity.

[Linking ventilators](#) to other objects can give the impression that the object is causing the wind (for instance to simulate the effect of a helicopter landing in a field).

To create a Ventilator, either select it from the **Particles Effector/Directional Ventilator** icon in the left toolbar () , or select the menu commands **Object | Create | [Omni / Directional] Ventilator**. The *Object Properties* panel is used to modify the [properties of ventilators](#).

Cameras

The last type of object is probably the most useful: it is the object that will turn your scene into a finished, colorful picture. So without a camera, you would never see your scene! This is why it is impossible to destroy the camera.

The camera may be moved, rotated and resized. However, it may not be twisted. Resizing it only changes the focal length. No materials can be assigned to it.


Camera Target

When you select the camera, a small cube appears in front of the camera. This is known as the camera target; it is a helper designed to facilitate aiming of the camera. It also shows the distance to the focus point: objects that are at the same distance from the



camera than the target will be in focus. You can attach the target to a given object so that this object stays in focus whatever its movement relative to the camera. You can read more about Camera and Camera Target options [here](#).

Advanced Options

If you double-click on the camera in the *3D views*, or if you click the **Edit** button () on the top toolbar or in the *World Browser*, or select **Edit Object** from the **Objects** menu when the camera is selected, the *Advanced Camera Options* is displayed. Using this dialog, you can adjust the aspect ratio of the picture, as well as apply post-processing effects to your renders.




Editing Objects

This section details how objects can be modified inside Vue.

Objects may be moved, rotated and resized interactively inside the orthogonal *3D Views*. Alternatively, these operations can be done in a more precise manner using the [Numerics](#) tab of the *Object Properties* panel.

Selecting Objects



Object manipulation handles

Before you make any modification to an object, this object has to be selected.

Selected objects are displayed in red, and are framed by black dots inside the active *3D View*. They are also highlighted in the *World Browser*.

To select an object, you can either click on it once in the active *3D View*, or click on the name of the object in the *World Browser*. Clicking while in an inactive *3D View* will activate it first, so you'll have to click again on the object to select it.



Inside the 3D Views

You can also select objects in the *3D Views* by clicking outside all of the objects, and dragging the mouse: a rectangle is drawn, showing the **Selection area**. All objects that have their center inside the rectangle will be selected.

If the object that you want to select is inside a group, clicking on it in the 3D Views will select the whole of the group. The only way to select that particular object is to use the *World Browser*.

You can extend the selection to other objects by clicking on them while the **Shift** key is down. Re-selecting an object that is already selected will deselect it.


Pressing **Control** when you click to select an object will cause all objects under the mouse cursor to be selected at the same time.

Inside the World Browser

Selecting objects using the *World Browser* will reposition the *3D Views* so that the selected objects appear centered. This feature can be turned off using the [Options](#) dialog.

You may select group members independently, provided the group is expanded on the list. You can extend the selection with the standard Windows™ method, using **Shift** and **Control** keys.

By Category

Selecting all objects that have either the same type, Preview color, or material as the selected object is achieved through menu commands or icons in the top toolbar (). The menu commands are found in the **Edit | Select by** sub-menu.

Also, clicking on a material inside the *Summary of Materials* or in the **Material** tab of the *World Browser* will select all the objects that use this material. This feature can be turned off using the General Preferences tab of the [Options dialog](#).



Walking through a Selection of Objects

When you have selected several objects, it becomes possible to select each of these objects, in turn. This operation is called “walking the selection”, and is achieved by the menu item **Edit | Walk Selection** or by pressing the **Tab** key. So, repeatedly pressing **Tab** will cycle you through all of the objects that you have selected, letting you pick out any one of them easily. For instance, you could select all objects under the cursor (pressing **Control** as you click), and then walk through those objects pressing **Tab** until you reach the one you want.

Deselecting Everything

To deselect everything, either click on an empty part of the *World Browser*, click outside the icons on the toolbars, press **Escape**, or select **Deselect All** from the **Edit** menu.

Moving Objects

Selected objects can be moved in three ways:

- By dragging them using the mouse or the [Position Gizmo](#) inside the *3D Views*
- By entering new values for **Position** in the [Numerics](#) tab of the *Object Properties* panel
- By using the nudge keys. The nudge directions are relative to the active *3D view*. One nudge is equal to 5 units of distance. Pressing **Shift** as you nudge divides the nudge distance by ten.

If you select one of the axis arrows, movement will only be possible along that axis. If you select the square between two axis arrows, movement will be constrained to the corresponding plane.



Rotating Objects

Selected objects can be rotated in two ways:

- By using the rotation handles or the [Rotation Gizmo](#) inside the 3D Views
- By entering new values for **Rotation** in the [Numerics](#) tab of the Object Properties panel.

Inside the 3D Views

Selected objects are framed by 4 rotation handles in the corners. Selecting any of the handles and move them around will rotate selected objects around the axis that is perpendicular to the view.

The angle of rotation, together with the number of revolutions of the object are displayed in the *Status Bar*.

Resizing Objects

Selected objects can be resized in two ways:

- By using the resize handles or the [Size Gizmo](#) inside the *3D Views*
- By entering new values for **Size** in the [Numerics](#) tab of the *Object Properties* panel.

Inside the 3D Views

Selected objects are framed by 8 square dots, 4 in the corners, and 4 in the middle. These dots are either black (one object selected) or white (multiple objects selected).

The corner dots are called the '**Resize globally** handles'. Clicking on one of these, and dragging the mouse away will resize the selected objects. If you press **Shift** while dragging, the objects will be resized equally in all directions, thus keeping the proportions. If you press **Control** while you drag the mouse, the selected objects will be resized equally only along the two directions of the view, leaving the third direction unchanged.



The middle dots are called '**Resize in this direction** handles'. Clicking on one of these, and dragging the mouse away will resize the selected objects in the direction indicated by the shape of the mouse cursor. The top and bottom handles resize the objects vertically (inside the view); the left and right handles resize horizontally (inside the view). This sizing is linked to the views, and not to the object itself, which means that rotated objects may be twisted when resized. However, since size values in the **Numerics** tab are relative to the object, using **Numerics** is a good way of avoiding this problem. In addition, pressing **Control** as you drag these handles will resize the objects along their own axes.

If you select the **Resize around opposite corner** option, the object will be resized relative to the corner that is opposite to the resize handle you are currently using.

Note:

By default, resizing is proportional along all 3 axes, but if you hold the **Shift** key down while resizing, the resizing will be free. If you hold the **Control** key down, resizing will be proportional only along the two axes of the view. If you hold the **Alt** key down, resizing will be symmetrical around the center point.

Twisting Objects

Selected objects can be twisted in two ways:

- By resizing rotated objects inside the *3D views*
- By entering new values for **Twist** in the **Numerics** tab of the *Object Properties* panel.

Inside the 3D Views

Twisting objects occurs whenever you try to resize objects that have already been rotated.

Here is a method to twist objects interactively: select the object you would like to twist and rotate it 45°. Then resize it horizontally, and rotate it back to the initial position. The object is twisted.



Using the Numerics Tab

Numerical values for twisting are a bit complex. Basically, they will twist one axis of the object towards another axis. This is rather difficult to visualize, so the best is to try it out. However, please understand that, due to complex matrix operations, twisting and untwisting objects in several directions may not restore the initial object conformation.


Hint: twisting objects can be a powerful method for achieving impressive terrain overhangs.

Gizmo Manipulators

The gizmo manipulator tools let you move, rotate and resize objects accurately. Gizmos were first introduced in Maya, and have now become an industry standard for manipulating objects.

Because gizmo manipulators may not be as intuitive as VUE's original way of manipulating objects, you will have to disable gizmos to use the old method. To disable gizmos, simply select the menu command **Display | Gizmos | Show Gizmos**. You can revert back to the default manipulation system using the same command.

Selecting the Appropriate Gizmo




There are three different types of gizmos, each being designed for one type of operation (moving, rotating or resizing). The gizmos let you either move/rotate/resize along one, two or all axes simultaneously. The X, Y and Z axes are identified by different colors (the X axis is red, the Y axis is green and the Z axis is blue).



Whenever an object is selected, a gizmo will appear at its center, in all 3D Views. You can change the current gizmo tool either by:

- Clicking on the gizmo tool swatches alongside the gizmo (see opposite – you can hide these swatches; see further down),
- Selecting the menu command **Display | Gizmos | [Position/Rotation/Size] Gizmo**, or
- Using the corresponding menu shortcuts.

The Position Gizmo




The position gizmo is used to move the selected objects. It features two or three arrows indicating the view axes. When you drag the mouse over one of these arrows, it will change color to yellow. If you click and drag the arrow, the selected object(s) will move accordingly along the corresponding axis.

Close to the point where the arrows meet, you will notice that the arrows are joined by a square area. If you move the mouse over that area, it will turn yellow; click and drag that area to move the selected object(s) accordingly along the two axes joined by the square.

You can also move objects by clicking and dragging inside the selection box, outside the gizmo (see below). This behavior can be disabled using the **Display | Gizmos | Allow Moving Outside Gizmo**.



The Rotation Gizmo




The rotation gizmo is used to rotate the selected objects. It features 3 concentric circles in 3D, each representing one axis of rotation. In the orthogonal views when in **Global coordinates** mode, you will only see one circle, as the circles for the two other axes are seen from their side and hence appear as lines.

If you move the mouse over one of these circles/lines, it will turn yellow. Click and drag to rotate the selected objects accordingly around the corresponding axis.

If you move the mouse away from the circles, but close to the center of the gizmo tool, a gray disk will appear. If you click and drag that disk, the selected objects will be rotated around the two axes of the view.

When the position gizmo is selected, you can also rotate the selected objects around the axis perpendicular to the view by moving the mouse to the outside of the selection box. The mouse cursor will change to a rotation cursor. This feature can be disabled using the **Display | Gizmos | Show Rotation Handles**.

The Size Gizmo



The size gizmo is used to change the size of the selected objects. Like the position gizmo, the size gizmo displays two or three arrows (terminated by a round dot instead of a cone).

When you drag the mouse over one of these arrows, it will change color to yellow. If you click and drag the arrow, the selected object(s) will be resized accordingly along the corresponding axis.

Close to the point where the arrows meet, you will notice that the arrows are joined by a triangular area. This is the proportional resize area. If you move the mouse over that area, it will turn yellow; click and drag that area to resize the selected object(s) globally (the resizing will be proportional along all three axes).

If you move the mouse away from the proportional area, you will notice a strip joining the axes two by two. This strip turns yellow when you move the mouse over it; if you click and drag this area, the selected object(s) will be resized proportionally along the two axes joined by the strip (and only those two).

You can also resize objects by clicking and dragging the corner dots of the selection box, outside the gizmo (see below). This behavior can be disabled using the **Display | Gizmos | Show Resize Handles**.

Gizmo Coordinates

Gizmos can operate in any one of four different coordinate systems:

- **Local coordinates:** the coordinates are that of the selected object, allowing you to manipulate the object according to its current orientation. For instance, in this mode, the position gizmo along Z will always move the pyramid in the direction in which it is pointing.
- **Global coordinates:** in this mode, the gizmos operate along the view axes, whatever the orientation of the object.
- **Parent coordinates:** in this mode, the gizmos operate in the coordinate system of the first object that was selected. This mode is only useful if you have selected several objects (if not, it is identical to the local coordinate gizmo).
- **View coordinates:** in this mode, the gizmos operate in the coordinates system of the view. This is useful when working in the camera view, because objects will always be moved, rotated or resized in a plane facing the camera.

You can switch from one coordinate mode to the other using either:

- The gizmo coordinate swatches alongside the gizmo (you can hide these swatches; see further down),
- Selecting the menu command **Display | Gizmos | [Local/Global/Parent/View]**



Coordinates, or

- Using the corresponding menu shortcuts.

Customizing Gizmo Behavior

By default, the VUE gizmos let you move, rotate and resize objects whichever the current gizmo is:

- To move objects, click and drag the mouse outside the gizmo but inside the selection box.
- To rotate objects, click and drag the mouse on the outer side of the selection box, near the border of the selection box (the cursor changes to rotation handles).
- To resize objects, click and drag on the corner dots of the selection box. By default, resizing is proportional along all axes. If you hold down the Shift key, resizing will be free, and will only take place along the two axes of the view. Press Ctrl on top of Shift to make the resizing proportional along the two view axes. Pressing Alt will toggle the resize center from the opposite corner to the object's center.

In order to offer identical behavior to other 3D software packages, you can disable each one of these options individually:

- Use the menu command **Display | Gizmos | Allow Moving Outside Gizmo** to forbid moving objects other than using the Position Gizmo.
- Use the menu command **Display | Gizmos | Show Rotation Handles** to forbid rotating objects other than using the Rotation Gizmo.
- Use the menu command **Display | Gizmos | Show Resize Handles** to forbid resizing objects other than using the Size Gizmo.

You can change the size of the Gizmo tools using the menu commands **Display | Gizmos | [Reduce/Increase] Gizmo Size**. This has no effect on the way gizmos operate.



You can hide the gizmo swatches that appear alongside the gizmo tool by using the menu command **Display | Gizmos | Show Gizmo Helpers**.

If you are used to the original VUE way of manipulating objects, you can disable gizmos altogether using the menu command **Display | Gizmos | Show Gizmos**.



Changing Object Material

This is achieved through any of the following:

- Clicking the **Load material** button () in the **Aspect** tab of the *Object Properties* panel; this opens the Material Browser. Select the new material to replace the existing material and press **OK**. If the material is animated, the material is entirely replaced, and not just the current frame. You can leave the Material Browser open without blocking the rest of the interface.
- Double-clicking the material picture; this will open the [Material Editor](#) for the selected objects.
- Using menu command **Object | Change Object Material**; this opens the Material Browser.
- Dragging and dropping the material from another object onto the object (inside the *3D Views* or the *World Browser*).
- Using the *Summary of Materials* to load () or edit the material used by the object. All objects using this material throughout the scene will be affected.

If several objects that use different materials are selected, the picture of the material in the *Object Properties* panel will display the first material in the list. A pair of arrows will appear at the bottom of the material to let you browse through the different materials.

You can edit all these materials simultaneously by selecting the **Edit All Materials** command from the popup menu that appears when you right-click on the material preview (Ctrl+Click on Mac). The *Material Editor* will appear, displaying the settings for the current material. Any changes you make to that material will be applied to all the materials (provided that the materials are compatible with such changes). This is particularly useful for instance, when, after importing an object, you want to change the shininess of all its materials.

If you load an [EcoSystem](#) material, VUE will ask you if you want to populate the selected object with EcoSystem elements according to the EcoSystem material's population rules.

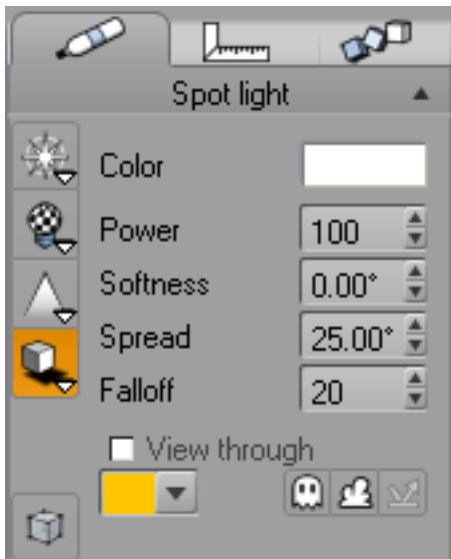
Underneath the material picture is a **Scale** control that lets you adjust the scale of the material when it is rendered inside your scene.



Changing Object Preview Color

The color used to draw the selected objects in the *3D Views* can be modified by picking a new one from the drop-down list. You may choose any color out of the 8 available. The preview color does not affect the color of the object when it is rendered. When a color is selected, the views are flashed to show the new color.

Editing Lights



Light properties panel

Once selected, lights are edited through the **Aspect** tab of the *Object Properties* panel. To the left of this panel, you will notice 4 icons that let you modify the behavior of your lights:



Lens Flare: click on this icon to activate [lens flares](#) on the selected light. Right-click to display the Lens Flare options menu. Double-click on the icon to display



the *Lens Flare tab* in the [Light Editor](#). With the **Light Editor** displayed, you have the option to modify all settings of this light.



Gel: click on this icon to add a [light gel](#) to the selected light. Right click on the icon to display the Gel options menu.



Volumetric light: click on this icon to activate volumetric light effects on the selected light. Right click on the icon to display the [Volumetric light](#) options menu.

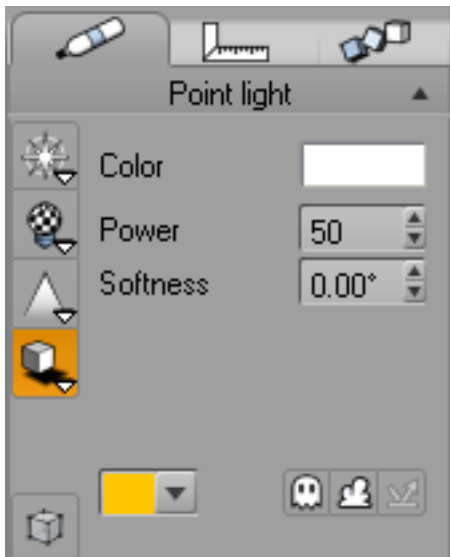


Shadow and Lighting: click on this icon to enable or disable shadows for the selected light. Right click on the icon to display the [Shadow and Lighting](#) options menu. Double-click on the icon to display the *Shadows tab* in the [Light Editor](#).

Unlike other objects that have a black default color, lights have a yellow default color.

Depending on the type of light that is selected, other controls will appear in the *Object Properties* panel.

Point Light and Quadratic Point Light



Point light properties



If the selected [light object](#) is a point light or a quadratic point light, the *Object Properties* panel displays as opposite.

At the top of the **Aspect** tab is the **Light color** control. Double-clicking this will open the *Color Editor*, letting you select a new color for the light.

Underneath this control is the **Power** control. This governs the intensity of light emitted by the light source. The more powerful the light, the greater the range in which it can light up objects. This control also visually affects the length of the rays cast from the light source in the *3D views*. Resizing a light source yields the same result. You can enter **Negative** values, in which case the light will “cast darkness” instead of casting light.

The next control is the **Softness** control. Turning this up to a non-zero value unleashes a powerful feature of VUE: soft, natural looking shadows. This truly realistic effect is possible because VUE can handle surface lights. The greater the value, the larger the surface of the light, and the more progressive the transition from light to shadow. A value of 5° usually yields nice results. However, since surface lights are much more computationally demanding than standard lights, it is recommended that you use such lights judiciously.



Switch On/Off: this option turns the light on or off without losing the light settings. Clicking on the light’s icon in the *World Browser* will also turn on or off the light.



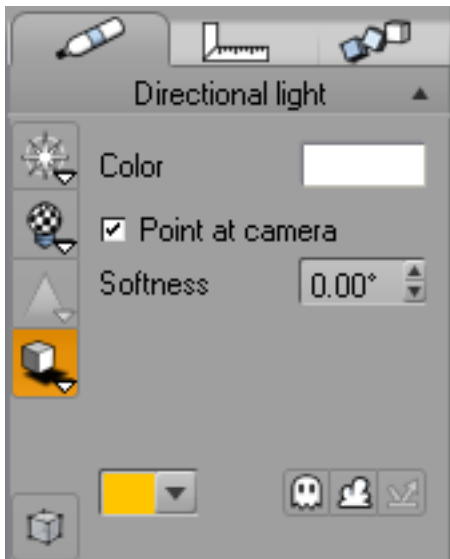
Influence Spectral Clouds: this option allows the light to influence spectral clouds.



Exclude from radiosity: this option is only available when the “Global Radiosity” lighting model has been enabled on the [Light tab](#) in the Atmosphere Editor. Select this option if you don’t want the influence of this light to be taken into account when processing the radiosity solution. The time it takes to compute the radiosity solution is directly connected to the number of lights in the scene, but not so much to the power of these lights. By excluding lights that do not contribute significantly to the illumination of the scene, you can speed up the rendering of Global Radiosity significantly.



Directional Light



Directional light properties

If the selected object is a directional light, the *Object Properties* panel displays as opposite.

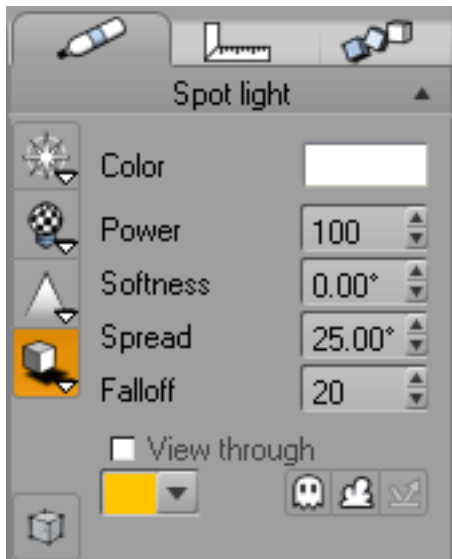
This is identical to the previous panel apart from the missing **Power** parameter (the power of directional lights is controlled solely via the color of the light) and the extra **Point at camera** checkbox. When selected, the position and direction of light are linked in such a way that it always points at the camera.

Note:

This doesn't mean the light is necessarily in front of the camera. Having this option selected is good practice, since it avoids misunderstanding the fact that only the orientation of directional lights is important.



Spot Light and Quadratic Spot Light



Spot light properties

If the selected object is a spot light or a quadratic spot light, the *Object Properties* panel displays as opposite. This is identical to the point light properties panel, apart from three extra controls: Spread, Falloff and View through.


- **Spread:** adjusts the spread of the light cone. The greater the value, the larger the angle of the cone of light. The maximum is 90°, which will spread light everywhere in front of the light source.
- **Falloff:** governs the speed of transition between light and dark on the edges of the cone. The greater the value, the more gradual the transition.
- **View through:** this option displays the scene in the *Main View* as if it were seen through the spotlight.

The area that is lit by the light is displayed as two concentric circles representing respectively the beginning and the end of the light falloff area. This is a very accurate and efficient way of adjusting the lighting from a spotlight.

The field of view of the *Main View* is automatically adjusted to match the spread angle of the light. Although the preview render displayed in the [Camera Control Center](#) still represents the view from the camera, the controls will act upon the currently selected spotlight instead of acting upon the camera. As soon as you deselect the spotlight, the *Main View* will flip back to the view as it was before selecting the spotlight.



Lens Flare

The **Lens Flare** icon () is a toggle icon. If it is orange or highlighted, it means that this light has a lens flare.


If no lens flare is defined, clicking on the icon will assign the default lens flare to the light (and the icon will turn orange). If a lens flare is assigned to the light, clicking on the icon will open the *Light Editor* on the **Lens Flare** tab (please read [here](#) for full details). If several lights are selected, any modifications will be applied to all the selected lights.

The Lens Flare icon is a [double action](#) icon. If you activate the icon's second action, a popup menu will appear. The options in this menu are:

- **Edit Lens Flare:** select this command to open the *Light Editor* on the **Lens Flare** tab.
- **No Lens Flare:** if this option is checked, the light has no lens flare; if not, select it to remove the lens flare from the light.
- **Default Lens Flare:** if this option is checked, the light has the scene's default lens flare for that type of light; if not, select it to assign the scene's default lens flare to the light. If the default lens flare is disabled, the word '(off)' will be appended to the menu label. Default lens flares are automatically assigned to the new lights you create. If you modify the [default lens flares](#), you will be modifying the lens flares of all lights that have a lens flare that isn't custom.
- **Custom Lens Flare:** if this option is checked, the light has a custom lens flare. This means that the light has a lens flare effect that is different from the scene's default lens flare. This is generally the case if you have modified the lens flare of a light. If this option isn't checked, selecting it will open the *Lens Flare Editor*.
- **Copy Lens Flare:** copies the lens flare settings of the light to the clipboard, so it can be pasted onto another light.
- **Paste Lens Flare:** pastes the lens flare settings in the clipboard to the current light. Use the Copy/Paste commands to transfer settings from one light to another.



Light Gel

In the real world, gels are colorful pieces of transparent plastic placed in front of light sources to give them colors or patterns. The **Gel** icon () lets you define an electronic counterpart for your lights.

The Gel icon is a toggle icon. If it is orange or down, it means that this light has a gel.

If no gel has been assigned to this light, clicking the icon will open the Material Browser, asking you to select a gel material for the light. You can load any [Simple material](#) to use as a gel, but only the material's color will be taken into account. Mixed, Layered and Volumetric materials cannot be used as gels. Once you have assigned a gel to the light, the icon turns orange.

If a gel is assigned to the light, clicking on the icon will open the *Light Editor* on the **Gel Tab**, letting you modify the colors of the gel as required.

The **Gel** icon is a [double action](#) icon. If you activate the icon's second action, a popup menu will appear. The options in this menu are:


- **Light Has a Gel:** if the light already has a gel, selecting this option will remove the gel. If no gel has been assigned to the light, selecting this option will open the Material Browser letting you select which gel you wish to use for the light.
- **Edit Gel:** this command is available only when the light has a gel. Click on this button to open the *Light Editor* on the **Gel** tab.
- **Flat Gel Type:** this option, together with the following, indicates what projection method will be used for the gel. If the Flat Gel Type option is selected, the gel is considered mapped to a plane placed in front of Spot lights, or mapped to a box placed around Point lights. Gel projection types are not available for Directional lights.
- **Spherical Gel Type:** when this option is selected, the gel is mapped to a sphere placed around the light source. Gel projection types are not available for Directional lights.
- **Copy Gel:** select this command to copy the gel to the clipboard. The gel is copied as a standard material.
- **Paste Gel:** select this command to paste the material currently in the clipboard into the light's gel. Use Copy/Paste to transfer gels from one light to another.
- **Load Gel:** selecting this command opens the Material Browser letting you select a material to be used as a gel. Keep in mind that only the color information is used for gels.
- **Save Gel:** select this command to save the current gel. The gel will be saved as a



standard material.

If several lights are selected, the modifications will be applied to all the selected lights.

Volumetric Light

The third icon to the left of the *Object Properties* () is the **Volumetric Light** icon. This is a toggle icon. If it is orange or down, it means that the light is volumetric.

If the light is not volumetric, clicking on the icon will make the light volumetric. If the light is volumetric, clicking on the icon will open the *Light Editor* on the **Volumetric** tab. This dialog lets you customize the volumetric behavior of the light (e.g. the intensity of the beam, whether smoke is visible in the beams, etc.).

Volumetric lighting applies to directional lights (e.g. the sun) only when the volumetric or spectral [atmosphere model](#) is used. If the current atmosphere model is the standard one, this icon will be disabled. Select a volumetric or spectral atmosphere to be able to create volumetric rays for the sun.


The Volumetric Light icon is a [double action](#) icon. If you activate the icon's second action, a popup menu will appear. The options in this menu are:

- **Volumetric Light:** this command has the same effect as clicking on the icon. If the light is volumetric, selecting this command will make it non-volumetric. If it isn't volumetric, it will become volumetric.
- **Edit Volumetric Settings:** select this command to open the *Volumetric Light Options* dialog.
- **Copy Volumetric Settings:** select this command to copy the light's volumetric settings to the clipboard.
- **Paste Volumetric Settings:** select this command to paste the volumetric settings from the clipboard to the light. Use Copy/Paste to transfer volumetric options from one light to another.

If several lights are selected, the modifications will be applied to all the selected lights. Keep in mind that volumetric lights are far more complex to render than standard lights.



Shadow and Lighting

The last icon to the left of the *Object Properties* panel () is the **Shadow and Lighting** icon. If it is orange or down, it indicates that the light casts some amount of shadow.

If the light doesn't cast any shadows, clicking the icon will enable shadow casting for that light. If the light casts some amount of shadow, clicking on the icon will open the *Light Editor* on the **Shadow** tab. This dialog lets you edit shadow and lighting options.

The Shadow and Lighting icon is a **double action** icon. If you activate the icon's second action, a popup menu will appear. The options in this menu are:

- **Casts Shadows (xx%)**: the xx value is the shadow density percentage of the light, where 0% means no shadows, and 100% means full shadows. Select this command to set full shadows for the light. If the light already has full shadows, selecting this option will have no effect.
- **No Shadows**: select this option to remove all shadows from the light. If the light already is non-shadowing, this option will have no effect.
- **Edit Shadows**: select this command to open the **Shadow** tab of the *Light Editor*, letting you select intermediate shadow densities and modify the lighting properties of the light.
- **Edit Lighting**: select this command to open the **Lighting** tab of the *Light Editor*, letting you modify the lighting properties of the light.
- **Edit Influence**: select this command to open the **Influence** tab of the *Light Editor*, letting you select which objects are influenced by the light.

If several lights are selected, the modifications will be applied to all the selected lights. Keep in mind that non-shadowing lights render much more rapidly than shadowing lights – even when the shadow density setting is low.

Light Panel


With the Light Panel, at the top of the **Aspect** tab is the **Color** control. Double-clicking this will open the Color Editor, letting you select a new color for the light.

Underneath this control is the **Power** control. This governs the intensity of light emitted by the light source. The more powerful the light, the greater the range in which it can light up objects. This control also visually affects the length of the rays cast from the light source in the 3D views. Resizing a light source yields the same result. You can enter **Negative** values, in which case the light will “cast darkness” instead of casting light.

For the Daylight Portal, there are no settings on the **Aspect** tab as its settings are controlled by the exterior light settings. There is a **Daylight Portal Editor** to increase




quality of lighting if desired and to invert its direction easily.



'Daylight Portal Editor'

Editing Terrains

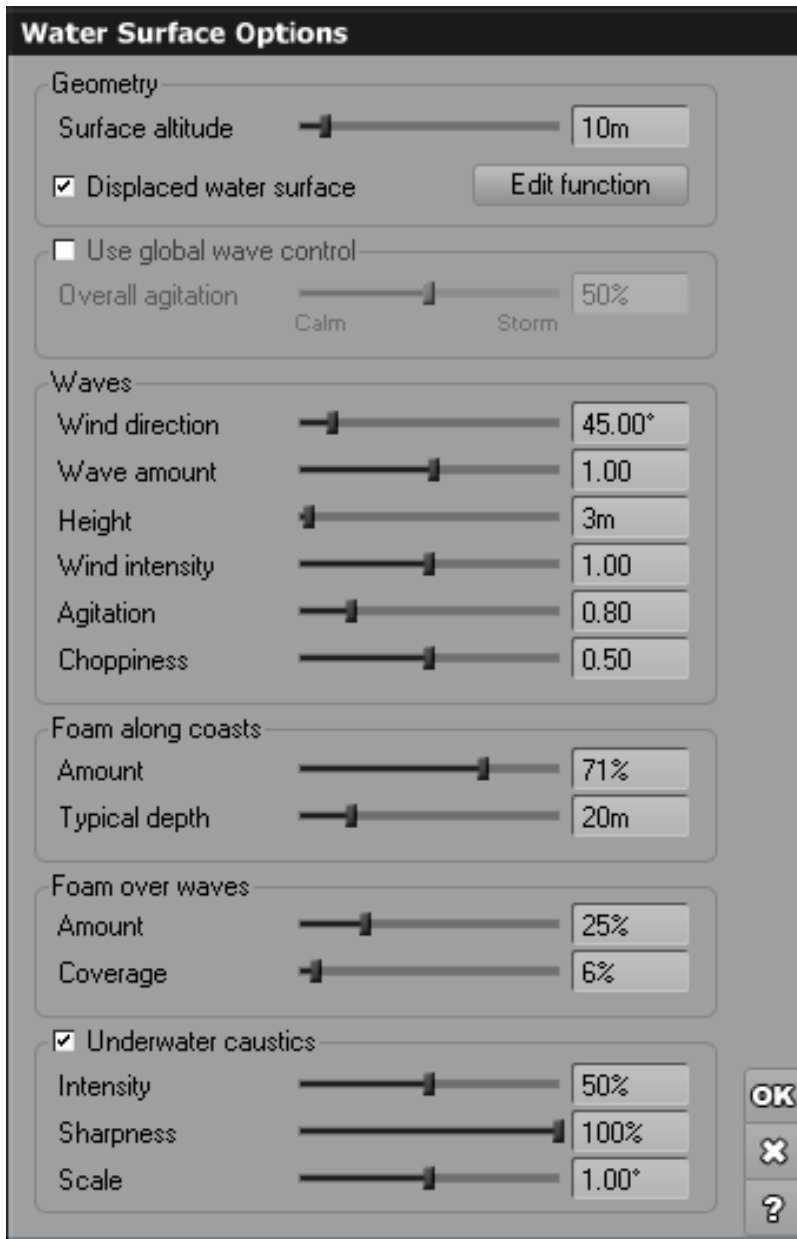
Terrains may be edited using the *Terrain Editor*. This is accessed by either:

- Double-clicking on the terrain in the *3D Views* or in the *World Browser*.
- Clicking on the **Edit object** button () on the top toolbar, when the terrain is selected.
- Using the menu command **Object | Edit Object**.

In the section under the Vue Editors, you will find detailed information for the [Terrain Editor](#).

Editing Bodies of Water






The Water Editor: a complete set of tools to model water

The *Water Surface Options* dialog provides a set of powerful tools designed to let you easily create and animate realistic water surfaces. It automatically distributes foam at



the surface of the water in a realistic manner, taking into account any surrounding terrains.

To access the *Water Surface Editor*, either:

- Double-click on the infinite water plane in the *3D Views* or in the *World Browser*,
- Click on the **Edit object** button () on the top toolbar, when the infinite water plane is selected,
- Use the menu command **Object | Edit object**.

Geometry

The **SurfaceAltitude** slider lets you easily adjust the altitude of the water plane.

By default, water planes are just a perfectly flat surface, with bump mapping added to fake the waves. If you want a more realistic water surface, where the waves are created out of real geometry, check the **Displaced water surface** box. This will automatically convert the flat water plane into a pseudo-infinite procedural terrain. During the conversion process, the bump mapping settings are ported to the [procedural altitude](#) production function, thus yielding similar visual results (with the added realism of true wave geometry).

The way the *Water Editor* works is that it implements and controls a complex graph in the [Function Graph](#). You can use the settings in this dialog to change the look of the water surface, but you can also further customize this look even further by editing the water material directly.

If the water surface is converted into a procedural terrain by ticking the **Displaced water surface** option, you can access the altitude production function of the underlying procedural terrain by clicking the **Edit Function** button. Obviously, this button is only available when the **Displaced water surface** option is checked and **Global wave control** is deactivated.

Use Global Wave Control

The **Use global wave control** box is selected by default. When this box is selected, you can adjust the overall aspect of the water surface by setting the **Overall agitation** slider to make the water calm or stormy. **Wind direction** can also be changed.

If this box is unchecked, the rest of the controls on this dialog become available. This is only if a metawater material is selected. If you are using another type of water material, these controls will not be available.



Waves

Wind direction: this parameter controls the direction in which the wind is blowing, as seen from above (the azimuth). A value of zero will make the wind blow from left to right in *Top view*. A value of 90° will make the wind blow from top to bottom in *Top view*. There is no relationship between this wind setting and the wind or breeze effects applied to plants.

Uncheck the **Use global wave control** box to gain access the other controls in this section:

- **Wave amount:** this parameter lets you adjust the overall amount of the waves created. Values greater than 1 will make more waves, while values less than 1 will make less waves.
- **Height:** this parameter controls the typical height of the waves. It is only available when the **Displaced water surface** option is checked.
- **Wind Intensity:** this parameter controls the intensity of the wind. Higher values will realistically lead to higher waves and rougher water surfaces.
- **Agitation:** this parameter lets you adjust the overall velocity of the waves created. Its effects are only visible in animations. Values greater than 1 will make the waves move faster at the surface of the water, while values less than 1 will slow down the waves.
- **Choppiness:** this parameter controls the shape of the waves. Small values will yield soft round waves, whereas high values will produce choppy waves that are sharp at their top.

Foam Along Coasts

If the water plane intersects a terrain, these controls will add foam near the shore.

- **Amount:** this parameter lets you adjust the overall amount of foam created.
- **Typical depth:** this parameter lets you adjust the overall depth that the water must have, in order to begin creating foam.

Foam Over Waves

- **Amount:** this parameter lets you adjust the overall amount of foam created on the wave.
- **Coverage:** this parameter lets you adjust the overall coverage of foam created on the wave.



Underwater Caustics

- **Intensity:** this parameter lets you adjust the intensity or brightness of the caustics. Higher values will produce brighter lighting over focusing regions and darker lighting elsewhere.
- **Sharpness:** this parameter lets you adjust the sharpness of the caustics. A lower setting blurs and softens the caustic effect.
- **Scale:** this parameter lets you adjust the scale, or size, of the caustic pattern.

Caustics will automatically appear at their maximum sharpness at some focus depth which depends on the caustics scale, while slowly going out of focus as depth increases or decreases.

Depending on the water material model, caustics will not only be generated over underwater surfaces, but also through the water medium, producing realistic beams of light through the water.

To enable this volumetric effect, the water material transparency should be set to use the **Physical transparency**, in either **Direct volumetric light** or **Indirect volumetric light mode**.

For more information about the Physical transparency, refer [here](#).

Changing the MetaWater Material

The *Water Surface Editor* is designed to present an easy-to-use front-end to a special type of water material, known as a MetaWater material. The MetaWater material is built using a special type of MetaNode in each material layer.

You can easily change the MetaWater material by loading any material from the **MetaWater** material collection and assigning it to the water plane.

Note:

if you wish to change the MetaWater material after having selected the **Displaced water surface** option, you will first have to deselect this option, and then re-select it after loading the new MetaWater material. If you don't do this, the water displacement won't correspond to the new material.

You can create your own MetaWater materials, but you need to ensure that the new materials are based on the same MetaNodes as other MetaWater materials, if not, it may not be possible to control the material using the *Water Surface Editor*.


You can also edit the MetaNodes that are used to construct the MetaWater material,



but again, great care must be taken not to change the interface of the MetaNode (do not remove or rename any published parameters).

Editing Plants

Applying Wind



Wind control, available in the Top view

When you select a plant, you will notice in the *Top View* a blue triangle inside a small circle. This is the Wind control. Simply click on the triangle and drag it to set the amount of wind that is applied to the plant at the current time. The longer the arrow, the stronger the wind; the direction of the arrow indicates the direction in which the wind is blowing. [Wind can be animated](#) like other *Object Properties* (e.g. size, orientation...).

Note:

Wind starts to be applied to the plant only when the wind arrow is dragged outside the blue circle. If you place the wind arrow back inside the circle, wind will be removed from the plant. This is designed for easier removal of the wind. You can monitor the intensity and direction of the wind in the *Status Bar*.

If several plants are selected, the new wind setting will apply to all selected plants. That



way, you can apply the same wind effects to several objects.

The length of the wind arrow represents the intensity of the wind. The length of the arrow is not affected by the zooming in the views. Just like in the real world, the effect of the wind on the plant depends on the type of plant. For instance, long and thin plants such as reeds will be strongly affected, while more solid plants such as trees will be less affected.

Note:

If the Wind arrow doesn't appear in the *Top View*, make sure wind is enabled in the [Wind](#) tab of the *Atmosphere Editor*.


Ventilators

If you bring a ventilator close to a plant, the plant will be influenced by the amount of wind produced by the [ventilator](#).

Editing Plant Geometry

You can edit the shape of the plant using the *Plant Editor*.

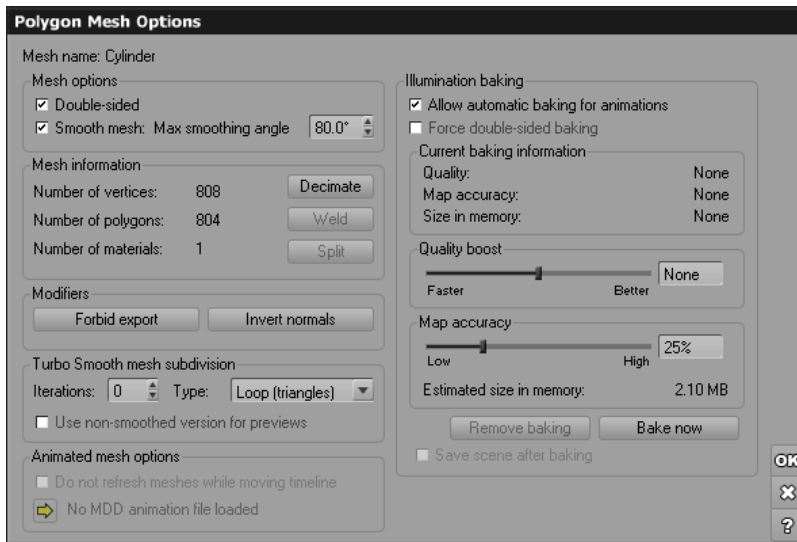
This is accessed by either:

- Double-clicking on the plant in the *3D Views* or in the *World Browser*.
- Clicking on the **Edit object** button () on the top toolbar, when the plant is selected.
- Using the menu command **Object | Edit Object**.

Use the Plant Editor for [editing plants](#). All of the basic parts of the plant can be modified.




Editing Polygon Meshes



Polygon Mesh Options

Polygon Mesh Options Dialog

Although the actual geometry of polygon meshes cannot be modified inside VUE, you do have access to some render options. These options are selected in the *Polygon Mesh Options* dialog, which is accessed by either:

- Double-clicking on the polygon mesh in the *3D Views* or in the *World Browser*,
- Clicking on the **Edit object** button () on the top application toolbar, or in the *World Browser* toolbar when the polygon mesh is selected,
- Using the menu command **Object | Edit Object** when the polygon mesh is selected.

Aside from displaying information relative to the complexity of the selected polygon meshes, this dialog lets you adjust rendering characteristics of the meshes.

The **Double sided** checkbox indicates that the polygons making up the mesh should be traced from both sides. This is generally the case, and keeping this option selected is recommended. However, if you are sure that your mesh will support tracing with one



sided polygons, unchecking it can slightly improve render speed. If it is not the case, you will notice holes in the rendered object.

When **Smooth mesh** is selected, the surface of the polygon mesh object is smoothed by averaging the normal vectors of all polygon faces, giving a less rough, “polygonal” aspect to the object. Sharp angles will be rubbed from the surface. You can adjust the maximum angle between polygons where smoothing should take place by using the **Max smoothing angle** box. For instance, a value less than 90° (e.g. 80°) would preserve right angles inside a cube, while still smoothing other, less angular features.

Mesh Information

This group displays information about the currently selected mesh(es), such as the total number of vertices, of polygons, as well as the number of materials used throughout the object(s).

This group also provides a set of buttons to modify the architecture of the selected object(s):

- **Decimate:** press this button to display the *Mesh Decimation Options* dialog. Decimation is a powerful feature that attempts to reduce the number of polygons in an object while maintaining as much as possible the original object geometry. The resulting object will be lighter than the original, and will render more rapidly. On this dialog, use the slider to indicate the level of decimation.
- **Weld:** this button only appears when several meshes are selected. If you click on this button, VUE will generate a single polygon mesh from all the currently selected meshes. This is useful when you have an object made up of lots of different parts. Welding them all together ensures two things: faster processing and certainty that relative positions/orientations of the parts won't be modified accidentally.
- **Split:** Use this button to split the mesh into a group of meshes according to the material assigned to each polygon (resulting in one mesh per material).

Modifiers

- **Forbid Export:** press this button to forbid exporting to other 3D applications the selected object(s). Be advised that once you have pressed this button, you cannot allow exporting again. This is useful when transferring data to other parties and you don't want this other party to be able to use your objects in other applications.
- **Invert Normals:** press this button to invert all the surface normals of the selected object(s). Pressing it again will revert to the initial situation.



Turbo Smooth

This option allows you to smooth a mesh. You have three options for turbo smoothing: **Catmull-Clark** (for quads), **Loop** (for triangles) and **Dynamic** (for stretched polygons only). This option can also be used for animated mesh in order to improve the rendering quality.

- **Type:** select **Catmull-Clark** when smoothing quads; select **Loop** when smoothing triangles. **Dynamic** only subdivides stretched polygons.
- **Iterations:** Since **Catmull-Clark** and the **Loop** algorithm subdivide all of the polygons, this setting is used to limit the level of subdivisions applied.

If you select **Dynamic**, a **Quality Boost** slider displays for you to set the smoothing quality.

- **Low poly in OpenGL and in the preview:** check this option to display the non-subdivided mesh in OpenGL display and in the preview renders which will speed up these rendering operations.

Illumination Baking

Illumination baking is only available when you have activated one of the lighting models that takes indirect lighting into account. Illumination baking is a complex process that involves the evaluation of indirect lighting over the entire surface of the object, and the creation of a separate texture channel to hold that illumination information. Only indirect lighting is taken into account in the baking process. Direct lighting (light received directly from light sources, or shadows cast by other objects) is computed separately at render time. For full details on illumination baking, please turn [here](#).

The controls in this frame let you customize the illumination baking process for that particular object.

- **Allow automatic baking for animations:** when this option is selected (the default), this object will be candidate for [automatic illumination baking](#) when rendering an animation. Uncheck this option to prevent the object's illumination from being baked when rendering an animation (e.g. for objects that move a lot).
- **Animated mesh imports:** If you imported a mesh with an accompanying MDD animation file, this will show up here. You have the option of not refreshing meshes while moving the timeline.
- **Force double sided baking:** when baking the illumination of an object, VUE will analyze the geometry of the object in order to determine if the illumination needs to be computed only on one side, or on both sides of an object. For instance, if the object is a single polygon, VUE will bake the illumination on the two sides of this polygon, as the polygon may be seen from one side or the other. Selecting the “Force double sided baking” option forces VUE to bake the illumination on the two sides of the object, even if it determines that illumination is only really needed on



one side. This is typically useful when you need to travel “inside” a closed mesh. The “Current baking information” group displays information about the current baking status of the object: the overall “quality” and “map accuracy” used at the time of last baking, as well as the total size occupied in memory by the illumination map.

Note:

Saving an object with illumination baking will save the illumination information together with the object.

- **Quality boost:** use this setting to adjust the rendering quality of the indirect lighting during baking. This setting is relative to the current [Advanced Effects Qualityrender](#) setting and the quality boost setting of the *Atmosphere Editor*’s lighting model controls.
- **Map accuracy:** use this setting to vary the size of the illumination map. Larger values will mean more detailed illumination, but at the expense of longer baking times and higher memory requirements. You should adjust the accuracy of the map according to the largest size on screen of the object during the entire animation.

Hint:

If you are baking the illumination on a mesh that has a complex geometry, the maximum setting for map accuracy may not be enough to achieve perfect results. In such cases, you should not hesitate to go beyond the maximum value by entering values larger than 100%.

- **Estimated size in memory:** this displays an estimate of the amount of memory that will be occupied by the illumination map after baking is completed. More memory may be required during the actual baking process.
- **Remove baking:** click this button to remove all baking information for this/these object(s).
- **Bake now:** press this button to begin the illumination baking process for the selected object(s). The baking of the illumination requires takes place in two steps: first the construction of the illumination map, and then the actual processing of the indirect lighting. The total time required to bake the object depends on the object’s complexity as well as the desired quality. This process can last several hours (but it’s an investment for the rendering of the animation). A progress bar will appear in the status bar to let you know the progress of the baking process. When an object has an illumination baking map, it appears yellow in the *World Browser*.
- **Save scene after baking:** select this option if you want VUE to automatically save the current scene after completing a bake operation. Because quality baking of objects can take a lot of time, this option is useful to ensure that you don’t lose the result of the baking.

Note:



Although baking the illumination of an object can take a considerable amount of time (several times the actual rendering time of a single frame in the animation), subsequent rendering of the animation frames can be accelerated in such a considerable manner that this “investment” in baking time will be recouped significantly. In some tests we ran (involving a flythrough of the Sponza atrium model), baking of the Sponza mesh took one day (at a very high quality setting), but subsequent rendering of the animation only took 2 additional days whereas the rendering of the “unbaked” animation would have required over one month!

Published Parameters

The **Published Parameters** feature copies specific settings from the **Object Graph** that you may need to change often and places them in a more convenient location for easier manipulation. With objects, these parameters display with the *Polygon Mesh Options* dialog.

Parameters available for publishing show as underlined when the cursor is moved over the field name. Click the underlined field name to select. A parameter name is supplied and a group name is asked to improve the display of the published parameter.

Decimating Imported Objects



From left to right: 29120, 13570, and 1895 polygons frogs!

Polygon mesh decimation, also known as polygon reduction, is a very powerful feature that will attempt to reduce the number of polygons in an object while maintaining as much as possible the original object geometry. The resulting geometry will be lighter than the original, and will render more rapidly. This feature is useful for example when you want to render a large object that is seen from a long distance – in such a case, you may not need to maintain all the details in the object geometry, and a decimated version may suffice. Another situation where mesh decimation is useful is when you want to do some quick test renders of very large objects (e.g. buildings). You can do the test with a decimated version, and then restore the full geometry for the final rendering.



Mesh decimation is controlled using the *Mesh Decimation Options* dialog. There are two ways of accessing this dialog:

- Once the object is imported, click the **Decimate** button in the *Polygon Mesh Options* dialog (see above), or
- At the time of importing the object, select the **Decimate on import** option in the *Import Options* dialog. The *Mesh Decimation Options* dialog will appear once the import is complete.

There is a single parameter in the *Mesh Decimation Options* dialog. This parameter, called **Level of decimation** is used to control the number of polygons in the resulting, decimated version of the object. If you choose no decimation, the object will be unaffected. The stronger the level of decimation, the less polygons in the resulting object, but also the stronger the distortion of the geometry. This parameter is somewhat similar in its behavior to say, the quality setting of JPEG compression.

Note:

Mesh decimation is a complex process that can take a long time to complete. A progress bar displays the progress of the decimation process in the *Status Bar*.

An estimate of the number of polygons and vertices in the mesh version of the object is displayed below the **Level of decimation** slider.

Once you have decimated a polygon mesh, you cannot regenerate the polygons that have been removed. The only way to restore the full object geometry is to re-import the object *without decimation*. See [here](#) for details about re-importing polygon meshes.

Baking Objects to Polygons

This is also a powerful feature that will convert any object in a VUE scene into a polygon mesh approximation. The word “approximation” is important here, as some objects used in VUE simply don’t have a polygonal equivalent (for instance, this is the case with spheres, plants...).

There are various reasons why you might want to convert an object into a polygon mesh. For instance, if you have a very complex Boolean object, converting it into polygons may speed up rendering considerably.

Note:

materials are not directly affected by the conversion. However, the slight modifications in the object’s geometry may cause some differences in the rendering of the material.

If the object uses [displacement mapping](#), the displacement information will be baked into



the polygon geometry.

Note:

Lights and cameras cannot be converted to polygons.

Baking to polygons is controlled via the *Mesh Baking Options* dialog. This dialog can be accessed by selecting the menu command **Object | Bake To Polygons**.

Note:

Once you have converted (baked) an object to polygons, you cannot revert later to the initial, “unbaked” version of the object.

The *Mesh Baking Options* dialog has a unique parameter called **Bake quality**. This parameter controls the accuracy of the conversion process. The higher the quality, the more precisely the polygon version of the object will match the initial object. But also the higher the number of polygons used in the object, and thus the longer it will take to render and the higher the memory requirements to handle this object. The default value for this quality setting corresponds to the *Background draw thread's* preview quality setting (see [here](#) for details on adjusting this preview quality). Please note that baking to polygons is a complex process that can take a long time to complete. A progress bar displays the progress of the baking process in the *Status Bar*.


An estimate of the number of polygons and vertices in the mesh version of the object is displayed below the **Bake quality** slider.

Note:

You can optimize the conversion of objects to polygon meshes by baking at a higher quality setting, and then applying some amount of decimation to the resulting mesh.



Direct Re-Posing of Rigged Meshes



Skeleton Editor

Rigged mesh objects (objects containing skeletons) can be re-posed directly inside VUE. You can create custom poses and movements with meshes that have been converted to the VUE Rigged Mesh format. Currently, VUE rigged meshes and Collada imports are supported.

Rigged meshes created in 3DS Max can be brought into VUE by converting them into *.vob* format in Max using the *Max to VUE exporter*. If the 3DS Max rigged mesh being imported does not have a skeleton, it will be imported as a standard mesh.

The *Max to VUE exporter* is available in the archive from which you installed your product: alongside the *Setup (Win).exe* folder, you will find an *Export Plugins* folder, with a



Windows subfolder. Read the *Max to Vue Exporter.pdf* document contained for instructions on how to install and use the plugin for your 3DSMax version.

Just load a rigged mesh into VUE – you can use *Dave* which comes with the VUE **Extra Contents** and can be found in the Objects Browser under **Characters**. Click on the figure and the *Skeleton Editor* displays. The bones should also be displayed in wireframe if you have that option checked (from the VUE menu, **Display | Rigged Meshes**).

To select a bone, click on it. If the parent object containing the geometry is already selected, press the CTRL key while clicking on the bone. The +/- icons beside each bone are for expanding/contracting the list of bones. To avoid expanding the hierarchy manually, you can double-click on a bone in the view and it will automatically expand the hierarchy until the selected bone.

Inverse Kinematics (IK) are applied to rigged meshes in VUE. IK does affect a chain of bones (for example two bones in an arm, two in a leg, or several in the neck or tail of a dinosaur). The bones that belong to an IK chain are specified in 3ds Max, when creating the rigged mesh. When the model is exported to VUE, a helper is visible at the end of the chain, allowing to control it.

The *MaxToVue* exporter also supports the “swivel angle” and “swivel target” features from 3ds Max. The swivel angle is defined relatively to the parent bone of an IK chain, and define a plane in which the chain will be contained. The swivel target is an helper, exported with the skeleton, that also define a plane (with 2 other points: the origin and the end of the chain) in which the bones of the IK chain will be contained (for example, if you have an IK chain for an arm, with the end of the chain located on the wrist, using a swivel target allows to control the elbow position).

The **Enable IK** box can only be checked when a helper representing the goal of an IK chain is selected. When this checkbox is unchecked, the bones belonging to the IK chain are not constrained so they can be transformed separately.

Click to change some options like **Double-sided** and, in VUE Pro versions, to modify **Turbo Smooth** settings which can be applied to all meshes rigged to the skeleton. The other options on this dialog are disabled as they do not apply to rigged meshes.

Bones can be modified two ways:

- **Manual method:** Bones can be modified like any other object using the **Gizmos**. Changing the position of a bone leads to the rotation of its parent (except if it's a helper).
- **Numeric method:** Numeric values can be entered into the *Skeleton Editor*, in the **Bone parameters** fields. Just select the bone in the editor or by clicking on the bone in the viewport and change the values as needed. Unlike moving the bones manually, using the numeric method will not affect its parent's rotation.



Morphers can also be accessed in the *Skeleton Editor*. After selecting a morpher in the **Skeleton hierarchy**, its list of target geometries appears on the right and an up/down button with a numeric field allows you to change the weight of each target.

Keep also in mind that the coordinates are always defined in the local space of the selected bone, and that the rotation is the composition of 3 separate rotations (around XYZ axes). Depending on the angles composition, changing one value around a specific axis may rotate the bone in an “unexpected” way, which is normal.

Poser imports cannot be modified using these techniques. You need to use the reposing techniques described in the previous section.

Animation Creation

In **Auto-Keyframing** mode, animation is created after modifying a bone at a non-zero time. Keyframes are then automatically added each time a bone is modified. If the **Auto-Keyframing** mode is not activated, bones’ keyframes are added by clicking on **Add Keyframe** in the *Timeline* menu. The whole skeleton animation can be deleted by right-clicking on the character in the *Timeline* and selecting **Destroy Object Animation**.


You can add animation to your character from pre-saved motion files. On the *Skeleton Editor* window, select the **Open** icon to display the Motion Browser and select any motions you may have saved there. Back on the *Skeleton Editor* window, use the **Duration** field to set the duration of the motion and click on the **Apply** button.

From the *Skeleton Editor* you can:

- **Load:** a new animation (which must match the existing bones hierarchy).
- **Save:** the current skeleton’s animation in a *.vom* file.
- **Clear:** the animation.
- **Scale:** the duration of the animation like in the *Timeline* except that it affects only the selected skeleton and not the whole scene. This would have to be changed using the animation toolbox.
- **Repeat:** changes the animation play mode (repeat or not).




Editing the Torus



Torus Options dialog

Aside from the standard rotate and squash fun, you can also modify the thickness of the torus rim (also known as the outer diameter). This is done using the *Torus Options* dialog. To display the *Torus Options* dialog, you can either:


- Double-click on the torus in the *3D Views* or in the *World Browser*,
- Click on the **Edit object** button () on the top toolbar, when the torus is selected,
- Use the menu command **Object | Edit Object** when the torus is selected.

Torus Options Dialog

This very simple dialog displays a single slider that lets you interactively adjust the outer diameter of the torus. Just slide the notch to the right or to the left, and watch as the torus swells or shrinks in the *3D Views*.

Editing 3D Text

Aside from the standard rotation, sizing and movement, you can also edit the shape and content of a 3D Text object using the *Text Editor*. This is accessed by either:

- Double-clicking on the text object in the *3D Views* or in the *World Browser*.
- Clicking on the **Edit object** button () on the top toolbar, when the 3D Text is selected.
- Using the menu command **Object | Edit Object**.

Note:



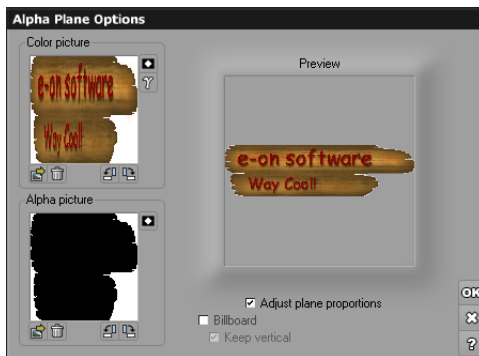
When you create a 3D Text, the text is placed in a special kind of group. If you ungroup the text object and then weld all the characters together, the rendering speed of the text may be slightly increased. However, after doing this, you can no longer edit the text.

Editing Alpha Planes

Alpha Planes may be edited using the *Alpha Plane Editor*. This is accessed by either:

- Double-clicking on the alpha plane in the *3D Views* or in the *World Browser*.
- Clicking on the **Edit object** button (🔧) on the top toolbar, when the alpha plane is selected.
- Using the menu command **Object | Edit Object**.

Alpha Plane Editor







Alpha Plane Options dialog

This dialog prompts you to select the pictures that will be used to map the object. The first picture (**Color picture**) will be used to produce the colors of the object, while the second (**Alpha picture**) will be used to generate transparency. Using this transparency picture, you can create objects with custom profiles. If the picture that you select for the colors has embedded transparency information, this information will automatically be loaded into the Alpha picture.

Click the **Load** icon (📄) below the picture previews, or double-click on the picture previews to open the Picture Browser and load a picture. You can rotate the pictures by



using the  and  arrows. You can also invert the pictures using the  button. This is particularly useful when the Alpha information is not encoded as expected. To remove a picture, click the **Remove** icon () below the picture preview.

The **Preview** displays a preview of what the Alpha plane will look like in your scene.

Check the **Adjust plane proportions** to automatically match the proportions of the Alpha plane to those of the pictures you loaded. Unchecking this option may result in unwanted stretching of the pictures.

Billboards

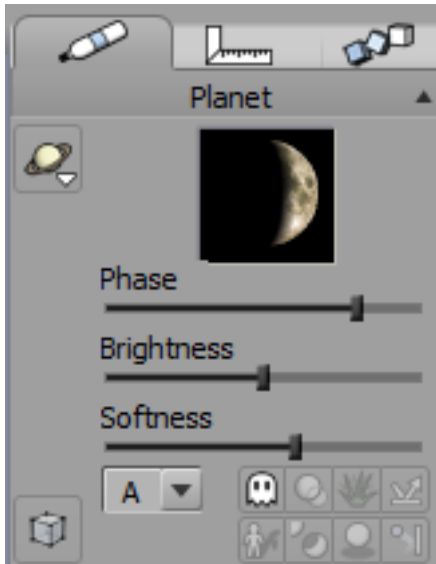
Billboards are automatically oriented to always face the camera (or whatever is looking at them). Billboards are particularly useful to easily add real-world photos to your scenes (e.g. to add people to an architectural rendering). They can also be used to create simple smoke or fog effects.

Billboard: simply check this option to turn the Alpha plane into a billboard. Now, the alpha plane will automatically be oriented to always face the camera. If the billboard is seen through a reflection, it will be seen as though it were facing the reflecting object. This is useful to avoid betraying the fact that the object is only a plane (the fact that the object is flat could show up in reflections).

Keep vertical: if this option is selected, the billboard will be oriented in such a way that it always faces the camera while remaining vertical. This is particularly useful when adding tree or character billboards to architectural projects – as you really want them to stay vertical under all conditions.

Editing Planets






Planet properties panel

To modify the aspect of planets, select the planet and use the [Aspect](#) tab of the *Object Properties* panel.

If the selected object is a planet, the *Object Properties* panel displays as opposite:

At the top of the panel is a preview of the planet. This preview is useful when adjusting the phase.

If you check the icon bar to the left of the panel, you will notice an extra icon (). This is the **Select planet** icon. Clicking this icon displays a list of planets that you can pick from to change planet. Selecting **Saturn** will appropriately create a planet with rings around it.

If you select the **Custom** option, the Picture Browser will appear, letting you select the picture of your choice to map the planet. If you want to create a planet with a ring from a picture, first select Saturn, and then change the type of planet to **Custom**.

Just below the planet preview, you will find three sliders that you can use to customize the look of the planet:

- **Phase:** use this slider to modify the direction of the side of the planet that is lit by the sun.

Note:

There is absolutely no relationship with the actual position of the sun in your scene.




So if realism is your goal, you should be careful to manually match the phase of your planets to the position of the sun in the scene.

- **Brightness:** this controls the brightness of the planet. Typically, if the sky is bright, the planet brightness should be low, and if the sky is very dark (e.g. by night), the planet should be bright. Planets, as opposed to all other types of objects, are placed behind the clouds. So they will be masked by clouds (which is quite realistic, after all, but can become a problem if your scene contains layers of clouds used to simulate stars – in which case the stars will appear in front of the planet).
- **Softness:** controls how gradual is the transition from lit up areas of the planet, to parts in the dark.

Planets are processed in the order in which they appear in the *World Browser* (that is the last ones are placed behind the first ones). You can change the order of the planets by moving them around in the *World Browser*.

As with other objects, you can use the resize and rotation tools in the 3D views to modify the size and orientation of the planet.

Boolean Objects



On the left a Boolean object, on the right the polygon version (OpenGL preview)

Boolean objects, also known as “Constructive Solid Geometry” objects, let you combine simple primitives (spheres, cubes...) into incredibly complex objects, using Boolean operations. For a nice example of a Boolean object, please take a look at the **Fortress** sample object. You can find this by selecting **Load Object** from the **File** menu.



Boolean objects behave like groups: they act as a “bundle” into which you can put objects. You can put as many member objects as you want inside a Boolean object. The difference between groups and Boolean objects, is that Boolean objects can combine their members using various Boolean operations.


Because VUE Booleans are computed at runtime, they retain all the accuracy of the primitives used in the Boolean operation (e.g. you won’t see polygons appear in a “Booleaned” sphere). The drawback to this computation at render time is that Boolean operations are slow to render.

If you need faster rendering Booleans but don’t require the accuracy, you can convert them to a polygon mesh using the **Object | Bake To Polygons** menu command (see [here](#)). This way, VUE will generate a polygon mesh representation of the Boolean operation.

Three types of Boolean objects are available, depending on the operation used to combine their member objects:


- Boolean union
- Boolean intersection
- Boolean difference

Boolean Union



Group of overlapping glass spheres





Boolean Union of glass spheres

Member objects of a Boolean union are “welded” together to form one unique object. The result is noticeable mainly with transparent materials (like glass), where a Boolean union will produce a continuous object with no internal edges (if you group two overlapping glass spheres, you will still notice an area where both spheres overlap).

Making a Boolean union from opaque objects yields exactly the same result as simply grouping the objects. Since Boolean unions require more computation, it is recommended that you use groups instead of Boolean unions where applicable.

All member objects of a Boolean union have the same role, so the order in which you place members inside the Boolean union is of no importance.

Boolean Intersection



Boolean Intersection of a cube and a cone


The result of a Boolean intersection is the object created where members overlap; Boolean intersections always yield smaller objects than any of their members.

The result of a Boolean intersection between a flattened cube and a cone will be a truncated cone (see opposite illustration).



All member objects of a Boolean intersection have the same role, so the order in which you place members inside the Boolean intersection is of no importance.

Boolean Difference



Boolean Difference: sphere minus cylinder



Combining materials in Boolean objects: Boolean Difference of a glass sphere and a checkerboard cube

A Boolean difference takes the first of its members and “subtracts” all subsequent members from it. This means that, unlike Boolean unions or intersections, Boolean differences give a different role to each of the member objects. The first member will be the base object from which subsequent members are “dug out”.

For instance, making a Boolean difference between a sphere and a stretched cylinder will give a sphere with a hole in it, where the cylinder used to be. If you inverse the order and place the cylinder first inside the Boolean difference, the result will be different: it will yield a cylinder, with the sphere missing.



This powerful tool lets you dig out of an object any number of other objects. For instance, making a tower with dozens of windows can be handled by one single Boolean difference!

When you make a Boolean difference, you have to pay attention to the order in which you select member objects before creating the Boolean difference: the first object you select will be the “solid” one; all others will be dug out. If you get it wrong, you can change the order of the members using the *World Browser*.

In VUE, Boolean objects are implemented in an extremely powerful way, letting you combine as many objects as you like inside one single Boolean object.

To have some noticeable effect, a Boolean object must contain at least two member objects. If you create a Boolean object with only one member, you will generate unnecessary computation.


If member objects in the Boolean objects don't all have the same material, the material of each object will be retained on the parts of the Boolean Object that pertain to this object (see opposite screenshot).

You may add, remove, or change the order of member objects inside a Boolean object by using the *World Browser*. Simply unfold the Boolean object, and drag objects into, or out of it. You may, of course, make Boolean objects that use other Boolean objects!

VUE will compute a polygonal preview of the result of the Boolean operation. This preview will be displayed shortly after creating or modifying the Boolean operation.

Metablobs





Blending the shapes of two spheres and a torus. Notice how the torus material is blended with the sphere material.

Metablobs “blend” together the shapes of the different primitives that are part of the group. You don’t need to have several primitives in the Metablob object in order to see the effects of the Metablob operation: Metablobs will remove all angular shapes and replace them with round, organic looking shapes (e.g. a cube will have all its edges rounded). Because there are no sharp edges in a sphere, there is no point in creating a Metablob from a single sphere (as this would simply create another sphere – only more complex to render).

Metablobs work with all the following primitives: sphere, cube, cylinder, cone, pyramid and torus. If you resize or rotate the primitives inside the Metablob, the shape of the resulting Metablob will be modified.

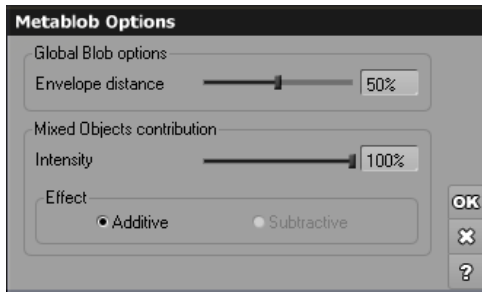
If member objects in the Metablob don’t all have the same material, the material of the different primitives will be blended together according to the contributions of each primitive.

You may add, remove, or change the order of member objects inside a Metablob by using the *World Browser*. Simply unfold the Metablob and drag objects into, or out of it. The order of the primitives inside the Metablob is not relevant. You can only drag basic primitives into Metablobs. If you create a Metablob from other Metablobs, all the primitives of the various Metablobs will be assembled together into a unique new Metablob.

VUE will compute a polygonal preview of the result of the Metablob. This preview will be displayed shortly after creating or modifying the Metablob.



Metablob Options



The Metablob Options dialog

On top of editing and adjusting the elements inside the Metablob object, you can also customize the way the Metablob blends its member objects using the *Metablob Options* dialog.

The Metablob options can be edited either globally, or on a per object basis. To edit the Metablob options globally, either:

- Double-click on the Metablob in the *3D Views* or in the *World Browser*.
- Click on the **Edit object** button (🔧) on the top toolbar, when the Metablob is selected.
- Use the menu command **Object | Edit Object**.

To edit the Metablob options of only one or several member objects, either:

- Double-click on the Metablob object member(s) in the *3D Views* or in the *World Browser*.
- Click on the **Edit object** button (🔧) on the top toolbar, when the Metablob object member(s) is(are) selected.
- Use the menu command **Object | Edit Object**.

There are three settings that you can act upon to modify the look of the Metablob:

Envelope distance: this setting is global to the entire Metablob. It controls the overall distance between the center of all the member objects and the envelope of the actual Metablob.

Object contribution: this setting can be adjusted on a per object basis. It controls the influence of the selected object(s) on the look of the final Metablob. This feature is very useful to fine tune the geometry of the Metablob. If not all member objects have the same contribution, the entry field will remain empty. If you enter a value or drag the slider, the new value will be assigned to all the member objects of the Metablob.



Effect: by default, all member objects contribute to the overall shape of the Metablob by adding their geometry to that of the other members. This is known as the **Additive** effect. However, if you turn one of the member objects **Subtractive**, the geometry of that object will be “removed” from that of the other member objects, resulting in a smoothly blended hole.

If you are editing a torus, the Metablob options will be displayed below the standard *Torus Options* setting **Torus thickness**.

Hyperblobs


A Hyperblob is a hypertextured Metablob that is baked at render time, removing any parts of the Hypertexture that are disconnected from the main object (an artifact of standard Hypertextures). Baking occurs at a resolution depending on distance to camera, to avoid building unnecessary details.

To create a Hyperblob:

- Use primitives to create a shape as you would a Metablob.
- Assign a [Hypertexture](#) to one of the primitives.
- Right-click on the Metablob icon to create the Hyperblob. You will be prompted if you wish to continue and assign the Hypertexture to all primitives. If there is no Hypertexture, a default Hypertexture material will be assigned.



Hyperblob Options



The Hyperblob Options dialog

On top of editing and adjusting the elements inside the Hyperblob object, you can also customize the way the Hyperblob blends its member objects using the *Hyperblob Options* dialog.

The Hyperblob options can be edited either globally, or on a per object basis. To edit the Hyperblob options globally, either:

- Double-click on the Hyperblob in the *3D Views* or in the *World Browser*.
- Click on the **Edit object** button (🔧) on the top toolbar, when the Hyperblob is selected.
- Use the menu command **Object | Edit Object**.

There are three settings that you can use to modify the look of the Hyperblob:

Global Blob Options – Envelope distance: this setting is global to the entire Hyperblob. It controls the overall distance between the center of all the member objects and the envelope of the actual Hyperblob.

Mixed Object(s) contribution – Intensity: this setting can be adjusted on a per object basis or on the entire Hyperblob. It controls the influence of the selected object(s) on the look of the final Hyperblob. This feature is very useful to fine tune the geometry of the Hyperblob. If not all member objects have the same contribution, the entry field will remain empty. If you enter a value or drag the slider, the new value will be assigned to all the member objects of the Hyperblob, or the selected object(s) if all aren't selected.



Effect: by default, all member objects contribute to the overall shape of the Hyperblob by adding their geometry to that of the other members. This is known as the **Additive** effect. However, if you change one of the member objects to **Subtractive**, the geometry of that object will be “removed” from that of the other member objects, resulting in a smoothly blended hole. The **Subtractive** option is not available if the entire object is being edited.

Hypertexture baking – Maximum resolution: A maximum resolution can be set between 20 and 250 to avoid baking excessive times and memory consumption. After baking, two post-processes apply:

Smooth mesh – Max smoothing angle: Mesh normals are smoothed according to a specified maximum angle.

Keep only largest single chunk: “Disconnected” components are removed. Only the largest compact subset of the mesh is kept.

Cache baked mesh between renders: Check this option to eliminate the baking of the Hyperblob each time you render. The Hyperblob will be baked the first time you render and the information saved with the Hyperblob. It should be noted that this can take up a lot of memory and increase the time it takes to save a scene. The baked mesh is updated when needed, like when a Hypertexture is modified.

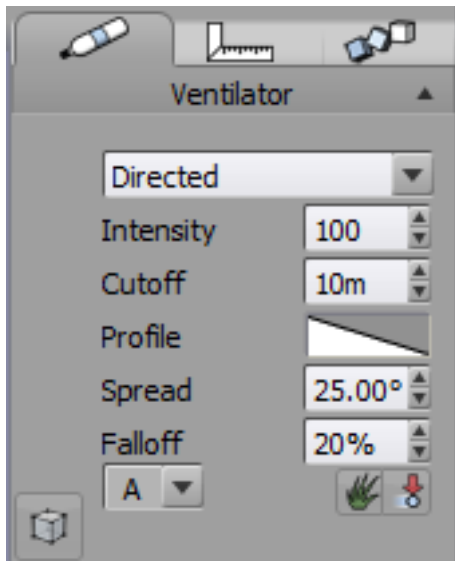
Using Hyperblobs in EcoSystems

When using Hyperblobs in EcoSystems, you should be aware that the position, rotation and scale of each Hyperblob instance is not accounted for when baking the Hypertextures as this would be incredibly resource intensive. The instances themselves are still rotated, however.

Ventilators



Ventilators



Ventilator properties panel

If the selected object is a Ventilator, the *Object Properties* panel displays as opposite. This is very similar to the light properties panels (see above).

Ventilator type: this lets you select the type of ventilator. There are two different types of ventilators in VUE: **Omni** and **Directional** ventilators. Omni ventilators will blow wind in all directions with equal intensity, whereas directional ventilators will blow wind in a specific direction only. Depending on the type of ventilator you selected, some of the controls below will become active.


Intensity: this setting controls the intensity of the wind generated by the ventilator.

Cut-off: this is the distance at which the ventilator ceases to affect plants.

Profile: this is a filter that lets you define how the intensity of the wind evolves with distance from the ventilator. By default, the intensity drops down linearly with distance. Double-click on the filter preview to load a preset filter, or edit the filter to create a custom intensity profile.

Influence EcoSystems (🌿): when this option is selected, the ventilator will also affect plants that are inside EcoSystems. Because the number of plants affected this way is potentially enormous, ventilators will affect in priority those plants that are closest to the ventilators.



Influence particles (): when this option is selected, the ventilator will also affect EcoParticles.

Directional Ventilators

On top of the above settings, directional ventilators also define the following parameters:

Spread: this setting adjusts the spread of the cone in which wind is blown. The greater the value, the larger the angle of the cone. The maximum is 90°, which will blow wind everywhere in front of the ventilator.

Falloff: this setting controls how suddenly the wind intensity drops near the edges of the cone. The greater the value, the more gradual the transition.

Replacing Objects



Work with simple objects, render with hi-res models

Vue lets you replace any object by another one. For instance, if you want to replace a pyramid with a cone, a light by another one of a different type, a primitive by a more



complex object such as a plant or an imported mesh, etc...

To replace an object with another one, select the object to be replaced, and then select one of the **Edit | Replace By** commands from the main menu, or from the popup menus in the *3D Views* or in the *World Browser*.

The Replace command comes in two flavors:


- **Replace By (Keep Proportions):** this version of the command replaces the object without deforming the new object. The selected object will be replaced by the new one by fitting the largest dimension of the new object to the corresponding dimension of the source object. Other dimensions are modified proportionally. This means that the original proportions of the new object are not altered. This is important if you want to preserve new object's general shape. However, scaling factors applied to source object will also be applied to new object. This provides you with a flexible way of replacing objects (if you want to replace an animated deformation object, for instance). It's ideal when you want to replace an object by a plant or a mesh, and you don't want the proportions of the object to be modified. You are also able to select more than one object for replacement. All objects selected will be replaced by the same object with the scaling of each replaced object maintained.
- **Replace By (Fit Object):** this version of the command modifies the dimensions of the replacing object so that it matches the replaced object exactly. This means that the new object will be scaled in such a way that its bounding box fits exactly the bounding box of the replaced object. You are also able to select more than one object for replacement. All objects selected will be replaced by the same object with the scaling of each replaced object maintained.

A maximum of information regarding the replaced object is transferred to the replacing object (as applicable). For instance, if you replace a terrain with a symmetrical terrain, the terrain geometry will be preserved. If you replace one type of light with another, color, power, etc. will be preserved. This is also true of animation and linking properties. This is very important as it means that you can setup complex animations with basic objects (such as cubes) and then replace these basic objects with elaborate meshes just before performing the final rendering.

If you choose to replace an object with a plant, the Visual Plant Browser will appear, letting you select the plant species to be used in the replacement. If you replace an object using the **Load Object...** or **Import Object...** commands, a Standard File Browser will also appear letting you select the object to be imported.



Saving Objects

You can save objects for future use in other scenes by activating the alternate action of the **Load object** icon () , or by selecting the menu command **File | Save Object**. You will be prompted for a filename; the VUE object file type is **.vob**. VUE will then render a preview of your object. By default, the new object is added to the **Personal** collection. The VOB file format cannot be used with other 3D applications. If you would like to export VUE objects to use them in other 3D applications, please read the next topic.

Working with Pixologic ZBrush

VUE and Zbrush can work together with the GoZBrush system. This Link adds this export option to the File menu. It opens ZBrush with current selected object loaded and changes made to this object in Zbrush should update the VUE copy.

Note:

Note that TPF plants cannot be exported from VUE to Zbrush. GOZ support for TPF plants is only available in the PlantFactory 2015 and up.

Installing GoZBrush for Vue

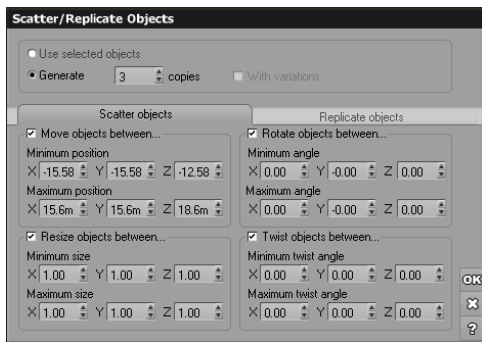
Once you've installed VUE 2015 and ZBrush (Zbrush should be installed first) on your computer, you'll have to run an install script. This is found inside **Zbrush**, in the menu **Preferences**, in the submenu **GoZ**, click on path to VUE. The script will try to automatically find VUE's installation folder on your computer. If it fails, indicate the location of VUE xStream/Infinite 2015.exe and select Install. If VUE is installed first, you'll have to install VUE again in order to use the GoZbrush link with VUE.



From VUE to ZBrush

To edit VUE's objects in Zbrush, select your object in VUE and go to the menu entry **File, Edit object in Pixologic ZBrush**. ZBrushs opens itself and allows you to edit your object there. Any press to the GoZ button in the ZBrush interface now updates the VUE's object consequently.

Scatter -- Replicate Objects



The Scatter/Replicate Objects dialog lets you create multiple copies of objects in one go.

The Scatter/Replicate Objects dialog lets you create multiple copies of the selected objects, while automatically moving, resizing, rotating and twisting the copies. This dialog is accessible through the **Edit | Scatter/Replicate Objects** menu command, or by long-clicking the **Duplicate** icon.

Check the **Generate copies** box if you want to generate copies of the objects that are currently selected. If you don't check this box, the selected objects will only be scattered. If **Generate copies** is checked, you can indicate the number of copies you request (e.g., if you have 3 spheres selected, asking for 4 copies will generate $4 \times 3 = 12$ spheres).

If at least one of the selected objects is either a plant, a terrain, a symmetrical terrain or a rock, the **With variations** checkbox will be enabled. If you check this option, VUE will generate variations of the selected objects (e.g. if you selected one Tropic plant and ask for 4 copies with variations, VUE will generate 4 new plants of the tropic species that are different from the original tropic). This is great for instance when you want to create a group of trees or rocks that are all different.



If you are trying various settings and wish to go back to the last setting you used, check the **Use last settings** box to undo the current settings.

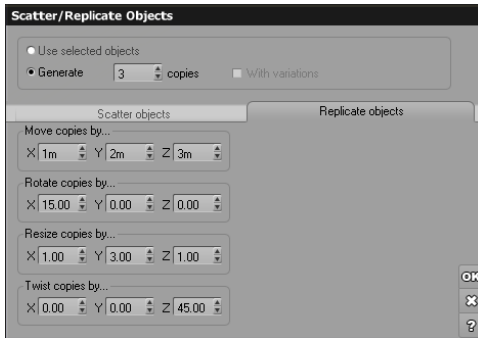
Aside from these controls, this dialog features two tabs, one for scattering objects randomly, and one for replicating them regularly.

Scatter Objects Tab

The objects and their copies can be moved, rotated, resized or twisted randomly depending on what boxes are checked.

For instance, if you want the copies to be moved about, check the **Move copies between...** box. The limits that are indicated are those of the selected objects. Copies of the objects will be moved inside these limits.

Replicate Objects Tab



Scatter/Replicate Objects dialog – Replicate tab

This tab lets you organize the copies regularly, by applying a constant translation, rotation, sizing or twisting between each copy. This is great for automatically building complex shapes from simple primitives.

Simply enter the different offsets before pressing **OK**.

If the **Generate copies** option is not selected, existing objects will be modified by the indicated values.

Note:



If you copy-paste, duplicate or Alt-copy objects, you can repeat that operation by using the **Repeat Operation Subdivide** and **Repeat Operation Extrapolate** commands in the **Edit** menu (shortcut * and /) to get a regular array of objects.



Camera Options and Framing

One of VUE's strengths is its camera. It works as a handheld camera would work, with aperture settings, focal lengths, and exposure to name just a few of its settings. When you create an image, you can “walk” around in your image, trying different camera angles, to come up with the perfect image.

You can also save a camera as a VUE object file (*.vob*) and reused in other scenes. Cameras cannot be saved in the same *.vob* file with other objects or lights. If a camera is linked to another object and you load it in a different scene, the link will be broken.

You can add cameras as you need them and switch between them for different views, rendering from each camera view as you wish. Each camera can have different settings. Adding and saving cameras can be done using the [Camera Manager](#).

Framing is a tool to aid you in finding the right camera angle and fine-tuning your image.

Framing



Framing strips due to a stretched picture format

As with any other object, the camera may be moved, rotated and resized directly inside the orthogonal *3D Views*. Although this is particularly welcome under certain circumstances, it is not necessarily the most intuitive way of framing a scene...

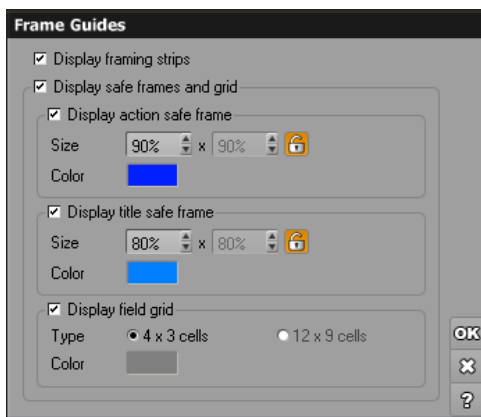


This is why framing can also be achieved directly inside the *Main camera view*. As you now know, orthogonal views can be moved around by clicking and dragging them with the right mouse button (Ctrl mouse on the Mac); you can also zoom into or out of them by pressing Control while you drag. In much the same way, the camera can be **rotated** up-down and right-left by dragging the main view with the right/Ctrl mouse button. Pressing Shift while you drag the view will cause the camera to **move** up-down and right-left, and pressing Control/Cmd while dragging will adjust the **focal** length of the camera lens.

If no objects are selected, and the active view is the *Main camera view*, pressing the arrow keys will nudge the camera **up**, **down**, **right** or **left**. Pressing the Page Up and Page Down keys will nudge the camera **forwards** and **backwards**. One nudge equals 5 units of distance. Shift-nudging nudges by 0.5 units.

You may have noticed two gray strips on the top and bottom, or right and left, of your main *3D View* (see above illustration for an example). They indicate the limits of the picture, according to the picture format you have selected in the *Render Options* or *Advanced Camera Options* dialogs. These stripes are here to help you get your framing just right. These stripes will vary in shape depending on the **Aspect ratio** you have set on your *Render Options* screen.

Frame Guides



Frame Guides dialog

The *Frame Guides* dialog lets you setup the visual guides that appear in the *Main camera view* to assist you in the framing of your scene.

To access this dialog, select the **Frame Guides...** command from the *Main camera view's View Display Options* menu.




Display Framing Strips

Because the aspect ratio of the final picture is not necessarily the same as that of the *Main camera view*, you can display in the camera view a visual clue of the size of the final picture, in the form of two semi-transparent gray rectangles that partially mask parts of the view that will not be featured in the final render. This clue is known as the framing strips. Deselect this option to hide these framing strips from the view.

Safe Frames

Safe frames appear as a set of rectangles in the main view, to help you keep the feature of animations in an appropriate position on screen. You can have up to two safe frames on screen (respectively known as “Action safe” and “Title safe” frames).

You can enable each one of these frames independently. For each one of these frames, you can define the size of the frame as well as the color used to display the frame in the *Main camera view*.

Enter the desired size in the **Size** fields. This is a percentage of the total width and height of the final picture. If the **Lock** icon () is selected, the vertical ratio is locked to the same value as the horizontal ratio.

You can change the color used to display the frame by double-clicking on the **Color** field.

Field Grids


Field grids appear as a set of vertical and horizontal lines in the main view that can be of help for framing. There are two types of field grids possible: 4×3 **cells** will split the screen into 4 cells horizontally and 3 cells vertically and 12×9 **cells** will further subdivide each one of these cells into 3×3 sub-cells (that appear in a paler shade of the grid color).

Double-click on the **Color** field to change the base color of the field grid.



Using the Camera

Managing Cameras




Camera properties panel

You can also use the [Camera Control Center](#) or the *Object Properties* panel (when the camera is selected) to set the position, orientation and framing of the camera.

If the selected object is the camera, the **Aspect** tab of the Object Properties panel displays as opposite.

Whenever you think you have come up with an interesting view of your scene, you can create a new camera based on the current camera by pressing the **Manage Cameras**

icon in the left icon bar (). This opens the [Camera Manager](#). Use this dialog to store the current settings, or replace/remove existing ones.

To add a camera, you can also just click the **Add Camera** option in the Object menu. That creates a camera in the World Browser with an identifying number.

A new camera object with the same settings as the perspective camera can be created as



well. This allows you to move a camera and when it is placed properly, you can create a new perspective camera object by selecting **Create from perspective camera** from the **Display** menu item.

Perspective views can be loaded into the current camera being used by using the icons on the **Title Bar** of the **Main Camera View**. The first icon after the **Quick Render** icon, **Store Camera Settings in Perspective View/Store Camera Settings in a new camera** is a dual action icon. A left click stores the current camera settings in the Perspective View and switches the Main Camera View to that view; a right-click stores the current camera settings in a new camera without switching to the other view.

The next icon is a **Toggle to Perspective View**. Now, if you've stored another camera setting in the Perspective View, that is what you will get when you click this icon.

Once you have created the camera, you can explore new framing options, knowing that you can instantly revert to the previously saved camera. By default, new scenes are created with 2 cameras. The first (**Ctrl + Num 0**) is the Main camera, and points North. The second (**Ctrl + Num 1**), called Top camera, looks down at the scene from above. You can create as many cameras as you like, but only the first ten are available through the **Display | Activate Camera** menu. Others should be activated using the Activate camera drop-down list of this panel.

- **Activate Camera:** this drop-down list shows all existing cameras. Select one from the list to activate the corresponding settings. You can also use the [Camera Control Center](#) to change the active camera. Cameras saved as *.vob* files are also included in this list and can be managed using the **Camera Control Center**.

Note:

You can change the current active cameras by unfolding the [camera group](#) in the *World Browser*, and selecting the new camera from the list of cameras, or by double-clicking on a camera in the *3D Views*.

Note:

If you [switch cameras in an animation](#) at a non zero time, a camera switch keyframe will be created in the *Timeline*.

- **Focal:** use this to adjust the focal length of the camera's lens numerically. The bigger the value, the greater the magnifying power of the lens. For landscape photography, values ranging from 24 to 35mm are often best suited. This [focal parameter](#) can be animated. If you are more familiar with camera Field of View rather than focal length, VUE can display the camera's horizontal FoV instead of its focal length in the *Object Properties* panel (disable the **Show camera FoV as Focal length** checkbox in the [Options](#) dialog).
- **Blur:** turn up this setting to activate depth of field. When depth of field is activated, only objects that are close to the Focus distance from the camera will appear sharp. Other objects will be out of focus and render blurred. The bigger the blur value,




the more rapidly objects get blurred as they move away from the focus distance. The limits of the “in focus” area are displayed in the camera’s field of view by two parallel planes. This [blur parameter](#) can be animated.

- **Focus:** when depth of field is active, this control lets you indicate the distance at which objects are in focus and rendered sharp. As objects gradually move away from the focus distance, they will get more and more blurry. The focus distance is displayed in the camera’s field of view by a cross. The [focus parameter](#) can be animated.
- **Exposure:** use this setting to vary the exposure of the scene. Positive values will make the scene brighter, while negative values will make it darker. The [exposure parameter](#) can be animated. The correction is expressed in diaphragms (a standard photographic unit of measure for the aperture of the lens). +1 diaphragm means the scene is twice as bright. Unlike a real camera, modifying the exposure has no influence on the depth of field. If the **Auto-exposure Film Settings** option in the *Camera Options* dialog has been enabled, this value indicates the correction of exposure that is to be applied to the exposure that was automatically computed for the scene. You can adjust exposure after the rendering completes in the *Post Render Options* dialog.

Note:

Varying the exposure is not the same as changing brightness in a post-processing pass using settings on the *Advanced Camera Options* dialog. Unlike the exposure setting found in the **Light** tab of *Atmosphere Editor*, this setting acts on the global exposure of the scene, and not only on the intensity of the lights.

- **Height:** use this to set the height of the camera. Clicking the **Lock** icon () will lock the height of the camera above ground. If the camera is dragged over a terrain, the camera will keep a fixed height above ground. Unlocking the camera will no longer keep the camera at the locked height. Whether the camera is either locked or unlocked, you can always move it manually in the *3D Views*. The **Height** setting will be automatically updated accordingly. Right-clicking on the **Height** icon, displays a menu with the options to ignore terrains, plants, and objects. You can select any, none or all to modify how the camera remains locked into position as it moves over the terrain. For example, if you select to ignore all three (object, plants, and terrains), the camera will be locked to the ground or any other infinite plane).

- 

Always keep level: if toggled, this option instructs VUE to make sure that the camera is always horizontal, resulting in a horizontal horizon. If (and only if) you deselect this option, you can add roll to the camera. This can energize pictures on occasion. However, for general purposes, we recommend that you leave this option on.





Lock camera attributes: This lock prevents you from selecting the camera in the *World Browser* and changing any of the camera's settings or *Object Properties*. All settings for this camera are then protected from accidental change including animation keyframes. This lock can be toggled off. The *Advanced Camera Options* window is still available if you need to make changes.




Backdrop: click this icon to load a backdrop image or animation into the background of the camera. Backdrop images or animations will appear behind all objects and will replace the sky. When you click this icon, the [Camera Backdrop Options](#) dialog will appear, letting you load the desired backdrop image or animation.



Switch to target: click this icon to switch selection to the camera's target – see below.

Camera Target




Properties of Camera Target

The camera target is a little box attached to the front of each camera and to which the camera is connected in view ports by a dotted line. This dotted line joining the camera



and its target always represents the direction the camera is pointing at. If you move the target, the camera orientation will be adjusted accordingly. This is a very useful and intuitive way to adjust camera orientation without having to tackle rotation angles. The length of the dotted line represents the focus distance of the camera. By moving the target away from the camera, you can adjust the focus distance graphically without having to enter numeric values.


Selection and Visibility

To select a camera's target, select the camera and click on its target object in the *3D Views* or click the **Switch to target** icon () in the *Aspect* tab of the *Object Properties* panel. Once selected, you can switch back to its owner camera by clicking the **Switch to camera** icon button in the attributes tab of the *Object Properties* panel.

You might have noticed that the target object becomes visible only when you select its camera owner. There is an option to make the target always visible, even when the camera isn't selected. To do this, select the camera target, and check the **Always visible** option in the *Aspect* tab of the *Object Properties* panel. When you deselect the camera, its target will remain visible.


Focusing on Objects

An interesting aspect of camera targets is that you can use them to **Focus on** any given object in your scene. Select the camera target and pick an object in the **Focus on** drop-down list box. From now on, the camera will always be focusing on this object.


You can also select the object in focus by using the **Pick object** icon () and then clicking on the desired object (or on an empty space to remove the connection). Focusing on an object does not affect the orientation of the camera, just its focus distance.





Camera Backdrop Options








Camera Backdrop Options

This little dialog lets you load a picture or an animation into the background of the camera, to be shown as backdrop where no other object is visible. This dialog is accessed by clicking the **Backdrop** icon () in the *Object Properties* panel, when the camera is selected.

To setup a backdrop for your camera, simply check the **Use backdrop** option. The other controls in the dialog become accessible:

Click the **Load** icon () to open the Picture Browser and load the image to be used as a backdrop. You can use any image, sequence of images or animation. You can even use different backdrops for different cameras. If you want to use a sequence of images, click the **Browse File** icon () in the Picture Browser to display a Standard File Browser and select all the images in the sequence.

If you load a sequence of images, or an animation, the **Animated backdrop options** icon () will appear under the picture preview. Click this icon to access the [backdrop animation settings](#).

If you need to rotate the picture, use the  and  buttons (90° increments). To invert the colors in the picture, click the **Invert** button (). Click on the **Remove** button () to delete the picture or animation.

Override atmosphere: If this option is unchecked, the backdrop will render, ignoring the



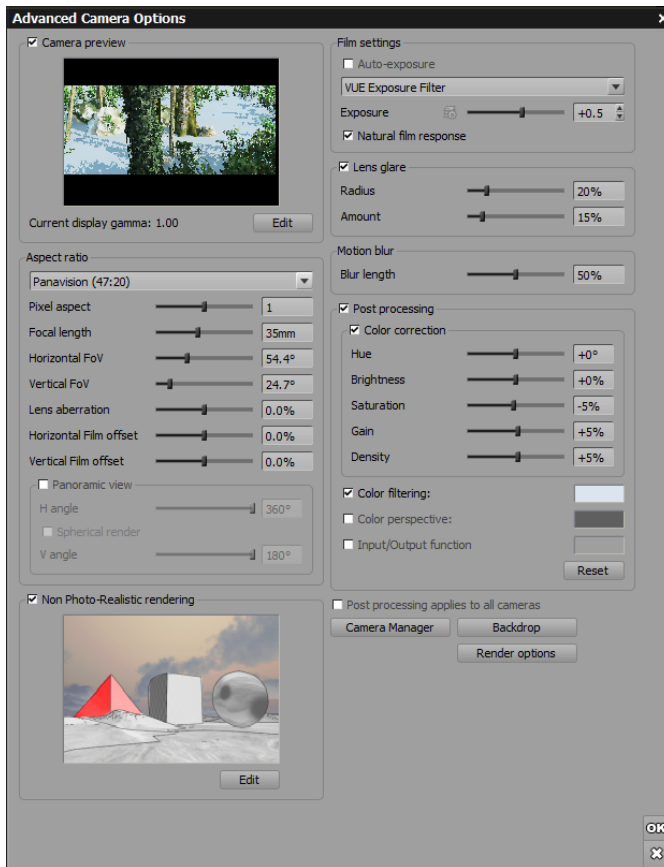
atmosphere completely. However, if checked, the atmosphere renders on top of the camera backdrop. VUE first renders the backdrop as an “outer space” value, and then applies atmospherics over it, making the backdrop less visible where atmosphere gets thicker (like objects fading out in the distance through the atmosphere, or like stars fading out beyond the atmosphere). As the atmosphere thins, the backdrop becomes more visible. Note that if you are in **Environment Mapping** atmosphere mode, unchecking the **Override atmosphere** option completely replaces the backdrop with the environment map, since the environment map has no transparency.

- **Zoom factor at render:** this setting lets you control the size of the image viewed in the background of the camera. The default setting of 1 will stretch the image appropriately so that it maps exactly to the background of the camera. Values smaller than 1 will result in the image not filling up the entire camera background. Values greater than 1 mean that the image will not be entirely visible in the camera background (the edges will be cropped).
- **OpenGL preview distance:** when you load a camera backdrop, it will appear in the OpenGL views. For your convenience, this control lets you adjust the distance from the camera at which the backdrop is displayed.
- **Animated OpenGL preview:** if you are using an image sequence or an animation as camera backdrop, this option lets you decide whether the OpenGL preview should reflect the animation when you change current time, or if it should always use the first frame (updating the animation can be a slow process, especially for large images or complex animation codecs).

Advanced Camera Options



Advanced Camera Options



Advanced Camera Options dialog

The *Advanced Camera Options* dialog lets you adjust the aspect ratio of your picture.

If you have several cameras in your scene, you can have a different aspect ration for each camera.

This dialog also lets you adjust the post processing effects.

Current Display Gamma: This field displays the current gamma setting. To change this setting, click on the **Edit** button to display the *Gamma Options* panel.

Changes made to the gamma settings here affect only the current image. To change the global gamma settings, access the **Gamma Options** button in the *Options* panel, [General](#)



Preferences tab.

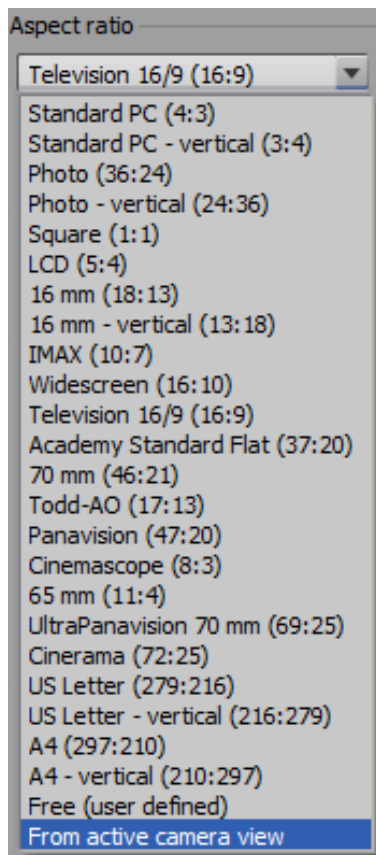
This dialog also lets you adjust the post processing effects.

Aspect Ratio

You can choose the format of your picture by picking a pre-defined **Aspect-ratio** from the drop-down list. Notice how the preview of the scene is updated when you change the picture's aspect ratio.

If you have several cameras in your scene, you can have a different aspect ratio for each camera, as long as the camera isn't used in animation for switching, as all cameras used in switching must have the same aspect ratio.


You can also adjust the picture's aspect ratio using the [Render Options](#) dialog.



Picture formats drop-down list




- **Pixel aspect:** this setting controls the aspect ratio of individual pixels.
The default value of 1 means that the pixels are square. For systems that display images stretched, it is necessary to use non-square pixel ratios in order to avoid any deformation in the final image. Pixel aspect ratios smaller than 1 will result in an image that is squashed vertically (for projection using devices that stretch images vertically when displaying them).
- **Focal length:** this setting controls the focal length of the camera. It is identical to the value displayed in the *Object Properties* panel (see [here](#)).
- **Horizontal FoV (Field of View):** this setting controls the horizontal viewing angle of the camera. It is connected to the focal length. Larger focal lengths result in smaller viewing angles. If you are more familiar with camera FoV rather than focal length, VUE can display the camera's horizontal FoV instead of its focal length in the *Object Properties* panel (disable the **Show camera FoV as Focal length** checkbox in the [Options](#) dialog. If you modify the value of the Horizontal FoV, the camera's focal length will be adjusted accordingly, as well as the Vertical FoV.
- **Vertical FoV (Field of View):** this setting controls the vertical viewing angle of the camera. It is connected to the focal length parameter, as well as the pixel aspect ratio (see above). Larger focal lengths result in smaller viewing angles. If you modify the value of the Vertical FoV, the camera's focal length will be adjusted accordingly, as well as the Horizontal FoV. If you modify the pixel aspect ratio, the Vertical FoV will be changed accordingly.
- **Lens aberrations:** this setting controls the percentage of lens aberration in the camera. In an ideal optical system, all rays of light from a point in the object plane would converge to the same point in the image plane, forming a clear image. The influences which cause different rays to converge to different points are called aberrations. Negative values translate into pin cushion distortion, whereas positive values translate into barrel distortions.




Barrel distortion (+80%)





no aberration (0%)



pin cushion distortion (-80%)

Lens aberrations

- **Horizontal and Vertical Film offset:** These settings controls the Film offset of the camera. Use the sliders to set the film offsets vertically and horizontally. Vertical file offset positive shifts the image down. Horizontal film offset positive shifts the image left. This movement can be animated.


Panoramic View

Real panoramic cameras are fixed to a handle that you have to hold when you take the shot. As it takes the shot, the camera rotates around the handle, thus embracing any required angle of view. Panoramic views can yield beautiful results.

This option is a numerical equivalent to the real panoramic camera. Here also, the camera is rotated as the render engine makes progress. All-around panoramas can be achieved this way, but beware: if the camera is not perfectly horizontal when you take the shot,



the horizon will undulate. This is not a numerical artifact: it happens also in the real world!



Undulating horizon due to an unlevel panoramic camera

When you select this option, the panoramic **H Angle** slider becomes available, letting you decide on the horizontal angle swept by the camera as it rotates.

The **Spherical render** option also becomes available when you activate Panoramic rendering. When you select this option, the scene will be rendered on a sphere (instead of being rendered on a cylinder). You can adjust the vertical angle swept by the camera using the **V Angle** setting. Spherical renders can be used, for example, for the rendering of an environment map. When this option is activated, the picture's aspect ratio is determined by the ratio of horizontal vs. vertical angles.

Because this option is only available when rendering to screen or to disk (see *Render Options* dialog – [Panoramic View](#)), activating panoramic rendering will automatically activate rendering to screen if it isn't already the case. You can also adjust the picture's panoramic settings using the *Render Options* dialog.

Non photo-Realistic rendering

VUE contains settings to create [Non-Photorealistic_Rendering](#) ([Documentation/Building_Scenes/Rendering](#)) with effects like “toon”. These settings are also available on the [Post Render Options](#) screen dialog to make adjustments to that particular render. The settings are also available from the top tool bar on the screen.

Film Settings

Photochemical films are made of tiny crystals of silver salt that react to light. When light reaches the surface of the film, it hits these crystals and triggers a chemical reaction that switches the state of the crystal (it becomes dark – this process is then inverted to result in a bright point). Once switched, a crystal will not be switched any further by more light hitting it (it can't be more black than black). It is the proportion of switched crystals that increases as light keeps on flooding in, making the point appear darker and darker. But, as more and more crystals have been switched by the incoming light, the chances of hitting an “unswitched” crystal go down. As a result, while points on the film will initially get dark very quickly, it will take more and more light to get them that much darker, resulting in a non-linear reaction to light. This non-linear reaction means that bright areas in the image will appear less bright, and dark areas less dark, resulting in a broader dynamic of light being visible in the final image.

- **Auto-exposure:** the difference in luminosity between noon and dusk is enormous,



but we are not necessarily aware of this fact, because the human eye automatically adjusts to the amount of ambient light. Auto-exposure simulates this behavior by automatically adapting the exposure of the camera to the amount of light in the scene. If this option is enabled, your images will be correctly exposed, even if you drag the sun from noon down to dusk. When the auto-exposure option is enabled, the camera re-evaluates its exposure continuously during the rendering process. This is why, when tile rendering mode is enabled, the overall exposure of the image may be adjusted as rendering progresses.

- **Exposure filters:** VUE has several filters you can use to change the effect of your images.
 - **Vue exposure filter:** this is the filter used in all previous VUE versions.
 - * **Exposure:** slider controls the general exposure for the filter.
 - * **Natural film response:** select this option to enable the non-linear reaction to light typical of photochemical films.
 - **Photographic exposure filter:** provides several settings to tweak the response curve.
 - * **Exposure:** slider controls the general exposure for the filter.
 - * **Shadows:** will control how dark areas are mapped (lower value will produce lighter shadowed areas, higher values will produce stronger shadows)
 - * **Mid-tones:** will control how intermediate intensities will be mapped (lower values will tend to flatten the response curve, effectively producing darker and less contrasted colors, higher values will make the curve steeper, increasing intermediate colors contrast and overall brightness)
 - * **Highlights:** control how bright intensities are mapped (lower values will flatten them to avoid saturation, while higher values will produce saturated bright areas).
 - **Linear Exposure filter:** this is the simplest of all of the filters.
 - * **Exposure:** exposure can be manually adjusted by this slider. If no contrast transformation is performed, the tone curve is flat, and the user can control how steep the line is. Either **Map entire range** is selected and the brightest rendered color will be mapped to full white, scaling all other colors accordingly or a custom linear coefficient can be specified to control the scaling of colors (1 meaning no scaling at all: raw colors are preserved).
 - **Reinhard tone filter:** there are two flavors of this filter.
 - * **A simplified one:** that gives only control over contrast enhancement
 - * **Reinhard 2 Exposure Filter:** that lets you control:
 - **Brightness:** an overall scaling factor for the output
 - **Chromatic adaptation:** color contrast



- **Light adaptation:** whether the tone operator acts rather locally (higher values) or globally (lower values) across the output picture.
both types of Reinhard filter allow you to manually adjust exposure.
- **False colors:** this can be used to visualize output HDR range of intensities of the rendered picture more easily.
- **Exposure:** This setting controls the overall exposure of the camera. Positive values will result in brighter images, whereas negative values will result in darker images. The value is expressed in diaphragms. If the auto-exposure option is enabled, this setting is relative to the automatic exposure value (think of it as a way of “touching up” the auto-exposure).

When using the **OpenGL shader engine** (set on the *Options* panel, **Display** tab), you can preview the results of the camera auto exposure factor and the scene exposure factor. The auto exposure is automatically updated at each refinements pass of the mini scene preview. The **Exposure** is only active in the camera view, and is altered by the camera exposure factor. This feature is not available for **OpenGL (fixed hardware pipeline)**.

Lens Glare

Lens glare is caused by imperfections in the lenses of real-world cameras. Instead of being perfectly refracted by the lenses of the camera, part of the light becomes diffused by little defects in the glass. This results in halos of light appearing around very bright points of the image.

Lens glare gives a soft, realistic look to the final images. The effect, sometimes also referred to as “specular bloom”, is particularly strong when the camera lenses are a little dirty (because light becomes diffused by the layer of dirt at the surface of the lens).

Lens glare is controlled via the following settings:

- **Radius:** this controls the average size of the halos of light that appear around bright points in the image.
- **Amount:** this controls the intensity of the glare effect.
- **Warning:** the larger the radius, the slower it becomes to compute the effect of the glare. When previewing glare on the full size image, a much faster approximation of the glare effect is used. This approximation can sometimes result in slight visual artifacts. These artifacts will disappear when the full-blown glare algorithms are applied to the image before saving.

Motion Blur Length

This simple control lets you adjust the amount of motion blur displayed in your render. Of course, the actual amount of motion blur depends on the speed of the moving objects, but this setting lets you adjust the overall amount of motion blur. It is somewhat similar to



the shutter speed (because the length of the motion blur depends on the distance traveled by the moving objects over the duration of the exposure).

Post Processing

Post processing is a special processing pass that takes place once the picture is completely rendered. Using this feature, you can adjust the colors and brightness of the final picture without having to use another specialized application. By post processing pictures inside VUE rather than using an external application, you ensure that the resulting colors retain all of their subtlety (when you save a picture, the colors in the picture are limited to 8 bits per pixel; artifacts and color banding can appear rapidly as soon as you affect anything but minor post-processing). In VUE, colors are computed and processed with a resolution that is literally several million times more refined than in an exported picture.

Post processing settings can also be animated for spectacular effects (see [here](#)).

To enable post processing of your picture, check the option. The post processing controls become available:

- **Color correction:** select this option to apply color correction to your picture. **Hue** shifts the color tones according to the angle indicated. **Brightness** will increase or reduce the overall brightness of the picture, while **Saturation** modifies the overall saturation of the picture. The **Gain** setting applies a smooth contrast to the picture. **Density** adds uniform density to all colors in the picture.
- **Color filtering:** this option lets you apply a color filter to the picture, as if it were seen through a colored gel. When you check this option, you can adjust the corresponding color by double-clicking on the color control.
- **Color perspective:** if you select this option, dark colors will be replaced by the indicated color. Black will be replaced with this exact color, while brighter colors will be blended according to the brightness of the color. When you check this option, you can adjust the corresponding color by double-clicking on the color control.
- **Input/Output Function:** this allows a very accurate tuning of contrasts and luminosity of each color channel (directly on the 32-bit float RGB components). Check the function to enable it. Right-click on the small function preview next to it and select **Edit Filter** in the pop-up menu. The *Color Input/Output Function Filter Editor* displays for changing the settings. See the [Editing Filters](#) section for more information about working with filters.

This editor has two tabs, **Profile** and **Channel**. The **Profile** tab is used to create key points. To create a new key point, you can either:

- double-click in the area where the curve is drawn. The new key point is created at the point you clicked. The curve is redrawn to use the new key point.
- click on the curve where you want the new key point; the coordinates of the clicked point appear in the **Position** boxes; you can edit them if required.



- press the **Add key point** button. The curve is redrawn.
- type the coordinates of the new key point in the **Position** boxes, then press the **Add key point** button. The curve is redrawn.

You can't create two key points at the same horizontal position.

The **Channel** tab contains the basic color channels. Right-click on the filter image to make changes. This way you can adjust each color channel separately.

If this function is enabled, it is recommended disabling **Natural film response** since it reduces the range of colors. This will give you better control over the final image.


Click the **Reset** button to reset all post-processing settings to zero.

The **Post processing applies to all cameras** option is checked by default, meaning that all cameras in the scene will use the same post-processing settings. If you deselect this option, you can assign different post processing settings to all your cameras. This can yield particularly interesting results when used in conjunction with the [Camera Switcher](#).

This parameter can be [animated](#).

Non-Photorealistic Rendering: check this option to use non-photographic rendering. This option must be checked here to be functional on the Post Render Options panel.



Camera Manager



Camera Manager panel



The *Camera Manager* is a very simple dialog that lets you name, select and delete cameras. To access this dialog, do one of the following:

- select the camera and click the **Manage cameras...** icon () from the *Aspect* tab of the *Object Properties* panel, or
- select the **Manage Cameras** menu command from the popup menu that appears when you long-click the **Store camera** icon () in the *Camera Control Center*.


When this dialog appears, it displays all of the cameras currently available in the scene:

- Double-click on a camera in the list to rename it.
- Click the **Add camera** icon to add a new camera. You can then rename it to whatever you wish.
- Click the **Load camera** icon to load a camera that has been previously saved as a .vob.
- Click the **Save camera** icon to save the highlighted camera as a .vob file.
- Click the **Delete camera** icon to delete the camera from the list and the scene.

You can also do a copy/paste of a camera in the *World Browser* to quickly create a new camera. You can switch between the two by double-clicking the camera name you want to use. These cameras will then also show up in the *Camera Manager* list as well.



Rendering

Once you are satisfied with the framing of your scene, click on the **Render** icon () on the top toolbar to start rendering the picture.

Rendering is an extremely complex (and time-consuming) process by which the computer converts the 3D geometrical description of the scene into a 2D picture you can look at.

With VUE, the Render settings have been slightly adjusted to increase the anti-aliasing threshold and to add a bit of texture filtering which will improve overall render quality. This is available from **Broadcast** quality on up. It may increase render times slightly, however.

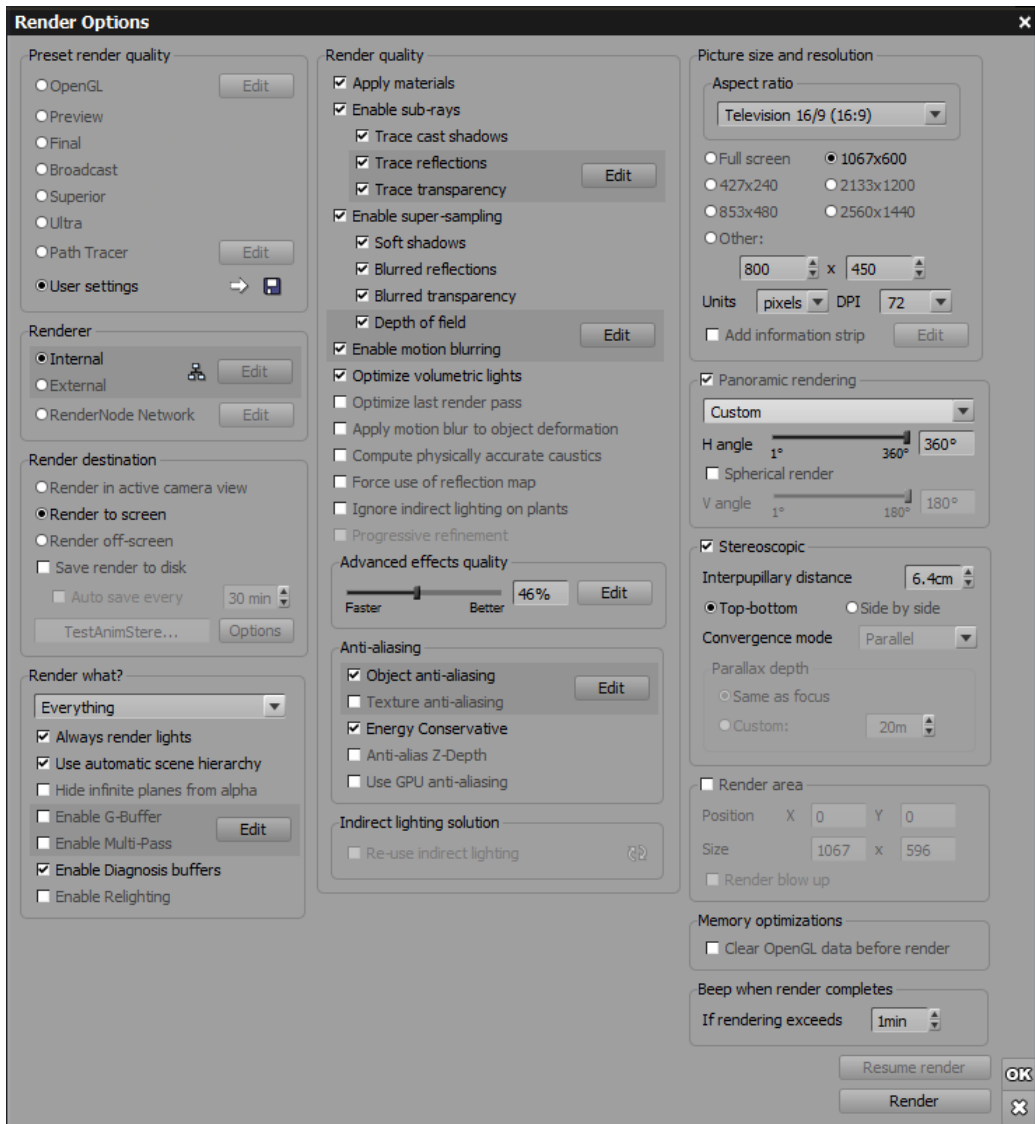
The Render icon is a [double action](#) icon. If you activate the icon's alternate action, the *Render Options* dialog will pop-up.

Bucket Rendering

Bucket rendering is a new way of organizing the rendering process that maximizes spatial correlation of scene geometry. The size of the buckets is automatically adapted to the overall render size (smaller buckets for small renders). This provides significant improvements when handling billion-polygon scenes, and results in a better optimization of memory resources as well as improved rendering speed.

Render Options





The Render Options dialog customizes the render engine

Because this dialog gives you full control over the render engine, it might look daunting at first sight. Don't worry though, you'll rapidly grow to understand the meaning of each setting. Besides, all the controls are not activated at the same time.



Preset Render Settings


On the top left corner of the dialog is a list of **Preset render settings**. These are predefined settings, useful because they let you quickly switch between whole groups of features. As you change render settings, notice how the render quality checkboxes are modified. These are disabled when you use a predefined setting. They become active if you select **User** settings, letting you pick your own options.

The available render settings are:

- **OpenGL:** this produces a very quick render of your scene using OpenGL, with no reflection, transparency nor cast shadows. It is useful when you want to quickly check positioning or motion of objects in your scene. In order to generate OpenGL renders as quickly as possible, EcoSystem instances that are sufficiently far from the camera are rendered as billboards instead of at full geometry (initial tests for rendering EcoSystems in OpenGL would take up to 2 hours per frame – much slower than the full-blown ray-tracing image!). You can adjust the quality of the OpenGL render by clicking the *Edit* button and changing the *Quality boost*. Due to issues with MacOS video drivers, mip-mapping is not available on the Mac platform for OpenGL rendering.
- **Preview:** this is the default setting whenever you create a new scene. It is a good working balance between picture quality and render speed. It traces reflections, transparency and cast shadows correctly, although it only mocks-up advanced features like soft shadows, blurred reflections / transparency and depth of field. The last render pass is optimized for speed, and the picture is not anti-aliased. We recommend you stick to this mode while you work on the picture, and only switch to Final settings when you have finished brushing up your scene.
- **Final:** as indicated by the name, this setting produces the final picture. It handles all features correctly, including advanced features such as soft shadows, and applies a reasonable quality anti-aliasing pass. Render times, however, are several times longer than in Preview setting. We recommend rendering pictures in this setting only when they are finished. If time is a critical aspect, you can use the *User* settings instead, and fine tune the render settings to achieve the best balance between quality and render time.
- **Broadcast:** this render setting was introduced for animation purposes. Basically, it adds motion blurring to the Final preset quality. However, it also features improved anti-aliasing quality, representing the optimal settings (in terms of render quality vs. render time) for rendering animations. Whenever the scene exhibits depth of field or motion blur, one single pass **Hybrid 2.5D technology** is done.
- **Superior:** this render setting is similar to Broadcast, with adjustments made to improve quality. Rendering with this setting is done in 5 Hybrid 2.5D technology



passes and is significantly slower than in Broadcast.

- **Ultra:** this is the best render quality available. It is also not very useful, since it takes several times longer to render than other settings, while not necessarily producing outstandingly better results. Use it only when you want to render very high quality pictures, at a not too high resolution. High DPI pictures for publishing usually render just as well in Final quality setting. Ultra setting adds superior anti-aliasing and improved advanced effects rendering.
- **Path Tracer:** a fast CPU+GPU render. Since it naturally simulates many effects that have to be specifically set with other methods, such as soft shadows, depth of field, motion blur, caustics, ambient occlusion, and indirect lighting, fewer settings are needed.
- **User Settings:** this not a preset. It grants you full access to customize the render engine as you like, by selecting the options you want from the render option checkboxes and anti-aliasing settings. These options will be detailed further down. The default User settings correspond to a faster version of Final render (with less super-sampling involved).
- **Load, Save:** When you select the User render quality setting, two small buttons () on the right hand side of this option become available. These buttons allow you to load or save your user settings. Pressing one of these buttons will display a Standard File Browser letting you load or save the selected file. The User Render Settings configuration files are stored in the *Environment* folder, and use the *.urs* extension. Files are supplied for all default preset render settings (Preview, Broadcast...). That way you can base your own user settings on an existing preset. You should avoid modifying any of these files.


Renderer

This option lets you select the renderer to be used for rendering the image or animation:

- **Internal:** select this option to use VUE's internal renderer. This is the best for quick renders that require interactive feedback (if you render to the active camera view or the screen, you will see the picture gradually appear on screen as it renders).
- **External:** when this option is selected, VUE will invoke an external rendering application that is installed together with VUE. This application will take care of the rendering. Because it is a separate application that is entirely dedicated to rendering, it doesn't have to deal with all the overhead of a graphical interface, and can consequently dedicate more memory to the actual rendering process. The caveat is that the scene has to be sent over to the external renderer so the time it takes to actually start rendering is longer than when using the internal renderer, and also, because rendering is done by a separate application, you do not see the



picture appear gradually on screen as it renders.

- **Batch rendering:** The batch rendering feature lets you schedule a number of rendering jobs that will be processed one after the other. Batch rendering occurs when you select the External renderer but do not select to use the network. You can also select this option from the menu using **Render | Batch render**. Batch rendering is handled through the *Batch Rendering* dialog.
- **Use network** (

To configure *HyperVue*, click the **Edit** button. Also, select this icon to make use of your network *RenderCows* when rendering to screen or in the *Main camera view*. When using this mode, rendering in VUE starts normally, without any delay. Then, while VUE is rendering, the *Interactive Network Rendering* controller silently contacts all the render nodes on your network and puts them to work on your scene. This happens in the background while your main computer continues rendering. As the other computers on your network start pitching in, the rendering accelerates. The *RenderCow* on the host machine is not used in this render as the program on the host machine is being used. You need to have your *RenderCows* added in *HyperVue* before using this render method.

When you configure your external renderer to use network rendering, the picture will automatically be divided up into sections. The *HyperVue™ Network Rendering Manager* will then assign each section to a *RenderCow™*. The network manager collects the resulting picture fragments and reassembles them into the final picture. When you press **OK** to begin rendering, the scene is added to the list of queued jobs.

- **RenderNode Network:** select this option to use a separate renderfarm administration tool and split the render load across a network of *RenderNodes*. Click the **Edit** button to access the *Rendernode Network Options*.

Note:

avoid using network rendering for quick renders, because the overhead of managing the render nodes and communicating over the network may actually result in longer render times. Use the internal renderer instead.

Please refer to [here](#) for details on the difference between *RenderCows* and *RenderNodes* and the *Rendernode Network Options* dialog.



Render Destination

This lets you decide whether you want the picture to be rendered inside the *Main camera view* (the picture resolution will be that of the *3D view*), if the picture should be rendered into a stand alone window, or if the picture should be rendered to disk.

- **Render in active camera view:** the picture will be rendered in the *Main camera view*, and the size of the picture will be that of the view.
- **Render to screen:** rendering will be done in a separate window that will appear when you start rendering (you will have to indicate the resolution of the picture). You can also save both the depth channel and the alpha channel renders when the render has finished. Previous renders are also displayed with the **Render to Screen** option. Renders are stacked; this means that they are saved and can be used for comparison or further editing (with **Post Render Options**). These are displayed on the [Render Display window](#) along with the current render.
- **Render off screen:** selecting this option instructs the render engine to save the picture as it renders, and not to display it. Saving pictures as they render is useful if you want to render pictures much larger than your screen. Selecting this option activates the **Options** button. Pressing it displays the *Render to Disk Options* dialog, letting you indicate which channels of information should be saved and the name of the file that they will be saved in. If the file already exists, VUE will ask for confirmation before starting the render.
- **Save render to disk:** This option is automatically checked if you select the **Render off-screen** option. But it can also be checked if you are rendering to screen or the active camera view. This automatically saves your render to disk as well as allows you to save after your screen render has finished. Press the **Options** button to indicate which channels of information should be saved and the name of the file and format to save it to.

Additionally, you have the option of an **Auto save** your render at a preset interval. You can set the interval in minutes for the auto save. You might want to do this so that you don't lose the entire render if your computer loses power (for example). You should be aware that auto saving does slow the render process down a bit, so it's probably not something you would want to do frequently.

If you then want to save it, you will have to select the **Save Displayed Picture** icon on the *Render Display* (farthest right icon under the rendered image). Using this method, you can also save the other channels of information in the picture.



Render What?

Click the drop box to display the following selections:

- **Everything:** will render everything in your scene. This is the default.
- **Only selected objects:** will only render the objects that were selected at the time you pressed render.
- **Only active layers:** will only render objects that are placed inside active layers.
- **Only visible layers:** will only render objects that are placed inside active or locked layers.

Selecting **Always render lights** will use all the lights defined in your scene, even if they are in layers that shouldn't be rendered. This guarantees the same lighting conditions for the rendered objects as that of the complete scene.

- **Use automatic scene hierarchy:** when this option is selected, VUE will optimize your scene so that it can be rendered faster. There are some cases, however, where optimizing the scene manually can yield even better results than the automatic optimization – and will save the preparation time. By unchecking this option, you will save the preparation time while retaining your optimizations.
- **Hide infinite planes from alpha:** this option will prevent infinite planes from appearing in the alpha channel of the picture. This is useful, for instance, if you want to treat the ground as part of the background.

G-Buffer / Multi-Pass Options

In Final or better preset render quality, select the **Enable G-Buffer / Multi-Pass** option to activate the collection of G-Buffer and Multi-Pass information.

Click the **Edit** button to open the *Multi-Pass Options (G-Buffer)* dialog and select the channels of information / rendering components / masks that you need.

If you enable G-Buffer rendering, you will have the option to save your render as a RLA or RPF File (these file formats preserve the G-Buffer channel information).

Note:

Generating the G-Buffer or Multi-Pass/Mask information increases memory requirements when rendering, and slows rendering down (especially if some objects have the **Render occluded** field set, see [here](#)). Also, if the **Force rendering of occluded objects** option is selected in the *G-Buffer / Multi-Pass Options* dialog, this will slow rendering down further, so you should only select this option when you actually require the extra information.



It is not possible to generate the G-Buffer or Multi-Pass/Mask information if the **Optimize last render pass** option is selected (see below). As a result, this button will be disabled in Preview and OpenGL render presets.

Enable Diagnosis Buffer

Select this option to activate the Diagnosis Render Passes information. This generates statistical information about the last render, giving access to normalized render time, anti-aliasing amount and global illumination samples placement for each rendered pixel. This can help you locate potential rendering bottlenecks across the image. For instance, parts that exhibit much longer render time compared with others might indicate the need for some optimizations on the local geometry and/or corresponding textures. Although not always possible, it is generally better to get a uniform rendering complexity over the image, especially when rendering on multi-core computers, so the rendering change may be ideally spread among all available processors.

The Diagnosis channels are available from the same icon on the bar above the Main camera view (or the render to screen window) that the other buffers are. These diagnosis channels display at the top of the menu.

These Diagnosis channels are coded in grayscale. A dark color corresponds to a low value (low render time, low anti-aliasing amount, or no global illumination sample), while a bright color corresponds to a high value (high render time, high anti-aliasing amount, or recorded global illumination sample). Pure white color indicates the maximum value across the image, thus all channels are normalized, giving relative information.

Isolated bright values may sporadically appear across the render time channel. This can correspond to internal precision errors from the involved high resolution timers used to evaluate each pixel render time, especially on multi-core computers. These isolated bright pixels should be ignored when considering render time issues.

The color-coded rule as well as the range of values can be hidden by using the little arrows that are usually used to navigate through the layers of additional channels. It may be useful in case relevant diagnosis information gets overlapped by the rule and values.



Enable Relighting

Relighting is a rendering feature that allows for the interactive modification of individual lights or groups of lights after render. It is possible to completely change lighting by modulating lights intensity or applying color filters to them without the need to re-render the scene.

To enable relighting, just check this option. This option is available for render quality Final and above. After render, the *Post Render Options* dialog will display with individual light or groups of lights controls, allowing you to tweak their intensity and to apply color filters to them.

By default, the relighting result will be displayed interactively in the little render preview of the dialog, but you can also enable the **Full interactive display** option on the *Post Render Options* dialog and see the full render display be updated during your changes. This option is also compatible with all other post render options like natural film response, automatic and manual camera exposure, and camera post processing effects. Furthermore, automatic exposure will be automatically adjusted according to your relighting settings, to consistently compensate for potential brightening or darkening of the render.

Relighting capabilities are compatible with almost all lighting features of the renderer: direct lighting, global illumination, atmospherics, physical caustics, volumetric lights, volumetric materials, subsurface scattering, reflections, refractions, lens flares and automatic exposure. The Ambient lighting contribution is also calculated and can be adjusted separately in the *Post Render Options* dialog. The only components that aren't affected by relighting are glow effects.

Just like the GBuffer and Multipass features, relighting isn't compatible with the **Optimize last render pass** render option, so this option must be disabled to be able to use relighting.

Relighting is also incompatible with Hybrid 2.5D effects; only distributed raytracing can be used.

For large renders with a lot of lights, or if lens glare is enabled, full interactive display can become too slow to be practical. In that case, it is recommended to just disable that feature and work with the little render preview in the *Post Render Options* dialog.

Any light that is switched off before render (using the **Hide from render** option) won't be accessible for relighting. Thus, all lights you plan to use for relighting must be enabled before rendering, and with a non black color so VUE can extract the proper lighting information during render.

VUE automatically defines relighting groups of lights depending on the first level of the



scene hierarchy. All lights having a common parent group object will be gathered into the same relighting entry. Therefore, group your lights according to your needs for relighting, to avoid the need of tweaking each light independently, which can be a long task if the scene contains many lights. Consequently, any light that isn't part of any group will correspond to a separate relighting entry.

Render Quality

The checkboxes in this group let you decide in detail what the render engine should trace, or not trace:

- **Apply materials:** deselecting this will replace object materials with uniform colors.
- **Enable sub-rays:** deselecting this disables tracing of any secondary rays (reflection, transparency or cast shadows).
- **Trace cast shadows:** selecting this is necessary, together with Enable sub-rays if you want cast shadows in your picture.
- **Trace reflections:** selecting this is necessary, together with Enable sub-rays, if you want reflections in your picture.
- **Trace transparency:** selecting this is necessary, together with Enable sub-rays, if you want transparency in your picture. Alongside the Trace reflections and Trace transparency options is a button labeled **Edit**. This button (only enabled when in “User” preset render quality) opens the [Sub-Ray Options](#) dialog to let you customize the ray recursion levels.
- **Enable super-sampling:** deselecting this disables any advanced features, such as soft shadows, blurred transparencies and reflections, or depth of field.
- **Soft shadows:** selecting this, together with Enable sub-rays, Trace cast shadows and Enable super-sampling will render soft shadows where applicable.
- **Blurred reflections:** selecting this, together with Enable sub-rays, Trace reflections and Enable super-sampling will render blurred reflections where applicable.
- **Blurred transparencies:** selecting this, together with Enable sub-rays, Trace transparency and Enable super-sampling will render blurred transparencies where applicable.
- **Depth of field:** selecting this, together with Enable super-sampling will render depth of field if some has been given to the camera.
- **Enable motion blurring:** selecting this option turns motion blurring on.

Note:

Memory requirements to render the scene increase considerably.



Alongside the Depth of field and Motion blurring options is a button labeled **Edit**. This button (only enabled when in User preset render quality) opens the [Blur Rendering Options](#) dialog to let you customize the blur rendering process.

- **Optimize volumetric lights:** when this option is selected, volumetric lights are rendered much more rapidly without any noticeable drop in picture quality. Except in very specific cases when volumes of light appear blurry, it is recommended that you always leave this option set.
- **Optimize last render pass:** selecting this will optimize the last render pass, making render up to 3 times faster, but leaving out minuscule details on occasion. You cannot generate G-Buffer information in this mode.
- **Apply motion blur to object deformation:** check this option if you would like the deformation of objects to appear with motion blur. For instance, when a plant is deformed by wind, the movement of the plant won't be blurred unless you select this option. While activating this option increases the quality of animations, be advised that it will dramatically increase the memory requirements when rendering objects with deformation. This setting is only applicable to objects that can be deformed (plants, terrains and animated meshes), and is of course only available when motion blurring is enabled. It increases significantly the amount of RAM required to render the scene.
- **Compute physically accurate caustics:** when this option is selected, [realistic caustics](#) are computed for the scene, including spectrum dispersion.

Note:

Computing realistic caustics adds a significant overhead to the rendering of the scene. If you don't select this option, caustics will be approximated using the much faster Fake caustics. Fake caustics, however, do not capture the effects caused by light bouncing off of reflective surfaces.

- **Force use of reflection map:** select this option to force all reflective materials to use [reflection maps](#), regardless of their respective reflection map settings. If no reflection map was used for a material, it will use the default reflection map.

Note:

This option doesn't modify the original materials; material settings will still indicate the use of true ray-traced reflections. It is only at the time of rendering that reflection maps will be used instead.

- **Ignore indirect lighting on plants:** because of the inherent complexity of plant geometry, rendering of indirect lighting on plants is extremely time consuming while not necessarily producing noticeably better results. This option lets you disable the computation of indirect lighting on plants altogether (however, the plant will still participate in the indirect lighting solution, e.g. by casting a dark shadow area beneath it).



- **Progressive refinement** (was called *tile rendering* before): This option is available in the User mode when Object anti-aliasing is disabled. Select this option to force the rendering engine to render the picture in tiles that get smaller and smaller as the render progresses. When this option is not selected, the rendering is done by the bucket renderer. This is the most efficient method of rendering.

The advantage of Progressive refinement is that you get a better overall impression of the picture right from the early stages of rendering. However, Progressive refinement may render the picture slower and requires large amounts of memory.

Advanced Effects Quality

- **Advanced effects quality:** this setting controls the overall quality of all the advanced rendering effects in the scene (e.g. volumetric lights, global illumination, procedural terrain, soft shadows, displacement mapping, etc.). The * **Edit:** button is accessible for all preset render modes, in order to get access to the **Optimize indirect lighting on plants** option. Of course, when not in User mode, only this option will be available, all the others (custom GI & photon map settings) will be grayed out. If you find that all the advanced rendering effects are rendered with artifacts (noise, splotches...), you can reduce these artifacts either by boosting the quality of each effect independently, or by increasing the quality globally using this slider.

Click the **Edit** button to open the [Advanced Effects Options](#) and gain advanced control over the rendering process.

Anti-Aliasing

[Anti-aliasing options](#) are automatically adjusted in the various preset render settings (see above). In the User render setting, however, you can control anti-aliasing options manually.

In addition to the standard Object anti-aliasing, VUE offers the possibility to use Texture anti-aliasing, for both bitmap and procedural textures.

- **Object anti-aliasing:** takes care of anti-aliasing in the geometry.
- **Texture anti-aliasing:** takes care of anti-aliasing in the textures.

The object and texture anti-aliasing are adjusted using the [Anti-Aliasing Options](#) dialog. This dialog is accessed by clicking the **Edit** button.

- **Energy Conservative:** This preserves bright details while limiting the amount of anti-aliasing needed to get a smooth rendered picture. Enabling this option is recommended whenever the rendered picture contains very small bright details like



narrow specular highlights or sun reflections over a perturbed water surface, especially in photometric lighting mode, which produces very high intensity variations between shadowed and lit areas.

- **Anti-alias Z-Depth:** Check this option to enable this form of anti-aliasing. Object anti-aliasing also needs to be enabled to access this feature, since depth anti-aliasing is performed in the same way as for color and alpha channels for consistency.

Note:

Depth anti-aliasing can produce undesirable effects depending on how z-depth information is used. When this option is disabled, VUE records the closest hit element distance within each rendered pixel. When enabled, all hit distances are averaged. This can lead to a resulting distance that doesn't correspond to any actually hit geometry. For example, when averaging a distance to a background mountain with a distance to a foreground object, resulting distance is between them effectively corresponding to none of them. Despite this incorrect result, depth anti-aliasing can be useful when using z-depth information in third-party post effect plugins (for depth of field simulation, for example), to avoid aliasing artifacts around objects.

- **Use GPU anti-aliasing:** Since this can, in some instances, increase render times, it is now an option you can change on a per-render basis. GPU anti-aliasing is not compatible with Standalone rendering since this type of rendering doesn't directly interact with any OpenGL processing. This is a technology that only work on meshes, and it allows to have a better anti aliasing on meshes with thin details such as cables or wires. Using this feature may slow down rendering in certain situations where, for example, you have a large mesh hidden behind another mesh. Both meshes would be sent to the graphics card. Regular raytracing would stop the rays at the mesh in front.

Background draw also has to be enabled for this feature to work.


Indirect Lighting Solution

When the **Re-use indirect lighting** box is selected, the radiosity calculation will not be performed again at the time of rendering. Instead, the last calculation will be re-used, and any subsequent lighting information gathered from further renderings will be added, appending new indirect lighting data to it at each new render, whenever needed.

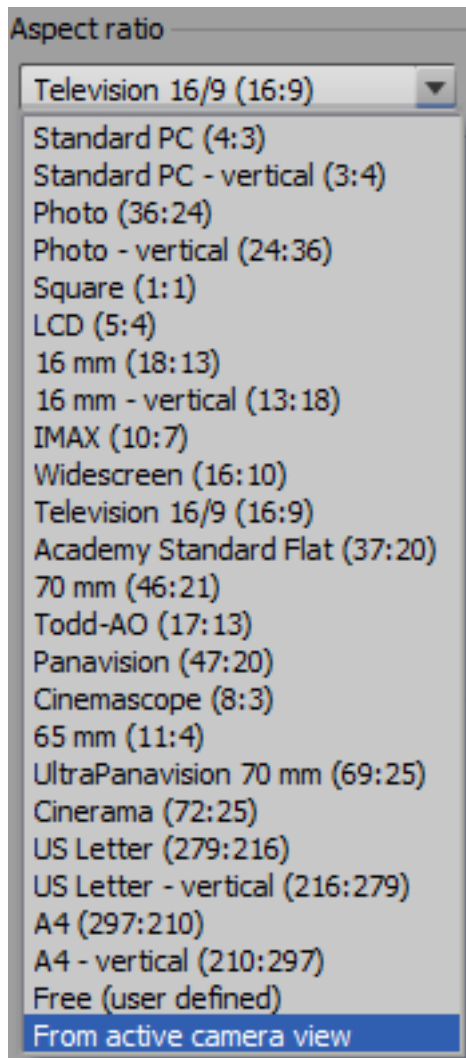
This great improvement is especially useful for walk-through animations, where the camera progressively discovers new parts of a scene while moving through it. Indeed, illumination caching will just compute any missing information at each frame, while reusing previous calculations wherever possible, significantly reducing render times while also reducing flickering artifacts.

Obviously, if the lighting conditions have changed, or if significant changes have been



made to the scene, the radiosity calculation may no longer be accurate. To update the calculation, press the **Update Indirect Lighting Next Time** icon () , or select the menu command **Render | Update Indirect Lighting Next Time**. This will update the indirect lighting the next time you render so that it matches any changes made to the scene.

Picture Size and Resolution



Picture formats drop-down list

You can choose the format of your picture by picking a pre-defined **Aspect-ratio** from the drop-down list. If no pre-defined aspect-ratio is suitable, select **Free (user defined)** and then type the size of your picture in the two **Other** boxes. You can also adjust the picture's aspect ratio using the [Advanced Camera Options](#) dialog.

A set of 6 boxes below the aspect-ratio list lets you select standard picture resolutions.

Alternately, you can enter any other resolution using the two boxes in the group called **Other**. If you have selected a pre-defined aspect-ratio, the vertical and horizontal resolutions of your picture will be linked together. These boxes are only available if you are rendering the picture to screen, since, if you render inside the main *3D View*, the resolution of the picture is defined by that of the view. Selecting **Full screen** guarantees the biggest possible resolution that fits inside your current display.

Once you change the aspect ratio, you may notice two gray stripes in the active camera view. They are here to show you the limits of the picture in the selected format, and help you optimize framing.

The **Units** drop-down list lets you select the units you want to work with. If you are planning to print the picture, you probably will want to switch to inches or centimeters. In this case, the **DPI** (Dots Per Inch) box becomes active, and you can enter the required DPI setting for the generated picture.

DPI (Dots Per Inch) indicate the number of pixels per inch in the picture once it is printed on paper (i.e. the definition of the picture). What you need to understand is that the only way to increase the number of DPI of a picture without reducing its size on paper, is to increase its resolution. It is of general acceptance that, for professional work, 300 DPI is a good compromise between definition of the printed picture, and size of the render. 240 DPI is probably enough for standard use. Select the required number of DPI for your picture (usual values are in the drop-down list but you can enter any value). The default is 72 DPI, which is the definition for screen viewing.

Note:


Changing DPI without changing the resolution of the image has no effect on the screen display size or quality of the finished image. The DPI setting is only made available as a convenience for those wishing to print.



Locking User Defined Aspect Ratio

If you select **Free (user defined)** as an aspect ratio, and enter a picture size in the **Other (size)** field, you have the option of locking this ratio. If you later change the width, the height will change automatically preserving that aspect ratio.

Add Information Strip



Render Information Strip dialog

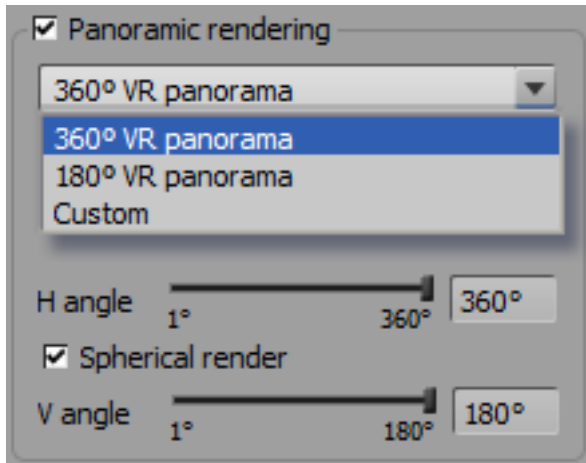
This is an information strip that displays at the bottom of an image rendered to screen or in any of the viewports. It can display a combination of the scene name, frame number, render time, image resolution and poly count. Select which items to display by clicking on the **Edit** button on the right.

When rendering to screen, the information strip can be toggled on or off.

If you wish to save this information, you have two options when saving the render. On the **Save As** dialog, you have the option of adding this strip to the image itself, or you can save it into a log file. This log file can be found in your render destination directory with the name *imagename.log*.



Panoramic View



Panoramic render settings

This option is available only when **Render to screen** or **Render to disk** is selected.

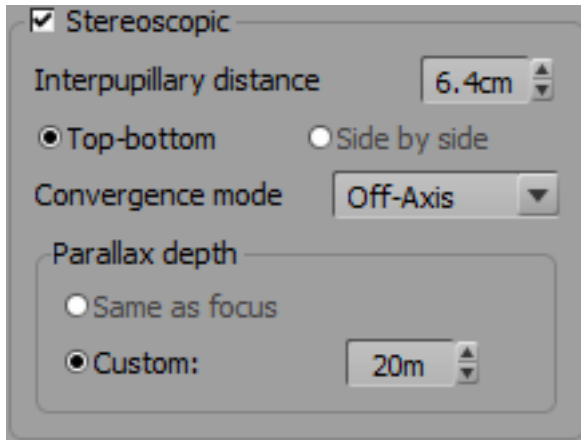
Two presets called 180 and 360 degrees Virtual Reality panorama are available. Instead, custom preset can be used to tweak settings individually.

Panorama static images exported as JPEG are now automatically recognized by Facebook as Panorama. When rendering a 360° animation and exporting it as .mp4 or .mov, VUE will automatically add a piece of metadata in the video file to make video-reading software (e.g. youtube) recognize the video as a 360° video, and display it as such.

When setting a camera to the **Panoramic View** mode and the camera is not perfectly horizontal, VUE will ask whether the user wants to level the camera. This is because panoramic renders are highly distorted when the camera is not horizontal and this is not a common use case.



Stereoscopic rendering




Stereoscopic render settings

Stereoscopic rendering allows the production of two images simultaneously, one as seen through the viewer's left eye and one seen through the right eye. The distance between the eyes (called *Interpupillary Distance*, is set to a value of 6.4 cm by default) can be reset as required by the individual user. The layout of where each of the renders appear in the final image can also be set: either one on top of the other (Top-Bottom) or one adjacent to the other (Side by side). If you export a 360° panoramic stereo animation as mp4 or mov, the metadata embedded in the file should also help a player, such as YouTube, to interpret it as stereoscopic video. YouTube needs the Top-Bottom layout.



Convergence modes



Stereoscopic convergence modes

Note:

In the case of Panoramic rendering, only the *Parallel* convergence mode is available.

There are three convergence modes available and the “Stereoscopic convergence modes” image helps to understand the difference between the three. These modes mainly differ on how the Parallax Effect is handled. Parallax effect helps the human brain understand how far objects are away from the viewer. Objects located at parallax depth will end up in the same position on both eye image. Objects in front will usually be offset on the left on the right eye image (and vice-versa) and thus be interpreted as closer by the brain and would usually give the feeling of being in front of the displaying device. In VUE, parallax depth can be set either to the same distance as the focus plane (this is the default behavior because our vision usually focuses and converges at the point of interest) or any value.

- **Parallel** convergence mode simply places the two cameras offset by the interpupillary distance but aligned. This mode behaves the same as the two others with



infinite Parallax Depth, this means that every object is subject to the 3D effect to a greater or lesser degree.

- **Convergered** mode rotates the camera toward the convergence point. In this mode, the zero parallax plane does not exist, instead there is a zero parallax vertical line where the two parallax depth planes intersect. Essentially this equates to the focal distance/point.
- **Off-axis** off axis convergence mode uses a film offset to keep a zero parallax plane and to make sure there is no horizontal offset (in the objects in the frame) at the zero parallax depth. Since this mode is the most natural one, this is the default behavior.

Render Area

This option lets you select a rectangular area in the picture outside of which the picture won't be rendered. You can also select a render area using the main menu command **Render | Select Render Area** and then drawing the render area with the mouse.

When you select this option, the render area controls become available:

- **Position:** these two settings let you define the top-left corner of the area to be rendered (in pixels).
- **Size:** these two settings indicate the width and height of the render area (in pixels).
- **Render blow-up:** this option is only available when rendering to screen or to disk. When it is selected, the render area will be rendered at the size of the picture indicated in the **Picture size and resolution** group. If it is not selected, the picture will be rendered at the exact size indicated. This option is useful when you want to render a close-up of a detail in your picture.

This render area can be locked by using the **Lock Render Area** option on the contextual menu. This will protect the selected area from mouse movements, allowing you to keep the setting while still working on the scene.

Rendering the Selected Render Area

When rendering, either in the viewports or to screen, you have the option of rendering just a selected area of the image. This is useful if you are working on a small area and you want to concentrate on just that area. Using the **Render blow-up** setting on the *Render Options* panel, you can enlarge this selected render area when rendering to screen. The Select render area does not work for animations and does not work with rendernodes. This feature is for stills only.



You can also opt to save a render of the full size image, with just the selected area rendered. For example, you have selected to render the lower right quarter of the image. After the render finishes, click to save, and on the **Save As** screen, there is an option to **Save full size image**. The image that is saved will be the rendered lower quarter of the image; the rest of the image will be black. This option is useful if you are planning to render a picture in sections, then overlay for the finished image.

Memory Optimization

- **Clear OpenGL before rendering:** when this option is selected, VUE will clear all OpenGL data and buffers in order to free up as much memory as possible for rendering. Depending on the complexity of the scene, this may free up a significant amount of memory for rendering. The drawback of enabling this option is that VUE will have to regenerate all OpenGL data after rendering, which can delay the refreshing of the views.

Closing the Dialog

Click **OK** to accept the changes and close the dialog. Click **Cancel** to cancel the changes.

To accept the changes and render the picture with the new settings, click the **Render** button.

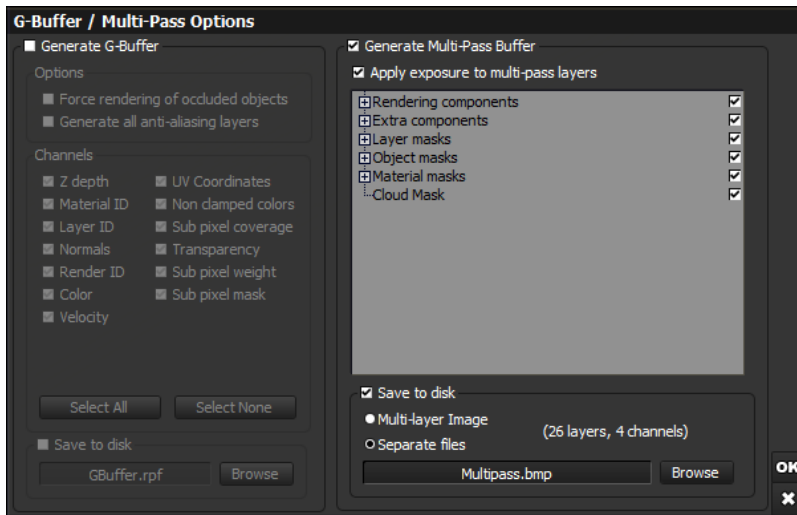
If you have interrupted a render in progress, the **Resume render** button will be active. Click on this button to resume rendering the picture.

Note:

Any changes to the render quality will make resuming a render impossible.



G-Buffer Multi-Pass Options



G-Buffer and Multi-Pass Options dialog

To access this dialog, open the [Render Options](#) dialog and press the **Edit** button alongside the **G-Buffer** and **Multi-Pass** options or go to the [Animation Render Options](#) dialog and press the **G-Buffer / Multi-Pass Options** button.

This dialog is separated in two frames. The first frame controls the G-Buffer rendering options, while the second deals with Multi-Pass rendering.

G-Buffer

Check the **Generate G-Buffer** option to enable the creation of G-Buffer information. When this option is selected, the controls in the G-Buffer frame become active.

There are two groups of controls in this frame, Rendering and Channels.

Rendering

You can activate two G-Buffer rendering features:

Force render occluded objects: if checked, every region occluded by objects will be



rendered in the G-Buffer. This allows for the possibility to remove objects from rendering during a post-processing phase or, for instance, to perform accurate motion blur effects without any missing information issues.

Note:

you can activate rendering of occluded regions on a per object basis using the *Object Properties* panel, *Aspect* tab. This is recommended, as rendering occluded regions slows rendering down significantly.

Generate all anti-aliasing layers: if checked, anti-aliasing information will be segregated from rendering information and placed onto separate layers. This can be useful in addition to the previous feature for extremely accurate object removals in the post-processing phase.

Channels

The G-Buffer is organized in a potentially unlimited number of layers. Each layer contains a number of channels of information.

All G-Buffer channels are supported in VUE. Here is the list of the different channels available:

- **Z Depth:** a floating point value representing the distance to the fragment.
- **Material ID:** an integer value that uniquely identifies the material assigned to the object hit in this fragment.
- **Layer ID:** an integer value that identifies the VUE layer that the object belongs to (see the description of [Layers](#) in the *World Browser* for details).
- **Normals:** an integer value representing the compressed normal vector to the surface of the fragment. This vector is stored in camera view space.
- **Render ID:** an integer value that uniquely identifies the object hit in this fragment.
- **Color:** 3 bytes representing the RGB color values of the fragment (after it is clamped to the visible spectrum).
- **Velocity:** two floating point values representing the velocity vector of the fragment in screen coordinates.
- **UV coordinates:** a pair of floating point values in the range of 0 through 1 representing the U and V coordinates of the textures mapped on the fragment.
- **Non clamped colors:** 4 bytes representing a Ward's Shared Exponent Format encoded color. This is the color that was actually rendered before it was clamped to the visible spectrum.
- **Sub-pixel coverage:** a byte representing the percentage of the pixel covered by



this fragment (255 meaning 100% coverage).

- **Transparency:** 3 bytes representing the RGB color values of the filtering applied to all fragments behind this fragment.
- **Sub-pixel weight:** 3 bytes representing the actual contribution of this fragment to the final pixel color (it takes transparency of all preceding fragments and this fragment's coverage into account). The final pixel color is the sum of all fragment colors multiplied by their respective sub-pixel weights.
- **Sub-pixel mask:** a 16 bit integer representing a 4×4 grid that indicates the portions of the pixel that are covered by the fragment.

You can specify which channels should be generated in the G-Buffer. If you don't want to generate them all, simply check the ones that are to be generated. Click the **Save to Disk** box to save the GBuffer.

Saving Pictures as RLA or RPF Files

If you have already rendered a picture (with G-Buffer information enabled) and would like to save the contents of the G-Buffer in a RPF multi-channel file, use any of the methods listed in the *Exporting Pictures* section below. You can also save the information using the RLA file format, but not all channels of information will be available.

If you want to save the G-Buffer information to file at the time of rendering, choose **Render to disk** in the **Render destination** field of the *Render Options* dialog, and click the **Options** button. Next to color picture name, click the **Browse** button and choose the **Run-Length Encoded (*.rla)** or **Rich Pixel Format (*.rpf)** picture formats. G-Buffer generation will automatically be checked for you if it wasn't already done, so all you have to do is edit the G-Buffer Options if needed. Then launch rendering and the result will be saved in the RLA or RPF file that you specified.

Note:

you won't be able to save in RLA or RPF file format if you didn't generate G-Buffer information during the last render. If so, then you will have to re-render your scene after activating G-Buffer information generation. You cannot save in RLA or RPF file format a picture rendered with the **Optimize last render pass** option set because it isn't possible to generate G-Buffer information in this case.

Saving Animations as RLA or RPF Files

In order to save an animation as a series of RLA or RPF files (one for each frame), choose the **Run-Length Encoded (*.rla)** or **Rich Pixel Format (*.rpf)** animation formats in the Advanced Animation Options (click **Browse** button of color channel) and launch the rendering of the animation.



Note:

the limitation on optimizing the last render pass (detailed above) also applies to the generation of G-Buffer information for animations.

Multi-Pass

Check the **Generate Multi-Pass Buffer** option to enable the creation of the Multi-Pass information. When this option is selected, the controls in the Multi-Pass frame become active. Simply add a check along each one of the layers/masks you would like VUE to generate. If you select a category, all the layers/masks of this category will be generated. To rename individual render passes click on the object pass that you want to rename and enter in the new name.

Note:

The more layers/masks you generate, the more system resources will be necessary to perform the rendering.

The **Apply exposure to multi-pass layers** option allows you to apply the render exposure (which combines the camera exposure and the automatic exposure) to relevant multi-pass layers, that is to the render components. This is allowed because exposure is a linear operation, thus still allowing a proper compositing of individually exposed layers without biasing the final result. When enabled, render exposure will be dynamically applied to multi-pass render components upon display and saving to disk. This is useful with HDR pictures, especially in photometric lighting mode.

Rendering Components

VUE's multi-pass rendering engine is capable of separating the following render information:

- **Diffuse:** this component contains the colors caused by light diffused by the surface of the object in all directions; it is saved as a normal (i.e. additive) layer in Photoshop documents.
- **Specular:** this component contains the light reflected by the surface of objects; it is saved as a normal (i.e. additive) layer in Photoshop documents.
- **Shadows:** this component contains the shadows cast by objects; it is saved as a product layer in Photoshop documents. This corresponds to shadow applied onto geometry visible in the Diffuse pass, which doesn't include clouds. If cloud shadows were included in this layer, a proper reconstruction of the full render wouldn't be possible, because multiplying diffuse pass by shadow pass would add wrong shadows onto geometry visible through those clouds.



- **Ambient:** this component contains the light created by the ambient lighting term; it is saved as a normal (i.e. additive) layer in Photoshop documents.
- **Refractions:** this component contains the colors refracted through objects; it is saved as a normal (i.e. additive) layer in Photoshop documents.
- **Reflections:** this component contains the colors reflected by objects; it is saved as a normal (i.e. additive) layer in Photoshop documents.
- **Background:** this component contains the background colors; it is saved as a normal (i.e. additive) layer in Photoshop documents.
- **Indirect lighting:** this component contains the lighting caused by other objects in the scene (when rendering with Global Radiosity); it is saved as a normal (i.e. additive) layer in Photoshop documents.
- **Atmosphere filter:** this component, together with the Atmosphere gain component contains the effects of the atmosphere on the rendering. It is necessary to hold this information on two layers, because of a limitation in the Photoshop layer handling (no true additive mode); this component is saved as a normal (i.e. additive) layer in Photoshop documents.
- **Atmosphere gain:** this is the second half of the atmosphere effect; it is saved as a product layer in Photoshop documents.
- **Post process:** this component contains the colors added in post-process (e.g. lens flares, glow) ; it is saved as a normal (i.e. additive) layer in Photoshop documents.

Additionally, if you render an animation with the **Show timecode** on frames option enabled, a **Timecode** layer will be added at the top of the list of layers in the Photoshop document (normal layer).

Extra Components

On top of the above rendering components, VUE can also produce the following additional rendering information (not part of the actual picture per se, but potentially useful when post-processing) grouped in the **Extra components** category:

- **Z Depth:** indicates the distance to the object at this point,
- **XY Normal:** indicates the direction of the normal vector to the surface of the object at this point, the X component of the vector being stored in the Red byte and the Y component being stored in the Green byte,
- **XYZ Normal:** indicates the direction of the normal vector to the surface of the object at this point in world coordinates, the X component of the vector being stored in the Red byte, the Y component being stored in the Green byte, and the Z component being stored in the Blue byte,
- **UVW coordinates:** indicates the value of the UVW texturing coordinates at this



point in world coordinates, the U coordinate being stored in the Red byte, the V coordinate being stored in the Green byte, and the W coordinate being stored in the Blue byte,

- **Diffuse lighting:** indicates the amount of diffuse lighting arriving at the surface of objects at this point, unaffected by object colors,
- **Specular lighting:** indicates the amount of specular lighting hitting the surface of objects at this point, unaffected by object colors,
- **Material color:** indicates the color of the object that was hit at this point, unaffected by light,
- **Object ID:** produces a color coded picture that indicates the ID of the objects at each point in the final image (this information is not anti-aliased),
- **Material ID:** produces a color coded picture that indicates the ID of the material at each point in the final image (this information is not anti-aliased).
- **Global alpha mask:** produces a picture that is black where no object was found, white if an object was hit at this point.
- **Distance to camera plane:** renders the distance from the camera to the main intersection point, but as if the intersection point is a plane perpendicular to the camera direction.
- **Screen space velocity:**
- **World Point Position:** represents the position (in the world) of the point seen under each pixel. This is used by compositors in multiple ways, from scene re-lighting directly in a compositing package (such as Nuke or After Effect), to creating quick atmosphere effects such as ground fog. You will want to save this layer with High resolution format to have more precision.

Unknown flavor xStream,

Layer Masks

The layer masks category lets you create a mask for all objects that are placed in a given layer. If you unfold the layer mask category, you will see one line for each one of the layers in the scene.

Layer masks create color and an alpha image that are designed to work together. The alpha image appears white where objects from the selected layer are directly visible in the final picture, black elsewhere. Layer masks are fully anti-aliased.



Object Masks

Object masks are similar to layer masks, except that they can be created for each object independently. If you unfold the object masks category, you will see a list of all the objects in your scene. Place a check alongside the objects for which you want VUE to generate a mask.

Object masks create color and an alpha image that are designed to work together. The alpha image appears white where the object is directly visible in the final picture, black elsewhere. Object masks are fully anti-aliased.

You can unfold group objects in order to access sub-parts of objects and generate masks only for some sub-parts and not others.

Objects that have an EcoSystem material assigned to them will also appear as groups. If you unfold the group, you will notice that you have the option to generate one mask for the actual object and another mask for the EcoSystem population placed on that object.

Material Masks

Material masks are similar to object masks, except that they are created based on material rather than object. If you unfold the material masks category, you will see a list of all the materials in your scene. Place a check alongside the materials for which you want VUE to generate a mask.

Material masks create color and an alpha image that are designed to work together. The alpha image appears white where the selected material is directly visible in the final picture, black elsewhere. Material masks are fully anti-aliased.

You cannot generate masks for sub-materials of mixed materials.

EcoSystem Material Masks

EcoSystem material masks are similar to material masks, except that they are created based on materials used in EcoSystem populations rather than materials used on objects in the scene. If you unfold the EcoSystem material masks category, you will see a list of all the materials used in the different EcoSystem populations in your scene. Place a check alongside the materials for which you want VUE to generate a mask.

EcoSystem material masks create color and an alpha image that are designed to work together. The alpha image appears white where the population of the selected EcoSystem material is directly visible in the final picture, black elsewhere. EcoSystem material masks are fully anti-aliased.

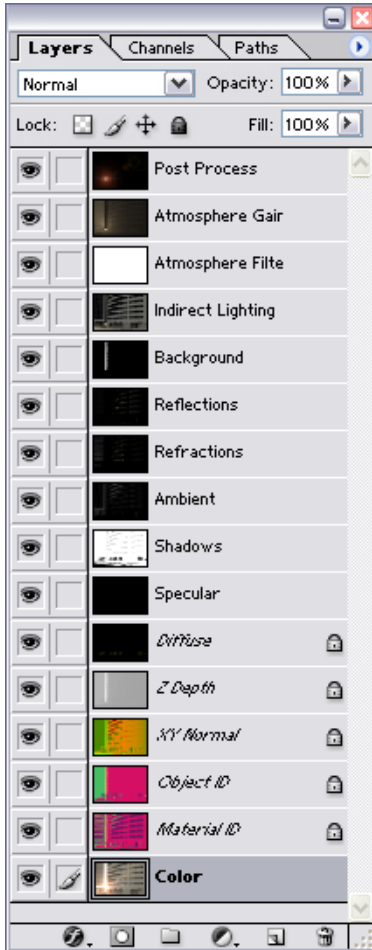
You cannot generate masks for sub-materials of mixed materials.



Cloud Masks

You can now create masks based on the Spectral clouds in the scene.

Saving as Multi-Layer Files



Sample Multi-Layer PSD Export

If you want to save the Multi-Pass Buffer information to file at the time of rendering, select the **Save to disk** option.

When this option is selected, you can either save the Multi-Pass information as separate files, or in a convenient multi-layer image.



Multi-layer Image (pre-combined): select this option to save the Multi-Pass information as a single PSD document.

The .exr format is also available for xStream. All components will be included as layers with the layer combination mode set so that the combination of all layers produces the most similar result as the final picture (the final picture is included on a separate layer for reference).

It is not possible to achieve a composite picture that is identical to the final render in Photoshop, because Photoshop does not support the “Additive” layer combination mode.

Masks are saved in the Photoshop document as both a separate layer for the mask color and a separate channel for the mask’s alpha. The number of layers and channels in the final Photoshop document will be displayed alongside the “Multi-layer PSD” option.

Note:

Photoshop only supports a limited number of channels per picture (24 or 56 depending on versions). You should check how many channels are supported by your copy of Photoshop before saving a picture with a lot of masks in it.

Unknown flavor xStream,

Separate files: each component/mask will be saved as a separate file. Masks will be saved as grayscale pictures.


Select the target file, format and location for the picture(s). If you select the “Multi-layer PSD” option, the file extension is automatically changed to **.PSD**. If you select the “Separate files” option, the actual file name for each layer/mask will be built from the name you entered plus a layer/mask identification.

High Dynamic Range: VUE generates all multi-pass renders in High Dynamic Range, including all object, cloud and layer masks, rendering components, shadows, reflections, atmospheric contributions, etc.

Multi-pass renders can be exported as single **.exr** 32 bit files or **.hdr** format containing all passes stored in high dynamic range format.



Sub-Ray Options






Sub-Ray Options dialog


You can access this dialog by pressing the **Edit** button alongside the **Trace reflections** and **Trace transparency** checkboxes in the *Render Options* dialog (see [here](#)). At least one of these options should be selected for the **Edit** button to be active.

This dialog lets you control the ray recursion level of VUE's ray-tracing engine. In order to avoid the ray-tracer becoming trapped in infinite inter-reflections (imagine two mirrors reflecting themselves to infinity) we need to instruct the ray-tracing engine to stop tracing reflections or transparencies after a certain number of recursions.

Max ray recursion depth: this setting controls the number of inter-reflections or refractions traced by VUE. When this number is reached, VUE stops tracing reflection and refraction rays.



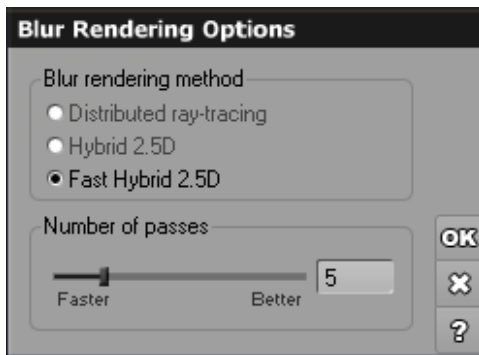




The effect of ray recursion depth on the rendering of two mirrors reflecting each other: from left to right, 2, 5, 10 and 32 ray recursion depths

Max total internal reflections: total internal reflections occur when light is trapped inside a refractive object: a ray of light originating from inside the object hits the surface of the object and is reflected inwards. This is typically what causes the glitter of diamonds: because of the high index of refraction of the diamond, light gets trapped inside and only manages to escape under very specific directions. This setting lets you control the ray recursion depth for total internal reflections.

Blur Rendering Options



Blur Rendering Options dialog



You can access this dialog by pressing the **Edit** button alongside the **Depth of field** and **Enable motion blurring** checkboxes in the *Render Options* dialog. At least one of these options should be selected for the **Edit** button to be active.

You have several [blur rendering options](#).

If you would like to use the **Distributed ray-tracing** blurring method, select the corresponding checkbox.

For **Hybrid 2.5D** rendering, you have two options:

- Hybrid 2.5D
- Fast Hybrid 2.5D

The **Fast Hybrid 2.5D** uses a new algorithm for depth of field generation. It is based on image blur like **Hybrid 2.5D** but uses a faster color spreading algorithm and works in conjunction with distributed ray tracing. Usually several passes are required to get all of the distributed ray-tracing noise smoothed out.

Systematic object anti-aliasing is incorporated inside **Fast Hybrid 2.5D**. Therefore anti-aliasing settings become linked to the depth of field settings. This means that only systematic anti-aliasing becomes available, and the minimum number of rays per pixel becomes equal to the number of depth of field passes (changing either of them changes both values).

Number of Passes

When you select either of the Hybrid 2.5 blurring method, the **Number of passes** setting becomes available. This is used to indicate the number of rendering passes used by the hybrid 2.5D blurring method. The higher the number of passes, the more accurate the result, but also the longer the render time.

The number of passes indicates to the render engine in how many “slices” the exposure time should be cut. The Hybrid blurring approach will then interpolate motion between each pass. For instance, if only one pass is specified, the rendering will take place at the middle of exposure time and the blurring will represent the entire motion covered during exposure time. If five passes are specified, the scene will be rendered five times and blurring between each pass will be computed progressively.

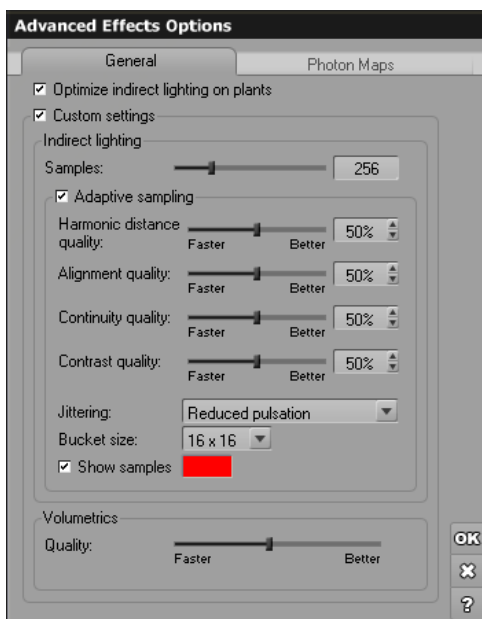
Rendering Motion blur usually requires less passes than Depth of field. Try to keep the number of passes as small as possible, as rendering time is directly proportional to the number of passes (10 passes will require 5 times more time than 2 passes). Suggested number of passes is 3-5 for an average scene and 10-15 for a scene with a very shallow DOF.



As you will see, the way the render in progress is displayed is different than for distributed ray-tracing. It is a good way of detecting the use of Hybrid 2.5D blurring effects.

If your scene does not contain any animation or any depth of field, the distributed ray-tracing method will be used regardless of whether you selected Hybrid 2.5D. This is because the results will be identical, but the distributed approach will render faster.

Advanced Effects Options



Advanced Effects Options dialog

You can access this dialog by clicking the **Edit** button alongside the **Advanced effects quality** setting in the *Render Options* dialog (see [here](#)). This button is only enabled if the **User** preset render setting has been selected.

This dialog offers you in-depth control over the rendering of advanced effects, such as the computation of indirect lighting as well as the processing of volumetric lights.

There are two tabs in this dialog:

- **General:** this tab takes care of general global illumination settings as well as the



rendering of volumetric effects.

- **Photon Maps:** this tab grants you in-depth control over the way photon maps are generated and used.

General Tab

- **Optimize indirect lighting on plants:** because of the intricate complexity of the geometry of typical plants, evaluation of indirect lighting on plants can be very slow. If you select this option, the processing of indirect lighting on the plants will be greatly simplified. As a result, the evaluation will be slightly less accurate, but also a lot faster. The results produced by the optimized evaluation are usually sufficient for rendering indirect lighting on plants. However, if you want perfectly accurate indirect lighting on your plants, you will need to deselect this option.

The controls in the **Custom Indirect Lighting** settings let you fine tune the way indirect lighting is evaluated in your scene. If you enable the **Custom settings** checkbox, the settings in this frame will override the *EasyGI™* “Advanced Effects Quality” setting of the *Render Options* dialog (see [here](#)).

- **Samples:** this setting controls the typical number of illumination samples that are processed to evaluate indirect lighting at each point in the scene.
- **Adaptive sampling:** when this option is checked, VUE will use a number of complex criteria to evaluate the frequency and accuracy at which the indirect lighting must be evaluated. If this option is disabled, the indirect lighting will be recomputed entirely at each sample. This will result in incredibly long render times and it is strongly advised that you do not disable adaptive sampling.
- **Harmonic distance quality:** this setting controls the way VUE evaluates the distance to the objects in the vicinity of a point in the image, and the way this distance influences the evaluation of the indirect lighting.
- **Alignment quality:** this setting controls the way VUE evaluates the alignment of the different lighting samples in space, and the way this alignment influences the evaluation of the indirect lighting.
- **Continuity quality:** this setting controls the way VUE evaluates the orientation of the different lighting samples in space, and the way this orientation influences the evaluation of the indirect lighting.
- **Contrast quality:** this setting controls the way VUE evaluates the contrast between the different sources of lighting and materials, and the way this contrast influences the evaluation of the indirect lighting.
- **Jittering:** this drop-down list controls the way the lighting samples are distributed in space. There are two options in the list:



Reduced pulsation: when this option is selected (the default), the samples are distributed in such a way as to reduce the low frequency pulsation that is typical of animation using adaptively sampled indirect lighting. This option is particularly useful when creating animations. If you are creating stills, the second option may be of interest.

Standard: this option ensures a better statistical distribution of lighting samples throughout the scene. This can result in slightly improved indirect lighting, but should be used only when rendering still frames. If you use this method when rendering an animation, you will notice a very unpleasant low-frequency pulsation in the indirect lighting.

- **Bucket size:** this drop-down list controls the base grid for the evaluation of the indirect lighting. You will have at least one sample for each bucket. Reducing the bucket size will increase the accuracy of the indirect lighting evaluation, but will also slow down renders quite significantly. This option can be useful if indirect lighting is consistently evaluated wrongly on small parts of your scenes.
- **Show samples:** if you check this option, the points at which the indirect lighting is evaluated will be displayed in the final picture as pixels of the indicated color. This is useful if you want to fine tune the evaluation of the indirect lighting solution and see the effects of the different settings above on this evaluation.
- **Volumetric settings:** this control lets you adjust the overall quality boost of the processing of volumetric effects (materials, lights, clouds).

Photon Maps Tab

The **Custom radiosity photon map options** frame lets you control the photon map that is used for the evaluation and rendering of radiosity.

- **Radiosity photons:** this setting controls the total number of photons that are sent into the scene in order to evaluate the radiosity illumination.
- **Maximum photon tracing level:** this setting controls how many times the light is bounced inside the scene. Higher values will result in a more accurate evaluation of the radiosity illumination, but also a longer processing time.
- **Custom photon gathering options:** when this option is checked, VUE will use custom options for the photon gathering.
- **Number of gathering photons:** this setting controls the number of photons that are used to evaluate the illumination at each point.
- **Maximum gathering radius:** this setting controls the maximum distance to a photon beyond which the influence of the photon will be ignored in the computation of the radiosity illumination.




- **Custom caustic photon map options:** the settings in this frame are identical to the ones in the radiosity photon map frame, except they apply to the caustic photon map instead of the radiosity photon map.

Anti-Aliasing Options

Anti-Aliasing Options



Hi-res render



Low-res render, no anti-aliasing





Low-res render, anti-aliasing

You can access the dialog [here](#) by pressing the **Edit** button in the **Anti-aliasing** group in the *Render Options* dialog. At least one of the anti-aliasing options must be selected for the **Edit** button to be active.

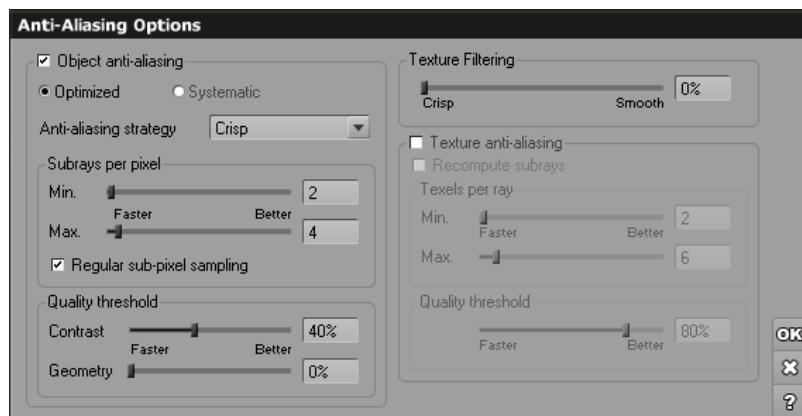
Anti-aliasing is a method used to reduce stair step effects (pixelization) on the edges of objects or textures. The method, called super-sampling, consists of tracing several sub-pixels for every pixel of the picture in order to improve transitions; the result being the creation of half tone pixels alongside the borders of objects/textures.

Anti-aliased pictures give the impression of having been rendered at much higher resolutions than that at which they were really rendered. The small drawback is that sometimes the picture appears slightly blurred. To achieve good results, many sub-pixels must be computed for each pixel, resulting in a considerable increase in render time. To optimize the method, more effort is concentrated on areas of transition.

Aliasing appears along the border of objects, as well as alongside sharp color transitions in texture maps. Object anti-aliasing improves the smoothness of the picture by re-sampling each pixel several times.



Object Anti-Aliasing



Object Anti-Aliasing Options

To enable **Object anti-aliasing**, check the corresponding box. Object anti-aliasing takes place at the end of the standard rendering pass.

VUE provides two ways of super-sampling object geometry:

- **Optimized:** method consists in super-sampling only the parts of the picture where transitions are found after the last render pass,
- **Systematic:** method will super-sample every single pixel in the picture during each render pass.

Anti-Aliasing Strategy

- **Anti-aliasing strategy:** this drop-down list lets you control how the different anti-aliasing samples are weighed into the final pixel:
- **Automatic:** when this option is selected, the renderer will use the most adapted strategy for each scenario, namely the Sharp method for rendering stills, and the Soft method for rendering animations.
- **Crisp:** this is the most accurate method, but also the method that requires the most samples in order to eliminate noise in the renders. It is the method used in prior versions of VUE.
- **Sharp:** this method is ideal for still renders. It produces relatively sharp results while efficiently eliminating noise.
- **Soft:** this is a slightly more blurry (and consequently less noisy) method of filtering, usually most suitable for rendering animations.



- **Blurred:** produces blurry results that could be suitable for certain types of animations.

Subrays

Super-sampling is handled in the following way: the render engine launches a first batch of rays and then, according to the results of this batch, decides if more sub-rays are required. When no more rays are required, it computes the average color and displays it. Systematic anti-aliasing yields slightly better results than optimized anti-aliasing, but at the expense of render times several times lengthier. It is usually not useful to use systematic anti-aliasing.

In the **Subrays per pixel** group you will find two controls that let you determine the minimum and maximum number of sub-rays computed for each pixel.

The **Min.** setting controls the number of rays initially sent inside a super-sampled pixel.

If the render engine decides that more anti-aliasing rays are required, it will keep sending new batches of rays until the total number of rays sent for that pixel reaches the **Max.** setting. For ultra-smooth results, you can bump this value up to 1024 (64 in Artist versions)! (although such high values will rarely yield better results than lower settings).

Regular sub-pixel sampling: when this option is selected, the rays in the first batch of sub-rays are placed exactly the same for all the pixels in the image. When it is not selected, sub-rays are cast randomly in each pixel. Although checking this option will usually produce better results, under certain conditions (regular patterns stretching to infinity), it may produce some visual interference.

The **Quality threshold** settings controls how the render engine decides whether more rays are required or not, after having computed the first batch. The higher the setting, the more often sub-rays will be sent into pixels.

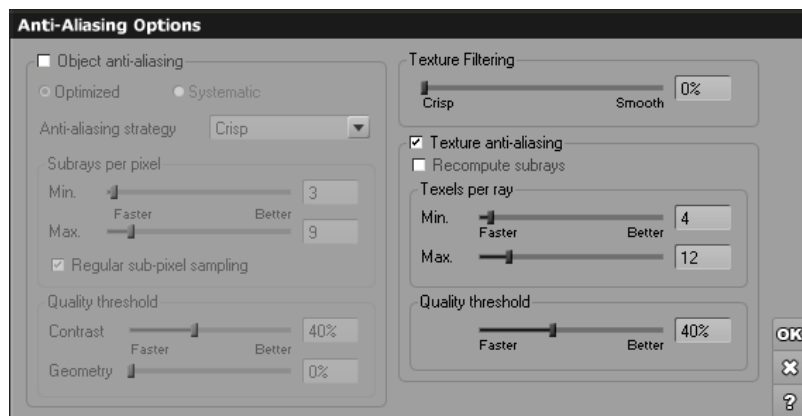
Contrast: this is a color-based anti-aliasing – basically the anti-aliasing that has existed in VUE in earlier versions. This compares colors: if the color difference in the corners is bigger than the threshold (**Contrast** setting), AA is applied.

Geometry: this is an edge-based anti-aliasing. It checks object IDs and depth.

Obviously, the higher these three settings, the better the quality, but the longer the render time...



Texture Filtering



Texture Filtering Options

Texture filtering controls the amount of automatic blurring that is applied to materials in the scene. This setting lets you control the overall “sharpness” of the render. For optimal results, this setting should be used together with the Anti-aliasing strategy setting (see above). Texture filtering is available from the **Broadcast** render setting on up through **Ultra**.

In the case of texture maps, the software automatically generates lower resolution versions of the images and uses them instead of the full-blown texture maps when they are seen from a distance.

While the results produced using some amount of filtering are generally smoother, you may occasionally find that your images are not as crisp as you would like them to be.

When rendering animations, it is recommended that you use some amount of filtering.

Note:

You can disable texture filtering on a per-image basis using [texture map nodes](#) and mip-mapping.

Texture Anti-Aliasing

Although Object anti-aliasing will take care of sharp color transitions as well, this comes at a high cost in terms of processing time. This is why VUE also offers a solution optimized for textures, known as “Texture anti-aliasing”.

This is a special form of anti-aliasing designed to reduce efficiently aliasing for both bitmap and procedural textures. Object anti-aliasing is good at cleaning up object and shadow



boundaries, but some textures might still display some moiré patterns or other unpleasant artifacts (for instance, in the distance because of a high frequency texture patterns, like when you render a ground plane with a checkerboard texture). In such cases, object anti-aliasing is not sufficiently efficient to clean-up rendering and eliminate these conspicuous artifacts in a reasonable time.

Texture anti-aliasing super-samples bitmap or procedural textures in order to properly integrate high frequency pattern repetitions. This very specific task is done much faster than object anti-aliasing, because it concentrates on the local properties of the texture rather than the entire scene. It is done by recomputing several texels (texture elements) for each pixel.

There is also the option of applying anti-aliasing on a per-texture basis in the *Advanced Material Editor* by clicking on the **Advanced Render Options** button.


To enable **Texture anti-aliasing**, select the corresponding option. This option must be turned on for anti-aliasing on a per-texture basis (in the *Advanced Material Editor*) to work. Remember that you can boost or reduce the quality of Texture Anti-Aliasing on a per material basis using the **Advanced Render Options settings**.

If the **Recompute subrays** option is checked, reflected and refracted rays will be traced for each texel. This can considerably improve anti-aliasing of reflection or refraction patterns, but will slow the anti-aliasing process down significantly. Except for specific cases (e.g. a reflective surface with strong bump mapping), this option is not recommended. If this option is unchecked, reflected and refracted rays will be computed without texture anti-aliasing.

The settings in the **Texels per ray** group are identical in their behavior to the **Subrays per pixel** settings described in Object anti-aliasing above.



Batch Rendering



Batch Rendering dialog

The *Batch Rendering* dialog features a list of all the jobs that are scheduled for rendering. You can add new jobs to the **List of jobs** by:

- Clicking the **Add** button and selecting a *.vue* scene file for rendering,
- Drag-dropping scenes onto the list,
- Clicking the **Add current** button to add the current scene to the list of render jobs, or
- Starting a new render with the [External Renderer](#) selected as the renderer.

Jobs are processed in the order they appear on the list. They are processed using the exact render settings stored in each individual scene, and saving the resulting images or animations to the location specified in the scene. If the scene is animated, the renderer will render the animation.

The order of jobs in the queue can be changed by highlighting the specific job and clicking on **Move job up** or **Move job down**.

You can remove scheduled jobs by highlighting them on the list and clicking **Remove**.

When a job completes, it stays in the **List of jobs**, so you can see which jobs were rendered last, and review the rendering statistics for these jobs. To remove the finished jobs from the list, press the **Clean finished** button.

When you schedule a scene for batch rendering, it is copied to the external renderer's scene folder. The scene will remain there until you clean the list of finished jobs or you remove the job from the list.


You can configure the external renderer to perform renders on your workstation, or, if you have access to a network of computers, the rendering can be distributed over your network of computers. Press **Edit** to display the *External Renderer Configuration* dialog and select the type of rendering.



The external renderer runs in the background. You can close the *Batch Rendering* dialog when you no longer need it. This will not stop render jobs from being processed. Render jobs run in parallel to VUE, so you can add jobs to the batch list, close the dialog and keep on working on the scene while the rendering takes place. You could even start other renders using VUE's internal renderer. VUE jobs will always run at a higher priority than the external renderer, so the latter will basically stop working while VUE is rendering.

Rendering To Screen -- The Render Display Window





Render Display Window

The *Render Display* window is made up of two sections. The top section is where the render takes place and the current render displays; the bottom section contains an area where previous renders can be stacked and options are available for manipulating them. If the stack option has been activated, a copy of the current render automatically moves to this **Stack** area when the render completes.

If you haven't just rendered, but wish to display previous renders, you can access this window from the **Render | Browse Previous Render** option on the VUE menu.



Current Render Display

This portion of the screen displays the current render. When the image does not fit the frame, you can drag the image with a left-mouse-button drag.

You can also zoom the image.

To the left, under the render screen are a row of buttons:



Compare: this toggles the [comparison mode](#).



Swap: if you have two images selected in the Stack below, clicking this button swaps the two images between the Stack area and the Render Display area.



Difference: this toggles an HDR difference view of the two selected renders for fine tuning differences. The slider adjusts the level of difference.

In the center under the render screen are five buttons:



Load: you can also load an existing image, previously rendered, into the stacked render for reference or comparison.



Merge: Rendered areas can be merged into full renders. This is useful if you render a large scene, make a small change and re-render only the area where that change is visible. By merging this area inside VUE, you can then readjust the post-processing settings (exposure,, contrasts, colors) directly, without the need of an external tool. The merged render can be saved.

To merge renders, first select the render that will be the “receiver” (surrounded by an orange frame), then click on the **Merge** icon and select the render that will be merged into the receiver (surrounded by a blue frame). You need to click on the **Merge** button again for the changes to take effect. The result is added to the stack.

Renders that were made in earlier versions of VUE are not mergeable. A render can be merged in another render only if its area is included in it and if the original full resolutions are the same.



Mirror display: This option flips the current render horizontally, giving you a different perspective of your render. This is helpful when evaluating composition and other aspects of the render. The flipped image cannot be saved, however.





Clear: this clears the stack. There is also a similar icon under each render to delete individual renders.



Options: this is where you turn the stacking feature on or off.

You can also select if you want to stack the Gbuffer, multi-pass buffers and diagnosis buffers.

Relighting buffers are also available.

You also have the options to disable stacking **Render Area** and **Preview** quality renders if you wish.

If checked, these are displayed in the **Stacked Renders** area below.

When the stack limit is reached a dialog will ask you if you want to stack the current render anyway. This deletes the first render in the stack and replaces it with the current one. You can always increase or decrease the stack size limit.

The option to use stacked renders can be turned on or off on the *Options* panel, **General Preferences** tab.

To the right, under the render screen, are a row of icons that become available when the render completes.



Zoom In / Zoom Out: use these to zoom in or out on the rendered image. Scroll bars become available if the image becomes bigger than the screen display.



Resize dialog to fit render: Click to remove any black space around the displayed render.



Full screen: select this icon to display the render full screen. Press **Esc** to return to the Render Display window.



Display Last Render (Color): displays the last color render.



Display Last Render (Alpha): displays the alpha channel of the last render.



Display Last Render (Depth): displays the depth channel of the last render.



Display Last Render (NPR): displays the NPR version of the last render.



Display Multi-Pass, Masks and G-Buffer: if these options were checked for the render, they can be displayed in the render area. Right-click on the icon to display the options that are available for viewing.



Current G-Buffer Layer: Click to page through the G-buffer layers rendered.



Post Render Options: click to display the Post Render Options dialog.



Save Displayed Picture: click to open the **Save As** dialog and save the render to disk.

The Render Stack

The Render Stack is the series of rendered images located under the current render. These are previous renders that have been saved. You can scroll through the renders and click a thumbnail to select it.

When you select a previous render, you can add a title on the thumbnail by right-clicking on it and selecting the command “Add/Remove editable thumbnail title”. It adds a gray area on top of the thumbnail. Click on that area to add a title or note to the render if you wish. This title can be removed with the same command, or simply by deleting all the characters with the keyboard which hides the gray area.

Right-click on the thumbnail to display a menu with the following options. Only options applicable to the render are available; other options are grayed out:

- Clone
- Delete clone
- ----- separator -----
- Delete render
- ----- separator -----
- Delete additional buffers
- Delete gbuffers
- Delete multipass buffers
- Delete relighting buffers



- Delete diagnosis buffers
- ----- separator -----
- Reapply camera settings to current camera
- Add/Remove editable thumbnail title

Icons under the thumbnail tell you what was rendered with the image (relighting, G-buffer, multi-pass). Click on the far right icon to display image information and any comments saved with the image. There is also a **Delete** icon available for this particular render.

In the *Render stacking options* dialog you find the following options:

- Auto-stacking of renders
- Disable stacking for render areas
- Disable stacking for preview renders
- Disable stacking for aborted renders
- Save a scene snapshot for each render
- Stack Gbuffer
- Stack multi-pass buffers
- Stack relighting buffers
- Stack diagnosis buffers
- **Stack size limit:** you can control how much can be stacked by limiting the size of the storage.
- **Stack folder:** The default folder for storing these stacked renders is `c:\user\username\appdata\Roaming\e-on\ software\VUE\ version\config\renderstack` on a Windows computer and on the Mac, it is `users/username/library/application\ support/e-on\ software/VUE\ version/config/renderstack`. On both types of computers, these are hidden directories. In Windows, this feature can be turned on in Window folders options. This location can be changed to another location if you wish by entering the pathname here.

There are three options for stacking your renders:

- **Use the same stack for all scenes:** the renders of all scenes are stacked in the same folder. This option is saved in the user preferences and reapplied for each new scene. You can use the default file location or change the location using the **Edit** button.
- **Use a dedicated stack for this scene:** the renders of the current scene are saved in a separate folder that you can choose by clicking the **Edit** button below the current **Stack folder on disk**. This option is only applied to the current scene and not saved in the user preferences.
- **Create a new stack for each new scene:** a new folder is automatically created



for each new scene. You can choose the root folder where all separate stack folders will be created by clicking the **Edit** button or use the default one. This option is saved in the user preferences.

Comparison

You can compare the current render with a previous render by first selecting the **Compare** button. This displays the renders available for comparison. Those not available will be marked.

Now move your cursor over the current render. You will see both renders with the cursor functioning as a horizontal separator (white line).

With the **Compare** button still on, if you activate the **Difference** toggle and use the slider to set a non-zero difference value, you can also see a difference display of the two renders. This is done in HDR so it's more powerful than your generic picture editor.

Reapplying Camera Settings to Current Camera

When viewing the **Stack**, if you see a camera view that you prefer over the current camera view, you can change the camera in the scene to the preferred setting by right-clicking on the render and selecting **Reapply camera settings to current camera** (or pressing **Ctrl+R**). This will change the camera settings in your current scene to the camera settings in the selected render.

This feature is only available for renders done in more current versions.

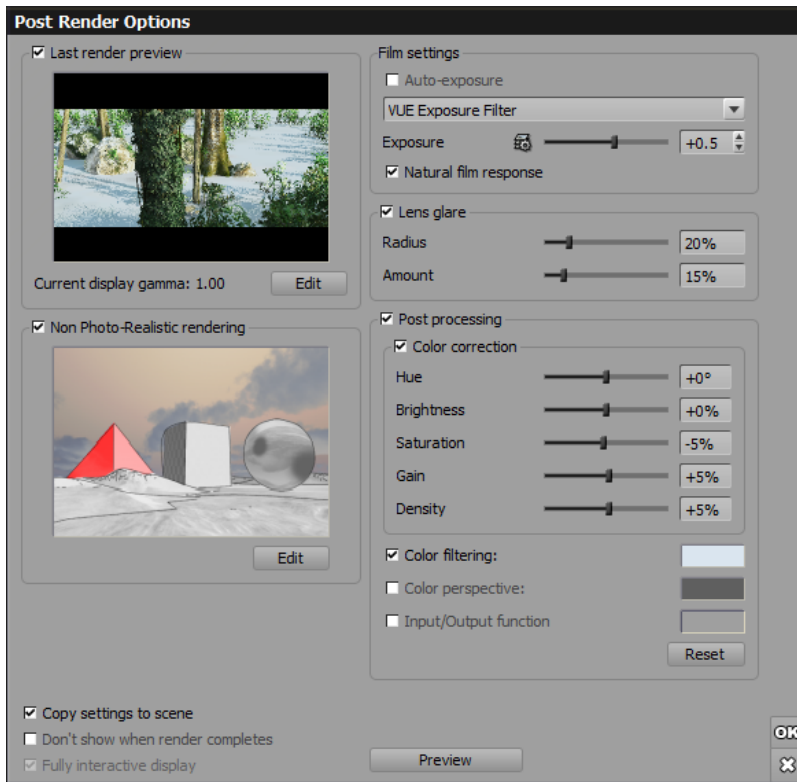
Hiding the Stack

If you aren't using the **Stack** feature, you can hide that part of the Render Display by clicking on the **Minimize** button in the upper-right corner, next to the **Close** icon.


If you wish to redisplay the **Stack** area, click on the now inverted **Minimize** button in the upper-right corner.



Post Render Options



Post Render Options

The *Post Render Options* dialog appears automatically when rendering completes. You can also display this dialog by clicking the **Post Render Options** icon () in the title bar of the *Render Display*.

This dialog lets you adjust post-processing options after the rendering completes – including adjusting the exposure. The settings in this dialog are a subset of those settings and work the same as the ones in the *Advanced Camera Options* dialog. The following interface elements are specific to the *Post Render Options*. Any changes made on this dialog are specific only to this render and do not affect the global settings in the *Advanced Camera Options* dialog.

- **Last render preview:** this picture displays the last render, with a preview of the



post processing effects applied to it.

- **Current display gamma:** You can change the gamma settings for this image by overriding the global setting in the program options.
- **Relighting:** This feature allows you to fine tune the lighting in your scene without re-rendering. You can instantly adjust the intensity and color of your different light sources in your scene in real time. This option has to be turned on in the *Render Options* dialog for the fields to appear here and is only available for a render quality of Final or above. You also have to enable **Stack relighting buffers** in the *Render stacking options* dialog.

Slide the bars to adjust the light (of the sun or any other lights you may have added to the scene) and you'll instantly see the results in the **Last render preview** window and your render, be it to screen or in the *Main camera view*. The **Full interactive display** option must also be checked for changes to be seen in your render. Be careful to not brighten the scene too much as this might produce unwanted artifacts. It's better to tone down light settings here rather than brighten.

There are separate adjustment sliders for **Sun Light** and **Ambient Light** control.

Non Photo-Realistic rendering

VUE contains settings to create [Non-Photorealistic Rendering \(Documentation/Building_Scenes/Rendering\)](#) with effects like “toon”. These settings are also available on the **Post Render Options** screen dialog to make adjustments. The settings are also available from the top tool bar on the screen.

- **Copy settings to scene:** when this option is selected, clicking **OK** to validate the changes will copy these changes to the scene.

Note:

If you are using Non Photo-Realistic Rendering and you use this setting, any changes that you make will be copied back to the Non Photo-Realistic Rendering in the Advanced Camera Options.

- **Show this dialog when render completes:** when this option is selected, this dialog will automatically appear when a render completes.
- **Fully interactive display:** Check this box to see any changes you make on this dialog in the **Main camera preview** window as well as the **Last render preview** window on this dialog.
- **Preview:** click this button to preview the effect on the full size image.



Film Settings

Photochemical films are made of tiny crystals of silver salt that react to light. When light reaches the surface of the film, it hits these crystals and triggers a chemical reaction that switches the state of the crystal (it becomes dark – this process is then inverted to result in a bright point). Once switched, a crystal will not be switched any further by more light hitting it (it can't be more black than black). It is the proportion of switched crystals that increases as light keeps on flooding in, making the point appear darker and darker. But, as more and more crystals have been switched by the incoming light, the chances of hitting an “unswitched” crystal go down. As a result, while points on the film will initially get dark very quickly, it will take more and more light to get them that much darker, resulting in a non-linear reaction to light. This non-linear reaction means that bright areas in the image will appear less bright, and dark areas less dark, resulting in a broader dynamic of light being visible in the final image.

- **Auto-exposure:** the difference in luminosity between noon and dusk is enormous, but we are not necessarily aware of this fact, because the human eye automatically adjusts to the amount of ambient light. Auto-exposure simulates this behavior by automatically adapting the exposure of the camera to the amount of light in the scene. If this option is enabled, your images will be correctly exposed, even if you drag the sun from noon down to dusk. When the auto-exposure option is enabled, the camera re-evaluates its exposure continuously during the rendering process. This is why, when tile rendering mode is enabled, the overall exposure of the image may be adjusted as rendering progresses.
- **Exposure filters:** VUE has several filters you can use to change the effect of your images.
 - **VUE exposure filter:** this is the filter used in all previous VUE versions.
 - * **Exposure:** slider controls the general exposure for the filter.
 - * **Natural film response:** select this option to enable the non-linear reaction to light typical of photochemical films.
 - **Photographic exposure filter:** provides several settings to tweak the response curve.
 - * **Exposure:** slider controls the general exposure for the filter.
 - * **Shadows:** will control how dark areas are mapped (lower value will produce lighter shadowed areas, higher values will produce stronger shadows)
 - * **Mid-tones:** will control how intermediate intensities will be mapped (lower values will tend to flatten the response curve, effectively producing darker and less contrasted colors, higher values will make the curve steeper, increasing intermediate colors contrast and overall brightness)
 - * **Highlights:** control how bright intensities are mapped (lower values will flatten them to avoid saturation, while higher values will produce saturated



bright areas).

- **Linear Exposure filter:** this is the simplest of all of the filters.
 - * **Exposure:** exposure can be manually adjusted by this slider. If no contrast transformation is performed, the tone curve is flat, and the user can control how steep the line is. Either **Map entire range** is selected and the brightest rendered color will be mapped to full white, scaling all other colors accordingly or a custom linear coefficient can be specified to control the scaling of colors (1 meaning no scaling at all: raw colors are preserved).
 - **Reinhard tone filter:** there are two flavors of this filter.
 - * **A simplified one:** that gives only control over contrast enhancement
 - * **Reinhard 2 Exposure Filter:** that lets you control:
 - **Brightness:** an overall scaling factor for the output
 - **Chromatic adaptation:** controls color contrast
 - **Light adaptation:** controls whether the tone operator acts rather locally (higher values) or globally (lower values) across the output picture.
- both types of Reinhard filter allow you to manually adjust exposure.
- **False colors:** this can be used to visualize output HDR range of intensities of the rendered picture more easily.

- **Exposure:** This setting controls the overall exposure of the camera. Positive values will result in brighter images, whereas negative values will result in darker images. The value is expressed in diaphragms. If the auto-exposure option is enabled, this setting is relative to the automatic exposure value (think of it as a way of “touching up” the auto-exposure).

When using the **OpenGL shader engine** (set on the *Options* panel, **Display** tab), you can preview the results of the camera auto exposure factor and the scene exposure factor. The auto exposure is automatically updated at each refinements pass of the mini scene preview. The **Exposure** is only active in the camera view, and is altered by the camera exposure factor. This feature is not available for **OpenGL (fixed hardware pipeline)**.

Lens Glare

Lens glare is caused by imperfections in the lenses of real-world cameras. Instead of being perfectly refracted by the lenses of the camera, part of the light becomes diffused by little defects in the glass. This results in halos of light appearing around very bright points of the image.

Lens glare gives a soft, realistic look to the final images. The effect, sometimes also referred to as “specular bloom”, is particularly strong when the camera lenses are a little dirty (because light becomes diffused by the layer of dirt at the surface of the lens).



Lens glare is controlled via the following settings:

- **Radius:** this controls the average size of the halos of light that appear around bright points in the image.
- **Amount:** this controls the intensity of the glare effect.
- **Warning:** the larger the radius, the slower it becomes to compute the effect of the glare. When previewing glare on the full size image, a much faster approximation of the glare effect is used. This approximation can sometimes result in slight visual artifacts. These artifacts will disappear when the full-blown glare algorithms are applied to the image before saving.

Post Processing

Post processing is a special processing pass that takes place once the picture is completely rendered. Using this feature, you can adjust the colors and brightness of the final picture without having to use another specialized application. By post processing pictures inside VUE rather than using an external application, you ensure that the resulting colors retain all of their subtlety (when you save a picture, the colors in the picture are limited to 8 bits per pixel; artifacts and color banding can appear rapidly as soon as you affect anything but minor post-processing). In VUE, colors are computed and processed with a resolution that is literally several million times more refined than in an exported picture.

To enable post processing of your picture, check the option. The post processing controls become available:

- **Color correction:** select this option to apply color correction to your picture. **Hue** shifts the color tones according to the angle indicated. **Brightness** will increase or reduce the overall brightness of the picture, while **Saturation** modifies the overall saturation of the picture. The **Gain** setting applies a smooth contrast to the picture. **Density** adds uniform density to all colors in the picture.
- **Color filtering:** this option lets you apply a color filter to the picture, as if it were seen through a colored gel. When you check this option, you can adjust the corresponding color by double-clicking on the color control.
- **Color perspective:** if you select this option, dark colors will be replaced by the indicated color. Black will be replaced with this exact color, while brighter colors will be blended according to the brightness of the color. When you check this option, you can adjust the corresponding color by double-clicking on the color control.
- **Input/Output Function:** this allows a very accurate tuning of contrasts and luminosity of each color channel (directly on the 32-bit float RGB components). Check the function to enable it. Right-click on the small function preview next to it and select **Edit Filter** in the pop-up menu. The *Color Input/Output Function Filter Editor* displays for changing the settings. See the [Editing Filters](#) section for more information about working with filters.



At the bottom of the dialog

- **Copy settings to scene:** when this option is selected, clicking **OK** to validate the changes will copy these changes to the scene.

Note:

If you are using Non Photo-Realistic Rendering and you use this setting, any changes that you make will be copied back to the Non Photo-Realistic Rendering in the Advanced Camera Options.

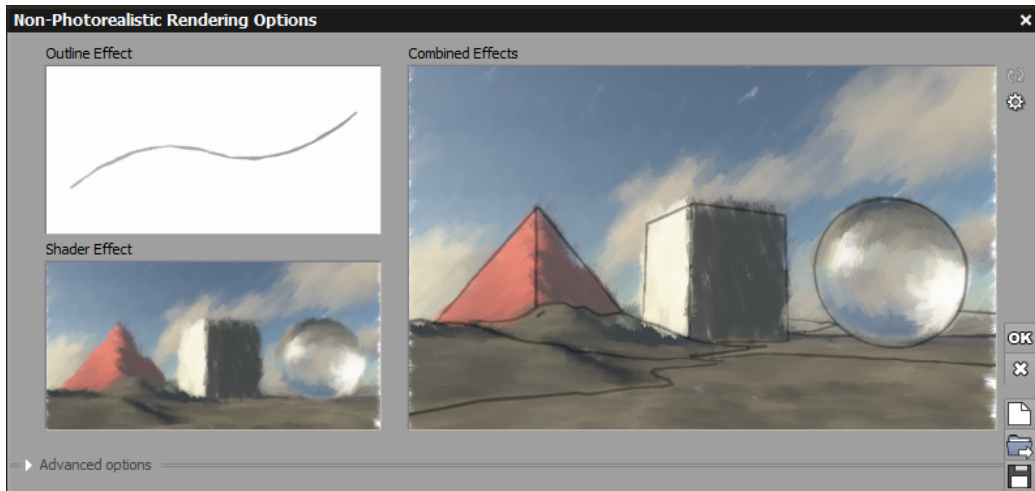
- **Show this dialog when render completes:** when this option is selected, this dialog will automatically appear when a render completes.
- **Fully interactive display:** Check this box to see any changes you make on this dialog in the **Main camera preview** window as well as the **Last render preview** window on this dialog.
- **Preview:** click this button to preview the effect on the full size image.

Non-Photorealistic Rendering

The NPR feature (Non-Photorealistic Rendering) transforms regular VUE renders into drawings, paintings, cartoons or stylized images.



Activating NPR and selecting style presets




NPR must be enabled before launching a render, not after it, because VUE needs to compute additional data to generate the different effects. At the first activation of NPR, a window displays some effects. Click on **Advanced options** to display the various tabs and settings.

The selection of existing presets is done with a style selector, a control containing a small rendering of the style and displaying a list of style names when clicking on it. When hovering the mouse cursor on names, their preview is displayed. A style can be selected by clicking on it.

From the main toolbar

The first and quickest way to enable NPR is to click on the button located in the main toolbar at the top of the VUE interface.



A right click simply enables or disables NPR depending on the current state. A left click enables it and opens the NPR Options dialog. On top of this dialog is the three styles selector, one for the outline style (top left), one for the shader style (bottom left) and one combining the two previous styles (right side). The advanced options of this dialog are described in the section “Editing NPR styles”.

From the Avanced Camera Options

The second way to activate NPR is to open the Camera Options (by double clicking on a Camera in the world browser), and to ckeck the option “Non-Photorealistic Rendering” in the bottom right corner. Under this option is a style selector listing the NPR presets.

Note:

NPR styles are in fact attached to cameras, like the exposure and post processing settings. When using the NPR button in the toolbar, it’s the settings of the active camera that are edited.

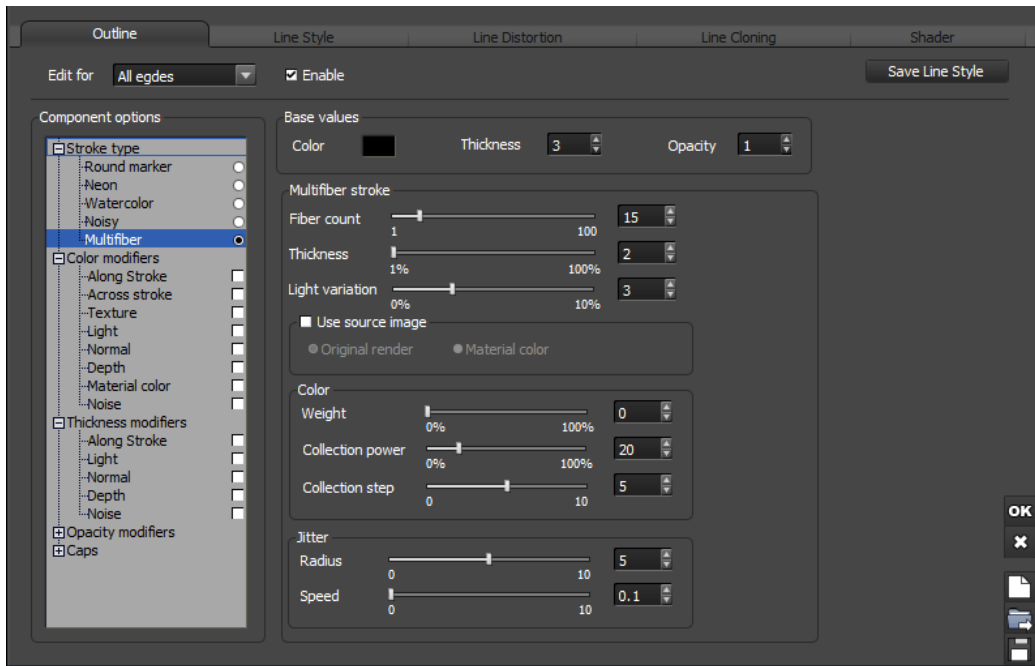
From the Post Render Options in the render stack

The third way to enable or disable NPR is to do it from the render stack (provided that it was first enabled before the render as said above), by clicking on this icon, which opens the Post Render Options dialog, in which “Non-Photorealistic Rendering” can be checked. The same style selector that you can find in the Advanced Camera Options.

Outline Tab

Outines are drawn on the edges of objects, materials, and where normals have sharp variations. You can control where they appear with the dropdown list on top of this tab. When selecting an a type of edge, the checkbox beside it enables the drawing. Below are the options for the currently selected type of edge.





- Base values :

In this frame, the 3 base values of the stroke are defined: color, thickness and opacity. They can be altered in various ways using the modifiers (see section below).

Stroke type

- **Round Marker Stroke:** The default stroke, simple and regular. It has a constant width and opacity.



- **Neon Stroke:** This stroke has a decreasing opacity, across its width. Several pa-



Parameters control how the opacity varies.



- **Watercolor Stroke:** This stroke imitates watercolors, with a slightly varying thickness and opacity.



- **Noisy Stroke:** This one imitates various types of strokes: chalk, charcoal, ink, gouache... It has a varying thickness and constant opacity.





- **Multifiber Stroke:** This stroke imitates oil or gouache paintings. It is composed of several small lines, like the fibers of paintbrushes.



Modifiers

Modifiers control 3 aspects of the stroke: color, thickness and opacity. It modifies them depending on various parameters, which are either data from the original render (like the objects' normal, depth or light...) or a local information (like the distance to the stroke's origin or median).

Color modifiers


Most color modifiers include a color map used to compute the output color.

- **Along stroke:** The color map is applied along the stroke. The picture below shows an example with a simple color map varying from green to blue.





- **Across stroke:** The color map is applied across the stroke. The picture below shows an example with a simple color map varying from green to blue.




- **Texture:** This modifier applies a texture file along the stroke. In the picture below, a image representing a star is loaded in the modifier's options.




- **Light:** In this modifier, the final color is taken in the color map, from its beginning (right) when the underlying pixel has low luminosity, to its end (left) when the underlying pixel has a great luminosity.



- **Normal:** This modifier applies a color map depending on the underlying pixel normal.
- **Depth:** In this modifier, the final color is taken in the color map, from its beginning (right) when the underlying pixel has a small depth, to its end (left) when the underlying pixel has a great depth. In the picture below, the terrain in the foreground and the sphere are blue because they are close to the camera. The pyramid and the horizon are green because they are distant.



- **Material color:** This modifier applies the color of the underlying material to the stroke. It is the brut material color, without any atmosphere effect, radiosity, reflection, etc...
- **Noise:** This modifier applies a noise to the stroke's color. It can have one or two dimensions, and one to eight octaves. It can be applied either to the hue, luminosity or saturation component of the base color. The picture below shows an example of a noise applied to the hue of the stroke (the base color is blue).




Thickness modifiers

Most thickness modifiers include a filter used to compute the output thickness.



- **Along Stroke:** The filter is applied along the stroke. The picture below shows an



example with a simple filter varying from zero to one.



- **Light:** This modifier applies a filter depending on the underlying pixel lightness.
- **Normal:** This modifier applies a filter depending on the underlying pixel normal.
- **Depth:** In this modifier, the final thickness is taken in the filter, from its beginning (right) when the underlying pixel has a small depth, to its end (left) when the underlying pixel has a great depth. In the picture below, the terrain in the foreground and the sphere have thick outlines because they are close to the camera. The pyramid and the horizon have thin outlines because they are distant.



- **Noise:** This modifier applies a noise to the stroke's thickness. It can have one or two dimensions, and one to eight octaves. The picture below shows an example of this modifier.









Opacity modifiers

Most Opacity modifiers include a filter used to compute the output opacity.

- **Along Stroke:** The filter is applied along the stroke. The picture below shows an example with a simple filter varying from zero to one.





- **Across Stroke:** The filter is applied across the stroke. The picture below shows an example with a simple filter varying from zero to one.



- **Texture:**



- **Light:** This modifier applies a filter depending on the underlying pixel lightness.
- **Normal:** This modifier applies a filter depending on the underlying pixel normal.
- **Depth:** In this modifier, the final opacity is taken in the filter, from its beginning (right) when the underlying pixel has a small depth, to its end (left) when the underlying pixel has a great depth. In the picture below, the terrain in the foreground and the sphere have opaque outlines because they are close to the camera. The pyramid and the horizon have almost transparent outlines because they are distant.




- **Noise:** This modifier applies a noise to the stroke's opacity. It can have one or two dimensions, and one to eight octaves. The picture below shows an example of this modifier.

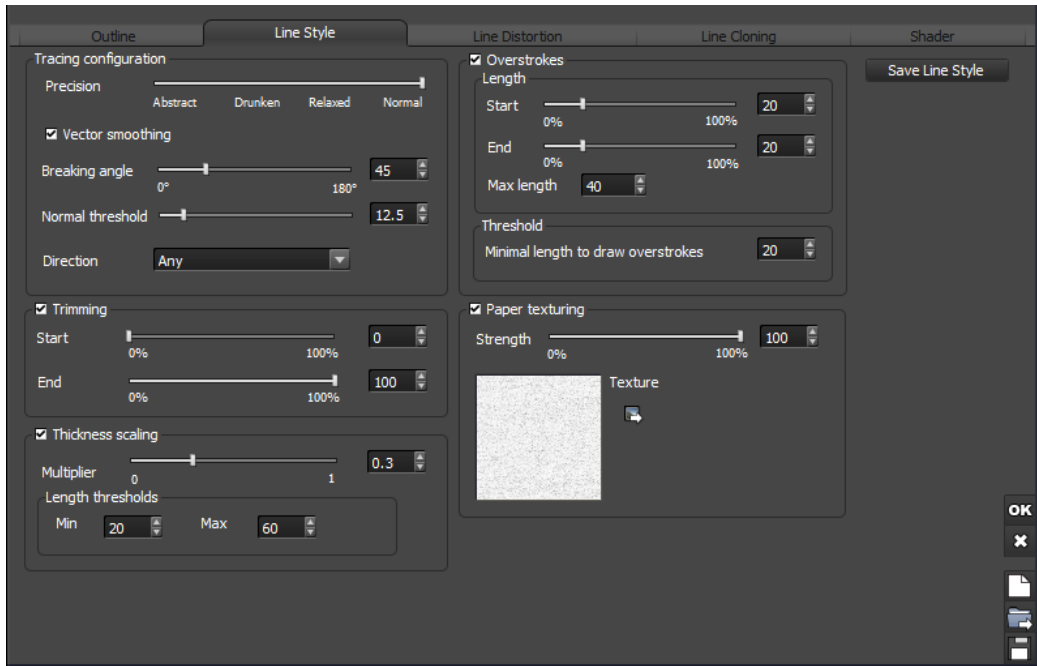




- **Pattern:** This modifier applies a pattern to the stroke's opacity. The pattern is defined by a row of black and white squares. The black ones are for parts with full opacity, and the white ones are for parts with a null opacity (transparent). The picture below shows an example of this modifier.




Line Style Tab



The **Line Style Tab** creates a style for your exterior lines, the way the lines are drawn.

Tracing Configuration: This defines how the line follows the chosen path.

Precision: Using the slider, set the line precision using a range from Abstract to Normal. The picture below shows the effect of this parameter, from left to right: Abstract, Drunken, Relaxed, Normal



Vector Smoothing: Check the box to use vector smoothing.

Breaking angle: Using the slider, set the breaking angle from 0 to 180 degrees.

Normal threshold: Set the threshold. Value is from 1 to 100.




Trimming: If using trimming, set the start and end limits.

Thickness scaling: Indicate the multiplier you are using.

Length Thresholds: Set the Min and Max of the Length thresholds.


Overstrokes:

Check to extend the lines at either side. The length of overstrokes depends on the parameters **Start** and **End**, which are a percentage of the original length of the stroke. This length can be limited by using the **Max length** parameter. To avoid drawing overstrokes on small details, use the **Threshold** parameter (**Minimal length to draw overstrokes**).

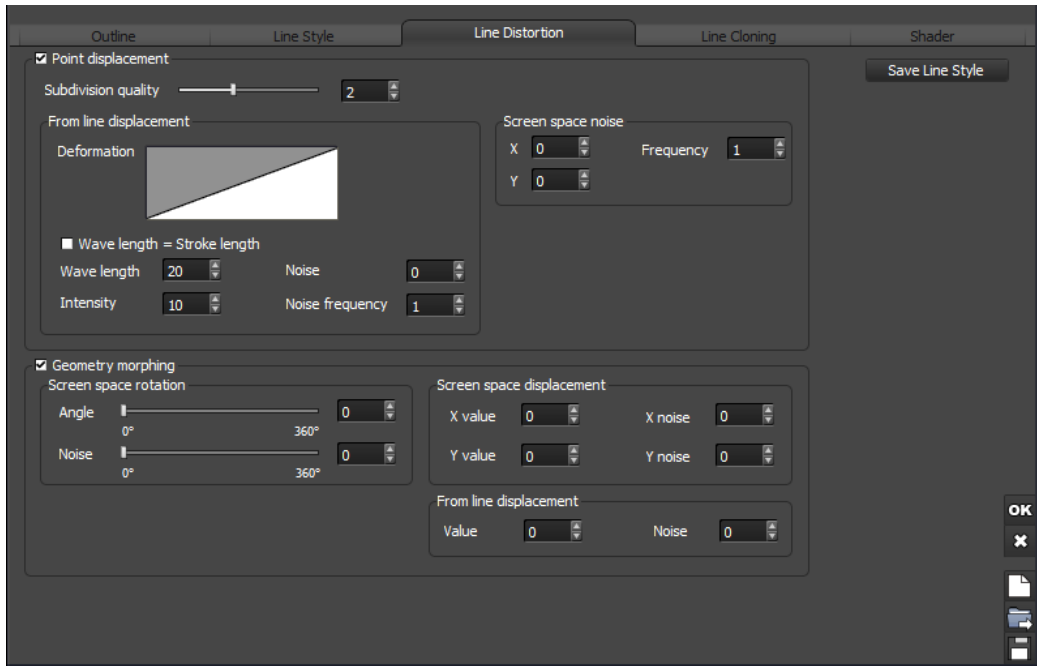


Paper texturing:

This is used to imitate a line drawn on paper. The line's color is tweaked by a bitmap. The bitmap is in fact used as an elevation map, to define the relief of the paper (black is low, white is high). When drawing the line on the relief, low areas will receive less color, whereas high areas will receive more color. The picture below shows the effect disabled on the left, and enabled on the right (with a **Strength** of 50%).



Line Distortion Tab



- Point Displacement

As the lines are built from several segments, it is possible to tweak each segment independently.

Subdivision quality: Use the slider to increase the quality of subdivision.

From line displacement: Use the graph to indicate deformation. Just click on the graph to select a filter for deformation. The Filters browser displays the existing filters available.

Wave length = Stroke length: Check to use this setting. Set values for **Wave length**, **Intensity**, **Noise**, and **Noise frequency**.

Screen space noise You can set the values for X and Y as well as the frequency.

- Geometry Morphing

Screen space rotation Defines the rotation without displacement of each segment. You can add some noise to have more random rotation. **Angle:** Set the rotation angle using the slider. **Noise:** Set the noise level using the slider.

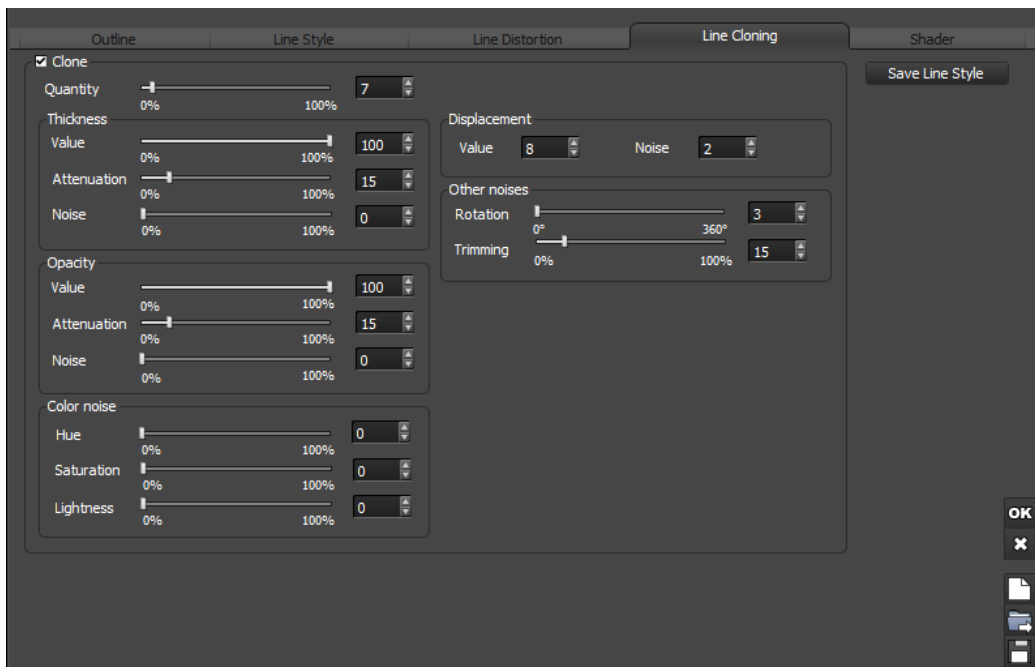
Screen space displacement is a joined vertical or horizontal, or both, displacement of

each segment. The noise to have a non joined displacement. Set an X and Y value for both **Value** and **Noise**.

From line displacement The From line displacement lets you scale each segment away from its initial position.

Set an X and Y value for both **Value** and **Noise**.

Line Cloning Tab



Use the **Line Cloning Tab** to multiply the lines, select their quantity, thickness, opacity and randomness in colour. You can also tweak their position spatially by applying a displacement or a rotation to them.

Quantity: Set the amount of cloned lines. The picture below shows an example with 5 clones.





Thickness: The **Value** parameter sets the base thickness of the cloned lines. It is a percentage of the thickness of the original line. The **Attenuation** parameter defines how much the clones' thickness will be reduced. It is applied symmetrically on both side of the main line. The **Noise** parameter adds irregularities in the attenuation, so that it looks more natural. The picture below shows an example with an attenuation set to 30% (and without noise).



Opacity: The **Value** parameter sets the base opacity of the cloned lines. It is a percentage of the opacity of the original line. The **Attenuation** parameter defines how much the clones' opacity will be reduced. It is applied symmetrically on both side of the main line. The **Noise** parameter adds irregularities in the attenuation, so that it looks more natural. The picture below shows an example with an attenuation set to 30% (and without noise).





Color noise: These change the value of the color noise.

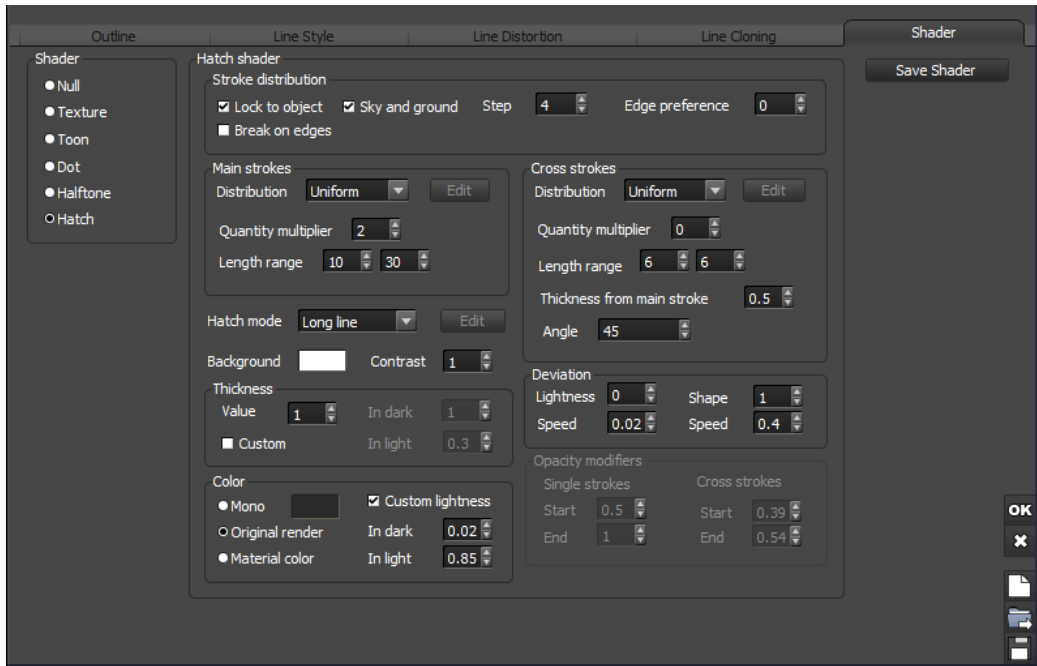
Displacement: The **Value** parameter defines how much space the cloned lines can take. The space between each cloned line is equal to this value divided by the number of clones. The picture below shows an example of 5 clones with a displacement set to 20 and a noise set to 10.



Other Noises: These set the value for other noise.



Shader Tab



The Shader Tab is where you can select the type of shader you wish to use. The shader options help you with the colours and materials of the scene. You can either keep the original ones (Shader: Null, Use original render), or replace them by:

- a unique color for the whole scene (Shader: Null, Fill with color)
- textures that you can load from a browser
- a color maps

Null Shader

The null shader gives you two options – to use the original render or to fill the space with color. You can select the color by clicking on the color square to open the Color Selector dialog.

Texture Shader

The texture shader allows you to load up to nine textures. Click the *Load* button to open the bitmap browser for texture selection.



Toon Shader

If you are creating a cartoon effect, this tab sets the coloring up for the shader.

Toon shader options:

Levels: Select the color map you wish to use. To change, click on the color map to open the color map browser and select a new map.

Count: Select the count

Smoothness: Adjust the color smoothness. Values go from 0 to 100%.

Dot Shader

The dot shader options consist of:

- **Pen color:** Indicate the color of the dots. Click on the color square to select the color.
- **Background:** Set the color of the background. Click on the color square to select the color.
- **Dot size:** Set the size of the dot.

Halftone Shader

This is used to add a pattern on top of the image. Halftone shader options are:

- **Pattern size:** Set the size or scale of the pattern.
- **Pattern angle:** Set the angle of the pattern from 0 to 360 degrees.
- **Pattern type:** Select the pattern type. Available types are stripes and dots.

Material color: You can select the Pattern Color and Background color for the pattern by unchecking the Material color. When Material color is selected, the pattern uses the underlying material color.

Hatch Shader

Hatch applies line strokes to the image. These are placed according to the parameters. These smaller lines have the same parameters as the Lines Options settings.

Stroke Distribution

- **Lock to object:** Check to lock the strokes to an object.
- **Sky and ground:** Check to include marks on sky and ground.
- **Step:** Indicate the step.
- **Edge preference:** Indicate the edge preference.



Main strokes

- **Distribution:** indicate whether stroke distribution is uniform, by lightness or by normal. With the lightness or normal options, the **Hatch Shader Options** dialog displays with fields to enter the **Input range**, **Output range**, and **Curve exponent**.
- **Quantity multiplier:** Indicate the quality modifier.
- **Length range:** Indicate the length of the stroke.

Cross strokes Distribution: Choices are: Uniform, By lightness, or By normal. With the lightness or normal options, the **Hatch Shader Options** dialog displays with fields to enter the **Input range**, **Output range**, and **Curve exponent**.

- **Quantity multiplier:** Indicate the quality modifier.
- **Length range:** Indicate the length of the stroke.
- **Thickness from main stroke:** Set the thickness.
- **Angle:** Set the angle. Range is 0 to 360 degrees.

Hatch mode: Values are **Simple line**, **Long line**, and **Procedural line**. If you choose **Procedural line**, click the *Edit* button to display the **Line Options dialog to further define the line if needed**. **Background:** Select the color of the background. To change, click the color box. **Contrast:** Enter the contrast setting. **Thickness:** Either enter a thickness value or select **Custom**. Then you can set the value for **In dark** and **In light**.

Color Check either **Mono** or **Custom lightness**. With the **Mono** setting you can change the color by clicking on the color box. With the **Custom lightness** setting, you can set the value for **In dark** and **In light**. You can also opt for **Original render** and **Material color**.

Deviation

- **Lightness:** Set the lightness value.
- **Shape:** Set the shape.
- **Speed:** Can be set for both **Lightness** and **Shape**.

Opacity modifiers For opacity modifiers, there are single strokes and cross strokes. Indicate the **Start** and **End** values for both types of strokes.

Saving styles

There are three kind of styles: Line Styles and Shader Styles, and Combined Styles (ie. Line + Shader). NPR styles are saved in .sty files.

- **Line Style:** The “Save Line Style” button lets you save the settings from these tabs (all together): Outline, Line Style, Line Distortion, Line Cloning. When choosing a destination file, you can put it either in the folder “NPR Styles/Listed/Outline” if



you want to access it quickly from the Line Style selector (top left image), or in the folder “NPR Styles/Personal” to access it from the browser.

- **Shader Style:** The “Save Shader” button lets you save the settings from the Shader tab only. When choosing a destination file, you can put it either in the folder “NPR Styles/Listed/Shader” if you want to access it quickly from the Shader Style selector (bottom left image), or in the folder “NPR Styles/Personal” to access it from the browser.
- **Combined Style:** The main “Save” button (with a floppy dick) lets you save all the settings in one .sty file. It is the combination of a Line Style and a Shader Style. When choosing a destination file, you can put it either in the folder “NPR Styles/Listed/Combined” if you want to access it quickly from the Combined Style selector (large image on the right), or in the folder “NPR Styles/Personal” to access it from the browser.

Path Tracer

The Path Tracer, a hybrid CPU / GPU renderer, naturally simulates many effects that have to be specifically set with other methods, such as soft shadows, depth of field, caustics, and indirect lighting. GPU acceleration allows for interactive rendering updates while editing your scene. The Path Tracer doesn’t need additional computation for features like various soft effects (e.g. soft shadows, blurry reflections and refractions, depth of field), nor any kind of pre-pass calculations for indirect lighting. The only artifact that can be generated by the Path Tracer is high-frequency noise, which is guaranteed to disappear with a sufficiently high amount of samples per pixel, as this renderer is unbiased.

Path tracing is available as an interactive rendering quality in the viewports, or as an offline rendering engine. Currently, only the main viewport (active camera view) can be replaced with the path tracer output in real time.

Several VUE features are NOT supported when rendering with the path tracer:

Volumetric lighting	Light gel effect
non Photometric lighting	Rain and snow
Procedural materials	Shadow maps
Multi-pass buffer	Planets
Relighting	Automatic exposure
G-buffer	Motion blur
Lens flares	Light dispersion
Spherical/planetary terrains	Dynamic EcoSystem population



Daylight portals
Ecosystem phasing

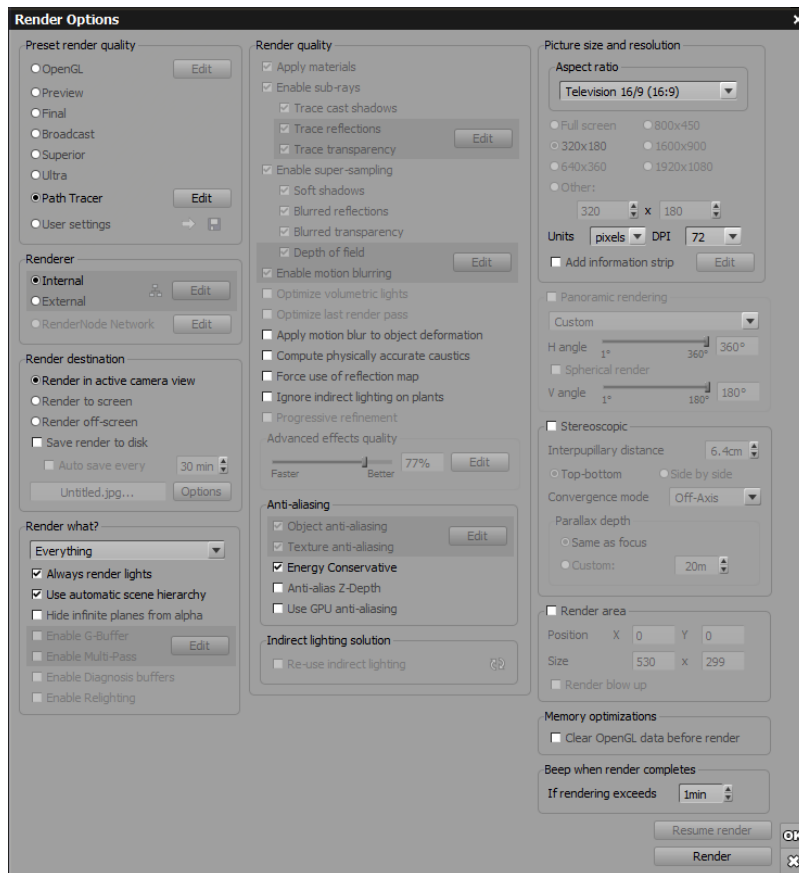
Two-sided materials
Network rendering

How to Use It

The **Interactive Path Tracer** is activated by default in the main view. Furthermore, the activation of the **Path Tracer** is remembered directly in the scene file, so that loading a scene that was used with the **Path Tracer** automatically activates it.

If you don't want this behavior, you can prevent it by setting the **Reset Interactive Path Tracer viewport option** to On. It is available in the **Options** dialog, the **Display Options Tab**.

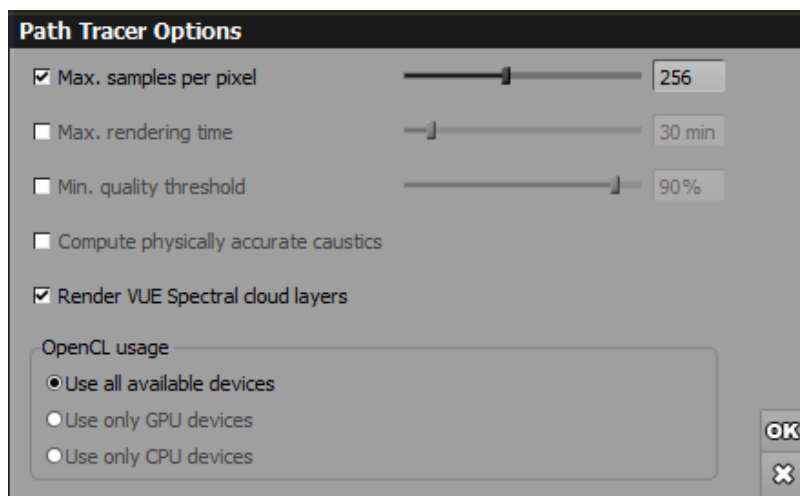
You can change the render quality of any viewport to **Interactive path tracing**. This option is available from the **View Display Options** icon in the viewport header.



Render Options dialog

Or, you can modify the parameters of the engine inside the *Render Options* dialog.

The Path Tracer options dialog can be accessed from the **Preset render quality** section, called **Path Tracer**. Beside it, a button labeled **Edit** displays a popup.



Path Tracer Options

This popup lets the user set the path tracer specific render options, detailed below:

Main render options

- **The maximum number of samples per pixel.** This must be a power of 2.
- **The maximum rendering time, in minutes.**
- **The minimum quality threshold.**
- All the parameters above define a different kind of limit to the end of the path tracing for a given frame.
- **Compute physically accurate caustics** when this option is selected, [realistic caustics](#) are computed for the scene, including spectrum dispersion.

Note:

Computing realistic caustics can add a significant amount of noise, depending on lighting conditions. Consequently, a higher amount of samples per pixel will generally be needed when this option is turned on.

- **Render VUE Spectral cloud layers** when this option is selected, the dynamic



clear sky model will be replaced with a progressive sky environment map, filled on demand by the path tracer, that will include VUE spectral cloud layers.

Note:

For performance reasons, this sky environment map will always be rendered from a low altitude point of view (regardless of the actual camera placement), making it suitable for most architectural scenes. For in or above clouds points of view, it is recommended to rather use the standard VUE renderer.

OpenCL usage

The **OpenCL usage** section lets the user choose what kind of devices should actually be used by the path Tracer. Currently, OpenCL devices can be of two kinds: CPU or GPU. VUE can be set to use all kinds of available devices, or devices of a single kind (either CPU or GPU).

- **Use all available devices** is the default option, using all devices regardless of their kind.
- **Use only GPU devices** will run OpenCL kernels only on available GPUs.
- **Use only CPU devices** will run OpenCL kernels only on available CPUs.

Note:

Changing the OpenCL usage mode will require VUE to be restarted

Tips

In order to get high quality images from path tracing, a large number of rays must generally be traced to avoid visible noisy artifacts, especially when the scene contains specular materials (glass, metals, mirrors, etc...).







For scenes containing a lot of specular materials, if caustics are of little interest, turning off the **Compute physically accurate caustics** option can significantly reduce noise, thus potentially needing less samples per pixel.

For scenes not using any Spectral VUE cloud layers, or if their rendering is not needed, turning off the **Render VUE Spectral cloud layers** option can significantly speed up the rendering.



Saving (Exporting) Images

Once rendering is complete, you can export the picture to other 2D applications (e.g. Photoshop™) using any of the following:

- Click the **Save displayed picture** icon () to the left of the channel buttons, in the *Main view's* title bar. This will save the currently displayed channel.
- Click the **Save color picture** icon (). This will save the color channel, regardless of which channel is currently displayed.
- Select the menu command **File | Export Picture** and select the channel you want to save.
- When rendering to screen, in the **Render Display** dialog, press the **Save picture** button (). The channel that will be saved is the one that is currently displayed (use the , , and  buttons to change channel).

Supported picture file formats are: **BMP**, **JPG**, **GIF**, **HDR**, **EPX**, **IFF**, **PCX**, **PNG**, **PSD**, **TGA**, **RLA**, **RPF** and **TIFF**.

If you save using the Piranesi EPX format, distance and material information will be included inside the file (you need to render the Material ID component, either as a G-Buffer channel or as a Multi-Pass component).

If you have enabled rendering of G-Buffer information, you can save this information as Run-Length Encoded (**RLA**) or Rich Picture Format (**RPF**) files using the above method and saving the Color channel.

When rendering in Multi-Pass, you have also the option to save the picture in a multi-layer, multi-channel Photoshop **PSD** document. However, you need to decide this before beginning the render. It is not possible to save as a multi-layer, multi-channel Photoshop PSD document after the rendering completes.

The .exr format is also available for xStream users.



xStream




Saving Animations

Animation file formats must be selected before the rendering of the animation begins. You cannot change the animation file format once the animation has been rendered. You should use an external application to do this.



Printing Pictures



Print Preview dialog

You can also print pictures directly from within Vue. Select the menu command **File | Print** to print the color channel of the picture that was rendered last. The picture will be sent directly to the active printer for printing. The print options are not available until you have rendered a picture.

To select the active printer and set it up, select the menu command **File | Printer Setup**. This displays the system's standard printer selection dialog. Please refer to your system's help for help on this dialog.

Print Preview

Use the menu command **File | Print Preview** to open the *Print Preview* dialog. This dialog shows a preview of the picture on the page, and lets you adjust the size of the printed picture.

The printable area of your printer is automatically retrieved and taken into account when positioning the picture. On some printers, the printout can go up to the edges of the sheet of paper, but this is not the case on all printers. This is why a white strip may remain around the picture if your printer doesn't support printing up to the edges.

If you select the **Full page** option, the picture will automatically be resized to fill the largest printable area on the page. The **DPI** setting indicates the resulting picture resolution on paper. It is generally admitted that optimal quality is achieved at 300 DPI.



Printing at a lower DPI setting will result in a blurred/jagged picture look. To increase the DPI setting, switch to **Manual zoom** and change the zooming ratio until the required DPI setting is achieved, or use the DPI setting directly. As you will notice, the size of the printed picture reduces accordingly. At 100% zoom, the DPI setting is 72 DPI by default.

To increase the DPI setting without reducing the size of the printed picture, you will need to render the picture at a higher [resolution](#).

You can adjust your printer's settings by clicking on the **Printer setup** button.

Once you are satisfied with the way the picture looks on the preview page, press **OK**, and select the menu command **File | Print** to actually print the picture, or press **Print** directly in the dialog.



Section 3

Importing and Exporting





Importing Objects

VUE provides a comprehensive set of import filters that can be used to import objects from other major 3D applications.

Supported file formats are:



Sample import from 3DS. Note how mappings are converted (including transparency maps!).

- **VOB:** VUE object file,
- **DXF:** Standard AutoCAD (raw geometry with groups, no textures),
- **OBJ:** Standard Wavefront (raw geometry with groups, texture UV mapping information),
- **3DS:** 3D Studio (raw geometry with groups, texture conversion),
- **LWO:** LightWave 5 (raw geometry with groups, texture conversion),
- **COB:** TrueSpace 5 (raw geometry with groups, texture conversion),
- **RAW:** RAW format
- **FBX:**

Note:

It is possible to import scenes with textured geometries and cameras, but other objects (like lights, characters with skin and bones, morphers, etc...) are not supported. Some materials may not be well imported (it can be caused when the exporter writes data specific to its application only).

xStream

- **3DMF:** 3D Meta File (raw geometry with groups, basic texturing information),
- **WRL:** VRML file format (raw geometry with groups, basic texturing information),
- **PZ3, PZZ:** Poser 4 through 11 animated and static model import filter (raw geom-



etry with groups, texture conversion). Importing Poser content requires that a valid license of Poser 4 and above (5 and above required for dynamic effects) or newer be installed on your computer. When importing native Poser content, the *Poser Import Options* dialog will appear, letting you configure the import/conversion process (see below for details). If you import content from Poser 6 and above, you will have the possibility to [re-pose the models](#) directly inside VUE and render the Poser materials using the Poser shader tree.

- **DAE Collada File Import:** Collada is rapidly becoming the standard for 3D file exchanges. Collada support includes the ability to import fully textured and animated objects. Vertex-based morph targets for animation (e.g. expressions, skin and muscle movement) are supported in VUE. If you are importing Collada animation files from DAZ Studio, you will need the “Animate Plus” plugin for DAZ Studio.
- **ABC:** Alembic format.

Note:

Any file created with any Alembic library < 1.7.3 is compatible with this module(ie, both v1.7.3 Ogawa and v1.0 HDF5 Alembic formats). VUE only imports Polymesh, XForm, and Camera Alembic Data structures. It ignores all others: Nurbs, points, lights, materials, and all specific data types (curvers, subD, faceSet).

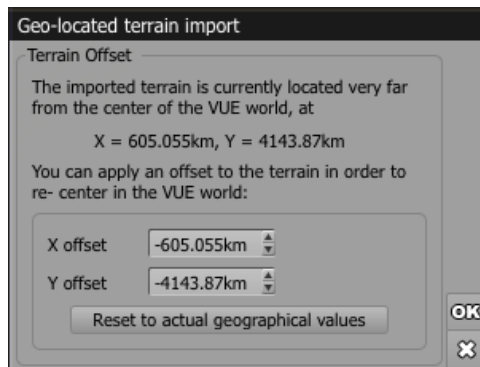
When importing a single object, the first object in an alembic file is selected. If you select to import an entire alembic scene file, it imports all objects in the file and they would import as separate files.

- **GoZ:** Pixologic Zbrush Object format

Note:

Working with ZBrush can be easier with the Link with ZBrush, see [Working with Pixologic ZBrush](#).

- **DEM:** USGS DEM elevation format (.ddf,.dted,.tif). When importing one of these terrain types, a dialog displays to relocate the terrain to a different center, instead of the physical world location.




Geo-Terrain Import dialog

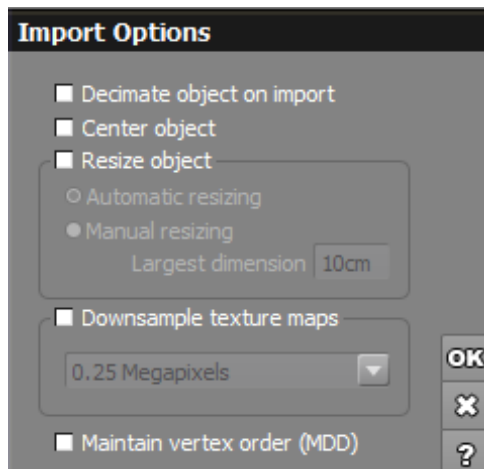
Terrain files are fully imported into VUE to be used as a heightmap on a terrain: be careful not to import files too big for your computer memory. Any one of the freely or commercially available G.I.S. (Geographical Information System) can easily be used to clip and/or resample huge terrain files before importing them into VUE. Each terrain file is imported as a new VUE terrain: adjacent terrain tiles will yield adjacent terrain

objects in your VUE scene.

Both the 3D Studio and LightWave import filters go incredibly far in the conversion of existing materials, converting such complex effects as color, bump, transparency and reflection maps, and even some procedural settings! The picture opposite illustrates the power of these filters; it displays a very nice 3DS fighter model from the Science Fiction Modeling Alliance, rendered in VUE, untouched from conversion.

To import an object from another application, select the menu command **File | Import Object**, or press **File** in the Objects Browser (you display this by pressing the  icon). A Standard File Browser appears that lets you select the file you want to convert. The conversion itself can be quite long (up to several minutes), especially when converting DXF or OBJ objects.

Import Options



Mesh Import Options dialog



When you import an object from another application (**File | Import Object** from the menu), the *Import Options* dialog displays only if you click the **Import Options** button on the object select dialog. If you don't need to display this *Import Options* dialog, just select the object to import and continue. This simple dialog lets you configure the way the object will be imported:

Decimate object on import: when this option is selected, the object will be automatically decimated in order to reduce its polygon count while preserving as much of its original geometry as possible.

Center object: when this option is selected, the imported object will appear at the center of the *3D Views*, regardless of the object position stored in the imported file.

Resize object: when this option is selected, the imported object will be automatically resized according to the following options:

- **Automatic:** the object is resized so that it fills up the viewports,
- **Manual:** the object is resized according to the indicated resize factor.

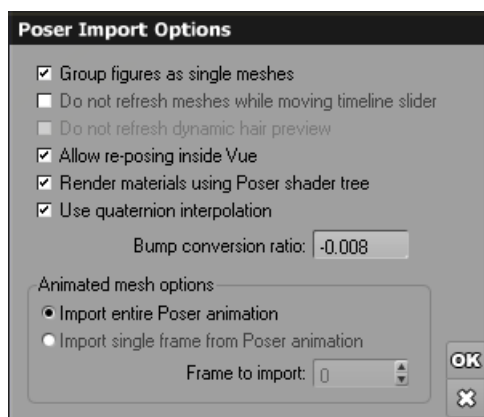
If the resize option is unchecked, the object will be scaled according to the information stored in the imported file.

Largest dimension: When the **Resize object** and **Manual resizing** options are checked, the object is resized so that its largest dimension is equal to the entered value.

Downsize texture maps: Check this option to automatically downsize all texture maps associated with the object you are importing. Overly large texture maps use a lot of computer resources and this is a good way to ensure that all texture maps are a reasonable size. You can select a size in megapixels for all of the texture maps using the dropdown options.



Poser Import Options



Poser Import Options dialog

Whenever you import content in Poser's native file format, the Poser Import Options dialog will appear, letting you configure the way the content is to be handled by VUE. If you are having problems with Poser imports, remember that there are two choices of SDK's when setting up your Poser on the *Options* panel.

- **Group figures as single meshes:** when this option is selected, the numerous parts in the Poser figures will be assembled into a single, multi-material mesh. This is both more efficient in terms of rendering speed and processing. It will also avoid cluttering the interface. However, if accessing the different parts of the mesh is something you need to do, you may elect not to group the figures.
- **Do not refresh meshes while moving timeline slider:** this option is only available if you are importing an animated Poser mesh. When you select this option, the animated mesh geometry won't be updated as you drag the *Timeline*. This will speed up refreshing of the scene as it avoids VUE having to communicate with Poser to update the mesh.
- **Do not refresh dynamic hair preview:** this option is only available when importing an animated Poser mesh that involves dynamic hair. When you select this option, the animated hair geometry won't be updated as you drag the *Timeline*. This will speed up refreshing of the scene as it avoids VUE having to communicate with Poser to update the hair geometry.
- **Allow re-posing inside VUE:** when this option is selected, VUE maintains an open communication socket with Poser in order to support re-posing of the Poser



meshes inside VUE. However, maintaining this open communication socket increases the memory requirements for handling the mesh. If you don't need this feature and want to avoid this overhead, uncheck the option.

- **Render materials using Poser shader tree:** when this option is enabled, the [Poser shader tree](#) will be used for the rendering of all the materials of the imported object. Like re-posing, this requires an open communication socket with Poser that increases the memory requirements for handling the mesh. If you don't need this feature and want to avoid this overhead, uncheck the option.
- **Use quaternion interpolation:** if you notice animation artifacts in your imported Poser meshes, check this option. However, depending on how the character was animated inside Poser, there may be some cases where checking this option affects the character's pose – in which case you should uncheck this option.
- **Bump conversion ratio:** this setting lets you control the default bump map gain that is applied when converting Poser materials to VUE materials. Because of the immense difference in between the render engines of these two applications, there is unfortunately no single setting that will work for all cases. The default value is the one that has been found to provide the best overall conversion – however, certain models may require different conversions ratios, and this setting will avoid having to modify the amplitude of the bump mapping of all the imported Poser materials.

Animation Import

These controls are only available when importing an animated Poser mesh.

- **Import entire animation:** if you select this option, the Poser mesh will be hosted by VUE as an animated mesh.
- **Import single frame from Poser animation:** if you select this option, the Poser mesh will be converted to a static VUE mesh. It will not be animated.
- **Frame to import:** when converting the animated Poser mesh to a static VUE mesh, you can use this setting to select which frame of the animation will be converted into a Vue mesh.
- **Warning:** the **Allow re-posing inside Vue**, **Render materials using Poser shader tree** or **Import entire animation** features require handling of the imported mesh by Poser “inside” VUE. The memory requirements are at least twice as large as for simple imports (could be a lot more, depending on the existence of morph targets and other Poser-specific features). Unless you absolutely need these advanced features, and especially if your scene is already complex or if the imported Poser mesh is heavy, you should consider disabling them. Because Poser will be running “inside” VUE to handle this mesh, consider as a rule of thumb that any scene that is “heavy” in Poser will possibly choke VUE (this limitation does not

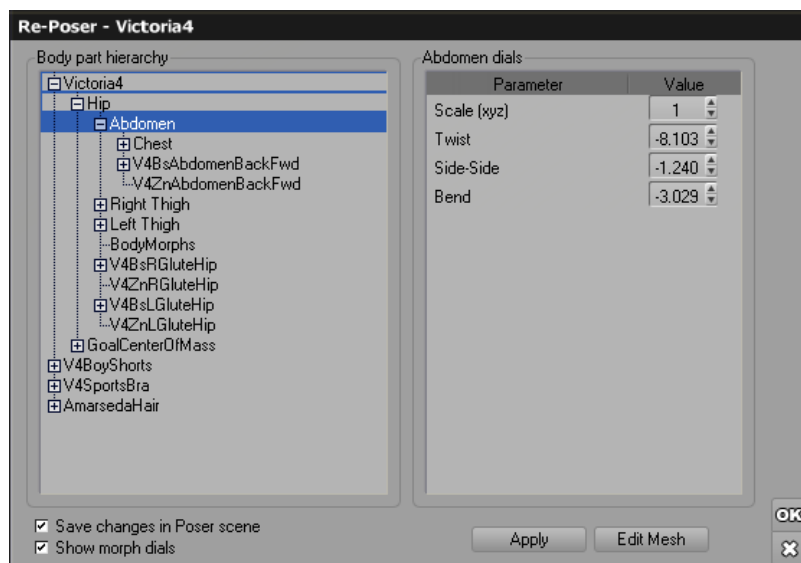


apply to 64 bit systems). As a reminder of the massive overhead required by these features, a warning message will be displayed each time running Poser inside VUE is required.

These features are not available when using VUE xStream in integrated mode (see [here](#)). |


xStream

Re-Posing



Re-Posing Poser meshes inside VUE

Thanks to the *Re-Poser* dialog, you can change the pose of your Poser characters directly inside VUE (this requires a valid license of Poser 6 or better). You open the *Re-Poser* dialog by either:

- Double-clicking on the Poser mesh in the *3D Views* or in the *World Browser*,
- Clicking on the **Edit object** button () on the top application toolbar, or in the *World Browser* toolbar when the Poser mesh is selected,
- Using the menu command **Object | Edit Object** when the Poser mesh is selected.

Note:

if you unchecked the **Allow re-posing inside Vue** option in the *Poser Import Options* dialog, your Poser mesh will be converted to a standard mesh and re-posing will not be possible.

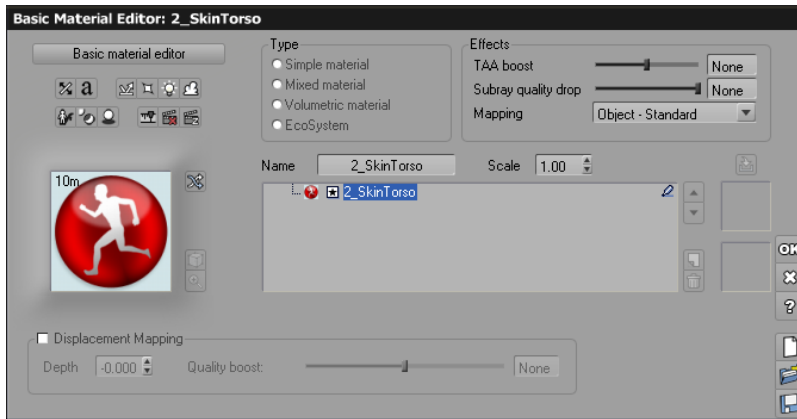


The *Re-Poser* dialog displays a hierarchy of all the body parts found in the Poser mesh, together with the Poser dials that correspond to the selected part.

- **Body part hierarchy:** this list displays a hierarchy of all the body parts in the Poser mesh. Unfold the hierarchy and select a body part to reveal the re-posing dials of that specific body part. If you double-clicked on the Poser mesh inside the *3D Views* to open this editor, the selected body part will be the one you clicked on.
- **‘XXX’ dials:** when a body part is selected in the hierarchy, the dials corresponding to that body part are displayed in this frame. If you modify one of the dials, the body part will be affected accordingly. To view the results of the change, click on the **Apply** button. For full details on the effects of the dials, please refer to the Poser documentation.
- **Apply:** click this button to update the Poser mesh inside the VUE scene.
- **Edit mesh:** click this button to open the standard [Polygon Mesh Options](#) dialog to adjust the underlying mesh options.
- **Save changes in Poser scene:** when this option is checked, the changes you make will be saved in the original Poser scene, so that if you re-open that scene in Poser, the changes you made in VUE will be visible there. If this option is not checked, the Poser scene will not be modified. The changes made to the Poser scene will be stored inside the VUE scene instead.
- **Show morph dials:** when this option is checked, you can access all the morph targets defined for the Poser character directly from within VUE. However, because some characters define a great number of morphs, and you don’t necessarily want to change the morphs of your characters inside VUE (this is possibly something you’d rather do in Poser), unchecking this option will reduce the number of accessible dials as well as the memory overhead required for processing the object.



Rendering Using the Poser Shader Tree



Poser Shader Tree material

In order to ensure that Poser characters imported into VUE look as much as possible the same inside VUE as they did inside Poser, e-on software and e-frontier have developed an advanced bridging technology that enables VUE to use Poser's internal shader tree when rendering Poser materials. This way, Poser materials no longer need to be converted into VUE materials (with the inevitable losses that result from any conversion process).

Such materials that are rendered by VUE using the Poser shader tree are identified by a Poser logo. Note that this feature requires a valid license of Poser 6 or better.

You cannot edit materials that are rendered using the Poser shader tree. You can however easily convert them to an equivalent VUE material: open the *Material Editor*, and select the **Simple material** type. When a Poser material is converted to a VUE material, it becomes fully editable (you may however observe some discrepancies between the way the materials renders in VUE and in Poser). You cannot convert a material back to the Poser shader tree material.

Note:

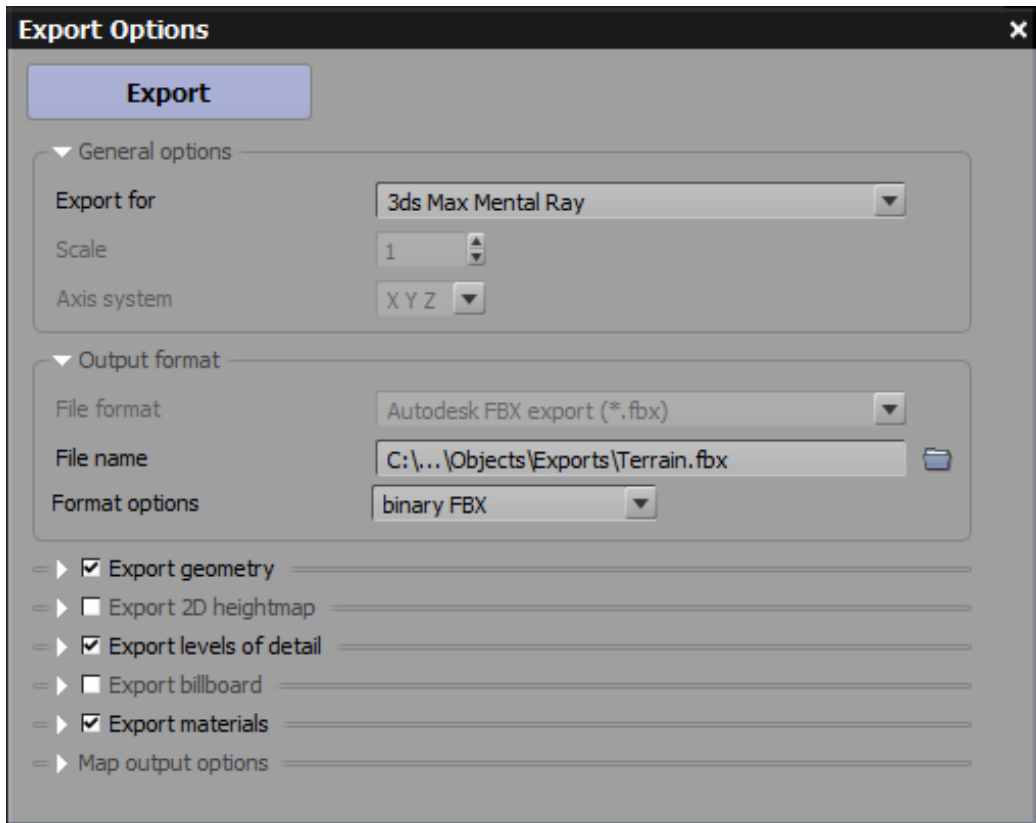
If you unchecked the **Render materials using Poser shader tree** option in the Poser Import Options dialog, the materials of the imported Poser mesh will be converted to equivalent VUE materials, instead of using the Poser shader tree.

Poser shader tree materials are compatible with VUE's displacement mapping. You can adjust the amplitude and the quality of the [displacement mapping](#) using the controls in the lower part of the *Material Editor* when a Poser shader tree material is selected.



Exporting Content

In addition to exporting pictures, VUE lets you export 3 different types of content: objects, atmospheres and entire scenes (Pro versions only). VUE rocks and metablobs can also be exported as objects.



Export Options dialog

The **Export Options** dialog lets you configure the details of the export. This dialog appears, with varying foldable sections, as soon as you want to export one of these elements.

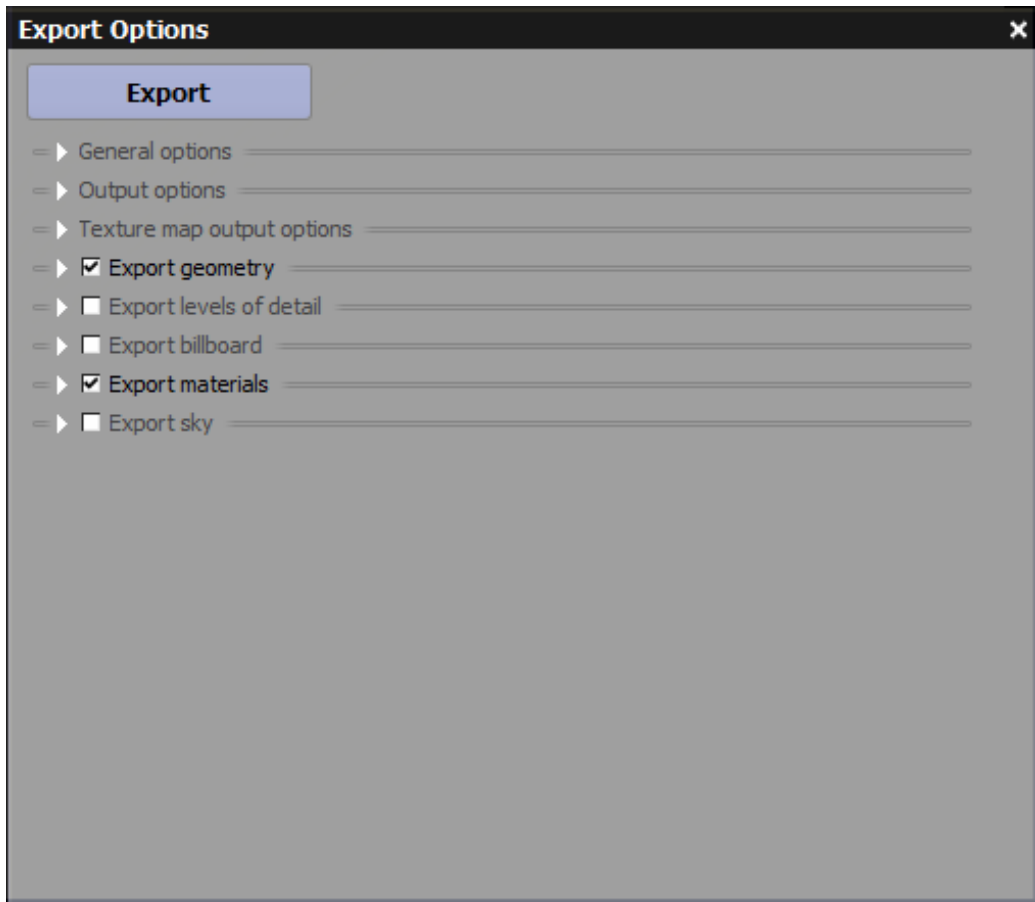
- Exporting Objects
- Exporting Skies
- Exporting Entire Scenes



- VUE xStream allows you to [Export EcoSystems](#) in Alembic or FBX formats.
- In VUE xStream, you can use the [Export Preview mode](#), to preview export directly in the real-time views to let you see what the objects will look like once imported in your target application. In this mode, export is also previewed in renders (render scene preview or main renderer).

xStream

Export Options

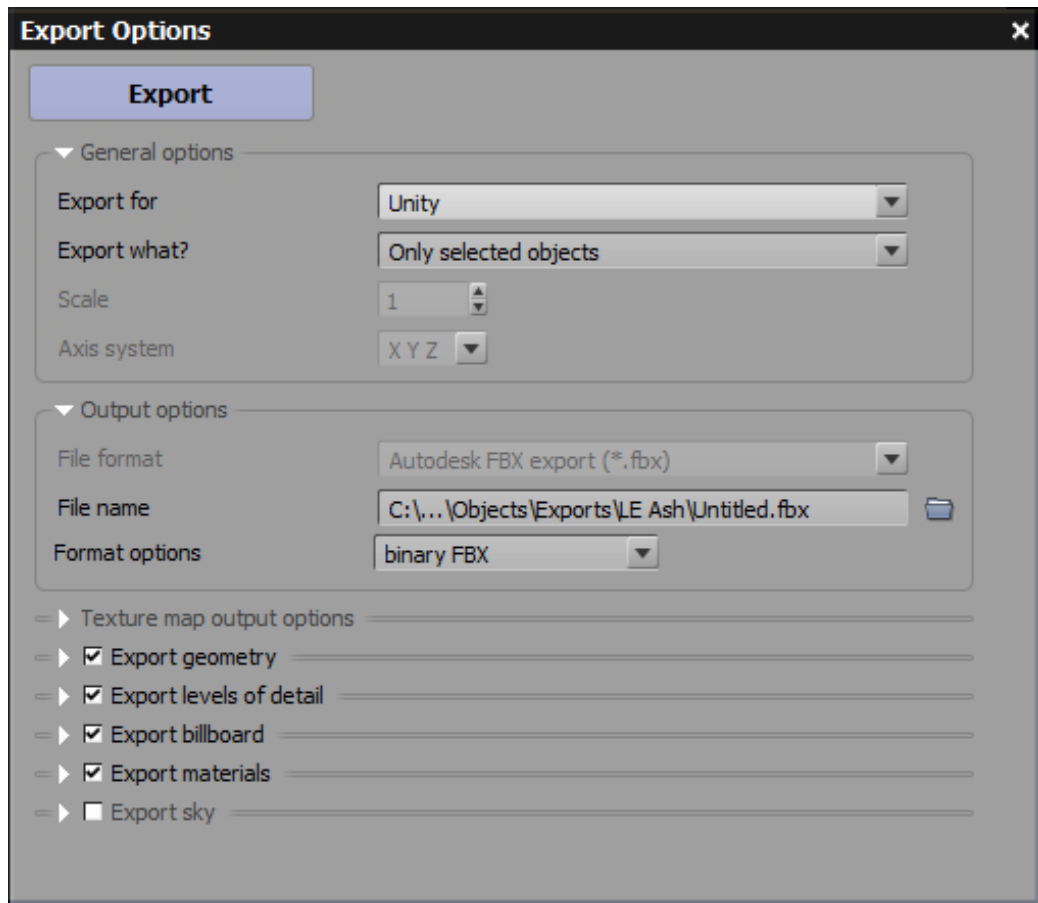


Export Options



This dialog is for configuring export options for objects, skies and entire scenes. This dialog appears, with varying foldable sections, as soon as you want to export one of these elements.

General options



'Output Format'

- **Export for:** Select your target 3D software/renderer from the dropdown list to define best export options. Available presets are:
 - Custom settings
 - 3DS Max Basic, Mental Ray and Vray

xStream...

- Blender
- Cinema 4D
- DAZ Studio
- Lightwave
- Maya Mental Ray and V-Ray
- Modo
- Pixologic ZBrush
- Poser
- SoftImage
- Unity
- Unreal Engine

To manually set all the fields, select **Custom settings**.

- **Export what?:** this option is only available when you are exporting an entire scene, or an infinite terrain. In such case, you can have up to 3 options:
 - **All objects:** (only for scene export) all objects of the scene will be exported
 - **Only selected objects:** only the objects currently selected will be exported (sky is not affected by this option)
 - **Restrict exported objects to zone:** available when exporting a scene, or a procedural terrain. It will let you define a 2D zone in the scene, and only the objects belonging to the zone (or, for a terrain, only the part of the terrain intersected by the zone) will be exported.

Note:

If your scene/object selection contains an infinite terrain, this is the only option, because of course VUE cannot export an infinite geometry in a finite 3D mesh format...

- **Scale:** when you do not select any preset, you have the possibility to apply a global scaling factor to the geometry. It can be useful for formats such as .OBJ, which do not embed any unit.
- **Axis system:** change the axis system for your mesh. Again, it is only available if you have not selected any preset (presets define the correct axis system for the targeted application automatically)



Output format

- **File format:**
 - **.3ds:** 3D Studio
 - **.c4d:** Cinema 4D v5 export
 - **.lwo:** LightWave export
 - **.obj:** Wavefront OBJ export
 - – **.abc:** Alembic export. VUE supports both v1.7 Ogawa export and v1.0 HDF5 export. With VUE, you can export static (non deformable) meshes with animated transformation (pos, rot, scale) in Alembic. You can also export cameras with animated transformation AND animated properties (focal length, focus distance, aperture and aspect ratio).
 - **.goz:** Pixologic ZBrush GoZ export
- – **.fbx:** Autodesk FBX export

Note:

Working with ZBrush can be easier with the Link with ZBrush, see [Working with Pixologic ZBrush](#).

xStream

- **File name:** Enter the name of the exported file. Browse to find the destination directory.
- **Format options:** some formats (Alembic, FBX) have specific format options (FBX exports may be written in binary or ASCII format, typically).

If you are exporting an entire scene, this file will either contain the description of the entire scene (3DS file format) or a list of all the other object files used in the description of the scene (LightWave LWS scene file referencing LWO object files). With alembic format, all files required for the converted scene will automatically be placed in a separate folder.

Export geometry

If enabled, the object will be exported as a 3D mesh file. You can control here how this mesh will be defined.

This group contains a unique setting, called **Resolution**. This setting controls the overall resolution of the objects when they are [converted to polygons](#) before being exported. The higher the setting, the more precise the conversion, but the larger the file and the longer the resulting processing times.

An estimate of the resulting number of polygons is displayed below this setting. Actual values may vary quite significantly.



Export 2D heightmap

When exporting a single terrain (either from the [Terrain Editor](#), or by selecting a terrain in the scene and exporting it), you can export it as a 2D heightmap (picture).


- **File format:** Choose the picture format. Format supporting 16-bit are recommended (TIF 16-bit will be the default). You can edit the format options with the icon.
- **File name:** path of the exported picture file.

Export levels of detail

If enabled, several versions of the geometry will be exported, each one corresponding to a specific LOD.

- **LOD count:** Select the number of versions you want to export.
- **Generate 1 texture per LOD:** check this if you want to export all corresponding resized textures.

Export billboard

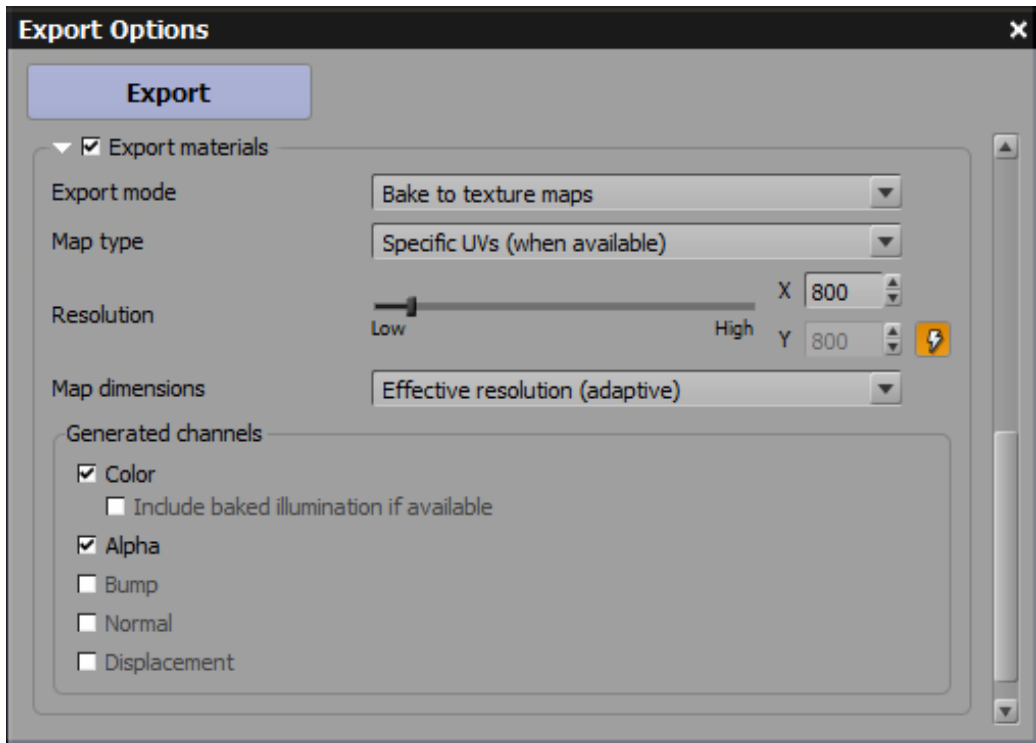


A banana plant (original version and billboard)

If enabled, a 3-plane billboard of the object will be baked and exported. **Resolution** parameter control the resolution of the baked texture. Note that using a high resolution may take a while to compute.



Export materials



'Material Export Options'

If enabled, materials will be written and set up in the exported object.

- **Export Mode:**

- **Automatic:** VUE will choose the best of the 2 options below, depending on the objects being exported.
- **Convert materials (using original texture maps):** This will write textures extracted from the object materials, setting up those in the output object. Define [here](#) how to write maps.
- **Bake to texture maps:** VUE will generate UV coordinates and bake textures to represent the materials. This mode can be used to export procedural materials. Beware it can take a long time to generate those maps.

Different types of map are available:

- * **Specific UVs (when available):** Produces UVs which are generally faster to compute and easier to edit



- * **Generic UVs:** Produces less distortion, but can take a very long time to generate
- * **Per polygon maps (Ptex):** Exports the material as a **Ptex** texture: draw a small texture for each polygon.
 - **Resolution:** The slider sets the number of pixel per unit. A higher pixel per unit gives an export with a better resolution.

xStream

- **Resolution:** Use the slider to set the quality of the material maps.
- **Automatic aspect ratio:** Check to enforce automatic aspect ratio for the maps.
- **Actual Dimensions:** Select to use the actual dimensions. All texture maps will be generated at the exact resolution indicated.
- **Effective resolution:** Select to use effective resolution. This is particularly useful when exporting entire scenes: in order to generate the texture maps, the 3D geometry of the objects has to be “unwrapped” onto a flat plane. But, because the geometry of the converted objects can be arbitrarily complex, some parts of the texture map may be unused; so the resolution of the parts of the texture map that are actually used may be less than the selected resolution. This can result in inconsistencies in the texturing resolution on different objects. *Effective resolution* attempts to generate texture maps where the resolution of the texture once it has been mapped on the supporting geometry is roughly the one indicated in the *Resolution* setting, and is more or less constant over all objects. Although this mode results in texture maps of varying resolutions, it is the one that will result in the most consistent texture mapping.
- **Generated channels:** Select any of the following to generate.
 - Color
 - Bump
 - Alpha
 - Normal
 - Displacement
 - Include baked illumination, if available

Map Output Options

You can tune here some options to make sure exported maps will work out of the box in your preferred 3D software.

- **Place textures in sub-folder:** places exported texture maps in sub-folder “Maps” of export folder.



- **Relative paths:** check to use relative paths for texture maps.
- **Invert alpha:** check to invert alpha.
- **Invert normal maps green:** check to invert green on normal maps (useful to switch between conventions).
- **Alpha treatment:** Select here how to export alpha :
 - **Unchanged:** this will export maps as you have defined them in the material editor.
 - **Extract:** this will export a separate alpha map.
 - **Merge:** this will export alpha as embed alpha channel of color map (thus merging into an RGBA map), this will try to export the texture in the same format the color map was originally.
- **Fallback merge format:** If **merge** alpha option is checked, chose here fallback format for the merged texture, in case original format does not support transparency. There are several options here : EXR, PNG, PSD, TGA, or TIF.
- **Apply gamma:** check to apply gamma settings used in PlantFactory to exported map.
- **File format (converted maps):** There are several options for the texture map: unchanged, BMP, EXR, GIF, HDR, IFF, JPG, PCX, PNG, PSD, TGA, or TIF. Unchanged option will keep texture maps original format.

Rigging and animation

This control is only available when exporting an object or an entire scene that contains animation, and when export is done as 3DS, LightWave or Alembic files. Check the box to include the animation in your export.

This is not available when exporting a sky.

Exporting Animated PlantFactory Plants

It is now possible to export a PlantFactory plant that is animated. The animation would be bone or point-level animation.

Export of PlantFactory animations requires a license of PlantFactory Producer.

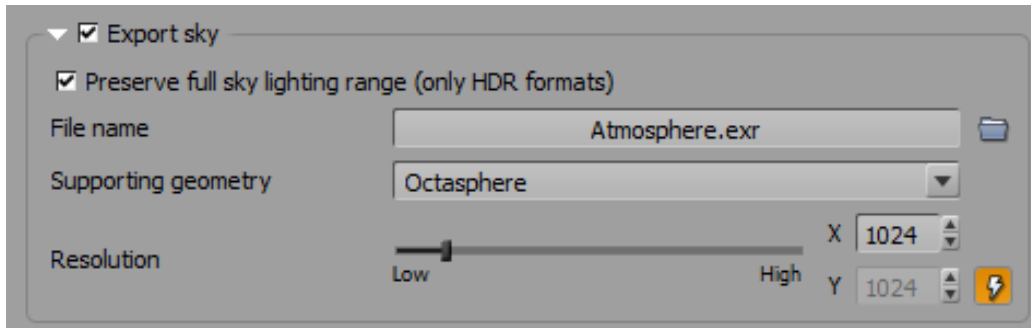
Only C4D, LWO and FBX formats give you the option of exporting rigging and animation. These options do not display for other format exports.

- **Animation data:** there are three options:
 - **Static Bones:** bones but no animation
 - **Animated Bones:** bones and animation



- **Animated Points:** points and animation
- **Duration:** length of the animation in seconds
- **Framerate:** frames per second
- **Cyclic animation:** Check to force the baked animation to be cyclic

Export sky



'Export Sky'

The controls in this group are only available when exporting a sky or an entire scene.

VUE can generate a backdrop of the sky in the scene for mapping on a variety of primitives. The result will be a stand-alone picture of the sky. To mimic VUE skies in other applications, this picture should be mapped on an appropriate object and placed around the scene.



Note:

This is just an approximation of VUE skies, as VUE skies are the result of complex volumetric interactions between objects and the surrounding atmosphere.


- **Export sky:** Be sure this option is selected when exporting the sky.
- **Preserve full sky lighting range:** This is only available when using HDR format for the sky file.
- **File name:** Enter the file name. All [supported picture file formats](#) can be used for exporting the sky. Use the **Browse** button to find the folder for the output file.
- **Supporting geometry:** Select the output geometry. Choices are: Cube, Octahedron, UV Sphere, Octasphere. See images below.
- **X, Y:** Enter the file size.


Here are the different types of output geometry available:





***Cube:** this is the simplest supporting geometry that will create a standard sky box for use e.g. in games. It presents important texture distortions in the corners.*







Octahedron: this is a simplified version of the sphere that represents an interesting compromise between supporting geometry complexity and texture map distortions.



UV Sphere: *this is a standard UV-mapped sphere. Because of the way the texture is mapped, some points on the sphere (near the top and bottom) will have a higher resolution than in the middle.*



Octasphere: *this is also a sphere, but this time it is mapped in such a way that the surface mapped by each pixel in the texture is more or less constant over the entire sphere. This is the best solution, although the resulting texture map may be hard to understand at first glance.*



Exporting Objects



A PlantFactory tree in Cinema 4D

When you decide to export an object from VUE, it will first be converted into data that can be processed by other 3D applications.

The geometry of the object will be converted to a polygon mesh, and materials will be converted to texture maps.

Conversion of the geometry is very similar to [Baking to polygons](#), except that the exported object in the VUE scene is not directly affected.

Also, on top of converting the object to polygons, VUE can also generate UV mapping information and converts materials into texture maps. You can learn more about material export [here](#)

To export an object, first select the object to be exported, and then select the menu command **File | Export Object**,

You can also do it from the Export icon, on the bottom of the World Browser.

The **Export Options** dialog will appear, letting you configure the export options for the selected object.

Note:

some objects may not be exportable, because the creator of the object does not want you to export them. You can, yourself, prevent other users from exporting your own models



by clicking **Forbid export** in the [Polygon Mesh Options](#) dialog.

Exporting Entire Scenes

On top of exporting independent objects and skies, VUE can also export an entire scene. Exporting scenes is a complex process that involves converting all objects in the scene into polygons (like when exporting independent objects), generating all corresponding texture maps as well as a sky preview, converting camera and lighting information and then saving all these elements to disk.

To export a scene, select the menu command **File | Export Entire Scene**.

The **Export Options** dialog will appear, letting you configure the export options for the scene (only relevant options will be enabled).

Note:

Some objects may not be exportable, because the creator of the object does not want you to export them. You can prevent other users from exporting your own models by clicking **Forbid export** in the [Polygon Mesh Options](#) dialog. Because authors could not expressly authorize or forbid object exports in VUE Esprit, no object imported from VUE Esprit can be exported.

Supported file formats for exporting scenes are:

- **3DS:** 3D Studio
- **LWS:** LightWave
- **DAE:** Collada
- **FBX:**

Note:

Cameras will also be exported into the FBX file. EcoSystems will be exported as FBX native instances.

xStream

- **ABC:** Alembic

Note:


All files required for the converted scene will automatically be placed in a separate folder. Cameras will also be exported into the Alembic file. EcoSystems will be exported as Alembic native instances (EcoSystems export requires xStream). As Alembic doesn't support sky map internally, you won't be able to export sky maps



while exporting the entire scene in Alembic.

In VUE xStream, scenes containing EcoSystems can be exported. You can configure how EcoSystem specimens are exported by right-clicking on the EcoSystem in the World Browser.

EcoSystem Export



'Exporting EcoSystems'

You can export EcoSystems in Alembic and FBX format.

If the object(s) selected for export has Ecosystem instances, the instances will be exported along with the object(s) provided you export it in .abc or .fbx format. The materials will not be exported in the case of .abc export, but will in .fbx export.

By default, EcoSystem instances are exported with global export options, but you can configure how EcoSystem instances are exported by right-clicking on them in the World Browser, Classes tab.



Limitations

For EcoSystem instances, export can only be previewed at render time (in Main Camera view or render scene preview) and not in real-time views.

Export Preview

xStream

In **Export Preview** mode, export is previewed directly in the real-time views to let you see what the objects will look like once imported in your target application. In this mode, export is also previewed in Xstream renders (render scene preview or main renderer).

To activate **Export Preview** mode, there are 2 different paths:



Export Preview Icon

- **From main menu:** *Display > Preview Export*. This can be linked to a shortcut like all menu commands.
- **Using the icon in the top toolbar:** If you have not set custom export options for any object of the scene, VUE will offer to apply default settings to all exportable objects.

In **Export Preview** mode, you cannot edit objects nor materials in their regular editor. You can only move/rotate/scale objects. You must switch VUE to normal mode again before editing objects/materials.

If you modify objects / materials and activate “Export Preview” again, VUE will offer to update the pre-baked export data for preview.

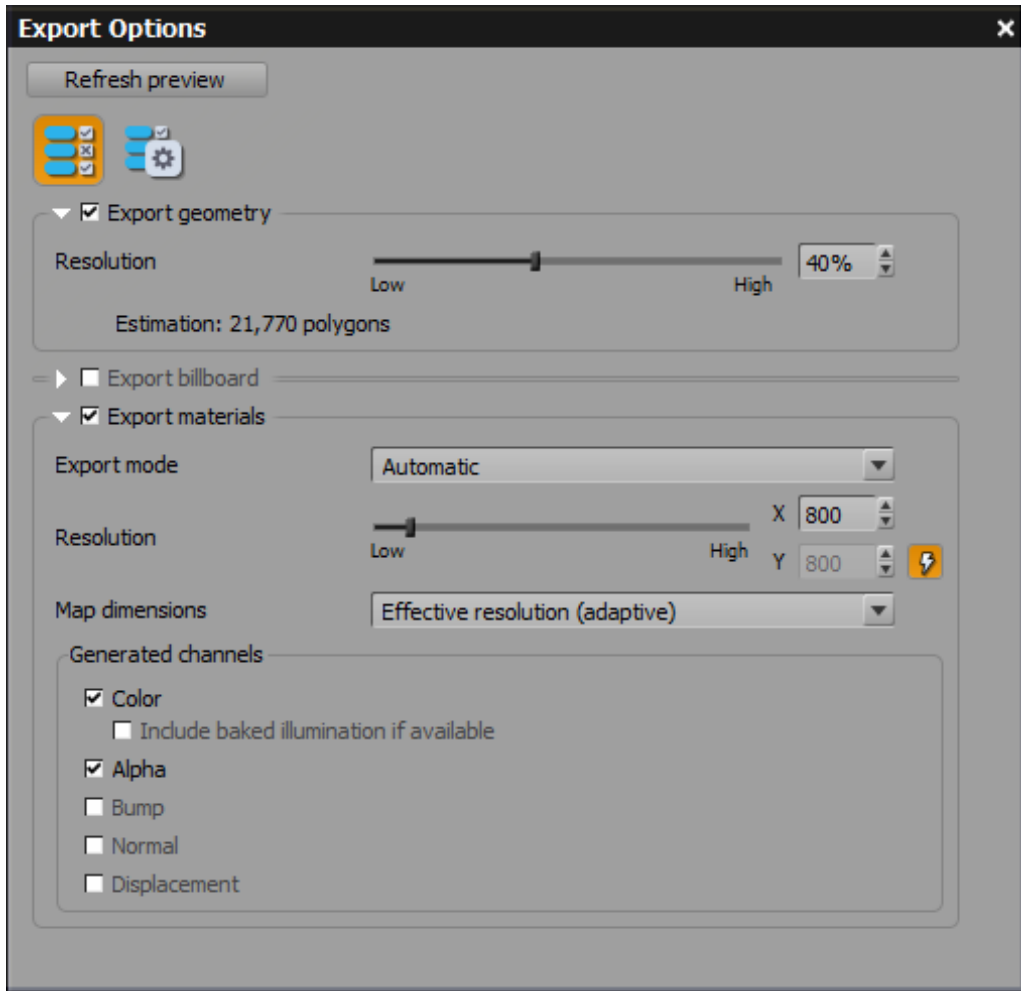
You can optionally save pre-baked export data with your scene.

xStream...



Export Objects on a Per-Object Basis

...xStream...



'Export Preview – Options'

You can change the export settings on a per-object basis. This is useful, for example, if only one object uses transparency in your scene, and therefore you don't need to generate alpha maps for all other objects. It also lets you boost the export quality (geometry or maps) for a specific object.

Select the object for which you want to change export options and right-click the export



icon (bottom of the World Browser). Or, you can access this from the VUE menu.

In **Export Preview** mode, when you double-click on an object, the *Export Options* dialog will appear, letting you adjust the object's specific options.

- **Use custom options:** you can toggle from specific to global options without losing the specific options you have defined for the object.
- **Copy from global options:** you can also choose to override the specific options with the global (per-scene) options.

If you are currently in Export Preview mode, you can click the **Refresh preview** button to see the result of your changes in the 3D views.

...xStream

PTex Export Preview

Unlike the other material export modes, which can be previewed in both the OpenGL real-time views and the renders (including PathTracer and scene preview), PTex can only be previewed in ray-traced renders (including scene preview).

xStream




Section 4

The Editors





Terrain Editor



The Terrain Editor: a complete set of tools to model terrains

Terrains are the primary building blocks of landscape imagery. Together with plants, they give VUE incredible modeling power for natural phenomena.

When you create a new terrain, a mountain is automatically generated using random fractal algorithms. These algorithms guarantee that no two mountains will ever have exactly the same shape (unless they use the same *seed*).

However, terrains can be made to capture other shapes than this basic shape. This is done by using the *Terrain Editor*. The *Terrain Editor* provides a set of powerful tools designed to let you easily model realistic terrains.



Terrain Types

Terrains come in four flavors:

- Heightfield terrains,
- Procedural terrains,
- Spherical terrains,
- Procedural terrains that can also be changed to Infinite terrains.

Both heightfield and procedural terrains can be made **Symmetrical** and/or **Skin only**. Symmetrical terrains are replicated negatively below their lower clipping altitude. Skin only terrains don't have any edges or flat bottom. This is useful to create thin elaborate surfaces, such as a flag.

Heightfield Terrains

Heightfield terrains use a fixed resolution grid to define the geometry of the terrain. They are the most basic and easy to use type of terrain. They are also the fastest to render.

Each heightfield terrain (previously called standard terrains) is now driven by a graph which is computed over the heightfield to make the final terrain. Each modification on the graph should trigger a re-computation of the terrain. All other modification (painting, old-style terrain effects, etc) are still possible and appear inside a **User Touch-Up** graph node. This is the last node connect before the output. If such node doesn't exist at this location, one is created. This node can be manipulated by the user in the graph to apply graph effects on user painting. However, additional painting can only modify the **User Touch-up** node; no additional nodes are added.

In order to speed up previewing, two low-resolution of the heightfield terrain are computed while working on the *Terrain Editor* : a 256*256 (or the terrain's resolution if the final terrain has a lesser resolution) and 512*512 (if the terrain is larger than 256*256); matches the terrain's resolution if the terrain's resolution is 512*512 or less.

Heightfield terrains allow the user to define custom dependencies and to work with them as usual (even if they are computed together with the terrain).

For maximum compatibility, the heightfield terrain uses the same properties as the procedural terrains (scale, filter, extension (automatic or forced), vertical gain, zero edge, etc).




Procedural Terrains

As opposed to heightfield terrains, procedural terrains are able to adapt their level of detail dynamically, so as to always ensure the same amount of detail in the geometry of the terrain, whatever the viewing conditions. You can zoom in onto procedural terrains indefinitely, you will always see new levels of detail – smaller and smaller details as you keep zooming in. This is made possible thanks to the fact that the surface of the terrain is defined by a complex fractal procedure – hence the name. As you can surely appreciate, the processing of a procedural terrain is much more complex than that of a heightfield terrain. Luckily, VUE implements advanced algorithms that ensure all this takes place with minimum memory requirements and maximum efficiency.

Procedural Terrain Presets

The downside to the power of these infinitely detailed procedural terrains is that setting them up can be tricky and time consuming. This is why VUE ships with a library of procedural terrain presets, with their associated materials.


To load a procedural terrain preset, click on the **Load Procedural Terrain Preset** icon () . This opens a Visual Browser, letting you select the preset you want to use.

Procedural terrain presets come under two categories: *Procedural terrains* and *Infinite terrains*. The *Procedural terrain* presets will create a standard procedural terrain of typically the same size as regular (non-procedural) terrains. The *Infinite terrain* presets create procedural terrains that are truly infinite. As you move your camera, the terrain will change but it will also move with the camera so you will not reach the edge of the terrain.

Select the style of terrain you want to create, and VUE will generate an each-time-different terrain of the corresponding style, complete with all associated materials.

If you create an infinite terrain, VUE will ask if you want to replace the ground plane, as the infinite terrain will typically become the new ground. Likewise, VUE will ask if you wish to replace any pre-existing infinite terrain in the scene.

Hint: if you don't like the infinite procedural terrain that was created from a preset, select that preset again to replace the terrain with a new variation.

You can create your own procedural terrain presets by clicking the **Save** icon () in the *Terrain Editor*. This will save the procedural terrain definition, together with any assigned materials, and add the required elements to allow randomization of the terrain each time you create a new one.



Touching Up Procedural Terrains

Unfortunately, the big drawback to procedural terrains (besides being slower to render than heightfield terrains, because of the necessary computation of sub-polygon detail everywhere) is that they are a lot trickier to use than heightfield terrains. That's because the altitude of a procedural terrain is defined using a mathematical function that is evaluated at each point. The result of that mathematical function is the altitude of the terrain at that point.

Usually, if you want to modify the shape of a procedural terrain, you have to delve into the mathematical intricacies of the altitude function. While this can be a lot of fun if you've got a strong mathematical background, or if you're of an adventurous type, it's not necessarily the kind of experience everybody wants when on a deadline!

That's why e-on software has developed a new technology to facilitate the edition of procedural terrains. Basically, the way it works is that you can use all the tools in the *Terrain Editor* to “touch up” the look of your procedural terrain. VUE uses these modifications to adjust the output of the altitude function. Polygonal artifacts are avoided by using powerful interpolation algorithms. As a result, you get the benefits of the virtually infinite terrain resolution of procedural terrains, with the ease of modification of heightfield terrains; customizing a procedural terrain becomes as easy as customizing a heightfield terrain. The user interface is extremely straightforward since, for the user, there is no difference in between editing a heightfield terrain and editing a procedural terrain. If you modify the altitude function later, the modifications will be applied to the new function. This is particularly useful when you want to add surface detail to the terrain without losing the overall shape.

The terrain resolution that is displayed below the top toolbar refers to the resolution of the touch-up data. Because of the interpolation technology, it is not useful to use large resolutions with procedural terrains.

Hint: For optimal results when touching up a procedural terrain, you should avoid painting small details and sharp/steep edges, but instead add gradual changes.

Of course, you can also modify the mathematical function, either by loading one of the many preset altitude functions, or by editing the function yourself using the [Function Graph](#).

Mapping Modes

The behavior of a procedural terrain depends on the mapping mode (refer [here](#) for details on the different mapping modes available in VUE) used for the terrain. By default, the **Object-Parametric** mapping mode is used. In this mode, the geometry of the terrain is not affected by resizing or moving.



However, if the terrain is mapped in **World-Standard** coordinates, the geometry of the terrain will change as you move the terrain about. In this mapping mode, the terrain should be understood as a window observing a particular area of the procedural altitude function. If you move that window, you see other parts of the function. But if you return to the initial location, the same part of the function will be observed and hence the geometry of the terrain will still be the same. If you enlarge the procedural terrain in the *3D Views*, you will be observing a larger area of the function: the features in the terrain won't be any larger, you'll just see more features. You can enlarge the terrain until it stretches up to the horizon, thus recreating the surface of an entire planet.

The other mapping modes represent different combinations of these two behaviors. For instance, with the **Object-Standard** mapping mode, moving/rotating the terrain won't change the terrain geometry (as in **Object-Parametric**), but resizing it will show more of the terrain (as in **World-Standard**).

Changing Mapping Mode

You change the terrain mapping mode using the drop-down list in the [Procedural Altitudes](#) tab. When you change the mapping mode, it is the coordinate system of the altitude function that is changed, resulting in a modification of the terrain geometry.

To avoid changing the geometry of the procedural terrain when changing the mapping mode, VUE will offer to add some nodes to the altitude function in order to preserve the shape of the terrain. You can examine these nodes in the [Function Scale](#). From then on, however, moving or resizing the terrain will be subject to the standard behavior of the new mapping mode.

Note:

You should avoid repeatedly changing the mapping mode of a procedural terrain, as new nodes will be added to the altitude function each time, and the resulting function graph may end up being uselessly complex and slow.



Spherical Terrains

VUE can create scenes in which all infinite planes and all infinite parametric terrains are spherical. Once you define a scene as being spherical, the existing and added infinite planes and terrains automatically assume that shape.

There are two kinds of spherical scenes:

- The basic spherical scene.
- The planet spherical scene.

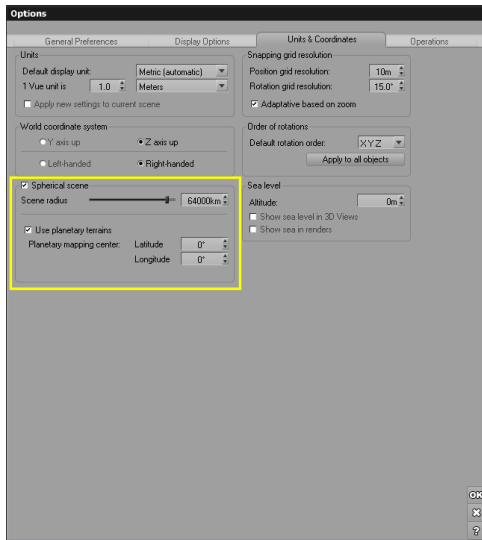
In a basic spherical scene, the scene is limited to a piece of a sphere. With this type of scene, terrains have the same look as a flat terrain with the same altitude function when viewed closer to ground level. In addition, you also have a mid-range view of a planet, as from a lower altitude orbit. This mode is more limited but allows you to have a spherical terrain with the same look as the flat associated terrain.

In a planet spherical scene, you have a whole planet drifting in space. However, the spherical terrains are a bit different from flat infinite terrains in that the altitude function is evaluated in three dimensional space to maintain continuity on the whole planet.

In both modes, the center of the world is set at the position (0,0,-radius). This means that the zero of the scene is at the “north pole” of the planet.



Creating and Manipulating a Spherical Terrain



Options dialog – Creating a Spherical Terrain

To create a spherical terrain, enable the spherical scene option on the **Units & Coordinates** tab in the *Options* dialog. If you are creating a planet, you also have the option of setting the scene radius, which is really the size of the planet.

Basic Spherical Scene


In the basic spherical scene you are working with a curved portion of a planet. Atmospheres and object placement follow the same rules as for infinite procedural terrains.

The easiest way to create a basic spherical scene is to open a new scene and add an infinite terrain. Then, on the *Options* panel, **Units & Coordinates** tab, select to create a spherical scene and enter the size.

Now, you can landscape your terrain just as you would an infinite terrain. Keep in mind that you will probably be viewing this terrain from a greater height than you normally would view a terrain.



Planet Spherical Scenes



Planet Spherical Scene

The planet spherical scene takes all of your infinite planes and creates an entire sphere, or a planet. This includes water planes and cloud layers, so you are truly simulating a planet. If you have an empty scene with just an atmosphere and a ground plane, the ground plane becomes the sphere. If you would then add an infinite procedural terrain to the scene, this infinite terrain replaces the ground plane and becomes the sphere. Of course, the terrain geometry will appear differently than it would as a flat infinite terrain. To increase the height of the terrain for a planet spherical terrain, use the numeric Z position field in the *Object Properties* panel. Please read the tutorial on [Creating a Planet](#) for an illustration of this topic.

Spherical terrains have the same properties as an infinite procedural terrain. However, when a spherical terrain is moved on Z-axis, its radius is also increased to keep the terrain attributes consistent.

In planet mode, the global translation gizmo has two modes. When you work on the whole planet, it can be easier to move the objects along the latitude and longitude axes of the planet. Therefore, a button has been added near the gizmo to switch to latitude-longitude mode instead of X-Y mode. This option is also available in the display menu (**Display | Gizmos | Spherical Coordinates**).



Planetary Mapping Mode

The planetary projected texture node allows you to use a planetary projected picture (Mercator projection) on a planetary terrain.

When added in a *Function Graph*, this node returns the value of the picture according to the position of the point on the planet. The projection of the map on the planet can be offset using the **Latitude** and **Longitude** parameters. These parameters give the position on the map of the point (0,0,0) in the scene.

The **Planetary mapping center** option can be found in the *Options* panel, **Units & Coordinates** tab.

Objects in Spherical Scenes

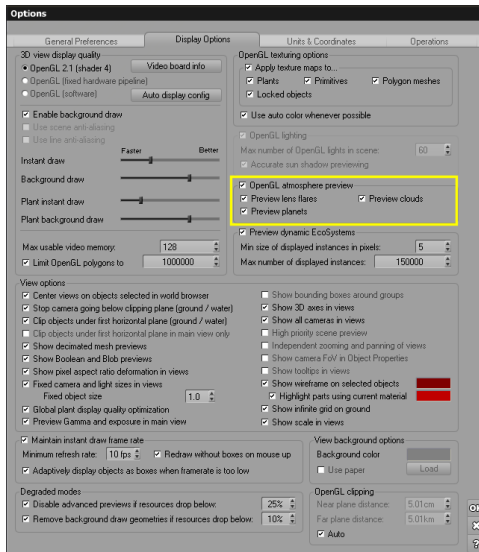
To make it easier to texture and manipulate objects, you can change the display mode from the **Display** menu, **Spherical Display** option. With this option unchecked, you can display your spherical terrain as if it was flat in the OpenGL views. This helps with object placement. The rendering remains spherical. When **Spherical Display** is checked, you will see the spherical terrain in all OpenGL views.

When you switch a scene from flat to spherical, the positions of all objects currently in the scene are moved to keep the same altitude and orientation from the ground. If using a group of objects, all the objects inside the group keep their relative position to the group. Only the position of the group is changed. However, if the group is ungrouped, the objects will move to follow the curve of the ground.

If you create a group in a spherical scene, the position of an object within the group will be the same as before being grouped. And, the object will keep its position if ungrouped.



Atmospheres in Spherical Scenes



Options dialog – Atmosphere preview

The atmosphere for spherical scenes includes the atmosphere over the planet (or part of the planet or curved terrain) and includes the dark areas of space.

The atmosphere can be previewed with or without clouds. This option is set on the *Options* panel, **Display** tab, under the **OpenGL atmosphere preview**.

Selecting and previewing cloud options is very resources intensive. This can cause some lags and unresponsiveness when editing your scene. Some atmospheres can need more than two minutes to reach their final resolution.


At any time, you can force the refresh of the atmosphere preview using the “refresh sky” option in the main view drop down menu. Be aware that moving the camera position doesn’t refresh the preview. When moving the camera too far away from the last refresh, the atmosphere will fade to the background to signal that the preview is invalid.

Note:

If you go higher in altitude than the radius of the planet, the preview of the sky will be deep black.



Editing Terrains



'Terrain Editor – Heightfield Terrain'

To access the *Terrain Editor*, either:

- Double-click on the terrain you want to edit in the *3D Views* or in the *World Browser*,
- Click on the **Edit object** button (⚙️) on the top toolbar, when the terrain is selected,
- Use the menu command **Object | Edit object**.

The Terrain Editor documentation will cover these topics in separate articles:

[Top Toolbar](#)



Predefined Terrain Styles

Sculpting Terrains

Brush Editor

Painting Materials

Procedural Altitudes

Zones

Effects

Exporting Terrain Geometry

Handling the 3D Terrain View

The terrain you are working on is displayed as a 3D map in the middle of the editor. This map is generated using e-on's unique *Solid3D™* real-time technology. It shows terrain altitudes by coloring the map with a gradation that depends on altitude. The colors of the map can be modified by double-clicking on the **Altitudes** below the map.

If you move the mouse over the terrain, you will see a pointer surrounded by a sphere that follows the mouse. This is the brush, and it is used to [apply local modifications to the terrain](#).



You can move, rotate and zoom the view using the *Pan*, *Orbit Around Selection* and *Zoom* triggers defined in the [shortcuts](#) list of the *Operations* panel.

Note that the trigger modifier *Slow Camera Controls* can be handy when panning close to the terrain.

Panning moves the view to the left or right as well as up or down (*vertical* panning). By pressing the *Space* bar, moving the mouse up/down will instead move the terrain forward/backward in the view (*horizontal* panning). Moving the view also moves the




rotation center, around which the view is rotated when using the *Orbit Around Selection* trigger. You can also relocate the rotation center without moving the view by Shift + double-clicking directly on the terrain.

You can also zoom in and out of the terrain using the **Zoom** icons ( and ). Note that zooming does not change the terrain resolution.









The triggers used are the ones defined in the *Triggers in Editor views* category. If one of the *Pan* or *Zoom* triggers are not defined in this category, the corresponding triggers from the *top-most* category are used.

If you moved or zoomed the view too much or simply want to frame the terrain again in

the view, use the **Reset view** button  in the top toolbar.

If you are more familiar with sculpting on a uniformly colored mesh or find the colors used on the terrain distracting, you can easily change the gradation color map by double-clicking on the **Altitudes** bar (below the terrain) and selecting a different color map to use.

Terrain Editor Top Toolbar

- 
New terrain: creates a new terrain in the Editor window.
- 
Save terrain: saves the terrain in .prt (VUE terrain) format.
- 
Export terrain: exports the terrain in various object formats.
- 
Terrain options: allows you to convert the current terrain to a different type of terrain:
 - **Symmetrical:**
 - **Skin Only:**
 - **Infinite:** for Procedural terrain only
 - **Procedural:** for Heightfield terrain only
 - **Heightfield:** for Procedural terrain only
- 
Reset View: resets your view to what you were using previously.
- 
Top view: changes the view from a perspective view to a view from the top (i.e.



as if your terrain were seen from an airplane or satellite). Dragging the view with the right mouse button (Ctrl drag on Mac) will rotate the map.



Show entire scene: displays the terrain with the current camera position as well as all objects and plants placed in the scene. It toggles back to the original view.



Show wireframe: shows the terrain as wireframe and toggles back to original view. This works best when zoomed in on a terrain.



Show specular: gives you a shiny surface to enhance your 3D perception of the terrain when in sculpting mode.



Show Texture Maps: displays a bitmap applied as a texture to a terrain in OpenGL render quality. The feature is particularly useful if you want to use a bitmap as a reference when sculpting.



Show clipping plane: displays the entire terrain area including what has been clipped.



Copy and Paste: exchange your terrain data with standard bitmap applications. You can Copy your terrain to the clipboard, then paste it into your favorite 2D app, modify it as you like, then copy it back to the clipboard, and Paste it into VUE.

Note:

The data copied to the clipboard is limited to 8 bit resolution, which is far less than the resolution of VUE terrains. This function copies only a heightmap, so no 3D displacement information can be included using this method.



Zoom: allow you to enlarge or reduce the terrain view in the editor. This does not change the actual size of the terrain; it just allows you more control when working on the terrain.

The **Terrain Resolution** is handled using the following icons:




Halve terrain resolution: to halve the resolution of the terrain.





Double terrain resolution: to double the resolution of the terrain.

A resolution of 512×512 will yield a very detailed terrain surface. 1024×1024 is a massive resolution that should only be used when an extremely detailed terrain surface is to be seen from close up. Please understand that such a terrain involves over 2 million polygons. Few 3D packages would even survive this. Making even larger terrains is possible but is usually unnecessary.



Resample Terrain dialog



Resample Terrain: is used to resample the terrain directly to any resolution. The *Resample Terrain* dialog pops-up, letting you select the new terrain resolution. This changes terrain resolution without actually changing terrain size in the scene.

Note:


The current terrain resolution is indicated just below that group of icons. The default terrain resolution is 256×256 .




Equalize: resamples the altitudes in the terrain so that they range from 0 to the highest standard altitude, which is 100. You can use this button when you have raised the terrain above the standard range. This option is only available for height-field terrains.



Invert: inverts all the altitudes of the terrain, making lower altitudes high, and vice-versa. In the case of procedural terrains, this function inverts the altitude filter.



The Altitude Filtering dialog lets you visually modify the profile of your terrains

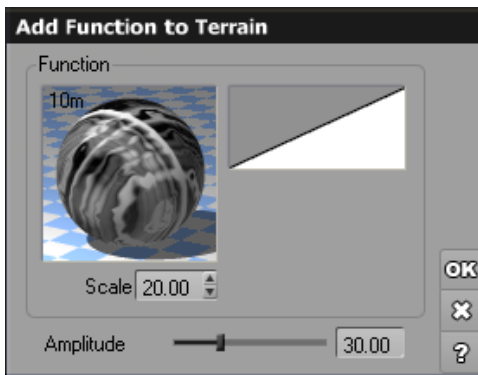
The **Filter Altitudes** icon () accesses a powerful feature that acts like the **Altitude distribution** filter of the fractal options. It lets you reorganize the altitudes of your terrain using a filter.

This option is not available when editing a procedural terrain, because it is superseded by the altitude function filter of the **Procedural Altitudes** tab.


Modifying the shape of the filter will modify the altitudes accordingly (e.g. adding ridges to the filter will create ridges in the resulting terrain). To edit the filter, Control click on it. This will open the *Filter Editor*.

Notice how the modifications in the filter are displayed in real time in the terrain preview. If you want to see a larger preview of the modified terrain, press **Preview**. This will apply the filter to the large terrain map in the *Terrain Editor*.

A special collection of filters is available for terrain altitude filtering (Select the Terrain profiles collection in the Visual Filter Browser).



Add Functions dialog

The **Add Function to Terrain** icon () lets you create terrains from any arbitrary function. Obviously, this option is not available for procedural terrains, as it is somewhat similar to the whole concept of procedural terrains. In the case of a heightfield terrain, the procedural altitudes are directly “baked” into the heightfield terrain’s altitude map. This means that you can use all of the elaborate tools in the *Function Graph* to customize every detail of your terrain. The “Dunes” predefined terrain style uses such a function.

To define the function that should be added to the terrain, Control-click the picture of the function to open the *Function Graph*.





To the right of the function is a filter that lets you change the profile of the function.


The **Scale** control lets you adjust the scale of the function when it is mapped onto the terrain.


Amplitude controls the intensity of the perturbations added by the function to the terrain surface. The greater the value, the more noticeable those perturbations will be.

Adding the function can take some time.

The **Retopologize** icon () works on heightfield terrains only, smoothing the entire terrain. A dialog opens to set the amount of smoothing done. If selected for a procedural terrain, you have the option of converting the procedural terrain to a heightfield terrain.

The **Extend terrain canvas** icon () extends the function of a procedural terrain, effectively enlarging it.

The **Force 2D** button () is an override for the entire *Terrain Editor*. The 3D brushes are disabled which avoids 3D displacements involving mesh creation and heavier computations. This also removes any 3D effects on the current terrain.

When the **Procedural material preview** button () is active, it plays as a mask. If you disable it, you will be able to paint anywhere and it will render as it appears in the *Terrain Editor*.

Template:NextButton Template:PrevButton

Terrain Editor -- Predefined Terrain Styles

On the left side of the *Terrain Editor* is a vertical set of icons representing predefined terrain styles. Clicking on any of these will generate a terrain of the requested style, based on the data from the current terrain (in the case of heightfield terrains only).

These terrain styles are:



Reset all: resets all terrain modifications or a selected type.



Reset 3D: resets all 3D terrain modifications.





Reset 2D: resets all 2D terrain modifications.



Reset Material: resets material modifications.



Zero edges: lowers terrain altitudes near the edges, ensuring that they gradually reach altitude zero on the edges. For procedural terrains, this is a toggle button.



Mountain: creates a terrain with higher altitudes near the center. The terrain is generated using a fractal terrain generation algorithm that captures the shape of natural mountains. This is the default style used to create terrains. In the case of procedural terrains, pressing this button replaces the altitude function with the default fractal. Randomness in the shape is achieved by randomizing the origin of the fractal each time you press the button.

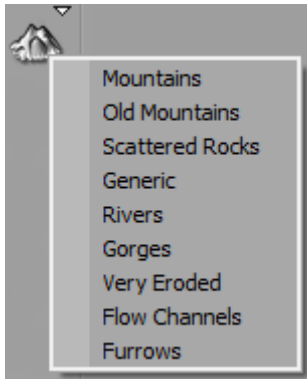


Peak: creates a terrain with higher altitudes near the center. The terrain is generated using a ridged fractal terrain generation algorithm that captures the shape of young mountain ranges. In the case of procedural terrains, pressing this button replaces the altitude function with a ridged fractal noise that produces similar results. Randomness in the shape is achieved by randomizing the origin of the fractal each time you press the button.



Eroded: with Heightfield terrains, uses a powerful erosion algorithm to generate a natural looking, eroded terrain from your existing terrain data. Several erosion presets are available, to achieve a wide range of terrain aspects!





In the case of procedural terrains, this button replaces the altitude function with a simple noise that looks like eroded mountains. Randomness in the shape is achieved by randomizing the origin of the function each time you press the button.



- **Canyon:** applies a filter to the altitudes of the terrain, generating ridges in the terrain profile. In the case of procedural terrains, pressing this button replaces the altitude filter in the [Procedural Altitudes](#) tab.



- **Mounds:** basically the same as the Mountain style, at higher frequency, thus generating several lower mounds that are added to existing terrain data (in the case of heightfield terrains only).



- **Dunes:** uses a function to add dunes to existing terrain data / replaces the altitude function by a dunes function in the case of procedural terrains.



- **Iceberg:** transforms existing terrain data into an iceberg, with a gently sloping, flat top surface. In the case of procedural terrains, the surface is flat instead of sloping, and the profile is achieved by replacing the altitude filter in the [Procedural Altitudes](#) tab.



- **Lunar:** uses erosion and crater effects to create a moonlike, crater-pitted surface from existing terrain data. This option is not available for procedural terrains.





Picture (for heightfield terrains only): lets you import a picture to be converted into a terrain (see *Importing Terrain Data* below). The brighter the picture, the higher the corresponding terrain altitude. These altitudes are blended with existing terrain data. The picture is re-sampled so that it fits exactly onto the terrain.




Options (for heightfield terrains only): displays the [Fractal Terrain Options](#) dialog which shows the typical horizon profile and allows you to change the seed number for random noise and adjust the filters for Noise distribution and Altitude distribution. The Amplitude of perturbations can also be adjusted in this dialog.




Bake to heightfield/Bake to selected layers: is available only when a spline for terrain effect has been applied to the terrain. To bake to selected layers, materials need to have been applied to the terrain and one (or more) selected for the baking process.

Importing Terrain Data






The Import Terrain Data dialog lets you mix existing terrain data with external data sources

VUE lets you create terrains (heightfield terrains, not procedural terrains) from existing data: press the **Picture** button () to open the *Import Terrain Data* dialog and press **Load** button or double-click on the picture preview to select the data file you want to convert. This data can be in any of the supported picture file formats:

{{BulletList Picture data is converted to terrain data based on the brightness of each



pixel. The brighter the pixel in the picture, the higher the resulting terrain altitude. The picture will be automatically resampled to fit the resolution of the terrain (turn to the next page for details on terrain resolution).

You can rotate the picture by using the  and  arrows. You can also invert the picture using the  button.

The *Import Terrain Data* dialog lets you mix existing terrain data with the data contained in the picture. You can indicate how the data is mixed by picking one of the following modes from the **Mixing mode** drop-down list.

Note:

These settings are not available until you load a picture

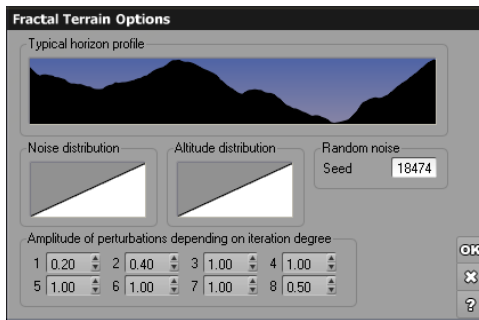
- **Blend:** the terrain data is blended with the data from the picture. The Proportions slider controls the blending ratio between both sources.
- **Add:** the data from the picture is added to that of the terrain. The Proportions slider controls the amount of picture data added to the terrain data.
- **Max:** the data from the terrain is replaced by that of the picture only if it is higher.
- **Min:** the data from the terrain is replaced by that of the picture only if it is lower.
- **Subtract:** the data from the picture is subtracted from that of the terrain.
- **Multiply:** both data sources are multiplied. High altitudes in both sources will remain high, others will be lowered.

Use the **Proportions** slider to vary the amount of mixing that takes place between the existing data and the picture. The result is displayed on a small preview of the terrain. If you press the **Preview** button, the changes will be visible in the *Terrain Editor's* map.


If you want to completely replace the existing terrain by the data contained in the picture, use a **Blend** mixing mode and drag the **Proportions** slider up to 100%.



Fractal Terrain Options



Fractal Terrain Options dialog

- **Options:**  this icon opens the *Fractal Terrain Options* dialog that lets you customize the algorithms used by VUE to generate the mountains. This dialog is not available for procedural terrains as the fractal options are set directly by editing the altitude function settings.
- **Seed:** this is a random number used by VUE to generate the terrain. The same seed will always lead to the same terrain, but you cannot predict the terrain that will be generated from a given seed.
- **Horizon profile:** gives you an idea of the profile of the terrain. A new horizon is displayed any time a parameter is modified.
- **Noise distribution:** specifies the distribution of random perturbations added to the terrain at each iteration step. By default, the noise is distributed randomly around 0.5, which means that bumps and ditches have the same statistical characteristics. By modifying this distribution, you can get interesting effects. Modify the distribution of the perturbations by changing the corresponding filter. If the specified distribution is not balanced (i.e. asymmetrical filter), the terrain may have a tendency to swell or shrink.
- **Altitude distribution:** is, by default, linear, which corresponds to a terrain in which bumps and holes have the same shape. By modifying this distribution, you can get interesting effects. Modify the altitude distribution by changing the corresponding filter. You will associate a new altitude to each existing altitude, proportionately to the value of the filter. Depending on the filter, you will get canyons, plateau... A special collection of filters is available for terrain effects (Select the *Terrain profiles* collection in the Visual Filter Browser).
- **Amplitude of perturbations:** the fractal process is iterative. For each new iteration, VUE adds perturbations to the terrain, the scale of which depends on the iteration step. The higher the iteration step, the finer the perturbations.



VUE lets you adjust the average amplitude of the perturbation at each iteration step. Values in the range of 0 to 3 are reasonable. To understand correctly the iteration process, indicate 0 for each iteration step, then, going from the smallest step (on the left) to the highest step (on the right), indicate 1. Watch the shape of the horizon profile as you adjust the values.

When you are done with adjusting the parameters, click **OK** to accept them. A new terrain will be generated, using the new settings. Any future terrain that you create will use these settings.

Template:NextButton Template:PrevButton



Terrain Editor -- Sculpting Terrains




Paint Tab – Terrain Editor

Sculpting tools are accessible from the **Paint** tab to the right of the terrain map. These tools, or brushes, let you modify manually the shape of your terrain by adding to or digging from it, and selectively applying given effects. The effect that the brush has on the terrain is defined in the Brush Presets group of controls (see below).

Using a Pressure-Sensitive Tablet

If you are using a pressure-sensitive tablet, the pressure can be used to control the parameters of each stroke. This makes for a much more natural and precise way of carving



terrains. Just press the *Pressure Sensitive*  icon at the right of the *Radius*, *Flow* or *Fall-off* parameters to link it to the exerted pressure.

Brush Presets

You now have the option of selecting either **3D** or **2D** brush settings. The **3D** settings are for terrain sculpting; the **2D** settings are for building and lowering the terrain. You can immediately see the 2D brushes by clicking the **Force 2d** button in the upper right of the editor

3D

- **Sculpt:** This option keeps the normal that is currently loaded in the brush and expands (or removes) terrain accordingly, perpendicular to the current surface. The extrusion will follow the direction that the brush is pointing. This direction doesn't change while you are brushing. Use the **Sculpt** brush to move the terrain in a precise direction.
- **Freeform:** This option picks the Normal during the mouse movement and so the extrusion direction changes while the mouse moves. The extrusion will follow the direction that the brush is pointing – this direction may change while you are painting.
- **Pinch:** This option pulls the vertices together.
- **Inflate:** With this option, the vertices are pushed along their normal, so the shape in the brush inflates. The polygons in the brush move away from each other creating a balloon effect.
- **Smear:** this option puts the geometry under the brush into the brush and moves it along with the brush.

Note:

With these 3D options, you cannot paint a material on just a portion of the extrusion, for example, a dot on the end of an extrusion. Extrusions can be painted as part of the terrain as a whole.

2D

- **Raise:** This raises the terrain altitude where it is brushed.
- **Plateaus:** With this option, all of the vertices in the brush are placed onto a horizontal plane. This plane is recomputed when you move the brush and is defined by the brush direction and position.
- **Altitude:** The brush brushes the terrain to the altitude set by the *Altitude Brush* dialog.
- **Smooth:** This smoothes the terrain geometry by adding 3D displacements. The



smooth brush relaxes the underlying polygon geometry in order to remove any distortions in the polygon mesh that might have occurred following painting and automatic subdivision. You can easily observe the effect of the **Smooth** brush by switching to wireframe mode.

- **Flatten:** With this option, all of the vertices in the brush are placed onto a plane. This plane is recomputed when you move the brush and is defined by the brush direction and position. The brush position is a point of the plane; the brush direction is the plane orientation.
- **UniSlope:** This option works like **Flatten**, except that the plane is computed at the first mouse click.

Mode

The first two icons in this section, **Sculpt** and **Material**, define what you are doing with the brush. If you are sculpting your terrain with either 3D or 2D brushes, you should have the **Sculpt** icon selected. If you are painting the terrain with a material, you need to have the **Material** icon selected. If you are painting with a material and sculpting at the same time, both icons should be selected.

The **Freeze** option locks a material layer to disable any further brush effects on that part of the terrain. The **Clear Freeze** option unlocks the layer. **Inverse Freeze** unlocks the currently frozen area and freezes the part of the terrain that was previously unfrozen.

Global Settings

You can set your brush using the settings in the **Global** section. These settings will be applied to the brush for any brush you might choose. These settings may be overwritten for a particular brush in the dialog for that specific brush. See [here](#) for more information.

- **Airbrush:** This controls whether the brush operates like a pen or like an airbrush. If the brush operates like a pen, passing over the same point several times in the same stroke has no additional effect, unlike the airbrush style. The airbrush style brush keeps adding effect as long as the mouse button is down. Non-airbrush mode for 3D painting works like 2D effects.
- **Invert:** Select this icon to subtract from the terrain, lowering altitude.
- **Radius:** Drag the slider to the right to increase the size of the brush. The size of the brush is reflected by the size of the pointer on the terrain map. If you increase terrain resolution the brush resolution will increase accordingly.
- **Flow:** This controls the amount of material added/removed by the brush per unit of time. The higher the flow setting, the more rapidly the terrain will be modified when you press the mouse button.
- **Falloff:** This controls the tapering off of the effect being painted from the center of



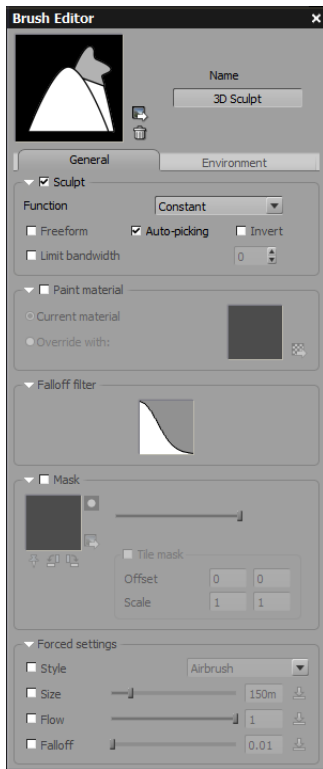
the brush. You can change the filter to change the brush effect.

- **Constrain to clipping range:** When checked, the brush cannot leave the clipping zone defined by the two clipping planes.

These parameters (except *Airbrush* and *Constrain to clipping range*) can be assigned triggers, for a smoother workflow while brushing; see the [Triggers in Editors](#) section of the *Triggers* documentation.

Template:NextButton Template:PrevButton

Terrain Editor -- Brush Editor



Brush Editor General

- **Sculpt:** this option is checked for all sculpting brushes. It is optional for a **Material** brush.
- **Function:** if you wish to change the default setting for this brush, select the type of function you wish to use with this brush from the drop-list. This field is most



useful when creating a new brush.

- **Freeform and Invert:** these further define the brush action.
- **Auto-picking:** this defines the behavior of the brush when you move the mouse after the first click. On the first click the mouse is on the terrain.
With auto picking on, the brush will stay on the terrain when you move the mouse. With auto picking off, it will move straight (in a way parallel to the camera plane) when you move the mouse.
For the most part, auto-picking should be left on. There are cases, such as with the Smear brush, where brush is more effective when it's off.
- **Invert:** Select to reverse the effect. If it is normally raised, this will dig out.
- **Limit bandwidth:** Use the slider to adjust the level of detail of the effect.
- **Paint Material:** check this option if you wish to associate a particular material to this brush. **Current material** is selected by default. If you click **Override with**, the Material browser opens so you can select another material to assign to this brush. This selection can always be changed by clicking the **Load** button next to the displayed material to re-access the Material browser.
- **Falloff filter:** use to pick a custom falloff filter for this brush.
- **Mask:** you can assign a bitmap mask to this brush and set subdivide limits. This image can be inverted and/or rotated.
- **Fix orientation:** Select the pin to fix the orientation of the mask no matter whether the terrain is rotated in the *Terrain Editor* or not. Deselect it if you want the mask to move with the terrain rotation.

The **Forced settings** take precedence over the **Global** settings on the **Paint** tab.

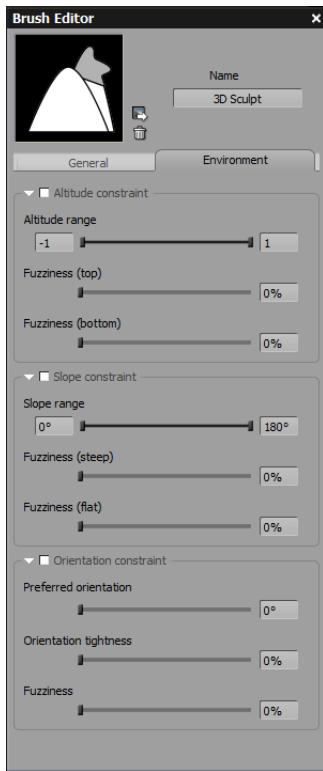
For **Custom** brush only: When you are creating a custom brush, two more fields display at the bottom of the *Brush Editor*.

The first, a dropdown, gives you the option of controlling the direction of an extrusion – along the Camera line, World Z, or Normal.

The first dial below allows you to modify this direction by orientation, the second dial allows you to modify by slope angle.



Environment Tab



Brush Editor Environment

The settings on this tab define **Altitude**, **Slope** and **Orientation** constraints, restricting the brush's effective area. All environment constraints set up on this tab are previewed in the *Terrain Editor*.

If these environment constraints are turned on for this brush you can see the active zones in the *Terrain Editor* where your painting will apply and they serve as a protective mask keeping you within the zone.

To save any changes you may have made to the brush setting, click on the **Save Brush preset** icon under the **Brush Presets**. If you wish to delete any of the **Brush Presets**, highlight the brush and select the **Delete** icon. You can always use the **Add** icon to add the brush back into the list at a later time.



Creating a New Brush

To create a new brush, just open any of the brush presets. As an example, we'll create a brush for the **Stairs** effect.

In the *Brush Editor*:

- Change the name of the current brush to the new brush stairs effect.

On the **General** tab:

- Check **Sculpt**. For **Function**, select **Effects 2D**.
- If you wish to set a default brush size, flow or fall off, these can be set in the **Forced settings** section.
- Skip down to the bottom, and select **Stairs** from the drop box.
- Click on the **Save** brush preset icon. On the screen that displays, key in *stairs_effect* for brush name.
- The new **Stairs** brush will now appear in the presets.


Template:NextButton Template:PrevButton

Terrain Editor -- Painting Materials

When editing a terrain, it is possible to manually paint the distribution of layered materials using the brush. This is available for both standard and procedural terrains, but not for infinite terrains. Zones can be painted separately from the main terrain as well. When painting, the terrain preview switches to a customizable multi-color display that will show where the materials will appear over the terrain. Painted material distributions are fully independent from material scaling and mapping modes, which means that tweaking these material settings in the *Material Editor* will not affect the painted mapping on the terrain.



How It Works



Terrain Editor – Painting Materials

Within the *Terrain Editor*, in the **Paint** tab, you will find the **Material** section. The material currently applied to the terrain is represented in the highlighted box. This is either the default material assigned to the terrain or the material you may have applied to the terrain previously using the *Material Editor*.

Be sure you have clicked on **Material** in the **Mode** section of the *Terrain Editor*. If you will be painting with 2D or 3D effects, you should have **Sculpt** selected as well.


You can add as many materials, or material layers, as you wish. To do so, click the **Add new material** icon (🎲). Mixed and multi-layered materials can be used. You can also remove a material at any time by highlighting it and clicking on the **Delete** icon (🗑️).

Painting a layer in the *Terrain Editor* does not destroy the original alpha of the material, but blends with it. The **Export alpha maps** icon (📄) below allows you to export that alpha information for use in the *Material Editor* if further editing is required.

Each material layer's preview color corresponds to the flat color accessible from the *Color Editor*. Click on the color box to change the representational color in the *Terrain Editor*.



to increase visibility while painting.

Next to the material box is an icon () that opens the *layer options menu*. See the **Layer options** section below for details.

Painting with Effects

To paint with effects, such as pebbles, stairs, cracks, select the pebbles effect brush. The **Sculpt** box should be checked and **Effects 2D** selected.

At the bottom of the *Brush Editor*, use the drop-box to select the effect you wish to paint with.

On the **Paint** tab, **Mode** section, be sure both the **Sculpt** and the **Material** icons are selected. Select the material you wish to paint the effect with. And paint away.

If you only wish to paint the effect and not use a material, just uncheck the **Material** icon in the **Mode** section.

Layer options

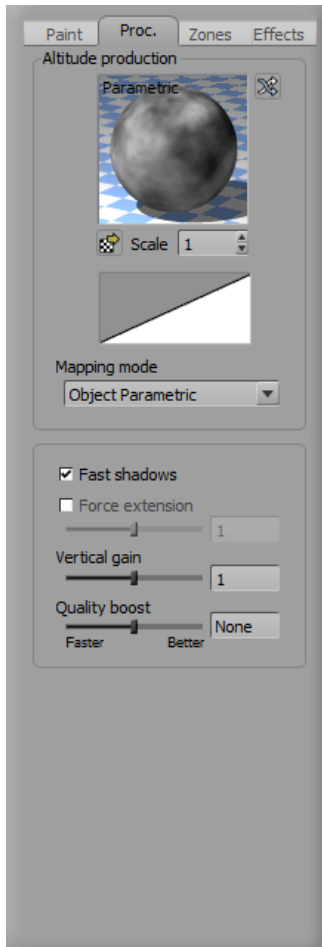
The layer options menu contains three entries: **Painted presence**, **Presence from graph**, and **Show in Preview**. The **Painted presence** and **Presence from graph** are two mutually exclusive modes. The **Show in Preview** menu entry can be used to show or hide layers as needed when painting.

Additionally to painting the layers directly on the terrain, you can instead drive the presence of a layer with a node from the graph. This disables the manual painting mode, but makes room for interesting effects like painting the material automatically on rough areas of the terrain.

To connect a layer to the graph, select the **Presence from graph** submenu. It displays a list of nodes and their possible outputs. If you select one, this layer will switch to the second mode, where the presence is defined by the graph. You can see the result of this connection in the main display of the *Terrain Editor*.



Terrain Editor -- Procedural Altitudes




Procedural Tab – Terrain Editor

If the terrain is a procedural terrain, a **Graph** tab is available in the terrain editing toolbox for altitude production. Each time you modify the definition, extension or mapping of the altitude function, VUE refreshes the 3D preview of the terrain. While this refresh is taking place, the heightfield terrain modification tools are disabled.

The controls in this tab are as follows:

- **Altitude production function:** this is the function that defines the altitudes of the terrain at each point. Double-click on the function preview, or click on the



Load function button () below the preview in order to load a new preset altitude function. The *Terrain altitude functions* collection contains interesting function presets for generating procedural terrain altitudes. However, any type of function can be used (including functions that output a color – in which case the color will be converted to a brightness value automatically). You can also edit the altitude function directly by opening the *Function Graph* to modify the definition of the function (Ctrl+Click on the function preview, or select **Edit Function** from the function popup menu). Turn [here](#) for details on the *Function Graph*.

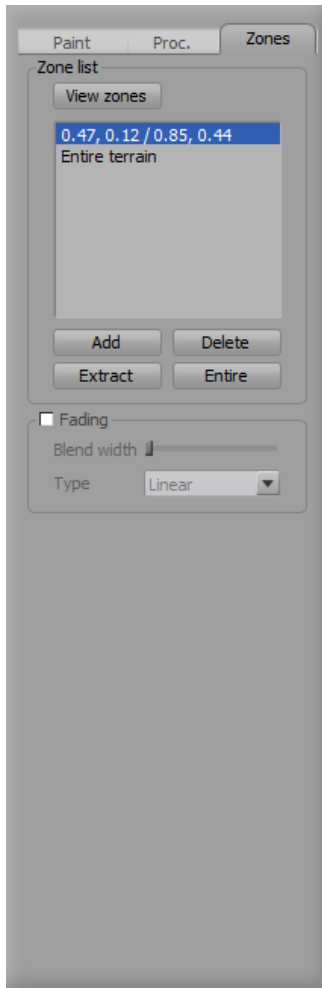
- **Scale:** this setting lets you adjust the scale at which the altitude function is mapped to the terrain. Increase the scale to create larger features, reduce it to see a larger portion of the function.
- **Filter:** this lets you specify a filter to modify the altitude values according to a user defined profile.
- **Mapping mode:** this lets you define the coordinate system that is used to generate the altitudes of the procedural terrain. The different coordinate systems are the same as the coordinate systems available in the Material Editor's **Mapping** list and described in the *Understanding VUE* section about [coordinate systems](#). If you change mapping mode, VUE will offer to insert nodes into the altitude function in order to maintain the look of the terrain despite the change.
- **Fast shadows:** when this option is selected (the default), procedural terrain shadows will be approximated using an extremely fast algorithm. However, in some cases, this approximation may not be satisfactory (e.g. in the shadows cast by a terrain onto very distant objects). If such artifacts appear, you should disable this option in order to enable a more complete (and significantly slower) processing of terrain shadows.
- **Force extension:** the terrain extension is a parameter that controls the vertical size of the region in which the terrain altitudes are evaluated. Any altitude beyond this extension will be clipped to the limits of the extension (similar in some way to the clipping options – see below). Terrain altitudes are mapped to the terrain altitude color map according to the value of this extension. The **Force extension** option lets you define the extension manually using the slider and edit field below the checkbox. However, it is usually recommended that you leave the **Force extension** option off, so that VUE computes automatically the best extension for the terrain (this is done in two passes). One drawback of letting VUE compute the extension automatically is that you don't get such a good understanding of the vertical size of the features in the terrain (you have to check in the *3D Views* to see that).
- **Vertical gain:** this lets you adjust the height of the terrain using the slider.
- **Quality boost:** this setting lets you adjust the precision with which the geometry of the terrain is evaluated at render time. This setting works in conjunction with



the **Advanced effects quality** setting in the *Render Options*. It is usually not useful to modify this setting, unless you notice unwanted artifacts – especially near crests – in the terrain rendering. Keep in mind however that increasing the render quality of the entire picture will also increase the rendering quality of the procedural terrain.

Template:NextButton Template:PrevButton

Terrain Editor -- Zones



Zones Tab – Terrain Editor

Zones are only available with procedural terrains. You can only work with zones using the **3D** settings, because 3D displacement accuracy is proportional to zone size, whereas 2D displacement is only proportional to the entire terrain size.

Zones are created by selecting **Add** on the **Zone** tab, then using the left mouse button to create the zone area directly on the terrain. When you release the mouse button, only the zone is displayed in the *Terrain Editor*.

The Zone name contains the coordinates inside the terrain. You can change this name if you choose.

This allows you to work with much higher resolution on that area of the terrain with a more detailed heightmap. There is also less subdivision occurring which speeds processing. And each zone can be painted/textured separately as well.

Fading: This controls the blending of the changes made in the zone into the rest of the terrain. You can define the size of the blending area and the type of blending to use.


Use the **Extract** button to create a new procedural terrain from a zone. You can extract the defined zone which then becomes a new procedural terrain. The new terrain is created in the same position and is the same size as the zone. It is created with an object mapping and no **Zero edges**. The newly created terrain can then be moved to a different position. You have the option of leaving the original terrain as it is or to actually subtract the new zone from the old terrain, leaving a hole in the terrain.

Terrain Editor -- Effects

All **Global** and **Erosion effects** can be applied multiple times for greater effects. Right click on the effect and a dialog displays with a slider to set the **Iteration count** and the effect will be applied that number of times.



Erosion Effects

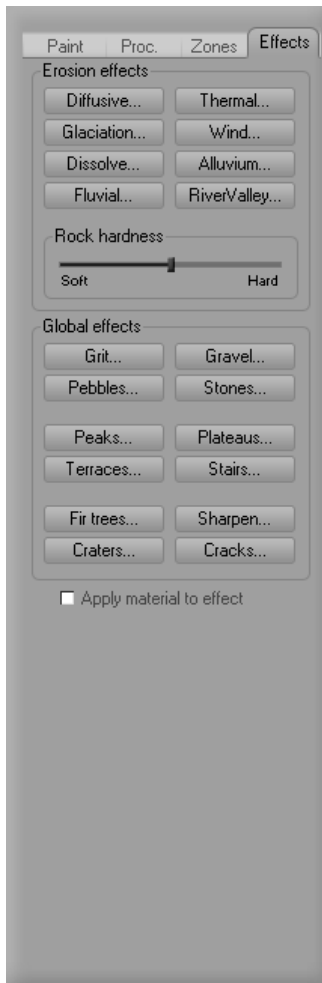


Using the controls of this section you can generate within seconds the effects of erosion that are achieved by nature in millions of years!

The **Rock hardness** slider influences all of the erosive processes. Unlike real rock hardness, this one may be modified at any time between successive applications of erosion.

There are 8 types of erosion to choose from:





Effects Tab – Terrain Editor

- **Diffusive:** erosion is the result of the application of numerous types of erosion (including vegetal growth, grazing animals...) over millions of years. It results in the rounding off of any sharp features on the terrain. Harder rocks are less subject to diffusive erosion than softer ones.
- **Thermal:** erosion results in the loosening up of substratum that subsequently falls down to pile up at the bottom of an incline. Thermal erosion creates scree slopes of a constant angle. It is caused by rocks bursting because of strong exposure to heat, or ice. Although this is not geologically correct, the **Rock hardness** slider here controls the angle of scree slopes: the harder the rock, the steeper the scree.



- **Glaciation:** is caused by glaciers tearing away parts of the terrain that are at low altitudes, resulting in the typical smooth valleys with rounded profiles. Harder rocks are less subject to this type of erosion than soft ones.
- **Wind:** erosion rounds off features of the terrain that are directly exposed to the wind. In VUE, this wind blows horizontally from left to right. Features that are sheltered by other parts of the terrain will not be eroded as much as ones that are fully exposed to the wind. Harder rocks are less subject to wind erosion than softer ones. You can change the direction of the wind by rotating the terrain in another 2D app (using Copy-Paste).
- **Dissolve:** erosion is caused by rainwater infiltrating the terrain and dissolving or flushing away parts of the terrain. The effect is particularly strong at low altitudes. As a result, numerous streams will appear at the surface of the terrain. If the rock is hard, the streams will remain narrow, but if the rock is soft, the streams will be wide and the entire surface of the terrain smoothed down.
- **Alluvium:** is the same as dissolve, except that the matter that gets torn off the surface of the terrain is moved down by the water streams onto flat areas where the stream slows down on and creates sediment deposits.
- **Fluvial:** erosion is caused by streams of rainwater. As rain falls over the terrain, it gathers in streams of increasing strength, tearing away pieces of rock as it goes by. When the rock surface is hard, streams tend to remain parallel longer, and slowly dig furrows with vertical sides. When the rock surface is soft, they join up more rapidly, and earth collapses on the sides of the streams, creating wide, gently sloping furrows.
- **River Valley:** erosion is a more accurate, geologically based erosion filter. Because it's more accurate, it's also quite a bit slower.

Maintaining one of the erosion buttons depressed will keep eroding the terrain until the button is released.

Global Effects

The second section on the **Effects** tab displays a group of geological effects that can be applied to the terrain. Holding any one of these buttons down will keep applying the effect until the button is released.

- **Grit:** adds random noise all over the surface of the terrain, resulting in a surface covered by little bumps and holes. Maintaining this button down creates higher bumps and deeper holes.
- **Gravel:** same as grit, except gravel concentrates on sloping areas of the terrain.
- **Pebbles:** adds randomly distributed pebbles all over the surface of the terrain. Useful for modeling pebble beaches... Maintaining this button down creates thicker pebbles.



- **Stones:** adds randomly distributed round, bulging stones all over the surface of the terrain (a good base for rocky terrains). Maintaining this button down creates thicker stones.
- **Peaks:** filters the altitude to emphasize high areas, while digging deeper valleys. Results in peaks separated by deep valleys.
- **Fir trees:** adds tiny, randomly distributed cones all over the surface of the terrain. Very useful for modeling distant forests. Maintaining this button down creates taller cones.
- **Plateaus:** very useful effect that causes high altitudes to swell, resulting in plateaus. Yields interesting results when used in conjunction with stones.
- **Terraces:** gradually transforms any part of a sloping terrain into terraces. Ideal for cultivating rice!
- **Stairs:** quantizes altitudes in your terrain, resulting in stair-stepped terrain structures. Very useful for modeling desert plateau structures. Keep the button pressed to reduce the number of steps.
- **Craters:** bombards the surface of the terrain with random meteorites. The distribution of meteorite size is varied realistically.
- **Sharpen:** this effect will increase the steepness of already steep parts of your terrain.
- **Cracks:** adds random vertical cracks to your terrain surface, not unlike those created by an earthquake.
- **Apply material to effect:** checking this box will apply the current selected material on the **Paint** tab to the effect selected.

Published Parameters (for Procedural Terrains only)

The **Published Parameters** feature copies specific settings from the *Function Graph* that you may need to change often and places them in a more convenient location for easier terrain manipulation. In the *Terrain Editor*, a new tab is created. For terrains, this is usually some parameter used for a procedural altitude function.

To select a parameter for publishing, just click the underlined parameter name. A parameter name is supplied and a group name is asked to improve the display of the published parameter. This parameter will then be available on a **Published Parameters** tab in the *Terrain Editor* for easy access.


Clipping Terrain Altitudes

The idea behind clipping altitudes is simple: anything beyond the clipping altitude will be left out of the terrain when it is rendered. This means you can make lower parts of the terrain actually become holes in the terrain, and higher parts become perfectly flat.



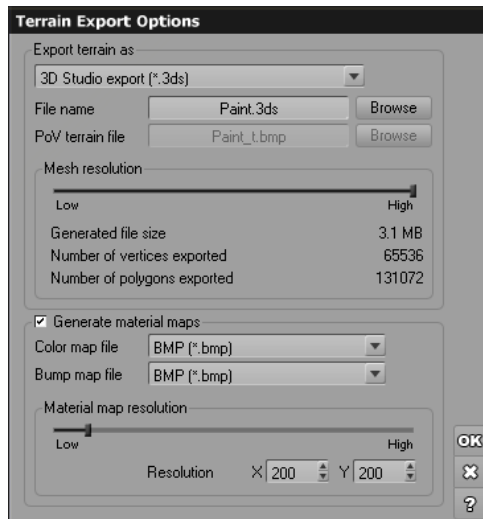
The creative power behind this is incredible. Please refer to the tutorial on [making stone arches](#) for an example using this feature. The result of clipping altitudes is the same as using a Boolean operation to remove low or high parts of the terrain (only much more efficient).

The clip slider has two entries, one at each end. Drag the lower end of the slider to adjust low clipping, and drag the upper end of the slider to adjust high clipping (you can also enter numerical values for the clipping altitudes). If you drag the slider from its center, both low and high clip values will be modified together. As you adjust the clipping slider, parts of the terrain that are clipped out disappear or become flat in the terrain map. These altitudes will be left out during render.

The altitude color map can be stretched to the clipping planes by clicking the small button  to the left of the color map.

Template:NextButton Template:PrevButton

Terrain Editor -- Exporting Terrain Geometry



Terrain Export dialog

Click on the **Export terrain** icon to export the terrain in a format readable by other 3D packages.

Along with the terrain, VUE can also automatically generate the corresponding color and



bump maps. These maps will fake the procedural materials used by VUE inside most other 3D packages.

The terrain can be exported using any one of the following file formats:

- **3DS:** 3D export in 3D Studio file format (automatic mapping with UV information),
- **LWO:** 3D export in LightWave file format (automatic mapping),
- **OBJ:** 3D export in standard Wavefront file format (UV mapping information),
- **C4D:** 3D export in Cinema 4D file format (UV mapping information),
- **DXF:** 3D export in standard AutoCAD file format (no mapping information),
- **Picture Formats:** terrains can be exported as 8 bit resolution bitmaps under all supported picture formats. Using the **TGA** format you can also export the terrain under the standard 16 bit height field encoding scheme (the red byte is the high order byte, the green one the low order byte, and the blue byte is ignored). The **PSD**(VUE Pro versions only) and **TIFF** file formats let you export the terrain data as a 16 bit file.

The topmost drop-down list lets you select the file format under which the terrain will be exported in.

The **Mesh resolution** control lets you adjust the resolution of the polygon mesh that will be generated by VUE when exporting the terrain as a 3D object. It also gives an estimate of the size of the files generated depending on the resolution you choose. The higher the resolution, the bigger the number of polygons in the file, and the larger the file.

If you select one of the picture file formats, the terrain data will be converted to grayscale pictures (high altitudes will appear as bright areas, and low ones dark areas).

If you would like VUE to generate the corresponding color or a bump maps, select the **Generate material maps** option. You will have to indicate the files that should be used to store the color and the bump maps. You can also indicate the **Resolution** at which these maps will be generated. The higher the resolution, the more detailed the maps will be, but the longer they will take to generate.

You can also export terrains using the generic **File | Export Object** menu command.

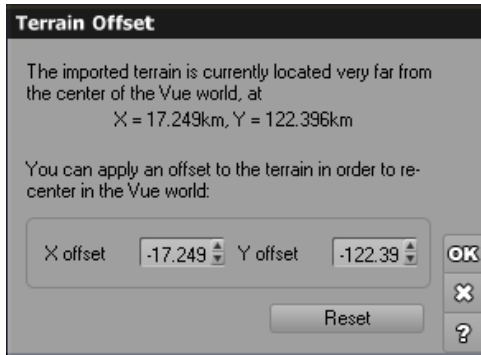
Please turn [here](#) for details on the “Export Object” command.

Template:NextButton Template:PrevButton



USGS Digital Elevation Model Data

You can import a large variety of terrain data using the *Terrain Editor*. If you import data into an already existing terrain, it will be resampled to fit the terrain's geometry.



The Terrain Offset dialog

Digital Elevation Models (DEM) that follow the file format specifications of the United States Geological Survey can be imported directly through the **File | Import Object** menu command. A terrain will automatically be created.

If the **Resize object** and **Center object** options are selected in the *Import Options* dialog, the terrain will be placed at the center of the views. If it isn't selected, the terrain will automatically have the same size and orientation as that described in the DEM file. The terrain will also be positioned according to the information in the DEM file. The immediate advantage of this is that you can import several DEM files and they will automatically be positioned correctly relative to each other. This topic is illustrated in the tutorial *Importing Multi-Part DEMs*.

Whenever the position of the DEM terrain is outside VUE's valid range, the *Terrain Offset* dialog will pop-up, asking you to enter an offset. The first time the dialog appears, the displayed offset will center the terrain in the views. Future imports will position the new terrains relative to the first one, letting you easily import multi-part DEM maps. Pressing **Reset** will compute a new default value for the offset so that the terrain is centered.




Hint

Applying effects at different terrain resolutions will result in these effects having different resolutions themselves.

For instance, a good way of achieving realistic erosion would be to start from a terrain with a resolution of 128×128 and adding dissolve erosion, then doubling terrain resolution, applying diffusive erosion, and then some more dissolve erosion. Then repeat this process again to raise the terrain resolution to 512×512 .

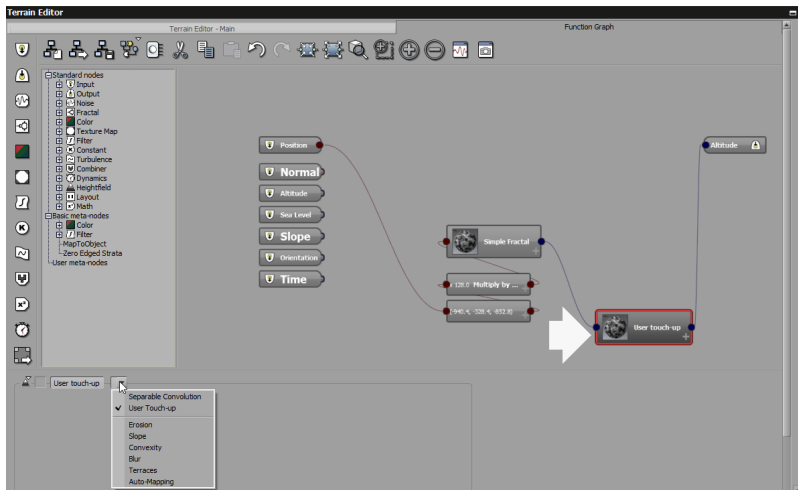
User Touch-up Graph Node -- Heightfield Terrains



User touch-up node

The User Touch-Up Graph Node is created when you paint on a heightfield terrain using a 2D sculpting tool (or brush) on a heightfield terrain.





This node appears in the Function Graph labeled **User Touch-up**. If you delete this node in the Function Graph, the changes you made to the terrain will be deleted as well. Selecting the *Reset 2D sculpting only* icon will also delete the node. If you then go back to the terrain and make more changes, you will get another **User Touch-up Node**.

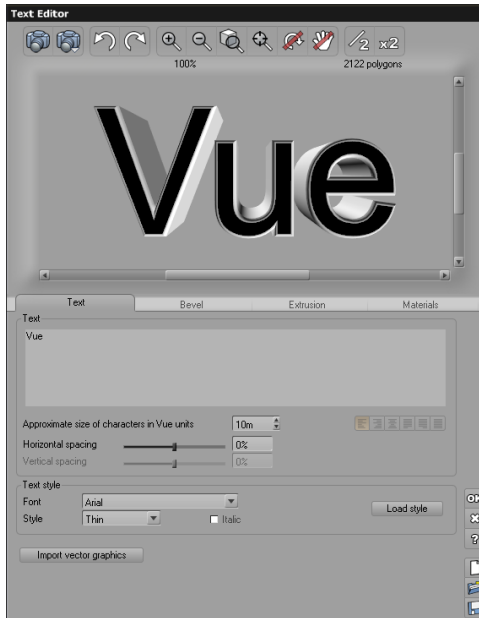
When you highlight this node, several options display in a drop-box below.

- Separable Convolution
- Erosion
- Slope
- Convexity
- Blur
- Terraces
- Auto-Mapping

This feature is only available for the heightfield terrain.




Text Editor




Text Editor – Text Tab

The *Text Editor* lets you create elaborate 3D Text effects. 3D Text can be used for creative titling, but it can be also used as a basic modeling tool.

To access the *Text Editor*, either:

- Double-click on the 3D Text you want to edit in the *3D Views* or in the *World Browser*,
- Click on the **Edit object** button () on the top toolbar, when the 3D Text object is selected,
- Use the menu command **Object | Edit Object**.

The *Text Editor* also automatically appears when you create a new 3D Text object by:

- Clicking on the 3D Text icon () in the left toolbar, or
- Selecting the menu command **Object | Create | 3D Text**.



Interface Elements

The *Text Editor* is comprised of the following elements:

- A toolbar that contains icons to manipulate the 3D Text,
- A preview area where a preview of the 3D Text is displayed, and
- 4 tabs: Text, Bevel, Extrusion and Materials.

Toolbar



New: creates a new empty file.




Open: opens a saved text file.



Save: saves a text file.



Render Preview: clicking this icon will start rendering a preview of the 3D Text (this preview uses the same settings whatever the scene setup). This rendering is done in the background, so you can still access settings as it renders. If you change a setting or click in the preview, rendering will stop. You can also use the options in the **Preview OptionsMenu** icon () to enable automatic rendering of the text (see below). When automatic rendering is enabled, this Render preview icon is disabled.



Preview Options Menu: the *Text Editor* preview options menu is accessed by clicking on this icon in the toolbar. Options in this menu are:

- **OpenGL Preview:** select this option to display a detailed OpenGL preview of the text as you edit it (the text is generated in a background thread).



- **Auto-Render:** select this option to automatically render the text as you edit it (the text needs to be regenerated before rendering can begin, hence the small delay – all processing is done in a separate thread for better response).



- **Zoom In:** click this icon to display a magnified view of the text in the preview, thus letting you observe the text in finer detail. The current zoom factor is displayed beneath the icon.



- **Zoom Out:** click this icon to display a reduced view of the text and get a more global view of it.



- **Reset Point of View:** if you rotate, pan and/or zoom in the interactive preview of the text, this resets all these settings to the initial ones.



- **Reset Zoom:** click this icon to reset the zooming of the text without affecting the orientation and pan of the view.



- **Reset Rotation:** click this icon to reset the orientation of the text without affecting the zooming and panning of the view.



- **Reset Pan:** click this icon to reset the panning of the text without affecting the orientation and zooming of the view.



- **Reduce precision:** click this icon to simplify the geometry of the text. The number of polygons in the text is displayed beneath the icon. The overall shape of the text won't be modified, but the level of detail will be reduced, resulting in less polygons and a faster – but coarser – representation of the text. This is ideal for text seen in the distance and that doesn't require the full level of detail.

Note:

Only curved sections of the text will be affected by the simplify/refine concept, and that the effect of the simplification will be to reduce the smoothness of these curved



sections. The more curved sections in the text, the larger the polygon reduction each time you press this icon. It is not possible to simplify the text geometry beyond a certain point – point at which the icon will become disabled.



- **Increase precision:** this is the opposite of the previous icon. What it does is increase the level of detail in the text (still without changing the overall geometry). This results in a better defined text with smoother curves, but a longer render time. Refining text is useful when you need to look at a text from up close, and you can see angular shapes in the curved sections.

Note:



Only curved sections of the text will be affected by the simplify/refine concept, and that the effect of the simplification will be to increase the smoothness of these curved sections. The more curved sections in the text, the larger the polygon increase each time you press this icon. You should keep an eye on the total number of polygons in the text as you can very easily end up with a massive number of polygons – this will result in slow rendering and heavy resource requirements.

Text Preview

The text you are working on is displayed in 3D in the middle of the editor (this display uses OpenGL). Each time you modify your 3D Text settings, VUE will regenerate a new preview of the text. Because this process can be time consuming, this regeneration is done in a background thread in order to avoid slowing you down. The different characters in the text will appear in the 3D preview as they are generated.

You can rotate this 3D preview by dragging it with the right mouse button (Ctrl drag on Mac), and you can pan it by holding the left mouse button as you drag the preview (or you can use the scroll bars).

When rotating the 3D preview, if you release the mouse button while the mouse is still in movement, the 3D Text will keep rotating automatically. To stop this automatic rotation, simply click again on the preview.

You may enlarge or reduce the zooming on the preview by pressing the **Zoom** icons ( and ) or by dragging the mouse up/down with the Control key and the right mouse button pressed. All this does is change the zooming at which the text is viewed in the preview. The zooming factor is displayed under the two Zoom icons.



Text Tab

The **Text** tab lets you define the characters in the text as well as overall text layout.

Text

The main field in this group is used to enter the text. If you enter several lines of text, the alignment and vertical spacing tools become active.

- **Approximate size of characters in VUE units:** this setting controls the approximate size of the characters as they will appear in the *3D Views*. This size is approximate because all the characters don't necessarily have the same size.
- **Horizontal spacing:** this setting controls the spacing between the characters. Negative values will make the characters closer, positive values will increase the space between the characters.
- **Vertical spacing:** this setting is only active if you have entered several lines of text. It controls the spacing between successive lines of text. Negative values will make the lines closer to each other, while positive values will increase the space between the lines.
- **Alignment tools:** these tools are only active if you have entered several lines of text. They control the way the text is aligned from one line to the next:



Align left: when this icon is active (orange) the lines will be aligned along their left edges.



Align right: when this icon is active (orange) the lines will be aligned along their right edges.



Align center: when this icon is active (orange) the lines will be aligned along their centers.



Justified left: when this icon is active (orange) the spacing between the char-



acters will be adjusted so that all lines except the last have exactly the same length. The last line will be aligned with the left edge of the other lines.



Justified right: when this icon is active (orange) the spacing between the characters will be adjusted so that all lines except the last have exactly the same length. The last line will be aligned with the right edge of the other lines.



Justified paragraph: when this icon is active (orange) the spacing between the characters will be adjusted so that all lines including the last have exactly the same length.

Text Style

The controls in this group let you select the style of the base characters in the text.

- **Font:** use this drop-down list to select the font that will be used to generate the 3D characters.
- **Style:** use this drop-down list to select the base thickness and style of the characters.
- **Italic:** this option, only available on Windows (on Mac OS X, the italic property is part of the character styles), lets you indicate whether the text should be italicized or not.

Using Vector Graphics

Instead of using regular text, you can also import a vector graphics file and use the *Text Editor* to bevel and extrude it. To do this, click on the **Import vector graphics** button and select the file to import using the Standard File Browser. Supported formats are:

- **PS:** Postscript,
- **EPS:** Encapsulated Postscript, and
- **AI:** Adobe Illustrator (up to version 3.2).

Note:

Because of the inherent complexity of these Postscript-based formats, it is impossible to ensure that all files will load correctly. In order to obtain best results, it is recommended that you use only the simplest form of Postscript language when exporting vector data. In the case of complex colored documents, only the vector contours are imported.



Once a vector graphics has been imported, the text field and text-specific controls become disabled. However, all other controls such as Extrusion, Bevel, etc. can be used to turn your 2D graphics into nice 3D logos.

Bevel Tab



Text Editor – Bevel Tab

This tab lets you control the beveling applied to the characters.

- **Bevel:** use this checkbox to enable character beveling.
- **Width:** this parameter controls the width of the bevel that will be added around the text. This is a percentage of the total width of the characters.
- **Depth:** this parameter controls the depth of the bevel that will be added around the text. This is a percentage of the total width of the characters.
- **Cap bevels:** when this option is enabled, the front and rear of the text will be capped in between the bevels. If it is disabled, the text will be hollow in between the bevels.



- **Bevel rear side:** if this option is selected, both the front and the rear of the text will be beveled.
- **Bevel inwards/Bevel outwards:** Check which way you would like the beveling to be done.
- **Bevel profile:** you can select a standard bevel profile using one of the presets, or you can define your own profile by loading or editing the **Custom** bevel profile (the profile is defined by a filter).
- **Interior bevel profile:** if you select this option, you will be able to define a different bevel profile for the holes in the characters (only applicable if the characters have holes, e.g. 'a, e, o...'). Select a standard bevel profile for the interior using one of the presets, or define your own profile by loading or editing the **Custom** bevel profile (the profile is defined by a filter).

Extrusion Tab



Text Editor – Extrusion Tab

This tab lets you control the way the characters are extruded.

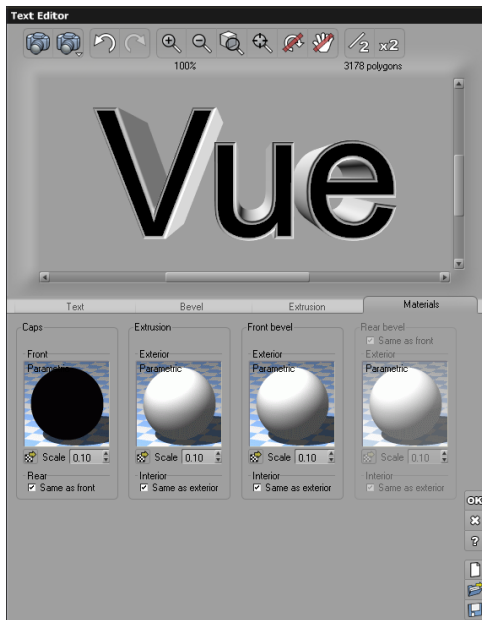


- **Extrusion:** use this checkbox to enable character extrusion.
- **Length:** this parameter controls the length of the extrusion in VUE units.
- **Amplitude:** this parameter controls the amplitude of the extrusion profile. If you extrude along a flat profile, this parameter has no effect and will be disabled. If the extrusion profile is not flat, the parameter controls the relative depth of the changes in the extrusion profile.
- **Scaled extrusion:** this option indicates that the scaling of the characters takes place relative to the center of the character.
- **Beveled extrusion:** this option indicates that the scaling is done in the same way as the beveling, that is relative to each stroke in the character.
- **Symmetrical extrusion:** when this option is selected, the extrusion profile is applied symmetrically, first from left to right, and then reversed. The extrusion length is doubled when you select this option.
- **Extrusion profile:** you can select a standard extrusion profile using one of the presets, or you can define your own profile by loading or editing the **Custom** extrusion profile (the profile is defined by a filter).

Hint: Using the extrusion tools, you can create entirely new geometrical shapes based on simple characters such as e.g. “o” or “H”.


Materials Tab





Text Editor – Materials Tab

This tab lets you select the materials that should be assigned to the different parts of the text. You can define separate materials for the caps, the extruded part and the beveled part, plus you can define separate materials for the front and rear bevels and caps and the interior or exterior bevels (8 materials altogether!).


Double click on the material preview to edit each material, or press the corresponding **Load material** button () to load an existing material. You can adjust the scale of the material using the **Scale** settings below each material.


- **Caps:** this group controls the materials assigned to the text caps. If **Same as front** is checked, the same material will be used for both the front and rear caps. Otherwise, a second material preview will appear letting you assign a different material to the front and rear caps.
- **Extrusion:** this group controls the materials assigned to the extruded parts. If **Same as exterior** is checked, the same material will be used for both the interior and exterior extrusion. Otherwise, a second material preview will appear letting you assign a different material to the interior and exterior extrusion.
- **Front bevel:** this group controls the materials assigned to the beveled parts on the front end of the text. If **Same as exterior** is checked, the same material will be used for both the interior and exterior bevels. Otherwise, a second material preview will appear letting you assign a different material to the interior and exterior bevels.




- **Rear bevel:** this group controls the materials assigned to the beveled parts on the rear end of the text. If **Same as front** is checked, the same materials will be used for the rear end as for the front end (the other controls in the group will be disabled). If **Same as exterior** is checked, the same material will be used for both the interior and exterior bevels. Otherwise, a second material preview will appear letting you assign a different material to the interior and exterior bevels.

Text Styles

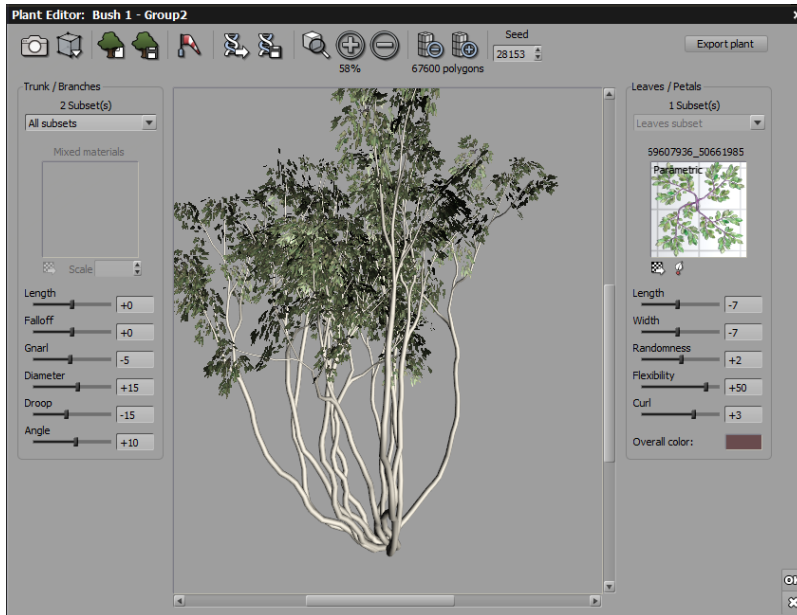
Use the **Load** icon () in the dialog bar to load a preset text style from the Text Styles Visual Browser (see [here](#) for details on using Visual Browsers). Text styles are previewed using the “Abc” text string. If the font used in a particular text style is not available on your system, the default font will be used instead.

If you have designed a text style that you are particularly happy with, you can save it for future use. Press the **Save** icon () and select a file, title and description for the new text style. A preview of the text style will be automatically generated.

You can reset all text settings anytime by clicking the **New** icon ()




Plant Editor



The Plant Editor dialog

Plants are the essential touch to turn those barren terrains into convincing natural scenery. Luckily, VUE has one of the most advanced plant generation technologies around. This technology, called *SolidGrowth™* lets you grow unique plants directly inside the application. Thanks to the *Plant Editor*, it is now possible to modify these plants and create whole new species.

To access the *Plant Editor*, either:

- Double-click on the plant you want to edit in the *3D Views* or in the *World Browser*.
- Click on the Edit object button () on the top toolbar, when the plant is selected.
- Use the menu command **Object | Edit object**.

Inside the *Plant Editor*, the plant that is being edited is viewed in 3D perspective. Modifications you make to the plant are reflected in real-time. You can also generate a rendered preview of the plant.





Plant Preview

The plant you are working on is displayed in 3D in the middle of the editor (this display uses OpenGL). Unlike the preview in the *3D Views*, the real-time 3D display of the *Plant Editor* attempts to reproduce the variations in leaf colors.

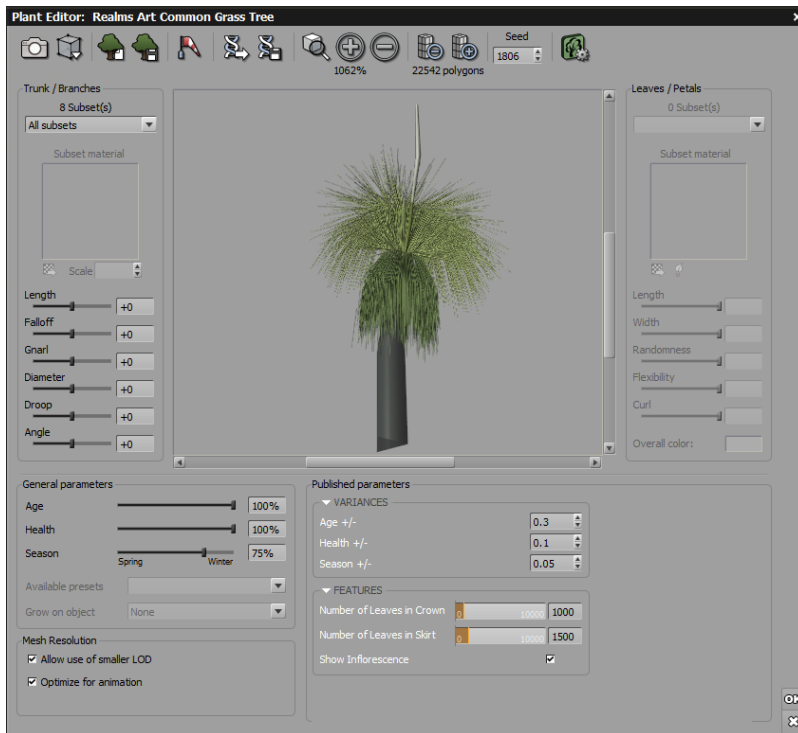
You can rotate this 3D preview by dragging it with the right mouse button (Ctrl drag on Mac), and you can pan it by pressing the left mouse button as you drag the preview (or you can use the scroll bars).

When rotating the 3D preview, if you release the mouse button while the mouse is still in movement, the plant will keep rotating automatically. To stop this automatic rotation, simply click again on the preview.

You may enlarge or reduce the zooming on the preview by pressing the **Zoom** icons ( and ) or by dragging the mouse up/down with the Control key pressed. All this does is change the zooming at which the plant is viewed in the preview. The zooming factor is displayed under the two Zoom icons.

Plant Preview for TPF Plants





Plant Editor for TPF Plants

Since plants created in PlantFactory may have a more complex structure than basic VUE SolidGrowth plants, the Plant Editor reflects the new parameters of the plant. In the example here, parameters that can be changed in the Plant Editor include Age, Health and Season parameters as well as a mainbranch count that can be varied.

Mesh Resolution

Meshing options are available in this Editor:

- **allow use of smaller LOD:** use automatically lower resolution when the plant is far from the camera
- **optimize for animation:** adapt the meshing for more realistic deformation




Published Parameters

If the plant has been created in TPF with published parameters, these parameters are available here so that you can easily adjust these parameters for the plant.

Toolbar



Render Preview: clicking this icon will start rendering a preview of the plant (this preview uses the same settings as for the previews of other plants shown in the Plant Browser). This rendering is done in the background, so you can still access settings as it renders. If you change a setting or click in the preview, rendering will stop.

You can also use the options in the **Preview OptionsMenu** icon () to enable automatic rendering of the plant (see below). When automatic rendering is enabled, this Render preview icon is disabled.



Preview Options Menu: the *Plant Editor options* menu is accessed by clicking on the icon in the toolbar. Options in this menu are:

- **OpenGL Preview:** select this option to display a detailed OpenGL preview of the plant as soon as you stop editing or moving the plant.
- **Auto-Render:** select this option to automatically render the plant as soon as you stop editing or moving the plant.
- **Show Wind Effects:** by default, the effects of wind are not viewed in the editor, as they can bias your perception of the plant. If you would rather like to view the effects of the wind on the plant, select this option. This option is, of course, only available when wind has been applied to the plant.



New Plant: click on this icon to build a new variation of the plant based on the same settings as the ones applied to the current plant. Clicking repeatedly will create a [different plant](#) each time thanks to the way *SolidGrowth* dynamically grows the plants inside the software. This is useful when you are trying to find a specific shape for a plant, or when you want to see the effects of your settings on a selection



of plants (prior to saving a new plant species, for instance).



Save Plant: click on this icon to open a Standard File Browser and save the plant under VUE's native **VOB** file format. If you would like to save the plant using another format, use the **Export plant** button instead.



Response To Wind: click on this icon to display the *Response To Wind Options* dialog. This dialog is used to adjust the amount of deformation of the plant under a wind (or breeze) of given intensity.



Load Plant Species: click on this icon to replace the current plant by a completely different plant of a new species. The Visual Plant Browser appears letting you select a new plant species. All adjustments made in the editor will be lost as they revert to the new species' settings.



Save Plant Species: this is one of the most powerful features in the *Plant Editor*. This icon lets you save an entirely new plant species that will appear in the Visual Plant Browser like any other plant. You can subsequently grow instances of this new plant species in future scenes. When you click this icon, the preview of the plant is rendered and a Standard File Browser appears, letting you enter a name and description for the **new plant species**.



Frame: click this icon to automatically adjust the framing of the preview so that the plant is centered and entirely visible.



Zoom In: click this icon to display a magnified view of the plant in the preview, thus letting you observe finer details of the plant. The current zoom factor is displayed beneath the icon.



Zoom Out: click this icon to display a reduced view of the plant and get a more



global view of it.



- **Simplify Plant:** click this icon to simplify the geometry of the plant. The number of polygons in the plant is displayed beneath the icon. The overall shape of the plant won't be modified, but the level of detail will be reduced, resulting in a faster – but coarser – representation of the plant. This is ideal for plants seen in the distance and that don't require the full level of detail.

Note:

Leaves are not affected by the simplify/refine concept, and that these operations are not stored in plant species.



- **Refine Plant:** this is the opposite of the previous icon. What it does is increase the level of detail in the plant (still without changing the overall geometry). This results in a better defined plant with smoother curves, but a longer render time. Refining plant is useful when you need to look at a plant up close, and you can see angular shapes in the branches.

Note:

Refining a previously simplified plant does not necessarily restore the initial geometry of the plant. Leaves are not affected by the simplify/refine concept, and that these operations are not stored in plant species.



- **Advanced Edition:** this icon replaces the **Export Plant** button when the plant being edited is a .tpf plant (The Plant Factory). When you click this icon, The Plant Factory opens with the plant displayed for you to make further changes to the plant. When you save the changes to the plant in The Plant Factory and re-export, the plant will be reloaded in the VUE Plant Editor.



Editing Plants

Editing plants is done by modifying the plant on a global basis, by acting upon subsets of a plant species. Although this doesn't let you act upon individual branches or leaves of the plant, nor create a new species from scratch, the fact that the plant is separated into subsets provides a very powerful means of modifying it, while maintaining a portable and generic alteration of the plant that can hence be systematized to create a whole new plant species.

Modification of the plant is done by acting upon “subsets” of the plant. Subsets are separated into two categories: trunk, branches and stems on the one side, and leaves and petals on the other side. Trunk and branch subsets appear to the left of the Editor, while leaf and petal subsets appear on the right hand side.

The number and the nature of the subsets depends on each plant species. For instance, coconut trees might define subsets for the trunk, the stump and the branches, whereas flowers might define a subset for the stems, the leaves and the petals. Some species may only have subsets in one category (e.g. the *Dead Tree* only has subsets in the branch/trunk category). Also, certain subsets may refer to parts of the plant that do not exist in a given instance of a plant. For instance, if you have a plant that randomly exhibits fruit, an instance of the plant that has no fruit will have nothing in the fruit subset.

Trunk and Branches

Description

You can modify the geometry of a plant's trunk and branches by acting upon the trunk and branch subsets of the plant. Trunk and branch subsets appear on the left of the editor. You can act on all of these subsets at the same time, or modify each one individually.

Trunk and branch subsets are modified using the controls in the **Trunk / Branches** group. You can switch from one subset to another using the subsets drop-down list, at the top of the group. The number of different trunk and branch subsets in the plant is indicated above this drop-down list. If there is only one trunk and branch subset in the plant, the drop-down list will be disabled.

Some plants have no trunk and branch subsets (e.g. the *Carex*). In this case, the controls in the group will be disabled.



All Subsets


When you first open the *Plant Editor*, the trunk and branch subset that is selected is the “All subsets” one. As its name indicates, this subset comprises all the different trunk and branch subsets of the plant. This means that any modification of that subset will apply to all the trunk and branches of the plant.

If you select another subset using the drop-down list, the modifications you make will instead only be applied to this selected subset.

Note:

Settings applied to all subsets are culminated with a given subset’s own settings to produce the expected result.

Empty Subsets

Some plants may contain subsets identified as “Empty Subset”. This indicates that the particular plant you are working on has no elements in this subset, but other plants of the same species might. Empty subsets can be modified like any other subset, but because there are no elements in it, the modifications will not be immediately visible on the plant. However, creating new variations of the plant (by clicking the  icon) may yield a plant that has some elements in the modified subset, and these elements will be affected according to the new settings. The Primrose is a typical example of a plant that features empty subsets, because different subsets are used for each basic primrose color.

Trunk and Branch Materials



Original





Falloff=100



Gnarl=300




Droop=30

Dead Tree

Immediately below the subset drop-down list is a standard material control. Using this control, you can replace or modify the current subset material. If the trunk and branch subsets of the plant use different materials and “All subsets” is selected, this control will



be disabled.

Click the **Load trunk/branchmaterial** button () to open a Material Browser, or double-click on the material preview to open the *Material Editor* and modify the material. You can also modify the scale of the material using the material **Scale** control. Please turn [here](#) for full details on editing materials.

Trunk and Branch Settings

The different settings available for trunk and branches will now be detailed:

- **Length:** this setting controls the overall length of the trunk and branch subset. Positive values will increase the length of the branches, whereas negative values will reduce this length. Zero setting means no modifications are applied to the length of the branches in this subset.
- **Falloff:** this setting controls the way the above length setting applies to the subset. The deciding parameter here is the “age” of a particular branch. The further down the tree hierarchy of the plant a particular branch is, the “younger” it is considered. Positive values for the falloff setting mean that the length setting applies mostly to young branches, whereas negative values indicate that the length setting applies mostly to older branches. Zero setting means that all branches in the subset are equally affected by the length setting. This setting is very important to fine tune the way the branches grow.
- **Gnarl:** this setting controls the amount of random curvature in the subset. Positive values will yield very twisted branches, whereas negative values will “straighten” out the branches.
- **Diameter:** this setting affects the overall diameter of the branches in the subset. Positive values will increase the diameter of the branches, whereas negative values will yield thinner branches.
- **Droop:** this is a very interesting setting in that it allows you to control the overall shape of the plant very easily. It affects the way the branches react to gravity. Positive values will curve branches towards the ground, whereas negative values will make the branches reach for the sky.
- **Angle:** this is also a setting that affects the global look of the plant. It controls the typical angle that branches make with the trunk, or the angle of the trunk with the ground. Positive values increase this angle, whereas negative values reduce it. For instance, with a *Coconut tree*, strong values mean that the trunk will start with a more horizontal direction (adding negative droop will then curve the tree upwards producing an interesting variety of coconut tree). On the opposite, negative values would cause the coconut tree to grow vertically from the ground.



Leaves and Petals

Description

You can modify the shape, color and aspect of a plant's leaves and petals by acting upon the leaf and petal subsets of the plant. Leaf and petal subsets appear on the right hand side of the editor. You can act on all of these subsets at the same time, or modify each one individually.

Leaf and petal subsets are modified using the controls in the **Leaves / Petals** group. You can switch from one subset to another using the subsets drop-down list, at the top of the group. The number of different leaf and petal subsets in the plant is indicated above this drop-down list. If there is only one leaf and petal subset in the plant, the drop-down list will be disabled.

Some plants have no leaf and petal subsets (e.g. the *Dead Tree*). In this case, the controls in the group will be disabled.

All Subsets

When you first open the *Plant Editor*, the leaf and petal subset that is selected is the “All subsets” one. As its name indicates, this subset comprises all the different leaf and petal subsets of the plant. This means that any modification of that subset will apply to all the leaves and petals of the plant.

If you select another subset using the drop-down list, the modifications you make will instead only be applied to this selected subset.

Note:

Settings applied to all subsets are cumulated with a given subset's own settings to produce the expected result.


Empty Subsets


Please refer to the section on trunk and branch subsets above for a description of what Empty Subsets are.

Leaf and Petal Materials

Immediately below the subset drop-down list is a standard material control. Using this control, you can replace or modify the current subset material. If the leaf and petal subsets of the plant use different materials and “All subsets” is selected, this control will be disabled.



Click the **Load leaf/petalmaterial** button () to open a Material Browser, or double-click on the material preview to open the *Material Editor* and modify the material. You can also modify the scale of the material using the material **Scale** control. Please turn [here](#) for full details on editing materials.

Another – easier way – of creating new leaf/petal materials is to use the **New leaf/petal** map button (). This button opens the *Leaf Editor* that provides a convenient way of mapping leaves as well as placing the point at which the leaf connects to the branch on these maps (see [here](#) for details on the *Leaf Editor*).


Leaf and Petal Settings

The different settings available for leaves and petals will now be detailed:

- **Length:** this setting controls the overall length of the leaves in the selected subset. Positive values will increase the length of the leaves, whereas negative values will reduce their length. Zero setting means no modifications are applied to the length of the leaves in this subset.
- **Width:** this is the same as the above setting, except it controls the width of the leaves instead of controlling their lengths.
- **Randomness:** this setting controls the amount of variation between the size of different leaves in the selected subset. Positive values mean that the larger leaves will become even larger, whereas smaller ones will become even smaller. On the contrary, negative values tend to reduce the difference in size of the various leaves in the subset. Zero leaves the relative sizes of the leaves untouched.
- **Flexibility:** this setting controls the overall flexibility of the leaves in the subset. Positive values increase flexibility, meaning the leaves will tend to droop towards the ground. Negative values will instead reduce the flexibility, and ultimately invert it so that leaves will reach for the sky instead of drooping to the ground.
- **Curl:** this is the same as flexibility, only acting on the width of the leaves instead of their length. Positive values will increase the curl, meaning the leaves will tend to curl around themselves and towards the ground. Negative values will instead reduce the curl, and ultimately invert it so that leaves will curl up towards the sky instead of curling down to the ground.
- **Overall color:** this is a very interesting setting as it influences the color of the leaves in the subset, letting you change the overall color of the plant in a wink. The control displays the current average color of the leaves in the subset. If you double-click on the color swatch, the *Color Picker* will appear, letting you select a new overall color for the leaves. Since this is an “average” color, color variations inside the leaves will be retained. This setting looks particularly good when animated...




Creating Variations of the Same Plant

It is often useful to be able to create a new plant of a given variety without losing other settings, such as the position of the plant, its size or its orientation. The **New variation of plant** icon () is designed for that very purpose. When you click on this icon, a new plant will be grown using the exact same settings as the current plant. But, due to the way plants are grown inside VUE, the result will be a different plant. By clicking repeatedly, you can create a set of variations of the same plant (and you can browse through that selection using the **Undo** and **Redo** icons). When you see a shape that you like particularly, press **OK** to replace the existing plant with the new variation.

Creating variations of a plant is also very useful when you are designing a new plant species. When you have found settings that produce interesting results, you can create several variations of the plant based on those settings. This is good to make sure that the settings consistently produce the results you are expecting.

Creating New Species

Once you have achieved an interesting new plant design, you can save that design as a new plant species. It is recommended that you test the new settings on a number of variations of the plant (see above), to make sure that the results are always satisfactory.

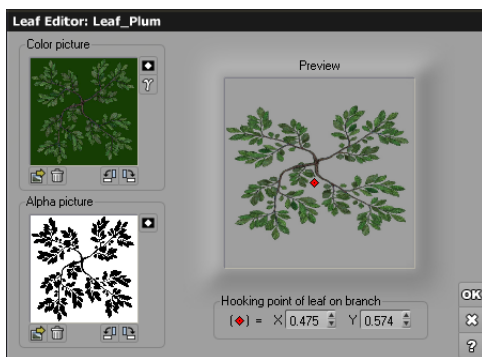
To save the new plant species, simply click on the **Save plant species** icon (). A Standard File Browser will appear, letting you select the name of the file for the new plant species, as well as its title and description. You can also opt to have all texture maps incorporated in the plant species file (convenient when you want to distribute your species). When you click OK, VUE will render a preview of the new plant species, based on the current *Plant Editor* preview. When saving the species is completed, you can check that it appears in the Visual Plant Browser, like any other plant species. Your new species is ready to breed!

Exporting Plants

You can export a plant in standard 3D file formats by clicking the **Export plant** button. This opens the *Export Options* dialog, letting you define the file format and export options, like for any other object.



Leaf Editor



Leaf Editor

This dialog prompts you to select the pictures that will be used to map the leaf. The first picture (**Color picture**) will be used to color the leaf, while the second (**Alpha picture**) will be used to define the shape of the leaf. If the picture that you select for the colors has embedded transparency information, this information will automatically be loaded into the Alpha picture.

Note:

You cannot create leaves that are partially transparent using the *Leaf Editor*. If this is something you want to do, and you are aware of the implications in terms of render times, you should use the *Material Editor* instead.

To load a picture, click the **Load** icon (📁) below the picture previews. You can rotate the pictures by using the ⬅️ and ➡️ arrows. You can also invert the pictures using the ◻️ button. This is particularly useful when the Alpha information is encoded the wrong way around. To remove the pictures, click the **Remove** icon (🗑️) below the picture previews.

The **Preview** displays a preview of what the leaf will look like.


Note:

it is important that you use a high definition picture for the alpha channel of your leaves. Using insufficient resolution will lead to aliases along the edges of the leaves. The color picture does not need to be such high resolution.




On the preview of the leaf, you will notice a small red diamond-like marker. This marker indicates the point where the leaf is attached to the branch that carries it. You can click and drag the marker to change the position, or you can use the numerical settings in the **Hooking point of leaf on branch** group. Be sure to set this point correctly, as results may look somewhat surprising otherwise...

Response To Wind Options



Response To Wind Options dialog

For VUE .veg files


The Response To Wind Options dialog appears when you click on the **Response To Wind** icon () in the *Plant Editor*. This little dialog is used to customize the way the plant reacts to wind. The settings in this dialog are not straightforward. They require trial and error before you can achieve satisfactory results.

Overall sensitivity of trunk/branches: this setting controls the rigidity of the trunk and branches to the breeze and wind. Higher values mean that the tree will be more deformed under a given level of wind.


Sensitivity of leaves to breeze: this setting lets you control the amount of movement in the leaves when subject to breeze. It has no effect on the movement of leaves in the wind. High values mean that the leaves move more in a given amount of breeze.



For PlantFactory .tpf files



Response To Wind Options dialog – for PlantFactory plants

The Response To Wind Options dialog appears when you click on the **Response To Wind** icon () in the *Plant Editor* when editing a PlantFactory plant. This dialog is used to customize the way the different plant parts react to wind.


Branch wind sensitivity: this setting controls the rigidity of the trunk and branches to the breeze and wind. Higher values mean that the tree will be more deformed under a given level of wind.

Blade wind sensitivity: this setting lets you control the amount of movement in the leaves or blades when subject to breeze.

Billboard wind sensitivity: this setting lets you control the amount of movement when billboards are used in the plant for leaves.


Using PlantFactory Plants With Presets

If a PlantFactory plant has presets and variations defined in PlantFactory 2016, such as size or shape, you can take advantage of these presets and variations directly inside VUE 2016. All presets are displayed in the Plant browser.


When loading a plant in VUE 2016 exported from PlantFactory 2016 with predefined variations a new icon button () is available in the plant editor toolbar to pick a predefined variation of the plant preset.



VUE 2016 Infinite & xStream — Reference Manual

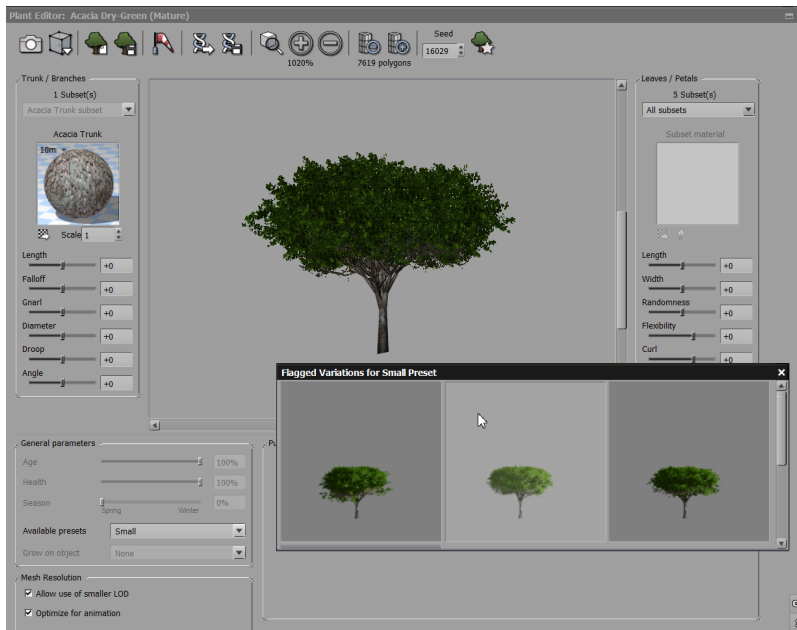


Plants with Presets



Plants with presets in the Plant Editor





Flagged variations for small preset



Atmospheres


The key to a successful picture is often the atmosphere it evokes.

The scenes you build inside VUE are part of a world. A world that extends far beyond your scene. A world with an atmosphere, with clouds, fog, haze and all... Real clouds, infinite ones, not just a backdrop picture of a sky. And this world is coherent. And best of all, you can act upon every aspect of it...

Parameters that describe the atmosphere are numerous. This section will show you through all of them. However, creating an atmosphere from scratch can be a complex and time consuming process. This is why VUE offers a list of predefined atmospheres. Choosing one of them will let you to start building your scene from a good basis. You may decide later that you want to fine tune such a setting to improve the atmosphere of your picture.

Basically, atmospheres include settings for the sky, clouds, sun, quality of light, fog and haze.

Loading an Atmosphere



The Atmosphere Browser

When you create a new scene, the Visual Atmosphere Browser pops up, prompting you to choose one of the predefined atmospheres. Select the one that is the closest to that which you want to achieve. You may load another atmosphere into your scene later, as work progresses.

How your atmosphere is previewed in the *Main camera view* can be controlled by options



on the *Display* menu. On the *Display* menu, select **Atmosphere Preview** and the submenu gives you the options to preview the following:

- OpenGL Atmosphere
- OpenGL Clouds
- OpenGL Lens Flares
- OpenGL Planets

These options can be toggled on and off.

After having worked on your scene for a while, you may suddenly realize that the atmosphere you chose at the beginning is no longer suitable. You can load another one into your scene at any time by choosing the menu command **Atmosphere | Load Atmosphere**.

If the atmosphere in your scene is basically correct, but you would like to modify certain parts of it, you will want to use the *Atmosphere Editor*.

The *Atmosphere Editor* is designed in a set of tabs. The number of tabs depends on the atmosphere model that is selected:

- **Standard atmosphere model:** this is VUE's traditional "workhorse" model. It has been widely used for a multitude of stills and animation projects. This model lets you control sky appearance through the use of color gradients. You can create an unlimited number of fully editable cloud layers, control fog and haze densities, or add special effects like twinkling stars, rainbows and ice rings. All atmospheric elements can be fully animated. The standard model's main advantages are ease of use and fast rendering.
- **Volumetric atmosphere model:** this model offers a good compromise between the standard and spectral models, giving you a higher level of realism, yet rendering faster than spectral atmospheres. Unlike the standard atmosphere, the appearance of the sky and sun is not defined by color gradients. It's directly affected by haze and fog settings and by the sun's position, much like in a real atmosphere. The volumetric model is especially suited for animations: simply moving the sun around produces beautiful color and lighting shifts.
- **Standard Spectral atmosphere model:** this is VUE's hyper-realistic model that accurately simulates the behavior of real-world atmosphere and lighting according to weather conditions. The appearance of sky, sun and clouds (both standard and spectral), as well as the character of direct and ambient light are all affected by the delicate balance between the elements that constitute the atmosphere: air, dust and water particles. The Spectral model provides its own set of controls that let you adjust each element's density and height and rendering quality. Using the Spectral model ensures a complete coherence of all the elements of a scene.
- **Environment mapping:** Especially suited for architectural visualization, this model lets you easily set up an environment based on panoramic photographs. By using



VUE's Image Based Lighting, HDRI support and global reflection mapping you can create a seamless integration between your scene and the background plates.

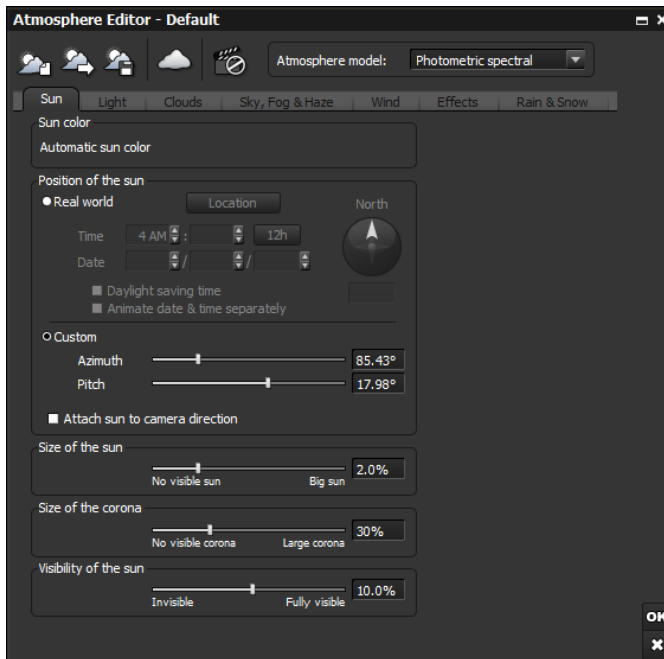
- **Standard Photometric:** Photometric lighting ensures that relative intensity of sunlight and artificial light sources is physically correct. The absolute intensity is also heavily affected, leading to much brighter daylight renders than before, thus the need for different exposure and natural film response settings that can be handled automatically by VUE internally. If you are changing between a photometric atmosphere and another type of atmosphere model, you are given the option via a popup to let VUE change your current settings to those more appropriate for this atmosphere model automatically. Or you can make adjustments yourself.

The *Atmosphere Editor* can stay open all the time without blocking access to other parts of the software. Modifications are taken into account immediately. You can reset, load and save an atmosphere by using the icons in the dialog bar (the icon bar on the lower right edge of the editor).

The following is a description of each of the tabs of the *Atmosphere Editor*.

Sun Tab





Atmosphere Editor – Sun tab

This tab allows you to adjust parameters related to the sun. It is mostly the same for all atmosphere models. If the scene doesn't contain any directional lights, this tab will not be available.

VUE considers the sun as being a pinpoint source of light located at infinity (directional light).

Sun color: this control lets you specify the color map that will be used to color the sun, from outside to inside as values increase. This option enables you to create a luminous halo effect around the sun, for example. To change the color map, double-click on it. In the Volumetric and Spectral atmosphere models, the sun color is defined automatically.

If you wish to set the sun position to exact geographic placement and time of day, the controls are available on this tab. This is available for all lighting types.

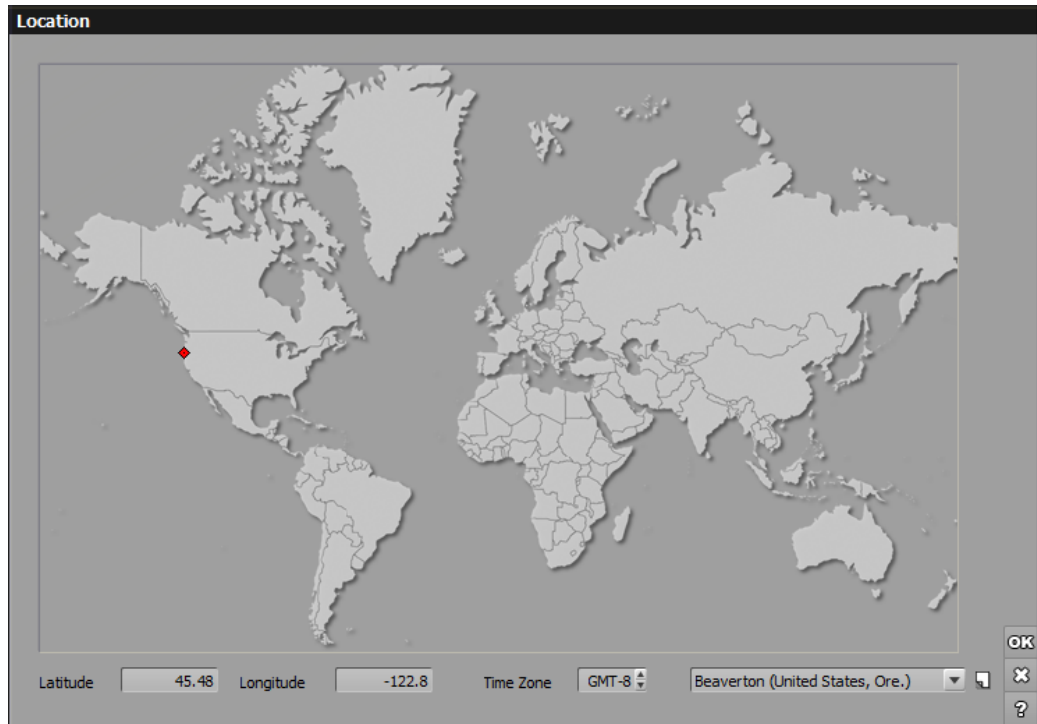
Under **Position of the sun**, click the **Real world**

- **Time:** Set the exact time using the **Hour** and **Minute** fields. In the last field, indicate whether you are using 12 hour time or 24 hour time.
- **Date:** set the date in DD/MM/YYYY format.
- **Daylight Saving Time:** click this option if applicable.



If you wish to change the position of **North**, do so by positioning the dial or enter the degree in the field below the dial.

Click the **Location** button to display a world map.



Atmosphere Editor – Sun tab-Map

On this map, click to select your location. You can also define location by precise **Latitude** and **Longitude** settings and **Time Zone**. There is also a list provided of locations. If you wish to save your location to the list, click the **Save** icon to the right of the list.

If you are not using the exact geographic positioning of the sun, check the **Custom** box. You can adjust the position of the sun using the **Azimuth** and the **Pitch** boxes. If you check the **Attach sun to camera direction** option, the position of the sun will be relative to that of the camera. Changing the heading of the camera will rotate the sun accordingly.

The **Size of the sun** control lets you adjust the size of the disk that represents the sun in the sky. If the value is non-zero, the sun will be visible, and the disk will be filled using the colors from the **Sun color** map. The color gradation ranges radially. If the sun size is zero, no sun will be visible in the sky, although it will still be emitting light. In this



case the sun color map is ignored.

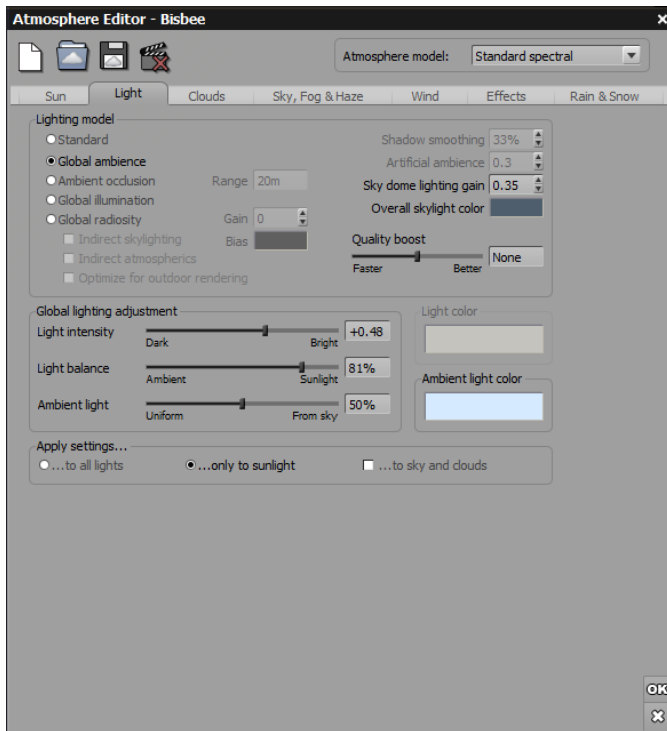
In the Standard and Volumetric atmosphere models, **Sunlight masked by clouds**, activates the masking of sunlight by low altitude clouds. The larger the value, the more the light is masked by clouds high in the sky. Because clouds diffuse light, high altitude clouds never get to mask sunlight. This effect is automatically catered for in the more advanced Spectral model.

Size of the corona: in some atmospheres, this setting is available to control the visible size of the solar disk that is added to the atmospheric glow.

The **Visibility of the sun** control, which is specific to Spectral atmosphere modes, lets you adjust the sun disc visibility as rendered within the atmosphere without affecting sky or geometry lighting. A physically correct sun disc visibility corresponds to 100%, which will render a sun disc that will produce a correct image-based lighting intensity when exporting high dynamic range skies (using HDR or EXR file format).

Light Tab





Atmosphere Editor – Light tab

There are 5 lighting models available in VUE:

- Standard,
- Global ambience,
- Ambient occlusion,
- Global illumination, and
- Global radiosity.

Each point of the scene receives light from the sun, the sky as well as from the environment (sky and surrounding objects). The different lighting models differentiate themselves by the way they estimate the amount of light coming from the environment.

In the **Standard model**, which is the most basic model available in VUE, the light coming from the environment is approximated to a constant term known as Ambient light. There is a slight subtlety in the way the standard VUE lighting model defines ambient light: you can define how much of the ambient light is actually coming from the sky (horizontal surfaces will get more of this ambient light than vertical faces) as opposed to ambient light coming from all directions. This is done using the Ambient light slider in the Global



lighting adjustment group (see below).

The **Global ambience** model offers a slightly more elaborate estimation of the light coming from the environment: this model takes into account the color of the sky in all directions. As a result, parts of the scene that look towards blue sky will take on blue shades of light, whereas other parts looking towards red sky will take on red shades of light. The global ambience model will add an interesting touch to your renders while requiring very little rendering overhead.

Ambient occlusion is an improved version of global ambience where each point on the sky dome is considered like a little source of light. Rays are traced towards each one of these lights, to see whether a neighboring object is occluding the light or not. This results in very subtle shadows appearing around objects that are close to one another. Obviously, tracing all these rays increases rendering times significantly; the effects of ambient occlusion are particularly noticeable and pleasing on areas of the scene that are not directly exposed to sources of light. Also, since ambient occlusion computes ambient lighting, it is usually recommended that you increase the contribution of ambient light in your scene when using this model. In order to speed up the rendering process, only neighboring objects that are closer than a given distance are taken into account in the occlusion process. Neighboring objects that are far away contribute less and less to the occlusion. This is a trick that enables the scene to be rendered much more rapidly than with Global illumination, without compromising too much on quality, because the renderer doesn't need to examine the entire scene to find occluders.

When the ambient occlusion model is selected, the ambient occlusion **Range** parameter becomes available. This controls the maximum distance beyond which objects will not contribute to the occlusion. The bigger this value, the closer you get to the Global illumination model and the slower the render. The smaller the value, the closer you get to global ambience (and the quicker the render).

The **Global illumination** model improves over the ambient occlusion model by tracing light rays all the way to the sky dome, thus ensuring that any object will cast ambient shadows onto other objects, whatever the distance. The result is usually darker than results achieved with the ambient occlusion model. Again, tracing all these rays increases rendering times significantly; the effects of global illumination are particularly noticeable and pleasing on areas of the scene that are not directly exposed to sources of light. Also, since global illumination computes ambient lighting, it is usually recommended that you increase the contribution of ambient light in your scene when using the global illumination model.

Note:

When using ambient occlusion or global illumination, it is essential that you increase the proportion of ambient light in your scene – otherwise you will barely see the effects of the advanced illumination model. You can actually even get very pleasing results with only



ambient light (this will result in a very foggy and overcast look).

The **Global radiosity** model is the ultimate model in terms of quality of illumination and realism. It propagates light in the scene, instead of propagating shadows as the ambient occlusion and global illumination models do. With this model, objects that are exposed to light will reemit some of that light in all directions, according to the optical properties of their surface. Light will thus “bounce around” repeatedly in the scene, as it would in reality. As a result, each point in the scene receives light from all the other objects in the scene. Obviously, this results in extremely complex computation, and, despite the numerous optimizations implemented in VUE, will lead to render times that are an order of magnitude slower than the standard model – but will also yield incredibly pleasing results. In this mode, the ambient vs. direct lighting slider controls the relative influence of light coming from the sky, versus light coming from light sources such as the sun.

VUE has the ability to preserve the indirect lighting calculation “in-between” renders, even if the scene has been modified. This speeds up dramatically the “tweaking” phase of scene preparation. When this mode is enabled, you can easily request the updating of [indirect lighting](#) next time you render, so that it matches any changes made to the scene.

Note:

When using radiosity, be aware that materials containing luminosity or that have non-standard (60:40) proportions of ambient diffuse light may cause strange lighting effects. These materials may have to be adjusted to achieve the atmosphere effects you desire.

When the global radiosity model is selected, some controls that are specific to this model become available:

- **Indirect skylighting:** when this option is selected, VUE will evaluate the amount of skylight that is received by each object and cast back onto the other objects in the scene. If this option is not selected, the Ambient light color will be used instead of computing the indirect contribution of the skylight. Evaluating the indirect lighting caused by skylight is a slow process. Using the ambient color instead usually yields good enough results.
- **Indirect Atmospheric:** If you wish to take into account the light being reflected from clouds onto the objects in the scene, check this option to account for this effect when calculating indirect lighting.
- **Optimize for outdoor rendering:** When this option is selected, VUE assumes that you are rendering an infinite outdoor landscape. While radiosity usually has a very strong influence for indoor rendering, due to light being trapped into a room and bouncing several times around, it is generally much less noticeable for outdoor scenery due to the light quickly escaping the geometry towards the sky and very rarely getting trapped enough to produce a high contribution. Therefore, this option will lower the order of indirection for radiosity calculations, effectively ignor-



ing highly indirect lighting contribution, thus producing a faster and more robust render.

- **Gain:** this setting controls the intensity of the light that is scattered in between the objects.
- **Bias:** if you define a bias color, this color will be added to the light that objects receive from their environment. For instance, if you add a slightly red color, the shadows and light will take a very slight reddish tone. This setting should only be used for very fine tuning of the effects of radiosity.

When one of the global lighting models is selected, the controls in the right half of the Lighting model group become available. These controls are used to fine tune the effects of the lighting model:

- **Shadow smoothing:** this parameter is available for all lighting modes above global ambience. It is designed to control the overall smoothness of global illumination shadows. Low values will produce sharper and more accurate shadows, but may require higher quality settings to avoid noisy shadows. On the other hand, high values will tend to smooth out shadows, leading to less accurate results, but without the need for high quality settings.
- **Artificial ambience:** this parameter is available in the ambient occlusion and global illumination models. It is designed to compensate for the fact that there is no inter-object light reflection in these models. The indicated amount of ambient light will be added to the sky's contribution to determine the total amount of light that each point receives from its environment. The color of this term can be controlled using the ambient light color setting.
- **Sky dome lighting gain:** this parameter controls the overall intensity of the light received from the sky. Increasing this setting adds more ambient light to your scene. It is somewhat similar to dragging the light balance control towards ambient and increasing the exposure of the scene globally. This control does offer an added level of flexibility, though.
- **Overall skylight color:** this color control represents the overall color of the light coming from the sky dome, and lets you adjust it in order to fine tune the ambient lighting of your scene (double-click on the color control to edit the color). For instance, if you feel that the parts of your scene in the shadows are taking a color tone that is too pronounced, you can reduce the saturation of the overall color. Because this color control represents the overall color of the sky, if you modify the settings of the sky, the color displayed in the control will change. However, the color correction that you indicated by modifying the previous color will still be applied to the new color.
- **Quality boost:** this setting is part of the *EasyGI™* technology that synthesizes the complex settings required to efficiently render global illumination into one single quality setting. As with other quality boost settings throughout VUE (such as



the volumetric atmosphere quality boost), this setting is used in conjunction with the **Advanced effects quality** setting in the *Render Options* dialog (see [here](#) for details). When you are putting the scene together and creating **Preview** renders, the quality of the global illumination evaluation is rather crude, but as you decide that the scene is ready for rendering in **Final** mode (see the preset render quality settings of the *Render Options* dialog), the quality of the computation of global illumination is automatically increased to produce nicer results. The **Quality boost** setting should only be used if you notice that there are some imperfections in the quality of the illumination in the final production render. Alternatively, if you are doing a lot of test renders where the quality of the global illumination is not essential, you can reduce the quality boost setting in order to accelerate the render process.

Global Lighting

The controls in this group let you adjust the distribution of light throughout your scene. If you are using a global illumination or global radiosity model, it is recommended that you increase the proportion of ambient light in the scene, in order to make the effects of this global illumination more visible.

You may adjust the overall luminosity of your scene using the **Light intensity** control.

Note:

Exposure only affects the sources of light, and, as such, is different from the exposure control available in the camera's *Object Properties*.

The correction is expressed in diaphragms.

- **Light balance:** lets you adjust the relative quantities of light coming from the sun and from the environment (ambient light). Scenes with a bright sky will have lots of ambient light, whereas sunset scenes should have little.
- **Ambient light balance:** lets you further customize ambient light by deciding how it is shared between light coming from the sky, and light coming from everywhere (uniform light). Scenes with fog will usually require larger amounts of uniform ambient light, whereas scenes with a bright sky will have lots of light coming from the sky.

Lastly, you get the choice to **Apply these settings** either to all the lights, or only to sunlight. Light color and exposure affect the color and intensity of light sources. If lights in your scene (other than sunlight) are used for the atmosphere, you should apply the settings to **all lights**. Alternatively, if you have for instance a house that is lit up from the inside, you will not want that light to be affected by changes in the exposure outside of the house, so you should select **Only to sunlight**.

If you check the **...to sky and clouds** option, the lighting adjustments will also affect the color of the sky and clouds.



Light Color

You can give a different color temperature to the sun light and to the sky / ambient light. Double-click on the colors to adjust them.

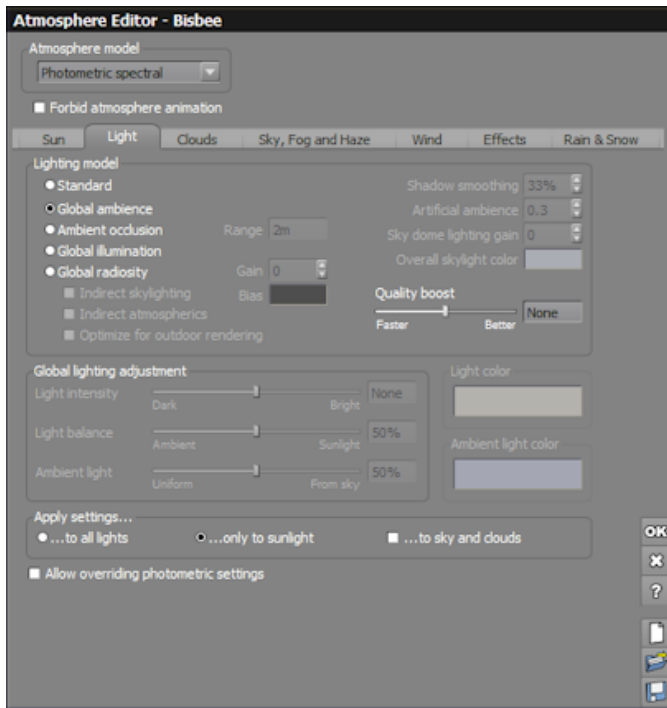
- **Light color:** acts as a filter for the color of the lights in your scene. Having an orange light color, with a sun color that is green will make the light coming from the sun a darker shade of green. It is a natural phenomenon that sunlight gets warmer (i.e. takes yellow to orange shades) as the sun gets closer to the horizon. Such warm lights can yield pleasing results when exposed objects take on warm shades. This setting does not apply to Spectral atmospheres. To change light color, modify the color of the sun light.
- **Ambient light color:** adjusts the color of ambient lighting, whether light comes from the sky or from the environment. Since ambient light is diffused by the clouds, it usually has a cooler tone (shades of blue) than that of direct sunlight.

Auto Decay Sunlight Color

This setting is available for Standard, Volumetric and Environment Mapping atmospheres. For Spectral and Photometric atmospheres, the auto decay color and settings are handled on the [Sky, Fog & Haze Tab](#). The options in this group control the way the color of light reddens as it gets closer to the horizon. This effect should not be disabled with the Volumetric atmosphere model, because with this model, light is affected by the atmosphere as it travels through it anyway...

Light decay color: this is the color that is used to make the color of light turn to red. By default, the color is R=218, G=171, B=71. Preferably, you shouldn't modify these values, because if you do, you'll probably get unexpected (and unrealistic) results. Entering a different color will completely affect the way light is colored by the atmosphere. Can be cool for alien planets!





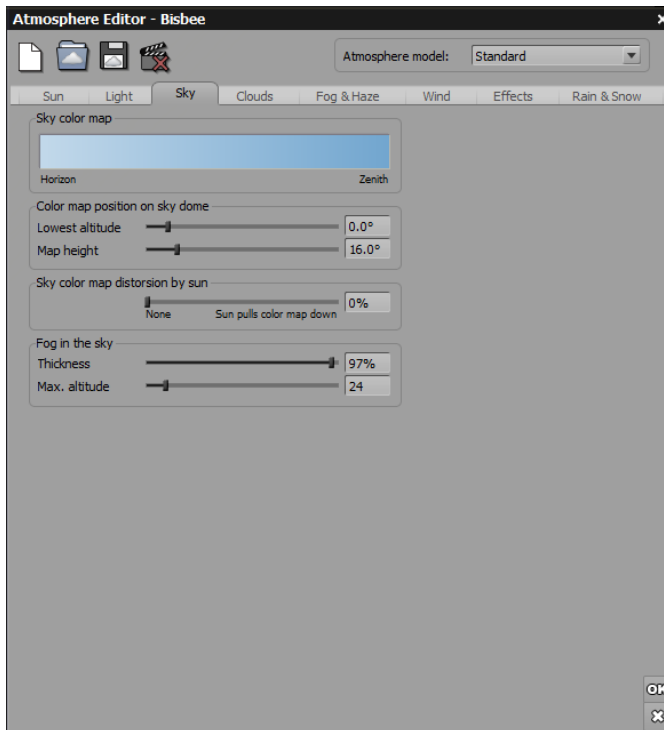
Atmosphere Editor – Light tab – Photometric Settings

Changes for Photometric Atmospheres

If the atmospheric model selected is Photometric and you select to let VUE control the settings for this model, several fields on the *Light Tab* that can produce unrealistic settings are grayed out and not selectable. A field has been added at the bottom of this called **Allow overriding photometric settings**. If you check this settings the grayed out fields become available for use again.



Sky Tab



Atmosphere Editor – Sky tab Standard model

This tab deals with the colors of the sky in the standard model of atmosphere only. It isn't available in the other models.

The most important control is the **Sky color map**. This is where sky colors are defined. To edit the color map, Control-click on it, which will open up the *Color map editor*.

The color of the sky is generated by vertically mapping these colors. The resulting gradation occurs inside a strip that is pulled down by the sun. This means that colors on the left side of the map will show close to the sun, while colors on the right side of the map show towards the zenith, and it also means that as you get higher in the sky, colors further to the right of the gradation appear.

The two **Color map position on sky dome** controls let you adjust the altitudes at

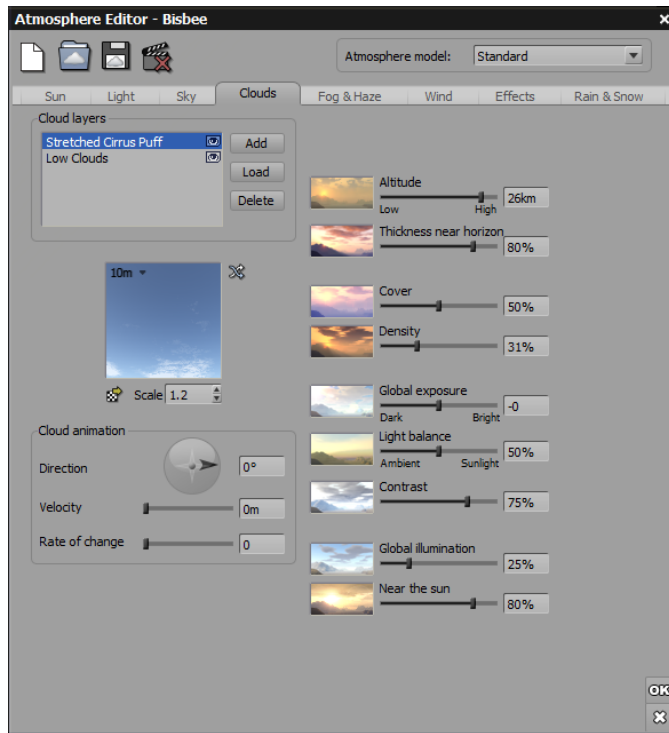


which the gradation starts and ends in the sky. Bear in mind that these altitudes will be modified by the presence of the sun, so probably some experimentation will be required here.

Sky color map distortion by sun controls the effect the sun has on the **Sky color map**. The higher the value, the more distorted the gradation becomes, eventually ending up in circles around the sun. As you get closer to the sun, colors to the left of the map get displayed.

Near the horizon, the sky is often a different color. It is the color of the fog (and haze), creeping up into the sky with distance. You can adjust this effect using two controls: **Thickness** and **Max altitude**. Thickness controls the maximum density of fog achieved at the horizon, and Max altitude controls how high up in the sky the fog is seen.

Clouds Tab






Atmosphere Editor – Clouds tab

This tab is available for all atmosphere models except Environment Mapping.

Clouds are generated inside flat layers, realistically positioned and mapped at different altitudes around the earth. You can create as many layers of clouds as you want. These cloud layers show up in the *World Browser* as objects. Each layer is displayed individually and can be grouped together for ease of handling.

At the top-left of the tab, you will find a list of all the cloud layers in the atmosphere. To edit one of the cloud layers, select it by clicking on the layer name, and use the controls in the tab to edit the cloud layer settings.

Alongside each layer name, you will notice a  icon. Click on this icon to hide the layer; the icon changes to  indicating that the layer no longer appears in the sky. Click again to return the layer.

To add a new cloud layer, click the **Add** button, or click the **Add Cloud Layer** () icon in the left toolbar; the Cloud Layer Browser appears to let you select an existing layer. Cloud layers are sorted by altitude.



To delete the cloud layer that is currently selected, press **Delete**.

The details of each layer are displayed one at a time.

For each cloud layer in the atmosphere, the **Clouds** tab offers a set of controls that you can use to customize the looks of the layer.

A small preview window shows the typical effect that the setting has on the cloud layer.

The two most important features about a cloud layer are the **Altitude** of the layer, and the material of the clouds. Altitude is pretty straightforward. The slider covers “usual” altitudes, but any value can be indicated.

The current material used for the cloud is displayed in the preview window. To modify it, you can either load a new one (use the **Load material** button ) , or edit the current one (double-click on the preview of the material to open the Material Editor). When you load a new cloud material, the settings for that layer of clouds will be modified to reflect the settings of the new cloud material. You can also click the **Randomize** () button to change the cloud distribution to vary the cloud effect.

The preview of the cloud material is refreshed each time you modify an aspect of the cloud layer.

An important aspect of cloud materials is transparency. If the cloud material is opaque,



you will never get to see the cloud layers above or the sky through it. So it should be transparent in places. Varying **Global transparency** of a cloud material is a good way of modifying the thickness of the clouds. In the same way, using the **Transparency filter** is a good way of varying the density of the clouds.

You can change the scale of the clouds by using the **Scale** control.

Close to the horizon, the clouds tend to fade out, dissolving in a general haze. If you are using the **Standard** or **Volumetric** atmosphere models, this effect can be captured using the **Thickness near horizon** control which handles the speed at which clouds disappear as they get close to the horizon. In the **Spectral** atmosphere model, this effect is achieved automatically, and is replaced by a **Thickness** setting that lets you control how thick the cloud layer is in terms of altitude (the total height of the layer).

The **Cover** setting controls the overall amount of clouds visible in this layer, and the **Density** setting controls the overall density of the clouds that are visible.

Density: this setting is available only when using the spectral atmosphere model. It controls how far light penetrates inside the cloud and how the light is scattered inside the cloud.

Opacity: this setting is also available only when using the spectral atmosphere model. It controls how far you can see objects through the cloud. If the cloud is very opaque, objects inside the cloud will rapidly become invisible.

Adjust the **Global exposure** of the clouds using the slider to darken or brighten the clouds, and use the **Contrast** setting to adjust the contrast between dark and bright areas in the clouds. The **Light balance** control is only available if the **Apply to sky and clouds** option is not selected in the **Light** tab. This control lets you adjust the balance between ambient and sunlight in each layer independently. These 3 controls are only available in the standard and volumetric atmosphere models.

If the spectral model is selected instead, the following controls are available:

Detail amount: this controls the amplitude of the cloud detail. High values will create clouds with a lot of inner density variations, while low values will keep the clouds smooth and compact.

Altitude variations: this setting controls the amount of variation in the altitude of the cloud layer. Altitude variations are particularly noticeable at the top of the layer. This setting affects the cloud layer on a large scale only.

In the sky, clouds close to the sun are usually more illuminated than clouds on the opposite side relative to the sun. You can change this illumination using the **Global illumination** and **Near the sun** illumination parameters. Maximum illumination is achieved for clouds that are close to the sun.



Ambient lighting: this setting is only available in the spectral atmosphere model. It controls the amount of ambient light inside the cloud. The higher the value, the brighter the cloud will appear.

Shadow density: this setting is only available in the spectral atmosphere model. It controls how dark the shadows cast by the clouds are. This influences both the density of the shadows cast by the clouds onto the scene, and also the density of the Godrays (if this option has been enabled in the **Sky, Fog and Haze** tab – see [here](#)). It should be set to 100% in order to achieve Godrays that are as visible as possible. If the cloud layer has been made to not cast shadows, this setting is disabled.

If the Spectral II model is selected, the following controls are also available:

Sharpness: this controls how sharp the edges of the cloud appear. High values will create clouds with sharp edges. Low values will create smoother cloud shapes.

Feathers: this controls how “feathered” the clouds look. High values will create clouds with lots of filaments. Low values will create smoother cloud shapes.

Simple Cloud Animation

The **Cloud animation** group provides straightforward controls for basic cloud animation (movement and evolution): using **Direction** and **Velocity**, you can make your cloud layer drift in the wind. The **Rate of change** control is used to set the evolution rate of the cloud layer (whether the shape of the clouds evolves slowly or rapidly over time).

Volumetric Clouds

Volumetric clouds are a special kind of volumetric material that is optimized to accurately simulate real clouds. They are edited mostly in the same way as standard volumetric materials, except that some settings are locked. Please turn [here](#) for details on editing volumetric materials and clouds.



Spectral Clouds

Spectral clouds are a special kind of cloud layers that have a thickness. What this means is that the cloud layer is, just like in reality, a thick mass of air and humidity that derives its look from the way the light travels inside the material and interacts with the cloud particles.

Spectral cloud materials are edited mostly in the same way as standard volumetric materials, except for a couple of additional settings:

Height: this setting controls the overall height of the cloud layer. Higher values means the cloud will be higher, and consequently darker.

Shadow density: this setting controls the density of the shadows cast by this layer of clouds, both onto the ground and in the atmosphere (to create the famous Godrays effect).

Note:

In order to view Godrays in your scenes, you need to enable the Godrays option in the **Sky, Fog and Haze** tab, as well as setting the Godray intensity in the *Advanced Cloud Material Editor*.

Volumetric materials for clouds can be also be edited further in the *Advanced Material Editor*.


Cloud Layers as Objects

When you create a cloud layer, it appears in the *World Browser* as an object. The layer is also visible in the viewports, for example, the *Side view* if you zoom out a bit. You can select it and modify the origin, altitude, the height, overall size and the rotation all within the viewports. You can also animate the cloud layer.

As with other objects in the *World Browser*, you can hide the cloud layer object from render. And cloud layers can be grouped together for ease of movement and altitude control.



Cloud Layer Zone




Cloud Layer Zone

Spectral cloud layers can now be restricted to a user-defined circular zone by selecting the cloud layer in the *World Browser* and setting the new controls in the *Object Properties* area. The zone actually removes all the clouds outside the zone, leaving only the portion inside it at render.

Once activated, you can switch to the zone from the object properties and manipulate it in the viewports (you will see a dotted line cylinder corresponding to the zone and its falloff region, and a manipulation gizmo at the center of it).

To use this feature,

- Select a cloud layer in the *World Browser*.
- In the *Object Properties panel*, a new **Move cloud layer/Move limiting zone** icon  is available.
- Click to activate the restriction zone and display the **Cloud Layer Zone**. You can enable or disable the zone and edit the **Falloff** value. Selection can be switched back to the cloud layer by clicking on the **Move cloud layer/Move limiting zone** icon.
- In the view ports, the zone can be scaled within the cloud layer.

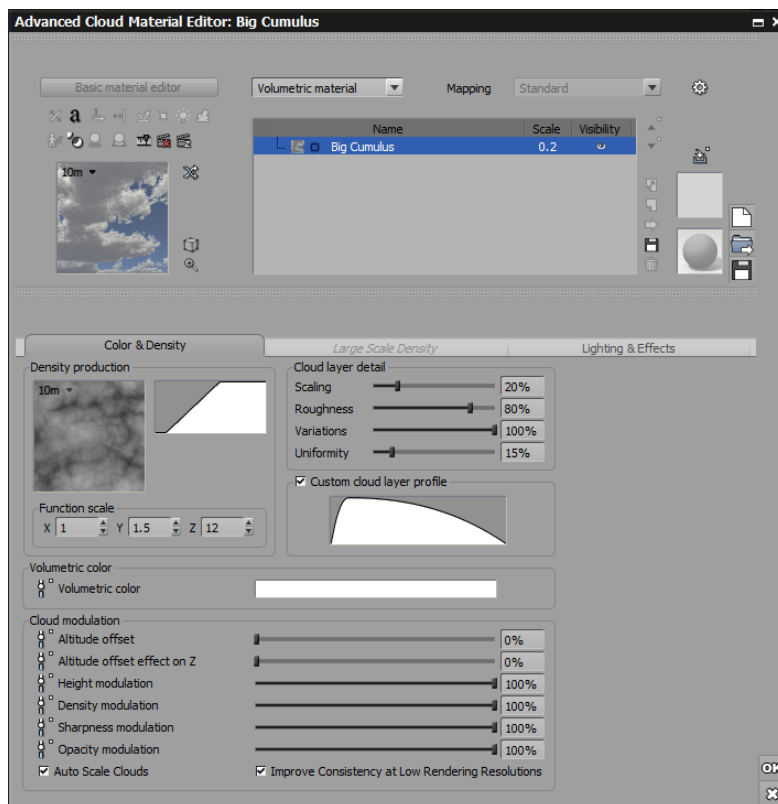


Use of these zones can greatly speed up rendering of cloud layers outside the zone since clipping can be performed at render time. It may be worth isolating the interesting part of a cloud layer with a limiting zone, to avoid rendering unnecessary parts of the layer.

Advanced Cloud Material Editor

In the Atmosphere Editor, the Clouds tab is where you can create cloud layers and edit clouds. When a cloud layer is added, a picture displays of the clouds being generated. If you wish to change the material and density of this cloud layer, click on the picture and the **Advanced Cloud Material Editor** displays.

Color & Density Tab



Color & Density Tab

Volumetric materials are based on a density production that indicates the local density



of the material over a cloud layer.

Density Production

This is how **Density production** works: for each point of the volume, the function returns a number in the range of 0 to 1 (0 appears black on the preview and 1 white). The number is then transformed by the filter into a density at this point (0 if the material doesn't exist, 1 if the material is solid). When rendering the material, VUE accumulates the density of material all along the ray of light that is traversing the material, and then computes the resulting color of the material.

To modify the function, double-click on the picture of the function. This will open the Visual Function Browser.

Use the scaling controls to scale the function along the **X**, **Y** and **Z** axes. The same for the filter.

Cloud layer detail


If the edited material is used for a Spectral cloud layer or a MetaCloud, the detail settings are available; these settings control the amount of detail that is visible in the cloud.

- **Scaling:** this parameter controls the overall scale of the details. Higher values will produce more elaborate cloud shapes.
- **Roughness:** this parameter controls the “feathering” in the cloud shape. Typical values should be below 0.5. Higher values will produce a lot of feathering around the cloud.
- **Variations:** this parameter controls how much the roughness effect varies throughout an entire cloudscape. The effect of this setting is not visible at small scales, but will add large-scale variety to your cloudscape.
- **Uniformity:** this parameter controls the variations in roughness according to altitude inside the cloud. High values mean that the cloud will be uniformly rough, whereas low values will create clouds that are rounder and less feathered underneath.


Cloud Modulation

Functions are available to adjust the properties of the layer. Contrary to **Density production** function which takes a local position impossible to translate to world, these functions could take a world position or a local position in entry. When you choose to link these function to the graph, you have a checkbox that allows you to choose between global or local coordinates.





Simple altitude offset over a terrain; x is the altitude variations value between 0 and 1



Altitude offset and height modulation over a terrain; x is the altitude variations value between 0 and 1; $h(pos)$ is the local height modulation between 0 and 1

There are five functions available:

- **Altitude offset:** The entry point in the cloud layer is offset by this function. If the altitude variation value is x , we divide the height H of the layer in two parts, one of $xH / (1 + x)$ on which the offset factor is applied and which gives the offset, and one of height $H / (1 + x)$ which defines the real height of the layer. The summary diagrams illustrate this.
- **Altitude offset effect on Z:** The position inside the layer has the form $(x, y, z + \text{altOffset})$. **Effect on Z** is a slider affecting the `altOffset` so as to replace the position by $(x, y, z + \text{altOffset} * \text{effectOnZ})$. The use of this slider is to negate the effect of the deformation given by offset, this way clouds will look as if they gently slide over the deformation when `EffectOnZ` is 0%. When it's 100% clouds will look



as if appearing from the contour of the deformation.

- **Height modulation:** The height of the layer is modulated by a $[0; 1]$ factor.
- **Density modulation:** This modulation can act as a helper for the density production. When the value is at 100% the cloud aspect matches the parameter entered in the Atmosphere Editor. When reaching 0% we move towards a soft and feathered cloud. It can help make the clouds realistically disappear when you want them to fade out.
- **Sharpness modulation:** This modulation can vary the sharpness of the cloud. Moving toward 1 gives you a sharper edged cloud.
- **Auto Scale Clouds:** This scales your clouds as the render size changes.
- **Opacity modulation:** This modulation can vary the opacity of the cloud. Moving toward 1 gives you a darker cloud. (it replaces the opacity parameter from the atmosphere editor.)
- **Auto Scale Clouds:** This scales your clouds as the render size changes.
- **Improve Low Quality Consistency:** This option allows the cloud layer to have a better consistency between low resolution render and FullHD renders. Unchecking this option will give you the old render. Every new scene will have this option activated by default on each layer. The option is here not to break the rendering of old scene. This option works on spectral atmospheres, both Photometric and Standard.

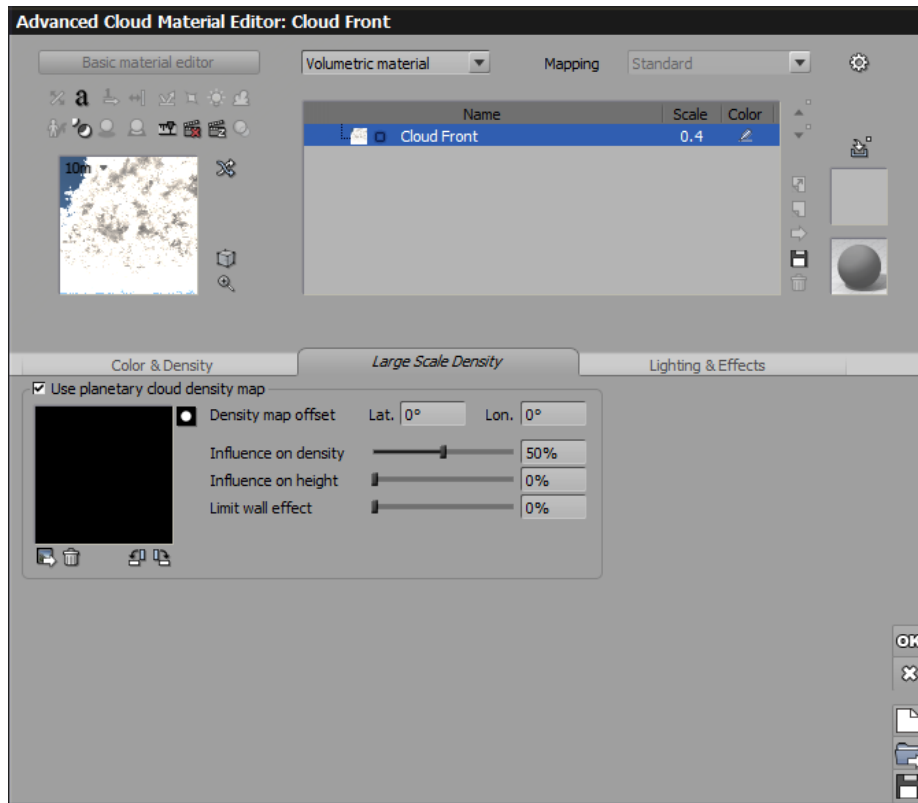
Other settings

Custom layer profile: When selected, this graph shows the profile of the cloud layer. Click on the graph to open the graph editor to make any changes you need.

Volumetric color: Volumetric color of clouds can be linked to a graph instead of using a constant color. This works with the Ambient color contribution from sky and sun found on the **Lighting & Effects Tab**. On this tab, you can drive this volumetric color with a function. If you want a constant color, click the color square to pick the desired color.



Large scale density Tab



Large Scale Density Tab

The cloud settings on this tab are for the use with planetary terrains . You select if you want to view the entire from space or closer to the surface of the terrain.

Use planetary cloud density map: Check the option and use the Load button to load a planetary cloud density map. A map must be loaded to activate the settings on this tab. For proper planetary cloud mapping, the picture used should have been generated using cylindrical or Mercator projection (where the poles are spread over top & bottom sides of the picture). This prevents severe distortions at render time.

In addition to static maps, it is also possible to load image sequences or animation files within the spectral cloud layer's planetary cloud density maps. Simply load the desired sequence or animation file in the density map field.

Influence on density: The higher this setting is, the more visible the clouds will be

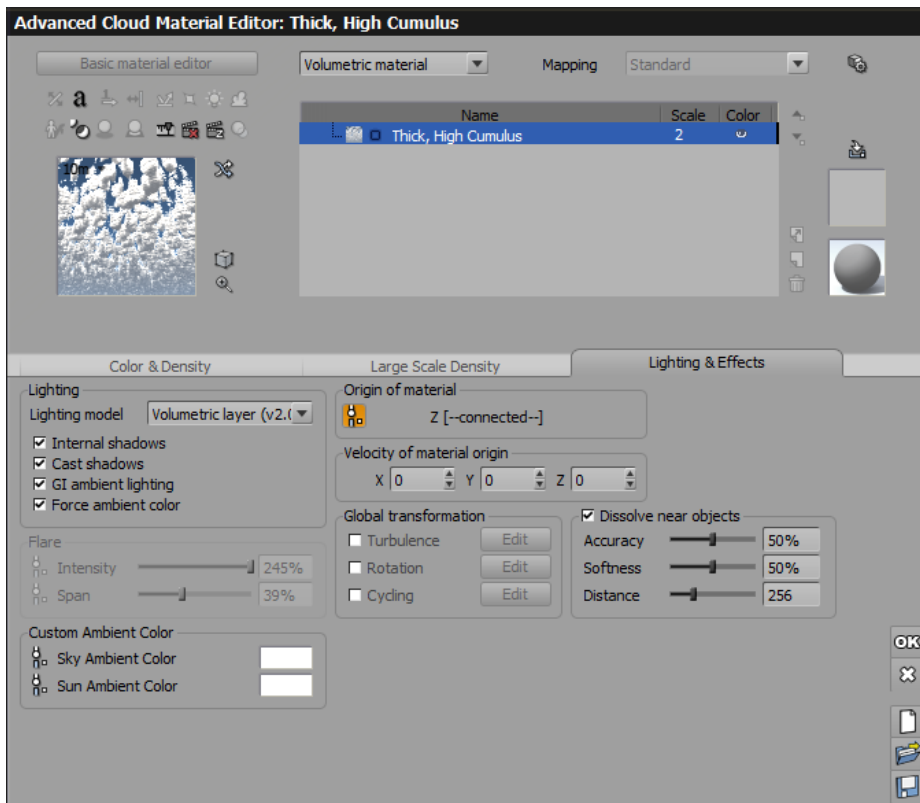


from space. Balance this with **Influence on height**.

Influence on height: This is used to drive altitude variations of the cloud layer using the planetary map. Bright regions being higher than dark ones. This works along with the base altitude variations setting of the cloud layer, both effects are blended according to the influence amount.

Limit wall effect: A wall effect can happen when you use a map to define the density of clouds. When you go from a density of 0 to a density of 1, this results in a “wall of clouds” suddenly appearing at the transition. The setting allows to limit this effect by automatically detecting such transitions and recreating a more realistic occurrence of clouds.

Lighting & Effects Tab



Lighting & Effects Tab



Lighting

The controls in this group let you customize the way the material reacts to light.

Lighting model: use this drop-down list to select the lighting model used for the volumetric material.

- **Flat layer:** this lighting model is the simplest. The color of the material is uniform, and only depends on the density. Light is not taken into account when computing the color. This is the default.
- **Volumetric layer (v.1.0):** in this model, the influence of light and shadows is computed at the surface of the material. The color of the volumetric material is added
- **Volumetric layer (v.2.0):** this is similar to the Volumetric (v.1.0) model.

If the selected lighting model is the Volumetric layer, the following checkboxes become available.

- **Internal shadows:** when this option is selected, the shadows inside the cloud will be computed, meaning that some parts of the cloud will cast shadows on other parts of the cloud. This option produces much more realistic clouds, but increases render times significantly.
- **Cast shadows:** when this option is selected, the cloud layer casts shadows in the atmosphere that can result in the appearance of Godrays if conditions are favorable. This option also increases render times significantly and should be used with caution, as it does not necessarily produce a noticeable improvement in picture quality. In order for Godrays to be visible, you must also enable Godrays in the *Atmosphere Editor*.
- **GI ambient color:** This turns on the GI ambient lighting.
- **Force ambient color:** Checking this option activates *Sky Ambient Color* and *Sun Ambient Color*. These colors can be set in the color squares next to the field. These also can be driven by a function. Be sure to set the Volumetric color on the **Color and Density Tab**. The actual color of the cloud is a result of complex light filtering, based on Volumetric color and Ambient colors. *Sky Ambient Color* and *Sun Ambient Color* settings modify the contribution of the sky and sun to the final color.

Flare

The controls in this group are only active in Flat Layer models.

When light is seen from behind a thin volumetric material, it will cause the material to become very bright. This is called flaring. Flaring doesn't occur when the material is either too dense, or too thin.

You control flaring through two settings: **Intensity** and **Span**. Flare span is the area



around the light that will flare-up. Larger values yield bigger flares.

Origin of Material

These fields let you offset the material in material coordinate space. This enables the precise positioning of clouds. By default, this settings is linked to the function graph and follow the cloud layer dummy so that you can move the cloud by selecting it in the world browser.

Velocity of Material Origin

These fields let you define a displacement over time of the origin of the material. As a result, the material will be changing as time passes.

Global Transformation

Selecting options in this group will apply global modifications to the material's density production. These options work the same as for Simple materials: when you select an option, the corresponding **Edit** button becomes enabled. Pressing this button displays a dialog that lets you adjust the effects. The editor dialogs for each type of modification are detailed in the section on [Simple Materials](#).


Dissolve Near Objects

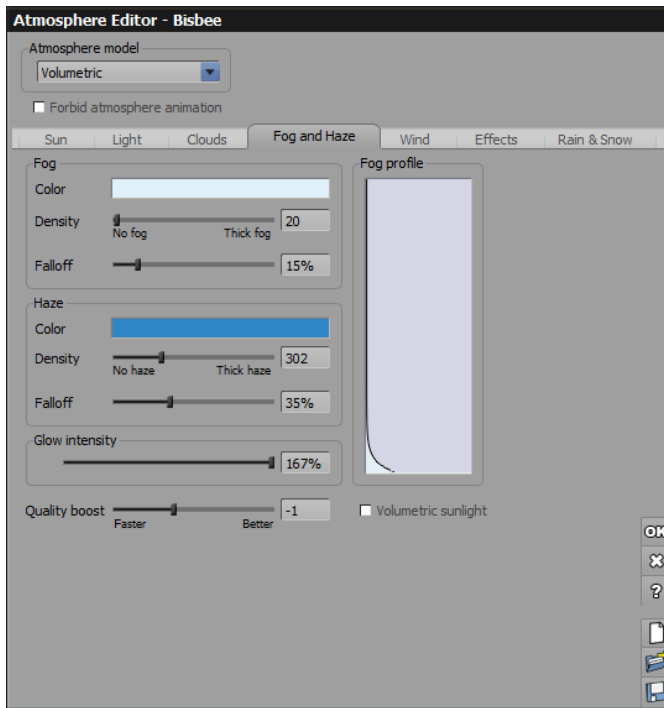
If the edited material is a cloud layer material, this group of controls is available. These controls let you automatically define how cloud layers react to the proximity of other objects in the scene – for instance to let high mountains peak through the clouds.

- **Dissolve near objects:** select this option to have the density of volumetric cloud layers drop automatically near the objects in the scene.
- **Accuracy:** this setting controls the precision with which the proximity to other objects in the scene is evaluated.
- **Softness:** this setting controls how gradually the cloud dissolves near foreign objects. Low values mean the clouds will vanish abruptly near objects.
- **Distance:** this setting controls how far away from the objects the clouds are influenced. You can control the amount of dissolving using a combination of both this and the Softness setting.



Fog and Haze Tab





Atmosphere Editor – Fog and Haze tab Standard (top) and Volumetric (bottom) models

The **Fog and Haze** tab is available in the standard, volumetric and environment mapping atmosphere models. In the spectral and photometric models, it is replaced by the **Sky, Fog and Haze** tab.

Although you may think fog and haze are used only on special occasions (e.g. to achieve particular photographic atmospheres), this is not the case. Whatever the weather conditions (unless you are out in space), you will find that fog and haze are always present. Fog and haze are what give its color to the sky. What changes is the distance at which they become significant. Fog and haze are important for fine tuning the atmosphere, because they give an idea of distance. This is why nearly all the predefined atmospheres have some amount of fog and haze.

Because fog and haze are responsible for the color of the atmosphere, they are an essential component of the volumetric atmosphere model. By adjusting the density of the fog and haze you will adjust the color of the sky.

This tab looks somewhat different depending on the selected model of atmosphere.



Fog

Fog is a generic term that covers all types of particles that you find in the atmosphere and that are large enough to reflect light (i.e. larger than the average wavelength of light). This is, in particular, the case of droplets of humidity, but also dust, crystals of ice, etc.

Objects tend to gradually disappear into fog as they move away from the camera.

The further the objects, the more the color of the objects will blend into the color of the fog. There are two types of fog in the standard atmosphere model:

- uniform fog that has constant thickness whatever the altitude,
- altitude dependent fog that has a density that varies exponentially with altitude.

In the volumetric atmosphere model, only the second type of fog is available.

Select the fog **Color** by double-clicking on the color box. The color editor pops-up, letting you select the new color.

Density is the distance at which objects totally disappear inside the fog, regardless of altitude.

Falloff controls the way that fog gets gradually thinner with altitude. The greater the value, the more rapidly the fog density decreases with altitude. In the standard model, the extra fog controls are enabled only if the fall off rate is non zero.

To help you in understanding how the thickness of the fog works, the curve on the right displays a **Preview** of fog thickness relative to altitude.

Standard Atmosphere Model Only

The following controls are available only in the standard atmosphere model:

- If you want the fog to be accumulated in the lower part of the scene (which is usually the case), select the **Fog gathers at low altitudes** box; alternately, if you want fog to accumulate at high altitudes (for instance to render mountains lost in high altitude fog or a smoke filled cave, for that matter), select the other box.
- Indicate the **Altitude** at which altitude dependent fog achieves maximum density. This controls the altitude of the “layer” of fog created by altitude dependent fog. In practice, it is somewhat tricky to use.



Haze

Haze is particularly strong on hot days. It is caused by light being scattered in all directions when it collides with the very small particles in the atmosphere (molecules of Oxygen and Nitrogen mainly). This is known as Rayleigh scattering, and is the reason why the sky is blue and the sun light turns red near the horizon. Other colors can be observed, depending on the various densities of particles in the air (e.g. the sky can sometimes be green after volcanic eruptions, because of the large quantities of very thin particles of smoke that get thrown into the air by the eruption).

Luckily, in VUE, it doesn't take a volcanic eruption to make the sky turn green.

Unlike fog, the effects of haze saturate with distance.

Haze controls are pretty straightforward: if the selected model for the atmosphere is the volumetric one, the controls for haze work exactly like those for fog; if it is the standard one, the fall off setting isn't available, because haze density is considered constant with altitude.

Volumetric Atmosphere Model Only

Glow intensity: this setting controls the intensity of the bright area around the sun, which is caused by the light being reflected on the larger particles in the atmosphere (droplets of water, dust...). The higher the setting, the brighter the atmosphere near the sun.

Volumetric sunlight: check this option to make the sun volumetric. Objects in the scene will start to cast shadows in the atmosphere. The result can be particularly impressive when the sun is low on the horizon. [Volumetric sunlight](#) should be used with great care, because they can dramatically increase render times, without necessarily having any noticeable effect.

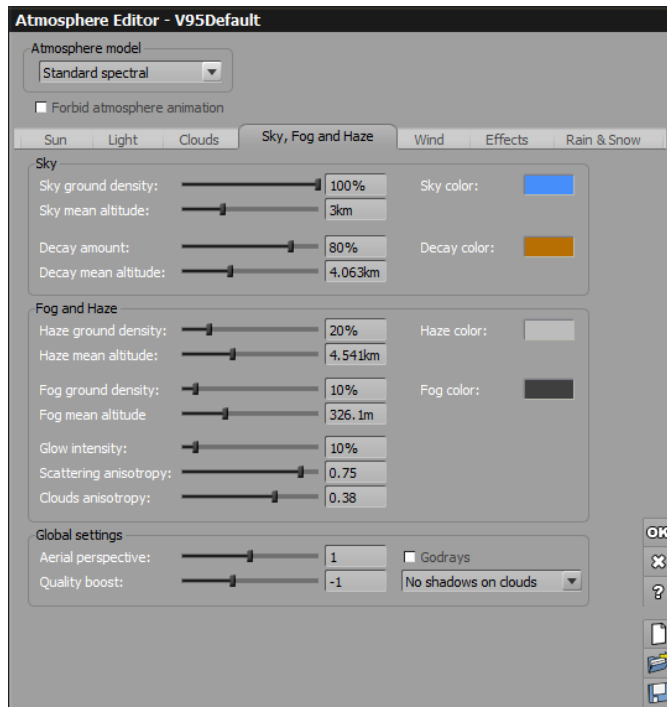
Quality boost: This setting is available only in the volumetric and spectral atmosphere models. It controls the number of samples that are taken throughout the atmosphere in order to compute the interactions of light with the air. Increase the Quality boost setting if you can see noise in the atmosphere (beware: longer render times will result). [Advanced effects](#) can be tuned as well.



Sky, Fog & Haze Tab

This tab is specific to the Spectral and Photometric atmosphere models.

Sky



Atmosphere Editor – Sky, Fog and Haze tab

The controls in this group let you adjust the density of the gasses that constitute the atmosphere (namely nitrogen and oxygen). On earth, these gasses are responsible for the blue color of the sky, and the reddening of the sun near the horizon.

The first set of controls in this group are related to the sky, and the way the blue color appears:

- **Sky ground density:** this setting controls the density of the atmospheric gasses at ground altitude.
- **Sky mean altitude:** this indicates the rate at which the density of the atmospheric



gasses drops with altitude. The lower the mean altitude, the more quickly the density drops (the atmosphere density is exponential with altitude).

Note:

Mean altitude: this is the altitude where density reaches half the density at ground level. Because density is exponential with altitude, this is usually a lot less than half the maximum altitude of the component. For instance, we know that the Earth's atmosphere reaches up to approximately 60 miles high, but its mean altitude is only 5.5 miles (at an altitude of 5.5 miles, the density of the atmosphere is half of that at ground level).

- **Sky color:** this lets you change the color shift caused by the gasses in the atmosphere. On earth, this color is blue, but you can imagine alien planets where the gasses in the atmosphere result in a different color for the atmosphere.

The other settings in this group are relative to the color decay caused by atmospheric gasses. This controls the way the color of light turns red as the sun gets closer to the horizon:

- **Decay amount:** this is the amount of reddening that occurs as the sun gets closer to the horizon.
- **Decay mean altitude:** like for sky, this controls the rate at which the decay disappears with increasing altitudes.
- **Decay color:** on earth, the atmospheric gasses result in a blue color in the sky, and a reddening of light near the horizon. However, the gasses in the atmospheres of other planets could result in a different base sky color and light decay. This setting lets you change the color hue taken by the sun as it gets lower on the horizon. For earthen atmospheres, you shouldn't need to modify the sky and decay colors.

Fog and Haze

The settings in this group control the other components of the atmosphere: small particles, such as dust, and humidity. Small particles are responsible for the haze while humidity is responsible for fog. The settings for Fog and Haze work like Sky and Decay above:

- **Haze ground density:** this indicates the density of particles of dust and pollution at ground altitude. Haze is typically responsible for the gray color that appears near the horizon when the sun is high up in the sky.
- **Haze mean altitude:** controls the rate at which the density of small particles in the atmosphere drops with altitude.
- **Haze color:** controls the color that is added to the atmosphere as a result of the small particles. Usually, this color is gray.
- **Fog ground density:** this indicates the density of water particles at ground altitude. These water particles create a strong glow effect when illuminated from behind. When there is a lot of humidity in the atmosphere, the atmosphere be-



comes gradually opaque.

- **Fog mean altitude:** controls the rate at which the density of water particles in the atmosphere drops with altitude.
- **Fog color:** controls the color that is added to the atmosphere as a result of the water particles. Usually, this color is a dark shade of gray.
- **Glow intensity:** glow is caused by water particles being illuminated from behind. They result in a bright glow around the sun. This setting lets you control the amount of glow in the atmosphere around the sun.
- **Scattering anisotropy:** this controls how “directional” the glow effect is. It influences the overall shape of the glow effect around the sun, and how bright the fog is depending on the direction you look at.
- **Clouds anisotropy:** this provides additional control over the way light is scattered inside clouds. This can make dramatic changes in sunset clouds nearest the sun, for example.

Global Settings

- **Aerial perspective:** this setting controls the overall “thickness” of the atmosphere. A value of 1 corresponds to the typical Earth atmosphere. If you increase this value, the effect, in terms of atmosphere, will be like increasing the scale of your scene.

Note:

In the preset atmospheres, this value is usually set to 10, so that the effects of the atmosphere can be seen without having to use “real-world” size environments. If you are looking for physical accuracy, you should reset this value to 1, which is the aerial perspective of the Earth’s atmosphere. You should also construct your environments at Earth scale (hundreds of miles).

- **Quality boost:** This setting is available only in the volumetric and spectral atmosphere models. It controls the number of samples that are taken throughout the atmosphere in order to compute the interactions of light with the air. Increase the Quality boost setting if you can see noise in the atmosphere (beware: longer render times will result). Please read the discussion on tuning [advanced effects](#).
- **Godrays:** when this option is checked, the clouds will cast shadows in the atmosphere, resulting in rays of light showing through the clouds. The result can be particularly impressive when the sun is low on the horizon. Rendering Godrays can dramatically increase the render times, without necessarily having any noticeable effect (just like in the real world, it takes very specific conditions to see Godrays shining through the clouds). Read [here](#) for more details on Godrays. This option is only available in the Spectral atmosphere model. You can turn Godrays on and off for each individual cloud layer (see [here](#)). You can also adjust the intensity of the

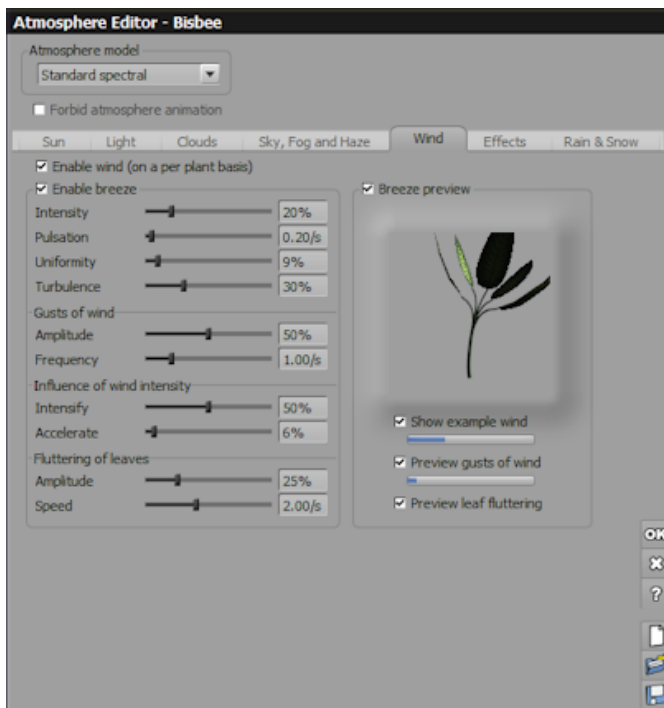


Godray effect using the [shadow density](#) setting in the **Clouds** tab.

The drop-list below Godrays has to do with shadows on clouds:

- **No shadows on clouds:** this option will not produce shadows on clouds from other clouds or any object above the cloud layer.
- **Projected shadows on clouds:** this selection will show shadows on clouds from other clouds or even from a plane flying above the cloud layer.
- **Volumetric sunlight:** this option casts full volumetric and projected shadows onto clouds and through the atmosphere. MetaClouds can cast Godrays only in the “Volumetric sunlight” mode.

Wind Tab



Atmosphere Editor – Wind tab

This tab lets you control the nature and the amount of breeze that is applied to the plants



in your scene.

Provided that breeze is enabled, all plants created in VUE will automatically move gently in the breeze.

On top of this global breeze that affects all plants, you can also define strong winds that will only be applied to given plants in the scene.

- **Enable wind:** if you uncheck this option, no wind will be applied to plants, even if some wind has been defined. Only breeze effects will be applied (if some breeze is defined).
- **Enable breeze:** uncheck this option if you don't want the plants in your scene to move in the breeze. Breeze is a global setting; unlike wind, you cannot disable breeze for given plants only.

Note:

do not confuse breeze and wind. Breeze is defined globally and applies to all plants. It is suitable for gentle, automatic movements of plants. Wind is defined on a per-plant basis, and is better suited for strong amplitude movement of the plants. Plants affected by wind are slower to render than plants that are only subjected to breeze.

Adjusting Breeze

In this section we take a closer look at the different settings that let you control the breeze effect.

Breeze Settings

- **Intensity:** this setting controls the overall intensity of the breeze. Low values mean very gentle breeze, while higher values will produce stronger movement of the plants.

Note:

when you vary the intensity of the breeze, you should also modify the other settings in order to capture realistic breeze movements.

- **Pulsation:** this setting controls the average speed of the plant movements created by the breeze. Use low values when recreating gentle breeze, but increase it if you are creating a stronger effect.
- **Uniformity:** the effect of the breeze is global throughout the scene; however, when you look at real plants moving in gentle breeze, you will notice that each plant seems to move independently. But you can also see an overall movement sliding across the landscape as stronger gusts of wind blow by. This effect is simulated by VUE breeze, and is controlled by the uniformity setting. Low values mean that the plants move independently, whereas high values mean that the plants move all together.



- **Turbulence:** the turbulence setting controls the amount of random movement of each leaf on the plant (as caused by turbulence in the air). Low values mean that all the leaves move together, and high values mean that all leaves move independently.

Gusts of Wind

Gusts of wind appear randomly on top of the overall breeze. They create sudden movements of greater amplitude. The controls in this group let you customize the influence of the gusts of wind:

- **Amplitude:** this setting controls the overall amplitude of the movement caused by the gusts of wind. Low values mean that the gusts of wind create very little noticeable effect, whereas high values mean that gusts of wind will cause strong movements within plants. Gusts of wind appear with random amplitude.
- **Frequency:** this setting controls the average rate at which the gusts of wind occur. Because gusts of wind appear randomly, this setting only indicates the average lapse of time between two gusts of wind. Also, because the amplitude of the gusts is random, not all gusts of wind will necessarily cause noticeable results.

Influence of Wind Intensity

If you look at the way a plant moves in the wind, you will notice that the amount of random movement increases with the intensity of the wind. This effect is captured by VUE's breeze model, and the settings in this group let you control the way the intensity of the wind influences the breeze:

- **Intensify:** this setting controls the overall relationship that exists between the intensity of the wind and that of the breeze. Low values mean that the intensity of the breeze increases only slightly as the intensity of the wind increases. This is appropriate if you want to simulate the deformation of a tree under wind without causing random “noise” in that deformation. High values mean that the intensity of the wind will be strongly influenced by that of the breeze. Strong winds will cause strong random movements of the plant around the wind direction.
- **Accelerate:** this setting controls the influence of the wind on the overall frequency of the movements caused by the breeze. If the value is low, the frequency of the random movements will be the same, whatever the intensity of the wind. If the value is high, strong winds will cause faster random movement of the plant.

Fluttering of Leaves

If you observe the way leaves move in gentle breeze, you will notice that randomly, the leaves suddenly start a burst of rapid movement. This effect is also captured by VUE and it is what we call leaf fluttering. The settings in this last group allow you to control the fluttering of the leaves:

- **Speed:** the speed setting simply controls the speed (frequency) at which the flut-



tering of the leaves happens.

- **Amplitude:** this is also a straightforward control that adjusts the amplitude of the fluttering of the leaves. Low values mean no fluttering, whereas high values mean sudden bursts of strong fluttering.

Breeze Preview

To the right of the breeze settings, you will notice a moving tropic. This tropic is used to preview the effects of the breeze.


Check the **Breeze preview** box to see an animated preview of the effects of the breeze on a typical tropic.


Underneath the tropic, you can see a set of 3 checkboxes and associated gauges. These checkboxes let you select which component of the breeze you want to preview:

- **Show example wind:** if you check this option, a wind of increasing intensity will be applied to the tropic. The wind is applied as a cycle where there is initially no wind, then the intensity of the wind increases gradually to a peak value and subsequently drops back down to zero and begins a new cycle. The current intensity of the wind can be monitored using the gauge. This option is interesting to observe the effects of the intensity of the wind on the behavior of the breeze. In particular, it is appropriate to understand the settings in the *Influence of Wind Intensity* group (see above).
- **Preview gusts of wind:** select this option if you want to view the effects of the sudden gusts of wind. The current intensity of the gusts of wind can be monitored using the gauge. Unchecking this option is like setting a 0 amplitude for the gusts of wind. Toggling this option on and off is a good way to appreciate the effects of the settings in the *Gusts of wind* group (see above).
- **Preview leaf fluttering:** select this option to view the effects of the sudden bursts of quick leaf movement. Toggling this option on and off is a good way to appreciate the effects of the controls in the *Fluttering of leaves* group (see above).



Effects Tab





Atmosphere Editor – Effects tab Volumetric (top) and Environment (bottom) models

This tab in the *Atmosphere Editor* is identical in all atmosphere models. It lets you add cool atmospheric effects such as stars, rainbows or ice rings to your renders.

In the case of the Environment mapped model, the **Effects** tab lets you define the picture to be used in the background.

Stars

Select this option to automatically add stars to your skies. When you turn on stars, the following controls become active:

Number of stars: increase the value to add more stars in the sky.

Brightness: increase the value to make the stars brighter. If the sky is dark, you will probably want to increase the brightness of the stars. If it is blue, you may want to reduce it, because stars are barely visible in the daytime.

Twinkle: this control is used to adjust the amount of twinkling of the stars in an animation. A value of zero means that the stars don't twinkle at all. A value of 100% means that the stars may be completely "turned off" during the animation process...



With lens flares: this option adds tiny cross-like lens flares to all the brighter stars.

Colored stars: checking this option makes the stars appear with random colors.

Rainbow

Selecting this option will automatically add a rainbow effect to your scene.

However, you must understand that rainbows only appear when the sun is shining from behind the camera. If this is not the case, the rainbow will appear outside the field of view. So if you don't see the rainbow, make sure the sun is placed behind the camera, close to the horizon. This is because rainbows are created by the light from the sun being diffracted inside drops of water (rain) and reflected back towards the source.

When you turn on the rainbow feature, the following controls become active:

Intensity: this setting controls the overall intensity of the rainbow effect. The lower the setting, the less noticeable the rainbow will be.

Size: this setting controls the thickness of the rainbow (i.e. the amount of angular spread between the two extreme colors, red and blue).

Falloff: this setting controls the way the intensity of the rainbow reduces with altitude. If the value is high, the rainbow will vanish near the top. The higher the value, the shorter the rainbow.

Secondary bow: turn this option on to display a secondary, inverted bow, larger and dimmer than the main bow. Notice how the sky becomes darker in between the two bows.

Realistic colors: select this option if you want the rainbow to display a realistic distribution of colors, rather than the regular, comic-style red-green-blue rainbow.

Ice Rings

As opposed to rainbows, ice rings are only visible in the atmosphere when looking straight at the sun. Ice rings are caused by tiny crystals of ice in suspension in the air. These ice crystals concentrate light into a ring around the sun, at a specific angle around the direction of the light. This angle (22°), and thus the size of the ring, is directly linked to the angle between the sides of the ice crystals.

When you turn on the ice ring option (this option is only available in the standard and volumetric atmosphere models), the following controls become active:

Intensity: this setting controls the overall intensity of the ring effect. The lower the setting, the less noticeable the ice ring will be.

Size: this setting controls the thickness of the ring (i.e. the amount of angular spread of



light). Low settings will make for less noticeable rings.

Parhelic arc: check this option to show a parhelic arc around the ice ring. This is a secondary, much dimmer ring that appears at an angle of 46° around the direction of the sun.

Sundogs: check this option to show the sundogs on either side of the sun. The sundogs are a horizontal flare of light that appear on either side of the sun, in between the sun and the ice ring.

Pillar: check this option to show the sun pillar. The sun pillar is a vertical flare of light that appears to extend the central sun flare to the edges of the ice ring.

Default Lens Flares

When you create a new light, it is assigned a default lens flare that depends on the type of light. Please turn [here](#) for an explanation of what Lens Flares are.

If the light is a directional light (e.g. the sun), it is assigned the default lens flare for directional lights.

If the light is a point or spot light, it is assigned the default lens flare for other types of light.

This avoids having to define lens flares for each light, and also offers the incredible opportunity to modify all lens flares at the same time. It is also close to reality, because since lens flares occur in the camera, there is no real reason why different lights should be causing different lens flares...

To remove lens flares, uncheck the appropriate options:

- If you don't want the directional lights to have a lens flare effect, simply uncheck the **...directional lights** option!
- If you don't want the other types of lights to have a lens flare effect, simply uncheck the **...all other lights** option!

This will remove the lens flares of all lights that don't have a custom lens flare.

You can also modify the lens flares of all the lights that don't have a custom lens flare by pressing the corresponding **Edit** button. This will open the [Lens Flare Editor](#). When you are done, close the *Lens Flare Editor*, and all the lights in the scene will use the modified lens flare!


If a light has a custom lens flare, it won't be affected by these settings. [Lens flares](#) can be customized for individual lights.

The settings for the default lens flares are saved together with the atmosphere.



Environment Map

This group of controls is only available when the **Environment mapping** model is selected. The main control in this group lets you define the image to be used as an environment map. The picture you select will appear in the back of your render, in place of the sky.





Click the **Load** icon () below the picture preview, or double-click on the picture preview to open the Picture Browser and load a picture. If the picture you load does not map seamlessly (meaning that a seam appears on the edge of the picture when it is mapped onto the background), VUE will detect this and offer to create a seamless joint.


When you load a picture as environment map, a message appears, asking you if you want to setup your scene for Image Based Lighting (read [here](#) for details on this type of rendering). If you click **Yes**, global illumination will be enabled and the lighting information in the picture will be used to illuminate the scene.

Note:

You don't have to use a HDRI image for Image Based Lighting. However, HDRI images produce the nicest results because they contain actual sources of light. If you use a standard picture, you will probably have to increase the sky dome lighting gain to compensate for the fact that there is no light in the map.

If you click **No** to the aforementioned message, the picture you loaded will simply be used as a background to your scene.

You can rotate the picture by using the  and  arrows. You can also invert the picture using the  button. To remove the picture, click the **Remove** icon () below the picture preview.

You can opt to animate the environmental map. The **Animated texture options** icon () is located directly under the picture.

Mapping mode


The **Mapping mode** parameter can have these values :

- **Spherical:** This mode can be also called *equirectangular* and matches the default previous behavior : the whole dome is mapped on the whole image rectangle with the Azimuth as U coordinate and the Pitch as V coordinate.
- **Hemispherical:** This mode matches the behavior of VUE before version 2016R5 with the check “Map upper atmosphere only” checked : like Spherical mode, the mapping is done via Azimuth and Pitch coordinates but only upper half dome is mapped.




- **LightProbe:** This mode allows the use LightProbe maps.
- **Cylindrical:** This mode allows to use the Panoramic Renders renders as cylindrical.
- **Cube Map and Compact Cube Map:** This mode allows to use Cube Maps; either with standard cube pattern or with compact one.

Here is the convention to be used for Cube Mapping :



Here is the convention to be used for Compact Cube Mapping :



Dome rotation


The **Dome rotation** controls let you fine tune the placement on the environment map on sky dome. The **Azimuth** parameter (same usage as previous Map Offset U coordinate) will rotate the picture around the vertical axis. The **Pitch** parameter will rotate the picture around the horizontal axis. The **Horizontal shift** (same usage as Map Offset V coordinate) will adjust the height of the horizon in the sky dome.





The **Exposure** and **Contrast** sliders let you adjust the exposure and contrast of the environment map. If the current environment map is a high dynamic range image, you can view the entire image's dynamic by sliding the exposure setting up and down.

- **Map ground plane:** when this option is selected, the lower half of the environment map is automatically mapped onto the ground plane. This will produce particularly nice results when the horizon in the environment map is exactly halfway up the picture.
- **Ignore atmosphere on map:** if you check this option, the effects of the atmosphere (i.e. fog and haze) won't be visible on the environment map. This is very useful when you need to match the atmosphere of the VUE scene with the atmosphere that is visible in the picture background. For instance, if the background picture show a rainy day, you will probably need to add fog to the scene – or else the objects in your scene will look fake and out of place.

Separate Illumination Map

This option is only available in the Environment mapping atmosphere model. In this model, the Environment map will be used by default to illuminate the scene. Using the **Separate illumination map**, you can however, specify a different image to use as the illumination source. This is particularly useful if you have a low resolution HDR image of an environment coupled with a high resolution non-HDR image.

Click the **Load** icon () below the picture preview, or double-click on the picture preview to open the Picture Browser and load a picture. If the picture you load does not map seamlessly (meaning that a seam appears on the edge of the picture when it is mapped onto the background), VUE will detect this and offer to create a seamless joint.

You can rotate the picture by using the  and  arrows. You can also invert the picture using the  button. To remove the picture, click the **Remove** icon () below the picture preview.

You can adjust the **Azimuth**, **Pitch** and **Horizon offset** parameters using this dialog (see [this](#) for more information).








Default Reflection Map

This is the reflection map that will be used for all materials in the scene that do not explicitly define another reflection map. Please turn [here](#) for further details on reflection maps.

If you are using the Environment mapping atmosphere model, by default the default reflection map will be the same as the Environment map. However, you can specify a different image to be used as reflection map by selecting the **Separate reflection map** option.

This setting is very useful since the environment is theoretically the same for all objects, hence all objects should use the same reflection map (this is of course not an obligation...).

To change the default reflection map, click the **Load**  icon underneath the reflection map preview or double click the reflection map preview to open the Bitmap Browser. Select the picture you want to use as reflection map and validate. A message should appear if your picture doesn't loop smoothly horizontally, and offer to create a smoothed junction between both edges. This is because the reflection map is mapped onto an imaginary sphere, thus looping horizontally. If you click **Yes**, then VUE will add a smooth transition strip from the right to the left border of the bitmap in order to avoid a sharp transitions in the reflection map. Of course, if you don't want to alter the bitmap, click **No**. Your bitmap should now be displayed in the reflection map preview.





You can rotate the picture by using the  and  arrows. You can also invert the picture using the  button. To remove the picture, click the **Remove** icon () below the picture preview.

You can also set the default [reflection map](#) using the *Material Editor* and pressing the **Set default** button.

You can adjust the **Azimuth**, **Pitch** and **Horizon offset** parameters using this dialog (see [this](#) for more information).

Use Environment Map Beyond Atmosphere

When this option is checked, you can select an environment map to use as an outer space atmosphere map. Just load the map using the **Load** icon in the lower left below the image area.

You can rotate the picture by using the () and () arrows. You can also invert the picture using the () button. To remove the picture, click the **Remove** icon () below the picture preview.

The **Dome rotation** controls let you fine tune the placement on the environment map on

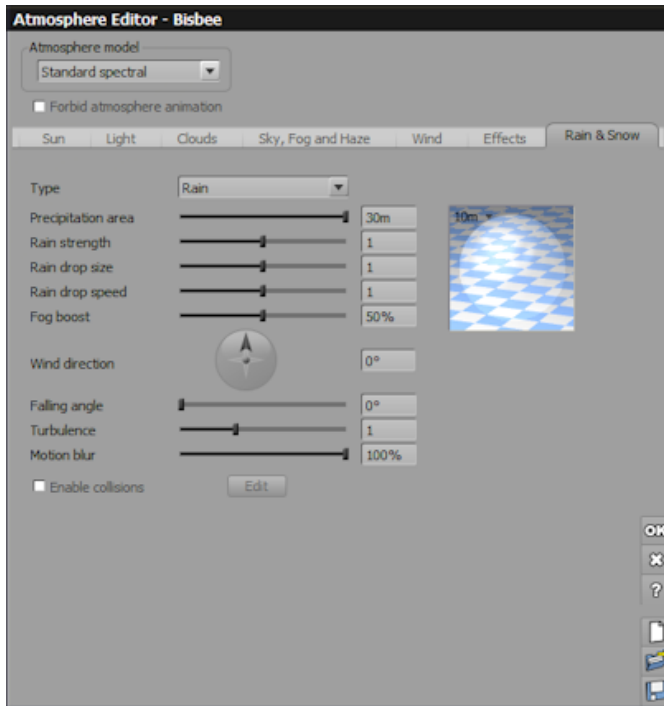


the environment hemisphere. You can adjust the **Azimuth**, **Pitch** and **Horizon offset** parameters using this dialog (see [this](#) for more information).

The **Exposure** and **Contrast** sliders let you adjust the exposure and contrast of the environment map. If the current environment map is a high dynamic range image, you can view the entire image's dynamic by sliding the exposure setting up and down.

You can adjust the **Mapping mode** parameter, see the [Mapping Mode](#) section for more information)..

Rain & Snow Tab



Atmosphere Editor – Rain & Snow tab

The **Rain & Snow** tab contains the settings for rain and snow weather systems. This tab is available for all atmosphere models.



This feature requires the *EcoSystem* module along with the *KronosFX* module.

Type: select either Rain or Snow from the dropbox.

Precipitation area: defines the size of the area affected by precipitation

Rain/Snow strength: slider allows definition of rain strength

Rain drop/Snow flake size: the size of the rain or snow EcoParticle

Rain drop/Snow flake speed: the speed that the snow or rain is falling

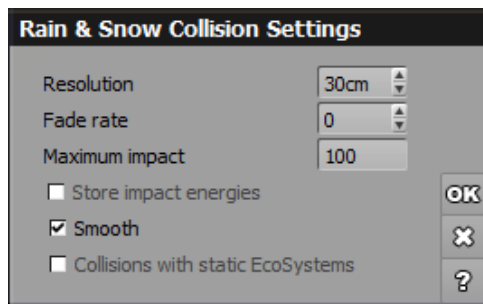
Fog boost: this boosts the fog levels in the precipitation area.

Wind direction: move the pointer to indicate the direction that the rain is falling

Falling angle: indicates the angle that the rain or snow is falling.

Turbulence: the amount of agitation of the snow (falling gently or whipping around).

Motion blur: this is the amount of motion blur applied to the rain or snow.



Collisions dialog

Enable collisions: check to allow collision between the snow or rain EcoParticles. Clicking on **Edit** displays an additional screen to set some additional parameters:

- **Resolution:** The size of the cells in the grid that makes up the precipitation area.
- **Fade rate:** the amount of gradual loss of intensity
- **Maximum impact:** the maximum force of an EcoParticle collision
- **Store impact energies:** when selected, the collision grid will store energy instead of collision count. For energy, the more the particle is fast and heavy, the more energy will be stored at collision.
- **Smooth:** the smoothness of EcoParticles impact



- **Collisions with static EcoSystems:** allows collisions with EcoSystem instances in a scene. This only applies to static EcoSystems. Collisions are not available for dynamic and EcoParticle systems.

Material preview: double-click the preview image to open the *Advanced Material Editor* to edit the weather EcoParticle material.

Saving an Atmosphere

This command lets you save your current atmosphere in a stand-alone file, for use in future scenes. Saved atmospheres will appear in the atmospheres Visual Browser like any other of the predefined atmospheres.

When you select this command, a Standard File Browser appears, letting you choose the name of the file under which the atmosphere will be saved. You may add a title and a comment (recommended).

By default, atmospheres are placed in the *Atmospheres* subfolder. This means that they will appear in the *Personal* collection inside the atmospheres Visual Browser.

Before writing the file to disk, VUE renders a small preview of the atmosphere. This preview will be used in the atmospheres Visual Browser to illustrate your atmosphere. You cannot act upon this preview (framing...).

Photometric atmospheres in xStream

xStream...

Because photometric atmospheres use a brighter sun (with an intensity matching the real sun), it is necessary to use exposure controls in the host applications to avoid overexposed renders. They allow the control of 3 parameters that define the exposure : shutter speed, film speed (iso) and f-number.

The default exposure value (EV) for a sunny exterior scene, recommended for the default VUE atmosphere, is **15 EV** which is obtained with:

- a shutter speed of 1/500s
- a film speed of 100 ISO



- an f-number of 8

You adjust the exposure by changing one of these parameters. For instance, you can add 1 EV by dividing the shutter speed by 2. This makes the render darker. Or subtract 1 EV by multiplying the shutter-speed by 2. This makes the render brighter.

Note:

This exposure value must not be confused with the exposure compensation of recent cameras, that works in the opposite way than exposure scale, and is relative to the automatic exposure value computed by the camera.

Here is a short description of how to activate exposure controls in the various host applications:

...xStream

3dsMax

xStream

V-Ray

With the V-Ray renderer, you have to use an Exposure Control that lets you change the parameters defined above. Select **Rendering > Exposure Control...** in the menu of 3ds Max, then select **Physical Camera Exposure Control** (with 3ds Max 2016 and newer versions) or **V-Ray Exposure Control** (with 3ds Max 2015 and older versions).

xStream...

Exposure Controls can be used in combination with a **Physical Camera**:

- To create a physical camera with 3ds Max 2016 and newer versions, use the menu **Create > Cameras > Physical Camera**. The exposure parameters are available in the right panel, in the “Modify” tab, in the “Physical Camera” and “Exposure” rollouts.
- With versions older than 3ds Max 2016, use the menu **Create > Cameras > V-Ray > V-Ray Physical Camera**. The exposure parameters are available in the right panel, in the “Modify” tab, in the “Basic parameters” rollout.

If you use a physical camera, make sure that the exposure applied to the scene is defined by the camera parameters, by checking **Use Per Camera Exposure** (or by selecting the “From VRay Camera” mode, if it’s a V-Ray Exposure Control).

Note:



There is a display bug in 3ds Max 2016 (independent from VUE): the exposure controls don't always affect the real-time display correctly. For some reason, the value "Exposure for Non-Physical Camera" is the one that affects the display even if the camera is a physical one. In 3ds Max 2015 and 2017 however, this bug is not present.

...xStream

Note:

There is another display bug in 3ds Max 2016 and 2017 (independent from VUE): with the Realistic display mode, the V-Ray Exposure Control doesn't work well with standard directional lights. The viewports appear fully white when using these features. The bug happens with VUE which uses a directional light for its sun. See the last note to avoid the bug. In 3ds Max 2015 however, this bug is not present.

Note:

With the Physical Camera Exposure Control, you should let the **Physical Scale** option enabled (it is by default), with a value of 1500 candelas / RGB unit. This option is in **Rendering > Exposure Control...**

Note:

If you encounter exposure problems in the real-time viewports (where the scene may appear too luminous), you can change their lighting settings. For this, right-click on the quality menu in the top-left corner of a viewport, select the submenu "Lighting and Shadows", then **check the option "Illuminate with Default Lights"**.

Mental Ray

With the Mental Ray renderer, select **Rendering > Exposure Control...** in the menu, then select **mr Photographic Exposure Control**. Once selected, the exposure controls can be modified just below.

xStream

Maya

With the V-Ray renderer, create a standard Maya camera, select it and select the associated shape tab in the *Attribute Editor*. In the menu, select **Attributes > V-Ray** and check **Physical camera**. In the **Extra V-Ray Attributes** rollout, check the option **Treat as V-Ray Physical camera** which enables the exposure settings.

xStream...

With the Mental Ray renderer, create a standard Maya camera, select it and select the associated shape tab in the *Attribute Editor*. Open the "Mental Ray" rollout, then click



on the icon to the right of the **Lens Shader** field. On the window that opens, search for **mia_exposure_photographic** and add this node. Once the lens shader is added, its parameters can be directly edited in the *Attribute Editor*. It is recommended to use a **Cm 2 Factor** of 1.

...xStream

Cinema 4D

xStream

First, select the physical renderer in the render settings. Create a camera, and select the **Physical** tab in the *Attributes Manager*, where you can then edit the exposure parameters. The option “Exposure” has to be checked otherwise these parameters won’t have any effect.

Lightwave

xStream

There is no real exposure control like in the other applications supported by VUE. You have to use an Exposer filter.

To use an Exposer filter, in the *Lightwave* menu select **Windows > Image Processing...** which opens a window. At the bottom of this window, click on the dropdown list **Add Image Filter** and select the Exposer. Once it is added, double click on its line to edit it, which opens a second window. Disable the **Auto-Iris** and **Auto-Exposure** options; these don’t give good results with VUE photometric atmospheres. It allows you to set your own values of **White point** and **Black point**. For instance, if you loaded the VUE photometric atmosphere named **Afternoon**, setting the white point value to 90 and the black point value to 3000% gives a correct exposure. If you loaded the **Sunset** photometric preset, setting the white point value to 10 and the black point value to 750% is more appropriate.

Note:

the effect of the Exposer is applied at the end of the render only, so the render shows an overexposed image while it is in progress.

An alternative solution to control the same settings is to use the **Image Viewer** in the menu **File > Image Controls**. In the window that is opened, check the option **Exposure** and adjust the **White Point** and **Black Point** values.



SoftImage

xStream

In the SoftImage Render menu, select **Render > Pass Options...** which opens a window. Select the **Pass Shaders** tab, and at the bottom of the window, in the **Lens** frame, select **Add After Camera Lens Shaders**. Then click on the **Add** button and select **mia_exposure_photographic**, which lets you edit the exposure settings. It is recommended to use a **Cm 2 Factor** of 1.


Natural Film Response

xStream

Alternatively, you can use the [Natural film response](#) option. This will emulate a exposure filter for all the renders done by VUE. However, the 3d host objects will not be treated. This option is better to do preview renderings before setting a proper exposure filter in your host application.

Automatic Sun Softness





Automatic Sun Softness

This option computes the softness of the shadows produced by the sun from other atmosphere settings, in order to reproduce realistic behavior in terms of soft shadows. More specifically, the softness depends on



- **the height of the sun:** (its pitch) This is on the Sun Tab.
- **on the fog ground density:** This is on the Sky, Fog & Haze Tab.
- **the haze ground density:** This is also on the Sky, Fog & Haze Tab.

The higher the sun, the softer the shadows are. The softness is also proportional to the fog ground density and the haze ground density.

How to use it

Select the sun of your scene, and look for the checkbox **Automatic sun softness** in the **Object – Aspect dialog**. If you check it, you won't be able to select manually the sun's softness: it will be determined automatically depending on the atmosphere's parameters (explained above).

By using this option, you can easily animate the lighting changes of a whole day.

Example:

- create two keyframes
- set the sun position to “Real world” in the *Atmosphere Editor*’ on the Sun Tab.
- set the time of the first keyframe be in the morning (e.g. 6 AM)
- set the time of the second keyframe be in the evening (e.g. 6 PM)
- animate between these two keyframes: the shadows will change automatically according to the time of day.



Light Editor

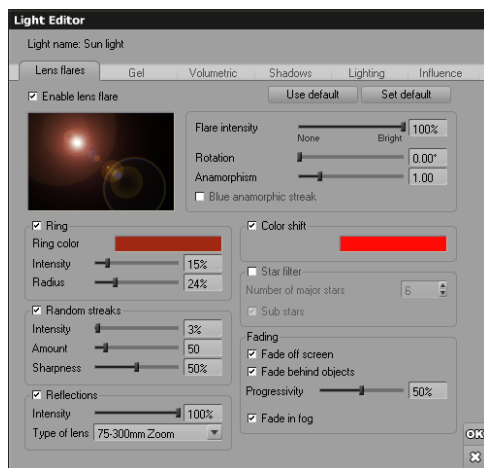
The *Light Editor* groups all the controls relating to the advanced lighting options. It is made of 6 tabs that each lets you control a specific aspect of lights:

- **Lens flares:** this controls the lens flares that appear on the lights
- **Gel:** this is used to have lights project varying colors or images
- **Volumetric:** this controls the visible rays of light caused by the light source
- **Shadows:** this controls the density of shadows as well as shadow mapping options
- **Lighting:** this controls how the intensity of the light evolves with distance from the source
- **Influence:** this controls how objects are influenced by the light source


The *Light Editor* is a non-modal dialog, meaning that it can remain open while you work on other aspects of the scene. If you select another object, the editor for that type of object will replace the *Light Editor*.

If several lights are selected when you open the dialog, the modifications will apply to all these lights. Settings that are not the same for all the lights will be displayed empty. The name of the light that is being edited is reminded at the top-left of the dialog.

Lens Flare Tab



Light Editor – Lens Flare tab

The **Lens Flare** tab of the *Light Editor* is directly accessed by either clicking on the Lens Flare icon () in the *Light Properties* panel (see [here](#)), or by clicking **Edit** in the **Effects tab** of the *Atmosphere Editor*.

The **Lens Flare** tab lets you customize the looks of the lens flares created by a light. When you have finished modifying the lens flare, press **OK** to close the editor. If several lights are selected when you open the lens flare editor, the modifications will apply to all these lights. Settings that are not the same for all the lights will be displayed empty. The name of the light that is being edited is reminded at the top-left of the editor.

The controls of this tab are:

Enable lens flare: select this option to show a lens flare for this light. If you deselect this option, all the controls in the editor will be disabled.

Use default: click this button to restore the lens flare to the default. This is the same as selecting **Default Lens Flare** from the Lens Flare icon menu. All modifications of the lens flare will be lost.

Set default: click this button to transfer the current settings to the default lens flare settings for this type of light. All lights that use the default lens flare will now reflect the modifications. This is the same as editing the lens flare through the **Effects** tab of the *Atmosphere* editor.

Flare intensity: this setting controls the overall intensity of the lens flare effect. Intensities above 100% are possible, when extremely bright lens flares are required. The size of the lens flare depends on the overall intensity, on the distance to the camera and on the brightness of the light source.

Rotation: use this setting to rotate the lens flare. This rotation is only visible on the stars and streaks (see below).

Anamorphism: this setting causes the lens flare to be stretched horizontally, as seen in motion pictures filmed in Panavision. The higher the value, the more stretched the lens flare will be. Values below 1 will make the lens flare narrower.

Blue anamorphic streak: this option creates a bright blue elliptical flare of light that always remains horizontal, and is caused by the use of an anamorphic lens.

Ring

Select this option to create a colored ring around the center of the lens flare. The ring settings become active when you select this option.



Ring color: double-click on this control to open the color selection dialog and modify the overall color of the ring. The default is red.

Intensity: this setting controls the overall brightness of the ring. Typical values are quite low.

Radius: this setting controls the distance between the center of the flare and the ring.

Color Shift

Select this option to gradually shift the color of the central flare towards the indicated color as light moves away from the center of the flare. The shift typical color is red. Double-click on the color to modify it.

If no color shift occurs, the central flare will be entirely white.

Random Streaks

This option adds dozens of small random streaks of light that emanate from the center of the light source. The random streaks settings become active when you select this option.

Intensity: this setting controls the average brightness of the streaks. Typical values are very low, because streaks usually appear as slight variations inside the central flare. However, you can achieve very nice effects by pushing up this value.

Amount: this setting controls the typical number of random streaks emanating from the light source.

Sharpness: this setting controls the average sharpness of each streak. Low values will create large, wide and soft streaks, whereas high values will create very sharp and narrow streaks.

If you click on one of the sliders without modifying the values, a new set of random streaks will be generated.

Star Filter

This option simulates the effects of a star filter placed in front of the lens. What it does is add a given number of regularly spaced, symmetrical streaks around the light source. The star filter settings become active when you select this option.

Number of major stars: this setting indicates the number of stars that appear around the center of the flare. Because the stars are necessarily symmetrical, only even numbers are allowed. You can create up to 10 stars around each light source.

Sub stars: check this option to add shorter and dimmer intermediate stars in between the major stars.



Reflections

Reflections appear as light is reflected on the surface of the different lenses inside the camera. They create rounded shapes that appear all over the picture, and are all lined up on a line that joins the center of the picture to the center of the light source in the picture. When the light moves in the picture, the lens flares move around accordingly.

The reflections settings become active when you select this option.

Intensity: this setting controls the overall brightness of all the reflections in the lens flare. The realistic value is 100% where the reflections have their nominal brightness, but you can reduce this value to create more subdued reflections. You can also increase the value above 100%, which will result in extremely bright, and generally detracting reflections.

Type of lens: the shape and distribution of the reflections in the lens flare are caused by the way the lenses are arranged inside the camera's lens. This drop-down list offers a selection of typical lenses that will each create a different type of reflection. Lens Flare Reflections files created using Vue Infinite's Lens Flare Reflections editor can be used in other versions of Vue. If you select the **Custom...** option at the bottom of the list, the *Lens Flare Reflections Editor* will appear, letting you customize the reflections of the lens flare.

You can save the [lens flare reflection](#) settings as a *.lfr* file, in which case the new reflections will appear in the drop-down list.

Fading

The options in this group control the general behavior of the lens flare inside the scene.

Fade off screen: turn this option on to make the lens flare gradually disappear as it moves out of the field of view. The lens flare will remain visible even when the light source does not appear in the picture. This is because, although the source isn't visible, some rays of light are entering the lens from the sides and still creating the flare.

Fade behind objects: select this option to make the lens flare disappear as the light source passes behind obscuring objects.

Progressivity: this setting applies to both of the previous fading modes. It is only available if at least one of these is turned on. If the progressivity setting is low, the flare will vanish suddenly as it exits the field of view, or as it passes behind objects. On the contrary, when the setting is high, it will disappear very gradually. As a result, part of the flaring effect will remain visible although the light has completely disappeared behind the masking object.

Fade in fog: when this option is selected, the brightness of the lens flare is affected by the fog in the scene.



Lens Flare Reflections Editor



Lens Flare Reflections Editor

The *Lens Flare Reflections Editor* is accessed by selecting the ‘Custom...’ option from the **Reflections** drop-down list in the *Lens Flare Editor* (see above). This dialog lets you create custom lens reflection effects, as well as save them for future use.

A global preview of the reflections is displayed at the top-left of the editor. This preview is automatically refreshed each time you change a setting. Below this preview is the list of reflections. This will be detailed below.

Polygonal Reflections

The shape of the reflections depends on whether the reflection is created on a lens placed before the camera’s diaphragm (such reflections will be circular), or if it is created on a lens placed after the camera’s diaphragm (such reflections will be polygonal, where the number of vertices in the polygon depends on the number of blades in the diaphragm).

Diaphragm blades: this setting lets you control the shape of the polygonal reflections on lenses after the diaphragm in the lens flare reflections. This number indicates the number of blades that constitute the diaphragm. Due to the way diaphragms are built, the reflections resulting from light are polygons, and the number of sides of those polygons is equal to the number of blades in the diaphragm.



Adding Interpolated Reflections

This frame provides you with the option of creating a large number of reflections quickly by interpolating two existing reflections. The settings in this frame only become active when 2 reflections are selected from the list of reflections. Enter the number of interpolated reflections that you want to create in the **Reflections to add** field, and press **Add reflections** to add the interpolated reflections. This will place in between the two selected reflections the requested number of additional reflections, with all new reflections settings being interpolated between the selection. If the two selected reflections are of a different type, the new reflections will be of the same type as the first selected reflection.

List of Reflections

This list displays all the reflections, together with their settings. If there are more reflections than will fit in the list, a scrollbar appears to the right of the list, letting you scroll to other parts of the reflection list.

The following settings are defined for each reflection:

- **Preview:** this displays a preview of the reflection. In order for the preview to be easily identifiable, the Intensity and the Size settings (see below) are not taken into account when building this preview. Obviously, depending on the Intensity setting, the reflection may appear much dimmer in the global lens flare reflections preview. You can make the previews show the actual lens flare reflection as it will appear in the global lens flare reflections by deselecting the **Saturated preview** option below the list

Note:

The size setting is never taken into account when generating the preview.

- **Position:** all reflections are positioned along the axis that joins the center of the picture to the light source in the image. This setting defines the position of the reflection on that axis. Positive values will mean that the reflection appears “on the same side” as the light source relative to the center of the image, while negative values mean that the reflection appears “on the other side”. A value of zero always places the reflection at the exact center of the image.
- **Type of reflection:** this drop-down list box lets you define the type of the reflection. These types are based on typical reflections observed in the real world, and are caused by the different types of lenses in the camera, or the position of these lenses. Available types of reflections are:



Circular, bright center: this creates a round spot with a bright center and dim edges.





- **Circular, bright center with ring:** this is the same as the previous type of reflection, with an added ring that appears around the center of the bright spot.



- **Circular, uniform:** this creates a uniformly colored disk.



- **Circular, uniform with soft edge:** this is the same as the previous reflection, only the edge of the colored disk fades gradually.



- **Circular, linear brightness:** this also creates a disk, but the brightness of the disk is dependent on the position along the axis. Brightness is maximum at the center of the disk, and drops as the point on the disk moves away from the center along the axis.



- **Circular, dim center:** this creates a “hollow disk” type of effect where the reflection is dull at the center of the disk, and bright on the edges.



- **Circular, dim center with ring:** this is the same as the previous reflection, with an added ring around the edge of the reflection.



- **Polygonal, bright center:** this is the equivalent of the ‘Circular, bright center’ reflection, only the shape of the reflection is polygonal instead of being circular, because the reflection is created by light that has already gone through the diaphragm.



- **Polygonal, bright center with ring:** this is the same as the previous type of reflection, with an added ring that appears around the center of the bright spot.





Polygonal, uniform: this creates a uniformly colored polygon.



Polygonal, uniform with soft edge: this is the same as the previous reflection, only the edge of the colored polygon fades gradually.



Polygonal, linear brightness: this also creates a polygon, but the brightness of the polygon is dependent on the position along the axis. Brightness is maximum at the center of the polygon, and drops as the point on the polygon moves away from the center along the axis.



Polygonal, dim center: this creates a “hollow polygon” type of effect where the reflection is dull at the center of the polygon, and bright on the edges.



Polygonal, dim center with ring: this is the same as the previous reflection, with an added ring around the edge of the reflection.



Rainbow ring: this creates a circular reflection, where the colors of the reflection go through the entire spectrum, ranging from blue on the inside of the ring through yellow in the middle and red on the outer edge. Usually, rings are created with a white color, but you could use other colors to tint the rainbow ring.

- **Color:** this defines the color of the reflection. In order to clearly see the preview of the reflection, as well as the reflection’s tint, you should use bright colors and adjust the actual brightness of the reflection using the intensity setting below.
- **Size:** this parameter defines the size of the reflection. This parameter is not taken into account when generating the reflection preview.
- **Intensity:** this parameter defines the intensity of the reflection. This parameter is not taken into account when generating the reflection preview, except if the **Saturated preview** option is deselected.

Note:



Keep in mind that it is the association of several basic reflection types that creates realistic reflections. In order to achieve realistic reflections, you should create many reflections of low intensity, rather than a few very bright reflections.

Underneath the list of reflections is the **Saturated preview** checkbox that lets you decide whether the reflection previews take into account the Intensity setting or not.

Click the **Sort** button to reorganize the list of reflections by position.

Click **Add** to add a new reflection at the end of the list, or after the selected reflection if a reflection is selected.

Click **Remove** to delete the currently selected reflection(s).


New, Load, Save

Pressing **New** will clear the list of reflections.


Press **Load** to load a Lens Flare Reflections file that defines a list of reflections.

Press **Save** to save the current Lens Flare Reflections in a stand-alone file, for use in future scenes. By default, Lens Flare Reflection files are placed in the *Environment/Lens Flare Reflections* folder, and have the extension *.lfr*. Saved Lens Flare Reflections will appear in the **Type of lens** drop-down list box of the *Lens Flare Editor* (see [here](#)), and will appear under the name of the file they were saved in.

Gel Tab



Light Editor – Gel tab

The **Gel** tab of the *Light Editor* is directly accessed by clicking on the **Gel** icon () in the *Object Properties* panel (see [here](#)).




The **Gel** tab lets you modify the gel material that is applied to the light.

You can edit the gel using the *Gel Editor* (double-click on the gel preview). In effect, the *Gel Editor* is nothing more than the *Simple Material Editor*, with only the **Color** tab available. When you have finished modifying the gel, press **OK** to close the editor.

Note:

you cannot create gels that are based on Mixed or Volumetric materials.

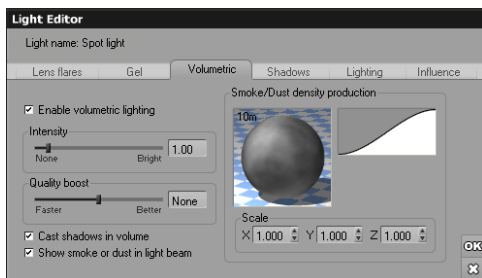
Uncheck the **Enable gel** option to remove the gel from the light.

Click the **Load** icon () to load a new gel material. You can select how the gel is projected. These options are the same as in the [Light Gel](#) options menu.


If several lights are selected when you open the editor, the new gel will apply to all these lights.

Please turn [here](#) for a detailed explanation of the *Material Editor*.

Volumetric Light Tab



Light Editor – Volumetric tab

The **Volumetric** tab is accessed directly by selecting **Edit Volumetric Options** from the Volumetric Light icon's menu () in the *Object Properties* panel (see [here](#)).

The **Volumetric** tab lets you customize the looks of the volumetric effect of lights. When you have finished modifying the volumetric options, press **OK** to close the dialog.

You can remove the volumetric effects from the light by unchecking the **Enable volu-**



metric lighting option.

If several lights are selected when you open the dialog, the modifications will apply to all these lights. Settings that are not the same for all the lights will be displayed empty. The name of the light that is being edited is reminded at the top-left of the dialog.

When turning on the **Volumetric light** property of one or several lights, use of shadow maps is automatically activated (this dramatically increases rendering speed). Although there is usually no reason for this, if you would rather use ray-traced volumetric shadows, remember to uncheck the [Use Shadow Map](#) option in the **Shadow** and **Lighting** tab.

Volumetric Light Controls

The controls of this tab are:

- **Intensity:** this setting controls the overall brightness of the beams of light cast by the light source.
- **Quality boost:** this is the quality boost setting for the volumetric light effect. The higher the value, the better the quality of the volumetric effect, but the longer it will take to compute. The [Advance Effects Quality](#) section has more information for using this setting appropriately.
- **Cast shadows in volume:** select this option if you want objects to cast shadows in the volumetric light. Turning it off will make the volumetric light render much more rapidly. So, when shadows in the volume are not required, it is a good practice to uncheck this option. The Cast shadows in volume option is not available when the light doesn't cast [shadows](#).

Smoke/Dust in Volumetric Light Beams

Select the **Show smoke or dust in light beam** option to add variations to the density of the volumetric beams of light. When you select this option, the Smoke/Dust density production controls appear.

This is how the **Smoke/Dust density production** works: for each point of the volume, the function returns a number in the range of 0 to 1 (0 appears black on the preview and 1 white). The number is then transformed by the filter into a smoke density at this point. The brightness of the light at this point is directly proportional to the smoke density (0 if there is no smoke, hence no light, 1 if there is a lot of smoke and the light is bright). When rendering the volumetric beams, VUE accumulates the density of smoke all along the ray that is traversing the beam, and then computes the resulting brightness of the volumetric effect.

To modify the function, double-click on the picture of the function. This will open the



Visual Function Browser.

Use the scaling controls to scale the function along the **X**, **Y** and **Z** axes. Ditto for the filter. Use a filter that presents a strong saturation to increase the contrast between the dark and bright areas of the beam.

Shadow Tab



Light Editor – Shadows tab

This tab lets you customize the looks of the shadows cast by lights, as well as the way these shadows are generated.

The **Shadow** tab is accessed directly by selecting **Edit Shadow** from the **Shadow and Lighting Options Light** icon menu (📷) in the *Object Properties* panel.

Enabling Shadows

If you want to remove all shadows cast by the light, uncheck the **Enable shadows** option.

The **Shadow density** setting controls the overall darkness of the shadows cast by the light. If the value is 0, it means that no shadows are cast by the light. If the value is



100%, it means that objects that are in the shadow of the light won't be receiving any light from it. Intermediate values indicate that objects that are in the shadow of the light do in fact receive some amount of light from it, as if the shadowing objects were partially transparent.

The **Softness quality** slider becomes available when the **Softness** of the light is set to a non zero value. Pressing the **Edit** button displays the **Custom softness options** dialog, letting you customize the softness quality further. You will find three controls:

- **Min:** this setting controls the number of rays initially sent to evaluate the softness of the shadow.
- **Max:** if the render engine decides that more shadow rays are required, it will keep sending new batches of rays until the total number of rays sent for that pixel reaches this setting.
- **Quality threshold:** this setting controls the severity with which the render engine decides whether more rays are required or not, after having computed the first batch. The higher the setting, the more often sub-rays will be sent into pixels.

Using Shadow Maps

To render the shadows cast by this light using an *AccuShadows™* shadow map, check the **Use shadow map** option. If this options isn't selected, shadows will be rendered using a standard ray-tracing algorithm.

Note:

when rendering a **volumetric light**, shadow maps will always be used to render the light volume (to increase rendering speed).

When you select this option, the controls in the shadow map group become available.

Projected Hard Shadows

This parameter is used when the light doesn't create soft shadows (the Softness setting of the light(s) is set to zero, thus producing hard shadows on the envrioning scene). You can specify how these projected hard shadows should be generated:

- **Ray-traced:** a shadow ray will simply be cast towards the light to determine any occluding objects. This produces the best quality results.
- **Shadow mapped:** a shadow map will be used to compute the shadows. Depending on the resolution of the shadow map, shadows will be computed much more rapidly than when using the ray-traced option. Although the shadows are less accurate, you may find that the quality is sufficient under certain circumstances.

Note:



When rendering a volumetric light (see [here](#)), shadow maps will always be used to render the light volume. The ray-traced versus shadow mapped setting above only acts upon the projected shadows – not on the light volume.

Quality

This is the most important group of controls for shadow maps. It is mainly with these parameters that you will be able to find a proper compromise between performance and accuracy of the shadow map.

- **Size of the shadow map:** the size of the shadow map controls the accuracy of the shadow map (i.e. the precision of the shadows generated – you will need higher sizes for sharp edged shadows, whereas a lower size would be enough for blurry shadows). If, for instance, you choose a size of 256, then your shadow map will consist of $256 \times 256 = 65536$ cells that will be used to map your light's field of view. With this parameter, you strongly influence the accuracy/performance compromise. When the **Auto size** option is checked, shadow map size is specified as a **Ratio** of the rendering resolution. For instance, if you set the **Ratio** to 0.5, your shadow map size will be half of the rendering resolution. When **Auto size** option is unchecked, you can directly specify the **Size** of the shadow map. Cell count is limited to 4096×4096 because of the high memory requirements involved. You will notice that default **Ratio** and **Size** can differ from a light type to another. This is because light's field of view can be more or less important, depending on light type. For instance, the field of view of a directional light is much greater than that of a spot light, therefore a bigger ratio is needed for directional lights to obtain an acceptable accuracy.
- **Bias:** is another very important parameter, but you may never need to change it. Because it is an approximation, the shadow map needs an error tolerance for the information contained in each cell. This tolerance is controlled by the **Bias** parameter. The default value of 1 is often acceptable, and changing it could have undesirable effects, as it is very sensitive. If you notice moiré patterns appearing on surfaces (this is caused by a lack of accuracy in some cases where the scene is ill-conditioned for shadow maps), you should try to modify the **Bias** and the Filter Bias to eliminate this undesirable effect. If it persists, try increasing the size of the shadow map instead; this will improve accuracy.
- **Filter Bias:** is another error tolerance controller used during filtering of the shadow map (described below) and is also very sensitive. Normally, you shouldn't have to modify this parameter, unless you encounter the previously mentioned moiré patterns artifacts.

Note:

These bias parameters are intended for advanced users only. They are useful only in very specific cases, and they are difficult to control. Generally speaking, simply increasing the size of the shadow map to improve accuracy is a safer approach for solving moiré pattern



issues.

- **Sampling boost:** this parameter deals with shadow map filtering. Filtering is used to improve the softness of shadow mapped soft shadows. It specifies the maximum number of samples used for filtering the shadow map. This parameter can strongly influence the quality/performance compromise: too few samples will result in noisy shadows, (especially if **Softness** is high) whereas too many samples will slow down rendering dramatically, as filtering is performed for every rendered pixel. If you notice noise in the shadows, you might want to increase the boost setting to improve quality. However, because the actual number of samples taken increases with the render quality setting (see [here](#) for details), this may not be necessary. You should first check the quality of the shadows with the better rendering modes.

Note:

The filtering radius of the shadow map is directly related to the **Softness** parameter: the higher the **Softness**, the greater the filtering radius. Thus, for high **Softness** settings, more filtering samples are necessary to reduce noise.

Softness

The settings in this group are used to control shadow dispersion with distance. When using ray-traced soft shadows, the softness parameter controls this dispersion. With shadow maps, it is only used to determine the filtering radius that controls the amount of softness. Dispersion with distance must therefore be simulated for improved realism of mapped shadows.

VUE offers three different methods for computing shadow dispersion:

- **Constant:** this is the most simple – and fastest – method for computing dispersion. In this mode, the softness doesn't vary with distance, thus dispersion is not taken into account. This option can be used when dispersion is not important. It avoids having to increase the Max samples value in regions that would otherwise display a high shadow dispersion.
- **Light based:** this is the most robust solution to simulate dispersion, although it is not physically accurate. The further away the shadow is from the light source, the more dispersed it becomes. With this method, you may notice that the shadows at the base of shadowing objects are not as hard as they should be... Directional lights are incompatible with the “Light based” mode.
- **Object based:** this method attempts to accurately capture the physical phenomenon involved with shadow dispersion. If you look closely, shadows are always hard near the object that casts the shadow. But as you move away, the shadow becomes softer. Unfortunately, this is a behavior that is extremely difficult to capture with shadow maps, and selecting this option may produce unexpected results (sudden transitions, shadow bleeding). In some cases however, the results can be satisfactory. If the physical accuracy of the soft shadows is essential, you should use ray-traced soft



shadows instead (this method accurately simulates the physical phenomenon behind soft shadows).

Note:

It is usually preferable to use the Light based rather than the Object based dispersion method. Light base is a more robust approach that usually produces acceptable results.

Directional lights are incompatible with the “Object based” softness mode.

- **Dispersion coefficient:** when used with the dispersion method setting, this parameter is used to modulate the dispersion over distance. For instance, if you feel that shadows are spreading too rapidly with distance, you can lower the value to reduce spreading with distance. Although this parameter has no physical justification, it is very useful because it allows you to control the dispersion over distance at whichever scale you work at.

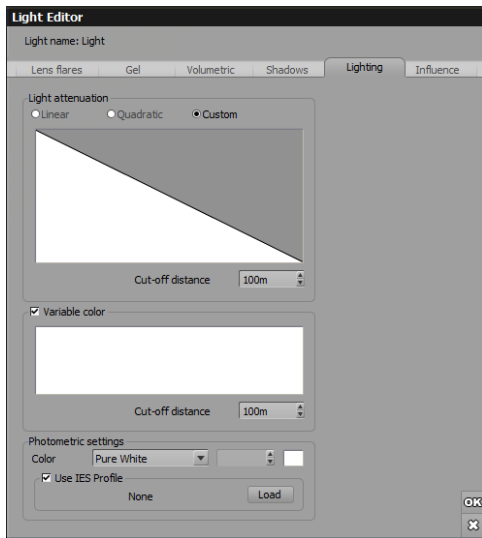
Additional Information

When turning on the **Volumetric light** property of one or several lights, use of shadow maps is automatically activated. Although there is usually no reason for this, if you would rather use ray-traced volumetric shadows, remember to uncheck the **Use shadow map** option.

Also, you may notice that with very simple scenes – e.g. made of only a few basic primitives – the use of shadow maps may actually be less efficient than the use of standard ray-traced shadows. This is because of the computational overhead involved in the creation and management of a shadow map. As mentioned in the introduction to this section, the benefits of using shadow maps increase with the scene’s complexity.

Lighting Tab





Light Editor – Lighting tab

The settings in this tab will let you modify the profile of light attenuation with distance, as well as customize the color of light based on distance.

Light Attenuation

The controls in this group are used to adjust the way the intensity of the light drops as distance to the light source increases.

The first set of controls are used to determine the attenuation profile of the light:

- **Linear:** this is a standard attenuation profile where the intensity of light is proportional to the distance from the light source. Although this is physically incorrect, it is useful for creating lights that “reach further”.
- **Quadratic:** this is the physically correct attenuation profile, whereby light intensity drops with the square of the distance to the light source. This is the default attenuation profile of Quadratic light sources. As mentioned above, quadratic attenuation is rather strong, and linear attenuation might be preferable in some cases.
- **Custom:** the attenuation profile is user defined, and controlled by the attenuation filter and the cut-off distance settings described below. This is useful for fine tuning of light intensity over distance.

When the Custom attenuation mode is selected, the following controls become available:


- **Attenuation filter:** this filter represents the attenuation profile used for the light(s) if the Custom attenuation mode is selected. The standard look of this filter – for a plausible attenuation profile – should be a decreasing curve, but you can achieve



interesting results by specifying other shapes. The profile you indicate will be applied to a linear attenuation of light intensity. This means that if you create a constant filter, the intensity of light will actually drop linearly with distance.

- **Cut-off distance:** this value specifies the range of distances to which the Attenuation filter applies. This means that the leftmost value of the filter is applied at zero distance from the light source, and the rightmost value of the filter is applied at the distance specified by this parameter. At a distance that is superior to this cut-off distance, the attenuation value corresponds to the rightmost filter value. Of course, this value will be used only if the Custom attenuation mode is selected.

Variable Color



Spotlight with variable color

The settings in this group let you control the color of the light emitted by a light source. If the **Variable color** option is selected, the color of the light will vary with the distance to the light source (of course, the intensity of the light varies independently from the color – you can adjust the way the intensity varies using the *Light Attenuation* controls described above). Although there is absolutely no physical justification behind this behavior, it can be used to create interesting effects on occasions (especially when using volumetric lights). If the **Variable color** option is not set, the color of the light will be the same, whatever the distance from the light (obviously, the intensity of the light changes with distance).

When the **Variable color** option is selected, the following controls become active:

- **Color map:** this color map represents the color of the light based on distance to the light source. Edit this color map to create custom lighting schemes.
- **Cut-off distance:** this value specifies over which distance range the color map is applied. This means that the leftmost color value of the color map is applied at zero distance from the light source, and the rightmost color value of the color map is applied at the distance specified by Cut-off distance. At a distance that is superior to the cut-off distance, the color of the light will be the rightmost color in the color map.



Photometric Settings

- **Color:** use the droplist to select the type of lighting you are using. For some types of lighting, the strength can be adjusted. Most types of lighting will have a default color that is indicated by the color square to the right.

The different types of lighting are:

Pure White

Candle flame

Tungsten

Halogen

Carbon arc

Sunlight

Xenon arc

Daylight

Warm fluorescent

Full spectrum fluorescent

Standard fluorescent

Cool white fluorescent

High pressure sodium

Sodium vapor

Mercury vapor

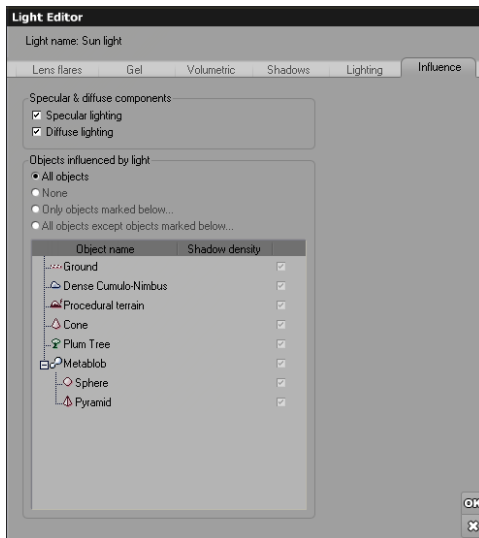
Metal halide

Custom

- **Use IES Profile:** click the checkbox if you want to use an IES profile. Then, click **Load** to select the file you wish to use.



Influence Tab



Light Editor – Influence tab

The settings in this tab let you specify which objects are influenced (or not influenced) by the light(s), as well as how they are influenced.

Specular and Diffuse Components

The options in this group let you decide which components of the light are applied to objects lit by the light:

- **Specular lighting:** if this option is selected, the specular (or highlight) component of the light will be applied to all objects lit by this light.
- **Diffuse lighting:** if this option is selected, the diffuse component of the light will be applied to all objects lit by this light.

Using these settings you can create light that do not exhibit any highlights on shiny materials – ideal when setting-up stage-like lighting.

Objects Influenced by Light

The options in this group let you decide which objects are going to be lit by the light.

There are four influence modes which can be used to specify how your objects should react



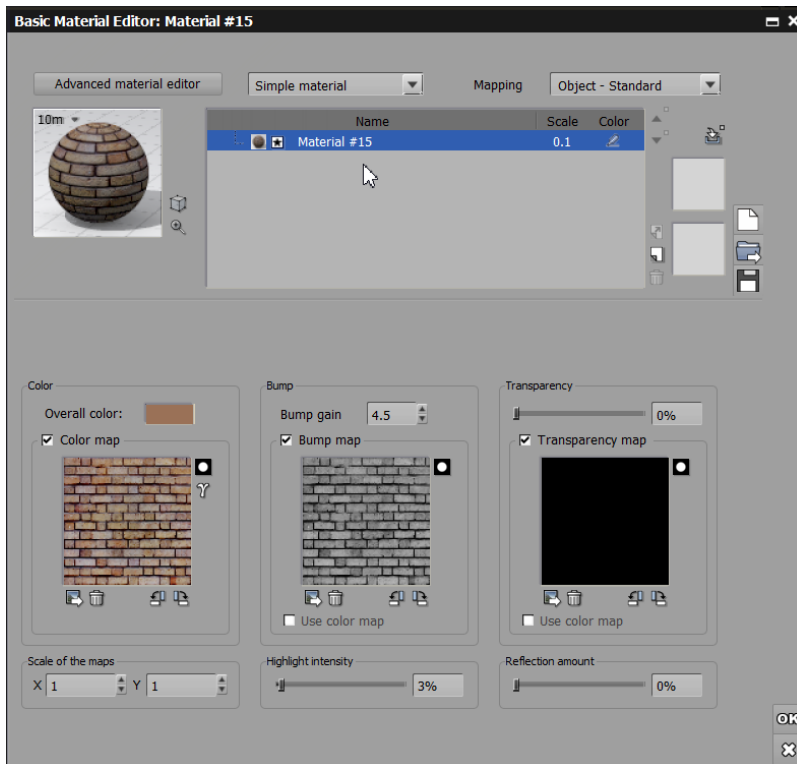
to light:


- **All:** all objects in the scene will be affected by the light. This is the default mode.
- **None:** none of your scene objects will be affected. Only the lens flare – if a lens flare was defined for this light – will reveal the presence of the light when rendering. This can be useful if you want to add a local lens flare without perturbing the lighting of your scene.
- **Only objects marked below:** only the objects selected in the object list below will be affected by the light. In other words, you specify which objects *should* be influenced.
- **All objects except...:** all objects in the scene *except* the ones selected in the object list below will be affected by the light. In other words, you specify which objects should not be influenced by the light.

Use the **Object list** to select the objects that are (or are not) affected by the light. All the objects in the scene appear in this list. Simply navigate through the list and check the ones you want by clicking the check boxes on the right of the object name. This list is only active when the **Only marked** or **All except** modes above are selected.



Material Editor





Material Editor – Basic (top) and Advanced (bottom)

Materials are the secret behind the quality of pictures generated by VUE. And the reason for this is twofold: VUE materials are not just 2D pictures mapped onto objects, they are truly three dimensional (which means when you carve into them, you actually carve into new parts of the material), and they are designed to respond to their environment (altitude, slope, orientation, etc).

Unfortunately, this visual quality has a drawback: creating materials can be a complex process. However, we have striven to keep it as simple and straightforward as possible, while maintaining full access to every aspect of material synthesis.

Each time you make a modification to a material, the material preview is redrawn by a multithreaded background task without slowing down the interface.

The *Material Editor* is accessed by double-clicking on the preview of a material, or by selecting **Edit Material** from the popup menu that appears when you press the right mouse button (Ctrl mouse on Mac) over the preview. It can stay open without restricting access to other parts of the software.

There are two types of *Material Editors* in VUE:

- the *Basic Material Editor*, ideal to setup basic texture mapped materials easily (see



[here](#), and,

- the *Advanced Material Editor* that gives you full access to all material parameters (see [here](#)).

You can switch from one *Material Editor* to the other anytime by clicking the large button at the top-left of the *Material Editor*.

Types of Materials

Materials are divided into 6 types:

- Simple materials,
- Two-sided materials,
- Mixed materials,
- Grouped materials,
- Volumetric materials (this type of material cannot be edited in the *Basic Material Editor*), and
- EcoSystem/Particles (this type of material cannot be edited in the *Basic Material Editor* either).

Mixed materials are built by mixing together 2 other materials, either simple ones or mixed ones themselves. You cannot mix together several volumetric materials, but you can mix EcoSystems.

Materials can also be layered to easily add e.g. stains to an existing material. Material layers work in a similar way to Photoshop layers in that they are added one on top of the other, and layers below are only visible in places where the layers above are nonexistent (see [here](#) for details).

Multi-Materials

Multi-Materials are created by saving all of the materials of a selected object or plant as one material. These materials can then be accessed as a **.mat** file from the Material Browser, the saved image in the browser showing the different materials in the file. These materials are a convenient way to quickly change materials of objects. A specific multi-material should always be used on a specific object, so that the number of materials matches.

These materials are saved or loaded from the menu displayed by right-clicking on the window in the *Object Properties* panel.



Common Material Controls

This section details the controls that are common to all types of materials.

New, Load, Save

In the dialog bar, on the right edge of the editor you will find these usual commands. These are found in the lower right corner:

- **Ok:** applies the changes you have made to the settings.
- **X:** cancels any changes you have made to the settings.

These are found on the right side of the screen in the upper part of the dialog:

- **New:** will reset all material characteristics so that you can start working on your material from a clean base.
- **Load:** lets you open and detail the characteristics of an existing material using the Material Browser.

Note:

If the material you are editing is animated, loading a material here will create a new [material animation](#) keyframe.

- **Save:** lets you save the current material in a stand-alone file, for use in future scenes. Saved materials will appear in the Visual Material Browser like any other of the predefined materials. By default, materials are placed in the **Materials** subfolder. This means that they will appear in the **Personal** collection inside the Visual Material Browser.

Note:

There are so many parameters involved in material creation, that it is usually easier to modify an existing material rather than to start from scratch.

Defining the Material

These icons are only available in the *Advanced Material Editor*:

- 

One sided: If the icon is selected, it indicates that objects using this material should be traced for only one intersection per ray. This option is only available in the *Advanced Material Editor*. Since opaque objects block all rays at their surface anyway, One sided will only affect transparent materials. Rays actually never pass through a one sided object, so this should not be used with materials that have some fading out. Some effects (like Fuzziness) will force One sided to be activated. One



sided objects can be very useful when rendering details on the surface of an object, while not wanting to see details on the opposite surface.



Disable anti-aliasing: If this icon is selected, it lets you selectively disable anti-aliasing on given materials. On a general basis, anti-aliasing increases picture quality. However, some materials may lose their grainy aspect when anti-aliased, and you may want to remove anti-aliasing in such cases.



Hide from camera rays: When this icon is selected, this material displays only through reflections or refractions.



Hide from reflected/refracted rays: When this icon is selected, this material displays only when being viewed directly through the camera.



Disable indirect lighting: Selecting this icon disables indirect lighting on the material being edited.



Disable caustics: Selecting this icon disables any caustics that might be used with this material.



Ignore lighting: Selecting this icon disables any influence that either sunlight or lighting sources such as a spotlight might have on this material.



Ignore atmosphere: Selecting this icon disables any influence that the sun, ambient lighting or any other kind of atmospheric effect would have on this material.



Don't cast shadows: Selecting this option will prevent the object from casting a shadow which can be useful for luminous objects. Deselecting this option when the shadow of an object is not needed (because it isn't visible) can also significantly improve rendering speed.



Don't receive shadows: When this icon is selected, objects made of this material will not be shadowed by other objects in the scene.

Note:



Materials that have no diffuse lighting never receive shadows anyway.

Note:

Since computing shadows is a time consuming process, you might want to turn shadows off where they are not required.



Only shadows: When this icon is selected, the object will not be directly visible in the rendered picture. It will however still cast a shadow on other objects. This is particularly useful when you want to create shadow-casting masks without actually seeing the mask.



Matte / Shadow / Reflection: When this option is selected, it will generate alpha masks that are proportionate to shadowing & reflected geometry at each shaded point. Global illumination shadowing is also taken into account and will also affect alpha masking accordingly.



Show in the timeline: This option adds this material to the animation timeline. When this option is selected, *Animated Material* is appended to the caption of the editor.



Disable material animation: This option blocks animation of this material.



Animate material surface (Z = time): This option turns material surface animation on for that material. Please read [here](#) for details on material surface animation (and other types of material animation). When this option is selected, *Time dependent material* is appended to the caption of the editor.



Influence by Vector Shapes:

Type

This is where you choose the type of material. The available types of material are:

- Simple material,
- Mixed material
- EcoSystem/Particles material.

These options are only available on the Advanced Material Editor


- Volumetric material



- Grouped
- Two-sided

Selecting one of these options will toggle between the different types of materials. Additionally, simple materials and EcoSystem materials can be [layered](#).

Effects



Effects

- **Advanced Render Options:** click to display the Effects dialog with these two fields:

- **TAA boost:** The Texture Anti-Aliasing boost offers the ability to adjust texture anti-aliasing quality on a specific material. Use the slider to raise or lower the amount of anti-aliasing compared with the global setting.

Note:

Anti-aliasing must be enabled in the Render Options for this feature to have any effect.

- **Subray quality drop:** This feature allows you to decrease secondary rays quality (i.e. reflected or refracted rays) of the render for this particular material.

Note:

Anti-aliasing must be enabled in the Render Options for this feature to have any effect.

Render times can be reduced using this feature for materials where reflections and/or refractions are costly to evaluate while their contribution to the image is relatively low. This can be useful for water surfaces with strong perturbations, as the reflections and refractions will be blurred by these perturbations. Thus, each secondary ray quality can be reduced without altering the final image quality.

- **Mapping:** this drop-down list lets you select the [mapping mode](#) that will be used for the material. Mapped pictures have their own specific mapping mode. They may be overridden inside a particular [function](#).




Material Hierarchy

The material hierarchy is the list that sits in the middle of the *Material Editor*. This list displays all the different sub-materials and layers of mixed or layered materials. Mixed or layered materials can be expanded to display the different layers and sub-materials. If you click on one of these items, the *Material Editor* will change to reflect the settings of that sub-material or layer. Using the material hierarchy, you can easily access all the different components of the material.

Click on the **Name** field to change the name of the material.

The material hierarchy operates exactly like the hierarchy displayed in the **Materials** tab of the *World Browser*, with the exception that only the hierarchy of the current material is displayed.

Alongside each line of the material hierarchy is the **Highlight** switch () . Click on this to highlight the corresponding sub-material or layer. Highlighted materials will be displayed using a solid color, thus letting you easily check where the material appears in a scene and adjust its contribution. When you are done adjusting the material, simply click this switch again to restore the normal colors of the material. You can adjust the highlight color by right-clicking (Ctrl+Click on Mac) on the color switch when the material is highlighted.

Layering Materials

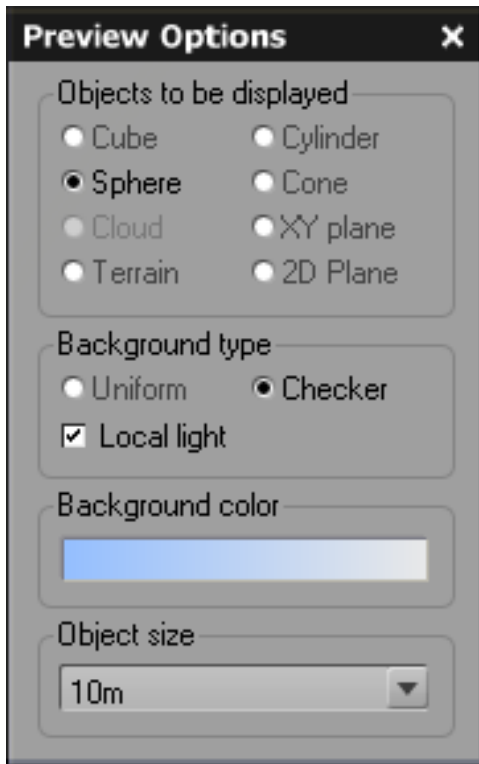
Click the **Add layer** button to add a layer to the current material. The Material Browser will appear, letting you select the material you want to add as another layer to the current material. Closing the Browser will add an empty layer. If the current material is a simple, mixed, or EcoSystem material, it will become a **layered material**, with at least two **layers**.

To remove a layer, select it in the hierarchy and click **Remove Layer**.

You can move layers up and down by clicking the Up and Down buttons alongside the material hierarchy.



Material Previews



Preview Options

In the middle of the Material Editor are square spaces that can contain materials. When you start, only the first square is occupied by a preview of your current material.

Randomize (🎲): clicking this icon makes a random change to all fractal and noise nodes used for a material. You can keep clicking until you find the effect you like. This is only for use with procedural materials.

Preview Options (⚙️): Clicking the **Options** icon displays the *Preview Options* dialog.


This dialog enables you to select which object should be used to preview materials or functions. **Sphere** is the fastest, and **Cloud** should only be used for cloud materials. **XY Plane** displays a 2D representation of the material in perspective, whereas **2D Plane** presents the material on a plane seen from above.

It also lets you choose a background type for the preview (**Uniform** or **Checker**), as well




as the **Background color** by modifying the color map (double-click on the map).


Check **Local light** to use a local light rather than a directional light.

Zoom: Clicking the **Zoom** icon () displays an enlarged view of the material. Click on **Render** to re-render the preview; press **Esc** to stop.

Store

Selecting the **Store** icon () copies the current material into the first available material preview (in the set of previews to the right of the material hierarchy), making it available for future retrieval. If you select **Restore This Version** (from the popup menu), or double-click on one of these stored materials, the corresponding settings are copied to the current material.

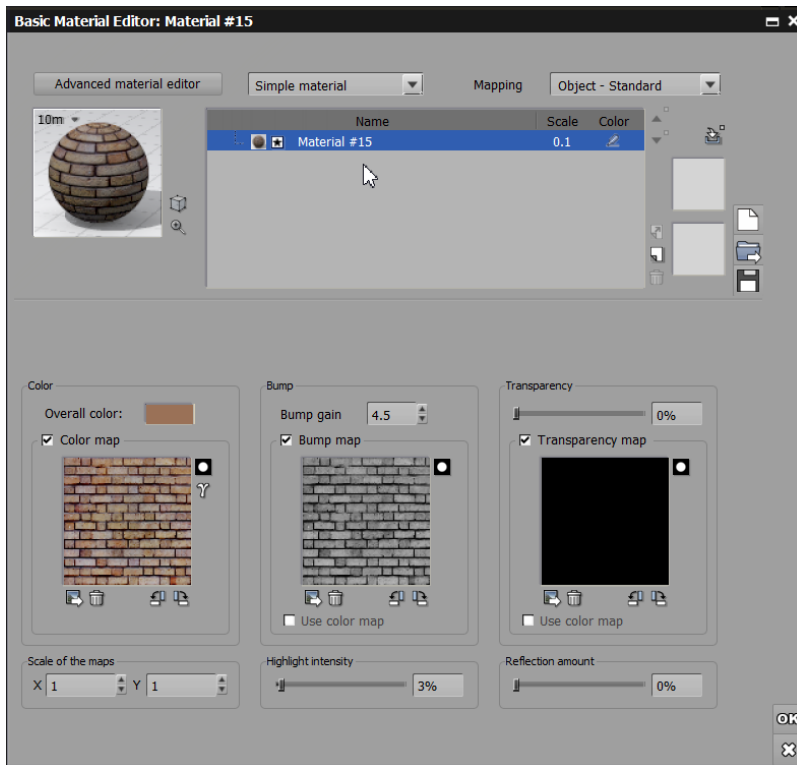
Tabs



The tabs contain further settings for the material. If the material does not contain any of the features on a particular tab, the tab will not be highlighted. You can still use that tab to make changes to the material; then it will become highlighted.

Basic Material Editor





Basic Material Editor

The *Basic Material Editor* is particularly convenient to easily setup texture mapped materials, or do basic modifications of existing materials. This version of the *Material Editor* won't let you access the entire range of effects possible in VUE, but it is a good way to begin with materials before you delve into the intricate complexity of the [Advanced Material Editor](#).

The *Basic Material Editor* is split in two halves. The top half has been detailed above. The contents of the lower half changes depending on the type of material (simple, mixed or layered). The following pages detail the lower part of the editor for each type of material.

If you try to load a Volumetric material or an EcoSystem material into the *Basic Material Editor*, a message will appear informing you that this requires switching to the *Advanced Material Editor*.




Simple Materials in Basic Editor


If the current material is a Simple Material, the *Basic Material Editor* displays the controls described below.

Color Frame





The controls in this frame let you adjust the colors of the material.

Overall color: this color control lets you modify the overall color of the material. Because the control displays an average color, this color may not be actually visible in the material. For instance, if the material exhibits a red and white checkerboard, the overall color will turn out pink – despite the fact that there is only red or white in the material. The overall color can be modified by double-clicking on the color control. All colors in the material will be modified in order to produce an average color that is the same as the one indicated by the overall color control.

Color map: check this option if you want the material to be colored by a picture. Double-click on the picture preview or click the **Load** icon () to load a new picture.

To create an animated texture map (Rotoscoping) , select an animation file from the Bitmaps Browser, or press the **Browse file** button () in the Bitmaps Browser to display a Standard File Browser and select multiple picture files.

The **Animated texture options** icon () will appear under the picture preview. Click this icon to display the [Animated Texture Options](#) dialog.

If you need to rotate the picture, use the  and  buttons (90° increments). To invert the colors in the picture, click the **Invert** button () . Click on the **Remove** button () to delete the picture or animation.

Bump Frame


The controls in this frame let you adjust the bumpiness of the material surface.


Bump gain: this control adjusts the amount of bump at the surface of the material. The bigger the value, the bumpier the surface.

Note:





If the material does not define any bumps (either through the use of a bump map, or procedurally), no amount of bump gain will make bumps appear on the surface of the material.



Bump map: check this option if you want the bumps at the surface of the material to be generated according to the grayscale values in a picture. Double-click on the picture preview or click the **Load** button () to load a new picture.

To create an animated texture map (Rotoscoping), select an animation file from the Bitmaps Browser, or press the **Browse file** button () in the Bitmaps Browser to display a Standard File Browser and select multiple picture files.

The **Animated texture options** icon () will appear under the picture preview. Click this icon to display the [Animated Texture Options](#) dialog.


If you need to rotate the picture, use the  and  buttons (90° increments). To invert the colors in the picture, click the **Invert** button (). Click on the **Remove** button () to delete the picture or animation.


Use color map: if this option is selected, the same map will be used for the bump map as the one used for the color map.

Transparency Frame





The controls in this frame let you adjust the transparency of the material.

The top slider controls the global transparency of the material. It is only available when no alpha map is applied to the material.

Transparency map: when this option is selected, the transparency of the material is defined according to the grayscale values in a picture (white areas will be transparent, whereas black areas will be opaque). If you select this option, the value of the global transparency displayed at the top of the frame will be bumped up to 100% (because the alpha map is just a modulation of the global transparency setting). Double-click on the picture preview or click the **Load** button () to load a new picture.

To create an animated transparency map (Rotoscoping), select an animation file from the Bitmaps Browser, or press the **Browse file** button () in the Bitmaps Browser to display a Standard File Browser and select multiple picture files.

The **Animated texture options** icon () will appear under the picture preview. Click this icon to display the [Animated Texture Options](#) dialog.

If you need to rotate the picture, use the  and  buttons (90° increments). To invert the colors in the picture, click the **Invert** button (). Click on the **Remove** button () to delete the picture or animation.



Use color map: if this option is selected, the same map will be used for the alpha map as the one used for the color map.

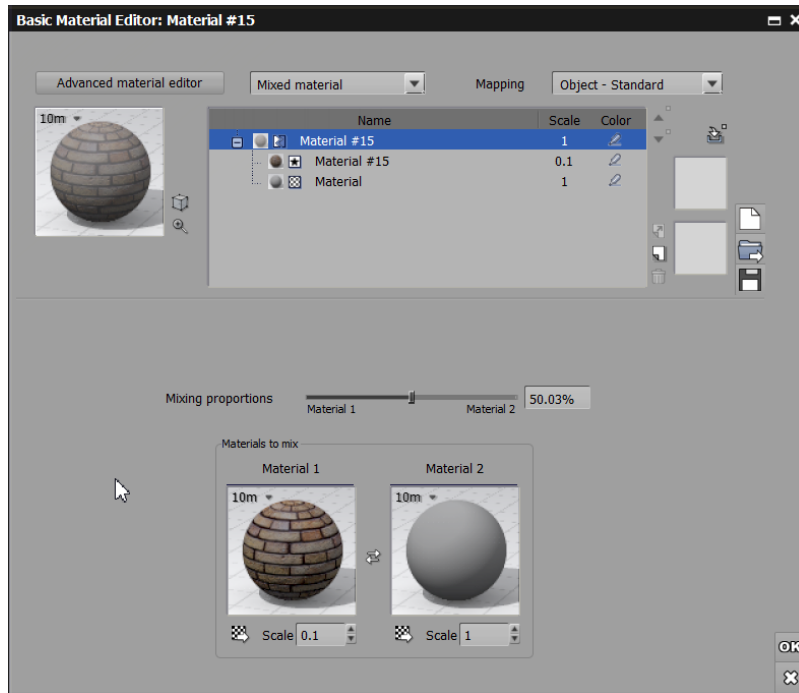
Other Settings

Scale of the maps: this control lets you adjust the scale of all the texture maps along the X and Y axes. If no texture maps are used, this control will be disabled.

Highlight intensity: this setting controls the overall intensity of the highlights that appear on surfaces that point towards the sources of light.

Reflection amount: this setting controls the overall reflectivity of the material.

Mixed Materials in Basic Editor



Basic Material Editor – Mixed Material


If the current material is a Mixed Material, the *Basic Material Editor* displays the controls described below.



You can switch from Simple Material to Mixed Material by selecting the appropriate option in the **Type** frame at the top of the *Material Editor*.

Mixed materials are created by mixing two other materials together. The rules for mixing the materials together can be very complex, but the *Basic Material Editor* only lets you select the two materials that are mixed as well as adjust the mixing proportions.

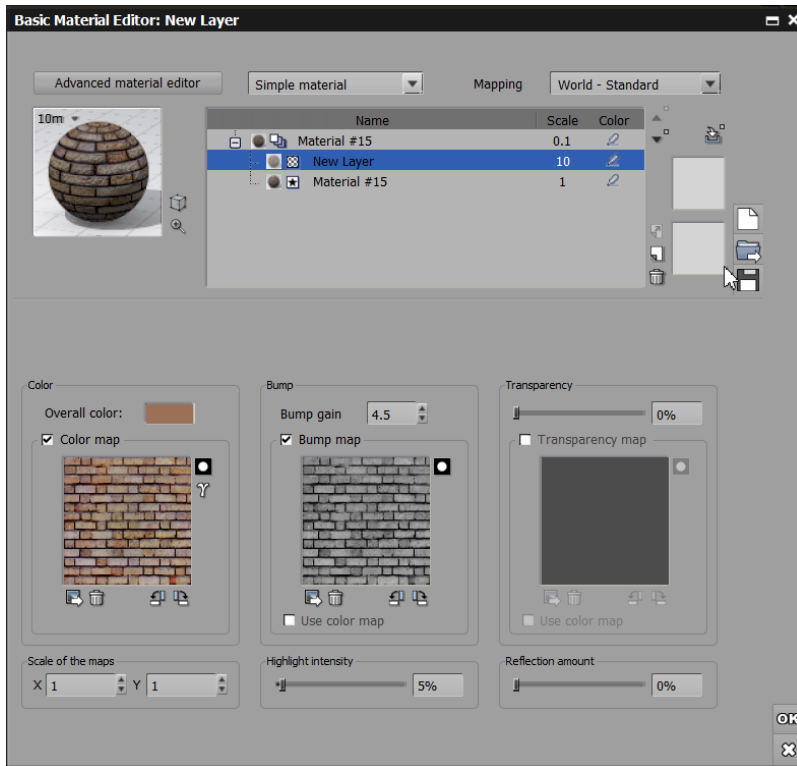
Mixing proportions: this slider lets you define “how much” of each of the two materials that are being mixed will be visible in the resulting material. If you drag the slider to the left, you will see more of Material 1, and if you drag it to the right, you will see more of Material 2. You may notice that all mixed materials are not mixed in the same way (e.g. some materials are influenced by slope or altitude). If you want finer control on the way the two materials are mixed, you will have to use the *Advanced Material Editor*.

Material 1 and 2: these are the two materials that are mixed together. Change the materials by loading materials that already exist on the disk with the  button, or by double-clicking on a material preview to edit it. Materials inside a mixed material can be scaled independently using the **Scale** controls. This only modifies the size of the material once it is applied to an object. A scale equal to 1 does not change the size of the material.

Swap: press this button to swap material 1 and 2.



Layered Material in Basic Editor

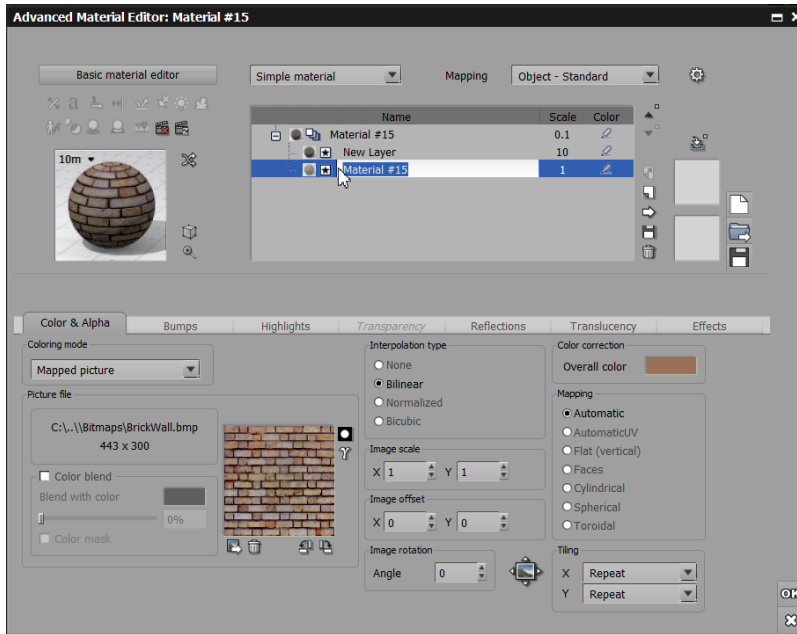


Basic Material Editor – Layered Material

If the current material is a layered material, the *Basic Material Editor* shows a list of all the layers of the current material. This list is very similar to the [Material Hierarchy](#), with the exception that it only displays the layers of the current material.



Advanced Material Editor



Advanced Material Editor


The *Advanced Material Editor* lets you define more precisely the look of your materials. The downside to this is that this version of the editor is significantly more complex to master. In the following pages, we will review and explain all the different options offered by this version of the *Material Editor*.



As with the *Basic Material Editor*, the lower part of the *Advanced Material Editor* changes according to the type of material being edited:

- *Simple materials*
- *Mixed materials*
- *Volumetric materials*
- *EcoSystems*



Driving Material Settings with Functions

In the *Advanced Material Editor*, a number of material settings are preceded by the **Drive with a function** icon () . If you click this button, the setting will be driven by a function: the **Function Graph** will appear with a new output node available. That output node corresponds to the material parameter. A constant node will be connected to it, and the value held by this constant node will be the same as that of the material parameter before it was extracted. At this point, the material is not yet affected by the operation (except under very specific cases where the extraction of the parameter changes the way that parameter is interpreted – such cases will be clearly documented in the corresponding parameter descriptions). However, now that the parameter is extracted, you may drive it with any type of function (e.g. a noise node!).

If you go back to the initial node, you will notice that the **Extract parameter** button () has been replaced by the **Disconnect parameter** button () , and instead of displaying input controls, an indication that the node is “connected” appears. If you click the disconnect parameter button, or if you disconnect that parameter’s output in the *Function Graph*, the parameter will be reintegrated into the node and restored to its initial constant value.

The underlying power of this simple feature is truly amazing! This can be used to create totally unique material shaders. For instance, by connecting the highlight color to a function, you can create unique iridescent effects.

Published Parameters

The **Published Parameters** feature copies specific settings from the *Function Graph* that you may need to change often and places them in a more convenient location for easier material manipulation. In the *Advanced Material Editor*, a new tab is created for these parameters.

To select a parameter for publishing, just click the underlined parameter field name in the *Function Graph*. For example, if you are using a variable noise fractal for a material bump, you might want to publish the **Roughness** parameter. A parameter name is supplied and a group name is asked to improve the display of the published parameter. This parameter will then appear on a **Published** tab in the *Material Editor* so that you can change the settings there.



Simple Materials

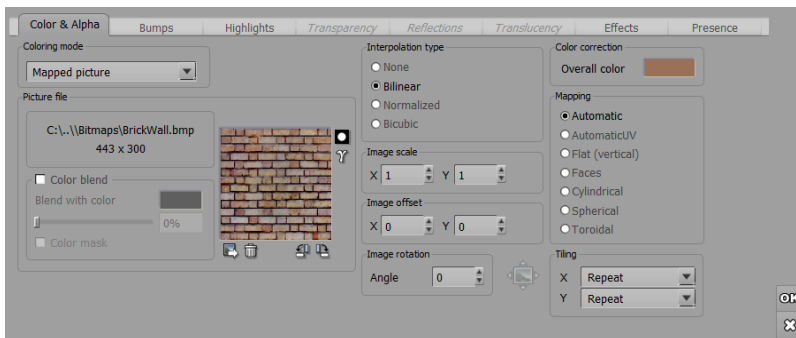
Simple materials are defined by 7 sets of parameters, each corresponding to a tab in the editor:

- **Color & Alpha:** color of the surface.
- **Bumps:** bumps on the surface (bump-mapping algorithm).
- **Highlights:** specular reflections: is the surface shiny or dull?
- **Transparency:** transparency / refraction of the material.
- **Reflections:** reflections on the surface of the material.
- **Translucency:** sub-surface scattering and translucency of the material.
- **Effects:** local surface lighting and special effects.
- **Environment:** used with layered materials.

Material Layer

If the material is a layer of a **multi-layer material**, the **Alpha boost** control will appear on top of the tab control. This setting lets you control the overall “presence” of the layer in the multi-layer material. If you drag the slider towards positive values, the layer will appear stronger (within the limits of the environment constraints that you set using the **Environment** tab).

Color & Alpha Tab



Color & Alpha tab – Simple Material – Mapped picture mode


The **Color & Alpha** tab defines the color of the surface of the material and the corresponding opacity (alpha).


You can choose from 3 types of surface coloration:



- **Mapped picture**,
- **Procedural colors**,
- **Natural grain**
- **Overall color**: this color control is common to the Mapped Picture and Procedural Color options. It lets you modify the overall color of the material. Because the control displays an average color, this color may not be actually visible in the material. For instance, if the material exhibits a red and white checkerboard, the overall color will turn out pink – despite the fact that there is only red or white in the material. The overall color can be modified by double-clicking on the color control. All colors in the material will be modified in order to produce an average color that is the same as the one indicated by the overall color control.





Mapped Picture

You can use any picture to color the surface of a material.

First, you must indicate the picture that you want to use by clicking the **Load** icon () and selecting a file from the Bitmaps Browser. You can change the name of the image (**Picture File**) by clicking on the image name in the *Material Editor*. This name can be changed in the *Function Graph* as well.

To the right of the image, you can select to set the **Gamma** correction () for this material, overriding the global settings.

To create an animated texture map (Rotoscoping), select an animation file from the Bitmaps Browser, or press the **Browse file** button () in the Bitmaps Browser to display a Standard File Browser and select multiple picture files. The **Animated texture options** icon () will appear under the picture preview. Click this icon to display the **Animated Texture Options** dialog.

If you need to rotate the picture, use the  and  buttons (90° increments). To invert the colors in the picture, click the **Invert** button (). Click on the **Remove** button () to delete the picture or animation.

It is possible to load image sequences directly into a single multi-image sample node and it will distribute the loaded images randomly over the texture. To load a sequence one



has to specify the path as a regular expression (for example, `c:\img*.bmp` will load all the *img1.bmp*, *img2.bmp* etc in `c:\`).

The advantage over creating several multi-image sample nodes and connecting them through an image combiner is that it's easier to use when one has many images and that the image overlapping order is not fixed. However, one cannot specify different distribution settings for these images (like density, rotation, scale etc).

To map the picture (2D by definition) onto a 3D volume, VUE must use one of the available mapping modes. Each of these mapping modes is best suited for some types of objects (e.g. spherical for Spheres).

Select one of the following:

- **Automatic:** The mapping technique is chosen automatically, depending on the object onto which the material is applied (e.g. Spherical for a sphere, cylindrical for a cylinder...).
- **Automatic UV:** This mapping technique is used for a 3D displaced textured terrain to allow for the generated mesh of the terrain at render time.
- **Flat:** Vertical projection / slide projector type, oriented so as to project the picture on the ground; values don't depend on altitude.
- **Faces:** Slide projector type of projection oriented along one of the three world axes. For each point, the projection axis is the closest axis to the normal vector of the object.
- **Cylindrical:** Mercator projection: the picture is wrapped around a cylinder around the vertical axis before being projected.
- **Spherical:** The picture is projected so that it covers exactly a sphere. Since the picture wraps around 180° vertically, and 360° horizontally, the scale seems to double vertically.
- **Torical:** The picture is projected so that it covers exactly a torus. A strange, and not very useful mapping mode, hum...

If you don't know which to use, select **Automatic**.

Note:

The shape of the object on which you project the picture does not have to be the same as the type of projection you choose.

You can control the way the picture is repeated along both axes using the **Tiling** drop-down boxes. Available options are:

- **Repeat:** this is the default. The image is repeated indefinitely along this axis.
- **Mirror:** in this mode, the image is also repeated indefinitely, however, it is mirrored



each time so that the repetitions join seamlessly.

- **Once:** the image is displayed only once along this axis.

If you want the picture to tile symmetrically horizontally, select **Mirror X**; If you want the picture to tile symmetrically vertically, select **Mirror Y**.

When the material is seen from very close, you may see pixels, due to the limited resolution of the picture. To reduce this effect, choose an **Interpolation type** method:

- **None:** No over sampling.
- **Bi-linear:** Bi-linear interpolation between pixels.
- **Normalized:** Values proportional to the distance to the corners of the pixel.
- **Bi-cubic:** Bi-cubic interpolation between pixels (continuous derivative).

Indicate the **Scale** of the picture along the X and Y axes.

You can position the picture precisely on the object by using the **Image offset** commands. This will move the picture around by increments of one pixel.


The **Color blend** group lets you blend the colors of the picture with a solid color. To activate this feature, check the corresponding checkbox.

The color is applied in product mode, and the slider lets you adjust the amount of blending. The higher the value, the more the solid color modifies the picture.

Click **Color mask** to apply the color in replacement of the bitmap as the setting increases. When set to 0%, the color is applied as a mask. When set at 100% the color completely replaces the bitmap.




Texture Placement Editor



Texture Placement Editor

This feature brings a new rotation parameter for image mapping and a new editor to act directly on your mapping with gizmos and live feedback. It supports standard terrains and meshes. For **Projected Texture Maps** only, there is a field for **Image Rotation** which allows you to rotate the material on the object.

Click on the **Gizmo** icon  to display the *Texture Placement Editor*. This dialog displays the selected object with a gizmo representing the scale/rotation center. You can show the isolated object by un-toggling the **Show Scene** button.

The gizmo you see is a representation of the scale/rotation center of your image on the edited object. Use **Alt** to move the pivot center.

Note:

Edition mode is only accessible in **Flat** and **Faces** mapping mode.

In **Faces**, the gizmo will position itself on the nearest mapping face regarding your view orientation. In **Flat**, the gizmo will be positioned in the top plane of the object's para-



metrical space.

Be careful to slow object's reconstructions. Changing the mapping necessitates that the UV are re-baked for the OpenGL display at every interaction frame. This can be computation heavy for large objects such as terrains. We recommend lowering the quality settings when editing mapping to improve the interactivity on the **Display** tab of the Options panel (**Instant draw** and **Background draw** faster instead of better).

The Pivot

The pivot is the rotation center for your manipulation. But it is also the position at which the manipulator will be displayed and reset after each movement. You can edit the pivot position using the **Alt** key, or the **Edit pivot** toggle button.

You can position your pivot in 3 dimensions in order to place the manipulator exactly at the location you are looking at. To simplify this task you can use the **Pick pivot position** toggle button. Just click that icon and click inside the editor window to select the position of the pivot on the object you are editing.

How It Works

There are two proxy objects, the Mapping manipulator (a little plane) and the Rotation center. When you manipulate the plane proxy, its movement is read and translated into change in the image mapping parameters: **Scale**, **Offset** and **Rotation**. These parameters are mirrored in the top of the editor. When you manipulate the pivot proxy, it only changes an internal position that is used to recalculate offsets when you are doing rotations with the plane. This is used to create the sensation of free rotation center. In reality, the rotation center is always at the (0.0) in UV space, but you will notice the “offsets” changing after each rotation to simulate the displaced center.

Limitations

There are inherent limitations of the above mentioned points. **This is not an UV editor tool**, it is only a mapping parameters editor. As such, for example, if you are in **Faces** mapping projection, you cannot edit independently the mapping parameters of each 6 sides.

The manipulator proxy object is trying to match at best the texture speed and movement, but in certain circumstances movement of the gizmo can very well be unrelated to the movement of the texture. Notably, when editing the mapping of a mesh with UV, we cannot predict the direction and speed that the texture will take when the mapping is changed.

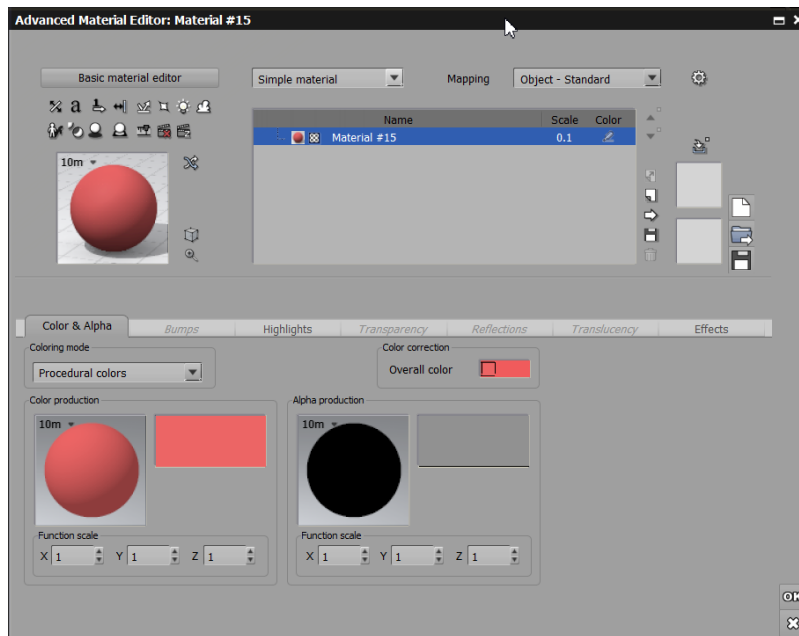
To avoid problems like this, avoid slants in object transformations, and use linear projections rather than UV mapping, like **Faces** and **Flat**.



There are also limitations in the previsualization. The VUE OpenGL is used to preview the modifications in conjunction with the small render scene preview at the right of the application main panel.

The limitations of VUE's standard viewports OpenGL also applies in the editor. It is impossible to preview multiple layer materials at once. On the performance side, it is impossible to change the mapping without reconstructing the entire object. (We plan to develop improvements in future versions). This limitation, however, does not prevent you from editing the mapping of an intermediate layer, in a multi-layer material. You can use the *Function Graph* for that. From the *Material Editor*, select the layer you want to edit, then in the **Color** Tab, right click on your image and choose **Edit Function**. Then, in the *Function Graph*, your **Projected Texture Map** node will be selected. Then access the **Show Manipulation Gizmo**. The editor's OpenGL will show a temporary material featuring only the layer you selected. You can do the same for any **Projected Texture Map** node in your graph.

Procedural Colors



Color & Alpha tab – Simple Material – Procedural colors

VUE can produce the colors of the material algorithmically, using a function, a filter and a color map.



This is how it works: for each point of the surface, the function calculates a value in the range of 0 to 1 (0 appears black on the preview of the function, 1 appears white). This value is then transformed by the filter into another value in the range of 0 to 1. The filter can be added in the *Function Graph*. From this last value, the color map produces the color of the surface (if this value is 0, the color will be the one at the left end of the map, if it is 1, it will be the one at the right end).

To modify the function, double-click on the picture of the function. This will open the Visual Function Browser.

Use the scaling controls to scale the function along the X and Y and Z axes. If necessary, use the filter to modify how function values are transformed into colors (Control-click on the filter).

Finally, indicate which colors are assigned to the values of the function by editing the color map (Control-click on the map).

Note:

If the color map is solid (only one color), whatever the function and filter values might be, the material will always yield a uniform color.

Alpha Values

As well as colors, this tab can be used to control the alpha value of the material. Alpha is the same as non-refractive transparency. It can be used to “cut out” parts of a material, and is especially useful when working with layers. For instance, if you wanted to create a label using a bitmap, you would create a bitmap layer and connect the alpha channel so that the layer is completely transparent outside of the label. Alpha can also be used in conjunction with refractive transparency to “cut out” parts of glass materials.

If the surface of the material is colored by a bitmap, the alpha output is automatically connected to the alpha value of the image (fully opaque by default). Outside the image (if the image is not tiled), the alpha value is automatically set to 0 (fully transparent).


If the surface of the material is colored using procedural colors, the alpha output is connected to the alpha value of the color map.

You can access the *Function Graph* by right-clicking on the sphere under **Alpha production**. Right click on the graph to the right of the alpha image to edit the filter.

Alpha values are especially useful to control the presence of a layer in a multi-layer material (see [here](#)).



Animated Texture Options



Animated Texture Options

The *Animated Texture Options* dialog can be accessed by clicking the **Animated texture options** button in the *Material Editor*. This button appears below the texture map previews, when a sequence of pictures or an animation has been loaded.

This dialog lets you customize the way animated texture maps are displayed.

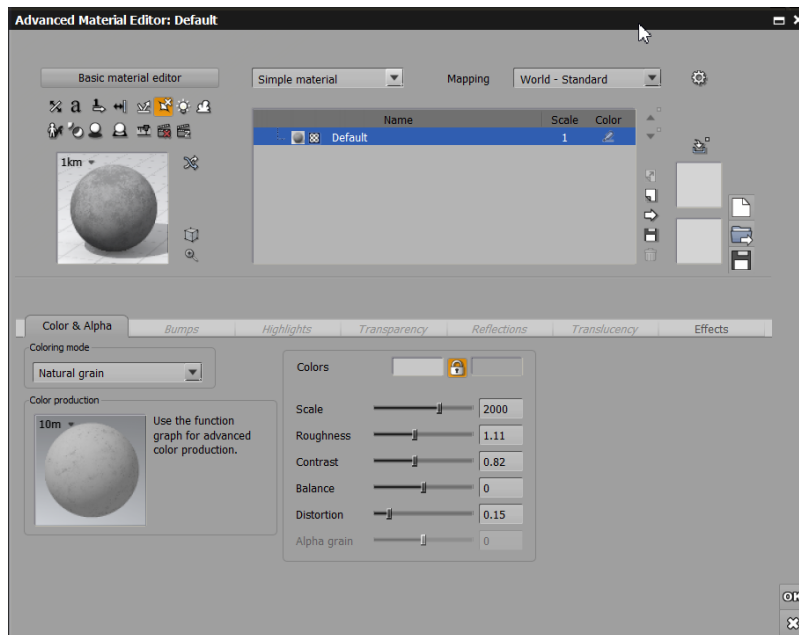
- **Image sequence:** this is the list of pictures to use in the animation. You can add new pictures by clicking the **Load** icon (📁). You can replace pictures in the list by selecting them and then pressing Load. To remove images from the list, select them and then press the **Remove** icon (🗑️).
- **Frame rate:** this defines the playback rate of the pictures on the list. Ideally, this



should at least be equal to the global animation frame rate.

- **Interpolate frames:** when this option is selected, in-between frames are interpolated by gradually blending the previous and the next frames. This ensures smooth playback and will avoid any jumps in the animated texture.
- **Animation filter:** use this filter to change the flow of time in the animated texture. Double-click on the filter to load a filter, or select **Edit** from the filter's popup menu to edit the filter.
- **Phase:** use this to adjust the start frame in the animation sequence. The value has to be set in seconds.
- **Image offset, Interpolation type, Mirror X & Y, Picture scale, Mapping mode:** these settings are identical to the settings in the **Color** tab of the *Advanced Material Editor* (see [here](#)). Changes made to these settings will be immediately reflected in the *Advanced Material Editor*.
- **Origin:** defines the point of origin of the projection – e.g., when mapping in spherical coordinates, defines the center of the sphere.

Natural Grain



Natural Grain Texture Options



The aim of this feature is to produce more natural or realistic color variations for terrains, ground, or any natural element in the scene, and provide an easy access to settings such as contrast, balance or roughness.

- **Coloring Mode:** Natural grain is selected.

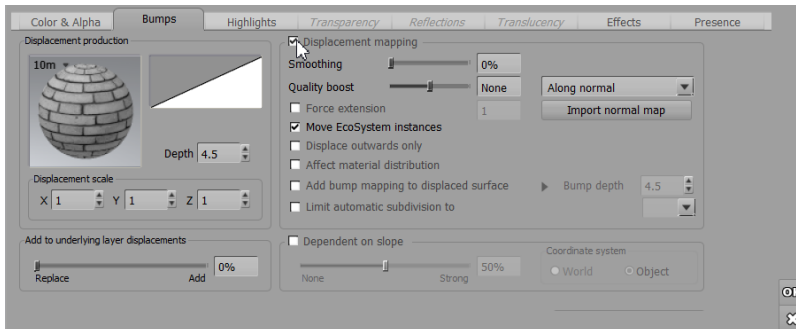
First you must indicate what is the basic color. Click on the first color box to use the color chart to select the color you wish. If you wish to mix with another color, add a second color. The lock icon toggles the second color on and off.

There are controls for you to manipulate the color(s).

- **Scale:** defines the overall scale of the noise. This should typically remain high (2000 by default) for a terrain or ground.
- **Roughness:** defines the roughness of the color.
- **Contrast:** defines the contrast of the changes made to the color.
- **Balance:** defines the balance of the two colors.
- **Distortion:** defines any distortion, if any.
- **Alpha grain:** this should be enabled only if the current material layer is not a leaf layer. In this case, it automatically adds ALPHA to produce the same kind of pattern as the color noise.

Now, if you right-click on the **Color Production** image, and select to edit the color function, the Function Graph displays with the fields listed above as well as other fields to further define the color.

Bumps Tab



Bumps tab

This tab describes bumps and holes that appear on the surface of the material. VUE generates bumps and holes on the surface of the material using a function and a filter.



You can select **Along normal** or **Normal Map**.

This is how it works: for each point on the surface, the function gives a value in the range of 0 to 1 (0 appears black on the preview of the function, 1 white). This value is then transformed by the filter into another value in the range of 0 to 1 that indicates the depth of the hole (or height of the bump) at this point (0 for a deep hole and 1 for a high bump).

To change the function, double-click on the picture of the function. This will open the Function Visual Browser.

Use the scaling controls to scale the function along the X, Y and Z axes.

If necessary, use the filter to modify the bump profile relative to the values of the function (double-click the filter).

Finally, indicate a **Bump gain** in the corresponding box. The bigger the value, the bumpier the surface will be.

Displacement Mapping

To activate displacement mapping for this material, check the **Displacement mapping** option. All settings become available.

The **Smoothing** slider is used to remove any high frequency artifacts caused by displacement settings.

Use the **Quality boost** slider to increase the amount of detail that is added to the geometry.

Very high values will result in adding micro-polygons that are not even visible in the final render. If the results look jagged, increase the setting (you should only do this when you are finalizing your work in high quality render modes, see [here](#) for an explanation). The higher the setting, the better the material will look, but the longer it will take to render and the higher the memory overhead.

Note:

Using displacement mapping is extremely easy in VUE: just check a box! However, you should be aware that this feature adds an incredible level of complexity and memory overhead to your scenes. Use displacement mapping with care – especially when creating very high resolution renders – because the amount of data added to the scene can become daunting. When creating such renders, you might like to consider reducing the quality boost of your materials, or baking your objects to a set resolution (see [here](#)).

Displacement mapping is designed to work with bump functions that output values in the standard range of -1 through 1. Any values outside this range will be clamped (i.e.



when using fractal nodes with large features). When using displacement mapping, make sure that your bump production functions do not output values beyond this standard bump range. Values outside the valid range will automatically be clamped. This does not affect the amplitude of the displacement. You can create arbitrarily large displacements by entering large values of bump gain.

- **Force extension:** check this option to set the displacement extension manually. The displacement extension is a parameter that controls the maximum amount of possible displacement. Any value beyond this limit will be clamped (saturated). By default (when the **Force extension** option is not selected), the extension is automatically evaluated so as to encompass all possible displacement values generated by your bump production function. However, it may happen that values “outside” the extension range are generated, which will result in flat displacement areas appearing on the displaced objects (these areas are saturating). This can be fixed easily by turning on the **Force extension** option and increasing the default extension by hand. Conversely, if you want to create flat areas in the displacement, you could force a lower value for the extension.
- **Remember:** if you notice that parts of your displacement are clamped (they appear as flat surfaces), this indicates that the bump production function is outputting values outside the extension range. You can fix this by forcing a greater value for the maximum extension using the **Force extension** option and increasing the default extension by hand.
- **Move EcoSystem instances:** when this option is selected, the EcoSystem instances that are attached to the material will be automatically repositioned according to the amplitude of the displacement (so that e.g. trees will always remain at the surface of the object, despite the fact that this surface is being displaced).
- **Displace outwards only:** when this option is selected, displacement values are adjusted so as to only produce positive values. As a result, the surface of the object will only be displaced outwards. This option is provided for compatibility with other applications that do not support negative displacement values (typically, in order to achieve similar results when using bitmap displacement maps created with such applications).
- **Re-evaluate material distribution after displacement:** when this option is selected, VUE will re-evaluate the contribution of environment-sensitive materials after displacement has been applied (typically so that the new displaced slope can be taken into account to determine material contribution).
- **Add bump mapping to displaced surface:** This feature makes it possible to render **Bump** in addition to **Displacement Mapping** to produce additional details. For example, you can create displacement mapping at a certain scale and add some bump mapping at a smaller scale. The **Additional Bump Mapping** channel can be edited from the *Function Graph* by clicking on the small right-pointing arrow



to open the *Function Graph* with the **Bump** output pre-selected.
Use **Bump Depth** to set the depth of this additional bump.

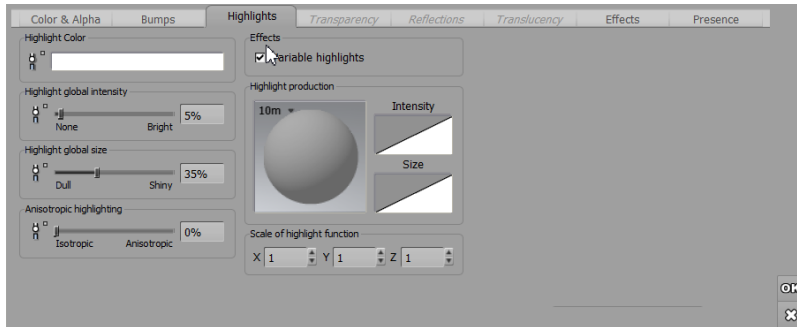
You can choose to **Limit automatic subdivision** from 1X to 32X. Check the option and choose the value in the drop box.

- **Dependent on slope:** when this option is selected, VUE produces a deeper displacement on the vertical surfaces than on flat surfaces – which is the case in the nature, typically on eroded terrains. Use the slider to indicate the strength of slope influence. You can also set the **Coordinate system** to **World** or **Object**.

Add to Underlying Layer Bumps

This option is only available when the material is a layer of a multi-layer material (and it isn't the bottom-most layer). Use this setting to control how the bumps of the current layer are added to the bumps caused by layers beneath it. If the value is 0%, the bumps of the underlying layer are replaced by the bumps of this layer. If it is 100%, the bumps are added.

Highlights Tab



Highlights tab – Simple Material

This tab describes the surface quality of the material (shiny or dull). The specular highlights create spots of light on the surface of the object, in the direction of the light sources. The smoother the surface, the more concentrated and bright the spots will be (e.g., think of polished marble).

The highlights are built with two parameters: the intensity of the light spots that appear on the surface and the size of the spots.


Highlight color gives a uniform color shade to highlights. This is useful for modeling pearl-like materials (where highlights take on a blue color).



The **Highlight global intensity** corresponds to the average intensity of the light spots. Indicate a brilliance percentage (0% = no spots, 100% = very intense spots).

The **Highlight global size** controls the average concentration (size) of the light spots. Indicate a concentration percentage (0% = big spots for dull materials, 100% = small spots for smooth materials).

Anisotropic highlights are used to simulate the special type of highlights that appear on woven or fibrous materials. They are particularly useful to create realistic hair effects. Anisotropic highlights appear around a privileged direction, known in VUE as the “Scratch direction”.

You can drive each one of these 4 parameters independently with a function by pressing the corresponding **Drive with a function** icons (). For global intensity and size, this can also be achieved using the Variable highlights option described below. However, by extracting the parameters, you can drive the two parameters each one by a function that is in no way correlated to the other. See [here](#) for further details on driving material parameters with functions.

Variable Highlights

If you want the characteristics of the specular highlights to depend on position, select **Variable highlights**.

VUE can generate variable highlights from a function and two filters, the first of which indicates the **highlights intensity** and the second the **highlight size**.

This is how it works: for each point of the surface, the function returns a number in the range of 0 to 1 (0 appears black on the preview and 1 white). The number is then transformed by the filters into an intensity and a size at this given point (0 for a dull surface, 1 for a shiny one). The maximum variable highlight value is the overall highlight value.

To modify the function, double-click on the picture of the function. This will open the Visual Function Browser.

Use the scaling controls to scale the function along the X, Y and Z axes.

Indicate a highlight intensity with the **Intensity** filter (double-click the filter).

Indicate a highlight size with the **Size** filter (double-click the filter).

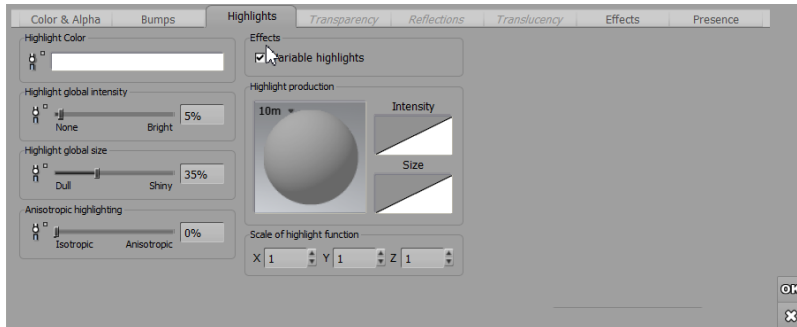
Bypassing the Standard Highlight Model

When editing material functions, the **Function Graph** defines an additional output that does not correspond to any specific material setting. This output is known as **Highlight**



Value. It expects a color value. If you connect a color node to that output, the color will be used as the highlight value. It will be evaluated for each light source, so that you can create a totally custom highlighting profile (e.g. create strong highlights at low angles of incidence).

Highlights



Highlights tab – Simple Material

This tab describes the surface quality of the material (shiny or dull). The specular highlights create spots of light on the surface of the object, in the direction of the light sources. The smoother the surface, the more concentrated and bright the spots will be (e.g., think of polished marble).


The highlights are built with two parameters: the intensity of the light spots that appear on the surface and the size of the spots.

Highlight color gives a uniform color shade to highlights. This is useful for modeling pearl-like materials (where highlights take on a blue color).

The **Highlight global intensity** corresponds to the average intensity of the light spots. Indicate a brilliance percentage (0% = no spots, 100% = very intense spots).

The **Highlight global size** controls the average concentration (size) of the light spots. Indicate a concentration percentage (0% = big spots for dull materials, 100% = small spots for smooth materials).

Anisotropic highlights are used to simulate the special type of highlights that appear on woven or fibrous materials. They are particularly useful to create realistic hair effects. Anisotropic highlights appear around a privileged direction, known in VUE as the “Scratch direction”.

You can drive each one of these 4 parameters independently with a function by pressing the corresponding **Drive with a function** icons (). For global intensity and size,



this can also be achieved using the Variable highlights option described below. However, by extracting the parameters, you can drive the two parameters each one by a function that is in no way correlated to the other. See [here](#) for further details on driving material parameters with functions.

Variable Highlights

If you want the characteristics of the specular highlights to depend on position, select **Variable highlights**.

VUE can generate variable highlights from a function and two filters, the first of which indicates the **highlights intensity** and the second the **highlight size**.

This is how it works: for each point of the surface, the function returns a number in the range of 0 to 1 (0 appears black on the preview and 1 white). The number is then transformed by the filters into an intensity and a size at this given point (0 for a dull surface, 1 for a shiny one). The maximum variable highlight value is the overall highlight value.

To modify the function, double-click on the picture of the function. This will open the Visual Function Browser.

Use the scaling controls to scale the function along the X, Y and Z axes.

Indicate a highlight intensity with the **Intensity** filter (double-click the filter).

Indicate a highlight size with the **Size** filter (double-click the filter).

Bypassing the Standard Highlight Model

When editing material functions, the [Function Graph](#) defines an additional output that does not correspond to any specific material setting. This output is known as **Highlight Value**. It expects a color value. If you connect a color node to that output, the color will be used as the highlight value. It will be evaluated for each light source, so that you can create a totally custom highlighting profile (e.g. create strong highlights at low angles of incidence).



Transparency Tab


This tab controls transparency and refraction over the surface of the material. Transparency can also be controlled via the alpha setting in the **Color** tab. This tab is not available in Simple materials that are layers of multi-layer materials (unless it is the bottom-most layer).

Note:

Alpha transparency and refraction are different in that transparency does not affect the direction of light, whereas refraction indicates that the direction of light is modified by the index of refraction of the material.

Incident light arriving on the surface of a material divides into 3 different lights:

- diffused light, sent by the surface in all directions, more intensely in the highlight direction; this makes the color of the surface,
- reflected light that bounces off the surface of the material, and
- refracted light (or transmitted); it is the light that goes through the surface and penetrates the material.

Some of the parameters in this tab can be driven independently with functions by pressing the corresponding **Drive with a function** icons (). For global transparency and blur, this can also be achieved using the Variable transparencies option described below. However, by extracting the parameters, you can drive the two parameters each one by a function that is in no way correlated to the other. See [here](#) for further details on driving material parameters with functions.

Global Transparency


The quantity of light diffused is equal to the quantity of incident light less the quantity of refracted and reflected light.

Indicate the amount of light that penetrates the surface of the object using the **Global transparency** control.

If you would like objects seen through the material to be blurred (because the material is impure or distorted), raise the value of **Blurred transparencies** up to non zero.



Refraction Index



Light subdivides as it hits the surface of a material

The refraction index (a.k.a. Index of Refraction) identifies the optical density of the material. It bends rays of light that cross the surface of the material, thus creating the magnifying glass effect, and giving the impression that a stick in the water is broken. Common refraction indexes are:

- **Air:** Index of Refraction = 1.00. This is the reference IOR.
- **Water:** Index of Refraction = 1.33.
- **Glass:** Index of Refraction = 1.52.

You can modify the Refraction index of the material using the **Refraction index** control.

Note:

refraction indexes less than 1 are seldom observed, and would correspond to materials less dense than air...

Flare

When light is seen from behind a partially transparent material, it will cause the surface of the material to become very bright. This is called flaring. Flaring is a bit to transparency what highlights are to reflectivity. It will not occur at the surface of perfectly transparent materials. It is maximum for a transparency amount of 50%. Flaring is particularly useful for clouds.

You control flaring through two settings: **Intensity** and **Span**. Flare span is the area around the light that will flare-up. Larger values yield bigger flares.



Turn Reflective with Angle

When a transparent refractive (i.e. IOR different from 1) material is seen at a low angle of incidence, it sometimes happens that it becomes reflective. Take a piece of glass, and look at it from the side. You will notice it acts as a mirror. The same thing happens with water: looking vertically, you see through it, but in the distance, it becomes reflective.

This behavior can be captured by VUE, using the **Turn reflective with angle** control. You can even fine-tune that effect using the slider. Zero cancels it. Values around 40% yield good results.

Effects

- **Fuzzy:** selecting this option will make the edges of the object become fuzzy (blurred) instead of being sharp. The **Fade out** control changes into a **Fuzziness** control, letting you adjust the strength of the effect. This, combined with variable transparencies, is the key to making realistic clouds (see tutorial on making clouds).
- **Additive:** When this option is selected, the color of the material is added to that of the background, yielding luminous, immaterial objects. This is an interesting effect for making light rays (see tutorial on making light beams).
- **Physical transparency:** this option allows for a physical simulation of light volumetric scattering and absorption through transparent media. It is particularly suited for realistic glass and water simulation.
- **Thin surface:** This setting is used to render thin, one-sided surfaces such as a window, with no refraction. You still have the reflectivity effect, but there no refraction, with distortion. To use, the material must be one-sided. The refraction index cannot equal “1” and **Turn reflective with angle** must be activated.
- **Variable transparency:** if you want transparency to be dependent on position, select the **Variable transparency** option.

VUE generates variable transparency from a function and a filter that indicates the amount of transparency depending on the value of the function.

This is the way it works: for each point on the surface, the function generates a number in the range of -1 to 1 (-1 is black on the preview of the function and 1 is white). This number is then transformed by the filter into a transparency value. The maximum variable transparency is equal to the overall transparency.

To change the function, double-click on the preview of the function. This will open the Function Visual Browser. Use the scaling controls to scale the function along the X, Y and Z axes.

Indicate the values of the transparency using the filter (double-click on the filter).



Absorption & Scattering

- **Mode:** specifies the desired accuracy of the water simulation. Three modes are available:
 - **Fast ocean:** an analytical model suitable for ocean simulation and should be used only for sea rendering. It is really fast but will only consider sun lighting and with no volumetric shadowing calculations. Using it over other objects might lead to unpredictable results.
 - **Direct volumetric light (default):** ambient lighting is roughly approximated within the medium, but direct lighting calculations are fully raytraced, allowing for the simulation of volumetric shadowing through the medium. It is therefore slower than the **Fast ocean** mode. This can be applied to any object.
 - **Indirect volumetric light:** all lighting is fully raytraced, thus faithfully simulating direct and indirect shadowing through the medium. However, it is slower than the other two modes. It should only be used when necessary, since volumetric ambient shadowing is generally subtle. Most often, **Direct volumetric light** mode should work fine, and will be faster to compute.
- **Depth:** specifies the distance beyond which objects aren't visible anymore through the medium. It therefore drives the overall optical thickness of the medium.
- **Absorption:** specifies the remaining fraction of light after absorption by the medium particles. It corresponds to the overall coloring of objects seen in the distance through the medium.
- **Scattering:** specifies the remaining fraction of light after being scattered in another direction through the medium. It corresponds to the coloring of light while being diffused through the medium.
- **Anisotropy:** specifies how light is statistically scattered through the medium. A forward anisotropy will mostly scatter light ahead, while a backward anisotropy will mostly scatter light back towards the incoming direction.
- **Quality boost:** specifies the simulation quality. A higher value will lower noise at render, but will take longer to compute.

Fading Out

When light travels through a material, it progressively fades out with distance. This is why deep water always looks blue. VUE captures such effects: indicate a **Fade out rate**, that is the depth at which light has completely disappeared and the color becomes that of the fade out color. If the value is small, the material will be clear, and you will see deep into it. If it is zero, no fading out will ever occur.

Indicate the **Fade out color** (double click on it), that is the color of the material when light has traveled deep into it.



Objects placed behind a transparent material, receive light of a color depending on the distance traveled through the transparent material. As light travels further through the material, it takes on a particular color that can be defined using the **Light color** control. This is how you make blue water look green when sand gets close to the surface.

Enable Dispersion

Dispersion simulates the spectral decomposition of light through refractive media, like when light gets refracted through a prism. It corresponds to a physical law which tells that the index of refraction actually varies with incident light wavelength. To make this option work, **Compute physically accurate caustics** render option must be enabled.

The dispersion value drives the amount of spectral dispersion when photons get refracted. Small values will tend to keep light quite concentrated, showing only a slight spectral decomposition, while large values will separate the spectrum more clearly.

Reflections Tab



Reflections tab – Simple Material


This tab controls how the surface of the material reflects light.

Indicate the amount of reflected light in the **Global reflectivity** box. Note that if the amount of reflected light + the amount of transmitted light exceeds 100%, the material will become “luminous”. If you are using radiosity it will actually be emitting light. **Minimal reflectivity at viewing angle** can be combined with the **Sensitivity to incidence angle** to modify the global reflectivity. If there is no sensitivity to incidence angle, the reflectivity is homogeneous. Also, a minimal reflectivity of 100% gives a homogeneous reflectivity as well.






Minimum reflectivity to angle 0%




Minimum reflectivity to angle 50%



Minimum reflectivity to angle 100%

Working with **Global reflectivity**, use the **Sensitivity to incidence angle** slider to change the angle of reflectivity.

If you would like the surface of the material to be imperfectly reflective, resulting in distant objects appearing blurred inside the reflections, push up the value of **Blurred reflections** to non-zero.

Some of the parameters in this tab can be driven independently with functions by pressing the corresponding **Drive with a function** icons (). For global reflectivity and blurred reflections, this can also be achieved using the Variable reflections option described below. However, by extracting the parameters, you can drive the two parameters each one by a function that is in no way correlated to the other. See [here](#) for further details on driving material parameters with functions.



Variable Reflections

If you want the reflection to be dependent on the position, select **Variable reflectivity**. You must have set some amount of reflectivity for this option to be available.

VUE generates variable reflections using a function and a filter that indicate the local amount of reflection depending on the value of the function.

This is the way it works: for each point on the surface, the function generates a number in the range of -1 to 1 (-1 is black on the preview of the function and 1 is white). This number is then transformed by the filter into a reflection amount. The maximum variable reflectivity is equal to the overall reflectivity.

To change the function, double-click on the picture of the function. This will open the Visual Function Browser.

Use the scaling controls to scale the function along the X, Y and Z axes.

Indicate the amount of reflection using the filter.

Reflection Map

If you would like to use a reflection map to simulate the reflections on this material, check the **Use reflection map** option (you must have set some amount of reflectivity for this option to be available). The settings in the Reflection map group become available, letting you define a custom reflection map for this material. It also lets you set the default reflection map. Please turn [here](#) for further details on reflection maps.

To define a new reflection map, press the **Load** button or double click the reflection map preview to open the Bitmap Browser. Select the picture you want to use as reflection map and validate. A message should appear if your picture doesn't loop smoothly horizontally, and offer to create a smoothed junction between both edges. This is because the reflection map is mapped onto an imaginary sphere, thus looping horizontally. If you click **Yes**, then VUE will add a smooth transition strip from the right to the left border of the bitmap in order to avoid a sharp transitions in the reflection map. Of course, if you don't want to alter the bitmap, click **No**. Your bitmap should now be displayed in the reflection map preview.

If you would rather use the default reflection map for this material, click the **Use default** button (the default reflection map is also accessible via the *Atmosphere Editor* – see [here](#)).

You can adjust the **Azimuth**, **Pitch** and **Horizon offset** parameters to rotate the map on the sky dome (see [this](#) for more information).

If the reflections on this material are blurry (**Blurred reflections** set to non zero), the reflection map will automatically be blurred accordingly. Also, the reflection map preview

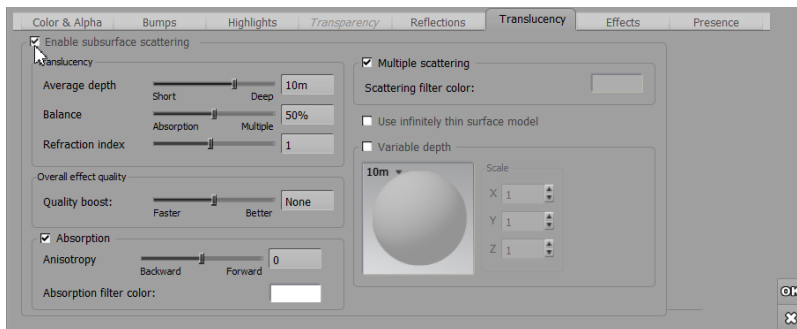


will appear more or less blurred according to the blurring amount.

If you click the **Set default** button, the current reflection map will become the default and will be applied to all materials that use the default reflection map. The U and V offsets will also be applied to all materials that use the default reflection map.

If you use the **Force use of reflection map** option in the *Render Options* dialog (see [here](#)), all materials with reflective surfaces will use reflection maps. If no reflection map was used for a material, it will use the default reflection map.

Translucency Tab




Translucency tab – Simple Material

This tab controls the translucent characteristics of the material. Translucent materials react to light in a very different way than “regular” materials.

With a regular material, incident light is either diffused, reflected, or refracted. With translucent materials, the light is also absorbed by the surface of the material and re-emitted at a point that is not the same as the point where it arrived. The technique used to capture this effect is known as subsurface scattering.



Translucency Settings



Light is scattered inside translucent materials

To enable translucency, activate sub-surface scattering for the material by checking the **Enable sub-surface scattering** box. When this option is selected, the sub-surface scattering controls become accessible. The Translucency group displays a set of controls that are common to absorption and multiple scattering:

- **Average depth:** this setting controls how translucent the material is. It indicates the average distance traveled by light inside the material. Typical “real-world” values are in the range of a fraction of a millimeter to a couple of centimeters (for wax-like materials). You must make sure that your translucent objects are compatible in size in order to see the effects of subsurface scattering (don’t expect to see anything – except desperately-long render times – if you assign a translucent material with an average depth of 1 inch to a square-mile terrain!).
- **Balance:** this setting controls the amount of absorption vs. multiple-scattering that takes place inside the material. The default is 50%, which means that all absorbed light is redistributed by the multiple scattering, but you can achieve interesting effects by varying this balance.
- **Refraction index:** this is identical to the [refraction index](#) in the **Transparency** tab. When you enable sub-surface scattering, the refraction index control in the **Transparency** tab becomes disabled.

Absorption

Check this option to enable absorption for this material.

The **Anisotropy** setting controls how directional the scattering is inside the material. A



value of 0 indicates that the light is scattered equally in all directions, a negative value indicates that light is scattered mostly backwards, and a positive value, that light is scattered forwards (the usual scattering).

- **Absorption filter color:** this setting controls the overall color that light picks up as it travels through the translucent material (the red color when you put your finger over a light source).

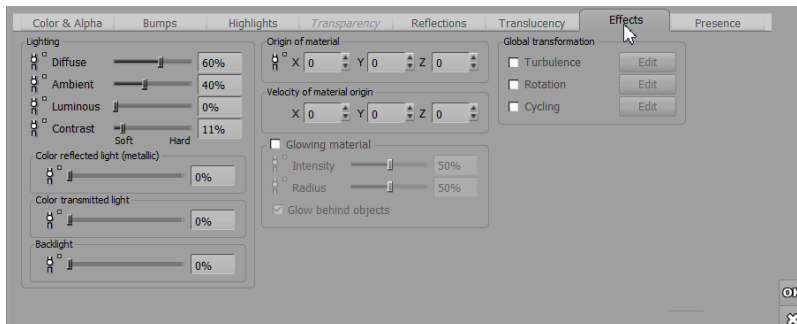
Multiple Scattering

- **Scattering filter color:** this setting controls the diffuse color of the material (the pink color of the skin). Because multiple scattering bounces light in all directions, there is no preferred direction for this effect (unlike absorption).
- **Use infinitely thin surface model:** select this option when rendering one sided translucent materials, such as planes.
- **Variable depth:** check for variable depth translucency. Reset the scale if necessary. Right-click to open the Function Graph to control the depth with a function.

Quality Boost

- **Quality boost:** use this slider to increase the number of samples taken to compute the translucency. If the results look noisy, increase the setting (you should only do this when you are finalizing your work in high quality render modes, see [here](#) for an explanation). The higher the setting, the better the material will look, but the longer it will take to render.


Effects Tab



Effects tab – Simple Material

This tab controls the lighting characteristics of the material, and miscellaneous effects.



Some of the parameters in this tab can be driven independently with functions by pressing the corresponding **Drive with a function** icons (). See [here](#) for further details on driving material parameters with functions.

Lighting

The surface of the material receives light from light sources (e.g. the sun) and from the environment, and may react differently to each of these types lights.

The **Diffuse** lighting parameter controls the way the material reacts to light coming directly from light sources.

The **Ambient** parameter controls the way the material reacts to ambient lighting.

By default, these values are respectively 60% and 40%.

It is usually not recommended that you modify these values for a material except under very special conditions (e.g. you could make a cloud more reactive to ambient light, because the cloud is far enough, and physically different from solid objects in the foreground). This is because it may cause a mismatch between the different materials of your scene. If you want to modify the diffuse and ambient settings, you're probably better off doing so on a scene level (see the *Atmosphere Editor*, [Light Tab](#) for details on how to do this). Also, the total amount of Diffuse + Ambient should always be equal to 100%.

If you want to create materials that seem to emit light, use the **Luminous** setting. Keep in mind that luminous objects do not cast real light, though (except when using the Global Radiosity lighting model – see [here](#)). If you wish to have a luminous object actually cast light, you could put a light source inside it and turn off **Casts shadows** for the material.

Luminous lighting is not affected by the global settings of the scene. This works particularly well when used in conjunction with Glow (see below), because it emphasizes the impression that the object is emitting light.

- **Contrast:** this setting adjusts the speed at which the material goes from light to shadow. This is useful for modeling fluffy materials.
- **Color reflected light:** to give a metallic aspect to a reflective material, select this option. This will give the color of the surface to highlights and reflections.
- **Color transmitted light:** selecting this option with a transparent material will give the color of the surface to the light crossing it. This is a great for colored glass and church windows.
- **Backlight:** use this option when a material is supposedly thin enough to let some light show through when illuminated from behind. This is typically what happens when the sun shines behind a leaf. The leaf isn't dark – light passes its surface, although it isn't transparent.



Origin of Material

These fields let you offset the material in material coordinate space. This enables the precise positioning of materials on objects.

If the material is completely animated (read [here](#) for details on material animation), VUE will automatically compute the corresponding velocity, and fill the **Velocity of the material origin** fields with the resulting values.

Velocity of Material Origin

These fields let you define a displacement with time of the origin of the material. As a result, the material will be changing as time passes. Defining a Velocity of material origin creates a Velocity Animated material (read [here](#) for details on the different types of material animation). The keyword “Time dependent material” appears in the caption of the *Material Editor*.

Changing the velocity of the material origin of a completely animated material will set the fields in the **Origin of material** group.

Glowing Material

Select this option to create a haze of light around the material (turn [here](#) for an explanation of glow). Keep in mind that glow is a post-processed effect added once the rendering pass is complete. So when the render starts, you won’t be able to see the glow. You need to wait until rendering is complete to be able to judge the effect.

When you select this option, the controls in the group become available. The **Intensity** slider controls the amount of light in the glow, and the **Radius** slider controls the average size of the haze of light.

Select **Glow behind objects** if you want the glow to be visible on objects that are placed in front of the glowing material. Uncheck it if you want the glow effect to be masked by objects in front.

The color of the glow is determined by the color of the material. Dark materials won’t glow much.

Adding some luminous lighting to glowing materials emphasizes the glowing effect (see above).

Remember that, although glowing materials give the impression that they are emitting light, no real light is actually cast by the material onto other objects in the scene. If you want a glowing object to actually cast light, you can add a point light inside it, and uncheck the **Casts shadows** option.

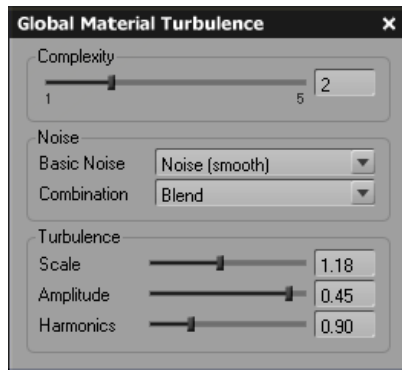


Global Transformation

Selecting options in this group will apply global modifications to the material.

When you select an option, the corresponding **Edit** button becomes enabled. Pressing this button displays a dialog that lets you adjust the effects.

Global Turbulence



Turbulence Editor

Press the **Edit** button to the right of the **Turbulence** checkbox to display the *Turbulence Editor*.

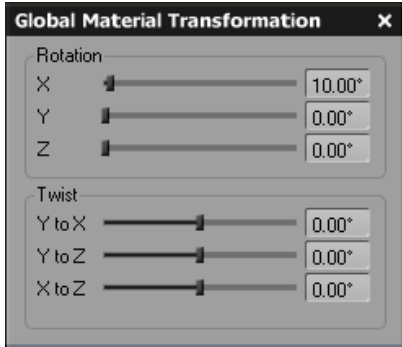
Using a noise, turbulence repeatedly displaces the location at which the material or the function is being evaluated. Turbulence is defined by 4 parameters (complexity, scale, amplitude and harmonics), a noise type and a combination mode.

- **Complexity:** defines the number of times the noise is repeated.
- **Amplitude:** is the average displacement applied by the noise to the material or to the function layer.
- **Scale:** controls the frequency at which the Noise functions vary relative to position.
- **Harmonics:** characterize the way the noise is scaled each time it is added: for each new addition, scale and amplitude are multiplied by the harmonic parameter. If the complexity is equal to 1, the harmonic parameter has no meaning.
- **Suggestion::** to understand correctly the effects of turbulence, watch the variation of a material made from simple functions (e.g. use a rectangular wave noise to drive the color channel).
- **Basic noise:** defines the type of noise that is added to the current position, and
- **Combination:** indicates how successive noise applications should be combined. For



full details on combination modes, refer to the section on the [Function Graph](#).

Rotation



Transformation Editor

Pressing the **Edit** button on the side of the **Rotation** checkbox brings up the *Transformation Editor*.

This dialog lets you indicate a rotation angle around each of the world axes, as well as a twisting angle of these axes one towards another.

Cycling

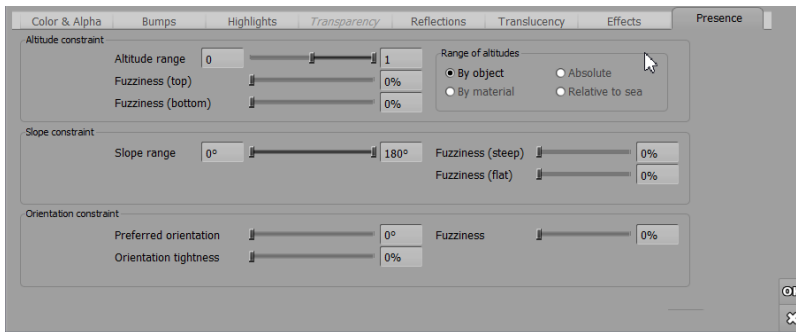
- **Cycling:** is a large scale perturbation of the material that helps to prevent it from looking too repetitive.

Presence Tab

If the current material is a layer of a multi-layer material, and if it is not the bottom-most layer on the stack, an additional tab called Presence is available. This tab lets you control how the environment affects the presence of the current layer.



Altitude Constraint



Presence tab – Layered Material

This group lets you control how altitude influences the presence of the layer:

- **Altitude range:** this dual slider lets you define the range of altitudes in which the current layer appears (provided it is not transparent at this point). The layer will not appear outside this range.
- **Fuzziness:** this setting controls how “suddenly” the changes to the layer presence are made in response to altitude. High values mean that the layer appears very gradually in its altitude range, whereas low values will result in the layer appearing as a solid strip.
- **Range of altitudes:** this lets you define in what coordinates the altitude range is defined:
 - **By object:** in this mode, the range is relative to each object to which the material is applied.
 - **By material:** in this mode, the range is relative to all the objects that use this material.
 - **Absolute:** in this mode, the range of altitudes is expressed in global coordinates.
 - **Relative to sea:** the altitude is computed from the sea level and not from zero.

Slope Constraint

This group lets you control how the local slope influences the presence of the layer:

- **Slope range:** this dual slider lets you define the range of slopes in which the current layer appears (provided it is not transparent at this point). The layer will not appear outside this range. Values to the right end of the slider indicate flat surfaces, and values to the left indicate upside-down surfaces. Intermediate values indicate vertical



surfaces. Slope values can range from -180 to +180 degrees.

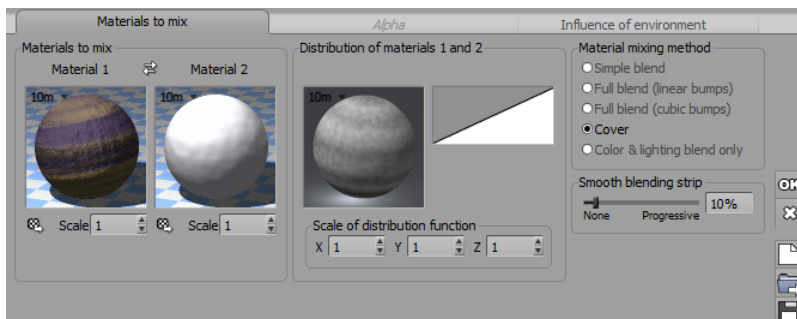
- **Fuzziness:** this setting controls how “suddenly” the changes to the layer presence are made in response to slope. High values mean that the layer appears very gradually in its slope range, whereas low values will result in the layer appearing as a solid strip on areas of appropriate slope.

Influence of Orientation

This group lets you control how the local orientation influences the presence of the layer:

- **Preferred orientation:** this setting controls the orientation of the surface that is the most favorable to the presence of the layer.
- **Orientation influence:** this setting controls the influence of orientation on the presence of the layer.
- **Fuzziness:** this setting controls how “suddenly” the changes to the layer presence are made in response to orientation. High values mean that the layer appears very gradually on surfaces of the preferred orientation, whereas low values will result in the layer appearing as a solid strip on areas of preferred orientation.

Mixed Materials



Material Editor – Mixed Materials

A mixed material takes two materials and mixes them together. Mixed materials supply a number of rules that you can use to define the way the materials mix together, including rules that depend on the environment.

Using these rules, a mixed material decides if, at any given point, it should be the first or the second material that shows, or a blend of the two.



You can mix together simple materials, materials that are themselves a mix of other materials (create nested material hierarchies for amazing effects!), or even EcoSystem materials (mix an EcoSystem of fir trees with an EcoSystem of rocks to place trees at low altitudes and rocks higher up!).


When you select the **Mixed materials** option at the top of the Material Editor, the tab part of the editor changes, displaying three tabs.

The **Mixing proportions** slider lets you adjust how much each of the two materials that are mixed together will be visible. Pushing the slider to the right will have more of the second Material showing, while pushing it to the left will show more of the first.

The **Alpha Channel** can also be edited for **Mixed materials**.

Materials to Mix

The first tab lets you choose the materials that will be mixed together, and the way they will be mixed. Volumetric materials cannot be mixed together.

Change the materials by loading materials that already exist on the disk with the  button, or by double-clicking on the material preview to edit it.

Materials inside a mixed material can be scaled independently using the **Scale** controls. This only modifies the size of the material once it is applied to an object. A scale equal to 1 does not change the size of the material.

Distribution of Materials

To decide if the mixed material should display the first or the second material, or a blend of the two, Vue basically uses a function and a filter.

This is how it works: for each point on the surface, the **Distribution function** generates a value in the range of 0 to 1 (0 appears black on the preview of the function and 1 white). This value is then transformed by the filter into another value in the range of 0 to 1, which is then compared to the **Mixing proportions** setting. If it is much less than this setting, material 1 is displayed. If it is far greater, material 2 is displayed. If the result is close to the Mixing proportions, inside a range indicated by the **Smooth blending strip**, a blend of the two materials is computed, in order to get smooth transitions from one material to the other.

The way materials mix can also be modified by local slope, altitude and orientation. See second tab.

To modify the function, double-click on the picture of the function. This will open the



Function Graph.

Use the scaling controls to scale the function along the X, Y and Z axes.

If necessary, use the filter to further adjust the distribution of materials.

Smooth Blending Strip

The width of the strip inside which materials are blended can be adjusted using the **Smooth blending strip** control. Pushing it to the right will make for smoother transitions, while pushing it to the left will yield fast transitions.

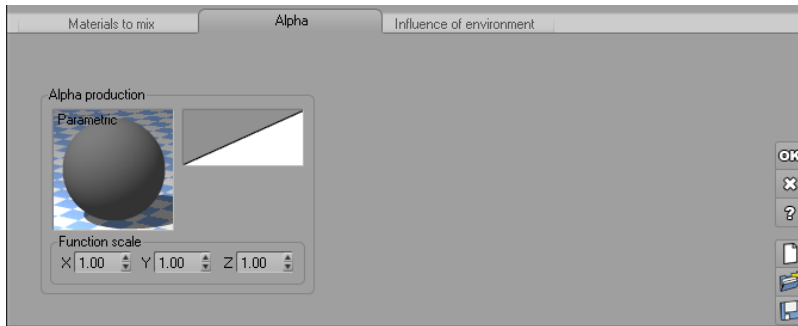
Blending Method

Inside the strip, materials are blended together. Bumps can be handled in several ways, depending on the result you are looking for:

- **Simple blend:** the surface aspect of the two materials are mixed together. This is the default. This method is compulsory to mix materials that are themselves mixed materials.
- **Full blend (linear bumps):** the characteristics of the two materials are blended together before the material is rendered. Surface heights are blended linearly, resulting in a chamfer between both materials (provided one of them has a higher surface bump than the other).
- **Full blend (cubic bumps):** same as blend bumps, except heights are blended following a cubic rule. The result is a rounded chamfer between the two materials (like snow on rocks).
- **Cover:** no smooth transition for colors, only for bumps. Material 2 seems to cover up material 1. Inside the transition strip, only material 2 is visible.
- **Color and lighting blend:** in this mode, only color and lighting (ambient and diffused) features are used from Material 2, retaining all other features of material 1. This is useful for shifting colors of a material, without having to duplicate it (for instance near the water surface).



Alpha

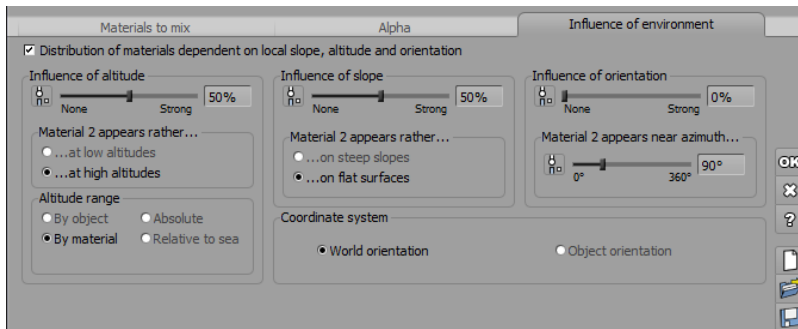


Alpha tab – Mixed Materials

Alpha production can be edited on this tab. You can load a new filter by right-clicking the displayed filter and selecting **Load Filter** to access the Filter Browser. Or, you can edit the current filter by right-clicking the displayed filter and selecting **Edit Filter** to open the *Material Alpha Filter* dialog.

You can also modify the alpha channel by using the *Function Graph*. Access the *Function Graph* by right-clicking the **Alpha production** picture and selecting **Edit Function**. You can also select to load a function from the Function Browser by selecting **Load Function** from the menu.

Influence of Environment




Influence of Environment tab – Mixed Materials

The third tab of the mixed materials editor lets you define the influence of slope, altitude,



and orientation on the way the two materials are mixed together. Controls in this tab become active once you select the option **Distribution of materials dependent on local slope, altitude and orientation**.

Some of the parameters in this tab can be driven independently with functions by pressing the corresponding **Drive with a function** icons (). See [here](#) for further details on driving material parameters with functions.

Influence of Altitude

Influence of altitude adjusts the influence that altitude has on the distribution of materials. Zero means that the distribution is not affected by altitude. Non-zero values mean that material 2 will appear more often at high (or low) altitudes. Indicate whether the second material should appear at high or at low altitudes by selecting the requested box.

Influence of Slope

Influence of slope adjusts the influence that slope has on the distribution of materials. Zero means that the distribution is not affected by slope. Non-zero values mean that material 2 will appear more often on steep slopes (or on flat surfaces). Indicate whether the second material should appear on steep parts of your scene, or on flat surfaces by selecting the requested box.

Influence of Orientation

Influence of orientation adjusts the influence that orientation has on the distribution of materials. Zero means that the distribution is not affected by orientation. Non-zero values mean that the second material will appear more often on parts of the scene that are facing the azimuth indicated by the **Azimuth** slider.

For a realistic distribution of snow on a landscape, you could for example indicate that snow appears mostly at high altitudes and that it tends to gather on flat surfaces. You could also make snow accumulate on faces of the landscape that are in shadow (using orientation).

Altitude Range

Altitude Range lets you define in what coordinates this range is defined:

- **Per object:** in this mode, the range is relative to each object to which the material is applied.
- **Global:** in this mode, the range is relative to all the objects that use this material.
- **Absolute:** in this mode, the range of altitudes is expressed in global coordinates.

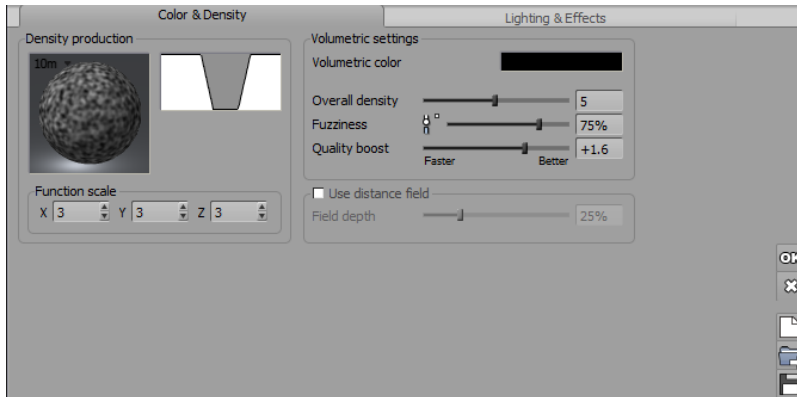


- **Relative to sea:** in this mode, the range of altitudes is expressed in global coordinates, but the altitude is computed from the sea level and not from zero.

Coordinate System

This group lets you indicate if the environment considered for mixing the materials should be linked to the object itself, or to the world. If the environment is linked to the object, rotating the object won't change the distribution of materials on the surface (the distribution moves with the object).

Volumetric Materials



Material Editor – Volumetric Materials

Unlike the other two types of materials that are only defined by their surface, volumetric materials are defined over the whole of their volume.

When you select the **Volumetric material** option at the top of the *Material Editor*, the tab part of the editor changes, displaying two tabs. Some of the options may change depending on whether you are editing a cloud material, or a material for an object.



Color and Density

Volumetric materials are based on a density production that indicates the local density of the material over a complete volume.

Density Production

This is how **Density production** works: for each point of the volume, the function returns a number in the range of 0 to 1 (0 appears black on the preview and 1 white). The number is then transformed by the filter into a density at this point (0 if the material doesn't exist, 1 if the material is solid). When rendering the material, Vue accumulates the density of material all along the ray of light that is traversing the material, and then computes the resulting color of the material.


To modify the function, double-click on the picture of the function. This will open the Visual Function Browser.

Use the scaling controls to scale the function along the **X**, **Y** and **Z** axes. The same for the filter.

Volumetric Settings

These settings are not available for volumetric spectral cloud layer materials in the Spectral atmosphere model. Instead, they are replaced by the **Cloud layer detail** settings (see below).

Most of the features in this section are available for the Artist products with the *Advanced Graphics* module.

- **Volumetric color:** this control indicates the global color of the material. Double-click on it to open the Color Selection dialog and modify the color.
- **Overall density:** use this slider to modify the overall density of the volumetric material. Overall density increases or reduces the average density of the material over its entire volume.
- **Fuzziness:** this slider controls the density of the material near its edges. The density of the material is automatically reduced as you get close to the edges of the material. If fuzziness is 0, the density of the material is not affected by the proximity of the edges. The greater the fuzziness, the thinner the material near the edges. This parameter can be driven by a function by pressing the **Drive with a function** icon. ().
- **Quality boost:** use this slider to increase the number of samples taken to compute the material. If the results look noisy, increase the setting (you should only do this when you are finalizing your work in **high quality** render modes. The higher the



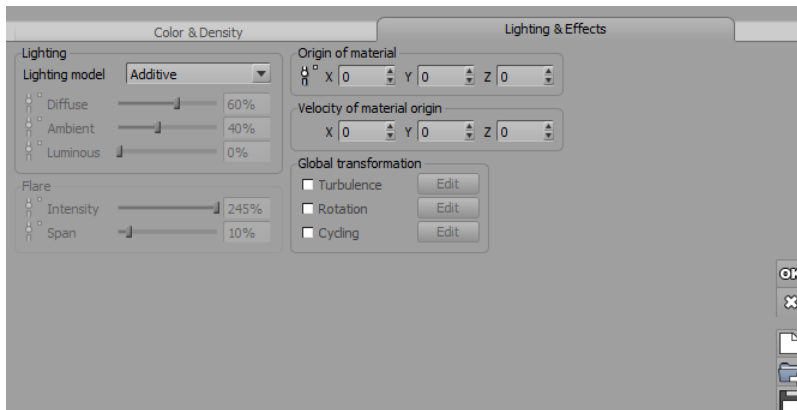
setting, the better the material will look, but the longer it will take to render.

- **Use distance field:** when this option is enabled, the volumetric density takes into account the depth inside the object. The density will automatically increase as you go deeper inside the object. When this option is checked, the **Field depth** parameter becomes active, letting you indicate the depth at which the density function reaches its maximum value.

Lighting and Effects

Some of the parameters in this tab can be driven independently with functions by pressing the corresponding **Drive with a function** icons. ().

Lighting



Lighting and Effects tab – Volumetric Materials

The controls in this group let you customize the way the material reacts to light.

Lighting model: use this drop-down list to select the lighting model used for the volumetric material.

- **Uniform:** this lighting model is the simplest. The color of the material is uniform, and only depends on the density. Light is not taken into account when computing the color. This is the default.
- **Shaded:** in this model, the influence of light and shadows is computed at the surface of the material. The Diffuse, Ambient, Luminous and Flare controls become active.
- **Additive:** this is similar to the Uniform model, except that the color of the volumetric material is added to the background instead of masking it. This is great for



fire balls, and other similar effects.

- **Volume shaded:** this is a more advanced model where the contribution of light is evaluated throughout the entire material instead of being evaluated only at the surface. The color of the material can be varied inside the volume. This is ideal for fireballs, thick smoke or explosions. Volume shaded materials are a lot slower to compute than the simpler shaded and uniform models.
- **Flat layer:** if you are editing a material for a standard cloud layer in the Spectral atmosphere model, the lighting model is locked to a model specifically designed for the rendering of flat cloud layers.
- **Volumetric layer** and **Cloud object:** if you are editing a material for a cloud layer or a MetaCloud, the lighting model is locked to a model specifically designed for the rendering of clouds.

The other controls in the lighting group are active only when a Shaded lighting model is selected. They perform the same as their equivalents in Simple materials:

The **Diffuse** lighting parameter controls the way the material reacts to light coming directly from light sources.

The **Ambient** parameter controls the way the material reacts to ambient lighting.

By default, these values are respectively 60% and 40%.

Also, the total Diffuse + Ambient should always be equal to 100%.

If you want to create materials that seem to emit light, use the **Luminous** setting. Keep in mind that luminous objects do not cast real light, though. Luminous lighting is not affected by the global settings of the scene.

If the selected lighting model is the Volumetric layer, the following checkboxes become available:

Internal shadows: when this option is selected, the shadows inside the cloud will be computed, meaning that some parts of the cloud will cast shadows on other parts of the cloud. This option produces much more realistic clouds, but increases render times significantly.

Cast shadows: when this option is selected, the cloud layer casts shadows in the atmosphere that can result in the appearance of Godrays if conditions are favorable. This option also increases render times significantly and should be used with caution, as it does not necessarily produce a noticeable improvement in picture quality. In order for Godrays to be visible, you must also enable Godrays in the *Atmosphere Editor*.



Flare

The controls in this group are only active in Shaded and Volume shaded models.

When light is seen from behind a thin volumetric material, it will cause the material to become very bright. This is called flaring. Flaring doesn't occur when the material is either too dense, or too thin.

You control flaring through two settings: **Intensity** and **Span**. Flare span is the area around the light that will flare-up. Larger values yield bigger flares.

Volume Shaded Material

The volume shaded lighting model is the most advanced volumetric material model. The color of the material is re-evaluated at each sample, as well as illumination and internal shadowing (when parts of the material cast shadows on other parts of the same material).

To let you define the color of the material inside the material, a third tab called **Volumetric Color** is added to the *Material Editor* when you select the Volume shaded model. This tab is similar in its behavior to the **Procedural Color** production of the *Advanced Material Editor* and lets you define the material color as a 3D volume.

Hypertextures

Hypertextures are a solid/volumetric hybrid that can be used to create porous materials such as corroded metal or sponge, and various special effects like water splashes. Hypertextures are also defined using a density function, however this density is used in a different way: instead of being interpreted as a gas density, the density is used to define the interface between the material and the absence of material. Wherever the density is higher than the **Overall density** setting is considered as being “inside” the material, and wherever the density is less than that value is considered as not being in the material.

If you select the Hypertexture model, a third tab called **Hypertexture Material** appears. This tab lets you define the material that appears at the surface of the hypertexture.

Origin of Material

These fields let you offset the material in material coordinate space. This enables the precise positioning of materials on objects.

If the material is completely **animated**, Vue will automatically compute the corresponding velocity, and fill the **Velocity of the material origin** fields with the resulting values.



Velocity of Material Origin

These fields let you define a displacement with time of the origin of the material. As a result, the material will be changing as time passes. Defining a Velocity of material origin creates a Velocity [Animated material](#). The keyword “Time dependent material” appears in the caption of the *Material Editor*.

Changing the velocity of the material origin of a completely animated material will set the fields in the **Origin of material** group.

Global Transformation

Selecting options in this group will apply global modifications to the material’s density production. These options work the same as for Simple materials: when you select an option, the corresponding **Edit** button becomes enabled. Pressing this button displays a dialog that lets you adjust the effects.

The editor dialogs for each type of modification are detailed in the section on [Simple Materials](#).

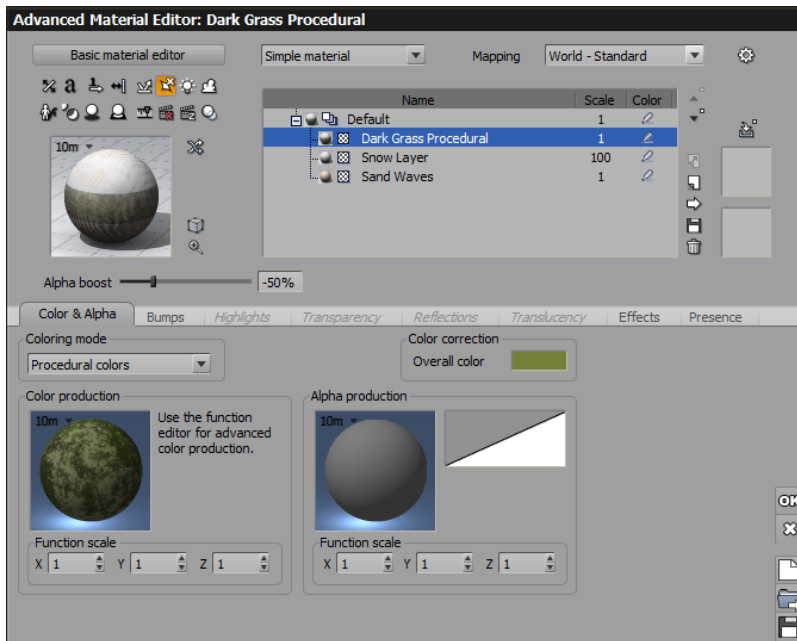
Dissolve Near Objects

If the edited material is a cloud layer material, this group of controls is available. These controls let you automatically define how cloud layers react to the proximity of other objects in the scene – for instance to let high mountains peak through the clouds.

- **Dissolve near objects:** select this option to have the density of volumetric cloud layers drop automatically near the objects in the scene.
- **Accuracy:** this setting controls the precision with which the proximity to other objects in the scene is evaluated.
- **Softness:** this setting controls how gradually the cloud dissolves near foreign objects. Low values mean the clouds will vanish abruptly near objects.
- **Distance:** this setting controls how far away from the objects the clouds are influenced. You can control the amount of dissolving using a combination of both this and the Softness setting.



Layered Materials



Layered Material – Layer List

Vue's *Material Editor* offers extended control over the creation of complex materials through a layered system. With material layers, you can:

- Add, delete, and rename layers on the fly.
- Rename materials and change scale in the layer list.
- Move layers up and down in the stack.
- Each layer has its own alpha channel.
- Each layer has its own independent reaction to altitude, slope and orientation.
- Create mixed materials with any number of sub-materials.
- Easily navigate even the most complex layered/mixed/nested materials.

Layers can be Simple, Mixed or EcoSystem materials. You cannot layer Volumetric materials. Mixed materials can be made up of layered materials.



Multi-Layer Materials

When a multi-layer (or layered) material is selected, the *Advanced Material Editor* appears as opposite. The Layer stack lists all layers used in this material.

When evaluating the layered material, Vue renders each layer in turn, starting with the top of the stack. If a layer is partially transparent (alpha is less than 1), or doesn't exist at that point (e.g. because of environment constraints), it moves on to rendering the layer beneath it on the stack. And so on, until total opacity is achieved or the bottom-most layer has been reached. Bumps at the surface of Simple materials are processed in the same way, except that bumps from a given layer are added to the bumps of the layer beneath it (unless the **Add to underlying bumps** option is set to 0 – see [here](#)).

If the layers in the layered material involve EcoSystems, the population of the EcoSystem is done from the bottom up. First, the bottom-most EcoSystem layer is populated, then the layer above it is populated according to the population of the previous EcoSystem layer.

Adding a Layer

You don't have to select a multi-material to add a layer to a material: simply select the line where you want to add the layer in the [Material Hierarchy](#) and press the **Add layer** button to the right of the hierarchy. The Material Browser appears, letting you select the new layer to add to the material. The layer is added immediately above the layer that was previously selected.

If you want to create a new layer without loading a preset material, just click **Cancel** in the Material Browser. A new “empty” layer will be added.

You can delete a layer by selecting it in the [Material Hierarchy](#) and pressing the **Del layer** button to the right of the hierarchy.

Changing the Order of Evaluation

You can change the order in which layers are evaluated by selecting a layer and pressing the **Up** and **Down** buttons to the right of the [Material Hierarchy](#), or by dragging it in the list with the mouse. If you move a layer up, it will be evaluated earlier on the stack, appearing “on top” of other layers. Using the left mouse button, this also affects the environment settings.

If you use the right mouse button to reposition the layers, the environment settings are not affected. Any distribution mapping you have done remains the same; only the materials are swapped.



Influence of Layers

Layers can be placed according to their environment, using the **Environment** tab of the *Material Editor* when the layer is selected. Using this tab, you can constrain the layer to appear only at given altitudes, on given slopes or at given orientations. Please turn [here](#) and [here](#) for details on the **Environment** tab of the Simple and EcoSystem materials.

The placement of the bottom-most layer on the layer stack cannot be influenced by environment (it has to be “everywhere”, since). Additionally, EcoSystem layers can interact with other EcoSystem layers beneath them through the affinity/repulsion settings (see [here](#)).

Material Snapshots


You can store a snapshot of a particular layer, by selecting the layer and clicking on the **Material Snapshot** icon, located to the far right above the last column. If you left-click another material or layer, it’s snapshot goes in the lower box; the next layer you select will go in the lower box and the previous material stored is moved to the upper box. Now, if you select a material and right-click you can select which box that snapshot will be stored in.

The Shared Material Layer

This feature is a way to allow you to create a layer in a material that can be shared by other materials in your scene. A modification on a shared layer in a material will modify this layer in any other materials it is used in. You can have many shared material layers.


To make a shared layer, right-click on the icon of the layer and select **Share layer**.

In order to add a shared layer onto an existing material, you can use the **Add shared layer** command from the contextual menu. This option is only available if you have previously created shared layer(s) in the scene. This operation is also accessible from any material line of the *World Browser*, from the *Summary of Materials* and from the selected object’s material preview in the *Object Properties* panel.

In the *Material Editor* (Advanced or Basic), you also have an **Add Shared Layer** icon () right above the **Add Layer** icon to perform this operation.

In order to “dissociate” a material layer previously added through the **Add shared layer** command, you have a **Make unique** command in the material menu (only in the *Material Editor*). By making a material layer unique, you dissociate it from the other materials previously sharing it. It means that further edition of this layer will not affect other materials anymore.

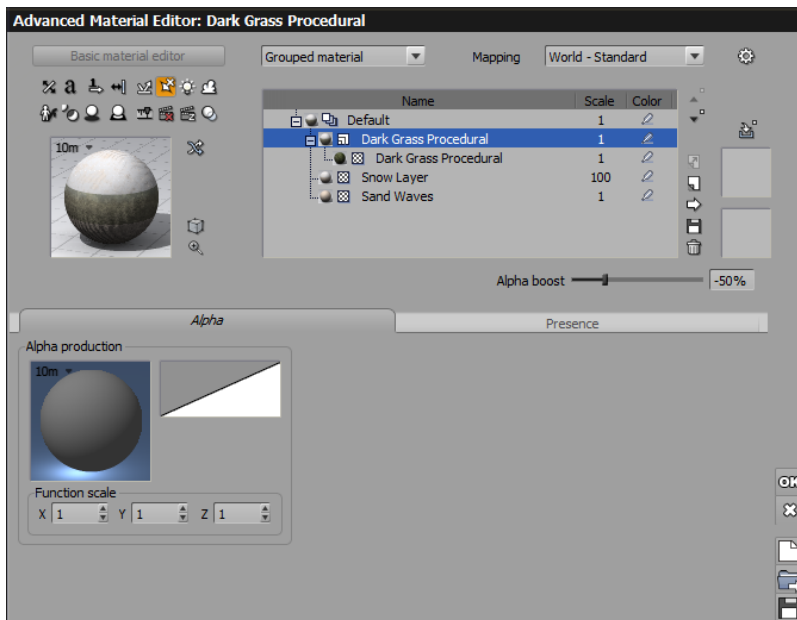


There is now a new pictogram to identify shared layers easily ()

Note:

Only Simple materials can be shared (neither Volumetric, Mixed nor EcoSystem/Particle materials).

Grouped Materials



Grouped Materials

Unlike Layered Materials, you can control multiple layer presence driven by one item in the materials list. Alpha is also controlled for all layers. You can, for example, control both an EcoSystem and it's underlying layer presence in one place with such material type.



Alpha Tab

Alpha production can be edited on this tab. You can load a new filter by right-clicking the displayed filter and selecting **Load Filter** to access the Filter Browser. Or, you can edit the current filter by right-clicking the displayed filter and selecting **Edit Filter** to open the *Material Alpha Filter* dialog.

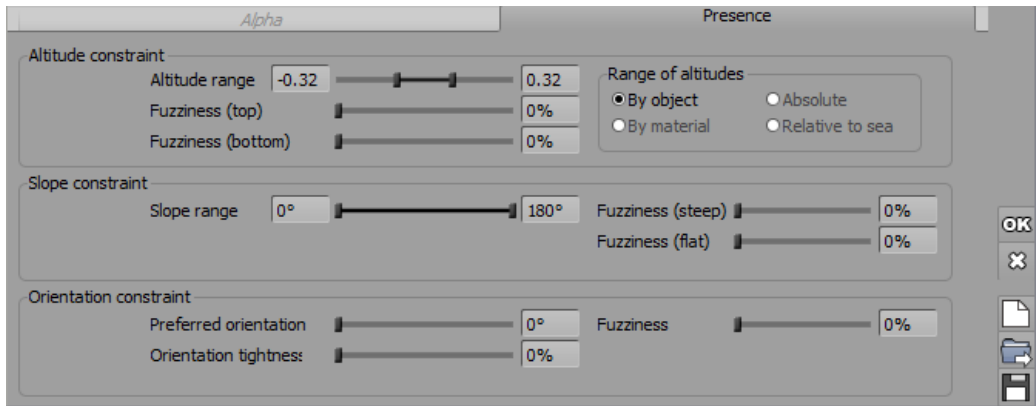
You can also modify the alpha channel by using the *Function Graph*. Access the *Function Graph* by right-clicking the **Alpha production** picture and selecting **Edit Function**. You can also select to load a function from the Function Browser by selecting **Load Function** from the menu.

Presence Tab

This group lets you control how altitude influences the presence of the layer:

- **Altitude range:** this dual slider lets you define the range of altitudes in which the current layer appears (provided it is not transparent at this point). The layer will not appear outside this range.
- **Fuzziness:** this setting controls how “suddenly” the changes to the layer presence are made in response to altitude. High values mean that the layer appears very gradually in its altitude range, whereas low values will result in the layer appearing as a solid strip.
- **Range of altitudes:** this lets you define in what coordinates the altitude range is defined:
 - **By object:** in this mode, the range is relative to each object to which the material is applied.
 - **By material:** in this mode, the range is relative to all the objects that use this material.
 - **Absolute:** in this mode, the range of altitudes is expressed in global coordinates.
 - **Relative to sea:** the altitude is computed from the sea level and not from zero.





Presence tab-Group Material

Slope Constraint

This group lets you control how the local slope influences the presence of the layer:

- **Slope range:** this dual slider lets you define the range of slopes in which the current layer appears (provided it is not transparent at this point). The layer will not appear outside this range. Values to the right end of the slider indicate flat surfaces, and values to the left indicate upside-down surfaces. Intermediate values indicate vertical surfaces. Slope values can range from -180 to +180 degrees.
- **Fuzziness (steep):** this setting controls how “suddenly” the changes to the layer presence are made in response to slope. High values mean that the layer appears very gradually in its slope range, whereas low values will result in the layer appearing as a solid strip on areas of appropriate slope.
- **Fuzziness (flat):** this setting controls how “suddenly” the changes to the layer presence are made in response to flat areas.

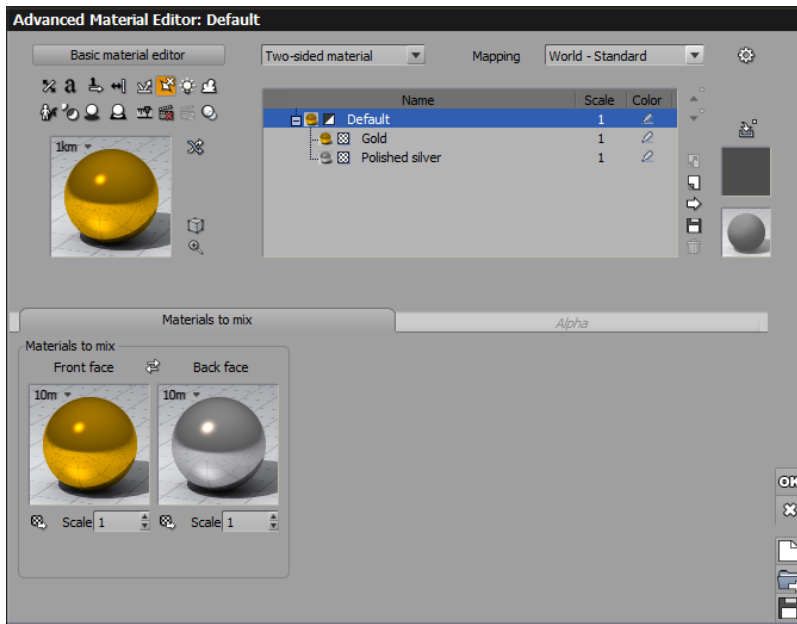
Influence of Orientation

This group lets you control how the local orientation influences the presence of the layer:

- **Preferred orientation:** this setting controls the orientation of the surface that is the most favorable to the presence of the layer.
- **Orientation influence:** this setting controls the influence of orientation on the presence of the layer.
- **Fuzziness:** this setting controls how “suddenly” the changes to the layer presence are made in response to orientation. High values mean that the layer appears very gradually on surfaces of the preferred orientation, whereas low values will result in the layer appearing as a solid strip on areas of preferred orientation.



Two-Sided Material



Material Editor – Two-Sided Materials

Two-sided materials allow you to define two different materials to each facet of a two sided object like a leaf. Only a few object types are supported: TPF plants, planes, alpha planes, terrains and meshes (imported or baked objects).

Materials to Mix Tab


This tab is used to select the materials to use.

Front Face – Back Face: Select the two colors to use, one for each face. You can switch the colors from front to back by clicking on the double-arrow icon.

Each of these materials can be edited by clicking on the material. You can mix together simple materials, materials that are themselves a mix of other materials (create nested material hierarchies for amazing effects!), or even EcoSystem materials.



Alpha Tab



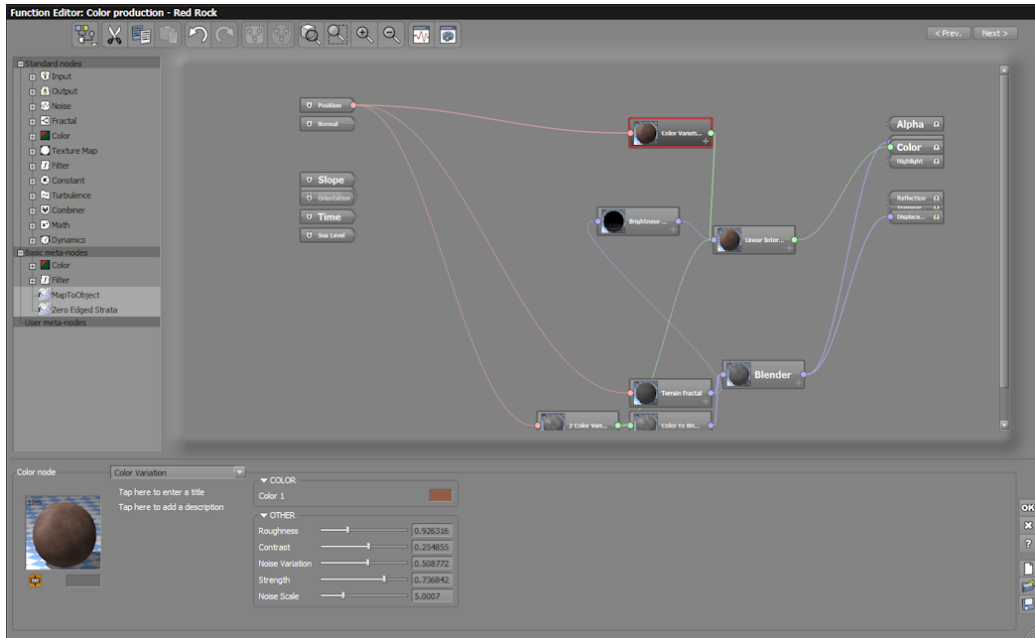
Material Editor – Two-Sided Material Alpha

Alpha production can be edited on this tab. You can load a new filter by right-clicking the displayed filter and selecting **Load Filter** to access the Filter Browser. Or, you can edit the current filter by right-clicking the displayed filter and selecting **Edit Filter** to open the *Material Alpha Filter* dialog.

You can also modify the alpha channel by using the *Function Graph*. Access the *Function Graph* by right-clicking the **Alpha production** picture and selecting **Edit Function**. You can also select to load a function from the Function Browser by selecting **Load Function** from the menu.



Function Graph



Function Graph

Description

Functions are the key to the visual quality of materials.

They are used every time it is necessary to generate a value depending on position (e.g. to indicate a transparency amount that depends on the position inside the material).

Basically, functions enable you to associate any point in space with a value in the range of -1 to 1.

The *Function Graph* is accessed by clicking on the preview picture of a function with the **Control** key pressed, or by selecting **Edit function** from the contextual menu. The editor can stay open without stopping other parts of the software from being accessible.



The Basics

What Is a Graph?

A graph displays a set of interconnected nodes that are used to generate output values based on the values of a given set of inputs.

Input and Output Nodes

The graph inputs sit on the left side of the graph. The graph outputs are placed at the right side of the graph. Input nodes are the points where data enters the graph, and output nodes are the points where the data exits the graph. Output nodes represent the value that is computed by the function.


Data enters the graph at the input nodes, flows through the different nodes and links in the graph, and exits at the output node. You cannot delete input or output nodes, and you cannot place other nodes to the left of the input nodes or to the right of the output nodes.

The default types of input are:

- **Position:** this input node produces a vector value representing the position of the point where the function is being evaluated. Obviously, the value of this input node depends on the mapping mode selected for the object's material.
- **Normal:** this input node produces a vector value representing the direction in which is pointing the surface at the point of evaluation of the function.
- **Altitude:** this input node produces a signal whose value is proportional to the altitude of the point where the function is being evaluated. The value of the input depends on the mapping mode, and can vary beyond the range -1 through 1.
- **Slope:** this input node produces a signal whose value is proportional to the local slope at the point where the function is being evaluated. If the surface is horizontal, the value of the input will be 1. If the surface is vertical, it will be 0. And if the surface is horizontal, only facing down, the value will be -1.
- **Orientation:** this input node produces a signal whose value varies between -1 and 1 according to the azimuth of direction in which is pointing the surface at the point where the function is being evaluated. If the surface is pointing up along the Y axis, the input value will be 0. The signal jumps from -1 to 1 as the normal turns from



south-west to south-east.

You can create additional input nodes by clicking the **Input Node** icon () in the left toolbar of the *Function Graph*. Please turn [here](#) for details on the different additional input nodes that can be created.

When editing an *Object Graph*, the input and output types depend on the type of object being edited. For instance, Light objects will expose a **Light Color** input and output. Please turn [here](#) for further details on object graphs.



Output Data

The name and type of output nodes depend on what values are expected from the function (i.e. what the purpose of the function is). For instance, if you are editing the terrain altitude function of a procedural terrain, the output will be labeled “Altitude”, and the value generated by the function will be used to generate the altitudes of the procedural terrain.

Usually, the type of data output by a function is a number (a floating point value), but there are some cases where functions can also export colors (e.g. when editing the color function of a procedural material). If you try to use a function that outputs a color where a number is expected, *SmartGraph™* (see below for details) will automatically add a node to convert the color into a brightness value.

Multiple and Master Outputs




In some cases, functions can output multiple channels of data. This is for instance the case when you edit a function from a simple procedural material. In that case, the *Function Graph* will display output nodes for all the different channels in the material (color, bump, transparency, etc.). You can reuse portions of the graph to generate outputs for several channels simultaneously instead of having to duplicate portions of the graph (e.g. you could plug the color and bump outputs to the same node).

You can create additional output nodes by clicking the **Output Node** icon () in the left toolbar of the *Function Graph*. The number and nature of additional output nodes that can be created depends on the context of the function. For instance, settings in the *Material Editor* that are “exportable” (they are identified by the  icon – see [here](#)) will have a corresponding optional output node in the corresponding *Function Graph*.

If there are several output nodes in the *Function Graph*, there will always be one of the outputs that is known as the master output. There can only be one master output in a



function graph. The master output is the output that corresponds to the channel from where the *Function Graph* was accessed. For instance, if you entered the *Function Graph* by editing the color channel production function, the master output will be the color output. But if you enter the *Function Graph* by editing the bump production function, then the master output will be the bump output instead of the color output.

The master output is displayed with a stronger contrast than the other output nodes – even when it is not selected (or a node connected to this output is not selected). If you press the **Save** icon () in the dialog bar, only the part of the graph that is connected to the master output will be saved. Saved functions will appear in the Visual Function Browser like any other of the predefined functions. By default, functions are placed in the **Functions** subfolder. This means that they will appear in the **Personal** collection inside the Visual Function Browser. In the same way, if you load a new function into the *Function Graph* using the **Load** icon () , or reset it by pressing the **New** icon () , only the part of the graph connected to the master output will be replaced (or removed).

Nodes

Nodes are represented by little boxes on the graph. A node receives a flow of data on its entries, affects a certain processing on that data according to its type and the values of optional parameters, and generates one or several flows of outgoing data. This outgoing data can be of the same type as the incoming data, or it can be of a different type.

There are 10 different categories of nodes: noise nodes, fractal nodes, color nodes, texture map nodes, filter nodes, constant nodes, turbulence nodes, combiner nodes, math nodes and dynamics nodes. For your convenience, the category of a node is identified by the shape of the box that represents the node on the graph.

Nodes can appear on the graph in two different sizes:

- a small version with a symbol that identifies the node category (on top of the shape of the box), or
- a large version with a preview of the node in it (the shape of the box still indicates the node category).

The size of the nodes is defined automatically according to the settings in the **Graph Options** menu (see [here](#)), but you can force any of the two versions using the **Show in graph** option in the node details section of the *Function Graph* (see [here](#)). That way you can make important nodes stand out by making them larger than the other ones.

It is possible to change the display color of any node (or groups of nodes) on the graph to improve readability. Just select the node; under the preview of the node there is a color box. Click on it to open the *Color Selection* and select a display color for that node.



Types of Data

The nodes in the *Function Graph* can process 4 different types of data:

- **Number:** this is a floating point value. It is the typical output of a function graph. Noise nodes and Fractal nodes (among others) produce numbers.
- **Color:** this is the typical output of the color nodes. If you are editing the Color channel of a material, the function may either output a number (in which case the number will be converted into a color outside of the function using a color map), or directly output a color.
- **Texture Coordinates:** this is a two-dimensional vector that typically indicates the texture coordinates of the point where the function is being evaluated. This is the typical output from the *Projection* node.
- **Vector:** this is a set of 3 numbers that indicate a position or a direction in space. Typically, the position and normal inputs are both vectors, where position indicates the position of the point where the function is being evaluated (converted into the appropriate coordinate system depending on the selected mapping mode), and where normal is the direction in which the surface of the object is pointing, at the point where the function is being evaluated.

Links

Links are the lines that connect different nodes together. Links represent the flow of data through the graph. The data always flows down, from top (inputs) to bottom (outputs). If a node is higher than another one, you know it is being processed before.

The color of the link indicates the type of data that is being transported by the link:

- **Blue link:** number (e.g. noise output),
- **Green link:** color information,
- **Purple link:** texture coordinates,
- **Red link:** vector data (e.g. position),
- **Gray link:** undefined data type.

When a link is selected (e.g. by clicking on it), it is drawn with a thicker line, and the two nodes that the link connects are displayed in the link properties part of the *Function Graph*.



SmartGraph™

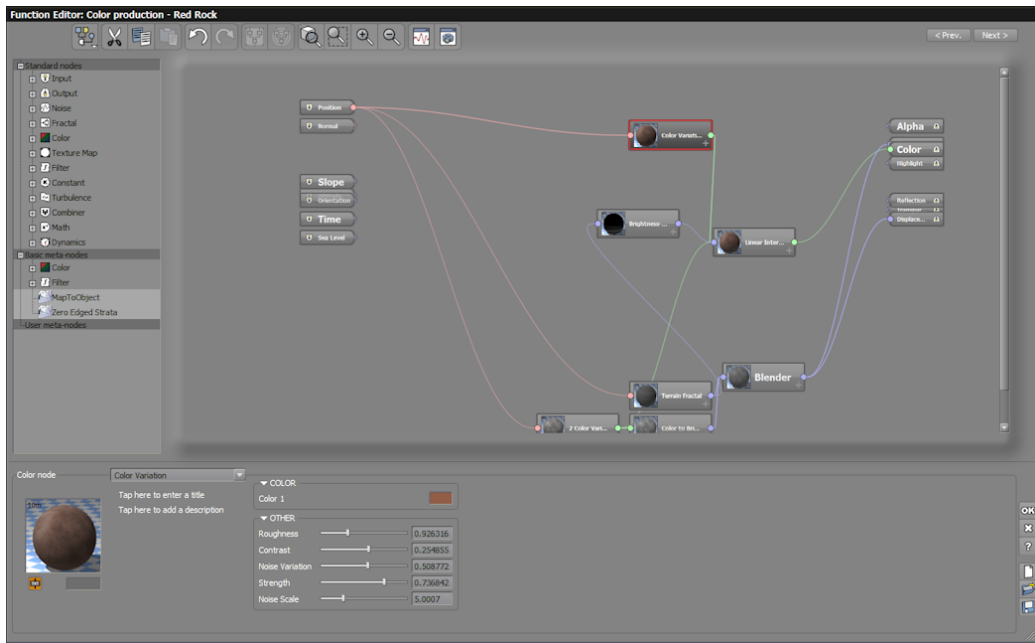
SmartGraph is a collection of unique and extremely clever technologies inspired by techniques used in artificial intelligence systems. The sole purpose of *SmartGraph* is to make the creation and edition of VUE function graphs both easy and fun.

As you add, replace or remove nodes or links, *SmartGraph* determines what can be done to simplify your task. For instance, if you delete a node, *SmartGraph* will attempt to reconnect the link that was broken. When you add a node from a given category behind another one, *SmartGraph* will try to find a type of node in that category that is compatible with the data generated by the previous node – and if the previous node was connected to another node, *SmartGraph* will look for a type of node that can be placed in between the two (compatible data types on entry and exit).

Presentation

The *Function Graph* is constituted of the following elements:





Function Graph

- Top Toolbar,
- Nodes Toolbar (the vertical toolbar),
- Function Graph (the main display area), and
- Node/Link Details.

Top Toolbar

Node and Function Previews

Node and function previews are used to visualize the results of the function at a given node. The function is represented by a black and white object (sphere, cube, cone...) with the value of the function indicated on its surface. If the node returns -1 at a point, this point will be black; if it returns 1, the point will be white.

If the node or the function returns a color, that color will be displayed at the surface of the preview object.

You can change the shape of the object used to preview the function with the *Preview Options* dialog (select the **Preview Options** command from the **Graph Options** icon



menu – see below). This is the same as the *Preview Options* dialog from the *Material Editor* (see [here](#) for details).

Toolbar Icons



Graph Options Menu: the *Function Graph* options menu is accessed by clicking on this icon in the toolbar. Options in this menu are:

Note:

The first four options on this menu are available only with the pre-2014 Function Graph format.

- **Show All Previews in Graph:** shows all node previews.
- **Preview Only Noises and Colors:** only shows previews for noise and color nodes.
- **Hide All Previews:** no previews are shown, only node boxes.
- **Show Arrows on Links:** when selected, shows arrows on the links to indicate data flow.
- **Show Links as StraightLines:** links can be shown as straight lines. Lines with right angles are the default.

Note:

These options are available for both the old and new Function Graph formats.

- **Enable Auto Graph Layout:** aligns nodes in graph automatically.
- **Keep Input and Output Nodes Aligned:** the input nodes are aligned in a column on the left; the output nodes are aligned in a column on the right of the *Function Graph*.
- **Constrain nodes to Grid:** enforces automatic spacing for nodes.
- **Automatically Map Connected Parameters:**
- **Show Real-World Units:** real-world units are shown in the *Function Graph*.
- **Preview Options:** select this option to display the *Preview Options* dialog (see above) and change the look of the node previews. These changes are global to the entire graph.



Show Parameter Connections: if you select this option, lines will appear on the sides of the nodes in the graph. These lines on the side correspond to parameters that can be connected to other nodes. You can grab the lines and drag them onto



another node in order to create a connection.

Note:

You do not have to show the parameter connections in order to establish this type of connection. Please read below for details on establishing connections between nodes.



Cut: this command is available when at least one node is selected. Press the Cut icon to cut the selected nodes out of the graph and onto the clipboard.



Copy: this command is available when at least one node is selected. Press the Copy icon to copy the selected nodes onto the clipboard.



Paste: this command is available when at least one node is selected. Press the Paste icon to paste the nodes in the clipboard onto the graph.



Undo: click this icon to Undo the last operation. You can undo multiple changes. When you undo an operation, the **Redo** icon becomes available.



Redo: click this icon to Redo the last operation that was undone. If you have undone multiple operations, you can redo them all (unless you make a change).



Create MetaNode from selection: this option is only available when you have selected several nodes. Clicking this icon will convert the selected nodes into a MetaNode. Connections are automatically preserved.



Ungroup MetaNode: when a MetaNode is selected, clicking this icon will place the content of the MetaNode back in the graph.



Frame: click this icon to automatically adjust the framing of the preview so that



graph is centered and all the nodes are visible.



Frame Selected Area: click this icon to automatically adjust the framing of the preview so that the selected nodes are centered in the graph view.



Zoom In: click this icon to display a magnified view of the graph in the preview, thus letting you observe the nodes in finer detail. The current zoom factor is displayed beneath the icon.



Zoom Out: click this icon to display a reduced view of the graph and get a more global view of it.



Function Node Preview: click this icon to open (or close) the *Function Node Preview* panel. This panel displays a detailed view of the output of the currently selected node along 3 axes. This is a toggle icon. If you close the *Function Graph* with this panel open, it will automatically open next time you open the *Function Graph*. Please turn [here](#) for details on the *Function Node Preview* panel.



Function Output Observer: click this icon to open (or close) the *Function Output Observer* panel. This panel displays a contextual preview of the function's output. If you are editing a material function, a preview of the material will be displayed. If you are editing the altitude function of a procedural terrain, it will display a terrain. This is a toggle icon. If you close the *Function Graph* with this panel open, it will automatically open next time you open the *Function Graph*. Please turn [here](#) for details on the *Function Output Observer* panel.



Nodes Toolbar

The nodes toolbar is the vertical toolbar on the left of the *Function Graph*. The nodes toolbar is used to add or replace nodes in the function graph. In order to make the creation of elaborate functions an easier process, the icons in the nodes toolbar will either add new nodes, or replace existing nodes, depending on the context. This unique e-on software technology, inspired by artificial intelligence systems, is known as *SmartGraph™*.

Here is the description of the different icons in the Nodes icon bar:



Input Node: click this icon to create a new input node. The Input nodes appear in a column on the left in the graph area.

When editing an **Object Graph**, the list of available input nodes depends on the type of object being edited. Most objects will have a **Position**, **Size** and **Orientation** input node, but other properties may also be available as inputs (e.g. **Light Color** when editing a light). You can also create **External Dependency** and **Recall Dependency** nodes that will let you import values from other graphs.

When editing material functions, the list of additional input nodes that can be added to your function graph is as follows (depending on context, some of these input nodes may not be present, or may not generate intelligent values – for instance, when editing the altitude function of a procedural terrain, the “Incident Light Angle” input will always return 0).

- **Position:** this is the ‘Position’ input.
- **Position Options:** this is an advanced version of the position node above that lets you select the coordinate system in which the position vector is expressed. This input node displays a drop-down list that lets you select the coordinate system for this instance of the node (you can create several “Position Options” nodes with different coordinate systems). The different coordinate systems are the same as the **coordinate systems** available in the *Material Editor’s Mapping* list. You can specify the **Distance unit** in VUE units, Display units or any other unit available in a drop-list for selection.
- **Normal:** this is the ‘Normal’ input node.
- **Normal Options:** again, this is an advanced version of the “Normal” node described above. This input node lets you define the coordinate system in which the normal vector is to be expressed. The node displays a drop-down list that lets you select the coordinate system for this instance of the node (you can create several “Normal Options” nodes with different coordinate systems). If you select the **Object** option, the normal vector is expressed in object coordinates, and hence is independent on the orientation of the object. On the other



hand, if you select the **World** option, the normal vector will be modified by the orientation of the object.

- **Slope:** this is the ‘Slope’ input.
- **Altitude:** this is the ‘Altitude’ input.
- **Orientation:** this is the ‘Orientation’ input.
- **UV Coordinates:** this input node returns the texture coordinates of the current point, according to the selected texture mapping mode. It is created automatically when you create a [Texture Map node](#). This node is typically used to drive the mapping of a texture. Please see [the options of the UV Coordinates node](#) for further details.
- **Time:** this input node returns the current time in seconds. This is used for animated functions. If you create a time-dependent node (such as the [Open Ocean](#) node in the Fractal category), a connection to this input node will automatically be created.
- **Angle of Incidence:** this input node returns the angle of incidence between the incoming ray (the ray coming from the camera, the reflected ray coming from a reflective surface, etc) and the surface of the object. If the surface of the object is exactly facing the ray (the incoming ray is perpendicular to the object surface), the input node returns a value of 1. At low incidence angles (when the ray is tangent to the surface), the return value is 0. If the ray hits the surface from the inside, the return value will be negative.
- **Ray Direction World:** this input node returns a vector that indicates the direction of the incoming ray expressed in world coordinates.
- **Ray Direction Object:** this is the same as the Ray Direction World input described above, except that the incoming ray direction vector is expressed in object coordinates instead of world coordinates. This is useful if you want to create a function that depends on the direction of the incoming ray, but is not affected by the orientation of the supporting object.
- **Position On Picture:** this input node returns a vector representing the position of the point in the final picture. The X and Y components of this vector respectively indicate the horizontal and vertical position of the point in the picture, where -1 is the left/top edge and +1 is the right/bottom edge of the picture. The Z component of the vector is always 0.
- **Distance to Camera:** this input node returns the distance between the point where the function is being evaluated and the camera, whatever the ray recursion depth. The distance to the camera is at most equal to the actual distance traveled along the ray. You can specify the **Distance unit** in VUE units, Display units or any other unit available in a drop-list for selection.



- **Distance on Ray:** this input node returns the total distance traveled by the ray from it's origin, including all recursions. This means that if the point where the function is being evaluated was hit by a reflected ray, the distance on ray would include the distance traveled by the reflected ray plus the distance traveled from the ray's origin to the point where it was reflected. You can specify the **Distance unit** in VUE units, Display units or any other unit available in a drop-list for selection.
- **Distance to Object Center:** this input node returns the distance between the point where the function is being evaluated and the actual center of the object that was hit by the ray. This would yield a constant value on a sphere, since, by definition, all points of the sphere's surface are at the same distance from the sphere's center. This is however not true for other geometries, and can also be useful when evaluating volumetric materials. You can specify the **Distance unit** in VUE units, Display units or any other unit available in a drop-list for selection.
- **Distance to Object Below:** this input node traces a ray downwards from the point where the function is being evaluated and returns the distance to the first object encountered. This could be useful for instance to evaluate the depth of an ocean and create foam (or waves) near the shore. Warning: because this input requires the actual tracing of a ray, it is very slow to process. You can specify the **Distance unit** in VUE units, Display units or any other unit available in a drop-list for selection.
- **Distance to Surface:** this input node returns the distance to the surface of the object. It is only defined for primitives and Metablob objects and is only really useful when computing volumetric effects (because it returns the distance to the surface, this input node will always return 0 when evaluating a standard material). You can specify the **Distance unit** in VUE units, Display units or any other unit available in a drop-list for selection.
- **Object Center:** this input node returns the coordinates of the center of the object that was hit by the ray. Obviously, this value is constant over the entire surface of the object, but it can be particularly useful to switch textures in an EcoSystem population (see the tutorial on [Varying Materials on EcoSystem Populations](#) for an example use).
- **Incident Light Angle:** this input node returns the angle of incidence between incoming light rays and the surface of the object. If the light hits the object's surface perpendicularly, the value returned by the input node is 1. If the light hits the object's surface at a very low angle of incidence, the value will be close to 0. This is useful e.g. if you want to create a custom BRDF (Bidirectional Reflectance Distribution Function) for your materials. Obviously, the value returned by this input node is usually different for each light source. This input is only valid when evaluating the specular contribution.



- **Specular Light Angle:** this input node returns the dot product between the direction of incident light and the direction of reflection of the viewing ray. This is useful for create custom specular reflection effects. Obviously, the value returned by this input node is usually different for each light source.
- **Light Direction:** this input node returns the direction of the incoming light expressed in world coordinates. Obviously, the value returned by this input node is usually different for each light source.
- **Light Color:** this input node returns the color and intensity of the incoming light. Obviously, the value returned by this input node is usually different for each light source.
- **Reflected Direction:** this input node returns the direction of reflection of the viewing ray, expressed in world coordinates.
- **Anisotropic Direction:** this input node returns the direction of the “scratches” used to compute anisotropic reflections, expressed in world coordinates.
- **Transformed Z Vector:** this input node returns the direction of the upwards vector transformed by the object’s transformation matrix, expressed in world coordinates. This basically tells you which direction is “up” on an object, taking into account the rotation of the object.

Note:

The resulting vector is not necessarily of length 1.

- **Transformed Left Vector:** this input node returns the direction of the left vector transformed by the object’s transformation matrix, expressed in world coordinates. This basically tells you which direction is “up” on an object, taking into account the rotation of the object. Please note that the resulting vector is not necessarily of length 1. This, together with the “Transformed Z Vector” provides all information about the object’s transformation matrix.
- **Sea level:** this input node returns the value of the current sea level. The value is set either in the *Options* panel, **Units & Coordinates** tab, or by moving the sea level directly in your scene.
- **Object specific (formerly known as *Impact specific data*):** This function is used when building materials for VUE rocks using local convexity information. It provides a measure, at any given point of the rock, of how much the rock is protruding at that point (or caving in). This information allows realistic customization of the rock material, taking into consideration that convex areas of a rock are more exposed to interaction with the environment (typically weathering) than the concave ones, and their aspect will therefore evolve very differently. There is one parameter, a scalar in $[0; 1]$ – Account for larger features. This controls whether the convexity information returned should focus on small scale details or on larger scale features. In other words, with a value



of 0.0, the node will return the local convexity at a small scale which helps identify small “pointy” bits on the rock. On the other hand, with a value of 1.0, the node will give information on the overall shape of the rock. When this node is called on something which is not a rock, the default value of 0 is returned.

The following input nodes are all EcoParticle system related. Functions related to collisions can be used for any object material, whereas other nodes can only be used for EcoParticle material.

- **Particles Collision Density:** this is a way to get data from the collision grid. You can either get a collision count or an energy clamped to a maximum. You have to select the EcoParticle system you want to sample from the node.
- **Atmospheric Particle Collision Density:** this is the same as **Particle Collision Density**, except it is for rain and snow.
- **Particle Speed (m.s – 1):** the velocity of the EcoParticle in meter per second.
- **Particle Mass (kg):** weight of the EcoParticle
- **Particle Adhesion Coeff [0; 1]:** how the EcoParticle will stick when it slides on a surface.
- **Particle elasticity [0; 1]:** how the EcoParticle will bounce when it collides.
- **Particle Age (s):** how long the EcoParticle has lived.
- **Particle Last Collision Time (s):** Particle Collision Count [0; ∞[
- **how much the EcoParticle has collided.:** Particle Radius at Birth (m)
- **the size of the EcoParticle at birth in meters:** Particle Last Collision Normal
- **The normal to surface the EcoParticle last collided (if any).:** Particle Drag Coefficient [0; ∞[
- **it is used to quantify the drag or resistance of an object in a fluid environment such as air or water. A lower drag coefficient indicates the object will have more aerodynamic or less drag. The drag coefficient is always associated with a particular surface area.:** Particle Birth Time (s)
- **when the EcoParticle was born in seconds. Time here is the same as in the *Timeline*.:** Particle Death Time (s)
- **the age of the EcoParticle when it will die. Time here is the same as in the *Timeline*.:** Particle Life Duration (s)
- **How long the EcoParticle lived (in seconds).:** Sub-particle Normalized Rank[0; 1[



- this is only useful for a child **EcoParticle** (emitted from another **EcoParticle**). It will be the normalized rank of the emission. For example a mother **EcoParticle** will emit 3 child **EcoParticles**. The first child will have 0 in this node (0/3), the second 0.33333... (1/3), and the third will have 0.66666... (2/3):: Particle Count
- the total **EcoParticle** count of the **EcoParticle** system in which the **EcoParticle** exists.:



Output Node: click this icon to create a new output node. The Output nodes appear in a column on the right in the graph area. The choice of possible output nodes depends on what the function is being used for. For instance, if it is being used to compute the altitudes of a procedural terrain, this will be the only output node. But if the function is used for a material, you can create new outputs for other material parameters. If the function is used as an *Object Graph*, there will be output nodes for all of the object's properties. Also, the **Custom Dependency** node lets you output any kind of data for retrieval in another graph.



Noise Node: click this icon to create a noise node. A noise node outputs a number between -1 and 1. If a fractal node is selected, it will be converted to a noise node of the same base noise as the fractal. See the [different types of noise nodes](#) for details.



Fractal Node: click this icon to create a fractal node. A fractal node is based on a noise that is repeated at several different frequencies in order to create much more elaborate patterns as the standard noise node. Fractal nodes create patterns that exhibit details over a large range of frequencies. If a noise node is selected at the time of clicking this icon, it will be replaced by a Simple Fractal node based on the same noise as the noise node. Please see [the different types of fractal nodes](#) for further details.



Color Node: click this icon to create a color node. Depending on the context, color nodes either output a color based on the value of a number, or converts a color into another color. If a node is selected at the time of clicking this icon, again depending on context, a color node of the appropriate type will usually be added behind the selected node. Please see [the different types of color nodes](#) for further details.





Texture Map Node: click this icon to create a texture map node. Texture map nodes are used to map pictures (texture maps) onto objects. The texture map node is also created together with a Projection input node. The projection input node converts the current position into mapping coordinates used by the texture map node to map the texture. Please see [texture map node](#) for further details.



Filter Node: click this icon to create a filter node. Filter nodes take a signal as input and output another signal. Clicking repeatedly on the Filter node icon will add as many filter nodes. Please see the [different types of filter nodes](#) for further details.



Constant Node: clicking on this icon will create a constant node. If another node was selected at the time of clicking, the selected node will be replaced by a constant node of the appropriate type. Please see [different types of constant nodes](#) for further details.



Turbulence Node: clicking on this icon will create a turbulence node. Turbulence nodes take a vector as input, and return a vector. They are usually plugged into the **Origin** noise parameter, as this is where they will behave as actual turbulence. Please see the [different types of turbulence nodes](#) for further details.



Combiner Node: click this icon to create a combiner node. Combiner nodes are used to combine together different values. Most of them work on all types of data, and output the same type of data as the one provided in input. Please see the [different types of combiner nodes](#) for further details.



Math Node: click this icon to create a math node. Math nodes are used to perform all sorts of operations and conversions between different data types. Please see the [different types of math nodes](#) for further details.



Dynamics Node: click this icon to create a dynamics node. Dynamics nodes are





mostly used to create dynamic connections between object properties. Please see the [different types of dynamics nodes](#) for further details.



Load MetaNode: click this icon to load a MetaNode from disk. A Standard File Browser will appear, letting you select the MetaNode you want to load. You can also create MetaNodes of a specific type using the popup menu. Please see the [MetaNodes](#) for details.

Function Graph

The function graph is the large area that sits in the middle of the *Function Graph* dialog. It is used to assemble the different nodes and links that will constitute the function.

You can zoom in and out of the function graph using the **Zoom** icons ( and ) , or by using the *Zoom* trigger defined in the [shortcuts list](#) of the *Operations (Documentation/Getting Started/Options and Preferences)* panel. There are limits to the amount of zooming in and out that you can do.

You can move the graph around by pressing space and dragging it, using the scrollbars, or using the *Pan* trigger defined in the [shortcuts list](#) of the *Operations* panel.

The *Pan* and *Zoom* triggers are the ones in the *Triggers in 1D & 2D views* category, if they are defined, or in the [top-most](#) category otherwise.

Node Selection

You can select nodes or links by clicking on them. Selected nodes appear with a red frame, while selected links appear bold. If you want to select multiple nodes, click on all the nodes you want to select with the Shift key pressed, or drag a marquee rectangle around the area that you want to select (all nodes in that rectangle will be selected).

You can also select an empty spot on the graph by clicking on it. A red square will appear around the selected area (clicking one of the node icons will place a node in that square).

In order to facilitate the understanding of the way data is processed in a particular function, whenever a node is selected, all the other nodes that are connected to it will appear with a drop-shadow.

Also, if the function has more than one output, the Master output together with all connected nodes will be displayed in solid text, whereas other nodes will appear in pale color.



Adding/Replacing Nodes

There are three ways of adding/replacing nodes:

- Click on an empty area of the graph; this will select a square in the graph. Click on one of the node icons, or pick the desired node from the popup menu. The new node will be placed at the selected location.
- Select a node and click on one of the node icons. Depending on which type of node was selected, and which icon was pressed, the selected node will either be replaced with a node of the new type, or a new node will be added beneath the selected node (e.g. pressing **Fractal** when a noise node is selected will replace the noise node by a fractal node of the same noise, but pressing **Filter** when a noise node is selected will connect a filter node on the noise output). If the selected node was already connected to another node, *SmartGraph* (see [here](#)) will attempt to insert a node of a compatible type from the desired category. If no such node can be found, a message will appear asking you if you want to break the connection.
- Select a link and click on one of the node icons or pick a node from the popup menu. The new node will be inserted in the link (the output of the first node will be connected to the entry of the new node, and the output of the new node will be connected to the second node on the link). If the node you are creating cannot be inserted in the link, nothing will happen. If you picked a node from the menu, the node will be created anyway – but it won't be placed on the link.

To delete a node, simply select it and press the **Delete** key. If possible, *SmartGraph* will reconnect the open link.

Connecting Nodes

When you move the mouse on top of a link, a little round handle will appear at each end of the link. Grab one end of the link by clicking on the appropriate handle, and drag it onto another node. As soon as you begin dragging an end of the link, all compatible nodes will be marked by a little circle. Drop the link on one of the compatible nodes, and the new connection will be established automatically. To cancel the operation, simply drop the handle on an empty part of the graph.

Under certain conditions, a small cross may appear at the center of the little round handles. This is a hint to help identify connections that are probably not appropriate in the current context. You may however still establish the link if you are sure that the connection is pertinent (a warning will appear when you establish such a connection).

To delete a link altogether, select the link and press **Delete**, or grab the link by its middle (not by one of its handles), and drop it outside of the graph.

Note:

It is possible to change the color display of every node in the *Function Graph* to improve



visibility of the graph. When a node is selected, under the preview of the node there is a color box. Click on that to select a display color for that node.

Published Parameters

The **Published Parameters** feature copies specific settings from the *Function Graph* that you may need to change often and places them in a more convenient location for easier access. If you have accessed the *Function Graph* from the *Terrain Editor*, the selected parameters will appear on a special tab in the *Terrain Editor*. If you accessed the *Function Graph* from the *Materials Editor*, a special tab will appear in the *Advanced Material Editor*.

To select a parameter for publishing, just click the underlined field name of the parameter. A parameter name is supplied and a group name is asked to improve the display of the published parameter.

Node/Link Details

The node/link details area is the area of the *Function Graph* that sits below the function graph. As its name indicates, this area displays details on the currently selected node or link.

Node Details

When a node is selected, the category of the noise appears as the title of the details area. Alongside this title, one or several drop-down list boxes will let you change the type of node in that category (e.g. when a Noise node is selected, two drop-down list boxes let you select the noise type used in that node).

To the right end of the details area title, you will see a pair of buttons that let you browse to the previous and next nodes.

A preview of the node's output is displayed on the left end of the node details area, below the title. Check or uncheck the **Show in graph** option below the preview in order to adjust the size of the node preview in the function graph (if the option is checked, the node will appear large).


You can add titles and descriptions to your nodes using the **Title** and **Description** fields. This is useful when designing complex shaders that should be used by other parties. If no title is provided for the node, the name of the node is used instead (or the value of the constant for constant nodes). The description will appear as a tooltip when you float the mouse cursor over the node in the graph.

If the selected node has options, these options will appear to the right of the node preview.






The nature and type of these options depends on each node, and will be discussed further down.

Extracting Constant Values


If the selected node has options, some of these options exhibit an **Extract parameter** button (). If you click this button, the node parameter will be extracted: a new constant node will be created and this node will be connected to the right hand side of the initial node (parameter connection). When you extract a parameter, the extracted parameter node is automatically selected.

The value held by this newly created constant value node is the same as that of the parameter before it was extracted. At this point, the output of the node is not affected by this extraction. But now that the parameter is extracted, you may replace it with any type of node (e.g. a noise node!).



If you go back to the initial node, you will notice that the **Extract parameter** button () has been replaced by the **Disconnect parameter** button () and instead of displaying input controls, an indication that the node is “connected” appears. If you click the disconnect parameter button, or if you destroy the parameter’s link, the parameter will be reintegrated into the node and restored to its initial constant value.

The underlying power of this simple feature is truly amazing! For instance, most noise nodes have an origin (which is the origin of the noise wave). If you extract this origin, and then click the **Turbulence node** icon () the origin will be replaced by a turbulent value. Or if you extract and connect the scale parameter of a noise node to the slope input, the scale of the noise will vary automatically according to slope! Please turn [here](#) for a few simple examples of how to use this feature.

Note:

if you clicked the **Show parameter connections** icon () small links appear to the right of the node on the graph. Each one of these links corresponds to a parameter. You can extract a parameter directly by grabbing the corresponding handle and connecting it to another node.

Link Details

When a link is selected, a preview of the two nodes that are connected by the link is displayed in the **Link Details** area, together with two small arrows that let you go to either one of the nodes. Click the left arrow button () to go to the upper node, and click the right arrow () to go to the lower node.



Multi-edition of nodes


If you select several nodes, parameters of the first selected node will be displayed. Edit one parameter: value of this parameter will be applied on all selected node (if selected node has the same parameter).


MetaNodes

MetaNodes are a special type of node that encapsulates a graph, or part of a graph.

You could think of them as the ability to group several nodes of a graph, but the concept behind MetaNodes is in reality a lot more powerful. Because MetaNodes can be saved and retrieved for future use, and because they give you the ability to easily create a simple user interface around them, you should rather think of them as a building block for more complex graphs.

Creating a MetaNode

To create a new MetaNode, simply select several nodes in a graph and click the **Group selection as MetaNode** icon (). The selected items are replaced by the metanode, and all connections to the items you selected are automatically re-connected to the MetaNode.

To remove the MetaNode and re-expose its content in the graph, simply click the **Ungroup MetaNode** icon (). The nodes that had been moved into the MetaNode will re-appear in the graph.

Editing a MetaNode

You can edit the content of a MetaNode by double-clicking on the MetaNode, or by clicking the **Graph** button in the MetaNode properties.

When you edit the content of the MetaNode, a new instance of the *Function Graph* will appear, displaying the MetaNode graph. You can edit this graph in the usual way.

Building a MetaNode Interface

While editing the MetaNode graph, you will notice that some of the parameters of the nodes inside the MetaNode show as underlined when the cursor passes over the field name. This indicates that the parameter can be *published* which makes the corresponding parameter directly accessible from the top-level MetaNode details panel (at the bottom



of the *Function Graph*, when you select the MetaNode in the graph where the MetaNode is located).

When you click the parameter, a little dialog will popup, prompting you to enter the name under which you wish to see the parameter appear in the MetaNode options. Enter a name and click **OK**. Now, if you close the MetaNode graph and select the MetaNode in the main graph, you will see that your parameter appears with the name you provided.


Using this ability, you can very easily create a simple interface to your MetaNode, by exposing only those parameters that are really useful for controlling the functionality of the MetaNode.

Note:

MetaNode parameters are listed on the MetaNode options panel in the order in which they were published. There is no way to subsequently change this order.

Saving and Re-using MetaNodes

You can save a MetaNode for future use with one of two ways:

- Click the **Save** button in the MetaNode details panel (at the bottom of the *Function Graph*, when the MetaNode is selected), or
- Click the **Save** icon () at the lower right corner of the *Function Graph* when editing the MetaNode graph.

When you save the MetaNode, a Standard File Browser will appear, letting you select the file under which to save the MetaNode. By default, the browser opens on the MetaNodes folder, where you will see a set of sub-folders corresponding to the different types of nodes. If, for instance, you save your MetaNode in the **Filters** sub-folder, the MetaNode will subsequently appear in the Filter node menu, for easy access.

Using this feature, you can rapidly create and enrich your collection of ready-to-use MetaNodes, and thus rapidly create extremely elaborate function graphs.

Of course, you can also save MetaNodes into your own folder, and retrieve them using the Load button on the MetaNode details panel, or in the *MetaNode Function Graph*.

Example Use

In this example, we will use a preset MetaNode that acts as a filter to create a zero output around the origin of the world. This is particularly useful to create a flat area around the camera in a procedural terrain.

Create a standard procedural terrain. Connect a Fractal node to the Altitude output. Now, add a filter node behind the Fractal node output, and, in the display panel, browse to the MetaNodes category, and select the **Flat area at origin** MetaNode. This zeroes



the altitudes of the procedural terrain around the origin. Notice that a single parameter is published, that lets you control the size of the flat area.

If you double-click on the MetaNode, you will access the MetaNode graph. You will notice how the MetaNode calculates the distance to the world's origin, and applies this to the input value, to zero the value around the origin.

Here, a MetaNode was used to create a new type of filter. Despite a high level of inherent complexity, it is incredibly easy to use this new “filter”, because that complexity is hidden away behind the MetaNode interface.

Locking MetaNode Content

You can prevent other users from viewing or editing the content of your MetaNodes by pressing the **Lock** button on the MetaNode details panel. Beware, however, that once a MetaNode has been locked it is impossible to unlock it.

Scene Graph Approach

The Scene Graph approach is a way of viewing your scene, from a graph-based point of view. This approach lets you define custom relationships between all the objects, terrains and materials of your scene, by visually connecting their properties using a set of nodes.


This is a powerful concept that lets you create incredibly elaborate scripts to control object, terrain or material properties based on other items.

For instance, control the distribution of rocks on terrains according to local roughness, assemble automatic piston rigs, make objects turn red when other objects grow large, turn lights on when doors open, etc.

You can even design scripts that will make objects interact dynamically with procedural terrains, to recreate, for example, the deep impact effect of an asteroid hitting a terrain (check out the tutorial [here](#)).




Object Graphs

Each object can now be controlled using a graph. This graph is unique to the object, and is accessed by selecting the menu option **Object | Edit Graph**, or by clicking on the **Edit object graph** button () at the bottom of the *World Browser* (see [here](#)).

The object graph features the different properties of the object both as input and output. The input value of an object property contains the current value of that property. For instance, the **Position** input contains the current position of the object.

If you connect something to one of the output properties, you are forcing the value of that property, and it will no longer be possible to change it in the user interface. For instance, if you connect the **Size** output directly to the **Position** input, you are forcing the size of the object to be exactly equal to its position (obviously not very useful, but this is just an example).

Now, if you close the *Object Graph*, you will notice that the size of the object is related to its position, and you are unable to change the size of the object. If you were to animate the position of the object, you would see its size vary according to its position, despite the fact that no size animation is defined. Actually, if you look at the *Properties Timeline*, you will notice that the **Size** property is marked as being **--connected--**.

The **Edit object graph** button of objects that have a graph attached to them appears toggled down () at the bottom of the *World Browser*.

To remove an *Object Graph* completely, simply delete all nodes and connections in the graph.

Connecting Graphs

The *Object Graph* is very useful to customize the behavior of a single object. However, its true power comes in the ability to connect different graphs together, so that you can cause some objects to react to other object's properties. You can also connect different types of graphs together. For instance, you can connect a material graph to an object's graph, or to the altitude production function of a procedural terrain.

Creating relationships between objects, materials or procedural terrains is done by importing/exporting parameters between graphs.

External Dependencies

The **External Dependency** input is the type of input that will let you import the value of a property of another object into a particular graph.



In the object graph, select **Add Input Node | External Dependency** from the popup menu. This adds an External dependency input to the graph. With this input selected, look at the node details. You will see a **Dependency** drop-down menu, where you can select which property of which object you want to import into the graph.

For instance, create a cylinder and a sphere. Select the cylinder, and edit its graph. In the cylinder graph, add an **External Dependency** input node on the sphere's position. Connect the cylinder size output to the sphere position dependency node, and close the graph. Now, notice how the size of the cylinder changes as you move the sphere around.

- **Dependency:** select the item for the dependency relationship from the drop-list. Depending on the item selected, **Size** may be available as a qualifier. If you then select **True** dimensions, you can specify the size in **VUE units**, **Display units**, or specific units such as meters or feet.
- **Relative to parent:** if the object is part of a group, a Boolean operation or a Metablob, selecting this option will express the dependency in coordinates that are relative to the group.

Note:

The relationship between the objects will be preserved, even if you group them together. Using this capability, you can, for instance, very easily create complex mechanical components such as the piston rig (see this [tutorial](#)).

You can also refer to object properties from inside a material graph, or a procedural terrain altitude graph.

Recall Dependency

The **Recall Dependency** type of node is incredibly powerful, as it behaves like the **External Dependency** node, but will actually remember the value of the dependency for the indicated time.

To create such a node, select the **Add Input Node | Recall Dependency** command from the popup menu.

- **Delay:** this parameter controls the amount of time by which the value of the dependency is delayed, so that the value that is returned by the node is actually the value that the dependency had “Delay time” before.
- **Dependency:** select the item for the dependency relationship from the drop-list. Depending on the item selected, **Size** may be available as a qualifier. If you then select **True** dimensions, you can specify the size in **VUE units**, **Display units**, or specific units such as meters or feet.

Note:

the **Recall Dependency** nodes only work in the context of *Object Graphs*.



Exporting Values

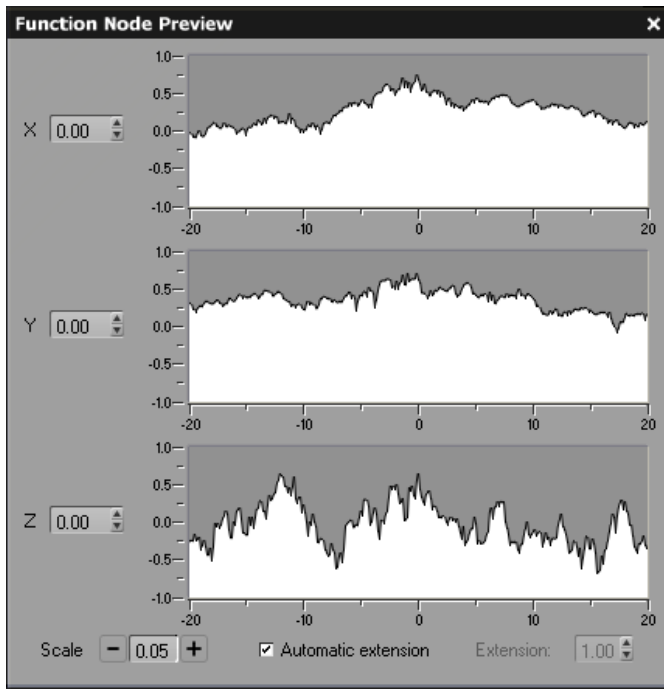
In the previous section, we saw how we can easily connect several object graphs together by referring to the properties of other objects. However, there may be cases where you wish to refer to something else than a property of an object. For instance, you could want to connect to an intermediate value that is calculated in another graph.

This can be done by exporting values. To do this, select **Add Output Node | Custom Dependency** from the popup menu. This creates a **Custom Dependency** output node. You can connect this node to whatever you wish to refer to in another graph.


You could use this feature to distribute materials on a procedural terrain according to terrain roughness. Let's create a new procedural terrain and edit the terrain altitude graph. Create a **Terrain Fractal** node and connect that to the **Altitude** output. Now, create a **Custom Dependency** output node in the terrain altitude graph, and connect this node to the Terrain Fractal node. Notice that, at the time of connecting, you are asked to select between two types of outputs (**Altitude** and **Rough areas**). Connect the dependency output to the **Rough areas** output of the **Terrain Fractal** node and close the graph. Now edit the material that is assigned to the terrain, and create a mixed material (see [here](#)). Edit the **Distribution** function, and, in the function graph, create an **External Dependency** input node. From the **Dependency** drop-down menu, select the **Custom Dependency** option of the procedural terrain, and connect this to the **Distribution** output. Close the graph. If you render the terrain, you will see that the distribution of materials now depends on the roughness of the terrain.

Function Node Preview





Function Node Preview

This panel is accessed by pressing the **Function Node Preview** icon () in the top toolbar of the *Function Graph* (see [here](#)).

If this icon is orange, this panel will appear automatically each time you open the *Function Graph*. This panel can be resized by its edges. By resizing the panel, you can enlarge the preview area.

The *Function Node Preview* displays a preview of the output of the currently selected node as a set of curves or colors. If the selected node outputs a number, the panel will display 3 curves showing a section of the node's output along each one of the 3 axes. If the selected node outputs a color, the panel will display 3 colored bands showing the output of the node along each one of the 3 axes. If the selected node outputs a vector or a texture, the panel displays nothing.

You can change the origin of the observation using the **X**, **Y** and **Z** input fields. The origin of the observation is the point at the exact center of the 3 curves/bands. Under each curve/band is an indication of the relative offset from the origin of observation.

Use the **Scale** setting to adjust the portion of the outputs that is displayed in the curves/bands – that is the size of the observation window along the axes that is viewed on



each curve/band. Press the **-** and **+** buttons to respectively reduce/increase the scale of the preview.

You can also change the origin of the observation by dragging the curves/bands: click on one of the curves/bands and drag the mouse to move the observation window. To change the scale of observation, press the **Control** key and drag the mouse up/down to zoom in/out.


Curves and Extension

When the selected node outputs a number, the *Function Node Preview* displays output values as 3 curves. A ruler indicating the amplitude of the output will be displayed on the left edge of each curve. These values are controlled through the **Extension** parameter.

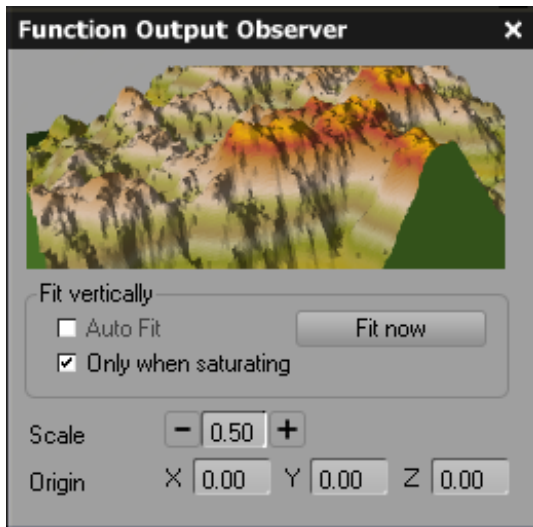
Automatic extension: select this option if you want the *Function Node Preview* to automatically select the appropriate extension so that all output values fit in the curves. This setting is only available when the output of the currently selected node is a number. The corresponding extension is displayed in the Extension field. When Automatic extension is selected, this field is disabled.

Extension: use this to control the minimum and maximum values displayed on the curves. This setting is only available when the output of the currently selected node is a number, and the Automatic extension option is deselected.

Function Output Observer




Material output



Terrain output



Function Output Observer

This panel is accessed by pressing the **Function Output Observer** icon () in the top toolbar of the *Function Graph* (see [here](#)).

If this icon is orange, this panel will appear automatically each time you open the *Function Graph*. This panel can be resized by its edges. By resizing the panel, you can enlarge the preview area.

The *Function Output Observer* is a simple panel that displays a view of the function's output. The nature of the view depends on the function's context:

- If you are editing a function that is part of the definition of a material, the panel shows a preview of the material.
- If you are editing the function that defines the altitudes of a procedural terrain, the panel displays a preview of the procedural terrain. You can rotate, pan and zoom the terrain preview like in the *Terrain Editor* (see [here](#)).

Scale: use this setting to zoom in or out of the preview. Press the  and  buttons to respectively reduce/increase the scale of the preview.

Origin: use this to enter the origin of the function observation. This point is the one at the center of the previews (center of the sphere if you are viewing a material on a sphere, center of the plane if you are viewing a material on a plane, center of the terrain if you



are viewing a terrain).

Auto Fit and **Fit Now** allow you to get a reliable feedback on the modifications performed on the function (the **Fit Vertically** option displays the output values in a range that fits the dialog, but, the updated terrain can be quite different from the preview because of the difference in the final range). By using **Fit Now**, the user can see the evolution of the function more easily.

Cyclic Nodes

VUE includes a set of noise and fractal nodes that are cyclic. This means that instead of an every varying pattern repeating over the mapped space, a single pattern will repeat itself periodically along all axes (3D space for 3D functions, and also along time for 4D functions).

The advantage of these nodes is that there is no seam between adjacent repetitions of the pattern.

Currently these are available for Noises and Fractals.

How to Use Them

Cyclic noises and fractals are located in a sub-menu of the standard noises and fractals. All parameters are exactly the same as those in the corresponding non-cyclic flavor of the noise or fractal. For technical reasons, not all noises and fractals have a corresponding cyclic version.

There are additional parameters to specify the repetition period over each of the 3 or 4 axis of the function. The period can be different along each axis, which leads to non-square patters (but still seamless). The period is expressed as a multiple of the wavelength.

Note:

you can also use a cyclic noise in a non-cyclic fractal, but it will lead to results much more predicable than a cyclic fractal, because the periodicity will be the same at each octave, whereas it is not in a cyclic fractal.



Noise Nodes

Common Parameters

Scale

The scale parameter is a number that controls the overall scale of the noise. Larger values mean that the noise pattern looks larger. This parameter works in conjunction with the Wavelength parameter to determine the final scale of the noise along each axis.

Wavelength

Whereas the scale parameter only lets you control the size of the noise pattern globally, the Wavelength parameter is a vector parameter that lets you adjust the scale of the noise along each axis. For instance, if you want the noise to vary only along the Z axis, enter 0 in the X and Y wavelengths.

Origin


The origin parameter is a vector that indicates the point at which the noise originates. By modifying this value, you can shift the noise pattern around. If you plug the Origin parameter into a Turbulence node, you will add turbulence to the noise.

Cellular Patterns


A cyclic version of **Cellular Patterns** is available. Refer to the previous section for more information.




Chipped, Crystals, Pebbles



Chipped noise



Crystal noise




Pebble noise


These nodes do not define any additional parameters.



Drought



Default crack width



Crack width = 0.25

This noise looks like the patterns created by wet soil that has dried out.

Crack width: controls the width of the cracks.

Voronoi

Voronoi noises produce patterns that are based on the distance to randomly positioned seed points on a grid.


Neighbor mode: determines what distance is taken into account to produce the noise pattern:

- **Closest neighbor:** the shortest distance to a neighboring seed point,
- **2nd closest neighbor:** not the shortest distance, but the 2nd shortest,
- **3rd closest neighbor:** not the shortest distance, but the 3rd shortest,
- **4th closest neighbor:** not the shortest distance, but the 4th shortest,




- **1st – 2nd neighbors:** distance to the closest neighbor minus distance to the 2nd closest,
- **2nd – 3rd neighbors:** distance to the 2nd closest neighbor minus distance to the 3rd closest,
- **3rd – 4th neighbors:** distance to the 3rd closest neighbor minus distance to the 4th closest,

Voronoi noises: neighbor mode




Closest neighbor (default)




2nd closest neighbor






3rd closest neighbor




4th closest neighbor




1st – 2nd neighbors





2nd – 3rd neighbors




3rd – 4th neighbors

Voronoi profile: determines the curvature of the noise over a fragment as the distance increases:


- **Flat:** creates fragments of uniform value, the distance to the closest neighbor being used on the entire fragment,
- **Spikes:** the noise amplitude varies linearly with the distance, creating pointy shapes,
- **Angles:** a little more rounded than spikes,
- **Round:** yet a little more rounded,
- **Smooth rounded:** the most rounded Voronoi profile.

Voronoi noises: profiles






Flat




Spikes




Angles





Rounded (default)



Smooth-rounded


Voronoi (Altitude)

Basically the same as the above Voronoi noise, except that the altitudes of the different fragments varies randomly. The Voronoi Altitude Flat noise is identical to the Voronoi Flat noise. You cannot select the neighbor mode for this type of Voronoi.


Voronoi profile: same as above.

Voronoi altitude: profile






Flat




Spikes




Angles





Rounded (default)



Smooth-rounded


Voronoi (Generalized)

The generalized Voronoi noise is yet another variation of the Voronoi noises where the curvature of the fragments is adjustable continually, and where you can adjust the amount of randomness in the size of the fragments.


Randomness: controls the amount of randomness in the size and shape of the different fragments that constitute the noise pattern. If 0 randomness is entered, the fragments will all be square.

Voronoi Generalized: Randomness






Randomness = 0




Randomness = 0.2



Randomness = 0.5







Randomness = 1

Voronoi profile: this controls the curvature of the fragments. It is similar to the Voronoi type described above, except that it lets you vary the curvature continuously.

Voronoi Generalized: profile




Profile = 1




Profile = 1.25





Profile = 1.75



Profile = 2

Distributed Patterns


These types of noises create a pattern by scattering a basic shape randomly in noise space. Warning: these types of noise are very slow to compute. Use the 2D counterparts wherever possible. A cyclic version of **Distributed Patterns** is available. Refer [here](#) for more information.

Round Samples and Round Samples (2D)


Distributes round patterns. These two noises are very similar. The only difference between the two noises is that the second version only scatters the patterns along the X and Y axes, resulting in much quicker evaluation (which is especially useful for procedural terrains).

Size: controls the average size of the patterns.






Size = 0.2




Size = 0.4




Size = 0.6






Size = 1

Randomness: controls the randomness in the distribution of patterns, both in terms of size and position.




Randomness = 0




Randomness = 0.2





Randomness = 0.5




Randomness = 1


Shape: controls the shape of the patterns in terms of altitude:

- **Cylinder:** the altitude of the pattern is constant all over its surface; the noise scatters tiny cylinders in noise space,
- **Cone:** the altitude of the pattern varies linearly with the distance to the center of the pattern; the noise scatters tiny cones in noise space,
- **Round:** the noise scatters hemispheres in noise space,
- **Smooth round:** the noise scatters little round bumps that connect smoothly with the underlying geometry,
- **Cone tower:** same as cone, except the cones are placed on tiny cylinders,
- **Round tower:** same as round, except the hemispheres are atop tiny cylinders.






Cylinder




Cone






Round




Smooth-round





Cone-tower




Round-tower


Samples: indicates the number of patterns that are scattered per grid unit:

- **1 sample per cell:** only one pattern will be mapped per grid unit,
- **2 samples per cell:** exactly two patterns will be mapped per grid unit,
- **3 samples per cell:** exactly three patterns will be mapped per grid unit,
- **4 samples per cell:** exactly four patterns will be mapped per grid unit,






1 sample per cell




2 samples per cell




3 samples per cell






4 samples per cell

- **0 to 1 samples per cell:** each cell will contain a maximum of 1 pattern, maybe none,
- **0 to 2 samples per cell:** each cell will contain anything from 0 through 2 patterns,
- **0 to 3 samples per cell:** each cell will contain anything from 0 through 3 patterns,
- **0 to 4 samples per cell:** each cell will contain anything from 0 through 4 patterns,




0 to 1 sample per cell




0 to 2 samples per cell






0 to 3 samples per cell




0 to 4 samples per cell

- **1 to 2 samples per cell:** each cell will contain either 1 or 2 patterns,
- **1 to 3 samples per cell:** each cell will contain anything from 1 through 3 patterns,
- **1 to 4 samples per cell:** each cell will contain anything from 1 through 4 patterns,




1 to 2 samples per cell






1 to 3 samples per cell




1 to 4 samples per cell

- **2 to 3 samples per cell:** each cell will contain either 2 or 3 patterns,
- **2 to 4 samples per cell:** each cell will contain anything from 2 through 4 patterns,
- **3 to 4 samples per cell:** each cell will contain anything from 3 through 4 patterns,




2 to 3 samples per cell






2 to 4 samples per cell




3 to 4 samples per cell

Random altitudes: this option, when checked, will assign a random altitude to each pattern.




Without random altitude






Without random altitude




With random altitude




With random altitude

Find maximums: if this option is checked, the noise will find the maximum of all the patterns that overlap the point of evaluation of the noise.






Without find maximums




Without find maximums



With find maximums





With find maximums


Square Samples and Square Samples (2D)

These noises are similar to the Round Sample noises, except they map square patterns instead of round patterns.

Size: same as the Round Samples noise.


Randomness: same as the Round Samples noise.

Scale variations: controls the amount of variation in the aspect ratio of the square patterns. If 0, all patterns will be square. If non zero, the patterns will be more or less stretched.




Scale variation = 0






Scale variation = 0.3




Scale variation = 0.6




Scale variation = 1

Angular variations: controls how well the patterns are aligned with the noise axes. If 0, all square patterns are aligned with the axes. If non zero, the patterns will be more or less twisted.






Angular variation = 0




Angular variation = 0.3



Angular variation = 0.6





Angular variation = 1


Shape: similar to the shape parameter of the Round Samples noise, except applied to square patterns:

- **Cube:** the noise scatters little cubes in noise space,
- **Pyramid:** the noise scatters little pyramids in noise space,
- **Round pyramid:** the noise scatters pyramids that have a rounded profile in noise space,
- **Pyramid tower:** same as pyramid, except the pyramids are placed atop little cubes,
- **Round pyramid tower:** same as round pyramid, except the round pyramids are placed atop little cubes.




Cube






Pyramid




Round pyramid





Pyramid tower



Round pyramidtower

Samples: same as the Round Samples noise.

Random altitudes: same as the Round Samples noise.

Find maximums: same as the Round Samples noise.




Flat Patterns


The noises in this category create flat patterns. They don't work so well for bumps, because they tend to create sharp edges. There are noises in other categories that also produce flat patterns. A cyclic version of **Flat Patterns** is available. Refer [here](#) for more information.

Varying Blocks, Clumps, Water Cress

These noises do not define any additional parameters.




Varying Blocks



Clumps






Water Cress

Line Patterns

The noises in this category create patterns that are mostly based on lines. A cyclic version of **Line Patterns** is available. Refer [here](#) for more information.


Lines, Fabric

These noises do not define any additional parameters.



Lines






Fabric

Cracks

Crack width: controls the width of the cracks.




Crack width = 0.01




Crack width = 0.03






Crack width = 0.20



Crack width = 0.70

Sparse Cracks

Crack width: controls the width of the cracks.




Crack width = 0.01






Crack width = 0.05



Crack width = 0.20



Crack width = 0.80




Math Patterns

The noises in this category define simple patterns based on mathematical functions. They are mostly used to combine other noises together or create special patterns that require the regularity of mathematical functions.


Onion, Wavelet, Step (Vertical), Step (Gradual), Tooth (Rectangular), Tooth (Triangular), Tooth (Gaussian), Radial Sine, Sine Wave, Triangular Wave, Leopard, Saw Teeth, Water Wave

These noises do not define any additional parameters.




Radial sine






Sine wave




Triangular






Leopard




Saw teeth






Water wave




Spiral






Rectangular




Onion





Wavelet




Step (rectangular)






Step (smooth)




Tooth (rectangular)





Tooth (triangular)




Tooth (Gaussian)

Spiral

Radial expansion: if checked, this option will make the wavelength of the spiral pattern increase as it moves away from its origin.






Without radial expansion




With radial expansion

Vertical warp: if set, this option indicates that the phase of the spiral changes with the altitude.



Without vertical warp







With vertical warp

Rectangular Wave

Step width: controls the steepness of the transitions between low and high values. 0 means perfectly vertical edges.




Step width = 0




Step width = 0.1






Step width = 0.3




Step width = 0.5

Up/down ratio: controls the size of the patterns when the output is high (up) versus when it is low (down). Similar to the pulse width. This parameter only has an effect if the step width is non zero.




Up/down ratio = 0.5






Up/down ratio = 0.6



Up/down ratio = 0.7



Up/down ratio = 0.9


Slope: controls which transitions are done abruptly and which ones are done smoothly. This parameter only has an effect if the step width is non zero.

Slope up and down: if the step width parameter is non zero, both transitions from up to down and from down to up will be gradual.




Slope up only: only transitions from down to up will be gradual. Transitions from up to down will be abrupt.


Slope down only: only transitions from up to down will be gradual. Transitions from down to up will be abrupt.



Slope up and down



Slope up only




Slope down only




Other Patterns

Dots, Water (Calm), Water (Rough), Granite


These noises do not define any additional parameters.




Dots



Water (Calm)



Water (Rough)



Granite

Perlin Noises

Noises in this category are all based on work by Ken Perlin. They produce repeatable patterns that look random and are the basis to most procedural textures.

There are 3 types of basic Perlin noises: Linear, Value and Gradient. Linear Perlin produces sharp edges, Value is a slightly better but slower version of the Perlin noise, and Gradient is the best (and also slowest version). Each type of Perlin noise has its pros and cons in terms of looks. A cyclic version of **Perlin Noises** is available. Refer [here](#) for more information.

Common Parameter


Ridged: this option creates ridges in the noise pattern. It also has the side effect of making the noise higher on average.

Animated: when this option is selected, the noise will be evaluated in 4 dimensions instead of 3, the fourth dimension being that of time. This will result in a noise that produces patterns that change over time. Whenever you select this option, a link will be automatically established with the “Time” input.


Linear, Value, Gradient

The basis Perlin styles of noises. No additional parameters – aside from the Ridged option – are defined for these noises.






Linear




Value



Gradient





Gradient + Ridged


Value-Gradient (Variable), Linear-Value-Gradient (Variable)

These two noises are combinations of the base Perlin noises. The different types of noises are blended according to a random pattern.


No additional parameters – aside from the Ridged option – are defined for these noises.




Value-Gradient (variable)



Value-Gradient (variable) + Ridged



Linear-Value-Gradient (variable)



Linear-Value-Gradient (variable) + Ridged


Value-Gradient, Linear-Value-Gradient

These two noises are combinations of the base Perlin noises. The different types of noises are blended according to the Ratio setting.

Ratio: controls the proportion of each type of Perlin noise in the final noise.

Value-Gradient:






Ratio = 0




Ratio = 0.2




Ratio = 0.5 (default)






Ratio = 0.9

Linear-Value-Gradient:




Ratio = -0.9




Ratio = -0.3





Ratio = 0 (default)



Ratio = 0.7


Square Patterns

A cyclic version of **Square Patterns** is available. Refer [here](#) for more information.


Random Altitudes, Squares, Squares (Pairs), Stones, Square Blobs, Square Stones

These noises do not define any additional parameters.






Random Altitudes




Squares




Squares (pairs)






Stones



Square Blobs



Square Stones




Fractal Nodes

Common Parameters

The following settings are common to all fractal nodes (some of them are not available in the “Basic Repeater”, “Fast Perlin Fractal” or “Open Ocean” nodes, because these nodes are simplified or degenerate forms of fractals).


Base noise: to create its output, the fractal node replicates the base noise at different frequencies, and with different amplitudes. This drop-down menu box lets you select the noise to be used by the fractal. If the noise defines extra parameters, you can access these extra parameters by clicking on the **Edit** button. This will open a *Node Options* dialog, letting you adjust the properties of the noise. If the noise has no extra parameters, the Edit button remain disabled. If you select a noise that is time dependent, a link will automatically be established with the “Time” input.

With rotation: check this option if you want the noise to be rotated in between each harmonic. This is useful if the base noise exhibits strong directional features and you want to minimize these directional features.



Without rotation






With rotation

Wavelength: this is the same as the wavelength parameter of noise nodes (see [here](#)).

Origin: this is the same as the origin parameter of noise nodes (see [here](#)).


Metascale: This is the scale of global variations in the noise. If you take a fractal that represents a mountain, the largest feature scale would be the size of mountains, the metascale would be the size of the entire mountain range, and the smallest feature would be the smallest detail.

Largest feature: this is the same as the scale parameter of noise nodes (see [here](#)). Generally speaking, fractal nodes should have features that are larger than the scale at which the fractal will be observed.




Largest feature = 0.1






Largest feature = 0.5




Largest feature = 2



Largest feature = 10

Smallest feature: by default, when computing a fractal pattern, VUE will keep adding detail until these details are so small that they cannot be seen in the final picture. This is the default behavior when the smallest feature setting is left at 0. There are cases where you may want to skip the smaller details in the fractal, in which case you should indicate the scale of the smallest details you want, using this setting.





Smallest feature = 0



Smallest feature = 0.02




Smallest feature = 0.05






Smallest feature = 0.10

Roughness: this parameter controls the overall roughness of the fractal pattern. Namely, the amplitude of each iteration of the fractal's base noise is multiplied by the Roughness parameter. The default value of 0.5 will produce a fractal pattern with the same level of detail at all scales. Smaller values for the roughness parameter will produce a smoother surface, whereas values greater than 0.5 will yield spiky patterns with lots of small details.




Roughness = 0.2




Roughness = 0.4






Roughness = 0.6




Roughness = 1

Gain: this parameter controls the overall amplitude of the signal output by the fractal. Because fractal patterns can have very large features, their output can be in a much larger range than the standard noise range of -1 through 1. You can use this parameter to tone down the amplitude of the fractal's output.




Gain = 0.5






Gain = 1



Gain = 4



Gain = 10




Stretch Damping: This setting is only available if the fractal is stretched along one or several axes (non uniform wavelength). Stretch damping will reduce the amount of stretching applied to the higher frequencies in the fractal, thus avoiding the entire fractal pattern looking as if it had been stretched.



The fractal's output is modulated by a user defined filter. The amount of filtering can be made to vary according to the harmonic. If no filter is defined, this processing is ignored. You can define the range of values in between which the filter is applied.


Filter: this is the filter that will define the profile of the altitudes. Double-click on the filter preview to load a new filter, or select **Edit** from the popup menu to customize the filter. Please turn [here](#) for details on the *Filter Editor*.

Function view






Function graph





Filter






Creep in: this parameter controls how much of the original (unfiltered) signal gets mixed back into the signal at each iteration. Higher values mean that the filter only affect a few large-size harmonics.


Min and **Max:** the range of values to which the filter applies is automatically defined according to the other settings in the fractal. Using Min and Max, you can adjust this range (for instance, if you want the filter to apply to an intermediate range of values only). The Min and Max values are given as percentages of the full range computed by the fractal.

For improved clarity, we will use the above cut-off filter in the examples below:




Creep in = 0






Creep in = 0.1



Creep in = 0.3





Creep in = 0.5

Outputs: most fractals are capable of outputting both an **Altitude** value (the default usage) and also a **Rough areas** value, which can be used to drive the distribution of materials according to the local roughness of the fractal pattern. When the second output is connected, a 2nd output: **Detect rough areas** option appears with the **Ref. feature size** setting. This is typically used to control the distribution of materials on the terrain according to fractal roughness.

Basic Repeater


The Basic Repeater is a special type of fractal that is in some respect “degenerate”. The reason for this is that basic repeaters only add a limited amount of detail to their patterns, whereas true fractals will add infinite details. What this means is that if you zoom in close onto a basic repeater pattern, you will begin to notice the lack of detail. There are cases when the basic repeater can be useful because it offers greater control over the harmonic behavior of the noise. Whenever possible, however, you should prefer true fractal patterns.

Note:


Although the Basic Repeater can have a filter assigned to its output values, the Creep in parameter defined above is not available in this node.

Repeat: this parameter controls the number of times the base noise is repeated in order to produce the final pattern. Higher values will produce very detailed patterns, but will take longer to render. It is rarely useful to use values higher than 4-6.






Repeat = 0




Repeat = 1




Repeat = 2






Repeat = 10

Scale: this parameter controls the scaling ratio that is applied to the base noise's wavelength in between each iteration of the noise. Values close to 0.5 produce the best results; values greater than 0.5 will enhance larger elements, whereas values under 0.5 will enhance the smaller details.




Scale = 1




Scale = 0.75






Scale = 0.5




Scale = 0.2

Amplitude: this parameter controls the amplitude ratio that is applied to the base noise's amplitude in between each iteration of the noise.




Amplitude = 0.25






Amplitude = 0.5



Amplitude = 0.75




Amplitude = 2

Combination mode: this drop-down list defines the method used to combine the noise iterations together:


- **Add:** values are added together.
- **Blend:** values are averaged.



- **Variable roughness:** values are added depending on the result of the first iteration. Low first iteration values mean lots of successive iterations being added in, high values mean little influence of successive iterations.
- **Variable roughness (abs):** same as Variable roughness, except the distance to 0.5 is considered instead of the value of the first iteration itself.
- **Max:** the biggest value is retained.
- **Max (abs):** the value that is the furthest from 0.5 is retained.
- **Min:** the smallest value is retained.
- **Min (abs):** the value that is the closest to 0.5 is retained.
- **Multiply:** values are multiplied together.




Add




Blend






Variable roughness




Variable roughness ABS




Max






Max ABS




Min



Min ABS






Multiply


Simple Fractal

This is the simplest type of fractal. It repeats the base noise uniformly.

The simple fractal node does not define any additional parameters.




Roughness = 0.5



Roughness = 1.0





Roughness = 1.5


A cyclic version of this **Simple Fractal** is available. A cyclic **Animated Simple Fractal** is also available. Refer [here](#) for more information.

Grainy Fractal

The Grainy fractal is particularly useful for color and bump patterns that exhibit a lot of detail at all frequencies.



Roughness = 0.75



Roughness = 1.0

With rotation: check this option if you want the noise to be rotated in between each harmonic. This is useful if the base noise exhibits strong directional features and you want to minimize these directional features.

Double noise: this option adds more interesting variations to the base noise. It is however more complex to compute.


Noise Variation: use the **Variation strength**, **Variation roughness**, and **Smooth area altitude** settings to control how the grain in the noise varies, and to create smooth and grainy areas.

Other: use the **Distortion** and **Filter Steepness** to add distortion to the overall fractal pattern, as if it had been smeared around randomly. **Steepness** controls the amount of contrast in the noise.

A cyclic version of this **Grainy Fractal** is available. Refer [here](#) for more information.


Terrain Fractal

This is the same as the grainy fractal, except that the noise/landscape type parameter can be made to vary according to the altitude of previous iterations of the base noise. This results in smooth areas at certain levels, and rougher areas away from this level. This node is mostly used for creating natural-looking terrains.




Ridges Roughness = 0.5






Ridges Roughness = 1.0




Plain noise Roughness = 0.5




Plain noise Roughness = 1.0






Billows Roughness = 0.5




Billows Roughness = 1.0




Ridge Mix Roughness = 0.5






Ridge Mix Roughness = 1.0



Billow-Ridge Mix Roughness = 0.5



Billow-Ridge Mix Roughness = 1.0

Distortion: adds distortion to the overall fractal pattern, as if it had been smeared around randomly.

Bump surge: causes bumpy areas to rise above or sink below the average surface.



Ridge smoothness: controls how much rounding is applied to the Perlin ridges/Billows. This setting is not available in Plain noise mode.

Noise/Landscape type: this drop-down list defines the shape of the base noise.

- **Ridges:** a modified version of the Perlin noise that creates sharp ridges.
- **Plain noise:** the basic Perlin noise.
- **Billows:** a modified version of the Perlin noise that creates billows.
- **Ridge-Mix:** a blend of different ridged Perlin noises.
- **Billow-Ridge Mix:** a blend of billowy and ridged Perlin noises.

Blend: this setting is only available for mixed noise types. It controls the method used to combine the noise and landscape iterations together.

A cyclic version of this **Terrain Fractal** is available. Refer [here](#) for more information.

Terrain Fractal 2

This node is a fractal function designed to create realistic terrain landscapes, similar to the **Terrain Fractal**. Differences are:

- This fractal has a better variability of shapes, and the rough areas simulating rocks and cliffs are more convincingly integrated in the relief.
- An optional stratification can be applied to create an effect similar to what a separate Strata Filter node would achieve if fed with the fractal's output, but with the added advantage of benefiting from knowledge of some of the fractal's internal value: For example, the strata follow the general relief of the landscape, to simulate the deformation of actual geological strata due to landscape movements after the formation of the strata themselves.
- Also, the stratification process is modulated to be much more visible on rough areas than on smooth areas. This is because the smooth areas represent parts of the landscape where sediments have covered the underlying, stratified rocks.
- Like most other fractals, **Terrain Fractal 2** also provides a **2nd output** which value reflects the terrain roughness at the evaluated point.

Terrain Fractal 2 comes with several groups of parameters. The first group contains generic parameters which are the same as those seen on other fractals. Please refer to previous documentation for details.

The **Overall aspect** parameters control the influence of the first few octaves of the fractal over the rest of the algorithms. These octaves will define regions with different density of rocks.

Turbulence: controls the overall distortion of the terrain.



Turbulence damping: controls the influence of the first octaves' turbulence on subsequent octaves of noise.

Large scale smoothness: controls the smoothness of the transition from regions of low rock density to others of high rock density.

Large scale contrast: defines the range in which the rock population density can vary.

Buoyancy: controls the balance between large scale noise octaves and smaller scale ones. A positive buoyancy means that the average altitude will be low and the rocky features will raise above it, whereas with a negative value the features will dig below a higher average altitude. A null buoyancy means that the average altitude will be around zero while some features will be above it and some below it.

This fractal tries to simulate rocks emerging from a sedimentary soil. These rocks tend to be gathered at specific places where the soil thickness is lower, whereas in thicker soil areas they are almost all hidden below the sediments. The **Ground aspect** parameters control this.

Bump surge: controls how much the rocks will spring up out of the ground.

Rock abundance: controls the quantity of rocks visible.

Soil thickness: controls the typical thickness of the layer covering the rocks. A thin layer will let more rocks show up, and most of the smoother areas will still retain a little bit of roughness. On the other hand, a thicker layer will cover more rocks, and most of the smoother areas will have almost no roughness at all.

Rock dispersion: controls how much rocks tend to be scattered in the landscape rather than gathered in specific areas.

The **Strata processing** parameters are similar to those available on the **Strata** filter located in the **Recursive** filters subcategory:

Processing strength: controls the influence of the strata filtering over the landscape.

Layer spacing: controls the height of the main layer.

Offset: allows fine-tuning of the vertical strata pattern positioning with respect to the underlying terrain.


Rocky Mountains Fractal and Eroded Rocky Mountains Fractal

This is a new type of fractal that produces terrain features typical of the tertiary geologic period.




Terrain features generated by this fractal are fully user adjustable. The fractal can also be used to drive material distributions and produce a wide variety of appearances.


The **Base settings** section of the fractal settings are the same as the Terrain fractal.



Separate mountains Roughness = 0.5




Separate mountains Roughness = 1.0



No Separate mountains Roughness = 0.5





No Separate mountains Roughness = 1.0

Two flavors of this fractal are available in the **Overall aspects** section. The **Separate mountains** checkbox drives this difference.

When this option is checked, the terrain appears as independent mountain “blocks” placed side by side. This is useful when in need of one or several big summits or to overlay on top of another relief. When the **Separate mountains** option is unchecked, the terrain appears as independent basins separated by irregular mountain ridges. This can be a useful basis to define interconnected or separate valleys, especially with proper distortion.

Number of iterations: this fractal is very specific in that it produces irregular ridges that appear at each iteration. This means the lower frequency components will not be as visible as they are in regular fractals. The trade-off is that it can be quite slow to compute with a lot of iterations.

Subdivision quality: the algorithm is in fact an approximation of an algorithm intrinsically much slower. Therefore, some faults (discontinuities) can appear in the fractal. This quality parameter allows some control over the performance/quality trade-off of the implementation.

Scale factor: each new iteration adds irregularities at a scale smaller than the previous iteration. This parameter defines how much smaller each new iteration will be. A higher scale factor will allow for smaller details with fewer iterations, but it will also be more predictable and less appealing.

Flat level (per iteration): each iteration applies some pattern which is made of some very smooth areas and some much rougher, ridged areas. This parameter controls the balance between the two types of areas. A high value will leave fewer ridges, while a lower value will yield much smaller smooth areas.

Ground level: this parameter, especially useful when **Separate mountains** is ticked, makes the fractal “sink” into the ground.



Stretch and distortion

Stretch factor: the pattern applied at each new iteration is stretched along some privileged direction, to reflect the way real ridge networks actually look like in a mountain range. This parameter controls the amount of stretching. This parameter is ignored when the **Separate mountains** setting is not checked, unless you are using **Eroded Rocky Mountains**.

Distortion: this parameter is quite similar to its namesake in Terrain Fractal. It distorts the input coordinates in order to perturb the fractal overall aspect.

Optional rocks: when activated, this feature overlays rocks on top of the fractal itself.

Rock correlation: This optional feature adds rocks in the rough areas, while preserving the smooth aspect of the flatter areas. To do this, it relies on the ridges seen at the iteration given by this parameter.

Rock roughness and **Rock height:** These allow for finer control over the aspect of the overlaid rocky fractal, and behave like **Roughness** and **Gain** would in a regular fractal.

2nd Output: Detect rough areas

Ref. feature size: The rough area output is divided into two subranges for easier filtering:


- Where there are no overlaid rocks, the underlying fractal's rough value is used, as if **Optionalrocks** were set to **None**. It is mapped to $[-1;0]$.
- On overlaid rocks areas, their height over the underlying fractal is used as “rough” value, mapped in $[0;1]$.

This is how to detour the rocks, at least when second output's **Ref. feature size** is 0. When it is not, rough area detection does not correlate with the terrain's aspect and rock detouring is no longer exact.


Fast Perlin Fractal

This is a highly optimized version of the Simple Fractal node, based on a standard Value-Perlin noise with rotation. The number of settings in this fractal is limited in order to maximize efficiency of the node. It is very useful for all cases where you need a basic – but good quality – fractal pattern.





Roughness = 0.5




Roughness = 1.0

Variable Roughness Fractal


This is the same as the simple fractal, except that the roughness parameter can be made to vary according to the altitude of previous iterations of the base noise. This results in smooth areas at certain levels, and rougher areas away from this level.

Smooth level: this is the reference level for minimum roughness of the fractal. The roughness increases according to the distance to the smooth level.






Smooth level = -0.6




Smooth level = -0.2




Smooth level = 0.1






Smooth level = 0.5

Influence: this parameter controls the influence of the altitude on the roughness. If set to 0, the Variable Roughness Fractal behaves exactly as a Simple Fractal.




Influence = 0




Influence = 0.4






Influence = 0.6




Influence = 1

Local influence: this parameter controls how the roughness is computed according to altitude. If set to 0, the roughness is modulated by altitude only. If set to 1, the roughness will be modulated by the altitude of the last iteration of the noise, resulting in local patches of “smoothness” appearing at different altitudes.




Local influence = 0






Local influence = 0.4




Local influence = 0.6




Local influence = 1

Creep in: this parameter controls how much of the original roughness gets mixed back into the local roughness at each iteration. Higher values mean that the variable roughness only affect a few harmonics.






Creep in = 0




Creep in = 0.02



Creep in = 0.1





Creep in = 0.5

Variable Noise Fractal


This type of fractal is able to vary its base noise according to altitude. The first noise is used to compute the first iteration. Subsequent iterations are computed by blending the two types of noise according to altitude: the first type of noise will appear at lower altitudes, whereas the second type of noise will appear at higher altitudes.

Variable Noise Fractals create very subtle variations in the surface properties. They are however very slow to compute.


Noise 1: this is the same as the base noise setting common to all fractals.

Noise 2: this drop-down menu lets you select the second noise to be applied at higher altitudes. You can edit the noise properties by pressing the **Edit** button.


Switch level: this parameter is similar to the smooth level parameter of Variable Roughness fractals (see above). It controls the point at which the fractal switches its noise. If the altitude is below the switch level, the fractal will use the first noise. If the altitude is higher than the switch level, the fractal will use the second noise. Around the switch level, the two noises are blended according to the altitude.




Switch level = -1



Switch level = 0




Switch level = 0.2




Switch level = 0.5

Switch speed: controls the speed at which the fractal switches noise around the switch level.






Switch speed = 0.1




Switch speed = 0.2




Switch speed = 0.5






Switch speed = 1

Local influence: this parameter controls how the fractal decides which noise to use according to altitude. If set to 0, the current altitude is used. If set to 1, the fractal will base its decision solely on the last iteration of the noise, resulting in local patches of one noise appearing at different altitudes.




Local influence = 0




Local influence = 0.5





Local influence = 0.75



Local influence = 1




Three Noise Fractal


This is a complex fractal that mixes different noises according to the scale of the noise. It also lets you control the fractal's roughness in the same way as the Variable Roughness fractal.

Variable roughness: all the settings in this group behave as the Variable Roughness fractal's settings.


Turbulence damping: this setting controls the influence of the turbulence (origin shift) according to the harmonic. If set to 0, turbulence will be applied to all harmonics the same. The higher the value, the less harmonics that are affected by the turbulence – only large scale patterns are affected by the turbulence.








Damping = 0




Damping = 0.15



Damping = 0.4





Damping = 0.5


Mid-scale noise: this is the second noise to be used by the fractal when the scale becomes less than the change-over setting below.

Change-over scale: this is the scale below which the fractal switches its base noise to the mid-scale noise.

Small-scale noise: this is the third noise to be used by the fractal when the scale gets very small and becomes less than the small-scale change-over setting below.


Change-over scale: this is the scale below which the fractal switches its base noise to the small-scale noise.

In the examples below, we will use the following noises:




Large scale noise: “Water cress”






Mid scale noise: “Rectangular” Change over scale = 8




Small scale noise: “Sparse cracks” Change over scale = 2



Zoom 0: water cress noise is mostly visible






Zoom x8: some small rectangles from rectangular noise start appearing



Zoom x80: mostly patterns from the rectangular noise, but the sparse cracks are starting to show up



Zoom x300: at large zoom levels, the sparse cracks noise dominates



Open Ocean


The Open Ocean node is a simple simulation of open ocean water surfaces. It will create a nice simulation of the surface of the water, but it will not take into account any surrounding objects – hence the name. This node works best when used to produce the altitude function of a procedural terrain. By assigning a “World – Standard” mapping mode to the terrain and resizing the terrain in the Top view so that it fills up the entire world will yield very nice “infinite” ocean surfaces. Turn [here](#) for a sample tutorial on how this could be achieved.

This node is not a fractal per se, because the shape of the waves is different depending on the size of the wave, and there is a wave size under which the waves stop appearing (due to water surface tension). However, at larger scales, it does exhibit a somewhat fractal behavior, hence its classification in this category.

Unlike other fractal nodes, the Open Ocean node does not use a base noise to create water patterns. The “With rotation” and “Roughness” parameters do not exist either. Because this node takes all of its parameters into account to create a simulation that is as accurate as possible, the actual “roughness” of the water surface is controlled through other settings:


Wind direction: this parameter controls the direction in which the wind is blowing, as seen from above (the azimuth). A value of zero will make the wind blow from left to right in *Top view*. A value of 90° will make the wind blow from top to bottom in *Top view*. There is no relationship between this wind setting and the wind or breeze effects applied to plants.

Intensity: this parameter controls the intensity of the wind. Higher values will realistically lead to higher waves and rougher water surfaces.




Wind intensity = 0






Wind intensity = 0.7



Wind intensity = 1




Wind intensity = 2

Wave agitation tweak: this parameter lets you adjust the overall velocity of the waves created by the Open Ocean node. Its effects are only visible in animations. Values greater than 1 will make the waves move faster at the surface of the water, while values less than 1 will slow down the waves.



Foam output: when you connect to the Open Ocean node, you get the choice between Altitude and Foam outputs. The Foam output represents the typical foam density at the top of waves, and can be exported for use in the *Material Editor* to realistically distribute foam on the water.



Agitation tweak = 0.2 waves move slowly in animations




Agitation tweak = 2: waves move quickly in animations


Choppiness: this parameter controls the shape of the waves. Small values will yield soft round waves, whereas high values will produce choppy waves that are sharp at their top.




Choppiness = 0



Choppiness = 0.5



Choppiness = 0.7




Choppiness = 1


Gain: like with other fractal nodes, this parameter lets you adjust the altitude of the waves without interfering on the other settings of the simulation. It is generally recommended that you leave this value to the default value of 1 as this creates a realistic water simulation.



This parameter may however come in useful, for instance if you have resized the supporting procedural terrain vertically.



Gain = 0.3




Gain = 1



Gain = 2





Gain = 10

Color Nodes

Color nodes all output a color. Depending on the type of node, they either convert a number into a color (the color creation nodes), or convert one color into another color (the color correction nodes). The Color Map node (see below) can also output an alpha value.

Color Creation Nodes

Color Map

This node basically converts a number in between -1 and 1 into a color or an alpha value using a color map.

Color map: this defines the color map used for the conversion of the input value into a color and alpha. You can load a color map by double-clicking on the color map control, and you can edit the color map by selecting **Edit Color Map** from the popup menu. Please turn [here](#) for details on editing color maps.

If you connect another node to an output of this color map node, you will have the choice of connecting to its color output, or its alpha output.



2 Color Output

This node produces only two colors, according to the value of the input.

Color 1: if the input value is less than the threshold, the node outputs this color.

Color 2: if the input value is greater than the threshold, the node outputs this color.

Threshold: defines the value at which the output color switches from the first to the second color.

Linear Interpolation 2

This color node blends the two colors according to the input value.

Color 1: this defines the color output by the node when the input value is -1.

Color 2: this defines the color output by the node when the input value is +1.

Spline Interpolation 2

This color node blends the two colors according to the input value. This is similar to the previous node, except that the interpolation profile favors the extreme colors (you will see more of the actual 2 colors you defined than you would using the linear interpolation node).

Color 1: this defines the color output by the node when the input value is -1.

Color 2: this defines the color output by the node when the input value is +1.

3 Color Output

This is similar to the 2 color output node, only this node can output any one of three colors, according to the value of the input relative to the values of the 2 thresholds.

Color 1: if the input value is less than the first threshold, the node outputs this color.

Color 2: if the input value is greater than the first threshold, and less than the second threshold, the node outputs this color.

Color 3: if the input value is greater than the second threshold, the node outputs this color.

First threshold: defines the value at which the output color switches from the first to the second color.

Second threshold: defines the value at which the output color switches from the second to the third color.



Linear Interpolation 3

This color node blends the three colors according to the input value.

Color 1: this defines the color output by the node when the input value is -1.

Color 2: this defines the color output by the node when the input value is 0.

Color 3: this defines the color output by the node when the input value is +1.

Spline Interpolation 3

This color node blends the three colors according to the input value. This is similar to the previous node, except that the interpolation profile favors the extreme colors (you will see more of the actual first and third colors you defined than you would using the linear interpolation node).

Color 1: this defines the color output by the node when the input value is -1.

Color 2: this defines the color output by the node when the input value is 0.

Color 3: this defines the color output by the node when the input value is +1.

Color Variation

This creates variation in roughness and noise using one color only.

Color 1: click on the color box to select the color from the color chart.

Roughness: defines the rough areas in the color.

Contrast: defines the contrast of the changes made to the color.

Noise Variation: use the slider to increase or decrease noise variation.

Strength: defines the strength of the noise.

Noise Scale: defines the overall scale of the noise.

Color Brightness Variation

This takes the variation from the Color Variation node and adds brightness and saturation to the mix.

Color 1: click on the color box to select the color from the color chart.

Roughness: defines the rough areas in the color.

Contrast: defines the contrast of the changes made to the color.



Noise Variation: use the slider to increase or decrease noise variation.

Color Variation: this is the amount of variation in the color.

Brightness Variation: this adjusts the brightness of the color.

Saturation Variation: this increases or decreases the strength of the color.

Noise Scale: defines the overall scale of the noise.

Two Color Variation

This creates variation in roughness and noise using one color only.

Color 1: click on the color box to select the color from the color chart.

Color 2: select a second color to mix with the first color.

Roughness: defines the rough areas in the color.

Contrast: defines the contrast of the changes made to the color.

Noise Variation: use the slider to increase or decrease noise variation.

Noise Scale: defines the overall scale of the noise.

Natural Color Blend 2

This produces more natural or realistic color variations for terrains, ground, or any natural element in the scene, and provide an easy access to settings such as contrast, balance or roughness.

Color 1: click the color box to choose color 1 from the color chart.

Use color 2: check to use a second color.

Color 2: click the color box to choose color 2 from the color chart.

There are controls for you to manipulate the color(s).

Noise Scale: defines the overall scale of the noise. This should typically remain high (2000 by default) for a terrain or ground.

Roughness: defines the rough areas in the color.

Contrast: defines the contrast of the changes made to the color.

Balance: defines the balance of the two colors.



Distortion: defines any distortion, if any.

Noise Variation: use the slider to increase or decrease noise variation.

'*Alpha grain*': this should be enabled only if the current material layer is not a leaf layer. In this case, it automatically adds ALPHA to produce the same kind of pattern as the color noise.

Color Variation Map

This node takes advantage of color gradient maps to produce a mixed material.

Color Map

Color Map: click the color box to select a gradient map to use.

Other

Roughness: defines rough areas in the color.

Contrast: defines the contrast of the changes in the color.

Noise Variation: use the slider to increase or decrease noise variation.

Noise Scale: set the scale of the noise pattern

Terrain Color Patterns

This node is a fractal function designed to create color patterns, typically for use in the *Material Editor's* functions. It produces a mix of smooth and rough variations of colors similar to the distribution of rocks on a sedimentary soil. It is in fact based on the new algorithm developed for Terrain Fractal 2.

The **Terrain Color Patterns** node provides several groups of parameters.

The **Roughness aspect** parameters are very much inspired from the “Ground aspect” parameters of the “Terrain Fractal 2” node.

- **Bump surge:** controls the contrast between rougher and smoother areas.
- **Roughness abundance:** controls the overall quantity of roughness present in the patterns.
- **Smooth Threshold:** controls the smoothness of the transition strip.
- **Roughness dispersion:** controls how much the roughness tends to be scattered in the patterns rather than gathered in specific areas.
- **Roughness trend at higher frequency:** controls whether roughness increases or



decreases at higher frequency.

The **Strata processing** parameters are similar to those available on the “Strata” filter located in the “Recursive” filters subcategory:

- **Processing strength:** controls the influence of the strata filtering over the color patterns.
- **Layer spacing:** controls the height of the main layer.
- **Offset:** allows fine-tuning of the vertical strata pattern positioning with respect to the underlying patterns.

Color Correction Nodes

The color correction nodes apply modifications to the color that is passed to them as input.

Common Settings

The following setting is common to quite a few color correction nodes:

Allow luminous colors: when this option is checked, VUE will generate colors that are brighter than white. Such colors actually generate light; they can produce very interesting lighting effects when used in conjunction with a radiosity rendering. If you don’t check this option, colors will be clamped at white.

Gamma

This color correction node lets you adjust the gamma setting for the color.

Gamma: the gamma color correction parameter. Higher gamma values will produce darker, more saturated colors.

Gain

This color correction node lets you adjust the gain setting for the color. Higher gain values boost the contrast of medium brightness colors.

Gain: the gain color correction parameter.

Brightness

This color correction node lets you adjust the brightness setting for the color.

Brightness: the brightness color correction parameter.



Contrast

This color correction node lets you adjust the contrast setting for the color.

Contrast: the contrast color correction parameter.

HLS Shift

HLS stands for Hue-Luminosity-Saturation. It is another way of working with colors than the standard Red-Green-Blue paradigm.

This color correction node lets you adjust the overall brightness (luminosity), color tone (hue) and saturation of your colors.

Hue shift: this parameter controls the amount of shifting applied to the color's hue.

Luminosity shift: this parameter controls the amount of shifting applied to the color's luminosity – in effect, a brightness setting.

Saturation shift: this parameter controls the amount of shifting applied to the color's saturation. Strong negative values will convert all colors to black and white (desaturation).

HLS Color Shift

This color correction node is similar to the previous one, in the sense that it also lets you adjust the overall brightness, color tone and saturation of your colors – however, adjustment is done via a color instead of independent parameters.

HLS shift color: this color is used to define the color correction that is applied to the colors. The default color is a pale shade of blue. If you edit this color, you will notice that it corresponds to a medium setting (128) for hue, luminosity and saturation. If you modify the hue for this HLS shift color, this modification of hue will be applied to all the colors passed to the node. In the same way, if you modify the luminosity or saturation, this modification will be applied to the colors passed to the node.

Filter

The filter color correction node multiplies all colors by the filter color – as if colors were seen through a colored gel.

Filter color: this parameter defines the color applied as a filter. Double-click to edit the color.



Perspective

The perspective color correction node replaces darker colors with the perspective color.

Perspective color: this parameter defines the color used to replace the darker colors. Double-click to edit the color.

Color Blender

The Color Blender node will blend the input color with a solid color.

Blending color: this is the color that will be blended in with the input color.

Blending ratio: this controls the amount of blending that takes place between the input color and the blending color. The higher the value, the more the blending color affects the input color.

Color mask: if this option is selected, the color is applied in replacement of the input color. When the blending ratio is set to 0%, the color is applied as a mask. When set at 100% the color completely replaces the bitmap. If this option is not selected, the blending color is applied in product (as a filter).

Texture Map Nodes

Mapping Nodes

Texture Map

The Texture Map node is used to map a picture onto objects. Its input is a texture coordinate, and the texture map node returns the color of the pixel in the texture map that is at the point indicated by the texture coordinate.

When you create a Texture Map node, *SmartGraph* automatically creates a “UV Coordinates” node and connects it to the node’s input. You can use the “UV Coordinates” node to define how the picture is mapped onto the object (see below for details on the “UV Coordinates” node).

Texture map nodes can be made to output any of the following values:

- **Color output:** the color of the pixel in the texture map that is at the point indicated by the texture coordinates.



- **Grayscale output:** the color of the pixel converted to a grayscale value.
- **Alpha output:** the alpha value corresponding to that pixel (if the point is inside the image, or 0 if it is outside the image).

When you connect a node to the output, a popup menu will appear so that you can select the desired type of output.




Image: double-click on the image preview to load a new picture to be used as texture map. You can rotate the picture by increments of 90° using the  and  buttons. If you need to invert the values, you can do so by pressing the  button.

Image offset: the image can be positioned precisely by using these parameters. These will shift the image around the origin (in pixel units).

Interpolation type: interpolation is used to reduce the pixelization effect when the texture map is seen from very close and the resolution of the map is insufficient. These interpolation methods are the same as that of the mapped picture material coloring mode (see [here](#) for details).

Tiling mode X: this is a drop-down list that lets you select the way the image is repeated along the X axis. Possible values are the same as that of the mapped picture material coloring mode (see [here](#) for details).

Tiling mode Y: this is the same as the above, only along the Y axis.

Note:

images mapped using the Texture map node are not mip-mapped. If you would like to enable mip-mapping for this map, you should use the Projected Texture Map node instead.

Projected Texture Map

The Projected Texture Map node is used to map a picture onto objects. It uses the coordinates of the current point to determine the color of the pixel in the texture map at that point. This node effectively combines the features of the Texture Map node (see above) and the UV Coordinates node (see [here](#)). Please refer to these nodes for a description of the settings available in the Projected Texture Map node. One additional setting is specific to the Projected Texture Map node: the “Allow mip-mapping” option.

Allow mip-mapping: mip-mapping is a process whereby the software automatically generates lower resolution versions of the image and uses them instead of the full-blown image as soon as it is seen from a distance. While the results produced using mip-mapping are generally smoother, certain images may actually look better without mip-mapping. This option is here so that you can forbid mip-mapping for a specific image, should you need to (just uncheck the option).



Note:

You can control the level of mip-mapping for the entire scene using the “Texture filtering” option in the Anti-Aliasing Options dialog (see [here](#)).

Images mapped using the Projected Texture Map node will be mip-mapped according to global scene settings and the “Allow mip-mapping” option.

Mapping position mode: options for this setting are: Automatic, World-Standard, World-Parametric, Object-Standard and Object-Parametric.

For **Projected Texture Maps**, a *Texture Placement Editor* is available for manipulation of the texture directly on the object. For more information about this *Texture Placement Editor*, please refer [here](#).

Animation Map



The Animation Map node is used to map an animated texture onto objects. Its input is a texture coordinate and a time, and the animation map node returns the color of the pixel in the current frame of the animated texture map that is at the point indicated by the texture coordinate.

When you create an Animation Map node, *SmartGraph* automatically creates a “UV Coordinates” node and connects it to the node’s input. You can use the “UV Coordinates” node to define how the animation is mapped onto the object (see below for details on the “UV Coordinates” node).

Texture map nodes can be made to output any of the following values:

- **Color output:** the color of the pixel in the texture map that is at the point indicated by the texture coordinate.
- **Grayscale output:** the color of the pixel converted to a grayscale value.
- **Alpha output:** the alpha value corresponding to that pixel (if the point is inside the image, or 0 if it is outside the image).

When you connect a node to the output, a popup menu will appear so that you can select the desired type of output.

Image sequence: this is the list of pictures to use in the animation. You can add new pictures by clicking the **Load** icon (). You can replace pictures in the list by selecting them and then pressing Load. To remove images from the list, select them and then press the **Remove** icon (.

Frame rate: this defines the playback rate of the pictures on the list. Ideally, this should at least be equal to the global animation frame rate.



Interpolate frames: when this option is selected, in-between frames are interpolated by gradually blending the previous and the next frames. This ensures smooth playback and will avoid any jumps in the animated texture.

Animation filter: use this filter to change the flow of time in the animated texture. Double-click on the filter to load a filter, or select **Edit** from the filter's popup menu to edit the filter.

Phase: use this to adjust the start frame in the animation sequence. The value has to be set in seconds.

Image offset, Interpolation type and Mirror X & Y are identical to the settings in the **Color** tab of the *Advanced Material Editor* (see [here](#)).

Warning: if several frames of the animation are required to render the texture correctly at a given time (e.g. after connecting the phase to a noise), memory requirements may increase and rendering may slow down significantly.

The Animation Map node doesn't support mip-mapping. If you would like your animation map to be mip-mapped, please use the Projected Animation Map node below instead.

Projected Animation Map Node

The Projected Animation Map node is used to map an animation onto objects. It uses the coordinates of the current point and the time input to determine the color of the corresponding pixel in the appropriate frame of the texture map. This node effectively combines the features of the Animation Map node (see above) and the UV Coordinates node (see [here](#)). Please refer to these nodes for a description of the settings available in the Projected Animation Map node. One additional setting is specific to the Projected Animation Map node: the "Allow mip-mapping" option.

Allow mip-mapping: mip-mapping is a process whereby the software automatically generates lower resolution versions of the image and uses them instead of the full-blown image as soon as it is seen from a distance. While the results produced using mip-mapping are generally smoother, certain images may actually look better without mip-mapping. This option is here so that you can forbid mip-mapping for a specific image, should you need to (just uncheck the option).

Note:

You can control the level of mip-mapping for the entire scene using the "Texture filtering" option in the *Anti-Aliasing Options* dialog (see [here](#)).

Images mapped using the Projected Animation Map node will be mip-mapped according to global scene settings and the "Allow mip-mapping" option.



Mapping position mode: this setting is available for this node as well as the **Projected Texture Map** node. Options for this setting are: Automatic, World – Standard, World – Parametric, Object – Standard and Object – Parametric.

Blended Image Node

This node is similar to the Texture Map Node, except that it blends the image into an existing color input, using a smooth blending strip. Outside the image, the input color remains unaffected. Inside the image, the input color is replaced by the image. If the image defines an alpha channel, this value will be used in the blending ratio.

This node outputs the following values:

- **Color output:** the color of the pixel in the texture map that is at the point indicated by the texture coordinate.
- **Grayscale output:** the color of the pixel converted to a grayscale value.
- **Alpha output:** the alpha value corresponding to that pixel (if the point is inside the image, or 0 if it is outside the image).
- **Blend ratio:** the proportion of the input color that was replaced by the image according to the blend profile and position in the image (not taking into account the image's alpha value).

The settings available for this node are the following:

Image offset and **Interpolation type** are identical to the settings in the Texture Map node.

Smooth blending strip lets you define how gradual the blending is. A value of 0 means that the image replaces the input color as soon as the point is inside the image. A value of 100% means that the image fully replaces the input color solely at the exact center of the image.

Blend profile: this setting controls how the blending is done. Possible values are:

- **Square:** the blend ratio is defined according to the distance to the nearest edge of the image.
- **Round:** the blend ratio is defined according to the distance from the center of the image.

Blended Grayscale Image Node

This node is identical to the “Blended Image” node, with the difference that it acts on a number instead of a color (the input value is a number instead of a color). This number is replaced by the grayscale value of the image at the current point, according to the same rules as with the “Blended Image” node. This is particularly useful when designing procedural terrain functions, and you want to add real-world data at some point: simply



use a “Blended Grayscale Image” node to replace the procedural altitudes with a DEM file at the desired location. Thanks to the smooth blending strip, the procedural altitudes will automatically blend into the DEM altitudes.

On top of the “Blended Image” node parameters, this node defines the following additional parameters:

Gain: this is a gain factor that is applied to the grayscale values in the image (in order to adapt to the range of input values the range of values defined by the image).

Offset: this is an offset that is applied to the grayscale values in the image.

Unlike the “Blended Image” node, this node does not define a “Color” output.

Image Sample and Multi-Image Sample Nodes

These texture map nodes are used with the Image Combiner node to create different texture effects. These need to be processed through the Image Combiner node or the transparency information won’t be processed correctly.

Each image sample node has the following settings, reflecting similar settings in the *Material Editor*. The image and pathname are displayed and various types of **UV Coordinates** can be selected:

- **Automatic**
- **Flat (vertical)**
- **Faces**
- **Cylindrical**
- **Spherical**
- **Torical**
- **Conical**
- **Automatic UV**

You can position the picture precisely on the object by using the **Image offset** commands. This will move the picture around by increments of one pixel.

When the material is seen from very close, you may see pixels, due to the limited resolution of the picture. To reduce this effect, choose an **Interpolation type** method:

- **None:** No over sampling.
- **Bi-linear:** Bi-linear interpolation between pixels.
- **Normalized:** Values proportional to the distance to the corners of the pixel.
- **Bi-cubic:** Bi-cubic interpolation between pixels (continuous derivative).



Density controls how many times the image is repeated.

In the **Rotation** section, you can select to rotate image samples in range and set the range using the slider.

You can also opt to **Flip** the image horizontally and/or vertically.

In the **Scale** section, you can indicate the **Global Sample Scale** of the picture along the X and Y axes with options to scale certain ranges on the X, Y axes. There is also an option to **Keep proportions** with scaling.

Use the **Image Sample** node to create even more special effects with image-based textures.

Ptex node

This node allows to extract two channels for a ptx file. Alpha channel (flagged inside the ptx file if present) can be retrieve with the alpha output. The three first other channel are merge into a color output.

xStream

Note:

For this node to work properly, the underlying geometry of the material should match exactly the geometry exported with the ptx file.

UV Coordinates Node

This node converts the current position into a texture coordinate. It is automatically created when you create a texture or animation map node.

Scale: defines the overall size of the texture map along its two axes.

Origin: defines the point of origin of the projection – e.g., when mapping in spherical coordinates, defines the center of the sphere.

Mapping mode: this setting defines the method used by the node to convert 3D coordinates into the 2D texture map coordinates. There are several mapping modes available, each of them better suited for some types of objects. If you don't know which to use, select **Automatic**. For details on the different mapping modes, please refer to the *Material Editor* section about material colors being mapped from a picture (see [here](#)).



Filter Nodes

Environment Sensitive Filters

Environment sensitive filters are able to adapt their response according to the local altitude, slope and orientation.

Altitude

The Altitude filter modulates its response according to altitude. For points at low altitudes, the filter output will 0.

Influence: this setting controls the percentage of the input signal that is modulated according to altitude.

Min effect altitude: this setting controls the altitude below which the response of the filter is uniformly 0.

Max effect altitude: this setting controls the altitude above which the filter's output is identical to the input. In between the two altitudes, the response of the filter is a blend of the two outputs.

If the two altitudes are inverted (i.e. max effect is actually lower than min effect), the behavior of the filter will be inverted (i.e. the filter will output 0 at all altitudes greater than min effect).

Slope

The Slope filter modulates its response according to slope. For points at low altitudes, the filter output will 0.

Influence: this setting controls the percentage of the input signal that is modulated according to slope.

Min effect slope: this setting controls the slope below which the response of the filter is uniformly 0.

Max effect slope: this setting controls the slope above which the filter's output is identical to the input. In between the two slopes, the response of the filter is a blend of the two outputs.

If the two slopes are inverted (i.e. max effect is actually lower than min effect), the



behavior of the filter will be inverted (i.e. the filter will output 0 at all slopes greater than min effect).

Altitude and Slope

This filter is a combination of the two above filters. It modulates its response according to the altitude and the slope.

The **Min** and **Max effect** settings are identical to the two previous filters. The Altitude and Slope filter also lets you adjust the relative influence of altitude and slope on the filter's response through the use of the **Importance** settings. The higher the importance of altitude, the stronger the influence the altitude will have on the filter's output. Ditto for slope.

Orientation

The Orientation filter modulates its response according to orientation of the surface on which the function is being computed. For points of the surface aiming in the opposite direction to the favored azimuth, the filter output will 0.

Favored azimuth: this parameter controls the azimuth of the direction in which the response of the filter will be unmodified. As the surface points away from this favored direction, the response of the filter gradually decreases until it reaches 0.

Tightness: this parameter controls the angular tolerance around the favored azimuth. If the tightness is 0, all points that are less than 90° away from the favored azimuth will get some filter response. Points that look in the opposite direction will get 0 response.

Transition speed: this parameter controls the speed at which the filter transitions from no response to full response as the surface points more towards the favored azimuth.

Environment

The environment filter is a combination of the orientation filter and the altitude and slope filter.

The parameters in the Environment filter are identical to those of these two filters (see above for a description of these parameters).

Patches

The Patches filter is a very special filter that automatically creates uniform patches on horizontal surfaces. The filter can output two values:

- **Patch value:** this is the standard filter's output,
- **Presence on patch:** this output is 1 if the current point is in a patch, and 0



otherwise.

When you connect a link to the filter's output, a menu will appear to let you select the desired output.

Altitude and slope settings: all the settings in the Altitude and Slope groups are identical to those in the Altitude and Slope filter.

Patch size: this parameter controls the average size of the patches.

Patch height: this parameter controls the average difference in height between areas that are on the patches, and areas that are outside the patches.

Noisiness: this parameter controls how uniform the edges of the patches are. Higher values mean that the patch edges are defined according to the variations in the underlying signal.

Transition speed: this parameter controls the speed at which the filter transitions from outside onto the inside of a patch. It affects the steepness of the patch edges.

Surface noise: this parameter controls the amount of underlying noise that remains at the surface of the patches.



Recursive Nodes




Diagram – Strata node parameters

The **Recursive** strata filters generate steps of a given size and orientation. The node is applied recursively until the setting limits are reached, exactly like a fractal.

The filter uses a pattern, repeated as many times as needed, spanning the whole (potentially tilted) Z axis in the standard **Strata** filter or restricted to the confining range for the **Confined Strata** filter. At each iteration, the filter is applied on the result of the previous iteration, with all distance parameters halved (i.e. spacing and thickness).

Though the parameter names are semantically related to the concept of rock strata in a terrain, the filter can, of course, be applied to any kind of scalar input to generate a complex banded pattern. One could also stratify a positions' coordinates before calling a fractal, to have a result along the x and/or y axis similar to what is obtained along the vertical axis when this filter is applied after a fractal's input.

The pattern is made of original input (marked **(1)** on diagram):

A rock layer, itself separated in two parts (marked **(2)** on diagram):

- Step from bottom to top of the layer (marked **(3)** on diagram)



- Plateau from the top of the layer to the invariant layer (marked (4) on the diagram)
An invariant layer, on which the filter does nothing.

Strata Processing Data Parameters

Processing strength: This indicates how much of the effect is actually taken into account in the output result. For example, if the value in this field is equal to .5, the output is half the input and half the stratified altitude.

Rock layer hardness (marked s on the diagram): The harder the layer, the steeper the filter's step between the bottom and the top of each rock layer.

Rock layer thickness (marked H on the diagram): The thickness can be smaller than the spacing between layers. Values higher than layer spacing are clipped.

Layer spacing: This is the distance between the repetition of two filtering patterns (marked SP on the diagram).

Plateau filling (marked f on the diagram): This option controls the slope of the plateau. At 0, the plateau is parallel to the underlying planes, whereas as the value rises, the plateau is raised, taking over more and more in the invariant range, which is correspondingly reduced.

Smooth edges (marked e on the diagram): This is the range which the filtered input is smoothed to avoid too sharp edges on the output result. To avoid cluttering, the figure only shows smoothing on the transition -> plateau edge (2) to (3), but the other two edges are smoothed as well (3) to (4), i.e. plateau -> unchanged range, and (4) to (2), unchanged range -> transition.

Smallest feature: This parameter tells how fine grained the filtering needs to be, i.e. how many iterations will occur.

These parameters are not measured directly on the input noise since the strata can be tilted. Rather, they are measured along the axis perpendicular to the potentially tilted strata planes.

Strata Positioning Features

Strata can be viewed as planes cutting through the terrain (or any other object). Parameters are:

Tilt heading (degrees): This defines the orientation of the axis around which the strata/planes will be tilted. In the diagram, this is represented in a vertical plane perpendicular to this axis.

Tilt angle (degrees): This defines the angle of rotation of the strata/planes.



Offset: This is an offset between the bottom of the “first” rock layer and the origin. Not shown on the diagram, this value would offset the point at which the strata crosses the origin R alone one the axis perpendicular to the strata, like the one marked **(5)**. It means the whole pattern of stratification (iterations included) is offset. Simulating strata deformation is then possible, by having the offset depend on **X** and **Y**.

2nd Output: Detect Rough Areas

Like a fractal, this recursive filter is capable of outputting some measure of the variation induced on the input noise. The output will vary roughly in $[-1:1]$. Here is the meaning of the value output when only one octave is considered:

At -1, the input was not affected at all by the stratification process (range **(4)** on the diagram, or outside the confinement area, in the case of confined stratification).

At 0, the input landed on a plateau (range **(3)** on the diagram).

At 1, the input landed on a transition (range **(2)** on the diagram).

When several octaves of stratification are included in the second output computation, the values at each octave are summed, with coefficients for each octave depending on the value of the parameters.

Confined Strata

The **Confined Strata** filter requires another set of parameters for stratification altitude.

Don't stratify below: This is the bottom of the stratification range.

Don't stratify above: This is the top of the stratification range.

Fade in/fade out height: This is the distance along which the stratification area is faded in/out inside the stratification range. This smoothes the transition with the unfiltered range.

Origin is also added to **Strata positioning**, replacing the **Offset** parameter of the **Strata** function. With **Confined Strata**, a 3D offset is more practical. For example, when the strata is tilted, the confined range will only cross the input range in a specific area. To allow for precise control of this area's positioning, the origin must be fully customizable.

3D Stratification

This node uses as input a vector (like a position), and the result on the output position will be the same as applying a “Strata” filter node on each of the coordinates of the vector. It is easier to use because no composer/decomposer node is required, and the parameters



for all three strata filters are gathered in a single, compact interface.

Since it filters a full 3D vector, a good place to insert it in a graph is between the input position and the fractal or noise node on which the effect is desired.

Input Filters

These filters are designed to modify the profile of the input values according to a simple filtering rule.

Filter

This filter uses a standard Filter control to determine the output profile.

Filter: this is the filter that is used to determine the output profile. Double-click on the filter control to load a new filter, or select **Edit** from the popup menu to edit the filter. Please turn [here](#) for details on editing filters.

Partial Filter

This is similar to the previous filter, except that you can modulate the amount of the signal that is actually filtered through the filter.

Filter: see above.

Filter ratio: this parameter controls the level of filtering of the signal. If set to 0, the output is unfiltered. If set to 1, the output is identical to the above Filter node. If set to 0.5, half of the signal will be filtered, and the other half will remain unfiltered.

Offset ($X + a$)

This is a very simple filter that adds an offset to the input signal.

Offset: this parameter controls the amount that is added to the filter's input.

Opposite ($-X$)

This filter simply returns the opposite of the input signal.

Multiply (aX)

This filter simply multiplies the input signal by a value.

Multiply by: this parameter controls the amount by which the input signal is multiplied.



Divide (a/X)

This filter simply divides the input signal by a value.

Divide by: this parameter controls the amount by which the input signal is divided.

Brightness-Contrast ($aX + b$)

This filter combines the effects of the Offset and Multiply filters into a single, convenient filter.

Brightness: this parameter controls the amount that is added to the input signal (in effect, this acts as a brightness setting).

Contrast: this parameter controls the amount by which the input signal is multiplied (in effect, this acts as a contrast setting).

Parabolic ($aX^2 + bX + c$)

This is a slightly more complex filter that creates a parabolic output profile.

a, b and c: represent the different terms used in the parabolic equation $aX^2 + bX + c$.

Absolute

This filter simply mirrors the input value around the threshold value.

Contrast: this parameter controls the amount by which the input signal is multiplied (in effect, this acts as a contrast setting).

Threshold: this parameter controls the value at which the input is reversed. As a result, the output value can never drop beneath this threshold.

Gamma

This filter applies a gamma correction to the input signal.

Gamma: this parameter controls the gamma correction applied to the input signal.

Bias

This filter applies a bias correction to the input signal.

Bias: this parameter controls the bias correction applied to the input signal.



Gain

This filter applies a gain to the input signal.

Gain: this parameter controls the gain applied to the input signal.

Power

This filter calculates the difference between the input value and a lower clip value, and raises it to a given exponent.

Exponent: this parameter controls the exponent applied to the input value.

Lower clip: this parameter controls the value below which the filter's output is uniformly 0. Above this value, the filter's output is the difference between the input value and this value, raised to the power of the exponent.

Gaussian

This filter passes the input signal through a Gaussian curve, in effect producing a response similar to a smoother version of the Absolute filter described above.

Contrast: this parameter controls the amount of contrast in the resulting output.

Threshold: this parameter controls the lower limit around which the signal is “mirrored” by the Gaussian profile.

Floor

This filter clamps any value below the **Floor** value to that value.

Ceiling

This filter clamps any value over the **Ceiling** value to that value.

Clamp

This filter lets you clamp the input signal to a given range.

Lower clip: this parameter controls the lower limit of the range to which the signal is clamped. Any input below this value will result in an output equal to this value.

Upper clip: this parameter controls the upper limit of the range to which the signal is clamped. Any input above this value will result in an output equal to this value.



Clip

The Clip filter combines the effects of the Brightness-Contrast filter with the effect of the Clamp filter.

Contrast and Brightness: identical to the Brightness-Contrast filter settings.

Lower and Upper clip: identical to the Clamp filter.

Smooth Clip

The Smooth Clip filter is identical to the Clip filter described above, except that the output values are smoothed around the extremes, in order to avoid sharp variations in slope near the lower or upper clip values. In effect, this filter produces a slightly more contrasted result as the standard clip.

Map

The Map filter maps a given input range of values to a given output range. When connecting a parameter to another node, this filter is particularly useful to adapt the range of the signal to the range of values expected by the parameter.

Lower input value: this parameter controls the lower limit of the expected input range.

Upper input value: this parameter controls the upper limit of the expected input range.

Lower output value: this parameter controls the lower output value. This value is achieved when the input value is equal to the lower input value.

Upper output value: this parameter controls the upper output value. This value is achieved when the input value is equal to the upper input value.

Clip out of range values: if this option is selected, values that are out of the input range will be clipped to the input range (similar in effect to applying a clamp filter on this filter's input).

If the upper output value is less than the lower output value, the signal will be inverted.

Smooth Map

The Smooth Map filter is identical to the Map filter described above, except that the output values are smoothed around the extremes, in order to avoid sharp variations in slope near the lower or upper input values. In effect, this filter produces a slightly more contrasted result as the standard map. Values that are beyond the input range are automatically clipped to the input range.



Quantize

The Quantize filter converts the input into a range of discrete values.

Steps: this parameter controls the number of different values that the filter can output. For instance, if set to 5, the output will be quantized to 5 different possible values.

Contrast and Brightness: these settings are the same as those of the Brightness-Contrast filter described above.

Saw Wave

The Saw Wave filter is equivalent to the fractional part of the input signal in the range of -1 through 1. When the signal reaches 1, it jumps back down to -1, creating a saw teeth like pattern.

Contrast and Brightness: these settings are the same as those of the Brightness-Contrast filter described above. Whenever the result of the brightness-contrast transformation exceeds 1, it jumps back down to -1.

Absolute Wave

The Absolute Wave filter is very similar to the Saw Wave filter, with the exception that the parts of the signal that are out of range are mirrored back instead of jumping back down. As a result, the Absolute Wave filter creates both up and down slopes, whereas the Saw Wave never inverts the slopes.

Contrast and Brightness: these settings are the same as those of the Brightness-Contrast filter described above. Whenever the result of the brightness-contrast transformation exceeds 1, it is mirrored back down.

Sine Wave

In effect, very similar to the Absolute Wave filter, except that this filter avoids the sharp changes in slope around the upper and lower limits. This version is usually preferred when the output is used to generate bumps.

Contrast and Brightness: these settings are the same as those of the Brightness-Contrast filter described above. Whenever the result of the brightness-contrast transformation exceeds 1, it is mirrored back down.

Threshold

This filter switches between two values depending on the input: if the input is less than **Threshold**, the node outputs the **Low value**. If it is greater, the node outputs the **High value**.



Smooth Threshold

This is similar to the **Threshold** filter, with the addition of a smooth transition strip, defined by the **Transition** parameter. Inside the transition strip, the node outputs a blend of both the **Low** and **High values**.

Constant Nodes

Constant nodes do not take any inputs. They output the value that is defined by the node.

Constant Number

Value: use this setting to define the number that is output by the constant node.

Constant Color

Color: use this setting to define the color that is output by the constant node. Double-click on the color to edit it.

Constant Coordinates

Value: use this setting to define the texture map coordinates that are output by the constant node.

Constant Vector

Value: use this setting to define the vector that is output by the constant node.

Random Constant Number

Value: use this to create a random seed for things like procedural materials.

Connectable Constant

Connectable constants are identical to regular constants, except that their value can be extracted. What is the point of extracting the value of a constant, you may ask? Indeed, in standard graphs, there is no point in doing so. However, connectable constants are very useful in the context of published MetaNode parameters, where a “published” connectable constant can be connected to other nodes at the higher MetaNode interface level (see [here](#)).

Note:



Connectable constants can also be used to add a “name label” to intermediate values inside a graph, so as to improve overall readability of the graph.

Turbulence Nodes

Turbulence nodes are very similar to fractal nodes, with the main difference being that turbulence nodes work in 3 dimensions to create vector displacements, whereas fractal nodes only work in one dimension.


Although turbulence nodes should be applied to the **Origin** of noises or fractals in order to produce the expected results, you can achieve interesting results by using turbulence on other parameters.

Turbulence will add interesting details to your functions, but this is at the expense of long processing times: in order to generate the turbulence, Vue has to compute several iterations of the noise along the 3 different axes, resulting in the long computation times.

Simple Turbulence


The simple turbulence node uses a Perlin style noise to generate a 3D perturbation. The following settings are available:

Wavelength, **Origin**, and **Largest feature**: these settings are the same as with the standard fractal nodes. Please turn [here](#) for details on the fractal nodes.




Largest feature = 0.1






Largest feature = 1




Largest feature = 5




Largest feature = 10

Amplitude: this parameter defines the amplitude of the perturbation created by the turbulence node. The stronger the setting, the more perturbed the signal to which is applied this turbulence.






Amplitude = 0.5




Amplitude = 1




Amplitude = 2






Amplitude = 5

Repeat count: this parameter defines the number of iterations of the base noise that are computed to generate the turbulence. Higher repeat counts will create more detailed turbulence, only at the expense of longer render times.




Repeat count = 1




Repeat count = 2






Repeat count = 4




Repeat count = 10

Scaling: this setting controls the frequency at which the noise varies relative to the current position.




Scaling = 0.5






Scaling = 1




Scaling = 2




Scaling = 5

Harmonics: this setting controls the way the noise is scaled each time an new iteration is added in: for each new addition, scale and amplitude are multiplied by the Harmonics parameter. If the “Repeat count” is equal to 1, this parameter has no influence. You should avoid values close to 1 as they tend to reduce the influence of additional iterations.






Harmonics = 0.25




Harmonics = 0.5




Harmonics = 0.9






Harmonics = 2

Combination mode: this drop-down list box lets you select how the successive noise iterations will be combined. For full details on combination modes, refer to the “Basic Repeater” fractal node (see [here](#)).




Combination mode = Add/Blend




Combination mode = Variable roughness






Combination mode = Variable roughness (abs)




Combination mode = Max




Combination mode = Max (abs)






Combination mode = Min



Combination mode = Min (abs)




Combination mode = Multiply

Misc Style Turbulence


This turbulence node is provided for compatibility with previous versions of Vue. It provides you with more control over the look of the turbulence but usually doesn't produce such nice results...



On top of the settings already defined by the Simple Turbulence node, this node lets you select the base **Noise** used to compute the turbulence.



Perlin - Gradient (default)




Perlin - Value




Perlin - Linear





Distributed – Round Samples




Cellular – Chipped




Cellular – Crystals






Cellular – Drought




Cellular – Pebble Noise




Line Patterns – Lines






Line Patterns – Sparse Cracks




Other – Granite




Other – Water Rough






Squares – Random Altitudes




Squares – Blobs



Squares – Stones Round






Squares – Stones Square


Advanced Turbulence

On top of the settings already defined by the other turbulence nodes, this node defines the following:


Roughness: this is similar to the standard fractal “Roughness” parameter.




Roughness = 0.25



Roughness = 0.5




Roughness = 0.75




Roughness = 1

With vortices: check this option if you want the turbulence to exhibit vortices.




Without vortices






With vortices



Without vortices



With vortices



Combiner Nodes

The combiner nodes take several inputs and combine them together into a single output. Most combiner nodes accept any type of input, with the exception of the Color combiner that only operates on colors, and the Combiner that only operates on numbers.

Combiner nodes that accept any type of data must receive the same type of data on all their inputs. This is why setting the first input locks the data type for other inputs.

Blender

The Blender node accepts two inputs and combines them together according to the combination mode and the ratio.

Ratio: this parameter controls the relative importance of each one of the inputs in the final node's output. Small values will favor the first input, whereas larger values will favor the second input.

Combination mode: this drop-down list defines the method used to combine the two inputs together:

- **Blend:** values are averaged,
- **Add:** values are added together,
- **Max:** the biggest value is retained,
- **Min:** the smallest value is retained,
- **Subtract:** the value of the second input is subtracted from the first input,
- **Multiply:** values are multiplied together.

Combiner

The combiner node can only operate on numbers. It can however combine an unlimited number of inputs. By default, the combiner node is created with 2 inputs, but as soon as you connect all inputs, a new input is added.

The combiner node combines inputs according to the overall combination mode as well as “per input” settings. The Node Details area displays settings that are relative each one of the inputs. If more inputs are added, new settings are added accordingly.

Combination mode: this drop-down list box lets you select how the different inputs will be combined. The different combination modes are the same as for the Basic Repeater fractal node (turn [here](#) for details).



Amplitude: this parameter controls the relative amplitude of each one of the inputs. The input is multiplied by the value of this amplitude parameter and offset according to the offset parameter below before being combined with the other inputs.

Offset: this parameter controls the relative offset of each one of the inputs. It is used together with the amplitude setting before the input is combined with the other inputs.

Color Combiner

The color combiner node only works with colors. It is capable of combining an unlimited number of colors according to a combination mode.

Combination mode: this drop-down list box lets you select the way the input colors will be combined:

- **Blend:** the colors are mixed in equal proportions.
- **Add:** the colors from different inputs are added together. The resulting color is necessarily brighter than the each one of the input colors.
- **Subtract:** successive input colors are subtracted from the first color, and clipped to black. The resulting color is necessarily darker than each one of the input colors.
- **Multiply:** all colors are multiplied together. The dark areas in each one of the inputs will be dark in the final color, and white areas will be the same as the other colors.
- **Divide:** successive colors are divided. Results can be unexpected...
- **Min:** the final color is the minimum of each color component, and thus necessarily darker than any one of the inputs.
- **Soft min:** this is the same as minimum, except that the color values are blended when they are close.
- **Max:** the final color is the maximum of each color component, and thus necessarily brighter than any one of the inputs.
- **Soft max:** this is the same as maximum, except that the values are blended when they are close.
- **Red filter:** the first input color is multiplied by the red component of all successive inputs.
- **Green filter:** the first input color is multiplied by the green component of all successive inputs.
- **Blue filter:** the first input color is multiplied by the blue component of all successive inputs.
- **Luminosity value:** the first input color is multiplied by the luminosity of the



successive colors.

- **Hue blend:** the hues of the different colors are blended. The saturation and luminosity of the first input color are retained in the final output.
- **Luminosity blend:** the luminosity values of the different colors are blended. The saturation and hue of the first input color are retained in the final output.
- **Saturation blend:** the saturation value of the different colors are blended. The hue and luminosity of the first input color are retained in the final output.
- **Hue shift:** the hue value of the first input color is shifted by the hue values of the successive colors. The saturation and luminosity of the first color are retained in the final output. Zero shifting occurs when the successive colors are Cyan (Hue=128).
- **Luminosity shift:** the luminosity value of the first input color is shifted by the luminosity values of the successive colors. The saturation and hue of the first color are retained in the final output. Zero shifting occurs when the successive colors have a luminosity value of 128.
- **Saturation shift:** the saturation value of the first input color is shifted by the saturation values of the successive colors. The luminosity and hue of the first color are retained in the final output. Zero shifting occurs when the successive colors have a saturation value of 128.
- **Slope blend:** the input colors are mixed in a proportion that depends on the local slope. The successive colors replaces the first color on vertical surface.

Add

This combiner node outputs the sum of all its inputs.

Subtract

This combiner node subtracts from the first input all subsequent inputs.

Multiply

This combiner node outputs the product of all its inputs.

Divide

This combiner node divides the first input by all subsequent inputs.



Image Combiner Node

The Image combiner node allows you to add bitmap images together, usually Image Sample Nodes, for a combined effect. An example would be a sand texture with rocks added. It takes the main (background) texture color as the first input, and then an arbitrary number of images, or image nodes in combination.

Math Nodes

The math nodes present utility operations that are not used in everyday graphs.

Conversions

Vector To RGB

This node receives a vector and outputs a color where the red component is equal to the value of the vector along the X axis, the green component is equal to the value of the vector along the Y axis, and the blue component is equal to the value of the vector along the Z axis.

RGB To Vector

This node receives a color and outputs a vector where the value of the vector along the X axis is equal to the red component of the input color, the value of the vector along the Y axis is equal to the green component of the input color, and the value of the vector along the Z axis is equal to the blue component of the input color.

RGB To HLS

This node is useful to convert colors from the Red-Green-Blue paradigm to the Hue-Luminosity-Saturation paradigm. It receives a color and outputs a vector where the value of the vector along the X axis is equal to the hue of the input color, the value of the vector along the Y axis is equal to the luminosity of the input color, and the value of the vector along the Z axis is equal to the saturation of the input color.



HLS To RGB

This node does the exact opposite of the previous one. It converts colors from the Hue-Luminosity-Saturation paradigm to the Red-Green-Blue paradigm. It receives a vector containing the HLS data and outputs a color based on that HLS data.

Color To Brightness

This simple converter node returns the brightness of the input color.

Vector Operations

Offset

This node simply adds an offset to the input vector.

Offset: defines the vector that will be added to the input vector.

Rotation And Twist

This node applies a rotation and twist transformation to the input vector.

Transformation: click on the Edit button to open the Transformation Editor. This dialog lets you indicate a rotation angle around each of the world axes, as well as a twisting angle of these axes one towards another.

Projection

This node transforms the input vector into the requested coordinate system.

Projection type: this drop-down list box lets you select the projection type of the node:

- **Cylindrical:** if this option is selected, the input vector will be converted to the cylindrical coordinate system.
- **Spherical:** if this option is selected, the input vector will be converted to the spherical coordinate system.

Matrix Transformation

This node lets you apply a user transformation matrix to the input vector.

Line 1..3: these 3 vector parameters specify the transformation matrix that will be applied to the input vector.



Decomposer

The Decomposer node takes a vector as input and outputs a number. It splits the input vector into 3 possible outputs that correspond to each one of the input vector's components. When you attempt to connect a link to a decomposer node, a popup menu will appear so that you can select the desired component.

Composer

The Composer node does just the opposite of the Decomposer node: it takes 3 numbers as inputs and outputs a vector constructed from these 3 inputs.

Length

This simple node takes a vector as input and returns a number representing the length of the vector.

Normalize

This node takes a vector as input and returns a vector pointing in the same direction, but with a length of exactly 1.

Dot Product

This node takes two vectors as inputs and returns a number corresponding to the dot product of both vectors. If the two vectors are normalized, the dot product is equal to the cosine of the angle between the two vectors. If the vectors point in exactly the same direction, the dot product is equal to 1, if they point in exactly opposing directions, it is -1, and if the two vectors are at right angles one with the other, it is 0.

Vector Product

This node takes two vectors as input and returns a vector that is the result of the vector product of the two input vectors. The result of the vector product is a vector that is at right angles with both of the input vectors.

Vector Quantization

Vector Quantization is a process where an input 3D vector is transformed into another one by choosing among a discrete subset of 3D space.

This is implemented by partitioning the input space in “cells”, choosing a privileged point in each cell, and always returning this specific point instead of the input vector when the input falls inside the cell.

To allow for a wider range of effects, this node implements a kind of smoothing of the resulting vector for some subset of parameters (see detailed description below) for which



the input vector will not be fully “snapped” to the privileged point, but only attracted by it to some extent.

Two cell partitions are currently implemented. Not all parameters shown apply to both partitions. Unavailable parameters are grayed out so as to avoid confusion.

Parameters

Origin: this parameter acts as an offset on the cell pattern applied on the input space. Only the pattern is affected; the output vector is not itself offset and will always be close to the input value (closeness depending only on cell sizes).

Scaling: this parameter acts on the size of cells along each of the X, Y, Z axis. Higher values mean bigger cells.

Quantization shape

Regular cells: these cells have a shape ranging from circular (influence = 0) to square (influence = 1), with all intermediate “rounded square” shapes in between. Smoothing is always applicable, even with an influence of 1.

Voronoi cells: these cells have an irregular polygonal shape and very different sizes. Influence defines the thickness of a border between cells, in which only partial attraction applies on the input vector. Smoothing only applies inside this border, ie. there can be no smoothing with an influence of 1.

Influence: this parameter defines how a cell’s privileged point attracts the input vectors falling inside the cell. Its exact meaning depends on the cell pattern used.

Smooth transition: this parameter defines how smooth is the transition between the invariant and attracted areas.

Number of cells to consider: this parameter only applies to Voronoi cells. It can be used to augment variety of the cellular partitioning by considering overlaps of 2 or more cells as distinct cells in themselves. On the other hand, influence and smoothing cannot apply when considering cell overlaps.

By blending the result with a non-quantized version of the same fractal, one can localize the effect following some pattern, eg. by driving the blending with another noise.

By adding a turbulence to the input vector, the cell pattern will also be affected, which means it is very easy to obtain irregular cell borders, which is even more interesting than a simple quantized-input fractal.



Other Math Nodes

Sine

Returns the sine of the input number.

Input as degrees: if checked, this option indicates that the input value is in degrees rather than radians.

Arc Cosine

Converts the input number into a number whose cosine is the input value.

Output as degrees: if checked, this option indicates that the output value is in degrees rather than radians.

Floor

Returns the round number that is just below the input number.

Fractional Part

The fractional part is the part of the number after the dot. It's equal to the number less the Floor of the input number.

Invert

This node inverts the input number (returns $1/x$).

Power

This node returns the first input raised to the power of the second input.

Square Root

This node returns the square root of the input number.

Multiply

This node multiplies its two inputs together.



Dynamics Nodes

The dynamics nodes provide a selection of nodes typically used for controlling the relationships between object properties. For instance, dynamics nodes are used to create loose link and track relationships between your objects (see [here](#)).

Note:

Dynamics nodes will only work in the context of Object Graphs.

Link Relationship

This node is used whenever a link relationship needs to be implemented between objects.

Offset: this setting controls the position offset between the input object position and the output position. In the case of linked objects, the offset is automatically updated when you move the linked object in the interface.

Link orientation, Link size and Link position: these settings replicate the linked object options available in the [Forward Dynamics Options](#) dialog.

Track Relationship

This node is used whenever a track relationship needs to be implemented between objects.

What it basically does is calculate the required orientation based on the difference between the two input positions.

Orientation offset: this setting controls the orientation offset with the tracked object. In the case of tracked objects, the offset is automatically updated when you rotate the tracking object in the interface.

Derivative

This node calculates the derivative over time of the input (for instance, the derivative of position over time is speed). The type of input and output is automatically defined as soon as you connect the node.

Integral

This node calculates the integral over time of the input. The type of input and output is automatically defined as soon as you connect the node.



Delay

This introduces a delay in the evaluation of the input. The value that is output by the node is the value of the input “delay-time” earlier. The type of input and output is automatically defined as soon as you connect the node.

Delay: the number of seconds by which the input value is delayed.

PID Controller

A PID controller is a standard type of automation controller that is used to try and achieve as best possible a desired value with an imperfectly responding system. PID stands for ‘Proportional, Integral, Derivative’.

The PID controller compares the current value of a parameter with a desired value and generates an output value based on the difference between them (the error), knowing that the output may not be exactly realized because what is being controlled is not a perfect mechanical system.

For instance, if a camera-man is tracking an object that suddenly moves out of sight, it will take a little time for him to re-track the object. This is caused by the time it takes for the camera-man to realize that the object has moved, the time it takes for his brain to send a signal to the arm, and then the time it takes for the arm to move. During this time, the tracked object may have moved again.

The PID controller takes two inputs:

- **Current:** this is the current value of the parameter being controlled.
- **Target:** this is the desired value that the parameter should have.

The PID controller uses the three following settings of standard automation controllers:

- **Proportional:** this setting determines the reaction to the difference between the desired value and the current value (error),
- **Integral:** this setting determines the reaction based on the accumulation of recent errors, and
- **Derivative:** this setting determines the reaction to the rate of change of the error.

It can be tricky to find the right settings for a PID controller, especially if the system being controlled exhibits non-linear responses.

For an in-depth understanding of how PID controllers work, we recommend that you refer to the Wikipedia article at http://en.wikipedia.org/wiki/PID_controller.



Distance Constraint

This node constrains the distance between the input value and a reference point to a certain range, so that, e.g. the position of an object cannot get closer or move further than a certain distance:

Center: this is the reference point.

Min distance: the minimum distance between the input value and the center (the “collision” size).

Max distance: the maximum distance between the input value and the center (the length of the leash).

Axis Constraint

This node constrains the input value to a given axis, defined by two points:

Point 1 and **Point 2:** these are the two points defining the axis to which the input value is constrained.

Grid Constraint

This node constrains the input value to a grid, the resolution of which is definable:

Grid size: the resolution of the grid.

Acceleration Limiter

This node calculates the double-derivative of the input (acceleration) at the current time and limits it to the indicated value.

Max acceleration: the maximum value allowed for acceleration. Speed changes will be constrained by this value.

Speed Limiter

This node calculates the derivative of the input (speed) at the current time and limits it to the indicated value.

Max speed: the maximum value allowed for the speed. Value changes will be constrained by this setting.

Low Pass Filter


This node averages the current input value with values at previous frames, thus eliminating any high-frequency variations.



Decay time: the delay between the occurrence of a value and the time when this value contributes exactly 50% of the output.

Examples

Creating Turbulence

In order to add turbulence to a noise, you need to **Extract** the **Origin** parameter (by clicking on the  button alongside the **Origin** parameter) and replace the constant node with a **Turbulence** node. You can adjust the look of turbulence by using the turbulence node parameters (see [here](#)).

Slope Dependent Scale

Create a noise node (e.g. a Value Perlin noise). Extract the **Scale** parameter and connect it to the **Slope** input. The scale of the noise now varies according to slope. However, if you are previewing the function on a sphere, you may notice that lower parts of the sphere are uniform. That's because scale is limited to positive values, and the slope varies between -1 and 1. So all parts of the sphere with a negative slope have a scale that is clamped to 0.

Select the link that connects the slope to the scale parameter, and click **Filter**. Change the type of the filter to “Absolute”. Now the scale of the noise bounces back up on the lower parts of the sphere and only reduces on vertical slopes.

If you want the scale of the noise to increase on vertical slopes instead of horizontal surfaces, you will need to invert the slope. To do this, select the “Absolute” filter node and click **Filter** again. This adds a second filter behind the absolute filter node. Change the type of that new filter node to “Brightness-Contrast” and indicate a **Contrast** of -1 and a **Brightness** of 1. The output of that node is now 0 on flat surfaces and 1 on steep slopes, hence smaller scale on flat surfaces.

You can also modulate the scale of the noise according to the altitude rather than the slope: simply plug the scale parameter into the **Altitude** input. However, altitudes vary significantly over the surface of the object, so you may want to reduce the amplitude of these variations, for instance by inserting a “Multiply” filter in between the Altitude input and the scale parameter, in order to tone down the variations of altitude (e.g. multiply



by 0.1 the Altitude input).

You could also plug another noise into the scale parameter. Create a new noise node and connect the scale parameter to that new noise node. If you enter 0 wavelength along the X and Y axes of this new node, you will create horizontal stripes resulting in horizontal stripes of varying scale. You can add the “Absolute” filter node to avoid negative values, and you can even control the influence of the horizontal stripes on the noise by adjusting the contrast and threshold (threshold will set the minimum scale, whereas contrast will control the variation in scale along the stripes).

Variable Color-Texture Mapping

You can easily blend the color of an object with a texture map according to the value of slope:

Edit the Color production function of a Simple Material. Select the color output and create a **Color** node. Make it a “Linear Interpolation 2” type of node. Plug the input of the node into the **Slope** input.

Note:

The color node’s output has automatically been plugged into the Color channel.

If you check the material, you will see that the color of the material varies according to slope. Now, in the color node, extract the “Color 2” parameter. Replace the constant color node with a **Texture Map** node and load a texture. The material is now solid color on flat surfaces, but mapped with a texture on vertical parts.

Of course, if you change the color node’s input from the slope to a noise, you can control the mapping with that noise.



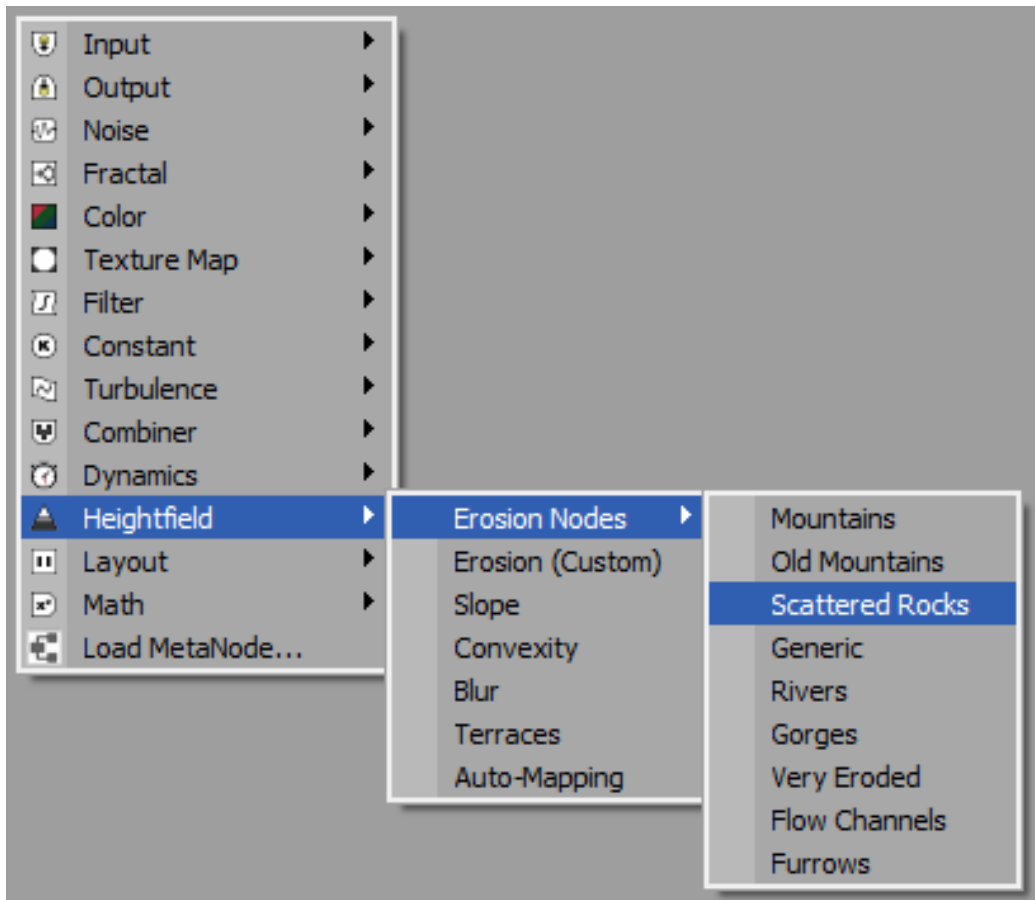
Heightfield Nodes

These nodes are for use in the graphs of heightfield terrain only. These create erosion effects that cannot be achieved procedurally.

Erosion nodes

These nodes simulate hydrological and thermal erosions of a landscape. Several pre-configured erosion presets are supplied as distinct nodes, and a fully customizable Erosion node is also available when it is necessary to precisely adjust all parameters. All nodes expose the same outputs and some of the parameters, but only the customizable node exposes the full parameter set.





The Erosion nodes in the Function Graph's contextual menu.

Available outputs

Several outputs are available, some more useful like the actual terrain altitude, while others are more suited to control material distributions or even to apply a 2D perturbation to some other nodes in the function graph.

- **Altitude:** ground altitude after the simulation.
- **Flow Surface:** when using this output, the terrain is “filled” to show the depth of the flow running over it. The flow surface is not horizontal because the flow simulation never really reaches an equilibrium.
- **Flow Speed:** this output shows the flow speed over the terrain at the end of the simulation.



- **Flow Vector:** similar to Flow Speed, but a 2D vector is output which direction and strength represent the actual flow over the terrain at the end of the simulation.
- **Flow Depth:** this output shows the flow depth (or humidity factor) over the terrain at the end of the simulation.
- **Transported Sediment:** this output shows the amount of sediment currently transported by the flow at the end of the simulation.
- **Displaced Soil:** this output shows the total amount of displaced soil at the end of the simulation.

Presets

The Erosion algorithm uses an extensive set of parameters in order to achieve a wide range of possible results. As a consequence, configuring the node has a somewhat steep learning curve. Presets help apprehending the roles and interactions of the parameters by supplying sets of parameters that work well together and produce distinctive results. You can either use the pre-configured Erosion nodes corresponding to each preset, or use a fully customizable Erosion node and apply the preset on it.

Note that when applying a preset in the fully customized Erosion node, almost all parameters of the node are changed to match the preset parameters, overriding any changes you may have made since the last preset applied! You can then tweak the parameters as you like as long as you do not apply another preset.

Note:


“Vertical Scale”, “Quality Boost” and “Evaporation Amount” are not modified by a change of preset.

Note:

A message will warn you when the node parameters are about to be overwritten by a change of preset.


Here are examples of the presets, applied on the following generic mountain flank with basic fractal detail:





Non-eroded terrain


Mountains: This preset leaves peaks and distinctly rougher areas on the terrain.



Mountain erosion


Old Mountains: Simulates how mountains are eroded after the passing of geological times.





Old mountains


Scattered Rock: This preset yields isolated rocks in the middle of smooth hill slopes made of sand/soil.



Scattered rocks


Generic: This generic preset applies to a wide range of terrains.





Generic Erosion


Rivers: More contrasted, this preset digs rivers into a terrain made of soil and rocks.



River-like Erosion


Gorges: Much more contrasted, this preset digs deeply into a terrain made of harder rocks.





Gorges-like Erosion


Very Eroded: This preset applies a heavy erosion on the terrain, erasing most of the rougher parts.



Very eroded terrain


Flow Channels: This preset applies moderate erosion and can also be used to determine where water would flow over the terrain.





Erosion channels

Furrows: A variant of “Flow Channels” where the erosion is allowed to apply more thoroughly, so that the terrain is more eroded along the water channels.



Furrows



Global Parameters

- **Vertical Scale (meters):** this parameter is expressed in physical units. The erosion processes simulate actual laws of physics, which rely among other physical quantities on the altitude variations over the heightfield. Obviously, the resulting terrain height cannot be automatically determined and used, since it itself depends on the result of this node. This is why the height of the simulated heightfield can be configured. It is also a way to tweak the effect of the simulation without changing the height of the resulting terrain in the scene.
The horizontal dimensions are also an important physical quantity on which the simulation depends. They are automatically determined by using the actual terrain dimensions in the scene, so take care to first scale your terrain before tweaking the erosion parameters.
- **Timescale:** this is a measure of the duration of the simulation. Raising or lowering this parameter will greatly change the aspect of the resulting terrain, but will also impact the duration of the computations proportionately.
- **Smallest Feature (meters):** this parameter is also expressed in physical units in order to be comparable with the Vertical Scale parameter. It is used to prevent the simulation process from caring about small details on the input heightfield. This parameter is important when you want to use the Erosion node on very high-resolution heightfields but you don't need to erode pixel-level details. With a low Smallest Feature parameter, the simulation will run on a grid the same resolution as your heightfield, which can lead to very long computing times.
- **Favor Larger Features:** this parameter has no particular unit. It influences the simulation algorithm in order to apply the erosion processes more at larger scales when the parameter value is higher. At or below a value of 1., on the other hand, smaller scale details are privileged.
- **Quality Boost:** as with all physical simulations, artifacts can appear because of the performance-realism trade-off taken in the algorithm. This parameter is a way to boost the quality of the simulation, by using values higher than 0., or on the contrary to speed-up the computation by using values lower than 0., when the simulation quality is not a limiting factor.

Terrain Properties

The hydrological part of the algorithm simulates dissolution, transport and deposition of soil along the water flow over the terrain. These parameters control the physical properties of the terrain.

- **Rock Hardness:** harder rocks are less eroded than softer ones.
- **Sediment Deposition:** this parameter is a way to control how fast soil is actually deposited while it is being transported over the terrain, ie. it is a way to tweak the



actual physical process at work.

- **Erosion Strength:** this is another non-realistic parameter to influence the speed of erosion. This can lead to unwanted artifacts though, so keep in mind that another way to erode your terrain more is to raise the Timescale parameter, at the expense of longer computing times.

Flow Parameters

The hydrological simulation is a compromise between physical realism and computation times. Simulating actual rain and the whole water cycle over the terrain is of course impossible, so the parameters below tweak the simulation using intuitive concepts:

- **Flow Level:** this is a measure of the minimum amount of water present over the terrain.
- **Channeling Threshold:** this parameter controls at which water level the terrain will begin to be affected by thermal and bank weathering (see Weathering Parameters below), digging river beds into the terrain. Though it can be seen as a weathering parameter more than a flow parameter, the effect of the Channeling Threshold value is closely linked to the value of the Flow Level. See how the various presets set up those two values to create different effects on the terrain.
- **Flow Depth Limit:** as a consequence of the performance-quality compromise, the simulation sometimes digs unrealistically deep into the terrain, as if an infinite height of soil was available to transport. This parameter can limit the digging depth and should be set up in relation with the Vertical Scale of your terrain.
- **Dampen Flow, Clamp Flow:** these non-realistic parameters are used to slow down the simulated flows in order to produce various effects on the terrain. Dampen Flow slows down the flows continually over the whole simulation, while Clamp Flow is an upper limit to the flow speed. See how the various presets set up those two values to create different effects on the terrain.
- **Erosion Emphasis:** this non-realistic parameter can be used to emphasize the effect of the hydrological erosion process without resorting to a longer Timescale.
- **Emphasis Threshold:** this parameter is used to adjust the shape of the Erosion Emphasis effect. Close to 0., the emphasized features are more spiky while at higher values the features are more rounded.

Weathering Parameters

- **Thermal Weathering:** some amount of thermal weathering can be added to the simulation. Thermal weathering is a physical process due to the alternating periods of colder and warmer weather. It erodes the steeper slopes by detaching little bits of rocks which roll down and accumulate to form conical talus. The actual amount of thermal weathering effective on your terrain will greatly depend on the steepness of



its slopes and hence, on the ratio of its Vertical Scale to its horizontal dimensions.

- **Bank Erosion:** when streams and rivers form over a terrain, its bank tends to erode quite fast because of the forces exerted by the running flows. This is a prominent effect in our simulation in order to create actual valleys on the terrain, whereas hydrological erosion alone would tend to create gorges instead.
- **Bank Transition:** this parameter blurs the boundary between the drier areas of the terrains and those affected by Bank Erosion.
- **Preserve Roughness:** this parameter can be used to preserve part of the roughness of the input terrain even over eroded areas. Less eroded areas will still preserve more roughness than more eroded ones.

Flow Surface Output

- **Evaporation Amount:** when using the Flow Surface output, this parameter can be used to adjust water depth considered by the algorithm to compute the output.

Slope node

This node computes the mean slope of the heightfield. The parameter “Account for larger features” hints at the prominent scale at which the heightfield slope is computed. Close to 0, the slope is computed by mostly considering small-scale detail, ie. using the full resolution of the heightfield. Closer to 1, the smaller-scale perturbation is smoothed out to leave only the overall slope of the heightfield.

Convexity node

This node computes the convexity of the terrain.


The parameter “Account for larger features” hints at the prominent scale at which the heightfield convexity is computed: close to 0, the convexity is computed by mostly considering small-scale detail, ie. using the full resolution of the heightfield. Closer to 1, the smaller-scale perturbation is smoothed out to leave only the overall convexity of the heightfield.



Blur node


This node smooths out the heightfield depending on a radius.

Terraces node



Levelled terraces






Irregular terraces





Rounded terraces





Square-ish terraces

This node uses an advanced algorithm to create customizable terraces on a heightfield terrain.

In addition to the usual altitude input, the node can take a special “Locate Terraces” input into which can be fed a noise which prominent features will influence the shape of the terraces. For example, a round noise will yield round terraces, a Voronoi noise will tend to yield Voronoi-cell shaped terraces, etc. which opens a lot of possibilities for customization.

The other parameters controlling the terraces are:

- Strength: overall strength of the terracing effect.
- Smooth contours: how much to try and smooth the terraces contours.
 - Note: using a high smoothing value will noticeably raise the node’s computing times.
- Slope Width: terraces are linked to each other by sloped areas, in order to avoid abrupt altitude changes. This parameter is the width a these slopes: the narrower the slopes are, the steeper they get.



- Max. Size: limits the dimensions of the terraces, to avoid having terraces spanning the whole terrain (depending on the “Locate Terraces” input).
- Max. Height: limits the altitude difference between two adjacent terraces.

This node has a secondary output named **Slope** which will typically allow you to apply different effects on the terraces’ flatter areas than on the slopes between distinct terraces.

Auto-mapping node

This node will linearly map the input range of the heightfield being fed into it into the desired output range. This is a real helper since thanks to it you no longer need to worry about the output range of a set of nodes (even noises and fractals, which are notably hard to predict). Whatever it is, you just have to insert an auto-mapping node to “normalize” the range to suit your needs.

For example, if you want to use the output of a node as a weighing factor, you will set your auto-mapper to output in the $[0 ; 1]$ range. Or if you need to output to a custom dependency for later reuse in a material graph, you will auto-map to $[-1 ; 1]$ in order to exploit the full range of values and no longer need to adjust the material graph filters each time you make changes in the terrain graph.

Compare to a similar situation in a procedural graph, where such a node is not available: you would have to evaluate yourself the range of values at a given node in the graph and manually tweak various filters to adjust the range to suit your needs, which usually takes some trial and error.

Apart from the output range parameters, the node takes an additional parameter named “Input Range Rounding”: this prevents the mapping from changing by a slight factor when the input range varies slightly. For example, if you set “0.1” in this parameter, the mapping factor will be the same whether the input range is for example $[0.02 ; 0.98]$ or $[0.0 ; 1.0]$.

This can be useful to maintain some stability in the output of a graph, for example to compare the overall effect of slightly different parameters in some other nodes. On the other hand, using a non-zero value here means the output range may not be exactly matched.



Layout Nodes


This category contains space related functions. That is, its nodes output values that depend on the layout of the input. For now, it contains only two nodes. The **Distance to Spline** node and the **Area Demarcation** node.

Area Demarcation Node

This node works in conjunction with zone extraction in the Terrain Editor for procedural terrains. When you extract a zone, you can create a hole in the original terrain (where the terrain was extracted). The **Area Demarcation** filter can be used to configure this hole.

You can also create this node and use it without the zone to get this type of hole effect.

The Area Demarcation node outputs one scalar, that can take one of only two values: one for points inside the area, another for points outside the area. For example, you can set the output to be 0.5 for points inside the shape, and -0.5 for points outside.



Area Demarcation example configuration (using “Rounded Square” shape) and preview in Terrain Editor and in a render



Available parameters

- **Area shape:** the shape of the area. Can be one of Circular, Rounded square, Rounded square(sharper) and Square.
- **Area center:** the position of the center of the area.
- **Area width:** the width of the area.
- **Area height:** the height of the area.
- **Demarcation smoothness:** specifies how much to interpolate at the transitions between inside and outside the area, in order to smooth the output.
- **Interior value:** the value outputted for a point inside the area.
- **Exterior value:** the value outputted for a point outside the area.

Spline Proximity node

The **Spline Proximity** node is located in the top-level category “Layout” in the *Function Graph*. Its goal is to drive terrains or materials based on a spline’s shape. This spline can be stored and edited inside the graph or can be a spline from the scene. Both types of spline can be converted into one another for editing purposes.

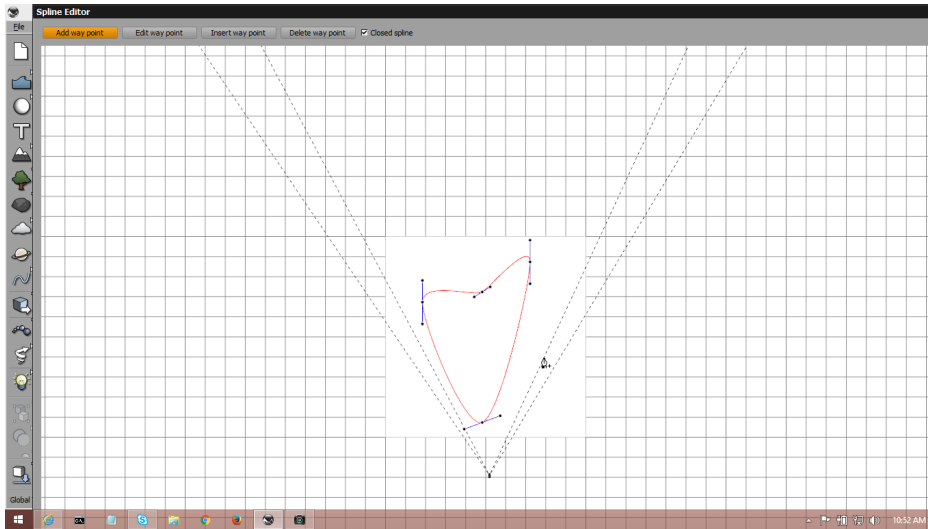
The **Spline Proximity** outputs a measure of proximity between a position and its closest point on a spline curve. The spline curve is user-defined: it can be a spline selected from the scene’s object list, or it can be created directly from the node’s interface, using a custom *Spline Editor* in the *Function Graph*. The Spline Editor lets you create a spline curve in world coordinates, by displaying the top view while placing the control points.

How to use it

You can either select an existing spline object from the scene’s objects list, or create a new spline curve from scratch using the VUE spline editor.

To create or edit the spline curve, click on the **Edit** button in the parameters area. The **Edit** function becomes available for the Internal Spline. The **Spline Editor** in the *Function Graph* will open.





Spline Editor

Inside the spline editor, you should see the current scene’s top view. To create the spline you want, you have four edition modes at your disposal:

- **Add way point:** When in this mode, left clicking on the view creates a new control point.
- **Edit way point:** This mode lets you move around control points that are already placed, by dragging them with the mouse on left click.
- **Insert way point:** When this mode is selected, you can insert control points directly on the curve. This means that the curve will have the same shape as before, but with one more control point.
- **Delete way point:** This mode lets you remove control points by left clicking on them.

You can also choose whether the spline should be closed or not, by checking the checkbox “Closed spline”. This adds or removes the link between the first and last control points, to make the spline cyclic or acyclic.

If you are not using the Internal Spline option, you can either use **Copy to Internal** to duplicate and edit the scene spline in the Function Graph, or use VUE’s usual spline edition tools to edit the spline in the scene. Note that the terrain is not automatically updated when the spline is changed in the scene until some other modification triggers the graph calculation.

The node has two parameters:

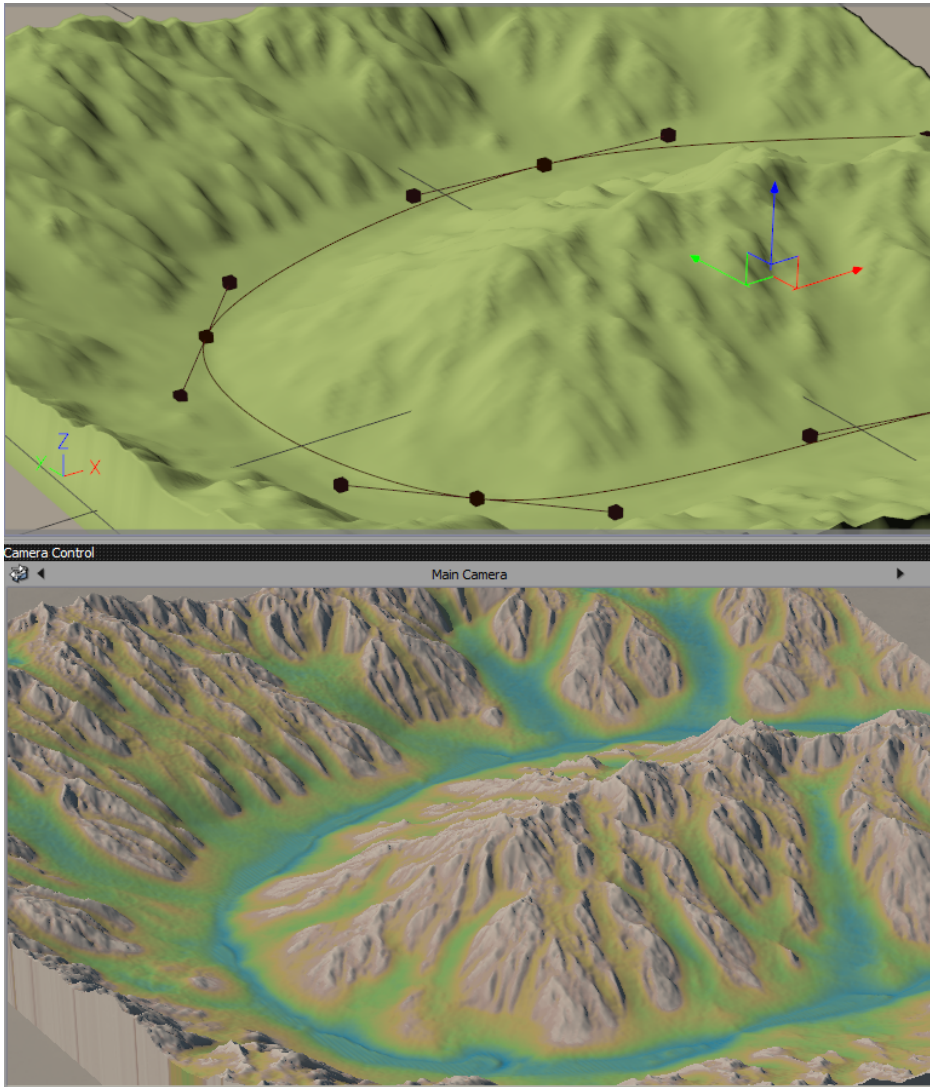
- *Invert Output:* when unchecked, the spline curve will be a “ridge” in the output



heightfield, ie. the points on the spline are actually the highest values of the output. When *Invert* is checked, the spline curve will instead the lowest values in the heightfield, like the bottom of a valley.

- *Invert Slope Inside Spline*: when the spline is closed (“cyclic”), it is possible to identify its interior, which can be useful for some terrain or material effects. Depending on your purpose, you may want to distinguish the values inside the spline from the values outside, by using a different output range. This also alters the shape of the output, hence the parameter name. For example, with a circular spline, you can check this parameter in order to have a volcano-shaped output (the spline would be the crater’s rim).





Example: Using a spline as a valley

Node outputs

The node has three possible outputs:

- *Distance*: this is simply the distance between each point at the base of the terrain and the closest position lying on the 2D or 3D spline.
 - When using a 3D spline, the output will also follow the spline's altitude profile,



which you can alter by leaving the terrain editor and moving the spline's control points inside the scene, using the gizmos available in VUE's four 3D views. Then go back to the terrain editor and recompute the terrain as the node output will not update automatically.

- *2D Distance*: similar to the *Distance* output, but only accounts for the horizontal distance (ie. the spline is projected on the terrain plane).
 - When using a 2D distance, the spline itself will remain on a level line in the node's output and, when *Invert Output* is checked, this level line will be the zero level.
- *Altitude Profile*: altitude of the closest spline position. Discontinuities may need to be blurred with a *Blur* node in order to make use of this output.

Scene Objects Mask Node

This node creates a heightfield which can be used as a mask to locate “foreign” objects lying over the terrain the heightfield represents. This is a powerful tool to create effects on the terrain easily, without having to design complex graphs. Examples of usage include:

- Digging the terrain to represent how it subsides slightly under the weight of objects.
- Proxy objects (not visible at render time) can be used to drive terrain effects, where several nodes would have been necessary using a traditional approach.
- Erosion or other effects can also be made to work around scene objects thanks to this node.



Spline Editor

A spline is an editable object in VUE. It can be used to create roads, define EcoSystems and geometric shapes.

It can also be used to configure “path effects”, which are various kinds of influence that a path can have on other objects of the scene, already used with spline curves in previous versions: material effects, geometry effects, and terrain effects. Geographical contours (actual line strings, but also the perimeters of geo-located polygons) can only be applied material effects (and, most importantly, generate EcoSystem populations).

A spline contains:

- Moving points that can be added or deleted
- Tangents that can be moved.

To this spline, effects can be applied such as:


- Material/EcoSystem effects
- A geometry effect
- A terrain effect, such as a road.

A spline is like a group; you have to Ctrl-click (Cmd click on a Mac) on one of its points or tangents to select it. Since a spline is listed in the *World Browser*, you can use the *World Browser* to select points and tangents.


- You can click to add points.
- You can switch to edit mode to select an existing point (or tangent).
- You can activate effects with tabs appearing for configuration.
- When you have more than two points, you can click **OK** to finalize the spline.

On the *Spline Editor*, there is a row of icons at the top.





Add point (): use to add a point to the spline.



Edit point (): use to change a point in the spline.

Broken / Aligned: Here are two ways a tangent can be moved. It keeps the angle.


Broken () moves a tangent but doesn't change the other side. **Aligned** () moves a tangent and moves the other side symmetrically.




Closed spline (): Use to join the first and last points of the spline to close the




spline.

Path finding (): Use to find the best path for the spline. Used for creating paths or roads through hilly terrains, for example.

Resample spline (): this creates new points for the spline, while keeping its shape. As it is done for path finding, new points are dropped on the underlying terrain.

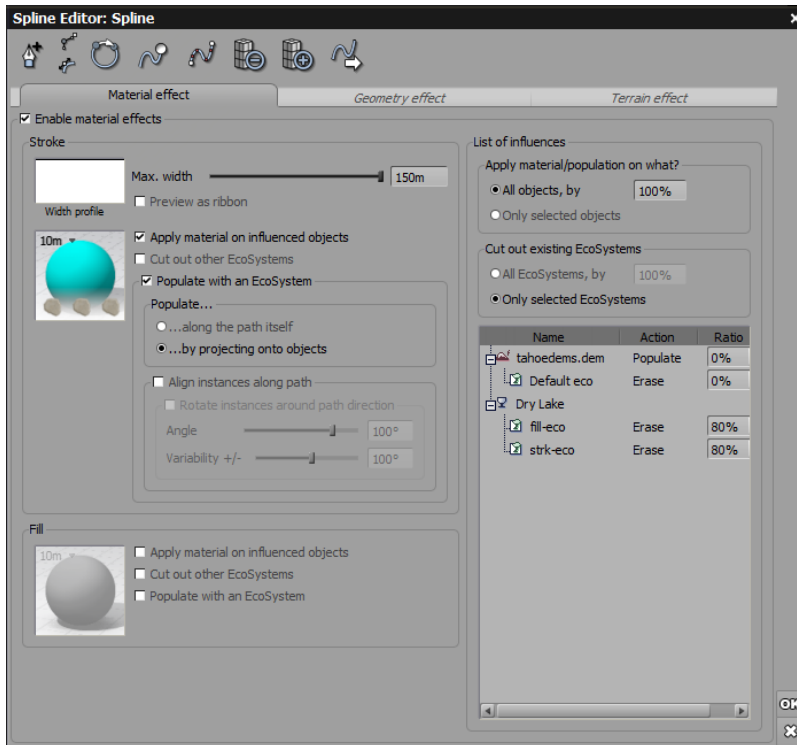
Use the **Halve** () and **Double quality** () to adjust the quality of the spline.

Import spline (): use to import a spline from another source. Valid sources are:

- Vector Graphics filter
- PostScript files
- Encapsulated PostScript files
- Adobe Illustrator (up to 3.2) files



Material Effect



Spline Editor – EcoSystem effect

Material effects allow you to apply materials and to populate EcoSystems along the spline and in the area defined by the spline, as well as the option to erase other EcoSystems existing under the spline.

Enable material effects: check this option to apply a material and/or EcoSystem.

In the **Stroke** section, you have the following settings:

Max. width: defines the width of the spline and defines a profile for the EcoSystem population using a filter.

Preview as ribbon: shows the spline as a ribbon in OpenGL but the ribbon doesn't render. This is useful for spline visualization if the **Geometry effect** isn't being used.



Apply material on influenced objects: the specified material layer will be applied on objects located directly above and below it in the scene. Click on the image to the left to access the Material Editor to create or load a material layer for the spline path.

Cut out other EcoSystems: cuts out areas of existing EcoSystems, for example, to create a path through a field.

Populate with an EcoSystem: allows creation of an EcoSystem on or below the spline. Click on the image to the left to access the *Material Editor* to create or load an EcoSystem.

Populate along the path itself: when checked, VUE always adds instances in the volume defined by the spline, considered as an area of the chosen width.

Populate by projecting onto objects: when selected, the instances are populated on the objects below the spline.

Align instances along path: EcoSystem instances can be aligned along spline direction using these settings.

Rotate instances around path direction: use the **Angle** and **Variability** settings to specify an angle of rotation around the spline normal.

The **Fill** section defines the area in the interior defined by the spline. Here you have the following settings:

Apply material on influenced objects: the specified material layer will be applied on objects located directly above and below it in the scene. Click on the image to the left to access the Material Editor to create or load a material layer for the spline path.

Cut out other EcoSystems: cuts out areas of existing EcoSystems, for example, to create a path through a field.

Populate with an EcoSystem: allows creation of an EcoSystem on or below the spline. Click on the image to the right to access the *Material Editor* to create or load an EcoSystem. Click on the image to the left to access the *Material Editor* to create or load an EcoSystem.


In the **List of influences** section, you have the following settings:

Apply material/population on what?: You can select **All objects** or **Only selected objects** for population. If you select all objects, you are able to specify the global ratio of influence. When applying a spline's material layer over other objects, the ratio will drive the blending of the spline's material with the object's other material layers. When populating an EcoSystem layer, the ratio will apply on the EcoSystem density. When selecting **Only selected objects**, the influence ratio can be set on a per-object basis in the list of objects displayed below this frame.



The **Cut out existing EcoSystems** section becomes active if this option is selected in either the **Stroke** or **Fill** sections. Here, select **All EcoSystems** or **Only selected EcoSystems** for the cut-out effect to apply on. If you select **All EcoSystems**, you are able to specify the percentage. If you select **Only selected EcoSystems**, the EcoSystems are listed in the box below for selection.

Geometry Effect



Spline Editor – Geometry effect

This effect extrudes a geometric shape along the spline, defined by these properties:

Max. Width: use the slider or key in the width of the geometry being created. Click on the Width Profile to the right to edit or load a filter.

Height: use the slider to indicate the height of the generated geometry. This value can be locked.

The filter also helps modify the shape of the geometry by defining the **Width profile**. Click to select a filter; right-click to edit the filter.

Offset z: use the slider to indicate the z offset.



Available geometric **Types** are:

- Tube
- Road
- Cobble (a square tube)
- Ribbon

Twist: There are two types of twist available – **Frenet twist** and the **Z twist**. These affect how geometries are created around each spline point. A **Z twist** generates normal paths/roads that go up, down, left and right following the underlying surface but they stay strictly horizontal on the surface. A **Frenet twist** generates paths/roads that follow the surface of the underlying structure. An example of this type of twist would be the tilted surface of tight curves of racing courses.

You can also add a manual roll by selecting any of the spline points and using the local rotation gizmo.

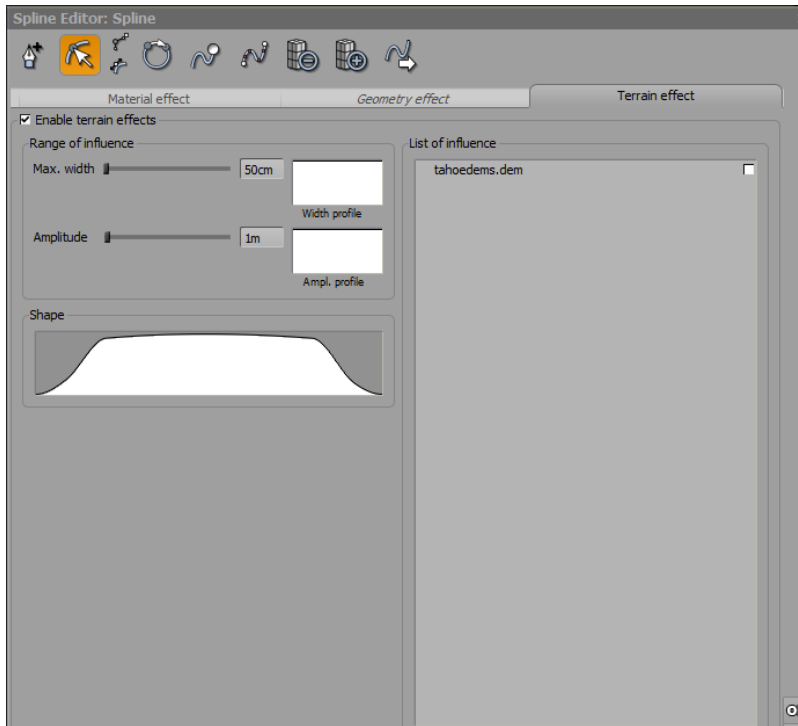
Limit: This defines a limit of the twist affect.

Caps: select to create covered ends on the geometry.

Click on the sphere in the **Material** section to assign materials to the geometric shape.



Terrain Effect



Spline Editor – Terrain effect

This effect relies on the filter or combination of filters to make an impact on the terrain. For example, you can cut grooves in a terrain. Or, when used in conjunction with the road spline, can be used to create a road bed.

Size and type

Max. Width: use the slider or entry field to define the size of the spline. Use the filter to define the shape of the terrain effect. Right-click on the filter image to edit the filter; left-click to open the browser and select a new filter.

Amplitude: use the slider or entry field to define the amplitude of the spline. Amplitude adds extra height between the spline position and terrain. Right-click on the filter image to edit the filter for amplitude; left-click to open the browser and select a new filter.

Shape: Use the filter here to define how the spline interacts with the terrain. Right-click on the image to edit the filter; left-click to open the browser and select a new filter.

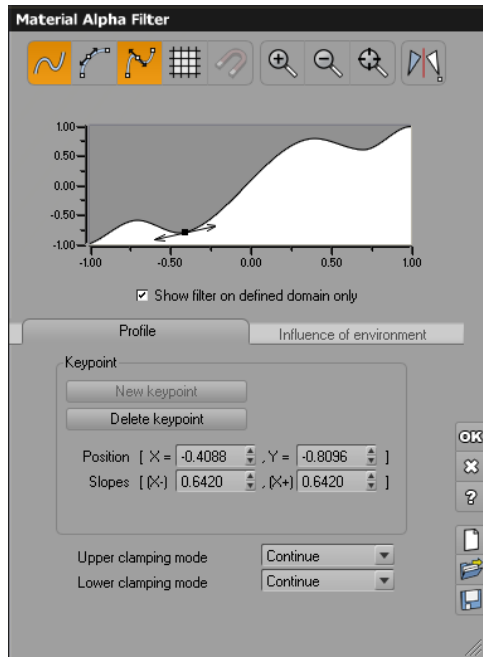


Editing Filters

Filters are used to modify profiles. The *Filter Editor* can be accessed in the *Terrain Editor* to modify the profile of a mountain by filtering the altitudes or from the *Material Editor* to modify certain aspects of the material, such as the transparency or bump definition.

The tools you use to modify filters are very similar to the ones you use to edit [Time Splines](#).

Description



Filter Editor – Profile tab

Filters enable you to transform any number in the range of -1 to 1 into another number, also in the range of -1 to 1, following a curve that you define. The value returned by the filter at a given position on the horizontal ruler can be read on the vertical ruler of the curve.



To open the *Filter Editor*, either click on the filter with the **Control** key pressed, or select **Edit Filter** from the filter's popup menu.

Inside the *Filter Editor*, you can zoom in and out, and pan the view using standard commands (Right mouse drag to pan, Ctrl + Right mouse drag to zoom). You can also resize this Editor if you need a more detailed view of the filter.

Some filters can also be influenced by the environment.

Below the *Filter Editor's* curve you will see a tab control with one or two tabs. The first tab controls the profile (shape) of the curve, while the second controls the influence of the environment. If the filter cannot be influenced by environment, the second tab is not visible.

The Curve

The curve is the large display that sits in the middle of the editor, just below the toolbar. This area represents the profile of the filter. You can zoom in and out, and pan the view using standard commands (Right mouse drag to pan, Ctrl + Right mouse drag to zoom).

Filters are built from **Key Points**, joined together by straight lines or curves. You can modify a filter by adding, moving or deleting key points. The key points are figured by small handles (📌) on the curve. These handles appear as soon as the mouse cursor is placed above the curve. All filters start from (0,0) and have a key point on the right edge (the corresponding handle can only be moved vertically).

Smooth Filters

VUE offers two types of filters: standard (linear) and smooth.

Linear filters are generated from segments while smooth filters are generated from cubic curves.

You can switch from linear to smooth filters, by clicking the **Smooth curve** icon in the toolbar.

The behavior of a smooth filter is identical to that of a linear filter except that you can change the slope of the curve around the key points, yielding a smoother -round- profile.

To modify the slope around a key point, select the key point by clicking on its handle (📌), or by typing its horizontal position in the **Position X** box. The **Slope** boxes now indicate the slope to the left and to the right of the key point. Type in new slope values.








If you selected the handle by clicking on it, the tangents to the curve. You can drag the ends of the tangents to modify the slope.

Selecting **Smooth joint** icon will ensure that the slope is the same on either side of the key point (the default). If you want to have a different slope on either side of the key point (e.g. to create a crease in the curve), you must deselect this option and then modify the slope.

Toolbar

The *Filter Editor's* toolbar is the collection of icons at the top of the editor. The meaning of these icons is as follows:

-  **Smooth curve:** this is a toggle icon. If the icon is not toggled, the filter will be built from straight lines; if it is toggled, the filter will be built from curves. Click on the icon to change the type of filter.
-  **Auto-tangents:** this is also a toggle icon; it is only available when the filter is smooth. If the icon is toggled, the tangents at newly added key points will be computed automatically in order to modify as little as possible the overall shape of the curve. If you drag a key point when this mode is active, the tangents will be modified dynamically so as to minimize the deformation of the curve.
-  **Smooth joint:** this is also a toggle icon; it is only available when the filter is smooth and a key point is selected. If the icon is toggled, the slope on either side of the key point will be the same, ensuring that the resulting curve doesn't exhibit any sudden changes in slope around that key point. If you deselect this option, the slope on either side of the key point can be modified independently, resulting in a crease in the curve.
-  **Show grid:** this is a toggle icon. When it is orange (enabled) a grid will be displayed on top of the curve. This grid can be used for reference when building a filter.
-  **Snap to grid:** this is a toggle icon, available only when the grid is displayed. When



snapping is on (the icon is orange), key points will be automatically “attracted” to the grid when you approach the mouse cursor from the grid. This is useful for setting up filters with “rounded” values.



- **Zoom in:** click this icon to display a zoomed view of the filter. This lets you edit detailed portions of the filter.



- **Zoom out:** click this icon to zoom out of the view of the filter. This lets you visualize a larger portion of the filter.



- **Reset pan/zoom:** click this icon to reset the view of the filter so that the filter fills up the entire graph exactly.



- **Flip Vertical Axis/Flip Horizontal Axis:** this flips the axis of the graph either horizontally or vertically.

New, Load, Save

Pressing **New** will reset the filter by deleting all key points.

Press **Load** to load one of the sample filters using the Visual Filter Browser.

Press **Save** to save the current filter in a stand-alone file, for use in future scenes. Saved filters will appear in the Visual Filter Browser like any other of the predefined filters. By default, filters are placed in the **Filters** subfolder. This means that they will appear in the **Personal** collection inside the Visual Filter Browser.



Profile Tab

This tab controls the general shape of the filter. Use this tab to add, modify or remove key points. If no tab is visible in the *Filter Editor*, then the controls that are displayed in the editor pertain to this tab.

Show filter on defined domain only: Check this field if you want to actually see the clamping effect of the filter. As you scroll, the clamping becomes evident.

When checked, the filter curve will only be displayed in the range (x and y) where the filter is defined. For example when editing a bump filter (in the *Material Editor*), x is only defined within $[-1, 1]$ and same for y. So when the box is checked, and the user uses the right mouse button to “scroll” the curve, the curve will not be displayed outside this range. It’s like a clipping rectangle.

You can scroll by pressing the right mouse button on the picture and moving the mouse (with the right mouse button pressed). So if you check the checkbox, and scroll to the right (for example), eventually the curve will be clipped. The definition domain depends on the use of the filter.

Adding Key Points

To create a new key point, you can either:

- double-click in the area where the curve is drawn. The new key point is created at the point you clicked. The curve is redrawn to use the new key point.
- click on the curve where you want the new key point; the coordinates of the clicked point appear in the **Position** boxes; you can edit them if required. To create the new key point, press the **Add key point** button. The curve is redrawn.
- type the coordinates of the new key point in the **Position** boxes, then press the **Add key point** button. The curve is redrawn.

You can’t create two key points at the same horizontal position.



Modifying Key Points

To modify a key point, you can either:

- click on the key point's handle (🖱️) and drag it with the mouse button pressed. If you press Control as you drag the cursor, the movement will be constrained to the closest axis. Each key point must stay between the previous one and the next one. When you select a key point, you can jump to the next one by pressing Tab (Shift Tab jumps to the previous). The selected key point becomes black. You can also modify the position of the key point by using the Up/Down and Left/Right arrow keys.
- click the handle (🖱️) of the key point you want to modify. The handle becomes black, and the **Position** indicated is now the position of the key point. Type the new position of the key point.
- type the horizontal position of the key point you want to modify in the **Position X** box, then indicate the new vertical position of the key point.

Note:

You can't move the horizontal key point position using this method.

Deleting Key Points

To delete a key point, click on the handle (🖱️) of the key point you want to delete, or type its horizontal **Position** in the box **X**, then press the **Delete key point** button. When you select a key point, you can jump to the next one by pressing Tab (Shift Tab jumps to the previous). You cannot delete the right-most key point.

Clamping modes

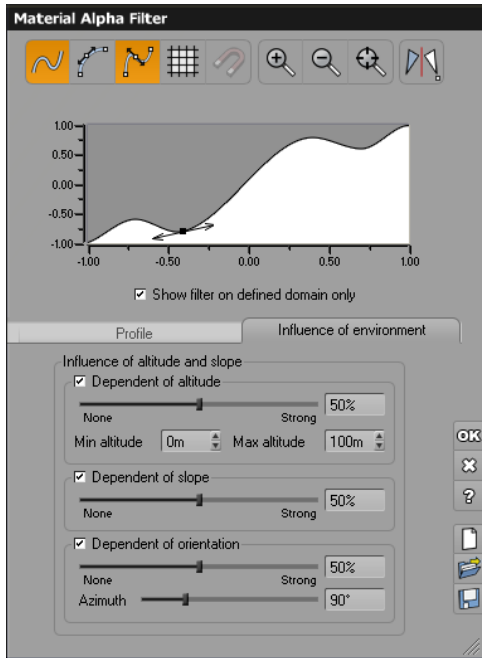
There are two clamping modes, **Upper** and **Lower**. There are four types of clamping modes that can be used:

- **Clamp:** values smaller than -1 will be forced to -1 for lower, and values larger than 1 will be forced to 1 for upper.
- **Continue:** values will be extrapolated from the inner curve (for instance, for linear filters, it is as if the line continued outside input range).
- **Mirror:** input value is transformed to simulate a mirrored repetition of the inner curve.



- **Repeat:** input value is cyclic. It's always shifted to fall back into $[-1, 1]$.

Influence of Environment Tab



Filter Editor – Environment tab

This tab is only available for filters that can be influenced by the environment (e.g. filters that are part of materials). If no tab is visible in the *Filter Editor*, then the controls that are displayed in the editor pertain to the **Profile** tab, and **Influence of Environment** controls are not available.

Influence of Environment works more or less like its equivalent in the *Mixed Material Editor* (see [here](#)). The filter can be influenced by the environment in three different ways: dependency to altitude, to slope and to orientation.

Note:

The notion of “environment” only has a sense when the filter is part of a material.



Dependent of Altitude

Check this option if you want the values of the filter to be influenced by altitude. The slider position indicates the importance of the effect. If the value is 0, altitude has no influence. If the value is 100%, the filter will always return 0 (whatever its profile) when the altitude is below Min altitude; it will always return 1 when the altitude is above Max altitude.

If this option is selected, the values returned by the filter will be saturated by altitude. This means that, if the altitude is close to the Min altitude, values will be lowered, but, as altitude increases, the values will be less and less lowered, and eventually start being raised as altitude gets nearer to the Max altitude (values will always stay clipped to 0-1).

Min altitude and **Max altitude**: these settings indicate the range of altitudes where the filter is influenced by altitude. Outside this range, the influence will be constant.

Dependent of Slope

Check this option if you want the values of the filter to be influenced by slope. The slider position indicates the importance of the effect. If the value is 0, slope has no influence. If the value is 100%, the filter will always return 0 (whatever its profile) when the surface is horizontal, and it will always return 1 when the surface is vertical.

Dependent of Orientation


Check this option if you want the values of the filter to be influenced by orientation. The slider position indicates the importance of the effect. If the value is 0, orientation has no influence. If the value is 100%, the filter will always return 0 (whatever its profile) when the surface points in the opposite direction to the value indicated by the Azimuth , and it will always return 1 when the surface points in the direction of the Azimuth.

Azimuth: use this setting to indicate the preferred orientation. When the orientation gets closer to this setting, the filter will return higher values. 0° corresponds to positive values on the X axis, 90° to positive values on the Y axis, 180° to negative values on the X axis, and 270° to negative values on the Y axis.



Editing Color Maps

Description



Color Map Editor

Color maps are used to produce variable colors and opacity (alpha).

Basically, they associate a color and an opacity to each of the values in the range of 0 to 1, following a gradation defined by the user.


To open the *Color Map Editor*, either click on the color map with the **Control** key pressed, or select **Edit Color Map** from the color maps contextual menu. This window is resizable.

In the editor, the color gradation displays the colors that will be returned by the color map for each value on the horizontal ruler. The black and white gradation above displays the alpha value that will be returned for these points.

Color maps are built from **Key Colors** and **Key Opacities**. Key colors define the color of the map at given positions. Key opacities define the opacity of the color map at given positions. VUE automatically generates a gradation to smoothly join the key values. You



can modify a color map by adding, moving or deleting key colors and key opacities.

Key values are figured by handles () on the horizontal rulers around the color gradation. Key colors appear under the color gradation, whereas key opacities appear above it.

Adding Key Colors

To add a new key color, you can either:

- Double-click in the central area where the color map is drawn; a dialog box pops up letting you choose the new color that will appear at the point you clicked. Select the color and click **OK**. The color map is redrawn to include the new color.
- Click in the central area where the color map is drawn; the position of the point you clicked appears in the **Position** box. To create the key color, press **Add key color** or double-click on the **Current color** square in the middle of the editor; a dialog box pops up letting you choose the new color that will appear at the point you clicked. Select the color and click **OK**. The color map is redrawn to include the new color.
- Type the position of the new key color in the **Position** box then press **Add key color** or click on the **Current color** square in the middle of the editor; a dialog box pops up letting you choose the new color that will appear at the point you clicked. Select the color and click **OK**. The color map is redrawn to include the new color.

You can't create two key colors with the same horizontal position, but you can drag a key color up to another one in order to superpose them and create a sudden change in color.

Adding Key Opacities


To add a new key opacity, you can either:

- Double-click on the black & white gradation above the color map; a new opacity key is added. This newly added key will not change the overall opacity of the color map.
- Click on the black & white gradation above the color map; the position of the point you clicked appears in the **Position** box. To create the key opacity, press **Add key opacity** or modify the **Current opacity**.
- Type the position of the new key opacity in the **Position** box then press **Add key opacity** or modify the **Current opacity** to create the key opacity.



You can't create two key opacities with the same horizontal position, but you can drag a key opacity up to another one in order to superpose them and create a sudden change in opacity.




Modifying Key Colors & Opacities

To move a key value, click the key's handle () and drag it with the mouse button pressed. Each key value is constrained by other keys on either side.

To modify the color of a key color, you can either:

- double-click the handle () of the key you want to modify; a dialog box opens letting you select a new color for the key color. Select the color and click **OK**. The color map is redrawn to include the new color. When you select a key color, you can jump to the next one by pressing Tab (Shift Tab jumps to the previous). You can also modify the position of the key colors by using the Left and Right arrow keys.
- click the handle () of the key you want to modify. The handle becomes black, and the **Position** indicated is now that of the key color. Double-click on the **Current color** square in the middle of the editor; a dialog box opens letting you select a new color for the key color. Select the color and click **OK**. The color map is redrawn to include the new color.

To modify the opacity value of a key opacity, click the handle () of the key you want to modify. The handle becomes black, and the **Position** indicated is now that of the key opacity. Use the **Current opacity** setting to adjust the opacity of the key.

Manipulating Multiple Key Colors


You can manipulate several key colors simultaneously. To select multiple key colors, press the **Ctrl** key and then click on each key color you want to select. You can also select a key color, then press the **Shift** key and select another key color, which results in the selection of all key colors between the two selected. Once several keys are selected, it is then possible to:

- Move them together by moving any of the selected keys.
- Delete them by pressing the **Delete** key.
- Copy them by pressing **Ctrl+C**.
- Paste the copied key colors (at the current cursor position) by pressing **Ctrl+V**.
- Change the color's values (hue, luminosity and saturation) by using the appropriate up/down buttons in the editor.


You can also flip the entire color map by clicking on the icon in the lower right under the color map itself.



Deleting Keys

To delete a key, click on the handle () of the key you want to delete, or type its horizontal **Position** in the corresponding box, then press the **Delete key color** or **Delete key opacity** button. When you have selected a key, you can jump to the next key by pressing Tab (Shift Tab jumps to the previous key).

Advanced Opacity Control

If you want to control the opacity of your color map precisely, right-click or control-click on the black & white opacity gradation above the color gradation, or click on the **Edit opacity filter** button () to the right of the opacity gradation. This will open the *Opacity Filter Editor*, letting you accurately control the profile of the opacity of your color map.

New, Load, Save

Pressing **New** will reset the color map to a fully opaque, black-to-white gradation.


Press **Load** to load one of the sample color maps using the Color Map Browser.

Press **Save** to save the current color map in a stand-alone file, for use in future scenes. Saved color maps will appear in the Color Map Browser. By default, color maps are placed in the *Color maps* subfolder. This means that they will appear in the *Personal* collection inside the Color Map Browser.



Selecting Colors

Quick Color Selection



Quick Color Selector

When you click on a color swatch, the *Quick Color Selector* appears. This panel is designed to let you select a color in just one mouse click. Simply drag the mouse cursor to the desired color and release it to select the color. If you release the mouse button without moving the mouse, the *Quick Color Selector* will be replaced with the *Color Selection* dialog (see below).

The panel displays 4 blocks of colors and a series of color swatches. The 4 blocks represent an approximation of the complete color space. Each block displays all possible values of luminosity and hue for a given color saturation. The top block represents 100% saturated colors, the second, 50% saturation, the third, 25% saturation, and the lower block, is 0% saturation (a gray strip).

The series of color swatches to the left of these color blocks are your “favorite” colors. They can be edited via the *Color Selection* dialog (see below). To select one of your favorite colors, just drag the mouse cursor above that color and release the mouse button. The favorite colors won’t appear if you disabled this option in the *Color Selection* dialog.

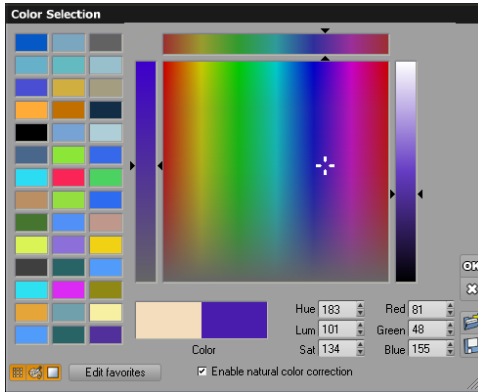
At the bottom of this panel, you will find a display of the current color, together with the corresponding Hue, Luminosity, Saturation, Red, Green and Blue values. If you drag your mouse over another color, the color display will be split in two, the left half showing the original color, and the right half showing the new color.

If you don’t want to use the *Quick Color Selector*, disable the **Use quick color selection**



tool option in the *Color Selection* dialog (see below).

Color Selection Dialog



Color Selection dialog

This dialog lets you select a color precisely. To display this dialog, you can:

- Control-click, Shift-click or double-click on a color swatch, or
- Click on a color swatch and release the mouse button without moving the mouse.

If you have disabled the **Use quick color selection tool** option, this dialog will appear with a simple click on a color swatch.

The large square palette at the center of the dialog represents all possible hue and saturation values for a given luminosity. To the right of this palette is the luminosity slider. Above the palette is the hue slider, and to the left, the saturation slider. When you use one of these sliders, the palette automatically displays all possible variations of colors around the value you just set. For instance, if you set a saturation using the saturation slider to the left of the palette, the palette will display all combinations of hue and luminosity for the saturation value you just indicated. Likewise, if you set a hue using the top slider, the palette will represent all saturation and luminosity possibilities for this hue.

At the bottom left of the dialog are three icons:



Show favorite colors: when this option is selected, the favorite color swatches are displayed on the left of the color palette.



Use quick color selection tool: when this option is selected, the *Quick Color*



Selector appears when you click on a color swatch. Uncheck this option if you don't want to use the *Quick Color Selector*.



Contextual palette: when this option is selected, the color palette automatically adjusts to the values selected in the sliders. If you uncheck this option, the palette will always remain the same, whatever the selections made in the saturation, hue and luminosity sliders.

At the bottom-right of the dialog, you will find a display of the current color, together with the corresponding Hue, Luminosity, Saturation, Red, Green and Blue values. You can set each one of these values manually. If you select another color, the color display will be split in two, the left half showing the original color, and the right half showing the new color.

You can resize the dialog in order to display a larger selection of favorite colors and bigger color selection tools.

Enable natural color selection: Checking this option allocates more space to the medium saturated colors which are the most natural hues to use.

Favorite Colors

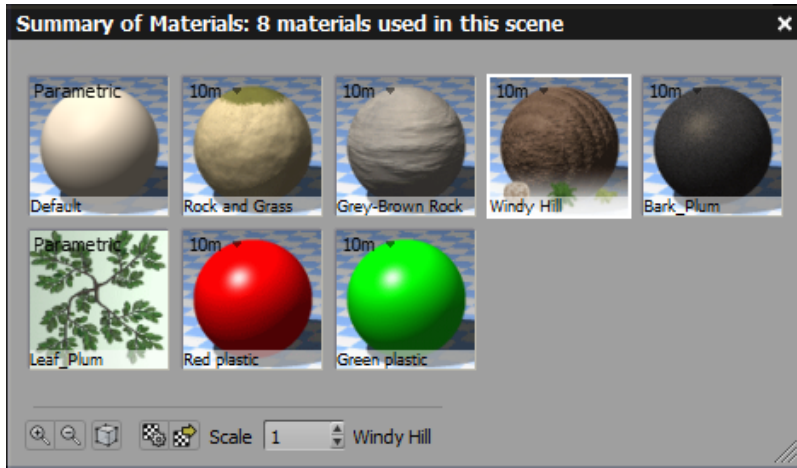
To the left of the dialog is a set of color slots that you can customize to store colors you use frequently. Simply drag a color and drop into one of the slots to set a favorite color. The favorite color selection is saved in between VUE sessions.

You can also edit the selection of favorite colors by clicking the **Edit favorites** button. Using the favorite color editor, you can set the number of columns of favorite colors that appear in the *Color Selection* dialog.

You can also save and retrieve your selection of favorite colors using the **Load** and **Save** buttons.



Summary of Materials



The Material Summary dialog displays materials currently used in the scene

The *Summary of Materials* dialog is a dialog that can stay open on screen without restricting access to other parts of the software. It displays a list of all the materials used in the scene at a given time. It does not, however, display materials used for clouds. It is displayed by selecting the menu command **Display | Display Material Summary**.

This summary of materials lets you approach your scene in a different way: clicking on a material selects all the objects of your scene that use this material (this feature can be turned off using the *Options* dialog). This is a powerful and supplementary method for navigating inside scenes. You may rapidly modify a material (e.g. change its scale), without worrying about all objects that use the material being updated.

And you can drag from one material onto another to make all the objects that used the old material use the new one (notice how the old material, having now become useless, disappears from the list).

The title bar of this dialog indicates how many materials are currently used in the scene. You may navigate through the list using the scrollbar at the bottom of the dialog.



Loading, Editing and Scaling Materials

Like anywhere, materials can be loaded, edited or scaled using the material summary. The display of materials on the *Summary of Materials* dialog can be resized by zooming. Use the **Zoom In/Out** icons at the bottom of the dialog.

The **Options** icon displays the **Preview Options** dialog. This dialog enables you to select which object should be used to preview materials. **Sphere** is the fastest, and **Cloud** should only be used for cloud materials. **XY Plane** displays a 2D representation of the material in perspective, whereas **2D Plane** presents the material on a plane seen from above. This option also lets you choose a background type for the preview (**Uniform** or **Checker**), as well as the **Background color** by modifying the color map (double-click on the map). Check **Local light** to use a local light rather than a directional light.

The **Edit material** icon opens the *Material Editor* for the selected material.

The **Load material** icon opens the Materials Browser so that you can replace the selected material.


You can change the **Scale** of the selected material as well.



Interface Colors

This dialog lets you customize the look of the VUE interface.

Flat Interface Colors



Interface Color Editor

When the **Flat** interface style is selected, you can modify the colors of the different interface items as follows:

- **Gradated background:** Check this setting for a two-toned background, lighter to darker.



- **Dialog background:** this is the background color of dialogs.
- **Text:** this is the color of the dialog text.
- **View background:** this is the color of the background of the *3D Views*. This setting is the same as that in the *Options* dialog.
- **Selection wireframe:** this is the color of a selected item's wireframe view.
- **Wireframe highlight:** if part of a selected wireframe item is selected, this color is used.
- **Rigged mesh bone:** this is the color of the bone of a rigged mesh.
- **Rigged mesh bone helper:** this is the color of a bone helper of a rigged mesh.
- **View caption (active):** this is the background color of the active view's title bar. Inactive view title bars have the same background color as the rest of the interface.
- **View caption text (active):** this is the color of the text in the active view's title bar.
- **View caption text (inactive):** this is the color of the text in inactive view title bars.
- **Caption (active):** this is the background color of active dialog title bars.
- **Caption (inactive):** this is the background color of inactive dialog title bars.
- **Caption text (active):** this is the color of the title text in active dialog title bars.
- **Caption text (inactive):** this is the color of the title text in inactive dialog title bars.
- **Radio/check text (checked):** this is the color of the text of all radio and check-boxes when the control is selected.
- **Radio/check text (default):** this is the color of the text of all radio and check-boxes when the control is not selected.
- **Tab (active):** this is the color of the background of the current tab (usually the same as the background color).
- **Tab (inactive):** this is the color of the background of the non-current tabs.
- **Tab text (active):** this is the color of the text of the current tab.
- **Tab text (inactive):** this is the color of the text of the non-current tabs.
- **Edit field text:** the color of text in edit fields.
- **Edit field text highlighting:** this is the color of selected text in the edit fields.
- **Edit field back (active):** this is the color of the background of active edit fields.



- **Edit field back (inactive):** this is the color of the background of inactive edit fields.
- **Button text:** this is the color of the text of buttons.
- **Button back:** this is the color of the back of buttons.
- **Button back (highlight):** this is the color of the back of buttons when the mouse is above the button.
- **Button back (toggled):** this is the color of the back of buttons when the button is toggled.
- **List item text:** this is the color of text in list boxes.
- **List item text (selected):** this is the color of selected text in list boxes.
- **List back (active):** this is the color of the background of active list boxes.
- **List back (inactive):** this is the color of the background of inactive list boxes.
- **List selected item back:** this is the color of the background of selected text in list boxes.
- **List column title:** this is the color of the background of the column titles in multi-column list boxes.
- **Curve (line):** this is the color of the lines in curve displays (filters and time splines).
- **Curve (below):** this is the fill color below lines in curve displays (filters and time splines).
- **Curve (above):** this is the fill color below lines in curve displays (filters and time splines).
- **Menu text:** this is the color of the text of menu items.
- **Menu text (selected item):** this is the color of the text of selected menu items.
- **Menu back:** this is the color of the menu background.
- **Menu back (selected item):** this is the color of the background of selected menu items.

Changes are previewed interactively in the *Interface Colors* dialog. If you want to preview color changes in the entire interface, click the **Apply** button.



New, Load, Save

Press the **New** icon to reset the interface colors to the default gray interface.

Click Load to load a preset interface color scheme. A Standard File Browser will appear letting you load the desired color scheme. Interface color schemes have the .ics extension.



Click Save to save the current settings for future use. A Standard File Browser will appear letting you set the name of the color scheme. Interface color schemes have the **.ics** extension.




Macros

Ever wished there was a way to automate repetitive tasks in VUE without having to delve into the intricacies of Python scripting? VUE's Macro recording and playback system is the easy solution!

Recording Macros

Record Macro: press the **Record Macro** icon () or use the **Automation | Record Macro** menu command to start recording a set of operations, and save it to disk as a macro. When done, press the **Stop Recording** icon () to stop recording. This will bring up the standard Save dialog where you can enter a title and a description for your macro.


Playing Macros

To play a macro, press the **Play Macro** icon () or use the **Automation | Play Macro** menu command. A Standard File Browser will appear letting you select the macro file that you want to playback.

Recently used macros are listed in the **Automation | Recent Macros** menu.



Macro QuickLaunch



Macro QuickLaunch Editor

You can setup the **Automation | Macro QuickLaunch** menu to list frequently used macros using the *Macro QuickLaunch Editor*. The macros listed in this menu can be assigned to keyboard shortcuts (see [here](#)).

To open the *Macro QuickLaunch Editor*, select the menu command **Automation | Macro QuickLaunch | Edit**.

Click **Add** to add a macro to the **QuickLaunch** menu. A Standard File Browser appears letting you select the macro you wish to add to the list.

To remove a macro from the **QuickLaunch** menu, select the macro on the list and press **Remove**.

You can replace an existing macro with a new one by selecting the macro to be replaced



on the list and pressing **Replace**. The Standard File Browser will appear, letting you select the new file. Replacing macros is particularly useful if you have assigned keyboard shortcuts and do not wish to lose this assignment.



Section 5

EcoSystems





EcoSystem™ is the name of e-on software's revolutionary set of patented technologies to distribute, manage and render millions of instances of plants or objects in your scenes. With this technology, you are able to recreate the millions of plants, trees and rocks that are required to create convincing environments.



Painting EcoSystems

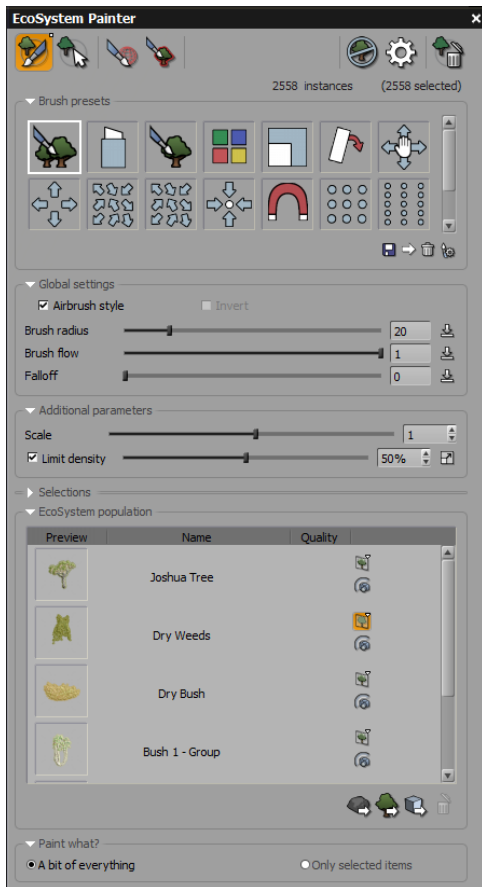
The EcoSystem Generation IV technology lets you paint EcoSystem populations directly onto elements of your scenes and tweak your EcoSystems manually for unprecedented control. Paint from any angle, switch to *Side view*, or spin around the object to paint from the other side.

Thanks to this new technology, you can add, remove and modify EcoSystem instances interactively.

EcoSystem Painter

Global EcoSystems are a way to create EcoSystem populations which do not lie on any specific underlying object. Instead, instances are populated over any object in the scene.





EcoSystem Painting Interface

The *EcoSystem Painter* is the general interface that controls the painting and selecting of EcoSystem instances.


The *EcoSystem Painter* is a tool to act on a subset of an EcoSystem's instances to change some of their properties (position, orientations, scale, color) in powerful ways:

- Any number of effects can be combined together, configured and applied by painting in any of the OpenGL previews of the scene (orthogonal or perspective).
- Painting can be restricted to only specific instances depending on their specimen, material layer or underlying object,
- The *EcoSystem Instance Selector* and its instance selection stack have been merged with the *EcoSystem Painter*, with the added ability that eco-painting can also be optionally applied only on the current selection of instances.



- Configured brushes can be saved in your own collection of EcoBrushes.

You can display this panel in one of two ways:

- Click on the **Paint EcoSystem** icon () in the *Top Toolbar* or select the menu command **Edit | Paint EcoSystem**: this lets you paint using the [Global EcoSystem](#), or
- Click the **Paint** button inside the *Material Editor* when editing an EcoSystem material. This feature requires the [EcoSystem](#) module. This lets you [modify a specific EcoSystem material](#) (and only that EcoSystem).

The *EcoSystem Painter* offers different tools to let you paint instances, as well as ways of controlling what is being painted. Painting can be done in any view.

When a view is active, a circle appears around the mouse pointer to indicate the area where the painting will take place (if no items are selected in the EcoSystem population, a black cross will appear instead).

When using the *EcoSystem Painter*, EcoSystem instances are depicted using a small representation of an object, a rock, or a plant. This representation always appears, even if the billboard preview of the instance is not displayed. The color of the instance indicates the overall color of the item being painted. These representations disappear when you close the *EcoSystem Painter*.

You can edit the EcoSystem population using the *Material Editor* to further define your EcoSystem. After adding the objects, plants and/or rocks to your EcoSystem population, just click to select the **Use EcoSystem population rules** and click on the **Edit** button next to this field. The *Material Editor* opens and you have all of the EcoSystem definition criteria at your disposal to use in defining this global EcoSystem.

EcoSystem Painter Tools

At the top of the *EcoSystem Painter* panel, you can see a series of buttons:



Painting mode on/off: left-click this icon to toggle painting mode. By default, this is turned on.



Select EcoSystem instances: this allows you to select certain EcoSystem instances to act upon. The *EcoSystem Painter* screen displays only the settings for the [selection of instances](#). Click the **Painting mode** icon to return to the regular *EcoSystem Painter* dialog.



This is in addition to the Selections section.



- **Restrict to selected objects:** click this icon to prevent the EcoSystem from being painted on anything other than the objects currently selected in the *World Browser*.



- **Restrict to selected instances:** when this icon is selected, the brush selects certain instances to act upon.



- **Hide from Render:** click this icon to hide the selected items in the EcoSystem from render.



- **Display Options:** this icon is only available when editing a Global EcoSystem. It displays the *Display Options* dialog. Here you can select the display quality of the EcoSystem elements and whether you want to display full quality near camera. If you select that option, you can set the radius limit for full quality.



- **Clear:** click this icon to remove all painted instances from the EcoSystem.



- **Show brush editor:** click this icon to automatically open the *Brush Editor* when that brush preset is selected. If unchecked, you can use this icon to display (or close) the *Brush Editor* as needed.

Fold/unfold the dialog: Click the **Minimize** button in the upper right corner next to the **Close** button to toggle the screen folding. The screen folds to keep it out of your way while painting. Another click restores the screen.

In addition, each section of the EcoSystem Painter dialog and the Brush Editor can be collapsed, reducing the size of the overall dialog. Just click on the downward arrow to the left of the title of each section to close the section. To reopen the section, click the right-facing arrow.



Painting With Brushes

These tools or brushes add instances, delete instances and act upon selected existing EcoSystem instances.

Building the EcoSystem Population

One group of tools is used to build the EcoSystem population. These are the effectors probably used the most in the *EcoPainter* and can only be used alone and not in combination with the other effectors.



Populate: this effector will spawn many instances randomly in the brushed area.



Erase: this effector is used alone and is used to delete EcoSystem instances.



Single instance: this effector, previously available as the “single instance painting mode” prior to VUE 11, can only be used alone. It is used to add or remove instances one by one.

Influencing the EcoSystem Population

These brushes are used to change or influence the current EcoSystem population.



Color: this shifts instances' color closer to the chosen color.



Scale: this lets you change the size of the instances



Lean: bends an instance so that its “up” (Z) axis aligns with the reference axis.



Drag: this drags the instances under the brush along with the brush movement.



Move: this moves a group of instances that are under the brush



Jitter Position: this randomizes instances' position.



Jitter Orientation: this randomizes instances' orientation



Attract: this attracts instances towards the center of the brush.



Magnet: aligns an instance's X axis along the reference axis.



Grid Align: this aligns instances in a grid pattern to the base's X and Y vectors.



Rake: aligns instances along lines parallel to the reference axis



Ripple: aligns instances in widening circular lines



Lower: this lowers instances in relation to the underlying object.



Raise: this raises instances in relation to the underlying object



Tornado: this moves instances upwards and in a spiral around the brush center.



Color & Scale: this combines the features of the **Color** and **Scale** brushes into one brush.



Save brush preset: if you have made changes to a brush preset, click this icon to save your changes to that brush if you wish.





Add brush to presets: if you have made changes to a brush preset, but don't want to change the brush preset, you can save your changes by creating a new brush. Use this icon to save and create a new brush preset.



Remove brush from presets: use this icon to delete any brushes.




Global Settings

- **Airbrush style:** when this option is selected, the number of instances added to a given area is proportional to the time spent painting on that area.
- **Invert:** Not available for all tools. When activated, it inverts the action taken on the instances.
- **Brush radius:** this setting controls the area around the mouse cursor onto which instances are randomly distributed. If you are using a pressure-sensitive tablet, you can connect this setting to the tablet pressure by clicking the **Drive with pressure** icon (). This will result in instances being scattered around further from the mouse when you press harder on the tablet.
- **Brush flow:** this setting controls the number of instances added to the EcoSystem per unit of time (this setting is only available if you are using the airbrush style of brush). If you are using a pressure-sensitive tablet, you can connect this setting to the tablet pressure by clicking the **Drive with pressure** icon (). This will result in instances being added more rapidly when you press harder on the tablet.
- **Falloff:** this is the rate of diminishing effect moving away from the center of the brush. This can also be defined on the *Brush Editor* for this brush using a **Falloff Filter**.
- **Selections:** This allows you to select EcoSystem instances for manipulation. Selection is done in all views. When the view is active, a circle appears around the mouse pointer to indicate the area where the selection will take place. When using the *Selection Tools*, selected EcoSystem instances are depicted using a red dot. This dot always appears, even if the billboard preview of the instance is not displayed. The dots disappear when you close the *Selection Tools*.

Click **Select all** to select all EcoSystem instances and **Deselect all** to deselect all selected EcoSystem instances.




EcoSystem Population


The EcoSystem population list displays a list of all the items that can be painted using the *EcoSystem Painter*. You can add new items to this list using any of the buttons beneath the list. If you select to add a plant () , the Visual Plant Browser opens. Use drag and drop to select a plant from the browser and place it in the *EcoSystem Painter* dialog. The browser will remain open for you to select another plant. Click **OK** in the Visual Plant Browser when finished adding plants. If you select to add an object () , use the same method to select objects for your EcoSystem. If you select to add a rock () , the Rock Template Browser opens. Rocks are added the same way plants are.

You are also able to select any object or plant from the *World Browser* and drag it into the population list to use when painting.

To the right of the EcoSystem item you will see two icons that let you choose how you wish your EcoSystem population to display in OpenGL.

Quality: for Plant Factory plants only, use this setting to control the quality of the EcoSystem specimens. The quality ranges from -4 to 4, the higher the quality value the more detailed the specimens.

Instance previewing mode () : right-click to select the preview quality for this item in the EcoSystem. If you are not in OpenGL Shader mode, the Shaded Billboard option will be grayed out.

Full Quality Near Camera () : click on this icon to allow full display quality of this item near the camera. If you are not painting a global EcoSystem, **Allow full quality near camera** must be selected in the *Material Editor* for this option to be available.

You can remove items from the list by selecting them and pressing **Remove**. If some instances of this item are already in the scene, they will be removed together.

Paint What?

When using the *EcoSystem Painter* tools, you can either act upon all items of the EcoSystem population, or on only a subset of these items. This is controlled using the **Paint what?** options:

- **A bit of everything:** when this option is selected, all items in the population of the EcoSystem will be affected. If you use the “Brush” or “Single instance” tools, the instances added will be randomly picked from the entire population. Likewise, the “Eraser” and ‘Color/Scaling’ tools will erase or modify all the types of elements in the EcoSystem.



- **Only selected items:** when this option is selected, only the items that are selected on the EcoSystem population list will be used when painting. For instance, the “Brush” and “Single instance” tools will only add instances of the selected type. The “Eraser” and “Color/Scaling” tools will erase or modify only the types of elements that are selected on the list. This way, you could have an EcoSystem of rocks and trees, and decide to erase only the trees in some areas.

Once you have set up the *EcoSystem Painter*, you can reduce the screen size clicking the **Minimize** button to keep it out of the way while painting.

If you want to paint, be sure to turn **Painting mode on**.

Global EcoSystem Object

When you have painted the first instance of a Global EcoSystem, be it a single instance or using the brush, a Global EcoSystem object appears in the list of objects in the *World Browser*.

These objects have some of the options available that regular objects do, such as **Hide from render**, **Replace by**, and layer manipulation (**Hide Current Layer**, **Lock Current Layer**, and so on).


This object can be deleted, which will delete all instances of the Global EcoSystem as well. When this object is selected, all instances of the Global EcoSystems are selected and can be manipulated (only when the *EcoSystem Painter* and manipulate dialogs are not open).

If all instances of the Global EcoSystem are deleted from within the *EcoSystem Painter* dialog, the Global EcoSystem object is removed from the object list.

This Global EcoSystem can also be rendered like a mask in the **Multi-Pass** rendering option. To do so, you must enable the **Global EcoSystem** object mask in **Object masks**.



Multiple Global EcoSystem Objects



Two Global EcoSystems

You can now create multiple Global EcoSystems in your scene. If you have already created a global EcoSystem in the EcoPainter, from the VUE menu, select **Object>Create>Global EcoSystem**. This will open the EcoPainter with no EcoSystem population. While other EcoSystems you have created still exist and show up in the World Browser and in renders, you have a chance to create a new one.

Add the objects you need for your EcoSystem, and the first time you paint with it you will see another entry in the World Browser for that new Global EcoSystem.

Each EcoSystem can be manipulated separately. You can see the contents of each EcoSystem listed in the browser.

Using the Paint Function to Modify EcoSystem Materials

You can use the EcoPaint function to modify an EcoSystem Material. This is typically used when you have populated an EcoSystem material, but would like to touch up the way it was populated.

To modify an EcoSystem material, click the **Paint** button inside the *Material Editor* while you are editing the EcoSystem material. The *EcoSystem Painter* panel will appear.

Any modifications made to the EcoSystem population will be limited to instances that belong to that EcoSystem. This way, you do not risk to affect other EcoSystem populations. Modified EcoSystem populations are still bound to the object they are attached to. Moving that object will move the instances along with it.



You can constrain the painting to the object that the EcoSystem is attached to, using the **Restrict to underlying object** option. This avoids adding instances “outside” of the object they are attached to.

If you click on the **Use EcoSystem population rules**, what you paint will correspond to the settings you made for the original EcoSystem.

If the painted EcoSystem is part of a complex material hierarchy (materials with layers, mixed materials), the individual density of the EcoSystem’s layer is computed so that the *EcoSystem Painter* behaves like the **Populate** button. You can easily restrict your EcoSystem painting to a given area by creating a material layer and editing the alpha or by using mixed materials and editing their distribution.

EcoSystem Display Options

Because EcoSystem instances can easily become numerous in a scene, you need to be careful about how they are previewed in the OpenGL displays. On the other hand, if the instances are displayed too crudely, you may have difficulties adjusting your scene.

In order to let you find the optimal balance between display quality and previewing speed, you can adjust the way EcoSystem instances are displayed on a per-EcoSystem basis.

- If you want to adjust the preview quality of an EcoSystem material, go to the **General** tab in the *EcoSystem Material Editor* and use the **Display options** group.
- If you want to adjust the preview quality of the Global EcoSystem, click on the **Display options** button in the *EcoSystem Painter*. The *EcoSystem Display Options* dialog will appear. The controls in this dialog are identical to the EcoSystem material display options found in the *Material Editor*, except for the option **Hide EcoSystem from render**. When this option is selected, the global EcoSystem will be hidden from the render.

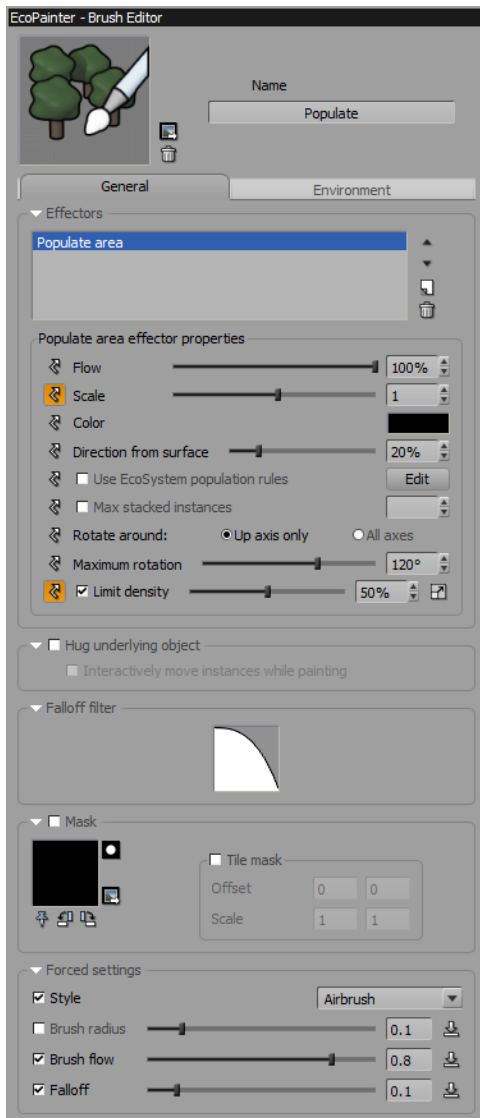
Brush Editor

Every brush has a corresponding *Brush Editor* dialog that provides more parameters for the tool. Not all parameters are available for each brush. These will change depending on the function.

Each section of the Brush Editor can be collapsed, reducing the size of the overall dialog. Just click on the downward arrow to the left of the title of each section to close the section. To reopen the section, click the right-facing arrow.



General Tab




Brush Editor – General tab

- **Effectors:** An effector is an elementary operator which will be applied on all instances found in the brushed area. Its effect will be weighted by a ratio depending on the brush's shape, fall-off, environmental influence, etc. Several effectors can be combined in a single brush to create an unlimited variety of effects. The parameters



displayed under the list of effectors change with the effector. If you have several effectors in the list, highlight the effector to display the fields.

- **Flow:** sets the movement of the brush. Its use may vary between the different brushes.
- **Scale:** sets the size of the instances being painted.
- **Color:** shifts instances' color closer to the chosen color.
- **Direction from surface:** this slider lets you indicate how the EcoSystem instances grow from the surface. If the slider is set to 0%, the instances will always grow vertically, whatever the slope of the underlying terrain is. A value of 100% means that the instances will always grow from the surface (perpendicular to that surface).
- **Use EcoSystem population rules:** this setting allows you to use any of the EcoSystem settings found in the *Material Editor* to define this Global EcoSystem.
- **Edit:** clicking this button opens the *Material Editor* for this Global EcoSystem so that you can set, for example, **Density**, **Scaling and Orientation**, **Color** and **Environment**. Items must already be selected in the EcoSystem population for this to be effective.
- **Max. stacked instances:** you can define how many instances can be stacked.
- **Rotate around:** this lets you define limited rotations of objects. This option is not available for **Area Population**.
- **Up axis only:** select this option if you want the rotation to only take place along the Z axis (vertical). This is typically the case for objects that grow from the surface, such as trees.
- **All axes:** if this option is selected, a random rotation will be applied to all axes of the instances. This is best used for objects that do not grow from the surface (e.g. rocks).
- **Maximum Rotation:** This lets you limit the random rotation that is applied to the instances in the EcoSystem population.
- **Limit Density:** this option lets you impose a limit to the density of instances. If it is enabled, instances will be added by the brush until the maximum allowed density is reached. If the **Grow over max density** icon () is selected, the instances under the brush will begin to swell if you keep painting when maximum density is achieved. This option is not available for **Single Instance** placement.

Note:

Any of these parameters can be moved to the main *EcoPainter* screen to make access easier or if you don't want to keep the brush screen open. Just click on the arrow icon to the left of the parameter and this parameter will display in an **Additional parameters**

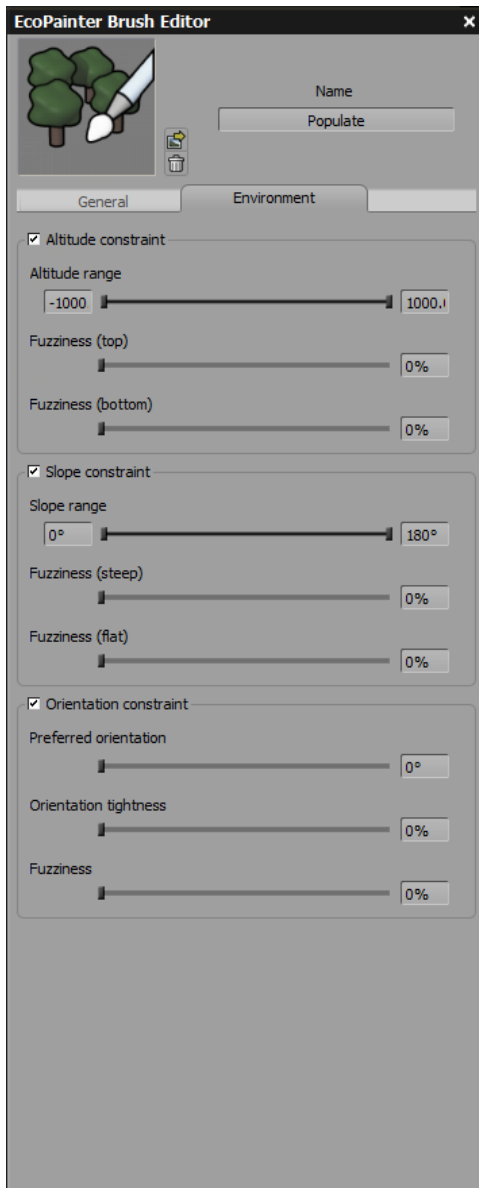


section on the *EcoPainter* screen directly under the **Global settings** section. Uncheck the parameter on the *Brush Editor* screen to delete it from the *EcoPainter* screen.

- **Hug underlying object:** Some effectors will move instances above or below the EcoSystem's underlying object. It is sometimes the case by design, but it can also be a side-effect of the brushing process. Therefore, there is an option to relocate instances on top of the underlying object. Since it can take some time to make these adjustments for all influenced instances, you have an option to handle in real-time or at the end of each stroke. Check **Interactively move instances while painting** to make changes in real-time.
- **Falloff Filter:** Fall-off curve is defined by a filter and tells how the influence (usually) decreases when the instance's distance to the brush center increases. The Filter Browser opens when you click on the filter image. Right-click on the filter to edit.
- **Mask:** You can assign a bitmap mask to this brush. It can be either stretched to exactly fill the brush area, or tiled over the whole underlying object (with custom **Offset** and **Scales** applying). This image can be inverted or rotated.
- **Forced settings:** These can differ from the global settings on the *EcoPainter* dialog.
- **Style:** Airbrush or Paintbrush
- **Brush Radius:** this is the brush size. Defines the area of influence.
- **Brush Flow:** sets the amount of effect applied by the brush. Its setting and use may vary between the different brushes.
- **Falloff:** this parameter (in the $[0; 1]$ range) is the radius ratio above which the fall-off is actually applied. For example, with a radius of 100 pixels, a fall-off of 0.0 means that the filter will be mapped from 0 to 100 pixels, while at 1.0, there is no fall-off at all. At 0.6, for example, the fall-off filter is mapped from a distance of 60 pixels to the maximum influenced distance of 100 pixels.



Environment Tab



Brush Editor – Environment tab

The settings on this tab define **Altitude**, **Slope** and **Orientation** constraints, restricting




the brush's effective area. Environment information like altitude/height, slope or orientation of the underlying object's geometry under each instance can be used to weight the brushing process's application.

- **Altitude constraint:** This group lets you control how altitude influences the presence of instances:
 - **Altitude range:** this dual slider lets you define the range of altitudes where the instances appear.
 - **Fuzziness:** this setting controls how “suddenly” the changes to the instances presence are made in response to altitude. High values mean that the instances appear very gradually in its altitude range, whereas low values will result in the instances appearing as a solid strip.
- **Slope constraint:** This group lets you control how the local slope influences the presence of instances:
 - **Slope range:** this dual slider lets you define the ranges of slopes in which the instances appear. The instances will not appear outside this range. Values to the right end of the slider indicate flat surfaces, and values to the left indicate upside-down surfaces. Intermediate values indicate vertical surfaces. Slope values can range from -180 to +180 degrees.
 - **Fuzziness:** this setting controls how “suddenly” the changes to the instances presence are made in response to slope. High values mean that the instances appear very gradually in its slope range, whereas low values will result in the instances appearing as a solid strip on areas of appropriate slope.
- **Orientation constraint:** This group lets you control how the local orientation influences the presence of instances:
 - **Preferred orientation:** this setting controls the orientation of the surface that is the most favorable to the presence of instances.
 - **Orientation tightness:** this setting controls the influence of orientation on the presence of instances.
 - **Fuzziness:** this setting controls how “suddenly” the changes to the instance presence are made in response to orientation. High values mean that instances appear very gradually on surfaces of the preferred orientation, whereas low values will result in the instances appearing as a solid strip on areas of preferred orientation.



Creating an EcoSystem Content Brush



Two global EcoSystems

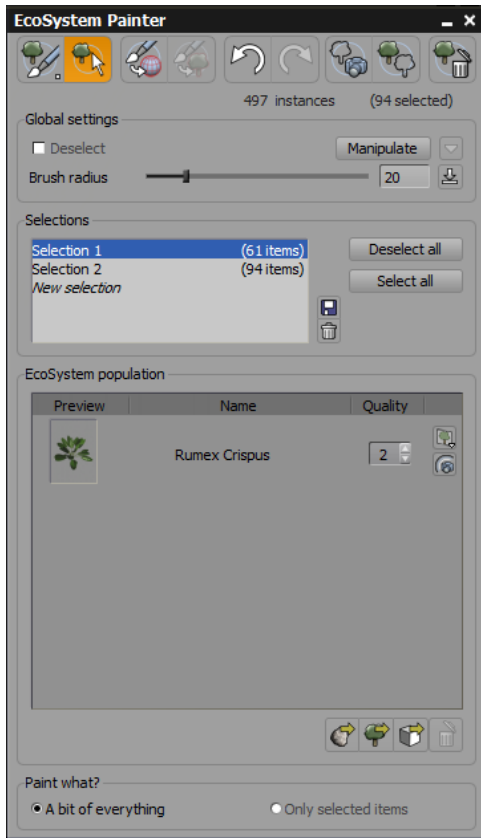
Once you have created a Global Ecosystem in the EcoPainter, you can now save it as a brush, complete with all contents, to be reused as you wish. Click on the **Save Brush Preset** icon.

It will then show in the World Browser as a Global EcoSystem. The contents of the EcoSystem will be displayed as well.

To create another Global EcoSystem, from the menu select Object>Create>Global EcoSystem. This will either clear the Population area of the EcoPainter window if you have it open, or it will open the EcoPainter dialog.




Selecting EcoSystem Instances



Selection Tools

The **Selections** section provides more options to manipulate EcoSystem instances. Selection can be done in all views. When the view is active, a circle appears around the mouse pointer to indicate the area where the selection will take place.


This feature becomes active by selecting the **Select EcoSystem Instances** icon () at top of the *EcoSystem Painter* dialog. Selection can be done in all views. When the view is active, a circle appears around the mouse pointer to indicate the area where the selection will take place.

Selected EcoSystem instances are depicted using a red dot. This dot always appears on selected instances, even if the billboard preview of the instance is not displayed. The dots



disappear when you close the *Selection Tools*.


You can invert the selection by selecting the **Inverse sel** button. This will deselect the current selected items and will select all the other items.


To return to the full *EcoPainter* dialog, just select the **Painting mode** icon () in the upper left of the *EcoPainter* dialog.

Note:

the *Selection Tools* make no distinction between the “origin” of EcoSystem instances (meaning that it will indistinctively select instances from all EcoSystem materials and the Global EcoSystem).

Selecting EcoSystem Instances

To select EcoSystem instances, choose the **Select EcoSystem Instances** icon () at the top of the *EcoPainter* screen. The brush will select all instances under the brush (depicted by a red circle in the view). Simply hold the mouse button down while you drag the brush over instances to select them, or click on individual instances. Conversely, checking the **Deselect** button will deselect all of the instances that the brush clicks on or brushes over.

Brush radius: this setting controls the area around the mouse cursor where instances are selected or deselected. If you are using a pressure-sensitive tablet, you can connect this setting to the tablet pressure by clicking the **Drive with pressure** icon (). This will result in instances being selected further away from the mouse when you press harder on the tablet.

In the **Selections** section, there are several buttons:

- **Select none:** click to deselect all instances in this EcoSystem.
- **Select all:** click to select all instances in this EcoSystem.
- **Save:** when you have selected some instances, you can select the **Save** button to store your current selection in that selection slot (and create a new slot).
- **Discard:** deletes the previously saved instances in this slot, but doesn't remove them from the actual EcoSystem.

Note:

selections are saved together with the scene. However, if you modify the EcoSystems in the scene, the selection may no longer correspond to the items originally selected.



Manipulating Selected Instances

Once you have selected some EcoSystem instances, the **Manipulate** tool becomes available. If you select this tool, you will be able to manipulate the selected instances as if they were standard objects. The standard manipulation tools will appear in the *3D Views*, so that you can move, rotate or resize the selected instances. You can also use the [Numerics](#) tab of the *Object Properties* panel.

Alongside the **Manipulate** option, the **EcoSystem operations** button displays a popup menu with a list of operations that can be made on EcoSystem instances.

Convert to Objects

If you select this option the selected instances will be converted into standard objects and will appear in the *World Browser* alongside other objects in the scene. When an EcoSystem instance has been converted into a standard object, you can edit it like any other object. In the *3D View* popup menu, when such converted objects are selected, you will have the option to put them back into the EcoSystem they belonged to, using the **Revert to Instances** command that appears at the top of the *3D View* popup menu. Using this pair of commands, you can move instances in and out of EcoSystems – e.g. when you need an accurate display for a given EcoSystem instance.

Moving Instances to another EcoSystem

The next command on the popup menu lets you move instances to another EcoSystem.

Some EcoSystem instances belong to the Global EcoSystem (see [here](#)) while others belong to an EcoSystem material, and are attached to specific objects in the scene. This command lets you change the EcoSystem to which the instances belong. For instance, if you placed some rocks on top of a terrain using the Global EcoSystem, but later find out that you would like to move the terrain, you will find out that the rocks don't follow when you move the terrain. To fix this problem, you can assign an EcoSystem material to the terrain and move the instances to that EcoSystem using this menu command. Now, when you move the terrain, the rocks follow. The popup menu lists all the different EcoSystems materials used in the scene, together with the name of each object that uses said EcoSystem material.

Changing the Type of EcoSystem Instances

The last command on the popup menu lets you change the type of the instances.


This command lets you change the type of the selected instance into any other item in any EcoSystem population throughout the entire scene. For instance, if you have an EcoSystem of trees, and another EcoSystem of rocks, you could select some of the trees and convert them into rocks! This is an extremely powerful command to touch up the automatic distribution of the different items in the population of an EcoSystem material.



Using the Brush Tools on Selected Instances


Once you have selected instances and saved them, you can now use any of the Brush presets on those instances, leaving unselected instances untouched.

The **Selections** frame is also available in **Painting mode** to make it easier to use that feature:

- make several selections you will want to paint on and save them,
- switch to **Painting mode**,
- select the **Restrict to selected instances** icon () at the top of the *EcoSystem Painter* dialog,
- choose a selection from the list of saved selections,
- paint on it with some brush,
- choose another selection,
- paint on it with some other brush.



The Vue EcoParticle System



Material Editor – EcoParticles

VUE's EcoParticle system is based on our EcoSystem Technology. While you find some settings of the EcoParticle system in various areas of the user interface, most of the creation and setup of the EcoParticle system is done in the *Material Editor*.




Setting up an EcoParticle System



EcoParticle systems use many of the same settings as other EcoSystems and you can refer to the [EcoSystems](#) section for any field definitions that you need, but there are certain fields for EcoParticle systems only.

General Tab of the Material Editor

Most of the settings for EcoParticle systems are found on this tab.

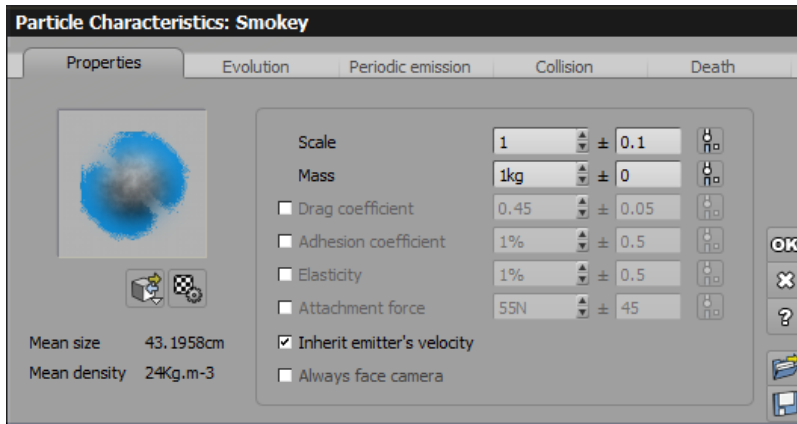
First of all, under EcoSystem population, when you add an EcoParticle, you have to check the **Particle motion** box. If the box isn't checked, then the EcoParticle is considered a regular EcoSystem instance. You have your setting for **Presence** just as a regular EcoSystem, but you now have a setting for **Velocity** of the EcoParticle. This can be connected to a function () created in the *Function Graph*.

You also have the **Instance Previewing** mode. For EcoParticles, this is set to **Billboard** for time and resources considerations. The **Full Quality Near Camera** mode is available as well.

Below the **EcoSystem population** box, the **Edit material** icon () is also available for EcoParticles. The fourth icon from the left is the **Edit specimen** icon (). Clicking this opens the *Particle Characteristics* dialog.



Particle Characteristics Dialog



Particle Characteristics dialog

This dialog defines the characteristics of the individual EcoParticle.

Note:

Many of the fields on these tabs can be connected to functions that can be refined in the *Function Graph*. Also many of the fields have two fields for data entry. The first field is the base value. The second field allows for slight variation, for example a value of 2s in the first box and 0.1 in the second would mean a time of 2 seconds, plus or minus 0.1 second.

Properties Tab

The Properties tab contains a preview of the rock, plant or object currently selected in the **EcoSystem population** list. To replace this specimen with another, click the **Replace** icon (🔄) under the preview. There is also an **Edit material** icon (⚙️) to open the *Material Editor* for any changes you might wish to make in the material of the EcoParticle.

The **Mean size** and **Mean density** of the EcoParticle are also displayed under the preview.

The following fields can be connected to a function in the *Function Graph*. The value in the second column is the “plus or minus” factor for the value in the first column allowing for slight variation.

- **Scale:** the size of the EcoParticle



- **Mass:** the weight of the EcoParticle
- **Drag coefficient:** defines how easily the EcoParticle will penetrate into a medium (0.5 for a sphere, 1 for a cube, 0.04 for a plane wing, for example). This value cannot be 0.
- **Adhesion coeff:** this shows the extent of adhesion variability in relation to the adhesion of the mean of the population.
- **Elasticity:** indicates whether the EcoParticle has the property of elasticity
- **Attachment force:** indicates the force to unstick an EcoParticle from another one
- **Inherit emitter velocity:** velocity will be determined by the value set in the emitter.
- **Always face camera:** If the particle is not a spherical-type object, for example, a billboard, check this to always keep it facing forward.

Evolution Tab

- **Life size variation:** linear size variation during EcoParticle's life
- **Grow by:** amount of growth throughout EcoParticle's life
- **Altitude Size variation:** linear size variation based on altitude
- **Grow by:** amount of growth
- **Every:** the distance the altitude change is to take place
- **Opacity variation:** used for EcoParticles that are dependent on an alpha channel for effect, for example, smoke and fire. Usually EcoParticles will be more dense at the bottom and more transparent at the top.
- **Final opacity:** amount of opacity at the end of the EcoParticle's life
- **Relative to size variation:** check if this change is dependent on size variation of the EcoParticle
- **Delete particle if opacity <:** If you wish to delete the EcoParticle as it becomes more invisible, indicate the percentage of opacity where it should be deleted.
- **Color variation:** used for size variation as EcoParticle changes from one color to another. The initial color is found in the *Material Editor*. Select the final color using the color field on this tab.

Periodic Emission Tab

- **Specimen:** select a specimen to be periodically emitted
- **Velocity:** local emission starting velocity.



- **Flow:** this is the frequency of new EcoParticle creation. For example, a value of 5 plus/minus 2 means that between 3 and 7 new EcoParticles are created each second.

Collision Tab


- **Die on collision:** if the EcoParticle should die on collision, check this option. If not, select the change in the EcoParticle.
- **Specimen:** select a specimen created after collision
- **Velocity:** this is the velocity of the EcoParticle after collision
- **Count:** this is the count of new EcoParticles after collision

Death Tab

- **Life duration:** enter the life duration of the EcoParticle in seconds.
- **Upon death clone into:** if you want the EcoParticle to change into another format, check this and select what you want the EcoParticle to change into
- **On instance end emit:**
 - Specimen
 - select a specimen created after death: Velocity
 - this is the velocity of the EcoParticle after cloning: Count
 - this is the count of new EcoParticles after cloning.:

Bake Particle Motion Dialog

Back on the **General** tab of the *Material Editor*, for EcoParticle systems the next field is **Particle motion**. There are two buttons under Particle motion:



Bake Particle Motion dialog

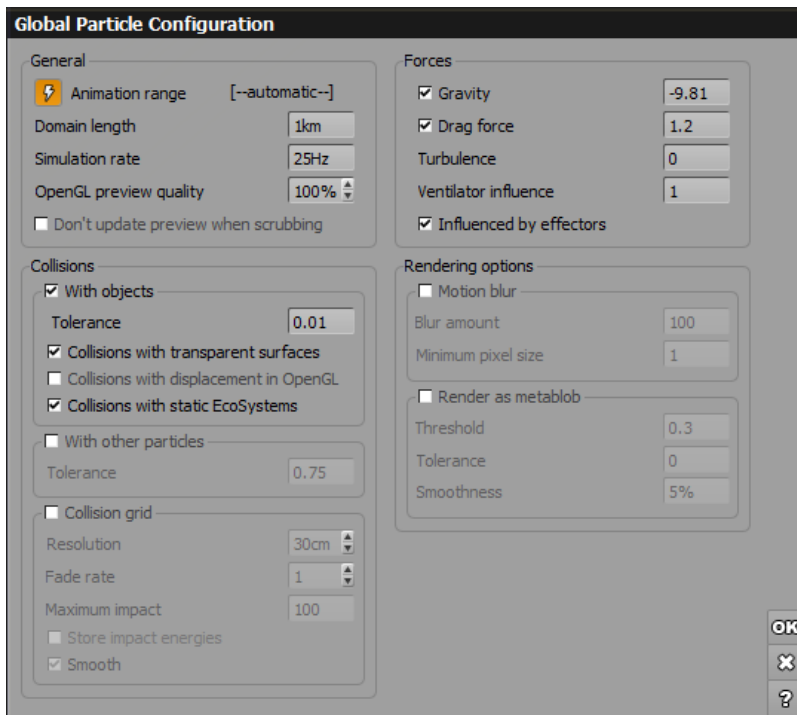
- **Bake:** click on this button to bake the physics of the EcoParticle system. While this is time consuming when selected, it will save time later during the animation phase. This opens the *Physics bake* dialog which has the following parameters:



- **Start time:** set the start time for baking
- **End time:** set the end time for baking
- **Frames per second:** set the rate of frames per second.
- **Edit:** clicking on this button displays the *Global Particle Configuration* dialog. This sets the general configuration for the actual EcoParticle.

Global Particle Configuration Dialog

General



Global Particle Configuration dialog

- **Animation range:** Automatic is checked by default. If unchecked, the default values are from 0 to 10.
- **Domain Length:** the defined area where the EcoParticle system will work.
- **Simulation rate:** this value is used to control strange or unstable behavior of the EcoParticle system. Raising the value will increase frequency of calculations and



stabilize the EcoParticle system.

- **OpenGL preview quality:** the higher the quality, the longer the preview times
- **Don't update preview when scrubbing:** this allows you to render out a preview without all of the physics calculations which speeds up your EcoParticles preview.



Collision

- **With objects:** this is collision with other objects in the scene that have collision enabled.
- **Tolerance:** the maximum amount of interpenetration; the precision of the collisions
- **Collisions with transparent surface:** defines the amount of collision with transparent objects
- **Collisions with displacement in OpenGL:** when colliding with an object, takes into account any displacement of the object.
- **Collisions with static EcoSystems:** allows collisions with EcoSystem instances in a scene. This only applies to static EcoSystems. Collisions are not available for dynamic and EcoParticle systems.
- **With other particles:** defines the amount of collision among the EcoParticles themselves
- **Tolerance:** the amount of interpenetration
- **Collision grid:** this is a world 3D grid encompassing the scene, storing collision density
- **Resolution:** the size of the individual cells in the grid
- **Fade rate:** the amount of gradual loss of intensity
- **Maximum impact:** this is the maximum force of an EcoParticle collision.
- **Store impact energies:** when selected, the collision grid will store energy instead of collision count. For energy, the more the particle is fast and heavy, the more energy will be stored at collision.
- **Smooth:** the smoothness of EcoParticle motion

Forces

- **Gravity:** the amount of gravity affecting the EcoParticles. -9.81 is the default value
- **Drag force:** the amount of drag on the EcoParticles. Default is 1.2
- **Turbulence:** the amount of turbulence applied to the forces of gravity and drag.

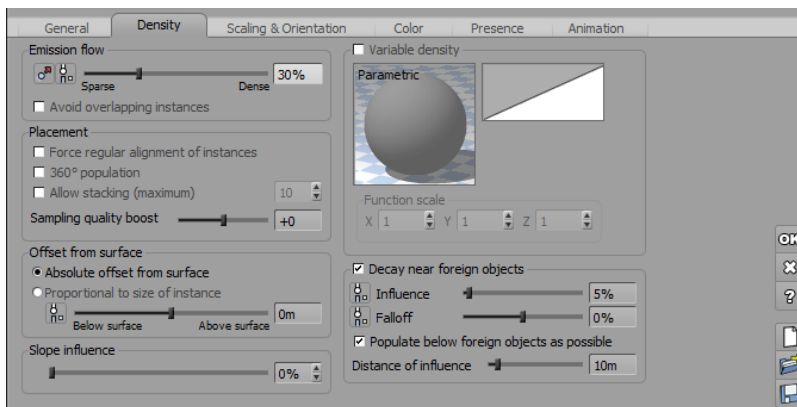


- **Ventilators influence:** the amount of influence the ventilators have. Directional ventilators are set from the icon () on the left toolbar. This feature requires the *Zephyr* module.
- **Influenced by effectors:** check to have influence from effectors. Particle effectors are set from the icon () on the left toolbar in the user interface. Once selected, the **Particle Effector** appears in the *World Browser*. When you click on it, it displays the *Effector Editor*, discussed a bit later in this section.

Rendering

- **Motion Blur:** check to enable motion blur.
- **Blur amount:** set the amount of blur.
- **Minimum pixel size:** the smallest an EcoParticle can be
- **Render as metablob:** used for effects such as flowing water
- **Threshold:** this is used to inflate or deflate the surface that is being rendered by “blobbing” the EcoParticles together. This smoothing allows more control on the final shape. The higher the threshold, the less transition between the individual spheres inside the blob.
- **Tolerance:** you can increase this value to speed up render but may cause errors.
- **Smoothness:** the smoothness of EcoParticle motion.



Density Tab of the Material Editor



Density tab – Particles



The first field on this tab is for EcoParticle systems.

- **Emission flow:** this setting increases or decreases the amount of EcoParticles, setting from **Sparse** to **Dense**. There are two icons to the left of this field:
- **Populate once only:** (): the EcoParticles are only populated once when checked, not continuously flowing.
- **Drive with a function:** (): click to open the *Function Graph* and set up a function to control emission flow.

These fields are basic EcoSystem fields, but have special considerations for EcoParticle systems.

- **Decay near foreign objects:** while this works for EcoParticle systems the same as it would for other EcoSystems, it is a good idea to turn this off for EcoParticle systems as it will speed up the population of the EcoParticles. Of course if it's necessary for an effect, it can be used.


Scaling & Orientation Tab of the Material Editor

- **Direction from surface:** to get a correct orientation on instances, one should set this to 100% perpendicular to the surface (vertical gives strange, unrealistic results).
- **Maximum rotation:** with EcoParticle systems, it's a good idea to set this to zero. This speeds up rendering for simple EcoParticle geometries.

For information on any other EcoSystem fields, please refer to the EcoSystems section of the *Material Editor* beginning [here](#).

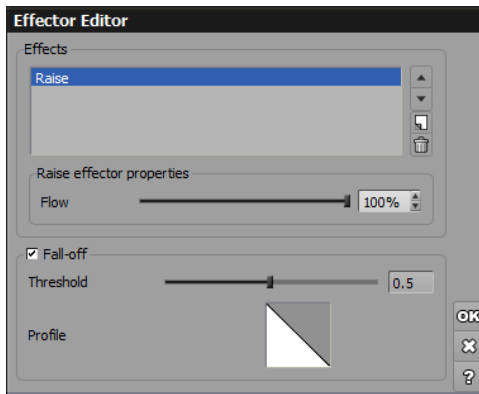
The Particles Effector

The **Particles Effector** quite simply effects EcoParticles. It is not necessary to use it but by choosing different effects or combinations of effects, you can force your EcoParticles to behave quite differently. If you do use the **Effector**, be sure to click that option in the *Global Particle Configuration* dialog.


The **Particle Effector** is created by clicking on its icon () in the left toolbar of the user interface. This creates an invisible proxy object in your scene and an entry in the *World Browser*. Click on either to display the *Effector Editor*.



Effector Editor



Effector Editor

- **Effects:** use the **Add a new effector...** icon () on the right side of this dialog to display the list of available effectors. Each effector has its own parameters.
- **Fall-off:** check for fall-off of the effector influence.
- **Threshold:** use the slider to set the amount of fall-off.
- **Profile:** right-click to open the *Filter Editor* to change the look of the fall-off. Left-click to open the Filter Browser to select a particular filter.

Animating an EcoParticle System

Animating a EcoParticle system is really no different than creating other animations, except for the computations.

On the *Global Particle Configuration* dialog, be sure to check the **Don't update preview when scrubbing** field. This will prevent physics calculations being made during the preview which will save a lot of time. However, it also makes the **Recompute** button unavailable. So if you want to make changes to your EcoParticle system, untick this box so the population can be recomputed.

If rendering another preview, check that box again. Leave in unchecked for final renders.

Limitations of the EcoParticle System

The following are limitations of the current EcoParticle system:

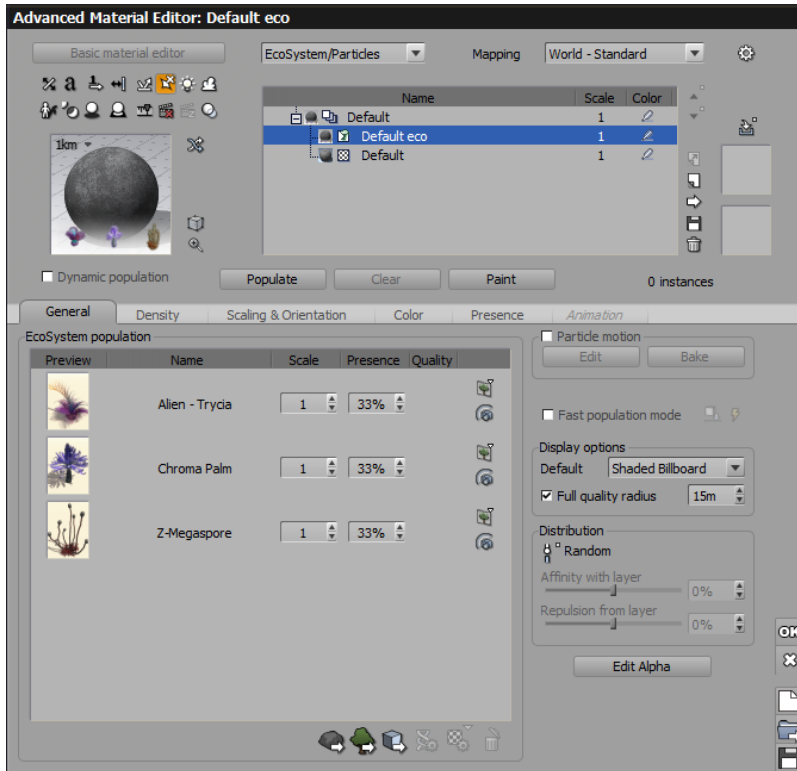
- Physical attributes are computed on the EcoParticle center (no way to rotate an EcoParticle at collision time) and collision takes only the spherical radius into account.



- There is no mesh triangle collision.
- There is no spring system for cloth simulation.
- There is no fluid simulation for realistic smoke or fluid movement computation.



EcoSystems in the Material Editor



Material Editor – EcoSystem Materials

EcoSystem materials are used to scatter instances of an EcoSystem population at the surface of objects. EcoSystem materials behave as standard materials. They can be mixed together using the [Mixed Materials](#) option, and will react to the environment just as other materials would (see [Vue's EcoSystem technology](#)).

EcoSystems can be layered like simple materials, and you can define the affinity between the different [EcoSystem](#) layers.

When you create an EcoSystem material, it will always appear as a multi-layer material, with the underlying material being the lowest layer on the stack, and the EcoSystem immediately above.

Just above the tab control, there are three buttons that are used to populate, paint or clear the EcoSystem population. The **Populate** and **Clear** buttons will also appear on mixed materials, if one or several materials in the material hierarchy is an EcoSystem:

- **Populate:** press this button to generate the instances of the EcoSystem population according to the EcoSystem material settings and the geometry of the underlying object. This button displays as **Preview** if the **Dynamic population** option is checked.
- **Clear:** press this button to remove all EcoSystem population instances.
- **Paint:** press this button to display the [EcoSystem Painter](#) tool. You can now use the **Paint** tool to apply the EcoSystem, or selected items from the EcoSystem, to the selected surfaces.
- **Dynamic population:** when you select the Dynamic population box, Vue will generate a few instances to let you preview the population close to the camera. Dynamic population actually takes place at render time. But it can be previewed at any time by selecting the **Preview** button next to the **Dynamic population** checkbox. The maximum number of instances and the minimum size in pixels previewed can be set in the **Display** tab of the *Options* panel. This option is extremely useful if you want to populate vast expanses of land (or even infinite planes). Vue features a number of very elaborate algorithms to automatically distribute a potentially infinite number of instances only on the areas that are really “seen” by the camera. This is an extremely efficient technique for handling very large EcoSystem populations.

Alongside these buttons is an indication of the current number of instances in the population.

Note:

Populating an EcoSystem usually involves a fair amount of randomness. This “randomness” is controlled in such a way as to avoid radical changes in the placement of the EcoSystem population after small changes in the material settings: pressing **Populate** after slight changes in the material settings will only cause slight changes in the actual population. If you are not happy with the way the EcoSystem population is distributed and would like to see another distribution for that EcoSystem, press **Clear** followed by **Populate**. This “new” population will still follow the rules of the EcoSystem material, but with a different distribution of the EcoSystem population.

EcoSystem materials are controlled through 6 tabs:

- **General:** this tab is used to define the EcoSystem population (objects that are placed at the surface of the underlying object), as well as the aspect of the surface of the underlying object.
- **Density:** this tab controls how the EcoSystem population is distributed at the surface of the object.



- **Scaling & Orientation:** this tab controls the size of the EcoSystem population and how this population is oriented relative to the underlying object's geometry.
- **Color:** this tab is used to define the variations in color of the EcoSystem population.
- **Presence:** this tab lets you control how the environment affects the presence of the current layer.
- **Animation:** this tab has setting to allow you to control the phasing of animated EcoSystem instances.

Temporary Global Settings for Quality Display of EcoSystems

If you notice a slowdown when further editing your scene, you can quickly uncheck the option to **Allow Full Quality Near Camera** which will override all settings made in the EcoSystems used in this scene. This clears all OpenGL data for the EcoSystems and reduce the scene overhead while you are working. This setting is available from the menu, **Display | EcoSystem Preview**.

Temporary limitations can also be set for the **Global Quality Limit**, overriding the settings made in the EcoSystems used in this scene, also available from the menu **Display | EcoSystem Preview**. You can choose to:

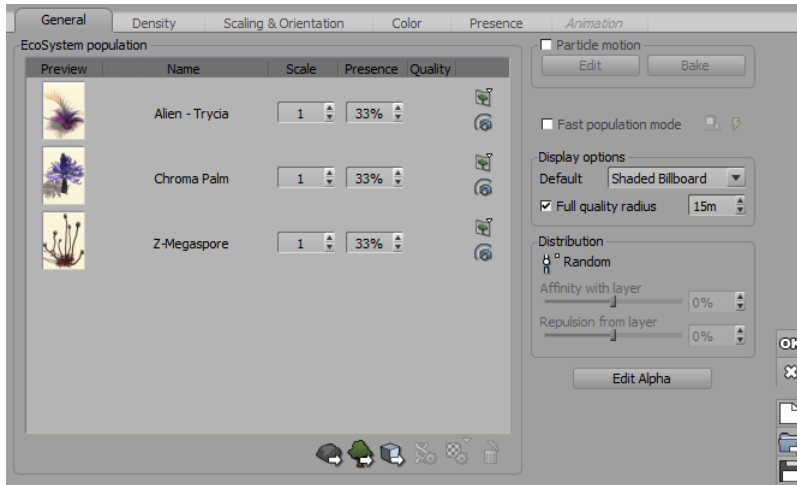
- Limit to None
- Limit to Flat Billboards
- Limit to Shaded Billboards (only available with the OpenGL Shader settings)

If you need to go back and work more on the individual EcoSystems, you can always uncheck these global options to return to the settings in the individual EcoSystems.



General Tab

EcoSystem Population



General tab – EcoSystem Materials

The large list in this tab is used to indicate what the EcoSystem population is made of. You can add as many objects, plants or rocks to this list as you want. When you populate the EcoSystem, Vue will select items from this list and add them to the surface of the underlying object.

If you are creating an EcoParticles system, many of the regular EcoSystem settings will apply. Some fields, however, are [EcoParticle specific](#).

To add a new item, press any of the buttons below the list;



Rock: select this option to add rocks to the EcoSystem population. The Rock Template Browser will appear, letting you select the type of rock you want using drag and drop to add it to the EcoSystem population list. The browser will remain open for you to select a rock. Click **OK** in the Rock Template Browser when finished adding rocks. When you populate the EcoSystem, variations of each type of rock will be created.



Plant: select this option to add a new plant species to the EcoSystem population. The Visual Plant Browser will appear, letting you select the desired plant species. You can use drag and drop to select a plant from the browser and place it in the EcoSystem population area. The browser will remain open for you to select another plant. Click **OK** in the Visual Plant Browser when finished adding plants. When you populate the EcoSystem, SolidGrowth will automatically create variations of the desired plant species, to avoid duplicated plants appearing in the EcoSystem.



Object: select this option to add a new object to the EcoSystem population. The Visual Object Browser will appear, letting you select the desired object. You can use drag and drop to select an object from the browser and place it in the EcoSystem population area. The browser will remain open for you to select another object. Click **OK** in the Visual Object Browser when finished adding objects. You can also import 3D objects from other applications for use directly inside the EcoSystem population.

If you have created a **.vob** file of a Vue rock or tree, you can use these in EcoSystems and Vue will create variations of that rock or tree automatically. For plants, you must have created a **.vob** file of a Vue tree. You can use these in EcoSystems and Vue will create variations of the tree automatically. You must have the original **.veg** plant that this **.vob** was created from however.

When the new item is loaded/prepared for use in the EcoSystem, it appears on the list. If you want to remove some items from the EcoSystem population, select them and press **Remove**.

Note:

If you want to create EcoSystems that use rocks with different materials, you first need to create Vue objects from those rocks: create a rock, map it with the desired material and save it as a Vue object. When you load this rock object into an EcoSystem, Vue will automatically detect the rock, and create variations of it using the desired material.

Among the icons found under the EcoSystem population list is the **Edit Material** icon. Highlight the item in the list you wish to edit and click on the icon to open the *Material Editor* for that item. If the item you have placed in the EcoSystem consists of several materials, right click on the item to select the material you wish to modify.

If this is an **EcoParticle system**, highlight the EcoParticle and the **Edit Specimen** icon becomes available.

Items are displayed on the list as a preview image, the name of the item, its overall scale and it's presence in the EcoSystem population.

Use the **Scale** setting to adjust the average size of instances of a given item in the EcoSys-



tem population. This is particularly useful to adjust the relative sizes of different items on the list. The overall size of the entire population can also be controlled using the **Overall scaling** parameter in the **Scaling & Orientation** tab.

The **Presence** setting lets you adjust how “often” the item is present in the final EcoSystem population. If you increase this setting, you will see this item more often in the population.



The **Quality** setting is included for Tree Factory plants. Since these plants can be more complex, this setting allows you set the quality for inclusion in an EcoSystem.

Note:

Because the presence setting is relative to the other items on the EcoSystem population list, increasing the presence of all items on the list does not increase the number of instances in the EcoSystem population.

If this is an EcoParticle system, there is also a field for **Velocity**.


If the Distribution is driven by a function instead of being random, the Presence setting indicates the output interval in which the item is present (see below).

- **Quality:** for Plant Factory plants only, use this setting to control the quality of the EcoSystem specimens. The quality ranges from -4 to 4, the higher the quality value the more detailed the specimens.
- **Instances Preview Quality:** () Right-click on this icon to select the preview quality for this particular item in the EcoSystem. This overrides the **Default** quality set up in the **Display options** on this same tab. If you are not running in OpenGL Shader mode, the **Shaded Billboard** option will be grayed out.
- **Full Quality Near Camera:** () If **Allow full quality near camera** is selected, you can click on this icon to deselect this EcoSystem item from that setting.

Underlying Material

This is the material that appears at the surface of the underlying object. For instance, if you populate a terrain with plants, you would probably have the underlying material set to some kind of soil material. In the [Material Hierarchy](#), the underlying material appears as a layer underneath the EcoSystem.

If you switch from another type of material to an EcoSystem, the underlying material will be the same as the material before switching.

Double click on the material preview to edit the underlying material, or press the **Load material** button () to load an existing material. You can adjust the scale of the material using the **Scale** setting below the material.





Fast Population Mode

This is a simplified (and much faster) population mode.

When this mode is activated, Vue can refresh the EcoSystem population interactively. The population is updated as soon as you change a value in the *Material Editor*, for example, when you change a value or change a parameter.

To activate, just check the option **Fast population mode** on the **General** tab of the *EcoSystem Material Editor*.

By default, the population is interactive, but you can disable it by clicking on the **Interactive Population** icon () . You might want to disable this if the population takes too long. This could happen if you populate a huge terrain and connect the density to a very complex fractal.

When you are satisfied with the current population, press the **Accurately reposition on surface** icon () to reposition the instances accurately. This does not change the count of instances, but it adjusts the altitude of the instances so that they match exactly the surface of the populated object.

If using this mode, some features will not be available because of time considerations. These unavailable features are:

- EcoSystem stacking
- Population over displaced objects
- Avoid overlapping instances
- Affinity/Repulsion with/from EcoSystem layer

Display Options



This group of controls let you select how the instances of this EcoSystem are displayed in the *3D Views*.

Note:

the display options only affect the way the instances are displayed in the *3D Views*. They do not affect the way the instances are rendered in the final image.

- **Default quality:** this controls the default display quality of the instances in the EcoSystem. The drop-down list offers the following options:
 - **None:** the instances in the EcoSystem are not displayed in the *3D Views*.
 - **Billboard:** the instances are displayed using 3 billboards aligned along each one of the axes – this is the default method for displaying EcoSystem instances. This is the preferred setting for EcoParticle systems




- **Shaded Billboard:** the instances are displayed in OpenGL as billboards with full preview lighting including shadows. This option is only available if you are using the **OpenGL 2.1 (shader 4)** display quality setting **Options | Display Tab**.
 - **Wireframe Box:** the instances are displayed as wireframe boxes.
 - **Filled Box:** the instances are displayed as solid boxes.
 - **Wireframe:** the instances are displayed with their full 3D geometry, as a wireframe.
 - **Flat Shaded:** the instances are displayed with their full, flat shaded 3D geometry.
 - **Smooth Shaded:** the instances are displayed with their full, smooth shaded 3D geometry. This is the best possible display quality.
- **Allow full quality near camera:** when this option is selected, the instances of the EcoSystem that are close to the camera may be displayed in full, smooth shaded 3D geometry. You can adjust the radius around the camera in which instances are displayed at full quality, using the **Radius** setting. Avoid using a large radius setting, as this will slow the preview down significantly.
If all of the instances are displayed in this quality, this might result in very slow refresh rates. You can selectively choose those items in your EcoSystem to preview in higher quality by selecting the first icon to the right of the item which is the **Instances Preview Quality** (). Right-click on this icon to choose the display quality of the item. Click the second icon to turn on **Full Quality near Camera** ().
Keep in mind that these quality settings can be resources heavy. Constant calculations are made to determine which instances are in the quality set range and to actually create the quality display. This is not recommended for large EcoSystems and may slow down your system.
 - **Edit Alpha:** This button allows you to access the alpha channel of the EcoSystem layer, which can be used to drive the presence of instances over the material. If some alpha is indeed defined, this button is toggled. Clicking this button will open the *Function Graph* to edit the alpha output.

Distribution


The distribution parameters control the way the items on the EcoSystem population list are selected when populating the EcoSystem, and how they are placed relative to items in other EcoSystem layers.

By default, the distribution algorithm is set to **Random**, indicating that the items on the list are selected randomly.



If you click the **Drive with a function** icon () , you can control the way items on the list are selected using a function. See [here](#) for further details on driving material parameters with functions. This feature requires the *Advanced Graphics* module. When the distribution is controlled by a function, the Presence setting on the EcoSystem population list indicates an interval for each item. To decide which item should be placed where, the EcoSystem populator evaluates the result of the function at the current point and finds the interval that contains the function output. The corresponding item is created. Values outside the valid range of $[-1;1]$ are clamped.

Connecting the distribution parameter to a function is one of the rare cases that immediately affects the way the material behaves: by default, the function outputs 0, so only the item whose presence range contains 0 will appear in the population.



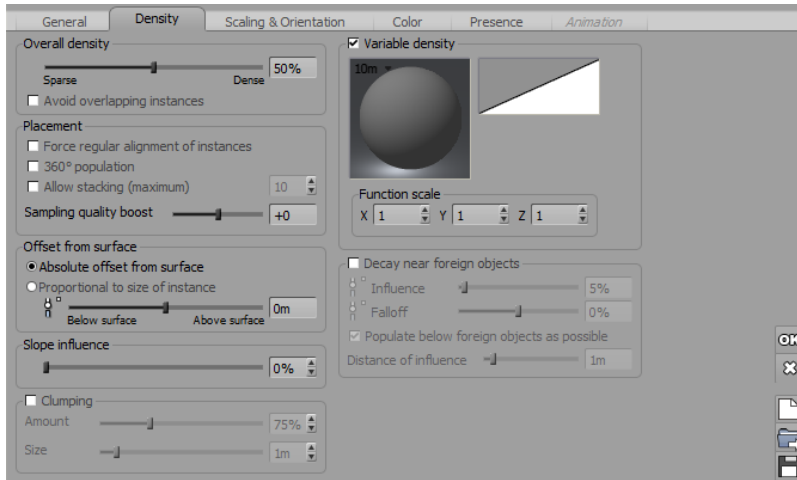
Effects of Affinity and Repulsion settings – left: positive affinity/positive repulsion – center: negative affinity/negative repulsion – right: positive affinity and repulsion/negative affinity and repulsion

- **Affinity with layer:** this setting controls how strongly instances from this layer are “attracted” to instances from other EcoSystems below it on the layer stack. For instance, if you have a layer of trees and want to have primroses around the trees, you’d add a layer of primroses on top of the tree layer, and enter a positive value for affinity. Higher values will make the primroses stick closer to the trees, and not appear anywhere else than near the trees. If you enter a negative value, the primroses will appear everywhere except near the trees. This setting is only available if there is another EcoSystem below this one on the layer stack.
- **Repulsion from layer:** this setting controls how close new instances can be added to instances from underlying EcoSystems. The effect of repulsion is a lot more “sudden” than that of affinity. For instance, if you have a layer of trees and want to have grass everywhere except around the trees, you’d add a layer of grass on top of the tree layer, and add some repulsion. Higher values will make the grass stay further away from the trees. If you enter a negative value for repulsion, you will have grass only near the trees. By using affinity and repulsion simultaneously, you could, for instance, have the grass appear near the trees, but not underneath them. This setting is only available if there is another EcoSystem below this one on the layer stack.





Density Tab

Overall Density



Density tab – EcoSystem Materials

Use the slider in the Overall Density frame to adjust the overall number of instances per unit of surface of the EcoSystem population throughout the entire EcoSystem. Higher values will mean more instances in the EcoSystem.

- **Emission flow:** (for EcoParticle systems) Click for the flow of emission for an EcoParticle system. Select the first icon () to enable a one time EcoParticle flow. Click the second icon () to control flow with a function.
- **Avoid overlapping instances:** select this option if you would like to avoid having instances of the EcoSystem population that overlap each other. Please note that enabling this option will limit the maximum density of the EcoSystem population, and does not mean that you will never have any overlapping instances. It does however minimize the number of occurrences of such overlaps. Reducing the overall density is another way to avoid overlapping instances.

Placement


- **Force regular alignment of instances:** select this option to remove the randomness in the placement of the EcoSystem population.



- **360° population:** this allows population of an EcoSystem to completely populate a material even though it covers sides and bottom of an object. For example, you can completely cover a sphere.
- **Allow stacking:** check this option to stack instances in an EcoSystem. If the density of EcoSystem population is high enough, Vue will stack instances instead of adding them at the same level. It works with layered EcoSystems too (i.e. instances of a top EcoSystem layer can now be added onto instances of the lower layers). It is compatible with dynamic population, with just a restriction of the maximum number of instances which can be stacked.
- **Max stacked instances:** this controls the number of instances you can stack in an EcoSystem.
- **Sampling quality:** this parameter controls the accuracy of the EcoSystem sampling process. If you need to closely follow a density profile (e.g. when using a bitmap to create an EcoSystem logo), you may need to increase this setting. Higher values will produce an EcoSystem population that follows the density profile more accurately, but population will be slower.

Offset from Surface

This frame lets you control how the EcoSystem population is placed relative to the surface of the underlying object.

Use the slider to adjust the offset between the base of the EcoSystem instances and the surface of the underlying object. A value of 0 (the default) will place the EcoSystem instances so that they rest on the surface of the underlying geometry. Positive values will make the EcoSystem instances “fly” above the surface, while negative values will sink the instances into the ground. This parameter can be driven by a function (thus letting you vary the altitude of the instances relative to the surface) by pressing the **Drive with a function** icons (). See [here](#) for further details on driving material parameters with functions.

The actual value of the offset from the surface depends on the options below:

- **Absolute offset from surface:** when this option is selected, the slider indicates the distance between the surface of the underlying object and the base of the instances in Vue units. All instances will be placed at the exact same distance from the surface (except if you drive the distance setting with a function).
- **Proportional to size of instance:** when this option is selected, the distance between the surface of the underlying object and the base of each instance is proportional to the size of the instance. A value of -50% will hence bury each instance halfway under the surface of the underlying object.



Slope Influence

- **Slope influence:** adjusts the influence that slope has on the distribution of instances. A value of 100% indicates that instances will automatically appear more sparsely on steeper slopes (like in Vue 6), whereas a value of 0% will indicate that the density of instances should be the same, whatever the slope.

Clumping

Check this option to group like species together, often as they appear in nature.

- **Amount:** use the slider or key in the amount of clumping in the EcoSystem (0% = no clumping, 100%=maximum intensity).
- **Size:** Use the slider or key in the size of the clumps.

Variable Density

Select this option to vary the density of the EcoSystem population over the surface of the underlying object.

Vue generates variable density using a function and a filter that indicates the local density depending on the value of the function. This is the way it works: for each point where density is evaluated, the function generates a number in the range of -1 to 1 (-1 is black on the preview of the function and 1 is white). This number is then transformed by the filter into a density. The maximum variable density is equal to the overall density. When you activate variable density, the default function returns a mid-range gray value, and the filter has no effect. This results in a final density that is exactly half of what it was before enabling the option.

To change the function, double-click on the picture of the function. This will open the Function Visual Browser. Use the scaling controls to scale the function along the X, Y and Z axes.

You can adjust the density according to the output of the function using the filter.


Decay near Foreign Objects

This is a very interesting feature that automatically adjusts the density of the EcoSystem population around objects that are placed on the underlying EcoSystem object. For instance, if you place a large rock in the middle of a terrain covered with vegetation, the density of the vegetation will automatically be reduced around the rock. Any object that is made of another material than the EcoSystem material will be considered a “foreign object”, unless the **Ignore object(s) when populating EcoSystems** option is selected for that object.

Because objects that are hidden from render still have an influence on the EcoSystem




density, you can use such objects to locally modify the EcoSystem.



Population below foreign object

- **Influence:** this parameter controls the influence of the foreign objects on the EcoSystem density. The higher the value, the larger the void around the foreign objects.
- **Falloff:** this parameter controls the profile of the decay around the foreign object. A value of 0 will create a linear decay profile allowing a smooth transition in the EcoSystem density around the foreign objects, while positive values will seemingly increase the void around the foreign objects while making that void much more sudden.

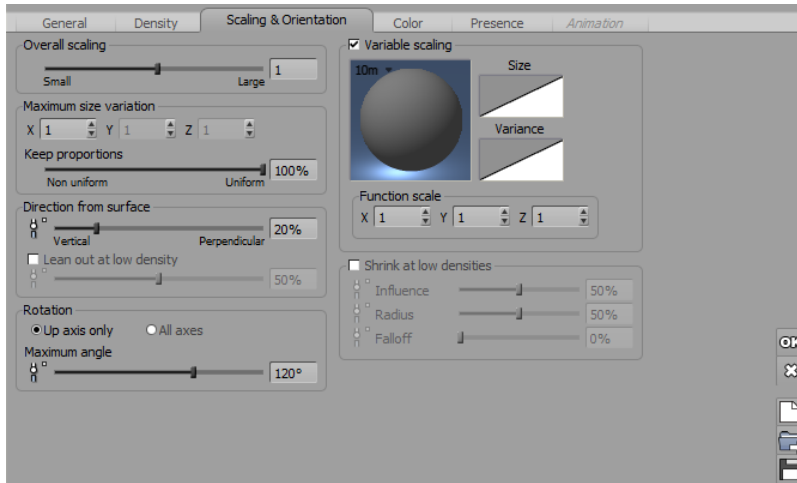
Both these parameters can be driven by a function by pressing the corresponding **Drive with a function** icons (). See [here](#) for further details on driving material parameters with functions.

- **Populate below foreign objects as possible:** this option works in conjunction with the “Decay near foreign objects” option (see below). Foreign objects create a void around them in the EcoSystem population. If this option is selected, and if the foreign object is above the EcoSystem surface, Vue will attempt to place instances of the EcoSystem population under the foreign object (see illustration).



Scaling & Orientation Tab

Overall Scaling



Scaling & Orientation tab – EcoSystem Materials

This setting lets you control the overall size of the instances in the EcoSystem population. A value of 1 leaves the size of the instances untouched, while a value of 2 will double the size of the instances of the EcoSystem population.

This setting works in combination with the **Scale** factor on the list of items in the [EcoSystem Population](#).

Maximum Size Variation

The controls in this frame let you indicate how the size of the instances in the EcoSystem population varies along each axis. The mathematics behind this variation are a little complex, but suffice to say that a value of 1 will create instances that are between one half and twice the size of the original item, that a value of 0 means no variation in size, and that the greater the value, the stronger the variation in size throughout the EcoSystem population. These settings along the 3 axes are also influenced by the “Keep proportions” setting detailed below. If the “Keep proportions” setting is at its maximum value of 100%, the size variation only needs a single value, the value for the other axes is simply ignored (hence only the X value can be set, the Y and Z controls are greyed-out).

- **Keep proportions:** this setting controls how stretched or squashed the instances




of the EcoSystem population will be. A value of 100% indicates that the proportions of the EcoSystem items are untouched (which doesn't mean the objects won't be resized), whereas a value of 0% indicates that the scaling along the 3 axes is not correlated, resulting in objects that can be strongly stretched or squashed.

For items in this next section, the Drive With a Function requires the *Advanced Graphics* module.

Direction from Surface


This slider lets you indicate how the EcoSystem instances “grow” from the surface. If the slider is set to 0%, the instances will always grow vertically, whatever the slope of the underlying “terrain”. A value of 100% means that the instances will always grow from the surface (perpendicular to that surface).

This parameter can be driven by a function by pressing the **Drive with a function** icon (). See [here](#) for further details on driving material parameters with functions.

Rotation

This frame lets you define the random rotation that is applied to the instances in the EcoSystem population.

- **Maximum angle:** this setting controls the maximum angle of the random rotation applied to the instances in the EcoSystem population. By limiting the angle of rotation, you can preserve the directionality of certain animation effects (for instance, if you wanted to apply wind effects to an entire forest, you could limit this angle of rotation so that all trees face in roughly the same direction). On the other hand, a larger angle of rotation means that the items in the EcoSystem population will be viewed under a greater variety of angles, resulting in a seemingly more diverse population.

If you click the **Drive with a function** icon () for the “Maximum angle” setting, you can control the exact angle of rotation that is applied to each instance in the EcoSystem population. When this parameter is connected, there is no more randomness in the angle of rotation. Expected values are in between -1 (-180°) and +1 (+180°).

Connecting the “Maximum angle” parameter to a function is one of the rare cases that immediately affects the way the material behaves: by default, the function outputs 0, so the entire population will be oriented the same.

- **Up axis only:** select this option if you want the rotation to only take place along the Up axis (vertical). This is typically the case for objects that “grow” from the surface, such as trees.
- **All axes:** if this option is selected, a random rotation will be applied to all axes of the instances. This is best used for objects that do not grow from the surface (e.g.



rocks).

Variable Scaling

Select this option to control the scaling of the EcoSystem instances using a function.

Vue generates variable scaling using a function and two filters that indicate the local scale and size variations depending on the value of the function. This is the way it works: for each point where scaling is evaluated, the function generates a number in the range of -1 to 1 (-1 is black on the preview of the function and 1 is white). This number is then transformed by the filters into a scaling and size variation. The maximum scaling ratio is equal to the overall scaling setting, and the maximum variation is equal to the Maximum size variation setting. The actual size of the instance along the 3 axes is determined randomly based on the size variation value and the “Keep proportions” setting.

When you activate variable scaling, the default function returns a mid-range gray value, and the filters have no effect. This results in a final scaling and size variation that is exactly half of what it was before enabling the option.

To change the function, double-click on the picture of the function. This will open the Function Graph. Use the scaling controls to scale the function along the X, Y and Z axes.

You can adjust the scaling and size variations according to the output of the function using the corresponding filters.


Shrink at Low Densities

When this option is selected, the size of the instances in the EcoSystem population will automatically be reduced when the density of the EcoSystem population is low. This could, for instance, be used in combination with the **Decay near foreign objects** option in the **Density** tab to automatically reduce the number and size of instances around a foreign object. You can also use this in combination with the **Decay color at low densities** option in the **Color** tab to simulate adverse growing conditions.

- **Influence:** this setting controls the amount of reduction in size of the instances as the density becomes low. If you enter a negative value, the size of the instances will increase at low densities.
- **Radius:** this setting controls the density value under which the density is considered to be low and hence have effects on the size of the instances. If the density function is gradual, this translates visually to a strip of influence on the edges of the EcoSystem population (although the relationship between radius and actual width of the strip is not straightforward).
- **Falloff:** this parameter controls the profile of the size reduction as the density becomes low. A value of 0 will create a linear size reduction profile, allowing for a smooth transition in the size of the EcoSystem instances around low density areas,




while positive values will seemingly increase the strip around the low density areas by making the size reduction much more sudden.

These 3 parameters can be driven by a function by pressing the corresponding **Drive with a function** icons (). See [here](#) for further details on driving material parameters with functions.


Lean Out at Low Density

When this option is selected, EcoSystem instances lean out in the zones where density is lower, simulating plants trying to get as much light as possible.

- **Influence:** Use the slider to set the intensity of the effect or click the **Drive with a Function** icon () to use the *Function Graph* to connect this to any parameter in the graph.

Color Tab

Color Correction



Color tab – EcoSystem Materials

- **Overall color:** this setting controls the overall color of the instances in the EcoSystem population. It represents the average color of all the items in the EcoSystem population. Because the color of these items can be very different, the overall color is often quite dull.




By modifying this overall color, you will modify the colors of all the instances in the EcoSystem population. For instance, if you make the overall color brighter, all the instances will become brighter.

Color at Low Densities

The “Color at low density” option works in a similar way to the “Shrink at low densities” option in the **Scaling & Orientation** tab. What it does is automatically alter the colors of the EcoSystem population when the density of the EcoSystem becomes low. Using this option, you could easily make the plants in your EcoSystem look pale and yellow where the density is low, thus simulating adverse growing conditions. This works particularly well in combination with the **Shrink at low densities** option described above.

- **Decay color:** this is the overall color of the EcoSystem population at low densities. It is not the actual color of the instances, because each instance can have differing colors, but think of it as the average color, in the same sense as the **Overall color** setting described above.
- **Influence:** this setting controls how strongly the colors of the instances are affected by the decay color.
- **Radius:** this setting controls the density value under which the density is considered to be low and hence have effects on the color of the instances. If the density function is gradual, this translates visually to a strip of influence on the edges of the EcoSystem population (although the relationship between radius and actual width of the strip is not straightforward).
- **Falloff:** this parameter controls the profile of the color changes as the density becomes low. A value of 0 will create a linear color change, allowing for a smooth transition in the EcoSystem colors around low density areas, while positive values will seemingly increase the strip around the low density areas by making the color change much more sudden.

These 4 parameters can be driven by a function by pressing the corresponding **Drive with a function** icons (). See [here](http://www.e-onsoftware.com/wiki/Vue/index.php/Documentation/EcoSystems/EcoSystems_in_the_Material_Editor/Color_Tab) for further details on driving material parameters with functions. http://www.e-onsoftware.com/wiki/Vue/index.php/Documentation/EcoSystems/EcoSystems_in_the_Material_Editor/Color_Tab

To use the Drive With a Function, the *Advanced Graphics* module is required.

Variable Color

These functions require the *Advanced Graphics* module.

Select this option to control the color of the EcoSystem instances using a function.

Vue generates variable colors using a function, a filter and a color map that indicate the local “average” color depending on the value of the function and the filter. This is the



way it works: for each point where color is evaluated, the function generates a number in the range of -1 to 1 (-1 is black on the preview of the function and 1 is white). This number is then transformed by the filter and the color map into a color. This color is not the actual color of the EcoSystem instance at that point, but rather the average color that **many** instances would have if they were all placed at that location.

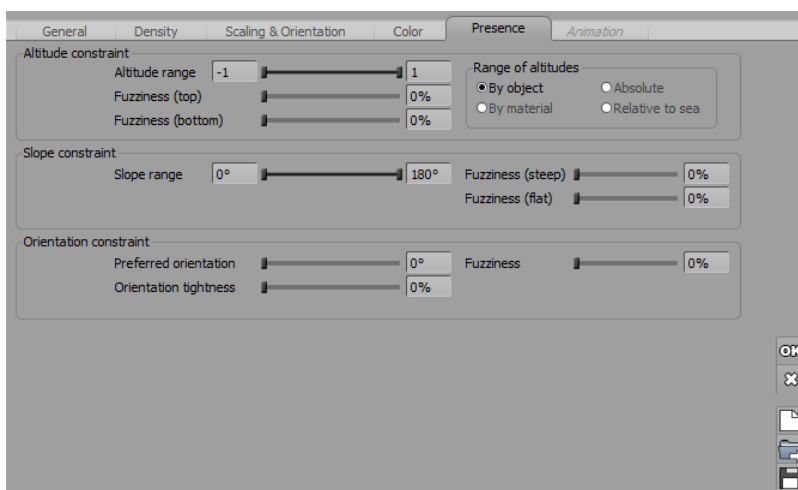
When you activate variable colors, the default function, filter and color map will yield the same color as the current overall color. Any changes that you make to the “Overall color” will be reflected in the color map. Likewise, any changes that you make to the color map will be reflected in the Overall color.

To change the function, double-click on the picture of the function. This will open the Function Visual Browser. Use the scaling controls to scale the function along the X, Y and Z axes.

Generally speaking, it is unwise to load a color map into the variable color slot, as this will result in dramatic (and usually unexpected) changes in the colors of the EcoSystem population. You will be better off editing the color map by hand to introduce very slight changes.

An intensity slider allows you to adjust the sensitivity to color variations in your EcoSystem population. Like other parameters, it can be driven by a function.

Presence Tab



Presence tab – EcoSystem Materials

This tab lets you control how the environment affects the presence of the current layer. The behavior of this tab is identical to that of the [Environment](#) tab of **Simple materials**.

Altitude Constraint

This group lets you control how altitude influences the presence of the layer:

- **Altitude range:** this dual slider lets you define the range of altitudes in which the current layer appears (provided it is not transparent at this point). The layer will not appear outside this range.
- **Fuzziness:** this setting controls how “suddenly” the changes to the layer presence are made in response to altitude. High values mean that the layer appears very gradually in its altitude range, whereas low values will result in the layer appearing as a solid strip.
- **Range of altitudes:** this lets you define in what coordinates the altitude range is defined:
- **By object:** in this mode, the range is relative to each object to which the material is applied.
- **By material:** in this mode, the range is relative to all the objects that use this material.
- **Absolute:** in this mode, the range of altitudes is expressed in global coordinates.
- **Relative to sea:** the altitude is computed from the sea level and not from zero.

Slope Constraint

This group lets you control how the local slope influences the presence of the layer:

- **Slope range:** this dual slider lets you define the range of slopes in which the current layer appears (provided it is not transparent at this point). The layer will not appear outside this range.
- **Fuzziness:** this setting controls how “suddenly” the changes to the layer presence are made in response to slope. High values mean that the layer appears very gradually in its slope range, whereas low values will result in the layer appearing as a solid strip on areas of appropriate slope.

Influence of Orientation

This group lets you control how the local orientation influences the presence of the layer:

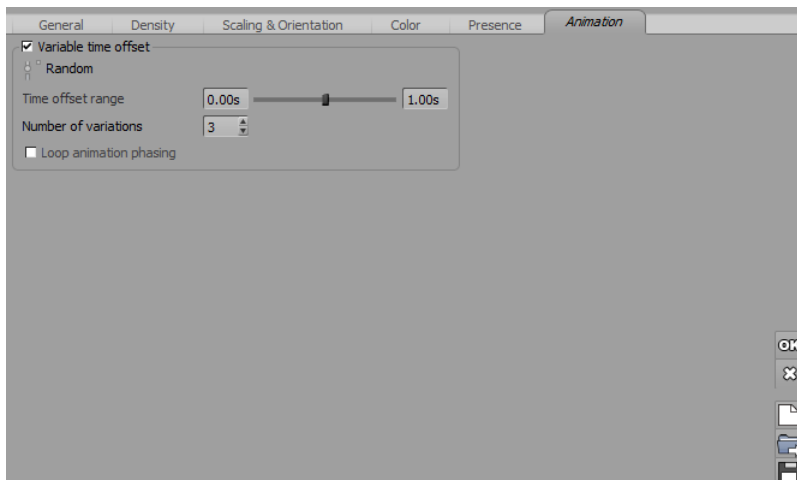
- **Preferred orientation:** this setting controls the orientation of the surface that is the most favorable to the presence of the layer.
- **Orientation influence:** this setting controls the influence of orientation on the



presence of the layer.

- **Fuzziness:** this setting controls how “suddenly” the changes to the layer presence are made in response to orientation. High values mean that the layer appears very gradually on surfaces of the preferred orientation, whereas low values will result in the layer appearing as a solid strip on areas of preferred orientation.

Animation Tab



Animation tab – EcoSystem Materials

This tab has setting to allow you to control the phasing of animated EcoSystem instances. The tab appears when at least one of the EcoSystem specimens contains animation. You can then animate the EcoSystem instances at different phases of the animation. This is useful for creating realistic animated populations, for example, a crowd rather than a marching army.

- **Variable time offset:** check this to enable the phasing features.
- **Random:** for better control of the phasing, you can use a function. Click to open the *Function Graph* and set further parameters for **Time Offset**.
- **Time Offset Range:** Adjust the slider to set the range of allowed phase shift. For instance, a range of [0; 1] means that each instance created will have a random phase shift of 0 to 1 second.




- **Loop animation phasing:** check to loop the animation phasing.



Converting Objects To EcoSystem Instances

You can convert all references matching a given master object to EcoSystem instances. These instances are added to an existing Global EcoSystem or you can create a new EcoSystem.

There are two ways to access the feature:



Select Master Object

From World Browser: select a Master Object (in Library Tab), right-click to open the menu, and choose **Convert to EcoSystem instances**

By drag and drop: drag one reference object from the **Objects Tab** of the **World Browser** to the **EcoSystem Painter** Population section.

This only works with Global EcoSystems.


If you are creating a new EcoSystem, once you have painted with it, you will see an entry for the EcoSystem in the **Objects Tab** of the **World Browser**.



EcoSystem Export

If the object you have selected for export has EcoSystem instances, the EcoSystem will be exported with the object, provided you export it in Alembic or FBX format. Materials are included with the FBX export.

xStream



Export Options

By default, EcoSystem instances are exported with global export options, but you can configure how EcoSystem specimens are exported by right-clicking on them in the World Browser / Classes tab.

For EcoSystem instances, export can only be previewed at render time (in main render or render scene preview), and not in real-time views.





Section 6

Tutorials





Introduction


If you only read one section of this manual, then let it be this one. Practicing (and understanding) these tutorials will give you in no time at all a good knowledge of how the software operates, and the incredible variety of results that can be achieved using it.

This section is divided into 4 parts:

- **Quick reference how-to's:** the first part of this section gives a list of quick and specific “how to” tutorials. It isn't necessary that you master all these tutorials to be able to use the program, because they are mostly designed as answers to specific questions. Consider them more like a reference.
- **Elaborate feature tutorials:** more complete tutorials, detailing ways of taking VUE to the limits. Just like the quick reference tutorials, you don't need to master all these tutorials to know how to use VUE.
- **Building a complete scene:** a detailed construction of a complete sample scene.
- **Animation tutorials:** a set of tutorials dealing specifically with the animation capabilities of VUE.

Quick Reference How-To's

Rendering and Saving a Picture

1. Create a new file by selecting the menu command **File | New**. Select an atmosphere from the Visual Atmosphere Browser, and click **OK**.
2. Create a Terrain by clicking the **Heightfield Terrain** icon (.
3. Select the menu command **Render | Render** and watch the picture as it is rendered.
4. You can stop the rendering process by pressing **Escape**.
5. To save the picture when it is rendered, click on the **Save Displayed Picture** icon on the *Render Display* (farthest right icon under the rendered image).




Rendering a Full Screen Picture

1. Select the menu command **Render | Render options**.
2. In the *Render Options* dialog, select **Render to screen**.
3. Select the requested picture resolution. Select **Full screen** for the largest possible picture on your machine.
4. Press **OK**. Watch the picture as it is rendered. Press **Escape** to stop rendering.
5. To save the picture when it is rendered, click on the **Save Displayed Picture** icon on the *Render Display* (farthest right icon under the rendered image).

Rendering a High Resolution Picture

1. Select the menu command **Render | Render Options**.
2. In the *Render Options* dialog, select **Render off screen**.
3. Select the required picture resolution (e.g. 3000 × 2000).
4. Click on the **Options** button to select a target location for the rendered file, a name for the file and a file format.
5. Press **OK**. You won't see the picture as it renders, but you'll see an indication on the progress in the *Status Bar*. Press **Escape** to stop rendering. The picture is saved in its final form when the render is complete.

Creating a Rock Material

1. Open the *Advanced Material Editor* by double clicking on the preview of the material that you want to modify (or select **Edit Material** from the popup menu of the material preview).
2. Press the **New** () button to reset the material.
3. Go to the **Colors** tab. Select a **Procedural colors** coloring type.
4. Replace the color map by right-clicking on the color map and select **Load color map**. Select the map named *Stone gray* from the *Rocks and Grass* collection.
5. To apply a texture, right click on the **Color Production** sphere and select **Load Function**.
6. When the Function Graph displays, select the *00_GrainyFractal* from the *Basic*





collection.

- Now, to add bump to the texture, select the **Bumps** tab. Right click on the **Bump production** sphere and select **Edit Function**.

- When the *Function Graph* opens, click on the **Bump** output and connect it to the **Grainy fractal** node. Select **Rough areas**. Click Ok to close the *Function Graph*.

You can change the effect of the rock by loading different functions in the bump editor. Experiment with the different functions.

Mapping a Material Using a Picture

- Open the *Basic Material Editor* by double clicking on the preview of the material that you want to modify (or select **Edit Material** from the popup menu of the material preview).
- Press **New** () to reset the material.
- Check the **Color map** option and load the picture of your choice by clicking the **Load** button ()
- Press the **Load** button, and indicate the name of the picture file you would like to use, then click **OK**.
- Select an **Automatic** mapping mode, and **None** for over sampling.
- To use the same picture for the bump map, check the **Bump map** option and check **Same as color map**. To use a different map, uncheck this last option and load the map of your choice.

Creating Fog at Low Altitudes


- Starting with the *Default* spectral atmosphere with no clouds, select the menu command **Atmosphere | Atmosphere Editor** to display the *Atmosphere Editor*.
- Go to the **Cloud** tab and select the **Add** button to add a cloud layer.
- In the Material Browser, select *Clouds/Spectral 2* collection and select the **Dense Cumulus** layer.
- Back on the **Cloud** tab, key in “0” for the **Altitude** value.
- Key in 4 meters for the **Height** value.
- The **Cover** slider is set to 97%. You can reset this to get the effect you want, but for now, you want to be sure to see the cloud layer.



7. Move the **Density** slider to at least 80%.
8. Move the **Opacity** slider to 75%.



Now, you will have a dense cloud layer hovering low to the ground. Use the settings mentioned above to raise the fog, or to change its density. You can also select the cloud layer in the *World Browser* and move it to a different height.

Making a Custom Canyon Terrain

1. Create a new standard heightfield terrain by clicking on the first terrain icon on the left of the user interface () or selecting **Object | Create | Heightfield Terrain | Standard Heightfield Terrain** from the menu.
2. Double-click on the terrain in the *3D views*, or in the *World Browser*. This opens the *Terrain Editor*.
3. Press the **Canyon** predefined terrain style. The map of the terrain changes to display the canyon.
4. Go to the **Effects** tab and click on **Fluvial** to add some erosion effect. Click **OK** to exit the editor.

You could also take a look at the tutorial on [Modifying a Terrain](#) for an example of another (interactive) way to create canyon terrains.

Creating Terrains with Snow Covered Tops

1. Create a scene with several terrains. Select all the terrains and the ground together (by Shift clicking on them in the *3D views*), and **Load** a rock material for these terrains (*Winnipeg* in the *Landscapes* collection for instance). Press .
2. Now we want to add snow to these terrains. Open the *Material Editor* for the *Winnipeg* material and click on **Mixed materials** at the top of the dialog. Notice how a box for an additional material is added.
3. Right-click on this material, select **Load material** , from the menu, and select *Snow* from the *Landscapes* collection.
4. Indicate a **Smooth blending strip** width of 20%.
5. The trouble is that the *Winnipeg* and *Snow* materials are equally mixed everywhere. We want snow at high altitudes only, and accumulated on flat surfaces...
6. Switch to the **Influence of environment** tab. Check the **Distribution of materials dependent on altitude, slope and orientation** box to activate environment



sensitivity.



7. Indicate that the altitude has an medium influence, by pushing **Influence of altitude** up to 50%.
8. Indicate that material 2 (the snow) gathers at **High altitudes**.
9. Indicate that slope also has an medium influence on the way snow accumulates by pushing **Influence of slope** up to 50% also.
10. Indicate that the material 2 (the snow) gathers on **Flat surfaces**.
11. The material is nearly ready. You could make the snow more dense on the sides of the mountains that are less exposed to sunshine. To do this, you'd have to indicate some **Influence of orientation**, and adjust the orientation using the **Azimuth** control.
12. To adjust the amount of snow that is visible in your landscape, push the **Mixing proportions** slider to the right to increase the amount of snow, and to the left to reduce it.
13. You can improve the material by making transitions from **Winnipeg** to **Snow** more realistic. To do this, you need to add noise into the mixing process. Go back to the **Materials to mix** tab, and edit the **Distribution of materials function**. Create a **Noise (smooth)** layer, and make it **Fractal** with a complexity of 2. You can adapt the influence of this noise by varying the extension of the function.
14. You could give the impression that snow is resting on the 'Winnipeg *by using a more elaborate blending algorithm: select the **Full blend (cubic bumps)** mixing method. The difference is subtle, but you'll notice that **Snow** now looks like it has a thickness.*

Adding a Sense of Depth to Scenes



1. To give the impression that your scene runs far away in the distance, you have to give it a sense of depth. To do that, create several **Terrains**, and move them progressively away from the camera. Keep in mind that, in order to see far away terrains, they must be very large. So you should enlarge them as they get further.
2. The use of an atmosphere that has strong fog or haze will help give the impression of depth.



Mixing Materials with Tilted Stratum

1. Open the *Material Editor* by double clicking on the preview of the material that you want to modify (or select **Edit Material** from the popup menu of the material preview).
2. Reset the material by pressing **New** () . Select the **Mixed materials** option.
3. Load rock materials of your choice into Material 1 and 2 using the  buttons.
4. Indicate a **Smooth blending strip** of 5%.
5. Open the *Function Scale* for the **Material distribution function**.
6. Create a **Perlin | Value** noise node.
7. Change the **Wavelength** along **X** and **Y** up to 5. Reduce **Z** scale to 0.5.
8. Select the **Position** input and click the **Math node** icon. This inserts a **Vector operation** node. Change the node type to **Rotation and Twist**.
9. Indicate a rotation around the **X** axis equal to 20° (in order to tilt the stratum).

Creating Undulated Furrows

1. Open the *Material Editor* by double clicking on the preview of the material that you want to modify (or select **Edit Material** from the popup menu of the material preview).
2. Reset the material by pressing **New** () . Select the **Simple material** option and go to the **Colors** tab. Select the **Procedural colors** option to create a gray material.
3. Go to the **Bumps** tab and open the *Function Graph* for the **Bump production function**.
4. Create a noise node of type **Math Pattern | Wave**.
5. Adjust the **Wavelength** along **Y** to 0.7. Zero the scaling along **X** and **Z** (to remove any variations). This will create the furrows along Y; to orientate the furrows differently, use the **Rotation** option.
6. The furrows are currently parallel. We need to make them undulate by adding turbulence. **Extract** the **Origin** parameter by clicking the extract icon () and click the **Turbulence node** icon to change the extracted origin into a turbulence node..




7. Indicate a turbulence **Repeat count** of 3, based on a **Perlin | Gradient** noise, with an **Amplitude** of 0.13 and **Roughness** of 0.6. Close the *Function Graph* by pressing **OK**.
8. Indicate a **Bump gain** equal to 0.5.

Adding Color Variations to Materials

There are three ways of adding color variations to materials:


- Mix two materials together that look similar (you can use the **Color and lighting only** blend mode).
- Introduce large scale color variations in a simple material with Procedural colors by using **Combination nodes** in the color production function.
- Create color variations directly by using different colors depending on Altitude/Slope and Orientation.

The following illustrates this last method by automatically adding color variations to a uniform material:

1. Open the *Material Editor* by double clicking on the preview of the material that you want to modify (or select **Edit Material** from the popup menu of the material preview).
2. Reset the material by pressing **New** (). Select the **Simple material** option and go to the **Colors** tab. Select the **Procedural colors** option to create a gray material.
3. Double-click on the **Color map** and load a colorful map (e.g. *Rainbow* from the *Colorful* collection). The resulting material should be uniformly cyan (bright blue).
4. Right click on the **Color map** and select **Edit function** to open the *Function Graph*. Select the **Filter** box and right-click on the filter chart below to open the *Filter* dialog.
5. Check the **Dependent of slope** option. Look at the material preview. It now displays the full range of colors in the color map. Colors to the left of the map are found on flat surfaces, and colors to the right on steep surfaces.
6. Check out the other types of dependencies. You can also combine them together. This is great for creating grass that automatically become yellow and dry at higher altitudes.



Modifying a Terrain

1. Create a terrain by selecting the terrain icon (). Double-click on the terrain in a view port to open the *Terrain Editor*.
2. On the **Paint** tab, select **Airbrush** for a brush tip and slide the **Flow** setting to the right. Select **Raise** in the **2D** section.
3. Drag the cursor onto the 3D view of the terrain and start painting by pressing the mouse button. Terrain altitudes increase gradually under the cursor. You can vary the size of the cursor using the brush tip **Size** setting.
4. Check the **Invert** option to dig into the terrain. Start carving a river out of the terrain.
5. Press the **Canyon** predefined terrain style.
6. On the **Effects** tab, select **Alluvium**. Watch the terrain as streams are dug out of the surface. If you increase the **Rock hardness**, the shape of the streams will be modified accordingly.
7. Play around with the other tools and start terraforming!

Modifying a Terrain Externally

1. Create a terrain by selecting the menu command **Object | Create | Terrain in Editor**. Indicate the size of the terrain you want to create and press **OK**. This brings you to the *Terrain Editor*.
2. Press the **Copy** button. In the 2D application of your choice (e.g. PhotoShop™), create a picture of the same resolution as the terrain, and **Paste** into it.
3. Now you can modify the terrain as you like. You can use filters on it to generate astounding effects. Remember that when you lighten a point of the picture it means that you are increasing the altitude of the corresponding point in the terrain.
4. When you have finished modifying the picture, **Copy** it back into the clipboard. In the *Terrain Editor*, press **Paste**. The modified terrain appears in the map.
5. When you are done, press **OK** to validate the new terrain.

The **Canyon** terrain object was created in the following way: a meandering black river was drawn on a white background; then, a gradation was introduced around the river; the picture was then filtered to create steps and imported into VUE; last, the terrain was slightly eroded (**Diffusive erosion**) to smooth the steps and dig water streams (**Fluvial erosion**).



Importing Multi-Part DEMs

Vue can import USGS DEM models that are composed of several files. However, there are a few precautions you should take when importing such multi-part DEMs.


1. In the **Import Options** dialog, you should uncheck the options called **Resize object** and **Center object**. If these options are checked, the models that you import will be automatically positioned at the center of the 3D views, and resized to a standard size, which means that relative positions and sizes will be lost.
2. Import the first DEM model by using the menu command **File | Import Object**. The DEM file contains information on the geographic position of the DEM slab. Unfortunately, because you can get DEMs from anywhere in the world, chances are that this position will be out of the Vue bounds. If this happens, the *Terrain Offset* dialog will appear, suggesting that you offset the DEM's origin by a value that will make it appear at the center of the *3D views*. Press **OK**.
3. Now import the other DEM models in the set. Each time the *Terrain Offset* dialog appears, be sure to press **OK** without modifying the values (so as to maintain relative positions).
4. When you have finished importing the set of models and you want to import another set, you can press the **Reset** button in the *Terrain Offset* dialog. This will compute new values for the offset, so that the new DEM model is centered in the *3D views*.




Elaborate Feature Tutorials

Boolean Objects


Building a Simple House



Simple house: step 3



Simple house: step 5



Simple house: step 6



Simple house: step 8

This tutorial will illustrate the power of Boolean objects by guiding you through the construction of a simple house, entirely made out of basic VUE primitives. The result of this tutorial can be seen in the Simple house object.

The house body:

1. Create a cube. Stretch it a bit horizontally in the *Top view* using the **Resize in this direction** handle. This is the main body of the house. Rename the cube as “House body”.
2. We will now add windows to it: create another cube, and resize it using the **Resize globally** handle so that it is approximately one quarter of the size of the House body. In the *Top view*, move the cube to one third of the length of the House body, and stretch it vertically using the **Resize in this direction** handle, so that it sticks out on both sides of the House body. Rename the cube as “Window”.
3. Now select the House body, and then Shift select the Window. Both objects should be selected (i.e. displayed in red, and highlighted in the *World Browser*). Click on Boolean difference. Make sure you got the order right by unfolding the newly created Boolean difference in the *World Browser*, and checking that the first member is the House body, and the second the Window. If not, change the order by dragging and dropping the House body to the first position. Rename the Boolean difference as “House wall”.
4. One window is obviously not enough, so we will add another one: in the *World Browser*, select the Window and drag it onto the House body, then, before dropping the Window, press Control. Then drop the Window. It is now duplicated (another way of doing this would be to **Copy-Paste** the Window, or to **Duplicate** it). Now, using the *Top view*, move the Window to the other side of the House body. If you test render, you will see that our house now has two sets of windows.
5. Now we will give a thickness to the House wall: in the *World Browser*, select the House body, and duplicate it, as above. Rename the new copy to “Wall thickness”,



and shrink it slightly using the **Resize globally** handle. Since this new object is also taken out of the House body, we now get walls that have a small thickness.

6. How about a door? We will make one with a nice rounded top. Create a cube and a cylinder, rotate the cylinder 90° so it runs horizontally in the *Top view*, then, using the arrows, nudge it up in the *Side view* until it's half way up the cube. Now group the objects, rename the group as “Door”. Stretch the Door vertically using the **Resize in this direction** handle until it has the height over width ratio of a plausible door, then position it on the side wall, letting it overlap on both sides. Finally, using the *World Browser*, drag the Door into the House wall, onto the last member Window. The Door will now be dug out of the side wall.
7. To build a roof for the house, we will use a Boolean intersection. Select the House body using the *World Browser*, and duplicate it as above. Rename it as “Roof – part 1”. Duplicate this, and rename it as “Roof – part 2”. Rotate it 45° in the *Side view* and lower it by half its height. Select both parts of the roof, and make a Boolean intersection; rename this as “House roof”. Move the roof up and resize it until it fits snugly on top of the body, jutting out slightly on both sides.
8. Group the House wall with the House roof, and rename this group as “Simple house”. The house is finished! You could add a **Point light** inside, and see it shine through the windows... Now think of what you can do with terrains!


Clipping Terrains

Vue offers the possibility to clip terrains under a given altitude. When you render such a terrain, all the parts that are under the clipping altitude appear as holes in the terrain. You can use this feature to make objects out of terrains that don't have square edges.

Making a Stone Arch

To illustrate what can be achieved using clipped terrains, we will examine a method for making stone arches.





Stone arch: step 7

1. First, create a terrain by selecting **Object | Create | Heightfield Terrain | Symmetrical Heightfield Terrain** from the menu. Rename it as “Stone arch”.
2. By default, VUE creates a fractal terrain. Since this doesn’t look at all like an arch, we will start off from scratch. Click **Reset**.
3. Select the **Paint** tab, select an **Airbrush** brush and its **Size** so that it is a little less than a quarter the width of the terrain map. Select **Raise** brush mode, set the **Flow** to halfway up, and **Softness** to one third up. Now draw a large inverted U (the arch).
4. Enter a **Low clip** altitude of 1.00, to clip out all parts of the terrain that are not on the inverted U you have just drawn. Everything should turn transparent but the arch. Transparent won’t be visible when the terrain is rendered; they are clipped out.
5. Close the editor. You can see the shape of the terrain in the *3D Views*. Tilt it up 90°, so that the arch is vertical.
6. If necessary, resize the terrain to make it more stocky.
7. Rendering will display a good basis for a stone arch. Now paint-in details like a wider base, or, using the same technique, create a big rock standing on two smaller ones; you could also create a layered-like stone by adding a tad of **Terraces** effect. You can also flatten out the top the terrain using the **High clip** control.
8. If you want to reduce the sudden change in slope where the terrain mirrors, filter the terrain altitudes using a **Round Mountain** -type filter.



Fuzzy Materials

Fuzzy material is an incredibly useful feature that can be turned on by selecting a box in the **Transparency** tab of the *Material Editor*.

A fuzzy object is one that has progressive (blurred) edges, instead of clearly defined ones. As a result, the frontiers of the object are unclear. This is particularly useful when trying to capture atmospheric effects. When used in conjunction with variable transparencies, objects of incredible visual complexity can be achieved at little expense.

Fuzzy objects are an interesting alternative to volumetric materials, because they render much more rapidly.

However, fuzzy objects are unfortunately rather difficult to use, and achieving good results may take quite a bit of experimentation. The first thing you have to know, is that the “fuzziness” of an object is influenced by the objects shape (but not global size), hence flattening an object will affect how fuzzily the object will be rendered. Consequently, it is important to get the shape of the object right before adjusting fuzziness.

Modeling Clouds



Clouds: step 4





Clouds: step 5



Clouds: step 7




Complex clouds made up of 4 spheres

Clouds are the perfect example of atmospheric effects that are incredibly complex to capture. However, using fuzzy materials, very pleasing results can be achieved at little expense. A method for achieving such results is now exposed (the final cloud, made up of 4 spheres, can be found in the sample objects, under the name **Cloud**):

1. Although any Primitive could be used as a basis for the cloud, it appears that spheres often yield the best results. Create a sphere, and flatten it using the **Resize in this direction** handle. Since clouds are enormous structures, enlarge it until it fills the entire view using the **Resize globally** handle, then move it up into the sky. Rename it to “Cloud”. Be sure you are satisfied by the shape of your cloud before proceeding.



2. Open the *Material Editor* for the Cloud and press the **New** icon () to start designing the Cloud material from scratch. Rename the material to “Cloud material”. Select the **Transparency** tab of the editor.
3. Select the **Fuzzy** box. If you take a close look at the edges of the preview, you will notice that it now has “blurred” edges. Move the **Fuzziness** slider up to watch this effect increase. Take it up to 70% and select **Don’t cast shadows** (unless you want shadows).
4. Go to the **Colors** tab, and select the **Procedural colors** box. Open the *Color Map Editor*, and create a uniformly white map.
5. For the time being, the result is a very dark and homogeneous cloud. Let us concentrate first on making it brighter. Light interacts differently with clouds as with “solid” objects: it travels inside the object, which means the side of the cloud opposite the sun still receives some direct light from the sun. As a consequence, the proportion of ambient light (which is constant all over the object) should be pushed up, and diffused light should be turned down. Go to the **Effects** tab, and turn **Diffuse** lighting down to 30%, and push **Ambient** lighting up to 100%.

Note:

Diffuse + Ambient should always be equal to 100%, unless you want to modify the overall luminosity of the object, which is what is done here.

6. Now we will concentrate on making the cloud less homogeneous. This usually takes quite a bit of experimentation before things look right, but here is a method that works well: raise **Global transparency** to 100% and click **Variable transparency**. Open the Transparency production *Function Graph*, and create a **Perlin | Value** noise node. Add complexity to this function by clicking the **Fractal** icon.
7. It may be necessary at this time to decrease **Fuzziness** to compensate for the transparency we just introduced. Reduce it to 50%. The result is a pleasingly complex cloud. It can be improved by adding bumps to the surface of the cloud, and adding yet larger scale variations to the cloud density.
8. Better results would be achieved by using multiple copies of this cloud, at different sizes, to capture the overall shape of the desired cloud (see sample object **Cloud**). Remember that varying the proportions of the clouds will affect their overall density. Also, it is usually good practice to place all the clouds in a layer different from that of the main scene, so you can easily hide them when they are fine tuned, thus avoiding visual clutter of the *3D Views*.

Of course, the best way to create clouds would be to use **MetaClouds**, but this tutorial demonstrates the use of fuzziness in this context. The above technique can be successfully used for many effects, such as smoke, fire, etc.



Additive Materials


Additive materials, often used in conjunction with fuzzy ones (see Fuzzy materials tutorial), allow to capture subtle lighting effects, such as the glow of a light in the dark, or an atmospheric halo around a distant planet, or the visible beam of light produced by a ray of sunlight breaking through the clouds.

Just like Fuzzy materials are an interesting alternative to Volumetric materials, additive materials are an interesting alternative to volumetric lights, because they also render so much more rapidly.

Additive materials always add their own color to what is behind them, which means that such a material will always be brighter than the background. Since a completely black material will not produce any light, black additive materials will be invisible. Turning them white will gradually increase the luminosity of the result.


Faking Volumetric Lights

Volumetric lights are a very computationally intensive method for generating visible beams of light. Additive materials provide an ideal, and incredibly efficient alternative for generating such effects:




Visible beam: step 5





Visible beam: step 7



Visible beam: step 8

1. Although any **Primitive** can be used for lighting effects, the cylinder is ideally suited for visible rays of sunlight (cones are best suited for spot lights, and spheres for point lights). Create a cylinder and stretch it using the **Resize in this direction** handle so it runs out of the scene on one side. Rename this as “Beam of light”.
2. Open the *Material Editor* for the Beam and press **New** (📄) to start designing the Beam material from scratch. Rename the material as “Beam material”. Select the **Transparency** tab of the editor.
3. Select the **Additive** box (notice the **Don’t cast shadows** icon is selected as beams of light don’t cast shadows). As expected, since the Empty material is black, the preview disappears except for the highlight.
4. Go to the **Colors** tab, and select the **Procedural colors** box, which results in an average gray color for the material. Now you are getting a very bright white circle with a darker area on the bottom left (the portion of the preview that used to be in the shadow).
5. We need to turn the brightness down, and make it constant along the beam: go to the **Effects** tab, and zero **Diffuse** lighting (diffuse lighting only appears on the side that is facing the light source, so is not homogeneous). Now, since the only lighting of the material is ambient, we are getting a homogeneous white surface.
6. If you find this is still a bit too bright, reduce **Ambient** lighting to 20%.



7. In order to make the limits of the light beam look less sharp, got to the **Transparency** tab and select the **Fuzzy** box. Push up **Fuzziness** to 70%. To compensate for the fuzziness, you will probably find out that you need to push ambient lighting back up.
8. Now we will add variations in the intensity of the light inside the beam. For this, we need to vary the color of the beam. Go to the **Colors** tab, and edit the **Color production function**. Create a **Perlin | Value** noise node, with a scale of 10 along Z. Variations of light intensity will appear in the preview (the same effect could have been achieved using the transparency function). Don't forget to select **Object-Standard** or **Object-Parametric** mapping to make sure that the material sticks to the orientation of the object it is applied to, letting variations follow the orientation of the Beam. If you find the variations to be too pronounced, make the black side of the color map brighter. You may also need to increase the global scale of the material.

9. Bear in mind that, to be realistic, such effects should remain subtle.

Ideally, you could add a light to illuminate the scene where the Beam hits the ground. Results are nice with cloudy gray landscapes. To vary the intensity of the light as it travels, add a layer to the color function that is oriented in the **Z** direction (**X** and **Y** scale 0).

Using Pictures inside Scenes

The following is an easy method for including 2D objects (i.e. pictures) inside your 3D scenes. Such objects can be used to achieve numerous effects (adding road-signs, statues, paintings...).

There is a **Primitive** that lets you do this easily: it is the **Alpha Plane** primitive. This primitive is designed in such a way that pictures can be easily and precisely mapped onto it, whatever the dimensions of the plane.



Making a Road Sign



Road sign: step 5




Road sign: step 8

1. When you want to include a picture inside a 3D scene, the first question you should ask yourself is: are all requested pictures available? That is, of course, the base picture, but also, a bump map if one is requested, and an alpha map if the object you want to create is not square (white areas of the alpha map will be transparent, black ones will be solid).
2. Once all pictures are available (use any bitmap application to generate them), you can start importing them into your scene. Create an **AlphaPlane** object. Rename it as “Road sign”.
3. The *Alpha Plane Options* dialog appears. Click on the **Load** button for the **Color**



picture, and select the bitmap to be used for the colors of the object (choose *Roadsign.bmp*) and press **OK**. If the bitmap contains alpha information, this information will be copied to the Alpha picture section.

4. Now we must take away the parts of the picture that should be invisible. Press the **Load** button for the **Alpha picture** and load the *Roadsign_a.bmp*.
 5. Take a look at the preview of the object to check that the cut-out areas (if any) appear in the right places. If this isn't the case, you might want to invert the picture by clicking the  button. Check the **Adjust plane proportions** option to automatically match the aspect ratio of the pictures in the Alpha plane.
 6. Click **OK** when you are satisfied with the settings. If you look at the **Aspect** tab of the *Object Properties* panel, you will notice that a new material has been created for the Alpha Plane, named after the color picture.
 7. Double-click on the material to open the *Advanced Material Editor*. If you look at the **Colors** tab, you will see that the material uses the picture you selected to generate its colors. It's the same for the Transparency function. If you edit the transparency function, you will see a mapped picture layer based on the Alpha picture.
 8. If we want to add bump to the sign, we'll switch to the Basic Material Editor and load the *Roadsign_b.bmp* map into the bump map group.
- The final result of this tutorial is available as a sample object named *Road sign*.

Glowing Neon Lights

The same technique can be used to create a glowing sign:


1. Create an **Alpha Plane** object and rename it as "Neon Sign". In the *Alpha Plane Editor*, load into the **Alpha picture** the *Neon.bmp* picture from the *Bitmaps* folder (or any picture created using your favorite 2D application – a picture with the text of the glowing sign).
2. It isn't necessary to load a picture in the **Color picture**. Press **OK** to close the editor, and double-click on the Alpha Plane's material in the *Object Properties* panel to open the *Material Editor*.
3. Go to the **Colors** tab and select **Procedural colors**. Double-click on the **Color map** and load the *Red* color map from the *Solid Colors* collection to create a uniformly red sign.
4. Switch to the **Effects** tab and raise **Luminous** to 100%. Click on the **Don't cast shadows** icon and the **Don't receive shadows** icon to select them.
5. Check the **Glowing material** option and set glow **Intensity** to 80% and glow




Radius to 40%.

The final result of this tutorial is available as a sample object named *Neon sign*.

Distant Planets



Distant planet: step 5



Distant planet: step 6

The following is a method to create moons and distant planets in your landscapes. It is based on the use of additive materials. The ***Moon*** object is the result of this tutorial. The big advantage of this technique is that you can make moons that have any shape and



color, instead of being limited to the predefined planets of the **Planet** primitive.

1. Create a **Sphere**. Rename it as “Moon”. Position it in the *Main camera view*, then move it away from the camera, so that it is very far away (and out of the rest of your scene, to avoid accidentally masking other objects, which would be disastrous for a distant planet!). Resize the Moon so that it has the right size in the main view.
2. Open the *Material Editor* for the Moon, and rename it as “Moon”. Go to the **Transparency** tab, and select **Additive**. Notice that the **Cast shadows** box is unselected. Leave it like that. The preview displays a very bright moon.
3. We need to get rid of the “disk-like” shape of the planet. Remove **Ambient** lighting using the **Effects** tab of the *Material Editor*. To avoid the moon being too bright, also reduce **Diffuse** lighting down to 20%.
4. We need sharper transitions from light to shadow. Push **Contrast** up to 60%.
5. Now go to the **Highlight** tab, and remove all highlights (you never see highlights on distant planets, do you?).
6. To add a crater pitted aspect to the surface (like that of a moon), we need to create a custom color map. Go to the **Colors & Alpha** tab of the *Advanced Material Editor*, and double-click on the color box to open the *Color Map Editor*. Select the **Black and white** color map from the *Rocks and Grass* collection.
7. Now, double-click on the **Color production** sphere to select a function to add texture to the planet. Select **00_GrainyFractal** from the *Basic* collection.

You now have a planet in the sky. You can experiment with different functions for a different effect.

Underwater Scenery


With just a few clicks you can get VUE to generate surprisingly realistic underwater scenery.

The sample landscape *Underwater* is a simple underwater scene that can be used as a good starting point for your own underwater scenery. You will find it in the tutorials collection. This is how it is generated:

1. First you need to create a **Water** plane (menu **Object | Create | Water**). Then, using the *Side view*, nudge this plane up several times so that you have sufficient space for your scene between the water plane and the ground plane.
2. Now you need to plunge the camera under the water (disable **Lock height above ground** first). Still using the *Side view*, nudge the camera down until it is beneath



the water plane.




Underwater: step 7


3. Underwater scenery is, like outer space scenery, one of the situations where fog and haze are undesirable. Open the *Atmosphere Editor*, and zero them both fog and haze density.
4. Light will fade out as it travels underwater. Open the water *Material Editor*. You need a convincing underwater fade out color. For instance, set the color to (0, 180, 190) for **Fade out color** using the **Transparency** tab. This will yield a nice greenish color. Green **Light color** is OK. It will give a green tint to all objects under the water surface.
5. Adjust the clearness of the water to get the right amount of fading. 30% works well.
6. To avoid an unnecessarily dark scene, make sure the **Overall transparency** of the water is close to 100%. Don't forget to reduce reflectivity accordingly, or else you will get luminous water. You'll notice the surface of the water is reflective, although no reflectivity has been defined for it: this is due to a natural optical phenomenon, called internal reflection, that, under certain circumstances, traps light inside objects that have a bigger Index of Refraction (IOR) than air.
7. To add the final touch, introduce a little bit of blurring to the camera: as light travels underwater, it often gets spread about by floating particles, et all. Push camera **Blur** up to 10%, keeping focus on the foreground.



Convincing Image Based Lighting




Standard IBL Render



IBL without atmosphere





IBL with atmosphere

In this tutorial we will see how Image Based Lighting is setup and more importantly, how you can adjust the atmosphere in VUE to match the background more accurately.

1. Create a new scene and add a sphere. Press **Drop** so that the sphere lies on the ground.
2. Open the *Atmosphere Editor* and select the **Environment map** model. Go to the **Effects** tab and load the HDRI image of a sky. You will find this in the Bitmaps Browser, **HDRI** collection, *skies.vim* file. Accept the offer to automatically setup the scene for Image Based Lighting. Rotate the camera so that it is facing the sun in the background (the camera should point West).
3. Remove all fog and haze and delete the sun. If you render the scene now, you will get a standard IBL render of your sphere. Adding a little reflection to the sphere will add that little pizzazz that you get in most IBL renders you see around!
4. Now create a set of new spheres and drag them away from the camera into the distance. Remove the reflectivity to get a better idea of the actual lighting. If you test render the scene, you will notice that the spheres further away don't blend in so well with the background image. They look as if they were collaged on top of the background picture. That's because your eye is sensitive to the overall atmosphere in the background, and doesn't understand why the spheres aren't subjected to this atmosphere. This is a very common problem with IBL. VUE solves this problem by letting you combine its volumetric atmosphere model with Image Based Lighting.
5. In the *Atmosphere Editor*, go to the **Fog and Haze** tab and increase the amount of fog and haze to match what is visible in the background picture (fog and haze densities of 80 should be about right). You can achieve interesting results by either applying or not applying the atmosphere effects to the background picture. Close the *Atmosphere Editor*.



6. Now reopen the *Atmosphere Editor* and when asked if you would like to create a sun, select **Yes**. In the *3D Views*, drag it so that it is located at the same point as the sun in the background picture (i.e. facing the camera). This is to ensure that VUE's volumetric atmosphere matches exactly the illumination in the background picture. If you enable **lens flares** for the sun, this adds the extra bonus of giving the impression that the sun in the background picture is actually creating a lens flare in the CG scene.
7. After a couple of test renders, you should be able to find the correct atmosphere settings that will give the impression that the spheres are really part of the same scene as the background picture.

The sample scene *IBLSpheres* illustrates this tutorial. You will find it in the *Tutorials* collection.

Varying Materials on EcoSystem Populations

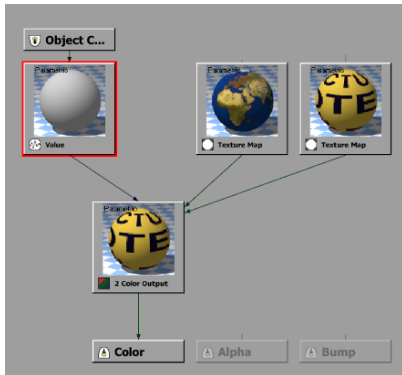


EcoSpheres: step 2






EcoSpheres: step 4





EcoSpheres: step 11




EcoSpheres: step 13

In this short tutorial we will see how you can change the materials of an object in an EcoSystem population without having to add the object several times to the EcoSystem population list.

1. Create a new scene and create a sphere. Double-click on the sphere's material preview in the *Object Properties* panel to edit the material. Map the sphere with a texture map of your choice (e.g. the **EarthMap** picture from the **Bitmaps** collection).
2. Save the sphere as a **.vob** for use in an EcoSystem.
3. Delete the sphere and create a plane. Drop the plane to the ground by clicking **Drop**. Double-click on the material preview in the *Object Properties* panel to edit the material.
4. Switch the material to **EcoSystem** and load the **Sphere.vob** into the EcoSystem population list. Populate the plane. If you render the picture, you will see that all the spheres have the same texture map. Now we will change that texture map on some of the spheres without adding a new object to the EcoSystem population list.
5. Go to the **Materials** tab of the *World Browser* and locate the sphere's material in the "EcoSystem materials" category. This category holds all the materials used by EcoSystem instances in the scene. Double-click on the material to open the *Material Editor*.
6. In the **Color** tab, switch the coloring mode to **Procedural colors** and edit the **Color production function**.
7. In the *Function Graph*, create a Color node by clicking on the **Add Color Node** icon (). Change the type of the node to **2 Color Output**.
8. Create a noise node (**Perlin Noise – Value** by default). Connect the input of the "2 Color Output" node to that noise node. Select the "2 Color Output" node and extract both colors by clicking the corresponding **Extract parameter** buttons (.
9. Connect the first color parameter to the texture map node that holds the **EarthMap** picture (this was created at the beginning when we mapped the sphere with the EarthMap picture).
10. Create another **Texture map** node and load another picture into it (e.g. the **Cactus Motel** picture). Connect the second color of the "2 Color Output" node to that new texture map node. Now, according to the value of the noise, the sphere is mapped using one texture or the other.
11. Currently, the texture used for each sphere changes on the surface of the sphere. We want both texture maps to be represented in the EcoSystem population, and we



want each sphere to have a unique texture map on its entire surface. For this, we will drive the noise node with the “Object Center” input. This input returns a constant value on the entire surface of the instance, equal to the center of the instance. Click the **Add Input Node** icon () and change the input type to **Object Center**. Connect the noise node to this input.

12. If you render the scene, you will see that the spheres in the scene have one of two different texture maps, although there is only one sphere in the EcoSystem population list. If you notice patterns in the way the textures are assigned to the population, you may need to increase the frequency of the noise (reduce the scale). If the scale of the noise is very small, its value will change many times over each instance, thus producing pseudo-random values at the center of each instance.
13. By using a **3 Color Output** node, or by using a combination of such nodes, we could easily increase the number of different textures without having to add objects to the EcoSystem population list. This is particularly useful when creating EcoSystems from very large polygon meshes. By avoiding the duplication of the object geometry, you save a lot of resources, but, by varying the texture maps used, you produce more variety in the population. For instance, if you are creating a crowd of people, you could have people with different hair colors based on the same character geometry.


The sample scene *EcoSpheres* illustrates this tutorial. You will find it in the *Tutorials* collection.

Creating a Piston Rig

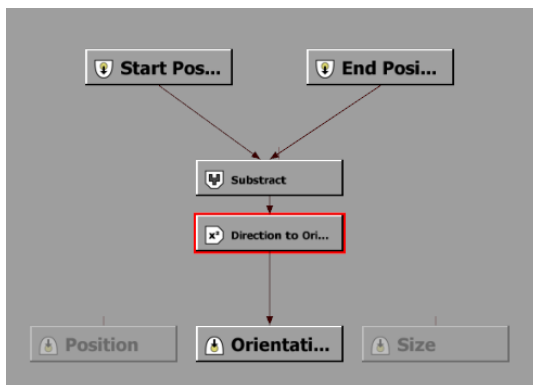
In this tutorial we will build an object with a complex behavior implemented using Object Graphs. This object is the *Piston Rig* that can be found in the *Miscellaneous/Dynamic Components* collection of the *Objects* library.

The piston rig will exhibit a piston that moves inside a sleeve within a constrained range.



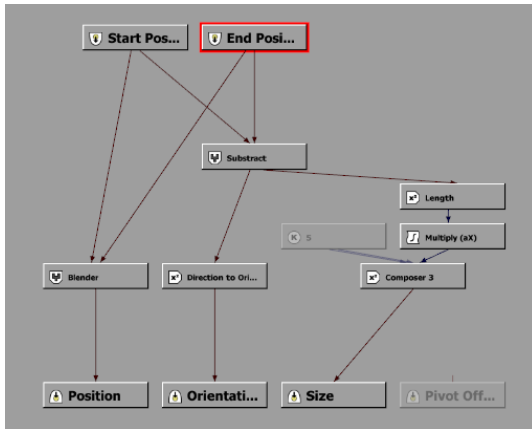


Dynamic Piston Rig

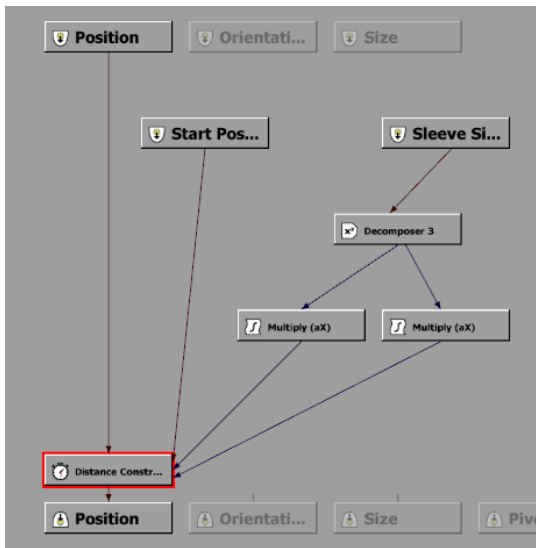


Piston Graph – Step 4






Piston Graph – Step 8



End Sphere Graph – Step 13


1. First, create a pair of spheres and move them apart by a few units. Reduce the size of the second sphere and rename both spheres as “Start” and “End” respectively. Assign them the **Black Porcelain** material from the **Basics** collection.
2. Create a cylinder and rename it as “Piston”. Assign it the **Mirror** material from the same collection. We will now create a graph to control the length and orientation of the piston so that it joins both spheres: select the Piston and create a graph for it by clicking the **Object Graph** icon (📐).



3. Create two **External Dependency** nodes using the menu command **Add Input Node | External Dependency**, one on the position of the the Start sphere, the other on the position of the End sphere (select **Start | Position** and **End | Position** in the respective **Dependency** drop-down lists).
4. We will set the orientation of the Piston so that it is always pointing at the spheres. Create a **Combiner | Subtract** node and connect its two inputs to the sphere position dependency nodes. Convert this into an orientation using a **Math | Vector Operations | Direction to Orientation** node. Close the **Object Graph** and check that the cylinder is now pointing at the two spheres.
5. Now we need to constrain the position of the Piston so that it is always in the middle of the two spheres. Create a **Combiner | Blender** node and connect the two sphere positions to the entries. Leave the **Ratio** at 0.5 and the **Combination mode** at **Blend**. This node outputs the middle point of the two entries. Connect it to the Piston **Position** output. Close the graph, and move the spheres around; notice that the Piston always stays in the middle, and always points at both spheres.
6. We will now adjust the length of the Piston so that it always joins both spheres exactly. Open the graph again and connect a **Math | Vector Operations | Length** node to the output of the **Subtract** node we created at step 4. This is the length of the piston. We need to force the **Z** length of the Piston to that value. Because the **Size** output expects a vector, we need to compose this value with a size along **X** and **Y**.
7. Create a **Constant | Constant Number** node and set its value to 5. This will force the diameter of the *Piston* to 5, which appears to be a nice value.
8. Now build the size vector using a **Math | Vector Operations | Composer 3** node that you will connect the first two inputs (**X** and **Y**) to the 5 **Constant** and the last input (**Z**) to the **Length** node we created at step 6. Connect this **Composer** node to the **Size** output, and close the graph. Notice how the Piston now joins both spheres exactly.
9. For now, there is no limit to the distance between the spheres, and the length of the Piston. We will now constrain the distance to a nice interval. Select the End sphere, and create a graph for it. In this graph, create a **Dynamics | Distance Constraint** node and connect it to the **Position** input and outputs. As it is, this constrains the position of the End sphere to 100 units from the World Center.
10. We want to constrain the distance to the Start sphere. Create an **External Dependency** on the Start sphere position. Extract the **Center** parameter of the **Distance Constraint** node by clicking the  icon and connect this output to the Start sphere position dependency. Discard the **Constant Vector** node that was created when you extracted the **Center** parameter.
11. For now, set the **Min Distance** to 20 and the **Max Distance** to 150 and check



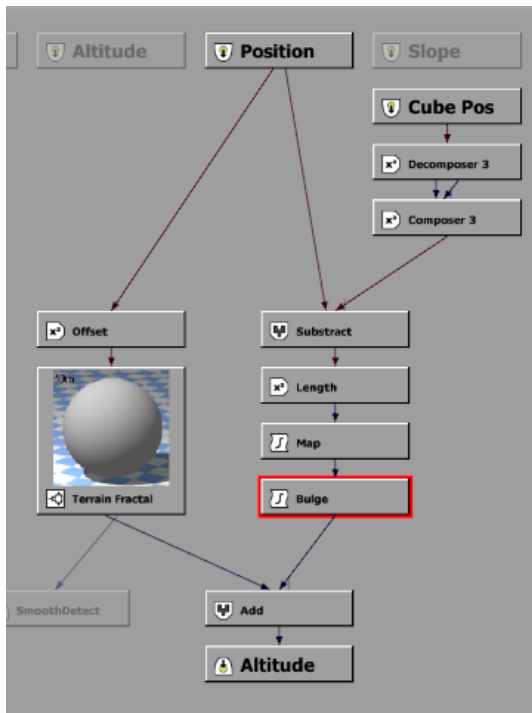
that the length of the assembly is indeed constrained to these values.

12. Now create a second cylinder, rename it as “Sleeve” and assign it the ***Black Porcelain*** material. Resize it so it is a little larger than the Piston cylinder. Open the Piston graph, select the two **External Dependency** nodes, the **Subtract** node and the **Direction to Orientation** and copy-paste them into the Sleeve graph.
13. Connect the **Orientation** output to the **Direction to Orientation** node, and connect the **Position** output to the ***Start*** sphere position dependency. This will ensure that the Sleeve is centered on the ***Start*** sphere and that it points at the ***End*** sphere.
14. The min and max length of the Piston was set to arbitrary values at step 11. In reality, they would be defined by the length of Sleeve. Re-open the End sphere graph and add an **External Dependency** to the **Size** of the Sleeve. Get the length of the Sleeve using a **Decomposer 3** node (the **Z** output is the Sleeve length). Multiply this by 1.5 using a **Filter | Multiply** node. This will be the max length of the Piston. Connect another **Multiply** node to the Sleeve length **Decomposer** node and set the **Multiply** by value to 0.5 this time. This will be the min length of the Piston.
15. Select the **Constrain Distance** node and extract both **Min** and **Max** distances by clicking the corresponding () icons. Connect these two extracted parameters to the nodes created above. The size of the Piston is now defined by the size of the Sleeve.

The sample scene ***Piston Rig*** illustrates this tutorial. You will find it in the ***Tutorials*** collection.

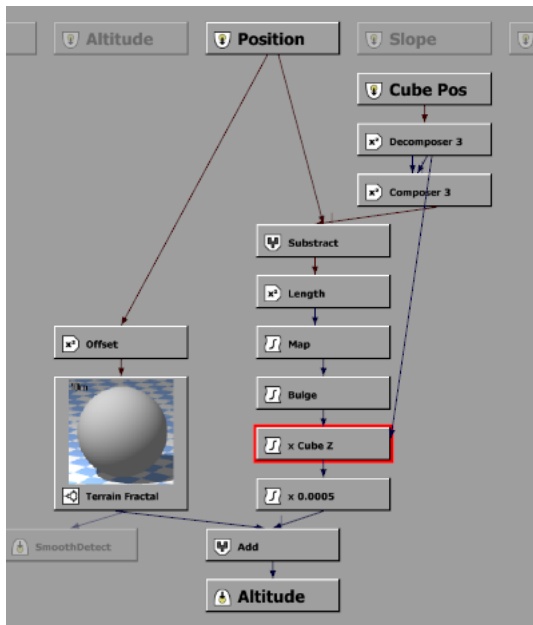
Interactive Procedural Terrain Modeling



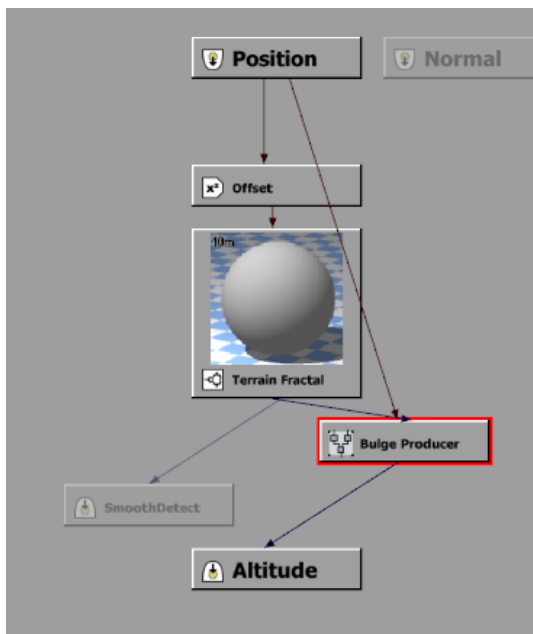


Adding the Bulge – Step 7







Bulge height control – Step 9




Bulge MetaNode – Step 11

In this tutorial we will see how the *Scene Graph* can be used to facilitate the interactive modeling of procedural terrains: we will use a set of cubes that will be used to “drag up” or “push down” the altitudes of the procedural terrain. It is recommended that you read and understand the previous tutorial before delving into this one.

1. Create a procedural terrain by clicking the **Load Procedural Terrain Preset** icon () and selecting one of the infinite terrain presets (you could make this work with a standard procedural terrain, but it would be more tricky since you’d need to take into account the position and scaling of the terrain).
2. Create a **Cube**. We will use this cube as a handle to control the procedural terrain altitudes. Hide the cube from render by clicking its icon in the *World Browser*, so it doesn’t interfere with the rendering.
3. Edit the procedural terrain altitude production function. We need to retrieve the position of the cube on the terrain map: in the function graph, create an **External Dependency** on the cube’s **Position**. Because the map is 2D, we only need the **X** and **Y** components of the cube’s position. Use a **Decomposer 3** node followed by a **Composer 3** node where you will only connect **X** and **Y** values.
4. We want the cube to create a “bulge” in the terrain. The “bulge” will be proportional to the distance to the cube. To compute this distance, use a **Combiner | Subtract** node to calculate the difference between the **Position** on the map and the **X,Y** components of the cube position (using the **Composer 3** node we just created). Calculate the length of this vector using a **Math | Vector Operations | Length** node. This is the distance to the cube in terrain map coordinates.
5. Based on this distance, we will create the bulge using a filter. Create a **Filter | Map** node to convert the length into a value that can be passed to a filter: use an **Input** range of 0 to 1000 (this will be the radius of the “bulge”). Because we want the “bulge” to be highest when the distance to the cube is 0, we need to invert the output of the map node: enter -1 and 1 for the **Output** range. Tick the **Clip out of range** values to ensure that no output from this node ever exceeds the -1;1 range.
6. This value can now be fed into a filter. Create a **Filter | Filter** node and edit the curve to define the profile of the “bulge”. Make the curve smooth using the **Smooth** curve option (). Create a key point at -1;0 so that the filter’s output range lays between 0 and 1. This is the height of the “bulge” that will be added to the terrain altitude. Rename the node as “Bulge”.
7. To add the “bulge” to the terrain altitude, use a **Combiner | Add** node, connected to the existing terrain altitude production node and to the “bulge”. Close the *Terrain Editor* and notice how a bulge appears in the terrain beneath the cube. Move the cube around, and watch the *Realtime Scene Preview* you will see that the bulge



follows the cube around (the OpenGL terrain preview is not refreshed – double-click the terrain to open the *Terrain Editor* and close it back to update the OpenGL preview).



8. It would be nice if we could control the altitude of the “bulge” by dragging the cube up and down. To do this, open the terrain altitude production graph and add a **Filter | Multiply** node after the Bulge. Extract the **Multiply by** parameter and connect it to the **Z** component of the Cube position **Decomposer**.
9. Because the typical altitude in the terrain production function is close to 1, while the **Z** component of the Cube can be very large, we need to scale down the effect of this multiplication by adding a second **Filter | Multiply** node to multiply the value by 0.0005.
10. Close the graph, and play around with the cube. Notice how the altitude of the cube now controls the height of the “bulge”.
11. In order to clarify the graph, you could select all nodes except the regular fractal terrain production nodes, and create a **MetaNode** by clicking the **Create MetaNode** icon (). Double-click on the **MetaNode** and select the **Map** node. Publish the **Upper** value of the Input range as “Bulge radius” by clicking the corresponding underlined field name. Select the **Filter** node and publish the filter as “Bulge profile”. Close the **MetaNode** graph and rename the **MetaNode** as “Bulge Producer”. You now have a high level interface to control the bulge, with direct access to the bulge radius and profile!

You could easily control the radius of the bulge with the same technique, using the size of the cube. The sample scene *Interactive Procedural Terrain* illustrates this tutorial. You will find it in the *Tutorials* collection.


Painting Materials on Terrains




VUE 2016 Infinite & xStream — Reference Manual





Painting Materials – Step 8





Painting Materials – Step 10

This tutorial leads you through the process of painting a terrain using two different textures – grass and rock.

1. Click on the Procedural Terrain icon () in the left toolbar of the VUE interface to create a procedural terrain.
2. In the *Object Properties* panel, select the **Load material** button () to load a terrain texture from the browser.
3. Choose the **Landscape** collection and select **Grass** (00a_Grass3.mat). This will cover the terrain with a grass material.
4. Now, double-click on the terrain to open the *Terrain Editor*. You will notice that the first color button reads **Grass** to correspond with the material selected in the *Material Editor*.
5. Click the *Top view* icon to position the terrain for painting.
6. In the **Mode** section, select the **Material** icon and deselect the **Sculpt** icon.
7. In the **Global** section, be sure **Airbrush** is checked and move the brush size slider to the middle. Move the **Flow** slider all the way to the right.
8. In the **Material** section, select the **Default rock** button, click on the **Material** brush and move your brush over the mountainous area of the terrain.

Remember, you can always reopen the *Terrain Editor* and repaint areas. If you see part of the rocky peak that you would like grass on, open the *Terrain Editor*, select Grass and repaint the area. Also, you can apply effects to the terrain, using the paint functions:

9. Open the *Terrain Editor* and select the **Pebbles Effect** brush. Be sure the **Sculpt** mode and **Material** mode is on.
10. Paint the stones in the grass areas of the terrain. You will now see brown rocks strewn through the grass areas.



Camera Mapping

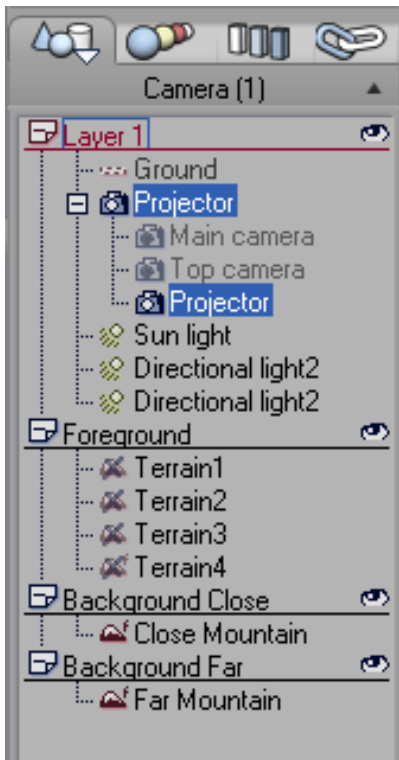


Basis scene – Cerro Verde

This tutorial is based on the *Cerro Verde* scene, which consists of two large distant mountains with many EcoSystem trees on them, and a dense EcoSystem in the foreground. Rendering an animation of this scene could be quite time consuming considering its complexity. If the animation just consists of changes in orientation or if the lateral camera movements are sufficiently small, camera mapping can be used to pre-render the distant mountains and project these renders onto simplified geometry (imposters), so that each animation frame will render quicker.

1. First, load the *Cerro Verde* scene and save it as *cameramapping.vue* .
2. We have to create a little order in the scene, so create a layer and move the first procedural terrain to it. Rename the layer **Background Close**.
3. Create another layer and move the next procedural terrain to it and rename the layer **Background Far**.
4. Create a third layer and move the remaining four terrains to it. Name the layer **Foreground**.
5. We will next hide all foreground elements which we will not pre-render since they may change during the animation in terms of position and orientation regarding viewing camera. It is important to keep in mind that good candidates for camera mapping are elements that are mostly seen from the same angle of view throughout your animation. Because of the parallax effect, these elements will generally be in the background. Let's hide all the elements in layer **Foreground** from render.





Camera Mapping – Step 5

Now a reference viewpoint must be chosen. The projector must match the viewing camera as much as possible to limit perspective distortions during the animation. We will proceed as follows:

6. Select the main camera at zero time, and click the **Store camera** icon button on the bottom left side of the render scene preview. This creates a new camera, **Camera 1**, that you can rename to “Projector”. This is this camera that will be used to project our pre-rendered background elements.
7. Make the **Projector** camera active by double clicking on it in the *World Browser*.
8. As the front distant mountain is much closer than the farther one, we will perform two independent renders for each mountain, and will project them on two separate geometries. This will preserve the overall parallax effect during the animation. However, before performing our reference renders, the animation of the Projector camera must be deleted from the *Timeline* so the projector will remain static throughout the animation.

When performing a reference render for camera mapping, all post render effects must be disabled, because they will eventually be applied during the final render including the camera mapped elements.




9. Edit the **Projector** camera, and make sure all post render options are disabled (natural film response, automatic exposure, lens glare and post processing). Also, any other lens effect such as lens flares, depth of field or motion blur must be disabled.
10. In the Render Options dialog, switch to **User Settings** mode, load the Broadcast preset, then uncheck **Depth of field** and **Motion blur** options.
11. Hide the closer mountain in layer **Background Close** from render. At this point, only the mountain in layer **BackgroundFar** should be visible at render. Launch a render to screen. Once the render is finished, save both resulting color and alpha channels to disk as *Far Mountain Color.bmp* and *Far Mountain Alpha.bmp*.



Camera Mapping – Step 11

12. Hide the **Background Far** layer element and unhide the **Background Close** layer element, then perform a render of the closer mountain, saving both resulting color and alpha as *Close Mountain Color.bmp* and *Close Mountain Alpha.bmp*.



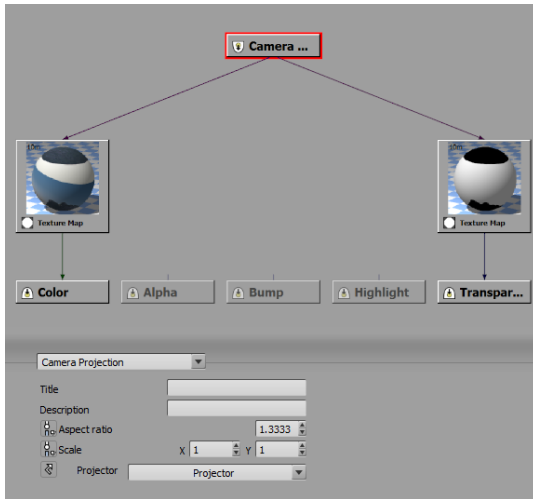
Camera Mapping – Step 12

13. Now let's create the impostor object for the far mountain. As we'll never get close to this mountain, we can afford projecting the render onto a simple alpha plane placed behind the close mountain. You will have to stretch the alpha plane so that it covers the whole visible surface of the far mountain, including trees. Reset the alpha plane's material.
14. The impostor for the close mountain must be more elaborate, because it is closer to the viewing camera, and the projection must be more accurate in order not to betray the effect. Moreover, we have to deal with the more complex transition between foreground components and the mountain. We'll use a simplified version of the mountain for this purpose. Just duplicate the close mountain, reset its material, edit the terrain and convert it to a heightfield terrain, which will be much faster to render.
15. Now the problem of the trees above the mountain remains because our impostor must cover any geometry, including trees above the mountain edges. One solution is to edit the impostor terrain, and paint extra altitudes so the terrain also covers the trees as needed.

The last steps consist of setting up the material for the imposters to use camera mapping with the pre-rendered mountains. We'll explain how to proceed for the far mountain, and the same procedure can be followed for the closer mountain.

16. Edit the alpha plane's material, and open the *Function Graph* for the color output. Delete any existing node and create a texture map node. Switch its type from **Projected texture map** to **Texture map** with an external UV projection node. Switch this projection node type to **Camera projection**. The camera projection node needs two pieces of information: the projected render aspect ratio (in our case: 1.33333), and the camera used for projection (in our case: Projector).





Camera Mapping – Step 16

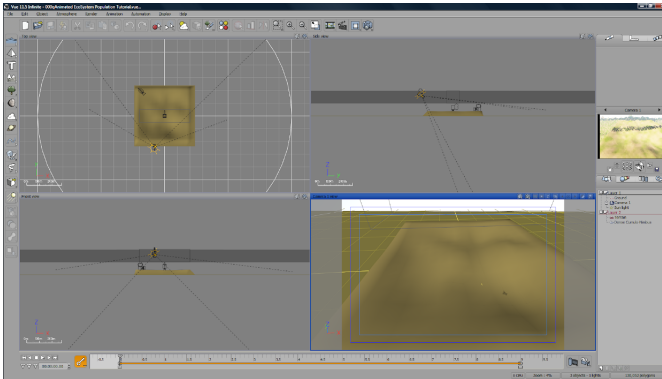
17. Switch back to the texture map node, and load ***Far Mountain Color.bmp*** . Change the interpolation mode to **Bicubic** to avoid aliasing at render.
18. Let's add another texture map node for the transparency output, so that the camera mapped impostor will be properly masked out using the alpha channel of our rendered mountains. You can use the same camera projection node as the input for both texture map nodes. In the texture map node used for transparency, load ***Far Mountain Alpha.bmp*** and change the interpolation mode to **Bicubic**. The transparency map must also be inverted because alpha represents opacity.
19. We are almost done. One final problem remains: our rendered mountains already contain atmospherics and lighting, so these components shouldn't be rendered again.
20. Materials now have two extra buttons in the the *Material Editor* to make them ignore atmospherics (☒) and/or scene lighting (☒). Let's check both buttons, so that the material color appears unchanged at render, effectively producing the very same color as the pre-rendered map.
21. Proceed the very same way for the close mountain impostor's material.
22. Finally, hide the original mountains from render so only the impostors are rendered. Activate the main camera and enable all of the needed post render and lens effects. You can render an animation at a fraction of the time it would have taken by rendering the whole geometry at each frame.



Animated Dynamic EcoSystems

This tutorial gives you a basis for creating an animated dynamic EcoSystem.

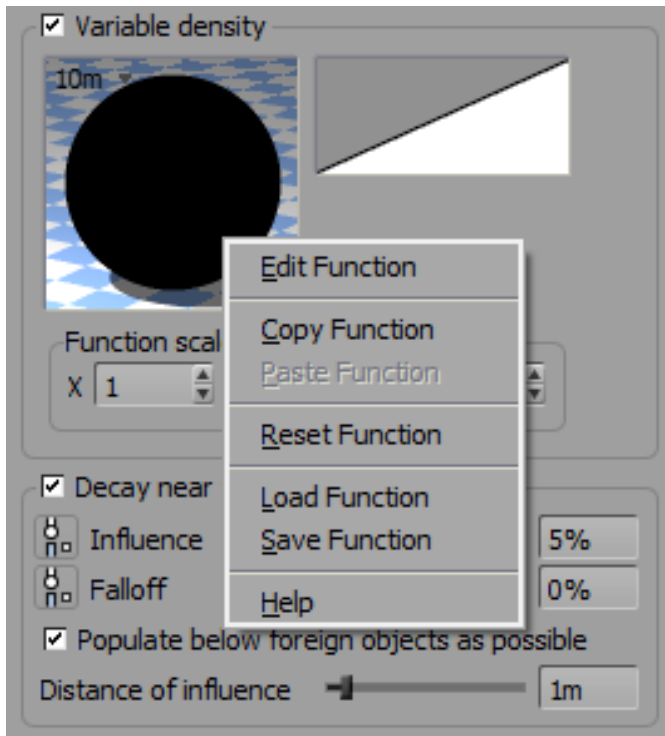
1. Start by loading the *Animated EcoSystem Population Tutorial.vue* scene.



Animated Dynamic EcoSystem – step 1

2. Select the terrain. In the *Object Properties* panel, right-click on the terrain material and select **Edit Material** from the menu.
3. In the *Material Editor*, select the **Density** tab, right-click on the picture and select the **Edit** function.



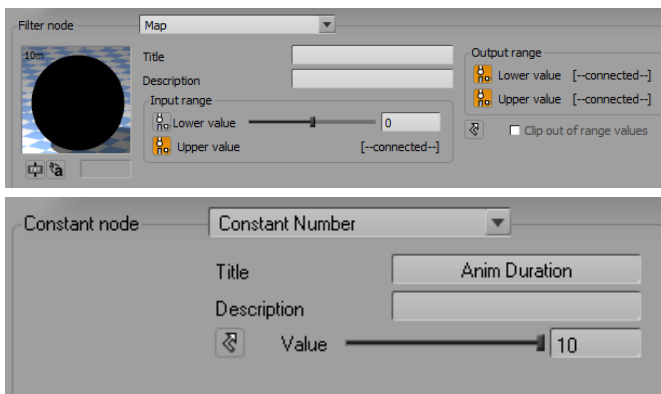


Animated Dynamic EcoSystem – step 3

4. The *Function Graph* displays with a simple material structure and a Density structure. We will analyze what each part of the Density function does to produce an automated population of this terrain.
5. We will start by looking at the **Threshold** node (**Add Filter Node | Threshold**). This node controls the movement of the EcoSystem along the terrain. The **Threshold** node takes 3 parameters: a **Threshold**, a **Low** value and a **High** value. For each value received by the node as input, the node will compare the value with the threshold. If the input value is inferior to the threshold, the node returns the **Low** value as output, and in the contrary case, if the input value is superior to the threshold, the node will return its **High** value. It should be noted that the **Threshold** creates a step which is a necessary part of this function. Animated EcoSystem populations will not work well if the density varies smoothly. Instead, it needs to vary suddenly, since strong flickering will occur in any areas of transition due to the fact that plants are being added and removed randomly.
6. As we want to move the plants along the **Y** axis, we will connect the **Threshold** node to the Y-Position of the image. To extract the **Y** component of a 3D position we use a **Decomposer 3** node (**Add Math Node | Vector Operations | Decomposer**



- 3). This is then connected to the **Position** input node.
7. We could directly attach the **Threshold** node to the **Time** input node, but we need to define it more. That is why we will add a **Map** node (Add Filter Node: **Map**) to apply a global scaling to time values (so that 1s in time is converted into more than 1 VUE Unit in space).
8. Let's take a look at the **Map** node's description: it takes 4 parameters, 2 values in **Lower value** and **Upper value** for the **Input Range**, **Lower value** and **Upper value** for the **Output range**. Basically, this node can be used to map values which are in range (**Input Lower value**, **Input Upper value**) to the second range (**Output Lower value**, **Output Upper value**) in a linear way (for instance, an input with a value of **Input Lower** will correspond to an output of **Output Lower**, an input with a value of **Input Upper** will correspond to an output of **Output Upper**. All intermediate input values will be mapped in a linear way (for example, the middle of the segment (**Input Lower**, **Input Upper**) will correspond to the middle of the output range (**Output Lower**, **Output Upper**). In our case, the **Lower** input for time is zero (we will start our animation at time 0). The **Upper** will be the total duration of our animation. A **Constant** node has been created to represent it, as this can then be modified in the future. This was created by clicking on the icon next to the **Input range**, **Upper value**. In its description shown below, **Constant Number** is selected from the drop box and **Anim Duration** was keyed in the **Title** so that it appears in the *Function Graph* box. A **Value** of 10.0 is entered for a ten second animation.

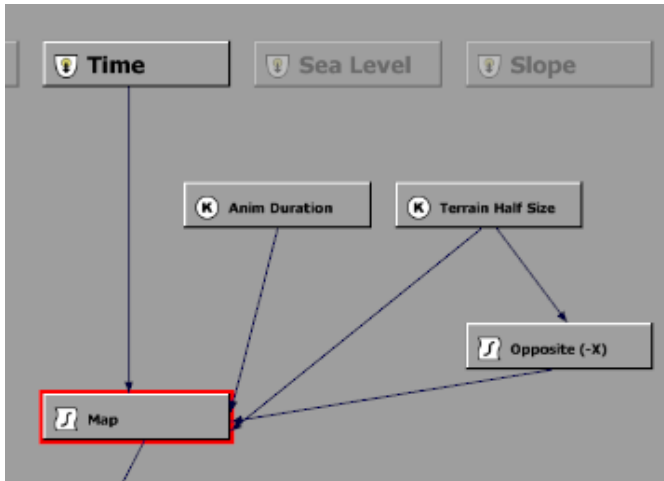


Animated Dynamic EcoSystem – step 10

9. Since we want the **EcoSystem** population to start at the bottom of the image close to the camera and continue to the end of the terrain, an **Opposite** ($-X$) node (Add Filter Node: **Opposite** ($-X$)) is needed. The 0 value of the map is reached at the middle of the terrain. In order to cover the entire terrain, the value of the map limits has to range from one end of the terrain to the other, which is **Terrain Half size** on one side, and minus **Terrain Half size** on

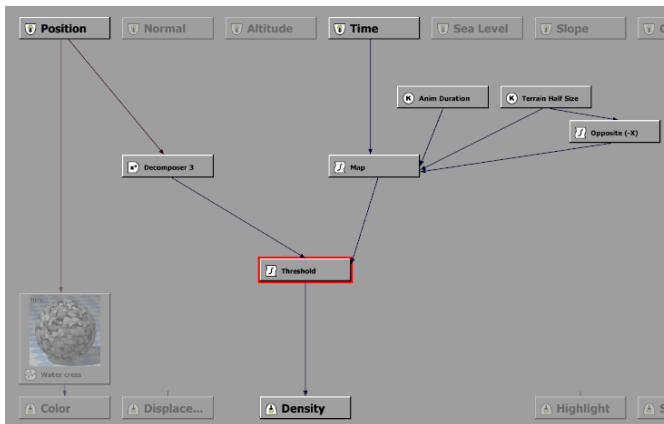


the other end.



Animated Dynamic EcoSystem – step 12

- The **Time** input node is now linked to the **Map** node whose value is now mapped from **Anim Duration** to **Terrain Half Size**. The **Threshold** node is then linked to the **Map** node. You can save this function so that you can reuse it, perhaps naming it *animated_threshold.fnc*.



Animated Dynamic EcoSystem – step 13

- Now, click **OK** to close the *Function Graph* and render the animation. Watch the EcoSystem slowly move up the screen.
Your results should look something like this:





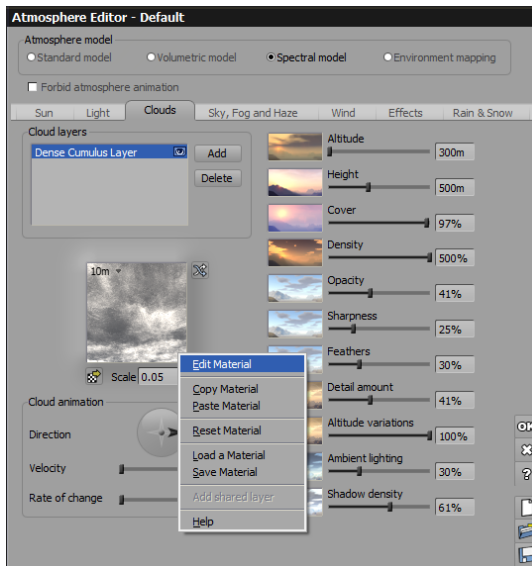
Animated Dynamic EcoSystem – at frame 0



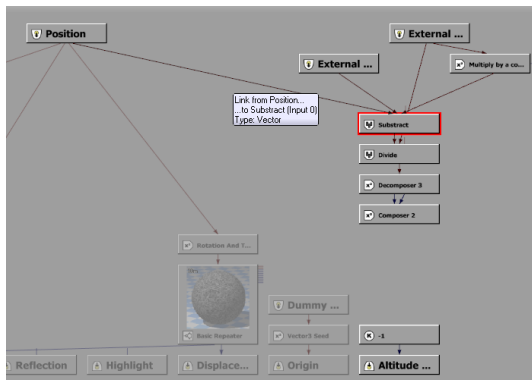
Animated Dynamic EcoSystem – at frame 192



Clouds Modulating with Landscapes



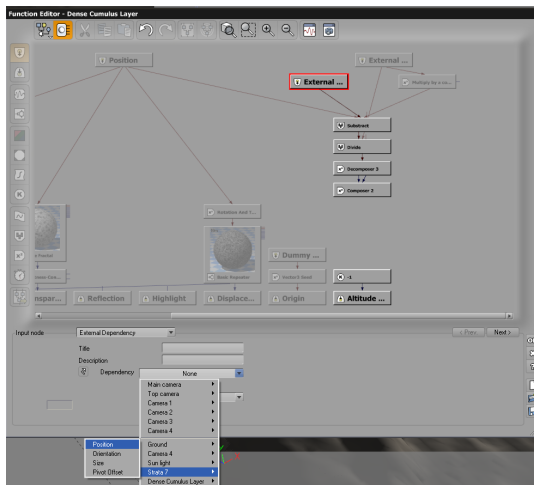
Clouds Modulating – step 3



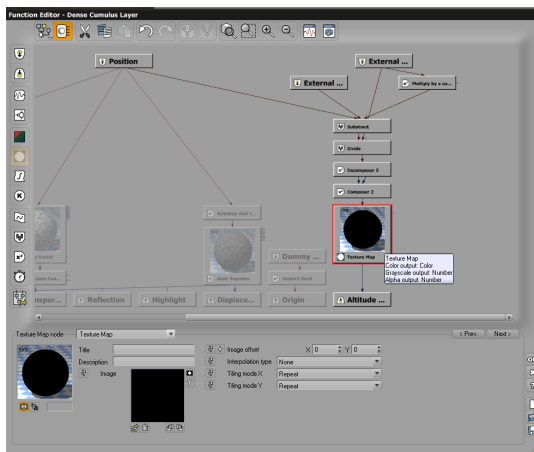
Clouds Modulating – step 6



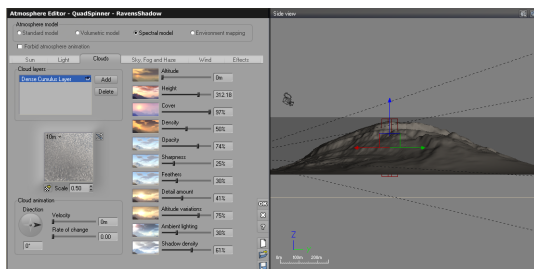
VUE 2016 Infinite & xStream — Reference Manual



Clouds Modulating – step 7





Clouds Modulating – step 9



Clouds Modulating – step 11

This tutorial shows you how to create low-hanging clouds that follow the contours of a landscape.

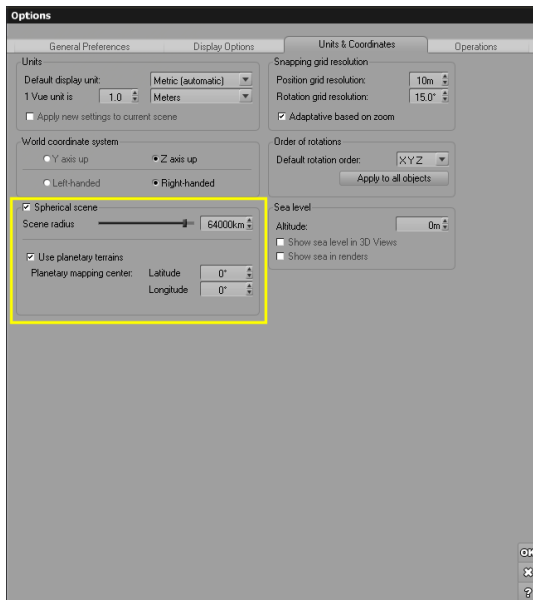
1. Create a terrain, either standard or procedural. We need to save the heightfield map created by this terrain as we will be using it later. Use the **Copy** function in the *Terrain Editor* to copy the heightfield. Then open a program such as Photoshop, create a new blank image and paste the heightfield into that image and save it.
2. In the *Atmosphere Editor*, select the **Cloud** tab and click to add a cloud layer. From the *Spectral 2* collection, select the *Dense Cumulus Layer*.
3. Right-click the picture of the cloud layer and select **Edit** to open the Advanced Material Editor.
4. On the **Color Density** tab, in the **Cloud Modulation** section, for the **Drive altitude offset**, select the icon to drive with a function (). This opens the *Function Graph*. In the *Function Graph*, we will be setting our altitude variation effect for the cloud to follow the geometry of the terrain below it.
5. Load the *MapToObject* metanode provided with VUE and **Ungroup** () it.
6. Connect **Input 0** of the **Subtract** node to the **Position** input.
7. Select the first **External Dependency** node (the one connected to **Input 1** of the **Subtract** node) and in the Dependency parameter, select terrain position.
8. Select the second **External Dependency** node (the one connected to **Input 2** of the **Subtract** node) and in the Dependency parameter, select terrain size.
9. Delete the **Constant** node that is connected to the **Altitude offset** output. Insert a **Texture Map**. Connect it to the **Composer 2** node and to the **Altitude offset** output, then load your heightmap previously saved (in step 1). Close the *Function Graph* by clicking **OK**.
10. Back in the *Atmosphere Editor*, set the **Altitude** and **Height** of your cloud layer to match your terrain. You need to adjust the **Height** parameter functions of your heightmap: i.e. if your heightmap contains pure black (0 0 0) and pure white (255 255 255), you will be modeling the cloud altitude within the full height of your cloud layer. If for instance your heightmap does not go up to pure white but is limited to mid grey (128 128 128), your cloud layer will be modulated within half of the cloud layer height, and you will have to adapt the height parameter of the cloud layer or modify your heightmap consequently.
11. Play with the **Altitude variations** parameters to adjust the thickness of your modulated cloud layer. Setting **Altitude variations** to 100% will result in an infinitely thin cloud layer. Setting it to 75% will result in a modulated cloud layer thickness of 25% of the cloud layer total height.



Now, when you're finished, you will have a great cloud layer that follows your terrain geometry.


For even better looking results, you can follow the same steps above to drive the 3 other modulation parameters: **Altitude offset effect on Z**, **Height modulation** and **Density modulation**, or use any type of other functions (instead of the metanode) to add more variation effects to your cloud layer.

Creating a Planet

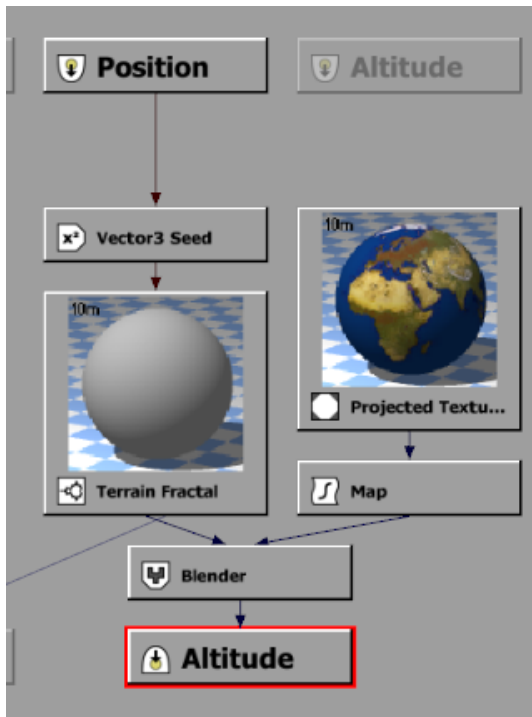


Creating a Planet – step 2






Creating a Planet – step 3



Creating a Planet – step 7





Creating a Planet – step 14



Creating a Planet – final render

This tutorial demonstrates how to use the new planetary spherical terrains to design an entire planet, as well as fully spherical cloud layers over the entire planet. Both terrain and cloud layer will use a mix of texture maps and procedural functions, in order to offer multi-scale components that can be viewed from very far away as well as from very short distances without loss of detail. Once such components are properly set up, everything is handled automatically and you just have to place the camera where you want to render from a given viewpoint.

1. We'll start by configuring a scene in planetary mode. Create a new empty scene.
2. Open the general *Options* dialog, **Units & Coordinates** tab. In the *Spherical* scene section, check both **Enable spherical scene** and **Use Planetary terrains** option. We'll leave the scene radius to the default value of 6400km, which is the actual radius value for planet earth. Click on **OK** to close the dialog.
3. The interactive display will switch to spherical mode, and if you zoom out enough in the viewports, you will see the entire 6400km radius sphere.



4. Now, we need to configure an atmosphere for the planet. We will configure one similar to that of earth.

5. Using the *Default* atmosphere in the *Spectral Sunshine* folder, make the following changes in the **Sky, Fog & Haze** tab of the *Atmosphere Editor*:

- **Sky ground density:** 100%
- **Sky mean altitude:** 8,800 meters
- **Decay Amount:** 40%
- **Haze ground density:** 20%
- **Haze mean altitude:** 1,200 meters

You can of course tweak these values as you wish, they are just listed here as a reference. For such an atmosphere which corresponds to a clear day, luminosity might be too strong by default at render. You can then lower the camera exposure, just like the human eye does naturally. Another option is to enable automatic exposure, if you don't want to manually tweak exposure value. You can load the *Planet Tutorial Start.vue* scene which is set to spherical mode with an earth-like atmosphere properly set up.

6. Now, let's create an earth-like procedural terrain. First, create an infinite procedural terrain by right-clicking on the **Procedural terrain** object icon in the vertical toolbar, and accept to replace the default ground plane. For now, say **No** to adding sea level to the scene.

7. Edit its altitude function by right clicking on the **Altitude production** sphere, building a mix between a fractal function for texturing and a bitmap (The map of the earth is available in the *Bitmaps* collection). You can load the *Planet Tutorial Step 1.vue* scene to see how it can be done. Earth elevation maps work best when creating planetary terrains.

8. Next, we add oceans to the planet. We begin by adding a water plane from the object vertical toolbar and move it vertically as needed to leave some land above the surface. A high viewpoint can be useful during this step to actually see how the water surface covers the planetary terrain. If you load the *Planet Tutorial Step2.vue* scene, it contains a planetary terrain covered with oceans.


9. Of course, the planet needs clouds. A volumetric spectral cloud layer works the best. We need to configure it to use a cloud texture map for large scale density distribution while keeping its original procedural density function at a smaller scale for close-up detail.

10. Starting with the basic default spectral atmosphere in the *Daytime/Spectral Sunshine* folder, open the *Atmosphere Editor* and on the **Clouds** tab, select the *Planetary Cloud Base.mat* layer from the *Clouds/Spectral Collection*. This adds a dense cloud layer that covers the entire planet.




11. Raise the layer up to roughly 7 kilometers so the layer is placed above the main camera. You can also raise the layer thickness to 8 kilometers to get richer variations from outer space.
 12. Then right click on the cloud picture and select **Edit Material**.
 13. In the *Advanced Cloud Material Editor*, **Large Scale Density** tab, select to **Use planetary cloud density map** and load the *planet_clouds_8k.jpg* map from the **Bitmaps** collection into the image slot. Make sure the **Influence on density** setting is set to 100%. **Influence on height** should be set to 50% and **Limit wall effect** to 50% as well. With these settings the cloud map will also be used to drive altitude variations without introducing cloud “walls” where the map suddenly switches from dark to light.
 14. At this point, you can switch your camera to the top camera and move it some distance from the planet so that you can view it from space.
 15. You can also load the *Planet Tutorial End.vue* scene to see the planet from outer space.
- A final render will give you a lovely planet with landforms, oceans and nice swirling clouds. Switch to the Main camera to view a nicely detailed planet landscape.

Creating a Road Using a Spline




Creating a Road – Step 3






Creating a Road – Step 4




Creating a Road – Step 6










Creating a Road – Step 8





Creating a Road – Step 10

This tutorial shows you how to create a road on a terrain. And add an EcoSystem as well.

1. Create a procedural terrain using the procedural terrain icon () on the left toolbar.
2. Right-click on the **Spline** icon and select the first option to create a spline ()
3. In the **Spline Editor**, add some points on the terrain using the **Add point** icon ()
4. Activate the terrain effect by clicking on the **Terrain effect** icon ()
5. Set the **Width** to 40m and keep the **Shape** as is. Be sure to check the **Procedural Terrain** in the **List of influence**. This creates a roadbed.
6. Now, activate the geometry effect by clicking the **Geometry effect** icon ()
. Since the roadbed we created was 40m wide, set the road width to 32m. Select **Road** in the **Type** dropdown with a Z-twist. Click the **Load material** button to open the *Material Editor* and select a material for your road. There are several available in the **Road** collection. Your terrain should look something like in the illustration.7. Or, if you wanted to generate an EcoSystem along that spline you created, click the **EcoSystem effect** icon ()
8. Set the width a little wider than the road effect – try 36m. Select **Populate with an EcoSystem** in the **Stroke** section and leave **Populate on the spline** unchecked. Double-click on the material preview to add a plant and populate. The result should look like this.
9. To get the trees out of the middle of the road, select the spline object in the *World Browser* and click the **Duplicate** button () in the top toolbar. In the new spline, close the **Terraineffect** and **Geometry effect** tab. In the **EcoSystem effect** tab check **Cut out existing EcoSystems** in the **Stroke** section and uncheck **Populate with an**



EcoSystem. Set the width equal to the road (in our case 32m).

10. In the *World Browser*, select the first spline you created, open its EcoSystem material in the *Material Editor* and repopulate the spline. That should do it – trees along the side of the road.



Building a Complete Scene




Vertigo Tutorial: final render

This tutorial will show you an example of how complete scenes can be designed with VUE. The scene is entitled “Vertigo” and features a wooden bridge disappearing in a dense forest. The instructions in this tutorial are much “looser” than in other tutorials to leave a wider share of personal interpretation.


Use the different paragraphs in this tutorial as a framework when designing your own scenes. A good understanding of the basics of operating the software is recommended before you delve into this tutorial...

Shaping out the Forest



Vertigo Tutorial: step 2




1. Create a new scene. You will be prompted to select an atmosphere for this new scene. Just pick one that corresponds more or less to the results you want to achieve. You can always select another one later, if you finally decide that you made a bad choice.
2. Most of the scene is going to be a forest. Create a large terrain, and stretch it out so that it occupies most of the *Main camera view*. Reduce its vertical amplitude by using the resize handles. Double-click on the terrain to open the *Terrain Editor*, and push up its resolution. Add bumps and a tad of erosion to give it a rough look.
3. Notice how the Terrain has been added to the list of objects in the World Browser (the list on the right side). Click on it, and change the name to “Distant forest”. Labeling your objects clearly is extremely important as it will enable you to find your way rapidly inside your scene as it grows more complex.
4. The *Object Properties* panel displays information relative to the selected objects. Select the terrain, and click on the **Load** material button (). Now choose something that looks more like a distant forest than the default material. We chose **Grass** from the **Landscapes** collection.

The completed scene is available as `Tutorials\Vertigo\ Tutorial\Step\ 1.vue`

Adding a Distant Background

5. In order to give more depth to the picture, we will now add a far background of a large and distant mountain. Add two terrains in the background. The first will be a kind of transition with the forest, the second a large mountain.




Vertigo Tutorial: step 5

6. Don't be afraid to make those terrains really large and distant. Resize them verti-



cally and stretch them laterally to increase their visible size on the picture. Assign them the **Grass** material too (just drag the material from the “Distant forest” terrain onto these new terrains).

7. Now let's make the grass turn into rock with altitude. Select the mountain terrain, and double-click on the material in the *Object Properties* panel to open the *Material Editor*. Select the **Mix materials** checkbox, and load **Clumpy** rock from the **Rocks** collection into “material 2”. This will mix the grass with the rocks. Using the **Influence of environment** tab, indicate that the rock material appears at high altitudes, and on steep slopes. Use the **Mixing proportions** slider to adapt the amount of rock visible on the mountain.



Vertigo Tutorial: step 7

8. Using the *World Browser*, drag the two terrains into Layer 2, and rename this as “Mountain range”. Layers let you instantly hide, lock or activate whole chunks of your scene, thus avoiding screen clutter.

The completed scene is available as **Tutorials\Vertigo\ Tutorial\Step\ 2.vue**.

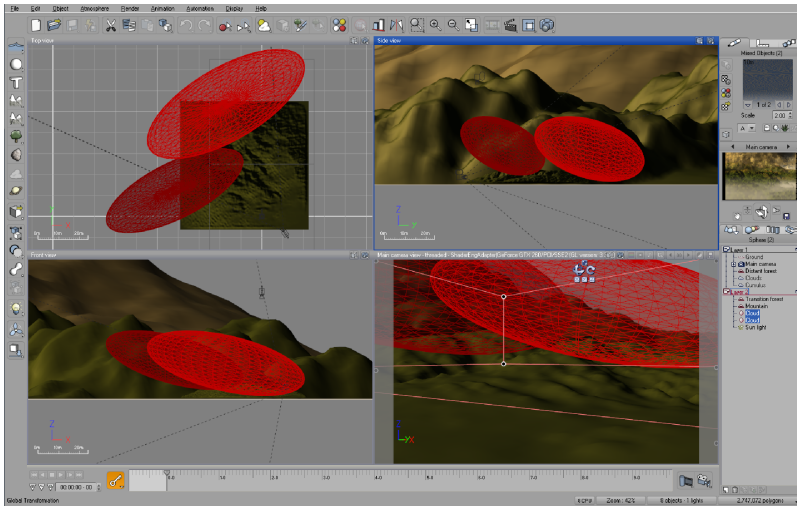


Tuning the Atmosphere

Before we make any further progress, we will improve the atmosphere of the picture.

9. We want the warm light from the sun to lick the scene from the right. Select the sun in the *Side view* and drag it down to the right of the camera. Choose an orange color for the sun color. Using the *Atmosphere Editor*, **Light** tab, modify the ambient light color to give it yellow tones too.
10. Now we will add fog at lower altitudes. Go to the **Fog and Haze** tab, and create altitude dependent fog using the **Altitude dependent fog** sliders. Select a pale gray color for the fog.
11. Next, go to the **Clouds** tab and increase the density of the clouds close to the horizon. Change cloud illumination boost to control how clouds are illuminated depending on their position relative to the sun. Reduce the size of the clouds using the cloud material **Scale** control.
12. We want to add clouds over the transition forest, in front of the large mountain. Create a couple of spheres and assign them the *Cloud Sphere #6* material from the *Other Clouds* collection. Make them **Wireframe** preview using the Preview options in the *Object Properties* panel. Drag the clouds above the transition forest and enlarge them until they are as large as the terrain. Having enlarged the spheres, we also have to increase the scale of the attached material. Open the *Summary of Materials* panel and push the scale of the cloud material up. The *Summary of Materials* is a very useful panel that displays all materials used in the scene at any given time. Modifying materials here ensures that all objects using that material will be modified accordingly.





Vertigo Tutorial: step 12

13. We've finished work on the background, so you can hide the “Mountain range” layer (click on the eye to the right of the layer title). Although the objects in that layer don't appear in the views, they will still render just the same.


The completed scene is available as **Tutorials\Vertigo\ Tutorial\Step\ 3.vue**.

Adding Vegetation

For this picture, bushes, plum trees and dead trees were used. If you have the Walnut or Sonneratia extra plants (that can be purchased separately), you may want to replace some of the plum trees by these.

14. Select the **Object | Load Vegetation Species** menu command and create a group of walnut and plum trees (ca 100 in all), and place them in front of the distant forest. Since Vue trees are very detailed, notice how the polygon count has jumped up to over 5 million polygons. This hardly slows the program though, thanks to the fully multi-threaded architecture.





Vertigo Tutorial: step 14


15. Position the trees on the terrain, then drag them up and press **Drop**, to position them precisely on top of the terrain. Now add a few bushes. Group plants that are close together. This will increase rendering speed significantly.

You could also assign to the terrain an EcoSystem material and populate it with the different plant species. The completed scene is available as `Tutorials\Vertigo\ Tutorial\ Step\ 4.vue`.

Importing the Subject


Now we are ready to import our subject: the bridge, and the dead branch and plum tree in the foreground.

16. The bridge was created in LightWave, but any application capable of exporting in a format supported by VUE could have been used. Select **Import Object** from the **Objects** menu and load the *Bridge.lwo* file. Position and resize it so that it starts out of the right corner of the picture, and leads to the trees that you have just added in the previous step. The material used for the bridge is a mix of wood and moss, created in much the same way as the grass and rock material from above. A distribution function was used to create the clumps of moss.



Vertigo Tutorial: step 16

17. Create a dead tree branch by clicking on the vegetation icon and picking **Dead tree** from the Visual Plant Browser. Resize and position it in the bottom right corner of the picture. Don't worry about placing it close to the camera: VUE plants look great, even up close! Place a plum tree in the top right corner, to fill up the empty space.
18. You can use render blow-ups to check that positioning is just right.




Vertigo Tutorial: step 18

The completed scene is available as `Tutorials\Vertigo\ Tutorial\Step\ 5.vue`.

Tuning the Light

19. This is one of the longest (and most important) steps. It's often a question of "test rendering" regularly.
20. One spot light was aimed at the subject to emphasize the dead branch in the foreground. It was given a yellow color.
21. Another spot was placed to add some light to the group of trees at the end of the bridge. This one was slightly green.
22. Finally, soft shadows were turned on for the sun light using the *Object Properties* panel when the sun is selected. This increases realism (but slows down rendering quite significantly if you use a better preset render quality such as **Final** or **Broadcast**).
23. The scene is now complete. Congratulations! It comprises about 4 million polygons, and 3 lights. It takes a minute to render 800 pixels wide, **Broadcast** quality with soft shadows. In **Final** quality, the same picture renders in just a few seconds...





Vertigo Tutorial: step 23

The completed scene is available as `Tutorials\Vertigo\ Tutorial\Step\ 6.vue`.

In this tutorial, we saw how easy it is to create realistic scenery, cleverly placing terrains and vegetation, and importing objects from other 3D applications.

We cannot stress enough that you organize your scenes into the layers. Hiding or locking these layers highly increases the response time of the software.






Animation Tutorials

This section provides a set of tutorials designed to teach you part of the art of animating in VUE. The focus is not on building impressive animations, but to concentrate on achieving one type of effect in each tutorial. By combining these effects into your own animations, you will achieve impressive results.

The tutorials are presented in “incremental” order, which means that some features explained in the preceding tutorial may be used without further explanation in the following tutorials. This is why it is recommended that you read and practice the tutorials in the correct order.

A Pursuit

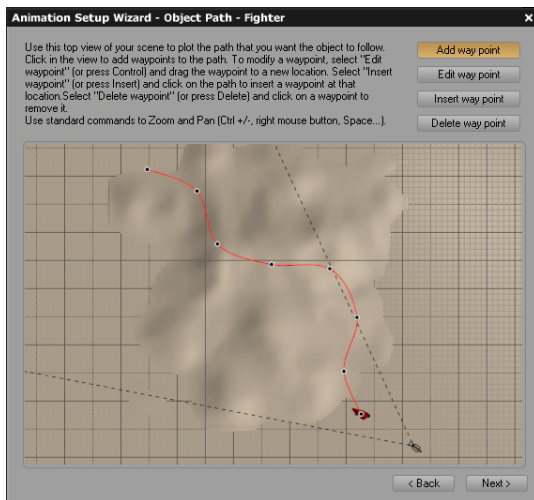
In this basic tutorial you will learn how to animate objects using the *Animation Wizard*. The tutorial features a race between two spaceships closely following the ground.

1. Create a new scene and load the **Default** atmosphere from the **Spectral Sunshine** collection.
2. Create a new Terrain and double-click on it to open the *Terrain Editor*. Double the resolution of the terrain using the **x2** button. This is because we will be flying close to the terrain, so it has to be detailed.
3. Press the **Mounds** button () to add many small mounds to the terrain, then press the **Eroded** button () to round off the terrain. Add some amount of **Diffusive Erosion** using a **Soft rock** setting to increase the influence of the erosion. This is because we need the surface we will be following to be relatively smooth, if not the animation will be much too shaky.
4. Using the Resize handles, increase the size of the terrain so that it fills up the *Top view*. Squash the terrain so that it is very flat (once again to reduce the amplitude of the bumps in the terrain).
5. Select the menu command **File | Import Object** and load the **Fighter.3ds** file in the **Objects\Swma\Fighter** folder. Resize the fighter so that it is real tiny compared to the terrain, and position it at the bottom-right corner of the terrain (in *Top view*).
6. Press the *Timeline* icon (). The *Animation Wizard* appears. In step 2 of the Wizard, select a **Motorbike** motion type. Now it may seem strange to select a Motorbike motion for a fighter... The idea is that the fighter should follow the




surface of the terrain, as if it were hovering above the ground. While selecting an Airplane motion would seem more logical, Motorbike motion maintains the object at a constant altitude above the ground. With Airplane motion, the fighter would fly straight above the terrain instead of following its surface (it looks good too!). The best is to experiment with the different types of motions until you become familiar with all the effects available.

7. In step 3 of the *Animation Wizard*, select a **–Y Main axis** so that the fighter points in the direction of travel.
8. In step 4 of the *Animation Wizard*, plot an undulating path that ends on the top-left corner of the terrain. In step 5, enter a total duration of the animation of 10 seconds. Finish the Wizard.





A Pursuit: step 8

9. Position the camera so that it looks at the fighter, and select a **Television Aspect ratio** from the **Render Options** dialog (alternate action of the **Render** button



A Pursuit: step 9

10. Now we will use the Wizard again to make the camera loosely follow the fighter. However, before doing that, we must make sure that the path followed by the fighter remains visible (for reference) when we deselect it: click on the **Persistent path** button () alongside the Fighter in the *Properties Timeline*.
11. Select the Camera and choose an **Automobile** type of motion in the *Animation Properties* panel. This opens the *Animation Wizard*.
12. Advance to step 4 of the Wizard and plot a path for the camera that follows (more or less) that of the fighter. In step 5, enter a total duration of the animation of 10 seconds. Finish the Wizard.
13. Play a 3D preview of the animation (press ) and edit the path of the camera so that it frames the fighter as much as possible (see the opposite screenshot for an example in *Top view*). To edit the path of the camera, click on a red way point on the camera path (when the camera is selected) and drag it to a new position. If you need to move the entire camera path, simply double-click on a way point (which selects all the way points of the path) and drag the path to another position. You can also rotate, resize, drop, flip, align... the path (or parts of the path) using the standard tools.



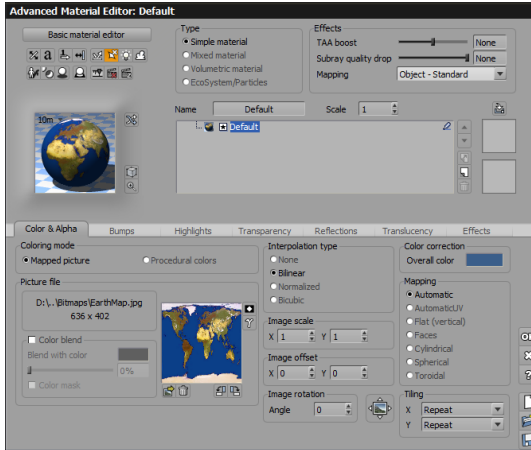
A Pursuit: step 13

The animation is complete. Press the  icon to do a preview render. The complete scene can be found in the ***Animation Tutorials*** collection.

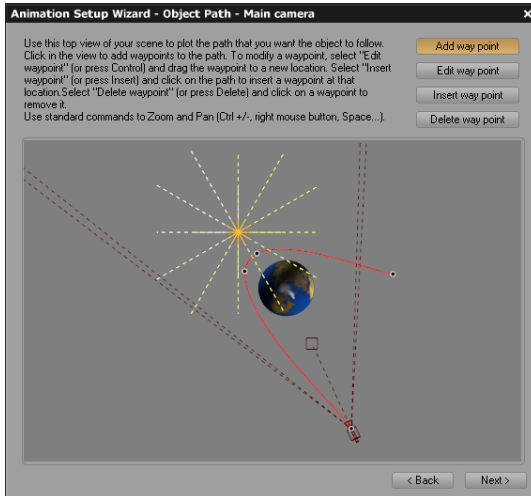


Getting a Spin

This tutorial will teach you how to create and use multiple revolutions by creating an animation of the earth as seen by a comet:



Getting a Spin: Step 4







Getting a Spin: Step 5

1. Create a new scene and select the *Outer Space* atmosphere from the *Others*



collection of the Atmospheres Browser.


2. Delete the Ground plane (in outer space, there is no ground!).
3. Now create a **Sphere**. Rename it as “Earth”.
4. Create the material map for the earth: double-click on the preview of the earth’s material in the **Aspects** tab of the *Object Properties* panel to open the *Material Editor* for the earth material. Go to the **Colors & Alpha** tab, select **Mapped picture**, and select **Load** to select a bitmap from the Bitmap Browser. In the *Bitmaps* collection, select *EarthMap.jpg*. Select **Bilinear** interpolation for improved quality. Check the **Object-Standard** box so that the material follows the object. Rename the material as “Earth” (see screenshot).
5. Now that we have setup the scene, we will proceed with animating it: select the Camera, and press the *Timeline* icon (). The *Animation Wizard* appears. In step 2, select a **Smoothed** motion type. In step 4, plot a path looking like that of a comet around the earth (see screenshot). In step 5, indicate 10 seconds for the total duration of the animation. Finish the Wizard to complete setting up the path.
6. We want the camera to always look at the earth. The easiest way to do this is to make the camera track the earth: select Earth from the *Animation Properties* **Track** drop-down list.
7. Place a **Point light** behind the earth (relative to the camera) so that the earth is slightly lit up when the camera passes behind it.
8. Now that camera motion is completed, we can give a spin to the Earth: select the Earth in *Top view* and press the **End** button () in the *Timeline* (this is to position the **Current time** at the end of the animation, i.e. 10 seconds, the total duration indicated above when setting up the camera animation).
9. Catch the  rotation handle. Now drag the mouse around the Earth to rotate it counter-clockwise. Keep going round the Earth until you have completed 4 revolutions (check the angle of rotation in the status bar). When the 4 revolutions are complete, release the mouse button. You need to complete the 4 revolutions in one go, if not the revolution counter will reset.

The animation is complete. You can do a preview render (press the  icon) to check that the Earth is now rotating rapidly as the camera flies around it. The complete scene can be found in the *Animation Tutorials* collection.






Opening a Window

In this simple tutorial you will learn how to use pivot points in animation.




Opening a Window: step 5

1. Create a new scene and load the *Noon* atmosphere from the *Physical* collection.
2. Press  and load the *Simple House* object from the *Miscellaneous/Boolean* collection.
3. Press  again and load the *Window* from the *Miscellaneous/Boolean* collection.
4. Move and resize the Window so that it fits precisely in the lower right window of the Simple House.
5. Move the camera so that it frames the window and select a **Square** Aspect ratio from the *Render Options* dialog (see opposite).
6. Create the window pivot. The pivot acts like a hinge around which the window will rotate: select the Window, go to the *Numeric Properties* panel, click the **Pivot position** icon and check the **Show pivot** button. The pivot of the Window now appears as a green square in the *3D views* when the Window is selected. Select the pivot in *Top view* and drag it to the left edge of the Window.
7. Now press the **Timeline** icon () to display the *Timeline*. If the *Animation Wizard*



appears, press **Close**.

8. Drag the **Current Time** slider to 2 seconds (the time for the window to open). Select the Window, and drag the  handle approximately 120° clockwise. This creates the animation of the window opening. Note how the Window rotates around its pivot.
9. Let's start opening the window slightly after the beginning of the animation, so that it remains closed for a fraction of a second: in the *Timeline*, drag the first keyframe a couple of frames to the right. The Window will start opening only from then.
10. Slide the **Current Time** back to 0 to close the Window. **Copy-Paste** the animated Window and nudge it up to the top right window of the Simple House using the **Nudge** keys.
11. Repeat for the windows on the left side of the house. Select these three new windows and pick **Not animated** from the *Animation Properties* panel to destroy their animation, so that only the first window opens.

The animation is complete. Press the  icon to do a preview render. The complete scene can be found in the *Animation Tutorials* collection.

Angular Paths

This tutorial demonstrates how angular motion paths can be achieved using either the *Animation Wizard* or the path editing tools available in the *3D views*. We will create a perfectly square motion.

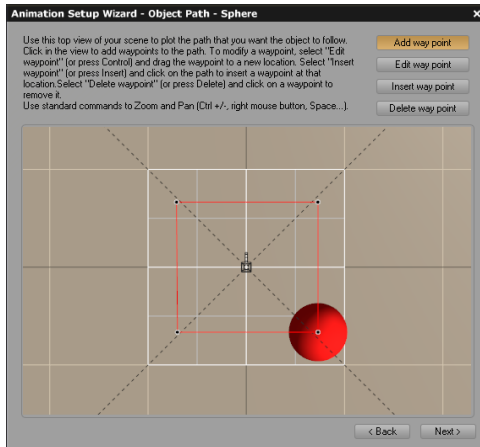
Because splines (the mathematics that underlie the creation of paths) are designed to produce the smoothest possible motion, it is difficult to create angular paths without knowing the trick that is exposed in this tutorial. The idea is to create several identical position keyframes:

1. Create a new scene and select the *Animated Arizona* atmosphere from the *Animated* collection.
2. Switch to Top camera using the menu command **Display | Activate Camera | Top Camera** to view the scene from above.
3. Create a sphere. Select **Standard** motion from the *Animation Properties* panel. This displays the *Animation Wizard*.
4. In the Object path step, create a way point upwards from the starting point. Now, without moving the mouse, click again to create a second way point precisely at the




same position as the first.

- Now move the mouse to the left and press again twice, at the same position. As you will notice, the path extends in a straight line. Complete the square using the same method, and finish the Wizard.




Angular Paths: step 5

- If you play the animation (▶) you will notice that the sphere pauses in between each straight segment. Using the *Timeline*, drag each keyframe marking the start of a segment so that it is positioned exactly at the same time as the one that marks the end of the preceding segment. This reduces the pause at each end of segment to zero. Click on the  button in the *Timeline* to open the Animation Toolbox. Enter 4 seconds as the total duration and press **OK**. This resamples the keyframes so that the animation lasts precisely 4 seconds.
- You can also create this path using the path editing tools in the *3D views*. The tools that are available make it a more precise job: restart from step 3, but this time, close the Wizard as it pops-up.
- Drag the **Current Time** slider to 1 second and **Nudge** the sphere up 15 times. This creates a new keyframe. Now move the **Current Time** a couple of frames to the right and select **Add Keyframe** from the *Timeline* popup menu. This adds a keyframe for all the animatable object properties. Don't worry about the extra keyframes, we'll get rid of them later on. The result of this is that we now have 2 keyframes exactly at the same position.
- Move the **Current Time** to 2 seconds and **Nudge** the sphere 15 times to the left. Repeat as above, then **Nudge** down and right to complete the square path.
- Get rid of all the extra keyframes that were created by the **Add Keyframes** command: drag the marquee rectangle over all unwanted keyframes and press **Delete**




to get rid of them.

11. If you play the animation, you will notice that the sphere pauses at the corners of the square. The duration of this pause is given by the space between the keyframes created by nudging the spheres and the ones created using the **Add Keyframe** command. Select the keyframes that were created using the **Add Keyframe** command, and drag them to the left so that they overlap precisely those that were created by nudging to reduce the pause to zero.
12. In the *Main camera view*, the sphere starts in the center, then moves out of the view. Let's move the sphere (with its path) so that it always remains visible: in *Top view*, select the sphere and double-click on one of its way points (the red dots). This selects all the way points of the sphere. You can now drag the complete path so that it is always visible in *Main camera view* (see screenshot).



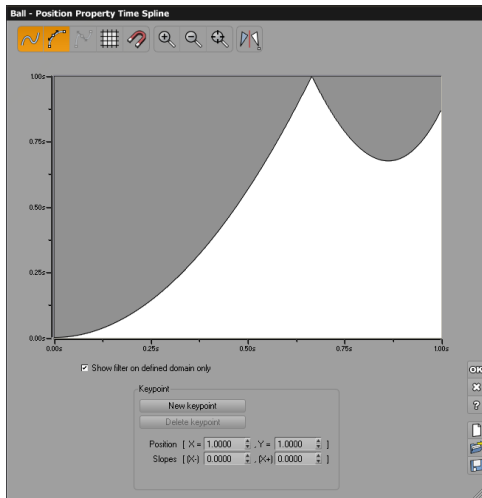
Angular Paths: step 12

The animation is complete. Press the  icon to do a preview render. The complete scene can be found in the **Animation Tutorials** collection, and the rendered animation is in the **Animations\Tutorials** folder on the CD (a red cube was added in the center of the path).





Drop and Bounce

This tutorial demonstrates the use of Time Splines to control the flow of time in an animation. We will animate a simple ball dropping and bouncing on the ground.



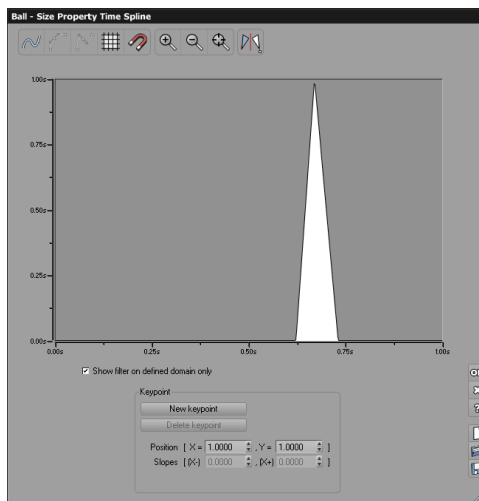
Drop and Bounce: step 9

1. Create a new scene and select the **Animated Bright Arizona** atmosphere from the **Animated** collection. This atmosphere features moving clouds.
2. Create a **Sphere**. Make it 3 times smaller using the resize handles, then drag it up to the top of the *Main camera view*. Assign it the **Radial Stripes** material from the **Basic** materials collection, and select the **Object–Standard** mapping so that the material follows the sphere as it moves. Rename the sphere as “Ball”.
3. Select a **Smoothed** motion, and **Close** the Wizard. The *Timeline* appears. We need a Smoothed motion so that the ball moves at a constant velocity along its path.
4. Drag the **Current Time** slider to ~0.5 second (this will be the time the object takes to fall). Now press the **Drop** icon () . This creates a new position keyframe. Check that when you play the animation, the ball moves down. But it doesn't look like it's dropping. Why? Because when an object falls, it keeps falling faster as it goes.
5. Time splines are the way to precisely control the speed of the Ball. Unfold the Properties Timeline (press ) and double-click on the Ball's **Position Property**



Time Spline. Load the *Simple Drop* filter from the *Time Splines* collection. This is a simple “power 2” spline that exactly captures the acceleration of motion as the Ball drops. Note how the Ball’s motion already looks more like it’s dropping.


6. We will now make the Ball bounce. Since there is going to be a bounce, let’s first increase the total duration of the animation: drag the last keyframe to 1 second.
7. To create the bounce we will revert the flow of time near the end of the animation: **Control-Click** on the Time Spline to open the *Time Spline Editor* (or select **Edit Filter** from the popup menu).
8. Drag the **Current time** slider back to approx. $\frac{2}{3}$ of a second. This will be the time when the ball first hits the ground. Notice how the black line indicating the current time follows in the *Time Spline Editor*. Check the **Smooth filter** option and create a new key point on that line, at $Y=1$. This means the Ball will be hitting the ground the first time at $\frac{2}{3}$ of the animation duration, then at the end again. The first part of the filter should keep the same continuous acceleration as before: select the key point you just created, and enter 3 (instead of 2, to compensate for the reduced time allowed for the first drop) as **Slope** to the left to create a nice increasingly steep slope.



Drop and Bounce: step 11

9. As it hits the ground, the Ball suddenly inverts it’s motion. We need to reflect this in our Time Spline by creating a sudden inversion in the flow of time: select the key point you just created, and enter -3 for the **Slope** to the right. Close the *Time Spline Editor* and press Play. The Ball bounces. The Time Spline you just created is available as *Drop and Bounce Once* in the *Time Splines* collection.
10. Now we will make the ball squash as it hits the ground. Move to the end of the



animation (press ) , and squash the Ball slightly. We have two problems: first, the Ball is no longer touching the ground in its rest position. Second, the Ball squashes gradually all along the duration of the animation instead of squashing when it hits the ground.

11. Press the **Drop** icon again so that the Ball touches the ground at the end of the animation. This fixes the first problem. We will now make the squashing occur as the Ball hits the ground; for this we will use another Time Spline for the Size Property of the Ball. Open the *Time Spline Editor* for the Ball's Size, and create a spike at 2/3, that is null anywhere else. Adjust the width of the spike by dragging the current time slider to the position where the unsquashed Ball first hits the ground (see screenshot of the Time Spline Editor). Press **Play** to preview the animation.
12. There is one last problem: when the animation completes, the Ball is partly under the ground (because it is no longer squashed). To correct this flaw, we need to stop the Ball before it goes back down all the way after the bounce. Open the Editor for the Position Time Spline again, and drag the last key point down slightly until the Ball rests on the ground in its final position.

The animation is complete. Press the  icon to do a preview render. The complete scene can be found in the *Animation Tutorials* collection.

Flickering Lights

This tutorial will teach you how to animate the color and intensity of lights.






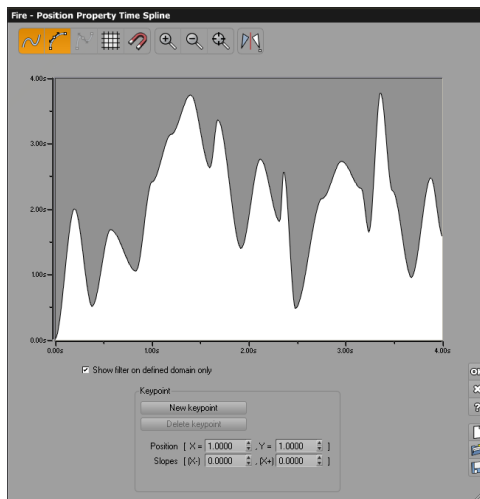
Flickering Lights: step 3

1. Create a new scene and select the *Blue Night* atmosphere from the *Effects/Others*






collection.

2. Press  and load the **Fortress** object from the *Miscellaneous/Boolean* collection.
3. Move and rotate the Fortress to achieve a tight framing of the entrance of the Fortress (see opposite).
4. Create a **Point light** and position it inside the Outer fortifications, just behind the entrance. This is going to be a fire, so it should lay close to the ground: press the **Drop** icon () . Rename the light as “Fire”.
5. Double-click on the light color in the *Aspect Properties*. Select a bright yellow color and press **OK**.
6. Now press the Timeline icon () to display the *Timeline*. If the *Animation Wizard* appears, press **Close**.
7. Drag the **Current Time** slider to 2 seconds and edit the light color again, this time giving it a more orange shade. Notice that a new keyframe has been added to the **Color Property** of the light.
8. Drag the **Current time** slider to 4 seconds, and edit the light color to give it a dark red color.
9. The color of our fire animates from yellow to orange to red. Doesn't look much like a fire... We will now use a Time Spline to create lots of sudden variations in the color of the light (the flickering): open the *Time Spline Editor* for the Color Property by **Control-Clicking** on the Time Spline.



Flickering Lights: step 10




10. Create a seemingly random Time Spline with many sudden variations (see screenshot). Make sure that the end position is identical to the first, so that the flickering will loop. The Time Spline used in this tutorial can be found as **Random** in the **Time Splines** collection. This will create the flickering of the fire.
11. Rewind the Current time by pressing **Start of animation** . **Copy-Paste** the light and drag it inside the Dungeon, near the top. This will also put a fire in the Dungeon.
12. If you render the animation, you will notice that the two fires flicker exactly at the same pace. Not very realistic... We will break that impression by simply dephasing the animations of the two fires: select the Dungeon fire, and select all its keyframes in the Timeline. Now drag the keyframes to the left 1 second. This introduces a 1 second dephasing in the intensities of the two fires.
13. When the animation passes the end of the Dungeon fire's animation, that light's color stops changing. Since the Color Time Spline was designed to loop, we'll make the animation of the light repeat endlessly. Press the  button alongside the Dungeon fire to open the *Animation Toolbox*, and select **Repeat** as **Repeat mode**. The two fires are now flickering endlessly, with a neat dephasing of one second.
14. Select the first fire again, and go to the end of the animation (press ). Now move the light slightly and create another pseudo-random Time Spline to give the impression that the fire is moving (this will be visible on the ray of light cast through the entrance).

The animation is complete. Press the  icon to do a preview render. The complete scene can be found in the **Animation Tutorials** collection.


Animating a Fish

In this tutorial you will learn how to build a hierarchy of objects and animate it using repeat modes.





Animating a Fish: step 4




Animating a Fish: step 9






Animating a Fish: step 11






Animating a Fish: step 21

1. Create a new scene and load the *Underwater* atmosphere from the *Effects/Others* collection.
2. Import the Angel fish model, supplied as a Wavefront OBJ object on the CD: select the menu command **File | Import Object** and pick the *Angelfish.obj* file from the *Objects* folder in the Extra Content files.
3. Double-click on the material and select **Mapped picture** in the **Colors** tab. Press **Load** and load the *Angelfish.jpg* picture from the *Bitmaps* folder. Because the model of the fish has UV mapping information, the picture will map precisely (and automatically) on the fish. Select **Bilinear** interpolation for the sake of quality.
4. Position the camera to frame the fish as shown in the opposite screenshot.
5. Drag the Angelfish into Layer 2 and **Ungroup** it. This is to enable us to define new linking options for the linking hierarchy. By dragging the fish into a separate layer, you make its selection easier once it is ungrouped. Rename the object called body1: 2 as “Angelfish”. This will be the master object to which all other parts of the fish are linked.
6. Select the body part body2: 2 and go to the **Numerics** tab, **Pivot position** sub-tab. Check the **Show pivot** button, and then, in *Top view*, drag the pivot handle so that it is positioned where body2: 2 joins up with the Angelfish object. Now link body2: 2 to the Angelfish object by pressing  in the **Animation** tab and selecting the *Angelfish* object.



7. Repeat for the other body parts, each time positioning the pivot on the joint with the previous segment, and linking to it (i.e. body3: 2 links to body2: 2, body4: 2 to body3: 2 and tail: 2 to body4: 2). Link the fins and the head to the Angelfish object. The hierarchy is complete. You have to be particularly careful when defining the pivot points, because if the pivot is not placed correctly, you may experience “cracks” in the body of the fish when animating it.
8. We will now make the fish swim by moving its tail. Select body2: 2 in *Top view* and rotate it a couple of degrees clockwise. As expected, the whole tail follows the body segment around its pivot. As this is the first segment of the body, it should be stiff, so keep the rotation angle low.
9. Keep progressing down the body, increasing very slightly the angle of rotation as the body gets thinner. You should end up with a nicely curved body, as shown in the opposite screenshot.
10. Press the **Timeline** icon () to display the *Timeline*. If the *Animation Wizard* appears, press **Close**.
11. Drag the **Current Time** slider to one second and rotate the body segments as above, only clockwise this time. You should end up with the fish’s body being smoothly curved the other way.
12. If you preview the animation, you will see that the body of the fish gently curves from one side to the other. We can now make the fish move.
13. Drag the **Current Time** slider to 10 seconds and move the Angelfish object forward so that it passes out of the right hand side of the *Main camera view*. As expected, the whole fish follows the Angelfish object (lucky fish!).
14. Unfortunately, the body of the fish only curves once. Let’s use a special Repeat mode called Pendular to improve this: pendular repetition means that the animation is repeated the other way round back to the start before restarting from the beginning again. Just what we need to make the tail curve back and forth! Select all the body parts that are linked to the Angelfish object and press the  button in the *Animation Properties* panel. The *Animation Toolbox* pops-up for all the selected objects. Select **Pendular** Repeat mode and press **OK**. If you play the animation, you will notice that the fish wags its body as it moves forward.
15. Let’s do the same for the fins: drag the **Current Time** slider down to 1/3 of a second (we’ll make the fins move faster) and rotate the fins in *Top* and *Front views*. Since Pendular Repeat mode has already been defined for the fins, they will keep wagging all along. This completes the animation of our first fish.
16. Revert to the start of the animation (press ). Click on the Layer 2 title in the *World Browser* to select all the parts of the fish, and **Copy-Paste** it into Layer 3.




17. Select the second Angelfish object and drag it behind the first so that it hardly shows on the left of the *Main camera view*. Go to the end of the animation (press ) and drag the Angelfish so that it moves parallel to the first fish.
18. If you preview the animation now, you will notice that both fish move their bodies exactly at the same time, and at the same pace. That's not very realistic... So we will dephase the two fish: select all the body segments of the second Angelfish and select all the corresponding keyframes in the *Timeline*. Drag them to the left approximately one half of a second.
19. Now select the last keyframe (the one corresponding to the body) and drag it to the left a couple of frames to make the body wag faster.
20. Repeat for the fins. The two fish are now out of phase. And the second moves faster than the first.
21. Obviously, placing the fish in a more “usual” media improves realism. Create a **Water** plane and drag it just above the fish. Place the plane in the *World Browser* just below the Ground. Load the *Underwater* material from the *Liquids* collection (press ) . Load the *Underwater* atmosphere from the *Effect/Others* collection (press the  button).

The animation is complete. Press the  icon to do a preview render. The complete scene can be found in the *Animation Tutorials* collection.


A Steam Power Train


In this tutorial you will learn how to build a complex “forward dynamics” hierarchy using different linking options. The focus will not be on achieving a detailed model of a steam power train, but on building a hierarchy of objects that closely mocks up the way the train operates.

1. Create a new scene and load the *Default* atmosphere from the *Spectral Sunshine* collection.
2. Create a Cylinder, and rotate it so that it rolls on the ground: using the *Numeric Properties* panel, **Rotation** sub-tab, enter 90° as **Yaw**. Flatten the cylinder in *Front view* to the dimensions of a train wheel and **Drop** it () so that it rests on the ground. Rename this as “Primary wheel” and assign it the *Radial Stripes* material from the *Basic* collection and map it as **Object-Standard**.
3. Create another Cylinder, and make it much smaller. This will be the pivot that




connects the wheel to the power rod. Move it up so that it is close to the edge of the Primary wheel. Rename it as “Pivot” and assign it the *Dirty Metal* material from the *Metals* collection and map it as **Object-Standard**.

4. Link the Pivot to the Primary wheel using the *Animation Properties* panel. We will be linking rods to the Pivot later on.
5. Select the Primary wheel and the Pivot, and make a copy of them. Select the new Primary wheel and rename it as “Secondary wheel”. Offset the secondary wheel to the front of the Primary wheel (the Pivot will follow).
6. Press the **Timeline** icon () to display the *Timeline*. If the *Animation Wizard* appears, press **Close**.
7. Drag the **Current Time** slider to 2 seconds, select the Primary wheel and rotate it one complete revolution. You can check that the Pivot follows the wheel.




A Steam Power Train: step 9

8. Now link the Secondary wheel to the Primary wheel, and uncheck the **Join** box so that the Secondary wheel rotates around its own center. If you play the animation, you will see that both wheels rotate together although no animation has been defined for the Secondary wheel.
9. Revert to the beginning of the animation by pressing . Create a **Cube** and resize and stretch it into a long rod that you will position in such a way that it links the two Pivots. Rename this as “Connecting rod”.
10. Link the Connecting rod to the Pivot. Uncheck the **Rotation** box so that the rod follows the Pivot, but doesn’t rotate with it.
11. Copy-Paste the Connecting rod and **Nudge** it to the left in *Side view* so that it has its end on the Primary wheel’s Pivot. Destroy that object’s animation by selecting **Not animated** in the *Animation Properties* panel. Rename it as “Power rod”.
12. Now position the Pivot point of the Power rod at the center of the Primary wheel’s Pivot cylinder.



13. Create another rod and position it at the same height than the center of the wheels. Rename this rod as “Cylinder rod”. This rod will be transferring power from the steam cylinder to the Power rod. Link it to the Power rod and uncheck the **Rotation** box.



A Steam Power Train: step 17



14. To create the cylinder, load the **Tube** object from the **Boolean** collection. Assign it the **Dirty Metal** material and rename it as “Cylinder”. Resize it and position it at the left end of the Cylinder rod, so that the rod seems to come out of the Cylinder.
15. Select the Power rod and rotate it so that its free end corresponds with the end of the Cylinder rod.
16. Drag the **Current Time** slider to 1 second (the time necessary for the wheel to go half a turn) and rotate it back up so that its end is in contact with the Cylinder rod again. Select a **Pendular** Repeat mode for that rod using the *Animation Toolbox*.
17. If you play the animation, you will notice that the Cylinder rod is moving up and down when it should be staying at the same height (this is because it follows the pivot of the Power rod). We will correct this by animation. Move the **Current Time** slider to 1 second and **Nudge** the Cylinder rod back up to its initial position. Make it also a **Pendular** Repeat mode.

The animation is complete. Press the  icon to do a preview render. The complete scene can be found in the **Animation Tutorials** collection.



Making Waves




In this tutorial you will learn how to animate the surface of the water using time dependent materials. Obviously, the most straightforward way of achieving this effect would be to use the *Water Surface Editor*, however, the technique described in this tutorial can be useful in many different contexts:

1. Create a new scene and load the *Sunny Seaside* atmosphere from the *Day-time/Spectral Sunshine* collection.
2. Add **Water** to the scene by pressing the  icon. Load the *Wavy Water* material from the *Liquids* collection (press ). Now select the Ground and **Delete** it.
3. Import the model of a statue of Cupidon, supplied as a Wavefront OBJ object in the Extra Content files: select the menu command **File | Import Object** and pick the *Cupidon.obj* model from the *Objects* folder on the CD. This will be used to emphasize movement of the waves as the statue reflects in the water.
4. Assign the *White Veins* material from the *Rocks/Miscellaneous* collection to the Cupidon statue (press).
5. Open the Render Options dialog and select a **Photo – vertical** aspect ratio.
6. Using the camera, frame the statue so that some water is visible in the bottom part of the picture (see opposite).



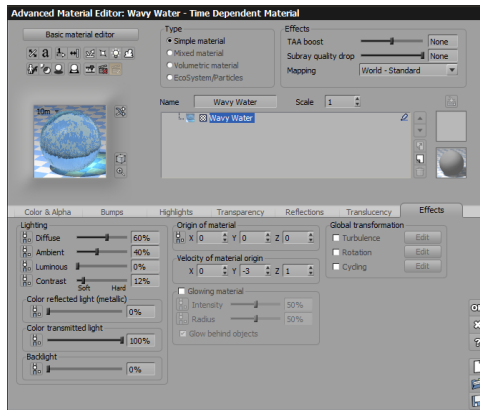


Making Waves: step 6

7. Select the Water plane and double-click on the Preview of the Water material to open the *Material Editor*.
8. Check the **Animate material surface** box to create a surface animation of the material (see [Animating Materials](#)).
9. Press the **Timeline** icon () to display the *Timeline*. If the *Animation Wizard* appears, press **Close**.
10. Notice that the Water material has been automatically added to the animated items of the scene. However, **Material Surface animation** and **Material Velocity animation** (see below) do not create any keyframes. This means that there are currently no keyframes defined in the scene. So you cannot preview it (because there is no Preview duration available). To define a Preview duration, click on the **Render animation** icon () and select **Render sequence** in the *Render Animation Options* dialog. Enter 10 seconds in the **End time** of the animation, and press **OK**. An animation duration (i.e. 10 seconds) is now available in the Timeline, and you can do a preview render (press )
11. If you render the animation, you will notice that the surface of the water is moving. Material Surface animation is the simplest type of material animation.
12. We will now try to add some sort of flow to the water. This effect will be achieved by “sliding” the material over the surface of the water plane. This is called **Material**



Velocity animation. Open the water *Material Editor* again, and go to the **Effects** tab. In the **Velocity of material origin** fields, enter -3 for **Y**. This means that the origin of the material will be moving along Y at the rate of 3 material units per second. It is also moving along Z at one unit per second because the material has a Surface animation. The movement we are introducing along Y will give the impression that the surface of the water is sliding from the left in the direction of the statue.



Making Waves: step 12


- Go to the **Bumps** tab of the water *Material Editor*, and edit the Bump production function (Control-click on the preview). As you will notice, there are two noise nodes defined for the water. The first (Noise) creates the small waves on the surface, while the second (Wave) creates a lower frequency wave front. The shape of these lower frequency waves should change very slowly over time, so that we can see them clearly progress towards the statue. To do that, increase the scale over the **Z** axis of the Wave layer to 20 (this means the shape of the wave will evolve more slowly along that axis).

The animation is complete. Press the  icon to do a preview render. The complete scene can be found in the **Animation Tutorials** collection, and the rendered animation is in the **Animations\Tutorials** folder on the CD.




Animating Clouds

In this tutorial you will learn how to animate clouds. This is similar to animating water, but will introduce the concept of Complete Material animation.






Animating Clouds: step 2


1. Create a new scene and load the *Brisbee* atmosphere from the *Daytime/Spectral Sunshine* collection.
2. In *Render Options*, select the **Photo – vertical aspect** ratio and click **OK**. Then close the dialog and turn the camera up to frame the sky (see opposite).
3. Open the *Atmosphere Editor* (press ) and go to the **Clouds** tab.
4. There are two layers of clouds in this atmosphere. Open the *Material Editor* for the first layer of clouds (double-click on the preview), go to the **Effects** tab and enter a **Velocity** of (1; 1; 0.5). This will create a **Material Velocity animation**, resulting in a movement of the clouds from right to left, with a slow evolution in



shape as time passes (read the previous tutorial for a better understanding of these settings).

5. Press the **Timeline** icon () to display the *Timeline*. If the *Animation Wizard* appears, press **Close**.
6. Notice that a folder entitled *Atmosphere* has been automatically added to the animated items of the scene, and that the *Cloud* material is placed in that folder. However, the *Cloud* material doesn't have any keyframes, because it is a **Material Velocity animation**. This means that there are currently no keyframes defined in the scene. So you cannot preview it (because there is no Preview duration available).

To define a Preview duration, click on the **Render animation** icon () and select **Render sequence** in the *Render Animation Options* dialog. Enter 10 seconds in the **End time** of the animation, and press **Close**. An animation duration (i.e. 10 seconds) is now available in the *Timeline*, and you can do a preview render (press )



7. Return to the *Atmosphere Editor*, **Clouds** tab, and select the second layer of Clouds. Click on the cloud preview to go to the *Material Editor*. Select the **Lighting & Effects** tab and enter a **Velocity** of (0; -1; 1). This creates a cloud movement in the Y axis direction. These clouds also evolve twice as fast as the first layer.
8. We will now make the clouds “dissolve” as time passes. This is achieved using a **Complete Material** animation: go to the end of the animation (press ) and reopen the *Cloud Material Editor*. Go to the **Transparency** tab, and double-click on the **Transparency production** Filter. Load the **Flatten 50%** from the **Other Filters** collection. A message appears asking if you want to animate the material. Click **Yes**. As you will notice, a keyframe has been added to the cloud material in the *Timeline*. Due to the new transparency filter, the clouds are now 50% thinner at the end of the animation than at the beginning.
9. Control-click on the **Transparency production** Filter to open the *Filter Editor*. Drag the first key point up to 80%. This will reduce the density of the clouds by 80% at the end of the animation, making the effect more obvious.
10. . Repeat operations 8 and 9 for the other layer of clouds.


The animation is complete. Press the  icon to do a preview render. The complete scene can be found in the *Animation Tutorials* collection.



Moving Caustics

This tutorial will show you how to create moving caustics. Caustics are the patterns of light that you observe on the bottom of swimming pools, or near the sea shore. They are caused by the waves on the surface of the water. The waves concentrate or spread out the light that hits the surface, depending on the concavity of the surface at that point. Caustics can be efficiently faked using the technique described below. They will add a great deal of realism to your underwater scenery.

1. Load the scene from the *Animating a Fish* tutorial.
2. Go to the start of the animation by pressing .
3. Create a **Plane** primitive. Rename it as “Caustic mask” and assign it the **Flat Black** material from the *Basic* collection (press ).
4. Drag it up right below the sunlight and resize it until it masks all of the sunlight visible in the scene (use test renders to decide). Now we will add transparent patterns to the mask.



Moving Caustics: step 9

5. Double-click on the material of the Caustic mask to open the *Material Editor*, and go to the **Transparency** tab. Push the **Global transparency** up to 100% and check the **Variable transparency** box. This will let us define transparent areas in the material.
6. Control-click on the **Transparency production** function to open the *Function Graph*, and create a **Perlin | Value** noise node.
7. Add a filter and load the *Caustic Patterns* filter from the *Other Filters* collec-



tion. This filter will create the actual patterns from the basic noise.

8. To add higher frequencies to the noise we've just created, replace the noise node with a **Fractal** noise, **Basic Repeater** type. Reduce the **Complexity** to 1 to avoid having too many high frequencies in the noise. Close the *Function Graph*.
9. Render the scene to check the caustics. Bump up the **Scale** of the material until you get a realistic size for the caustic patterns (see illustration).
10. Open the *Atmosphere Editor* and go to the **Light** tab. Push the **Global exposure** up by $\frac{1}{4}$ of a diaphragm to compensate for the loss of light caused by the Caustic mask.
11. Now we are ready to animate the caustic patterns. Open the caustic *Material Editor* again, and go to the **Effects** tab. Enter a **Velocity** of (0; 0; 10). This will cause the caustic patterns to change rapidly with time. Cool animations guaranteed!

The animation is complete. Press the  icon to render a preview. The complete scene is in the *Animation Tutorials* collection.

Stoned Frog

In this tutorial you will learn how to use complete material animation (i.e. animated blending of materials). The animation shows a frog that is turned to stone by a flash of light.



Stoned Frog: step 3









Stoned Frog: step 6




Stoned Frog: step 9

1. Create a new scene.
2. Press  and load the **Frog** object from the **Animals/Terrestrial** collection.
3. Move the camera so that it frames the Frog and select a **Square** Aspect ratio from the *Render Options* dialog (see opposite).
4. Create a **Point light** and place it just above the Frog. Turn the light off by reducing its **Power** to 0 (with the *Object Properties* panel).
5. Press the **Timeline** icon () to display the *Timeline*. If the *Animation Wizard* appears, press **Close**.
6. Drag the **Current Time** slider to frame #6 and increase the power of the light to 100. A keyframe is automatically added to the light.



7. Drag the **Current Time** slider to frame #3 of second 1 and turn the light off again by reducing its power to 0. This creates a flash of light that grows more rapidly than it goes out.
8. Select all the keyframes of the Light using by dragging the marquee selection rectangle, and drag all the keyframes up almost one second. This means we can admire the Frog in the final animation for almost one second before the flash of light occurs.
9. Now we turn the Frog to stone; drag the **Current Time** slider to 3 seconds and open the *Summary of Materials* panel (press ). This panel displays all the materials used by the different objects of the scene. Double-click on the Frog skin material to open the *Material Editor*. The Frog skin material is a mix of two procedural materials. Press the **Load** button and load the *Gray Clumps* material from the *Rocks* collection. A message appears asking if you want to animate that material. Click **Yes**. Notice how a new folder has been added to the animated items of the *Timeline*, and the Frog skin material placed in it. Note that if you press the  button in the *Summary of Materials*, the material will be replaced by the one you load instead of becoming animated.
10. Repeat for the Frog's eyes material.
11. We want the Frog to start turning to stone after the flash. Grab the first keyframes of the Frog skin and eyes material, and drag them up to a few frames before the end of the flash. The frog will start transforming from then.
12. We want the animation to continue a bit after the Frog has completely changed to stone. Drag the **End of Animation** slider up to 4 seconds, which leaves one complete second of stone frog.


The animation is complete. Press the  icon to do a preview render. The complete scene can be found in the *Animation Tutorials* collection. Select **Render sequence** from the **Animation Options** dialog to render the complete animation (if not, rendering stops after the last keyframe instead of going up to the **End of Animation** slider).



Animating Groups and Boolean Objects

In this simple tutorial you will learn how to animate Groups and Boolean Objects. The tutorial details the construction of an animated Boolean Difference, but the same technique applies to other Boolean operations, as well as to Groups.


1. Create a new scene and select the *Animated Bright Arizona* atmosphere from the *Effects/Animated* collection. This atmosphere features moving clouds.
2. Create a Torus and double-click on it to open the Torus Options dialog. Raise the **Torus thickness** to 0.40 and press **OK**.





Animating Booleans: step 4

3. In *Top view*, squash the Torus vertically using the small resize handles.
4. Move and rotate the camera to achieve the framing shown in the opposite screenshot.
5. Create a Cylinder and rotate it by entering a **Pitch** of 90° in the **Rotation** sub-tab of the **Numeric Properties**. Stretch it vertically and resize it to make it approximately half the diameter of the Torus ring. **Nudge** the Cylinder to the left 6 times so that the Cylinder is placed in the middle of the Torus ring. Rename the cylinder as “Hole”.
6. Now we will replicate the Hole so that there is one hole every 45° in the Torus. Go to the **Pivot position** sub-tab of the **Numeric Properties**, and check the **Show pivot** box. Enter 0, 0, 50 (the exact center of the Torus) as coordinates for the Pivot point.





Animating Booleans: step 8

7. Now **Copy-Paste** the Hole, and enter a Yaw angle of 45° using the **Rotation** sub-tab of the **Numeric Properties**. Keep Copy-Pasting and rotating the holes by increments of 45° until there are holes all around the Torus. You could also use the *Scatter/Replicate Objects* dialog.
8. Now we do the actual “Punching-out” of the holes in the Torus: select the Torus in the *World Browser* and Shift-click on the last Hole to select all objects in between (the Torus plus all the Holes). Click the **Boolean Difference** icon () to create the Boolean operation (see opposite).
9. Select all the Holes using the *World Browser*, and, in *Top view*, nudge the holes down three times, so that they don’t overlap the Torus any longer.
10. Press the **Timeline** icon () to display the *Timeline*. If the *Animation Wizard* appears, press **Close**.
11. Drag the **Current Time** slider to 1 second and select the first Hole in the Boolean Difference (using the *World Browser*). In *Top view*, nudge the Hole up twice so that it completely overlaps the Torus.
12. Drag the **Current Time** slider up by 5 frames (one third of a second) and select the second Hole. **Nudge** it up the same. Repeat for all the other Holes, incrementing the **Current Time** by 5 frames each time. As expected, running the animation, will show that each hole is punched out in turn.
13. We want each “punching-out” to last the same amount of time. So we need to delay the beginning of the movement of the Holes. Select each Hole in turn, and drag the first keyframe in the *Main Timeline* so that the duration of the movement is exactly



one second.

The animation is complete. Press the  icon to do a preview render. The complete scene can be found in the *Animation Tutorials* collection.


Using Spin and Vibrate Effects

In this short tutorial we will build an outer-space animation of a Soyuz spaceship rotating around the earth. As the Soyuz ship flies by the camera, the camera will vibrate strongly.

To build this animation, we will make extensive use of the Spin and Vibrate effects.


1. Create a new scene based on the *Orbit* atmosphere from the *Effects/Others* collection.
2. Select the ground plane in the *World Browser* and delete it.
3. Open the *Atmosphere Editor* and enter 0 as Fog and Haze densities. Drag the **Twinkle** slider in the **Effects** tab to 30% to make the stars twinkle slightly.
4. Drag the sun down so that it is level with the camera.
5. Create a very large sphere and rename it as “Earth”. Map it with the *EarthMap.jpg* and make it **Object Space** so it follows the sphere as it spins.
6. Create another sphere, centered on the Earth, and make it just a tad larger. This will be the Earth’s atmosphere. Rename it as “Atmosphere” and map it with a suitable cloud material. Make the material **Object Space** so it follows the sphere as it spins.
7. Select the Earth, and from the **Animation** tab of the *Object Properties* panel, select “Standard” animation. Close the *Animation Wizard*.
8. Open the *Animation Toolbox* and check the **Spin** box. Enter a spin speed of 1.5 degrees per second around **Z axis**. This will make the Earth spin slowly.
9. Now select the Atmosphere and **Link** it to the Earth. This will make the atmosphere spin together with the earth. You could add a slight spin to the atmosphere itself to make it slide very slowly relative to the Earth (very small values here).
10. Load the *Soyuz* ship from the *Vehicles* object collection. Add a red light on the tip of the ship and drag it inside the Soyuz group.





Step 11 – Placing the Soyuz

11. Using the *Top view*, drag the ship so it is not far from the surface of the Earth, pointing down.
12. Create a sphere a tiny at the center of the earth. This will be the Ship's pivot. Link the Soyuz to the pivot, and make the pivot Spin around its Z axis at 5.5 degrees per second.
13. Drag the Time slider in the *Timeline* and check that the Soyuz rotates around the Earth.
14. Now apply Spin to the Soyuz itself, at a speed of 180 degrees per second around the ship's main axis (X). Check that the Soyuz spins on itself as it spins around the Earth.
15. Animate the camera so it moves slowly for 20 seconds on its orbit, traveling in the opposite direction as the Soyuz. Give the camera a **Look ahead** property using the *Animation Toolbox*.



Step 16 – Rotate the Camera

16. Check that the Soyuz passes by real close to the camera (change it's orbit at frame

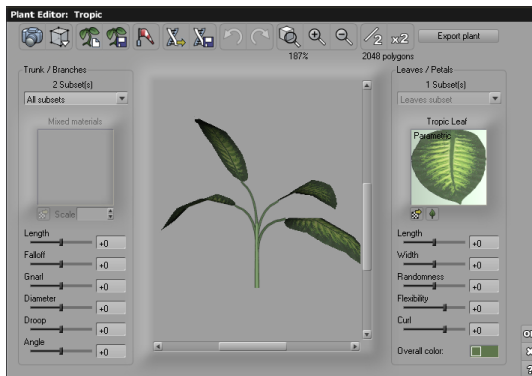


0 if adjustment is required).

17. Rotate the camera slightly at frame 0 so it is pointing at the Earth's horizon (see screenshot).
18. Now we will make the camera vibrate as the Soyuz passes by. Open the *Animation Toolbox* for the camera and activate **Vibrate** along all axes.
19. Currently, the camera vibrates all through the animation. We want it to vibrate only when the Soyuz passes by. This will be achieved using the **Variation of intensity** filter. Move the Time slider and make a note of the time when the Soyuz passes by the camera (should be at approx. 16 seconds).
20. In the Animation Toolbox, edit the **Variation of intensity** filter. Make the output value zero all along, except for a bulge at position $16/20 = 0.8$ (see screenshot). This will ensure that the camera only vibrates just after the Soyuz passes it by.
21. Animate the sun so it moves slowly from behind the Earth to the top left corner of the screen. We added a second directional light pointing in the opposite direction to the sun to brighten up the scene slightly.


The animation is complete. Press the  icon to do a preview render. The complete scene can be found in the *Animation Tutorials* collection.

Dying Plants

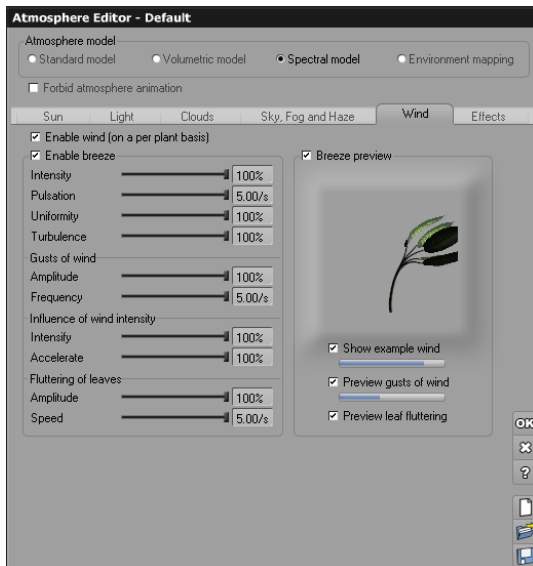


Dying Plant: step 2





Dying Plant: step 6





Dying Plant: step 8


Here you will learn how to animate the geometry of a plant using the *Plant Editor*. We will build a small animation of a dying plant.

1. Create a new scene based on the **Default** atmosphere in the **Daytime/Spectral Sunshine** collection. The focus of this animation is only on the plant, so we just want an atmosphere that lets us see the plant well.
2. The plant that we are going to animate will be a simple plant derived from the **Tropic**, with a longer trunk. Create a **Tropic** plant and double-click on it to open the *Plant Editor*. In the trunk and branches subset, select the trunk subset, and



increase the **Length** setting to +30 (only for the trunk, thus making a long stalk).


3. Close the *Plant Editor*, and use the camera controls to frame the plant up close.
4. Click on the **Timeline** icon () to display the *Timeline* (close the *Animation Wizard* if it appears). Drag the Current Time slider to 8 seconds. This will be the time that the plant takes to die (pretty much instant death, huh!).
5. Now double-click on the plant to open the *Plant Editor* at this new time. We will modify the shape of the plant so that it looks like it is almost dead, and create a new animation keyframe. Increase **Droop** to +25 and **Angle** to +5, reduce **Diameter** to -20 to make the trunk thinner.
6. Reduce the size of the leaves by entering -30 as leaf **Width** and **Length**, increase **Curl** to +100. Double-click on the **Overall color** setting, and select a pale brown color. Edit the trunk and stem materials to make them brown. Click on the **Render Preview** icon () in the *Plant Editor* to preview the look of the dead plant.
7. Click **OK** to close the *Plant Editor*. A message appears asking if you want to animate the plant geometry. Click **Yes** to animate the geometry of the plant. If you play the animation, you can see the shape evolve over the 8 seconds of the animation.
8. What we will now do is create a strong shaking of the plant using breeze. With the Current Time slider at 8 seconds, open the *Atmosphere Editor* and raise all the breeze settings to maximum.
Click **OK** to animate the breeze settings in the atmosphere, from very mild breeze up to this very strong breeze we have just created.
9. Now go to the **Light** tab and select a dark gray color for diffuse and ambient lighting. This will create a storm-like lighting to emphasize on the strong breeze.
10. We want the sudden gust of breeze to appear around 2 seconds and last for a couple of seconds, the death of the plant following this sudden shaking. In the *Timeline*, open the Time Spline for the atmosphere, and set the output to 0 all along, except between 2 and 4 seconds, where the output will be 1. This means that the breeze settings will be mild all along, except in between the second and third second in the animation where it will be very strong.
11. Browse to the first keyframe in the animation of the plant geometry, and drag that keyframe to 3 seconds, around the middle of the strong shaking. This way, the “death” of the plant will take place after the shaking and last 5 seconds. You can drag the end of animation marker up to 10 seconds to show the dead plant for a couple of seconds longer.

The animation is complete. Press the  icon to do a preview render.





Animating an Ocean

In this tutorial, we will see how to create a seemingly infinite open ocean surface.




Animating an Ocean: step 10

1. Create a new scene and select the **Animated Bright Arizona** atmosphere from the **Animated** collection. This atmosphere features moving clouds.
2. Add a water plane by clicking the **Water** icon (). This creates our ocean.
3. With the water plane selected, click the **Edit** icon () on the top toolbar, or by select **Edit Object** from the **Objects** menu to open the *Water Surface Editor*.
4. Convert the waves at the surface of water into true geometry by ticking the **Displaced water surface** option. This converts the water surface into an infinite procedural terrain.
5. Move the global wave **Overall agitated** slider up to 70% to create a more strongly agitated surface.
6. Press the **Edit function** button to open the *Function Graph* on the procedural terrain altitude function. Notice how an **Open Ocean** node is used to create the base wave pattern. Turn [here](#) for full details on the **Open Ocean** node. Close the *Function Graph*.
7. Now we will add a small island in the middle of the ocean. Create a terrain (pro-



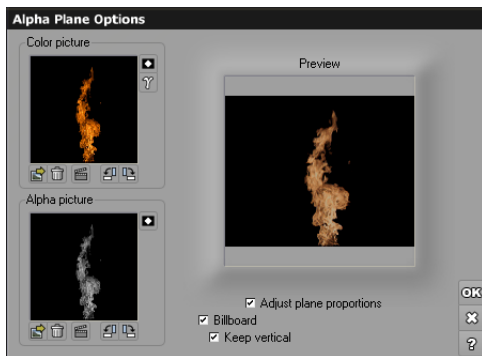
cedural or standard) and place it below the surface of the water so that the top of the terrain protrudes from the water and creates the island. Lower (or delete) the ground plane so that the ocean looks very deep.

8. Assign a rock or beach material to the island and edit the material. Change the material type to **EcoSystem**, load the *Coconut tree* into the EcoSystem population list and populate the island to add a little life to it.
9. Move the sun so that it faces the camera to emphasize the details at the surface of the water by creating interesting reflections.
10. Click on the **Timeline** icon () to display the *Timeline* (close the *Animation Wizard* if it appears) and select **Set Animation End** from the *Timeline* popup menu. Drag the animation end time to set a duration for the animation.


The animation is complete. Press the  icon to do a preview render. The complete scene can be found in the *Animation Tutorials* collection.

Using Animated Billboards to Create a Fire


In this tutorial, we will see how billboards can be used in conjunction with EcoSystems to create simple yet interesting fire effects.





Fire tutorial: step 3

1. Create a new scene and select the *Blue Night* atmosphere from the *Effects/Others* collection.
2. Create a new billboard by clicking the **Alpha Plane** icon (). Load the *Flame* animation from the *Animations* collection into the color and alpha channels. Invert



the Alpha channel by clicking the corresponding  button.

3. Select the **Billboard** option, and the **Keep vertical** option so that the flame always stays vertical. Close the editor.
4. Rename the billboard as “Flame” and edit the flame’s material. Disable **Cast shadows** and **Receive shadows**, and, in the **Effects** tab, make the material at least 100% luminous, with no ambient or diffuse lighting.
5. Save the flame as a **VOB** by pressing the **Save object** icon (). This is required in order to use the flame in an EcoSystem. Delete the Flame object.
6. Create a fireplace (e.g. a plane) and drop it to the ground. Create a point light and place it above the fireplace. Give it a yellow color, with a very strong power (approximately 500). This will create the glow around the fireplace.




Fire tutorial: step 7

7. Edit the fireplace material and switch the material to **EcoSystem**. In the EcoSystem population list, load the **Flame.vob**. Populate the fireplace with flames (adjust the flame size and density as required – probably even more than 100% density).
8. If you render the animation, you will notice that all the flames are in perfect sync. Not really convincing. We need to de-synchronize the flames throughout the animation.
9. In the *World Browser*, locate the Flame material in the EcoSystem materials cate-



gory. Double-click on it to open the *Material Editor*.


10. In the **Colors** tab, right-click (Ctrl-click on Mac) on the flame picture to edit the color production function.
11. In the “Animation map” node options, extract the **Phase** parameter by clicking the corresponding  icon. Replace the extracted constant by a “Value Perlin” noise with a very small scale to create pseudo-randomness. Warning: doing so means that the frame number of the flame animation changes at every point, resulting in very high memory requirements.
12. We need to apply the changes in phase to the transparency channel too. We could duplicate the graph, but the best here is to connect transparency to the color texture: delete the alpha animation node and connect the Transparency output to the color animation node. The picture needs to be inverted to drive alpha correctly: select the link between the animation node and the Transparency output and click the Filter icon. Select the **Opposite** filter type to invert the grayscale colors in the animated texture map. This will put colors and alpha back into phase.



Fire tutorial: dephased flames

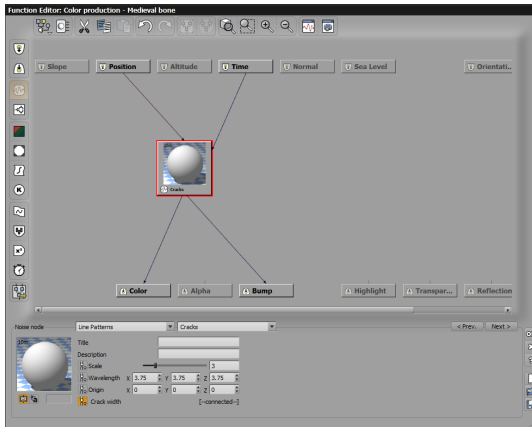
13. At this point, the phase parameter changes all the time over the surface of each flame. We need to make it different for all flames, but constant over each individual flame. For this, we will drive the noise node with the “Object center” input. Right-click in the graph and select **Add Input Node | Object Center** from the popup menu. This creates a new Object center input. The Object center input will return a constant value over each flame. Connect the noise node input to the Object center input. The Phase is now different for all the flames, but constant over each single flame.
14. In order to further improve the realism of our fire, we will vary the speed of each single flame animation: for this, we will multiply the Time by a different value for each flame. Select the link between the “Time” input and the “Animation map” node. Create a filter node and change its type to **Multiply**.



15. Extract the **Multiply by** parameter of the filter node by pressing the corresponding  icon and connect the parameter to another filter, of type **Map**.
16. Change the **Output range** of the “Map” filter to [0.9; 1.1] and select the option to **Clip out of range values**. Connect the Map filter’s input to the noise node. This makes the velocity of the flames vary of $\pm 10\%$ over the EcoSystem population.

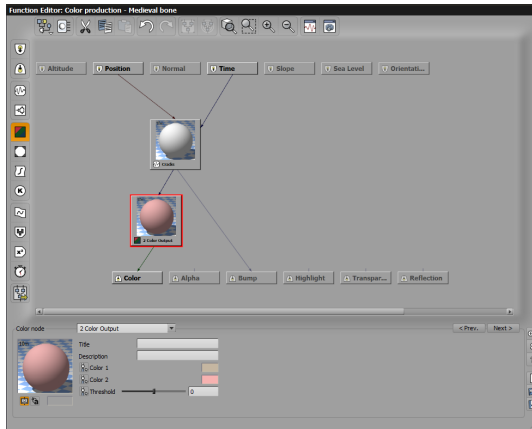
The animation is complete. Press the  icon to do a preview render. The complete scene can be found in the *Animation Tutorials* collection.

Dying Flesh




Dying flesh: step 4





Dying flesh: step 6

This tutorial is a simple illustration of *SmartGraph* material animation. The purpose of this tutorial is not to achieve a realistic effect (the result doesn't even look like dying flesh – whatever that may be!), but rather to show a basic example of what can be achieved using this type of animation:

1. Create a new scene and add a sphere. Double-click on the sphere's material preview to open the *Material Editor*.
2. Edit the bump production function in the **Bumps** tab. Create a **Line Patterns | Cracks** noise node.
3. Extract the **Crack width** parameter by clicking the corresponding  icon and connect the parameter to the **Time** input. This creates an animation of the material where the cracks reach maximum width within one second.
4. So that the material is initially smooth, we will also scale the depth of the cracks with time. Click on the **Filter** icon to insert a filter node after the crack noise. Change the filter node type to **Multiply**. Extract the **Multiply by** parameter and connect that to the **Time** input as well.
5. Now we would like the material to be initially pink, and then see some gray appear behind the cracks. Select the **Color** output and create a **Color map** node. Switch to **2 Color Output** type. Put a gray color in **Color 1** and a pink color in **Color 2**, and connect the node's input to the Crack noise node that is driving the bumps. Color 1 will appear behind the pink flakes.
6. Of course, you could improve the effect tremendously by mixing two different materials and driving the mix with an animated function. You could also use the animated noise nodes to create some gurgling effects on the flesh...

The animation is complete. Press the  icon to do a preview render. The complete





scene can be found in the *Animation Tutorials* collection.

Loose Dynamics




Loose Dynamics

In this tutorial we will see how loose dynamics can easily add a strong “real-world” feel to your animations.

1. Let's create a standard animation where an airplane flies over the camera that is filming it: load an airplane model, for instance the **FW190A3** from the **Vehicles/Aerial** collection. Animate it using the *Animation Wizard* (give it an **Airplane Dynamic Motion Reaction** preset) and plot it's path so that it flies in from the distance, over the camera, and away.
2. Select the camera, and make it track the airplane using the **Pick track** icon (). Add some clouds in the sky and an infinite procedural terrain to make things a little nicer.
3. Render a preview of the animation by clicking the **Animation Preview** icon (). The camera follows the airplane very meticulously, even when the airplane flies over the camera. In real life, there is no way the cameraman could be able to track the airplane so accurately. Besides, the fact that the airplane always remains exactly at the center of the frame is rather boring.




4. Let's spice things up a bit by adding some loose dynamics. In the **Animation** tab of the camera *Object Properties*, move the **Response** slider up to roughly $\frac{1}{4}$ of the available range (typical human reaction times). Render the animation again, and notice how the camera is not tracking the airplane as accurately as before, resulting in a much smoother overall movement. Notice also how the airplane moves around in the frame.
5. You may notice that, at the time when the airplane flies over the camera, the camera is a little "slow" to react, and it "loses" the airplane for a short time. To make the camera "anticipate" the movement of the airplane a little better, we'll have to fine tune its reaction using the *Forward Dynamics Options* dialog. Long-click on the **Pick track** icon () to open this dialog. We will fine tune the **Proportional**, **Integral** and **Derivative** terms of the controller in order to achieve the desired effect.
6. Adding a little bit of **Derivative** term will improve anticipation, but may also result in oscillations if you go over the top. The correct values are usually found in a trial and error process. Add a little to all the terms and render a preview. Repeat until you are satisfied with the result.

The complete scene can be found in the *Animation Tutorials* collection.

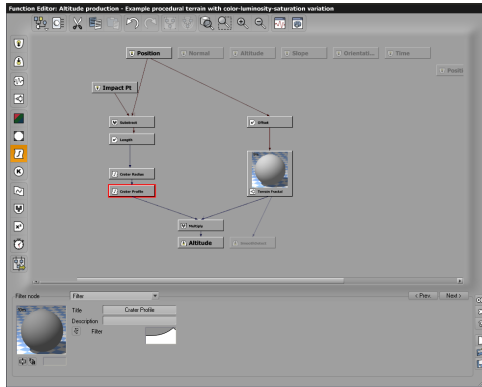
Crashing Meteor

In this tutorial we will see how we can setup a procedural terrain to react to the "impact" of a meteor. In this tutorial, it is assumed that you have a good working understanding of the *Function Graph* and *Object Graphs*.

1. Create a procedural terrain by clicking the **Load Procedural Terrain Preset** icon () and selecting one of the infinite terrain presets.
2. Create a pair of cubes, one high up above the terrain, the other just beneath the surface of the terrain. These cubes will be our guide for the motion of the falling meteor. Since they are only guides and shouldn't appear in the animation, make them both **Hidden from render** by clicking their icons in the *World Browser*. Rename both cubes as "Start Pos" and "Impact Point" respectively.
3. Create a sphere and rename it as "Meteor". Open the meteor's *Object Graph* and create two **External Dependencies** on both cubes. Blend these two values using a **Combiner | Blender** node, the **Ratio** of which will be controlled by time. Connect the Blender's output to the meteor **Position** output.
4. Create a **Filter | Divide** node and connect its input to **Time**. Divide the value by



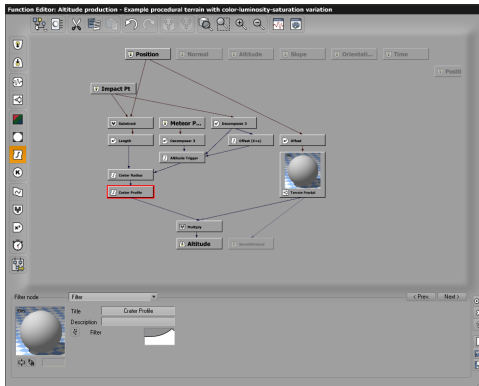
range. This is the parameter that controls the radius of the crater. We will drive this using the altitude of the meteor.



Terrain Altitudes – Step 11

11. Retrieve the **Position** of the meteor using an **External Dependency** node. Since we're only interested in the Z component (the altitude) of the meteor, use a **Decomposer 3** node and connect the Z output to a new **Filter | Map** node. This node will control the radius of the crater. Rename it to “Altitude Trigger” for the sake of clarity. Map its range from -500:0 to 10,000:0 and **Clip out of range values**, so that the crater has a radius of 0 when the meteor is above 0, and reaches its maximum radius of 10,000 when the meteor is at the impact point.
12. If you render a preview of the animation, you will probably notice that the crater does not appear exactly at the right moment. That's because the values we entered for the Altitude Trigger were “guessed”. In order to fine-tune this setting, we will improve our level of control over this value by connecting it to the altitude of the impact point: the Z component of the Impact Point's **Position** dependency node.
13. Extract the **Lower** and **Higher value** of the **Input range** of the Altitude Trigger map node, and connect the **Lower value** to the Z component of the Impact point's **Position** dependency. As a result, the crater will now reach its maximum radius when the meteor reaches the altitude of the Impact point.






Terrain Altitudes – Step 14


14. Let's consider that the meteor will sink 100 m deep into the terrain (1,000 Vue units). We want the crater to start growing when the meteor is less than a 100m above its point of impact. Use a **Filter | Offset** node to add 1,000 to the Z component of the Impact Point's **Position** dependency node. Connect the previously extracted **Higher value** of the **Altitude Trigger** to this node. The size of the crater is now entirely controlled by the position of the Impact point.
15. Close the editor and move the Impact point up or down until the crater starts appearing exactly when the meteor hits the terrain. Make sure that the camera's height isn't **Locked** to a fixed altitude above the ground, if not the altitude of the camera may change when the crater appears.

The animation is complete. Obviously, this is a very crude rendition of a meteor impact, but you could easily make it better. For instance, why not add shaking to the terrain by driving the terrain offset with a time-dependent function triggered by the meteor altitude? Or control the material based on the crater by making the crater affect the **Rough area Custom dependency**?


Press the  icon to do a preview render. The complete scene can be found in the *Animation Tutorials* collection.



EcoSystem Phasing



EcoSystem Phasing – Step 3



EcoSystem Phasing – Step 10





EcoSystem Phasing – Step 13

This tutorial demonstrates how to use EcoSystem Phasing for varied effects in your scene. This can be used to vary movement in crowds of people. Or, to vary movement in plants, such as a bed of seaweed, moving with the different currents.

If you don't have the *Broad Leaf Straight Trunk* tree used in this tutorial, just substitute with another tree, like the *Springtime White Birch*.

1. Open Vue and load a plant. Here we are using a *Broad Leaf Straight Trunk* tree.
2. Open the *Timeline*.
3. Move the time slider to second 2 and add some wind on the plant. If asked say no to animating the atmosphere.
4. Move the time slider to second 4 and remove the wind (to make an animation loop).
5. Save the animated plant as a *.VOB* file.
6. Open a new scene.
7. Add a plane, stretch it to a rectangular shape and create an EcoSystem on the plane's material.
8. Load the new animated plant.
9. Select the **Scaling & Orientation** tab and set the **Maximum rotation** to 0.
10. Select the **Animation** tab and check **Variable time offset**.
11. Set the **Time offset range** from 0 to 4 seconds (the duration of our plant animation).
12. Set the **Number of variations** to 100, which is 4 seconds of plant animation x 25 frames per second of animation. That is the number of positions that your instances can take. If you put 50, your instances will move every 2 frames.
13. Populate the EcoSystem at frame 0. You will see that some instances have the initial position of the animation and some the position of the end of the animation.
14. If you launch an animation render, all trees will be independently animated even though they are all based on the same single animated plant.



That illustrates the EcoSystem phasing feature.





Section 7

Animating Scenes





Animating from Scratch

In this section you will learn how to create a simple animation starting from a still scene. As you will see, VUE offers a powerful set of tools that turns part of the art of animation into a child's game.

Animation Properties Tab

Before you create an animation, you must first decide what objects in the scene will be animated. Let's say you want to animate a sphere. All you need to do is select that sphere, switch to the [Animation](#) tab, and select a type of **Motion** other than **Not animated** (which is the default). You will find further down a detailed list of all the available types of motion.

By selecting a motion type, you are turning animation on for the selected object (i.e. the sphere). However, no animation has been defined for that object yet.

To define an animation for the object you can use any of two following methods:

- use the *Animation Wizard* that was designed to help you setup your animations easily (see [here](#) for details), or
- use the *Timeline*.

Types of Motion

VUE features a set of elaborate algorithms called *Dynamic Motion Reaction™* that simulate the dynamic reactions of a number of predefined vehicles. They are called types of motion. By selecting one of these types of motion you instantly define major aspects of your object animation (e.g. airplanes bank as they turn). It highly simplifies the achievement of convincing animations by automating what can be a considerably tedious and time consuming process if done by hand.

VUE offers 10 different types of preset motions. Some of these can be further customized using the [Motion Options](#) dialog.

Some types of motion (i.e. vehicles) are airborne. They can vary their altitude relative to the ground beneath them. Others are grounded and will follow the surface of the ground they move on.



The available preset motion types are:

- **Standard:** the object moves from way point to way point with a near constant velocity. There may be sudden changes in object velocity when passing way points. Motion has no effect on object position or orientation. This type of animation is found in most 3D applications.
- **Smoothed:** basically the same as standard, except that the velocity of the animated object is automatically processed to ensure smooth acceleration/deceleration between way points. This type of animation is also available in some 3D applications. This sets the [Smoothed Velocity](#) property.
- **Look ahead:** the animated object is oriented in such a way that its main axis always points in the direction of travel. You can set which axis will be pointing in the direction of travel using the *Animation Wizard* (see below), or using the [Animation Properties](#) tab. Like in Smoothed motion type, velocity is processed to ensure smooth motion. This motion sets the [Look Ahead](#) properties. This type of animation is also available in a few other 3D applications.
- **Airplane:** now we get to the juicy ones! Airplane adds automatic banking to the Look ahead motion. What this means is that animated objects with an airplane motion type will bank automatically as they enter a bend, proportionately to the tightness of the curve! There's some complex physics going on behind to produce mechanically accurate banking... You can adjust the sensitivity to bend tightness using the [Motion Options](#) dialog.
- **Helicopter:** like with Airplane, objects animated with this type of motion will bank (slightly) as they turn. But they will also dip down as they accelerate! You can adjust the sensitivity to bend tightness and acceleration using the [Motion Options](#) dialog.
- **Missile:** basically the same as Airplane, except that objects with this type of motion will bank almost 90° as soon as they turn. You can adjust the sensitivity to bend tightness using the *Motion Options* dialog.
- **Automobile:** this is the first “grounded” type of motion. Objects animated with this type of motion will closely follow the ground they move on. The orientation of the objects is at all times given by that of the ground they are moving on.
- **Motorcycle:** also follows the surface of the ground, only banking as it turns, and looking up when accelerating. The bike eventually lifts on its rear wheel as acceleration gets strong enough! You can adjust the sensitivity to bend tightness and acceleration using the [Motion Options](#) dialog.
- **Pedestrian:** with this type of motion, the object also follows the surface of the ground. But it always looks straight ahead (in the direction of travel), whatever the slope of the ground it moves on.

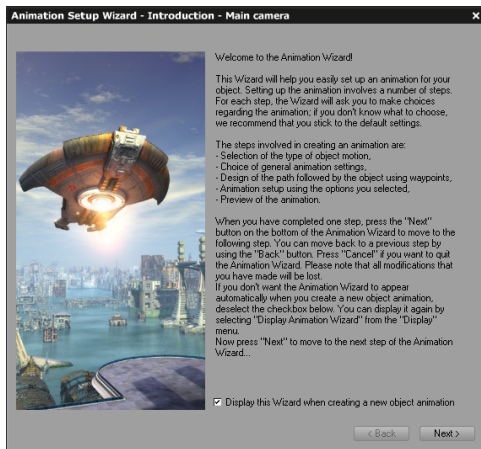


- **Speedboat:** the object follows the surface of the water (or the ground). It banks slightly as it turns, and looks up as it accelerates. You can adjust the sensitivity to bend tightness and acceleration using the *Motion Options* dialog.

To produce realistic motion, VUE uses accurate dynamic algorithms. This is your assurance of always obtaining realistic motion. Unfortunately, the drawback is that motion is dependent on scale: the larger the scale, the greater the speed of the object traveling from one way point to another. And the greater the accelerations it withstands, thus the greater the banking.

What this means is that if you feel your object is over-banking (or not banking enough) you may have to adjust its sensitivity to accelerations. This is done using the *Motion Options* dialog.


Animation Wizard



The Animation Wizard, step 1: Introduction


The *Animation Wizard* was designed to help you easily setup simple animations of your objects. All you need to do is follow each step of the Wizard, selecting any required options, and leaving others to their default values.

To display the *Animation Wizard*, do any of the following:

- Click on the Timeline icon () , or select the menu command **Display | Display Timeline** if it is not already displayed; the *Timeline* will be displayed after you are



done with the Wizard (see [here](#) for an introduction to the *Timeline*),

- Select a new type of motion using the [Animation Properties](#) tab; the *Timeline* will also be displayed when you are done with the *Wizard*,
- Activate the alternate action (see [here](#) for details) of the **Timeline** icon () or select the menu command **Animation | Animation Wizard**; the *Timeline* will not be displayed after you are done with the Wizard, or
- Press the *Animation Wizard* button in the [Animation Toolbox](#).

The name of the object that is currently being handled by the Wizard is displayed in the Wizard title. It is the object that was selected when you called the Wizard. If several objects were selected, then the Wizard looks for the first animated object in the selection. If no animated objects are selected, it picks the first object that was selected. If no objects were selected at all, it will operate on the camera.

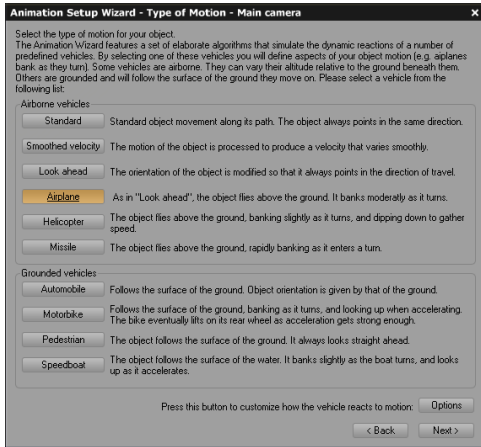
The Animation Wizard comprises 8 steps. Press the **Next** > button to move to the next step of the Wizard, and press < **Back** to return to the previous step.

Step 1: Introduction

This step merely displays information on operating the Wizard. It also gives you the opportunity to prevent the Wizard from appearing each time you create a new object animation (uncheck the **Display this Wizard when creating a new object animation** to prevent this from happening).



Step 2: Selecting a Motion



The Animation Wizard, step 2: Selecting a Motion

In this step you decide what type of Dynamic Motion Reaction the animated object will have. You will find a complete description of the different types of motion available in the section on [Types of Motion](#).

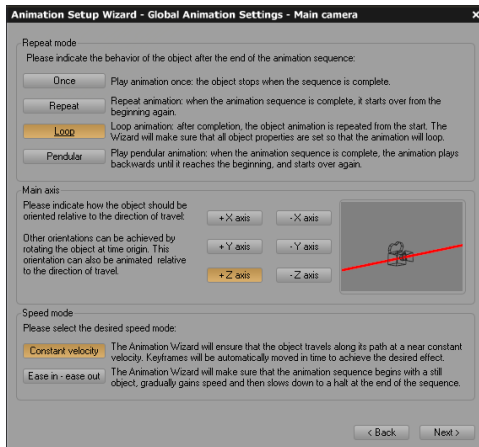
Click on the button corresponding to the type of motion you want to assign to the animated object.

If you have already defined a type of motion for the object, the corresponding button is selected. So you don't need to select it again.

If necessary, you can customize the sensitivity of the object to its motion using the *Motion Options* dialog (see [here](#) for details). Press the *Options* button to display this dialog.



Step 3: Global Animation Settings



The Animation Wizard, step 3: Global Animation Settings

In this step you choose settings that will affect the entire animation of your object, that is the Repeat mode, Main axis and Speed modes.

Repeat Mode

The repeat mode indicates how the object behaves when its animation is complete. By default, it simply stops, but you can instruct it to repeat the sequence in one of several ways:

- **Once:** this is the default setting: the object stops when its animation sequence is complete.
- **Repeat:** when the animation sequence is complete, it starts playing back from the beginning again.
- **Loop:** like repeat, except VUE does some extra processing to ensure that the first frame always corresponds to the last, thus ensuring a perfectly smooth and undetectable jump as it loops back to the start of the sequence. VUE will automatically add keyframes to ensure a smooth loop.
- **Pendular:** when the animation sequence is complete, it reverses, playing back until it reaches the start, and then starts playing normally again.

If you select a repeat mode other than Once, the animation repeats indefinitely. You can start a repeating animation anytime in the animation of the scene by dragging the first keyframe to the requested starting time (read [here](#) for an example).



Note:

Although the repeat mode applies to all object properties, repeating is done on a “per animated property” basis, which means that you can have an object with orientation being repeated faster than position!

You can also change the Repeat mode of your object using the [Animation Toolbox](#).

Main Axis

The Main axis setting is only available if you have selected a type of motion that makes the animated object look in the direction of travel (i.e. has the **Look ahead** property set). Standard and Smoothed motion types don’t give you access to the Main axis setting (because the setting isn’t applicable).

This setting lets you select which axis of your object will be pointing in the direction of travel (or which axis of the object will be pointing at the tracked object if the object is tracking another one).

The axes correspond to that of the object in the Top *3D View* when all object rotations have been zeroed (you can do this using the **Numerics** tab of the *Object Properties* panel, see [here](#) for details). So if in this view your object points to the right, then its main axis is **+X**; if it points to the left, it is **-X**. If it points upwards it is **+Y**, and if it points downwards, it is **-Y**. Finally, if it points out of the screen, it is **+Z**, and if it points into the screen, it is **-Z**.

The *Animation Wizard* displays a small preview of your object sitting on a large red arrow that indicates the direction of movement. This preview should help you decide which is the main axis of your object. If for some reason, none of the available axis seems to be right for your object, you might want to read the [Look Ahead Objects](#) troubleshoot.

You can also change the **Main axis** of your object using the **Animation** tab of the Object Properties panel (see [here](#)).

Speed Mode

This setting is not available if your object doesn’t have the [Smoothed Velocity](#) property. All motion types except Standard have this property set.

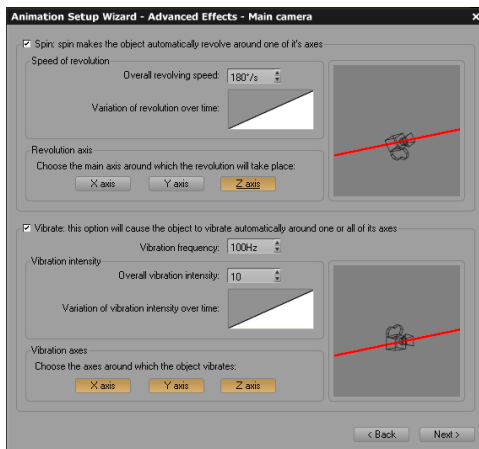
The speed mode is a powerful feature that automatically processes the path followed by your object (you will define this in the following step) so that the object either moves at a constant velocity, or stands still at the beginning and at the end of the animation sequence:

- **Constant velocity:** way points will be automatically moved in time so that the object travels at a near constant velocity,



- **Ease in – Ease out:** the animation sequence will start with the object standing still, then it will accelerate smoothly until it reaches its maximum velocity half-way through; it will then gradually slow down until it reaches a stop at the end of the animation sequence. This option sets a Time spline (see [here](#) for the position property).

Step 4: Advanced Effects



The Animation Wizard, step 4: Configuring Spin and Vibration

This step lets you enable and configure advanced automatic animation effects such as Spin and Vibration. By default, both effects are disabled.

Spin

The Spin effect lets you easily make any object spin precisely around one of it's axes, however complex be the path followed by the object. Using Spin, you can also easily and precisely animate the speed at which the object spins around its axis.

This effect, combined with the various types of motions available in Dynamic Motion Reaction, can create very complex movements.

If you want to apply spin to your object, select the **Spin** checkbox. The corresponding configuration controls will become active:

Overall revolving speed controls the number of degrees the object spins by each second. The default is 180 degrees per second, which means the object performs a complete revolution every 2 seconds. Increase the value if you want the object to spin faster. Enter



a negative value if you want to reverse the direction of spin.

The **Variation of revolution over time** graph is a filter that indicates the variation of the angle of the object over time. The steeper the slope, the faster the object will spin with time. By default, the filter exhibits a straight line meaning that the speed at which the object spins is constant. You can [modify the filter](#) by double-clicking or selecting **Edit** from the popup menu.

The range of time covered by the filter starts at the first position keyframe, and ends at the last position keyframe. If no position keyframes are defined, the range starts at time 0 and ends at the end of the animation.

Loading a *Power 2* filter from the *Other Filters* collection will result in the object gradually spinning faster and faster as time passes.

The **Revolution axis** lets you select the axis around which the revolution will take place. These buttons are exclusive because you can only spin around one axis at a time. If the object has an advanced type of motion, the axes will be contextual to the general movement of the object.

The small display to the right shows a preview of the Spin effect as applied to your object. The total duration of an animation cycle in the preview is 5 seconds (useful for calibrating the variation of revolution speed).

Note on Spin and Pivots

If you define a pivot position that is different from the position of the object, spinning the object will be done relative to the pivot point. You cannot however modify the pivot point once an animation is defined for the object, so you need to do so before applying Spin to the object.

Do not rotate and spin such an object simultaneously, as its path would become enormous.

Vibrate

Vibration is an extremely difficult and time consuming effect to achieve by hand. It can however lead to much more realistic movies than a perfectly stable motion. Vibration is to motion what texturing is to surfaces. Without it, they appear way too perfect.

Without the Vibrate effect, creating vibration requires the addition of numerous keyframes that are not only tedious to setup, but also a real nightmare to edit after hand.

Thanks to the Vibrate effect, you can automatically apply vibration to any given object, however complex its motion. You can also indicate on what axes of the object the vibration takes place, and this vibration will follow the object's orientation as it banks around a tight turn. Using the Vibrate effect, you can also easily and precisely animate the intensity



of the vibration of the object.

An important thing to understand about vibration is that the effect of the vibration depends on the type of motion selected. If the animated object doesn't have the [Look Ahead](#) property set, the vibration will cause a small displacement of the object around its trajectory, without affecting the direction in which the object points.

If the “Look ahead” property is set, the vibration will also cause small variations in the orientation of the animated object, resulting in an effect relevant of turbulence.

If you want to apply vibration to your animated object, select the **Vibrate** checkbox. The corresponding configuration controls will become active:

Vibration frequency controls the frequency of the vibration. Lower values will result in the object wandering randomly around its position, whereas high values will result in very quick and jerky movements of the object.

The **Overall vibration intensity** setting controls the amount of vibration that takes place. Lower values mean little vibration, whereas high values mean strong vibration.

The **Variation of vibration intensity over time** graph is a filter that lets you animate the amount of vibration over time. The higher the value of the filter, the stronger the vibration. By default, the filter exhibits a flat line meaning that the intensity of the vibration is constant over time. You can modify the filter by double-clicking or selecting **Edit** from the popup menu. Please read the section on *Editing Filters* for full details.

The range of time covered by the filter starts at the first position keyframe, and ends at the last position keyframe. If no position keyframes are defined, the range starts at time 0 and ends at the end of the animation.

Loading a *Tooth 10 filter* from the *Other Filters* collection will result in the object suddenly vibrating around the middle of the animation, and not vibrating the rest of the time.

The **Vibration axes** let you select around which axes the vibration will take place. Vibration can occur on any and all axes. These axes will follow the orientation of the object, including orientation caused by Dynamic Motion Reaction.

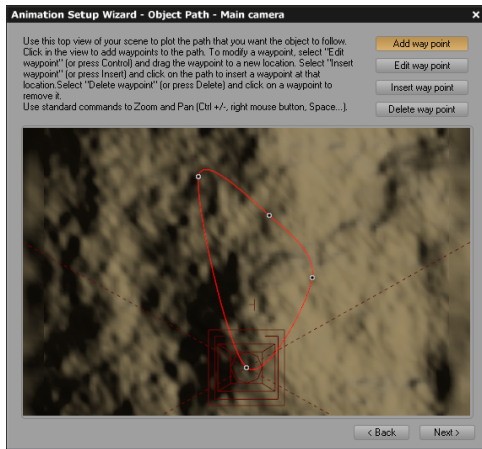
Note:

Applying vibration to the main axis of a Look ahead object can lead to unexpected effects...

The small display to the right shows a preview of the Vibrate effect as applied to your object. The total duration of an animation cycle in the preview is 5 seconds (useful for calibrating the variation of vibration intensity).



Step 5: Object Path



The Animation Wizard, step 5: Defining the Object Path

This is the fun part. In this step, you plot on a *Top view* of your scene the path that will be followed by your object.

Adding Way Points

This mode is active when the **Add way point** button is selected (the default).

Initially, there is no path defined for your object. All you see is a tiny black diamond that marks the initial position of your object. However, each time you click on the left mouse button, a new way point is added to the path. A way point is a point through which the animated object will necessarily pass. What happens in between way points is managed by the program. Way points are connected with a solid red line that indicates the path followed by the object. The path is automatically curved to produce the smoothest possible motion.

Note:

The way points are appended to the end of the path. If you need to add a way point somewhere in the path, you will need to use the Insert tool (see below).

At this time, you have no control over the altitude of the object at the way point.



Editing Way Points

If you need to modify the position of a way point, click the **Edit way point** button. When you move the mouse cursor over a way point, the cursor will change to the edit shape. Just click and drag the way point to its new position, then release the mouse button.

This tool lets you touch up the path followed by the object.

Inserting Way Points

If you need to add a way point somewhere in the middle of the path, select the **Insert way point** button. The cursor changes to the insert shape. Click on the path where you want to add the way point. If you move the mouse before releasing the mouse button, the newly inserted way point will follow. In this way you can locally adjust the shape of the path to fit your requirements.

Deleting Way Points

To delete any unwanted way points, select the **Delete way point** button. The mouse cursor changes to the delete shape. Just click on a way point to delete it. The shape of the path is automatically redrawn to fit the newly defined path.

Scrolling/Zooming the View

You can move around in the view using the standard controls used in the *3D Views* (right/Ctrl mouse drag, or Space + drag; read [here](#)). You can also zoom into or out of the view using the standard controls (Ctrl/Cmd + right/Ctrl mouse drag, or Ctrl + and Ctrl -).

Restrictions

Using the aforementioned tools, you can easily draw the path that will be followed by your object. However, due to the fact that you have no control over the altitude of the object as it passes the way point, there are some effects that cannot be achieved solely with the Wizard. You will need to touch up the path in the *3D Views*, when you are done with the Wizard. Please read section on [Editing Paths in 3D Views](#) for details on how to modify way points in the *3D Views*.

It is important to understand at this time how the Wizard expands the 2 dimensional path you have just drawn into a full 3D motion. Basically, it looks for the highest object encountered at each way point, and positions the animated object at the same altitude above that object as its altitude at the initial position. Some more processing is done later to avoid hitting any objects.



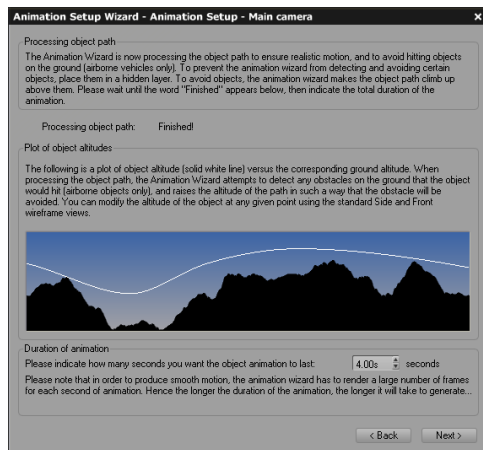
The Tunnel Case

This is typically an effect that cannot be achieved solely with the Wizard. To fly an object through a tunnel, or under an arch, you will have to modify the path manually in the *3D Views*. This is because the Wizard will detect the ceiling of the tunnel, and automatically position your animated object on top of that ceiling. You will have to drag the way points down, back into the tunnel.

Note:

Other animated objects don't move in this *Top view*, which can make path edition a matter of trial and error if you want to animate an object relative to another one. You'd probably better off doing this directly in the *3D Views*.

Step 6: Animation Setup



The Animation Wizard, step 6: Animation Setup


In this step the Wizard processes the path you have defined, attempting to maintain the same altitude above the ground as the initial altitude of the animated object (read details about this process in the preceding section). It also attempts to avoid hitting any objects lying in the way of the animated object. This is an iterative process that results in a path that more or less maintains the same altitude relative to the ground, but can get closer or farther in some areas.

Processing can take some amount of time. When it is finished, a plot of the animated object altitudes compared to “ground” altitudes is shown. You cannot act upon this plot at this time. You will have to do it in the *3D Views*.



To complete this step, just enter the total duration of your objects animation. The Wizard will automatically resample your path so that it completes in the requested time. Obviously, the longer the animation, the longer it will take to render...

Step 7: Animation Preview



The Animation Wizard, step 7: Animation Preview

All you have to do in this step is watch your object animation and decide whether it is satisfactory. If not, you can switch back to the previous steps, and modify the animation path.

Note:

the animation is played as seen from the camera point of view. For some reason, your object may not be visible at this time. You can always play back the animation later, inside the *3D Views*.

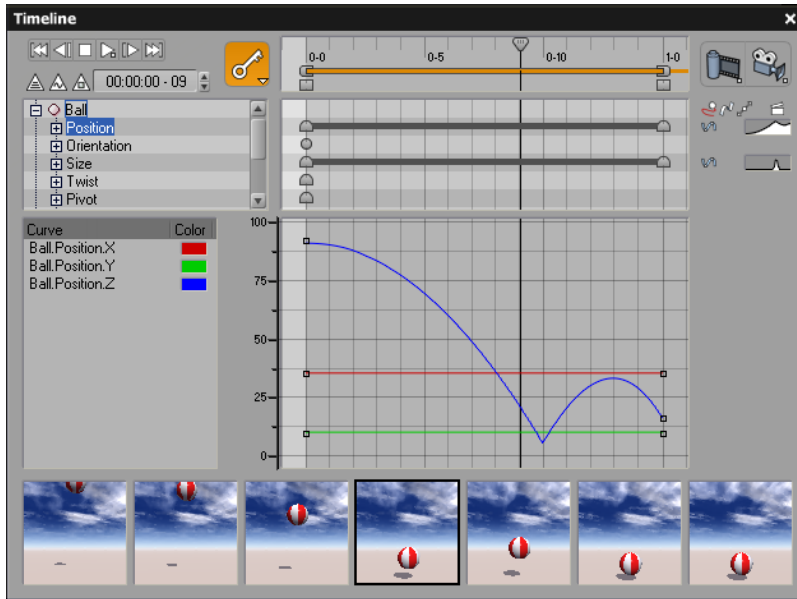
Step 8: Conclusion

This is the last step of the Wizard. You have nothing to do here (apart from reading the summary information on using the *Timeline*).


When you close the Wizard, the *Timeline* appears on the screen (if it was not already there). Use the *Timeline* to fine tune your animation. You will find details on how to do this in the next section.



Animating with the Timeline



The complete Timeline (Properties and Preview Timelines unfolded)

To display the animation *Timeline*, select the menu command **Display | Display Timeline** or click the **Display Timeline** icon (). This displays the *Timeline* together with the *Animation Wizard* (you can disable this feature). The *Animation Wizard* helps you easily setup an animation of your scene.

As explained in the [Interface Overview](#) section the *Timeline* is subdivided in 3 sections: the *Main Timeline*, the *Properties Timeline* and the *Animation Preview*.

In this section, you will learn in detail how to use the *Main Timeline* and the *Properties Timeline*. For greater clarity we will consider that the *Properties Timeline* has been expanded.


The *Timeline* can be either docked at the bottom of the interface, or placed in a floating panel. To dock or undock the *Timeline*, select the menu command **Display | Dock Timeline**.



Navigating Inside the Animation

Use the navigation controls () to easily find your way through the animation.



Start of animation: moves the current time to the start of the animation sequence (as defined by the  slider).





Previous keyframe: moves the current time to the first previous keyframe available. If there is no keyframe before the current time, the speaker will beep.



Stop: press this button to stop playing the animation.




Play animation: press this to start playing a 3D preview of your animation in the *3D Views*. The **Animation Preview** is also played (if it is visible). The  button has an alternate action () that plays only the Animation Preview. This ensures smooth playback, since the program doesn't have to redraw all the *3D Views*.



Next keyframe: moves the current time to the first following keyframe available. If there is no keyframe after the current time, the speaker will beep.



End of animation: moves the current time to the end of the animation sequence (as defined by the last keyframe in the scene or by the  slider, whichever is encountered first).



Keyframes

Keyframes indicate a change in one or more object properties (position, orientation, size...) at a given time. VUE automatically interpolates the property to ensure smooth animation between the keyframes.

Past the last keyframe, the property stays identical to what it was at the last keyframe (unless you have defined a **Repeat mode** other than **Once**; read [here](#)).

Only the keyframes belonging to the currently selected objects are displayed in the main *Timeline*. If no objects are selected, the keyframes of all objects, materials and atmosphere in the scene appear.

The time lapse between keyframes indicates the time that the object property will take to evolve from its state at the current keyframe to its state at the following keyframe. For instance, if you consider the object position property, the time lapse between the keyframes indicates the time that the object will need to travel from the current way point to the next way point.

By varying the time lapse between keyframes, you can make the object accelerate or slow down by dragging the keyframes to the left or to the right.

The *Animation Wizard* includes very elaborate algorithms to produce constant velocity motion. So if you have defined the animation path of your object using the Wizard, keyframes will have been automatically positioned in time to produce the smoothest possible motion. You can also achieve this effect by using the *Animation Toolbox* dialog (press the **Make constant velocity motion** button).



Animation Properties

An animation property is a property of an object that can be animated. Each type of object has different animation properties. The following is a list of these properties, and how they are processed.

Standard Primitive, Polygon Mesh, Group, Boolean Object, and Metablob and Plant

- **Position:** position keyframes define the position of the object at the keyframe time. They are interpolated using splines (the mathematical implementation that ensures the smoothest possible movement).
- **Orientation:** orientation keyframes define the rotation of an object at the keyframe time. They are interpolated using quaternion arithmetic to produce the best rotation paths possible, and support multi-spins (read [here](#)).
- **Size:** size keyframes define the size of the object at the keyframe time. They are interpolated linearly.
- **Twist:** twist keyframes define the twist of the object at the keyframe time. Twist animation can produce strange results...
- **Pivot position:** pivot position keyframes define the position of the object at the keyframe time. They are also interpolated using splines.
- **Material:** material keyframes are not directly linked to the object itself, but rather to the material that the object is made out of. If you add a keyframe to a material, all objects that use this material will display the new keyframe automatically. You can move or delete material keyframes here. If necessary, a new material will automatically be created by VUE for that object. [Animated materials](#) are displayed at the bottom of the list of animated objects.

Underneath their own properties, Groups and Boolean objects display a list of their members.



Torus

- **Position, Orientation, Size, Twist, Pivot position, and Material:** as above.
- **Thickness:** thickness keyframes define the thickness of the rim (also known as the outer diameter) of the torus at the keyframe time. They are interpolated linearly. Read more about animating the thickness of a torus [here](#).

Terrain

- **Position, Orientation, Size, Twist, Pivot position, and Material:** as above.
- **Geometry:** geometry keyframes define the map of altitudes of the terrain at the keyframe time. Use this amazing feature to morph the shape of the terrain! Read more about animating terrain geometry [here](#).

Plant

- **Position, Orientation, Size, Twist, Pivot position, and Material:** as above.
- **Wind:** wind keyframes define the strength and direction of the wind at the keyframe time.
Wind can be animated to create sudden gusts of wind on a given plant (values are interpolated linearly between keyframes).
- **Geometry:** geometry keyframes define the overall shape of the plant at the keyframe time.
Use this amazing feature to morph the shapes of plants, thus simulating plant growth or transformation! Read more about animating plant geometry [here](#).

Directional Light

- **Orientation:** as above. **Position, Size, Twist, Pivot position, and Material:** are not relevant for Directional lights (position is linked to orientation).
- **Color:** color keyframes define the color of the light at the keyframe time. They are interpolated linearly to produce smooth changes.
- **Softness:** softness keyframes define the softness of the shadows cast by the light at the keyframe time. Yes, you can even animate this property!



Point Light and Quadratic Point Light

- **Position, Color, Softness and Pivot position:** as above. Orientation, Size, Twist and Material are not relevant for these types of lights (they cast light in all directions).
- **Power:** power keyframes define the power of the light at the keyframe time. They are interpolated linearly.

Spot Light and Quadratic Spot Light

- **Position, Orientation, Color, Softness and Pivot position:** as above. Size, Twist and Material are not relevant for Spot lights.
- **Spread:** spread keyframes define the angle of the cone of light spread at the keyframe time. They are interpolated linearly.

Ventilators

- **Intensity:** intensity keyframes define the intensity of the wind generated by the ventilator at the keyframe time. Intensity keyframes are interpolated linearly.
- **Cut-off:** cut-off keyframes define the cut-off distance at which the ventilator ceases to affect plants. Cut-off keyframes are interpolated linearly.
- **Spread:** spread keyframes define the angle of the cone in which wind is blown at the keyframe time. They are interpolated linearly.
- **Falloff:** falloff keyframes define the rate at which wind intensity drops near the edges of the spread cone at the keyframe time. They are interpolated linearly.

Camera

- **Position, Orientation and Pivot position:** as above. Size, Twist and Material are not relevant for cameras.
- **Focal:** focal keyframes define the focal length (or angle of view) of the camera at the keyframe time. Focal keyframes are interpolated exponentially to produce the zooming effect the eye is used to seeing.
- **Blur:** blur keyframes define the amount of blur at the keyframe time. High blur settings mean that the depth of field is reduced. Blur keyframes are interpolated



linearly.

- **Focus point:** the focus point distance defines the distance from the camera at which objects are seen in focus at the keyframe time. The distance ahead and behind that point where objects are still in focus depends on the blur (depth of field) setting. Focus point keyframes are interpolated exponentially to reproduce the effect that the eye is used to seeing.
- **Exposure:** exposure keyframes define the exposure setting of the camera at the keyframe time. Exposure is adjusted either using the *Object Properties* panel when the camera is selected, or inside the *Atmosphere Editor*. Exposure keyframes are interpolated linearly.
- **Motion blur length:** motion blur length keyframes define the amount of motion blur (as a ratio of the shutter opening time over the duration of a frame) at the keyframe time. Motion blur length is adjusted using the *Camera Options* dialog (see [here](#) for details on animating motion blur length). Motion blur length keyframes are interpolated linearly.
- **Post processing:** post processing keyframes define the post processing settings applied to the camera at the current time. Post processing is set using the *Camera Options* dialog. Post processing keyframes can be defined for each camera independently, or can apply to all cameras simultaneously.

Key frames for material, atmosphere and post-processing settings also appear in the list of animated properties.

Animating Objects

In this section you will learn how to build object animations using the *Timeline*.


You don't have to use the *Animation Wizard* to create an animation of your objects. In fact, there are many effects that cannot be achieved using the Wizard. This is why it is important to understand how to build animations without the help of the Wizard.




Creating the Animation

To create an animation in the *Timeline* you must modify a property of the object at a different time than the time the object was created at. That is, if you create your object when the Current time slider points to 1 second, you will have to modify it at some other time than 1 second to animate it. The object “remembers” the time it was created at (its date of birth?).

The cool thing about creation time is that you can freely modify an object at its creation time without it becoming animated. Great. So what happens if you forget what time you created your object at? And do you have to keep changing the current time each time you want to modify an object without it becoming animated? No, of course! If the current time is 0, all modifications will be considered done for the object at the time of creation.

If you want an object to never become animated, click the **Forbid animation** icon () in the **Animation** tab of the Object Properties panel (see).

Say you want to create a movement for a sphere: create the sphere, then drag the time slider up 1 second. Now drag the sphere to a new location. The sphere gets automatically animated, and will move from the first position to the second in a 1 second time lapse.

You can check this by pressing the play button (). By default, the object gets the **Standard** type of motion (see [here](#)). You can change this using the **Animation** tab of the *Object Properties* panel.

This works for all properties, including orientation. Better still: you can rotate an object relative to its direction of travel (**Look ahead** objects), so that you can animate an airplane that looks down 30° all the way. You can even animate this relative rotation (make your airplane look up 30° at the beginning of the animation, and down 30° at the end)!

Published Parameters for Animated Objects

If you have a published parameter for an object and the value of that parameter has changed when time is other than zero, this parameter will appear on the timeline and a keyframe will be added.

During animation, the parameter is interpolated, working the same as an object position animation, for example. You can change the spline describing the evolution of the parameter in the *Timeline*. Refer [here](#) for more information about object published parameters.



Published Parameters for Animated Textures

When you use published parameters for animated textures, VUE will do an interpolation of the texture which will achieve better animation results.

Working with Keyframes

Keyframe Types

The shape of a keyframe indicates the type of interpolation on both sides of the keyframe (see [Keyframe Tangents](#) for details).

If the type of interpolation is different on both sides of the keyframe, the shape of the keyframe will be different on both sides in order to reflect the difference in interpolation.

Selecting Keyframes

To select a keyframe, click on it in the *Timeline* rulers.

To deselect all keyframes, click on an empty part of the rulers.

If you want to select multiple keyframes at the same time, click on an empty part of the ruler, and drag a marquee rectangle to encompass all the required keyframes. You can extend a selection of keyframes by pressing Control at the same time as you click on the new keyframe. If you Control select an already selected keyframe, the said keyframe will be deselected. By pressing Shift instead of Control, you will also select all keyframes situated in the interval between the last selected keyframe and this new one.

You can also select a property by positioning the Current time slider at the time of the keyframe, and clicking on the property of the object you want to select the keyframe for.

To select all keyframes of a given animation property, double-click on the said property.

Moving Keyframes

To move a keyframe, click on it then drag it with the mouse button down. Release the mouse button at the desired location.


To move a set of keyframes, select the keyframes (as explained above), then click on one of the selected keyframes and drag it with the mouse button down. All other selected keyframes will move with it.

If you press **Control** as you drag a keyframe, all the keyframe that follow will be dragged with it.



Adding Keyframes

When auto-keyframing is enabled (the default), keyframes are automatically added each time you modify an object property at a time where no keyframe is defined for that property. To add a keyframe, position the **Current time** slider at the requested keyframe time, and modify the object. The new keyframe is automatically created for the property that was modified (e.g. if you move the object, you will create a new Position keyframe).

If auto-keyframing is disabled, click the **Add Keyframe** icon () to manually add a keyframe at the current time. You can access further options by long-clicking this icon. Please turn [here](#) for a discussion on auto-keyframing versus manual keyframing.

You can also select the menu command **Add Keyframe** from the Timeline popup menu (click on the right/Ctrl mouse button to display this). Keyframes will be created for all object properties (except material, which is handled somewhat differently).

Keyframes can also be added by clicking on the appropriate spot on the **Property Timeline** or in the Animation curves.

Copy-Pasting Keyframes

You can copy keyframes in the *Timeline* by selecting one or several keyframes and pressing **Ctrl + C** (**Cmd + C** on Mac), or selecting **Copy Keyframe** from the *Timeline* menu.



If you drag the Current Time slider to a new time and press **Ctrl + V** (**Cmd + V** on Mac), or select **Paste Keyframe** from the *Timeline* menu, the keyframes that you copied to the clipboard will be pasted at the current time. If you selected several keyframes that are not all at the same time, they will be pasted with the first keyframe being placed at the current time, and others being placed subsequently, with the original difference in time.

Deleting Keyframes

To delete a keyframe, select the keyframe and then press **Delete**.

You can also delete the keyframe by positioning the Current time slider at the keyframe time, selecting the required property and then choosing the menu command **Delete keyframe** from the *Timeline* popup menu (click on the right/Ctrl mouse button to display this).

Modifying the Value of a Keyframe

To modify a property at a keyframe, first move the current time to that of the keyframe (use the time slider, or better still, use the  and  controls to jump directly to the keyframe). Now set the object property to the new value. The keyframe is automatically



modified.

For instance, to modify the path of an object, move to the required position keyframe, then drag the object to a new location in the *3D Views*.

Keyframe Tangents

When a keyframe is selected, the tangent to the curve at this point will be displayed in the *Animation Graph*. You can drag the end of the tangent to change the shape of the curve around the keyframe. If you press Control while dragging the end of a tangent, only that half of the tangent will be modified. A “break” in the curve will be created around the keyframe.

On top of changing the orientation of the keyframe (which modifies the curve “velocity” around the keyframe), you can also change its length. This controls how close the curve sticks to the tangent around the keyframe, and is known as the “tension”. The longer the tangent, the closer to the tangent the curve stays.

Note:

You can also modify position keyframes and tangents in the *3D Views*.

Keyframes can be of a number of predefined types. These types basically control how the tangents around the keyframe are constructed. You can change the type of keyframe using the **Keyframe Options** menu. The different types of keyframes are identified by the following pictograms:



Smooth (Constant): in this mode, the tangent to the keyframe is created in such a way as to minimize the deformation to the curve. Subsequent modifications of the keyframe will not affect the tangents.



Smooth (Weighted): in this mode, the tangent to the keyframe is created in such a way as to minimize the deformation to the curve. If you modify the keyframe, the tangents will be recomputed to keep the deformation of the curve minimal.



Ease In/Ease Out: in this mode, the tangents are always flat. This avoids “jolts” in the animation and results in slow transitions around keyframes, with smoother overall animation.



Linear: in this mode, the tangents are created to ensure linear keyframe interpola-



tion. The tangents on both sides of the keyframe are usually different, resulting in a break in the animation curve.



Step: in this mode, the value of the animation property is constant in between keyframes. The value remains equal to that of the keyframe until the next keyframe is reached.



Custom: in this mode, the tangents are user defined. This becomes the active mode as soon as you modify a tangent manually.

You can define a different tangent mode on either side of the keyframe. In such a case, the pictogram identifying the keyframe is different on both sides of the keyframe. This is done using the other options on the **Keyframe Options** menu:



Smooth In (Constant): this makes the current keyframe mode the same as **Smooth (Constant)** on the left side of the keyframe only.



Smooth In (Weighted): this makes the current keyframe mode the same as **Smooth (Weighted)** on the left side of the keyframe only.



Ease In: this makes the current keyframe mode the same as **Ease In/Ease Out** on the left side of the keyframe only.



Linear In: this makes the current keyframe mode the same as **Linear** on the left side of the keyframe only.



Custom In: this makes the current keyframe mode the same as **Custom** on the left side of the keyframe only.



Smooth Out (Constant): this makes the current keyframe mode the same as **Smooth (Constant)** on the right side of the keyframe only.



Smooth Out (Weighted): this makes the current keyframe mode the same as **Smooth (Weighted)** on the right side of the keyframe only.



Ease Out: this makes the current keyframe mode the same as **Ease In/Ease Out** on the right side of the keyframe only.



Linear Out: this makes the current keyframe mode the same as **Linear** on the right side of the keyframe only.




Custom Out: this makes the current keyframe mode the same as **Custom** on the right side of the keyframe only.

Note:

By default, tangents are created as Smooth (Constant).

Keyframe Values



Keyframe Values Editor

When a keyframe is selected, you can edit the keyframe parameters using the popup menu command **Numerical Input**. This opens the *Keyframe Values* editor, letting you input precise numerical values for all meaningful keyframe parameters (time, value, velocity and tension).

Quaternion vs. Euler Rotations

By default, VUE handles orientation animation is processed using quaternion arithmetic. Thanks to quaternion arithmetics, you can easily create animations that smoothly interpolate orientation keyframes.



Quaternions are both easy to use, and produce smooth results. The downside of quaternions, however, is that you cannot break-up their components into anything intelligible.



So you cannot control quaternion animation as precisely as other animation properties.

In order to gain precise control over the orientation animation, you need to change the underlying orientation animation model to Euler orientation.

Euler orientation is based on combining rotations around the 3 different axes. These rotations are applied in a specific order (you can set this order using **Default rotation order** in the *Options* dialog – see [here](#)).

To switch to Euler orientation, simply expand the orientation property by clicking on the  /  symbol in front of the property name. A message will appear, informing you that the orientation model is about to change. VUE will compute the Euler angles that correspond to the quaternion orientation keyframes, but interpolation of the keyframes will be affected by the change – the animation will look different. Orientation properties that use the Euler model are followed by ‘(XYZ)’ in the list of animation properties.

Note:

Once you have switched from Quaternion to Euler orientation model, you cannot switch back.

Editing Paths in 3D Views

The case of the position property is slightly different from other properties. The reason for this is that position is very often animated. So some extra tools are supplied to make the modification of paths easier and faster.

Selecting Way Points

In the active 3D View (see [here](#)), position keyframes (i.e. way points) are depicted by tiny red dots on the object’s path. If you drag the mouse over one of these dots, the cursor will change to the **Move way point** cursor.

Clicking over one of these dots makes it turn white. It is now selected. Notice how it also becomes selected in the *Timeline*.

You can select multiple way points using the standard Control and Shift commands (Control extends the selection, Shift extends the selection, selecting all way points between the current way point and the last one selected). You can even select way points that belong to different objects!

Double-clicking on a way point will select all the way points of the object.



Gray Way Points

Some way points are depicted in gray instead of red. They cannot be selected. As you will notice, these are the way points that are close to the object itself. This is to avoid confusion between the object and its way points.



To access a gray way point, first select a red way point. All the way points of the object now become selectable (and consequently turn red). You can now select the desired way point by clicking on it.

Moving Way Points

To move a way point, just click on it and drag it to a new location. The path is automatically processed to smoothly travel through the way point.

Editing Way Points

The power of the way point editor really turns on when you select several way points. All the standard object modification controls become available!



So you can **Rotate** a whole group of way points using the  and  rotation handles (it's so easy to modify the global orientation of the path using this feature)! You can **Resize** groups of way points using the standard object size controls too (the black dots on the corners of the group of way points)! You can use the flip tools, the alignment tools, and you can use the Drop command to drop whole groups of keyframes (also works for single way points)!

Changing Rigged Mesh Motion

When you double-click on a rigged mesh, the *Skeleton Editor* becomes active (see [here](#)). Click on the **Open** button to select an alternate *Motion* file.

Multi-Spins

Although VUE always attempts to find the shortest possible rotation path when interpolating orientation keyframes, it is possible to create a rotation of several revolutions. This is called a multi-spin.

To create a multi-spin drag the current time slider to the time of the end of the multi-spin, and start rotating the object using one of the two rotation handles ( and ). You will notice that an indication of the angle of rotation and the number of rotations is displayed in the *Status Bar*. If you keep rotating the object past the full turn, the



revolution counter will indicate one revolution. Keep “winding” the object for as many revolutions as required, and release the mouse button when you are done. Playing back the animation will show the object rotating for the required number of revolutions.

However, you cannot “add” revolutions to an existing rotation. When you modify the orientation property, the counter for the number of revolutions is reset. You need to get this right in one go (you can always start again, but you will have to “wind” the object all the way). Also, you cannot create a multi-spin by typing the rotation angle in the **Numerics** tab of the *Object Properties* panel.

Animating Plants

On top of the standard animation possibilities, there are three other ways of animating plants:

- **Breeze:** each plant that you create is automatically subject to the global breeze. Global breeze is adjusted using the *Atmosphere Editor* (turn [here](#) to for details).
- **Wind:** you can define a per-plant wind level, and animate it.
- **Geometry:** you can also animate the geometry of plants. Please read below for details.

Breeze

You don’t need to do anything for a plant to move in the breeze; the plant doesn’t even need to be explicitly animated! Simply create a plant and render an animation of it, and you will see that it moves in the breeze (provided that breeze is enabled).

Note:

You cannot create strong wind effects with breeze alone. For such effects, you will have to use wind.

Animating Wind

On top of the default breeze animation, you can also animate the wind that is applied to the plant. To animate the wind property, drag the current time slider to the time where you want to create the new keyframe, and simply modify the intensity or the direction of the wind using the wind control in the *Top View* (see [here](#) for details on setting wind intensity and direction).

The plant automatically becomes animated, and, if you look at the *Timeline*, you will notice that a keyframe has been added to the Wind property of the plant.

Note:



The movement of the plant subject to wind animation does not stop on the last wind keyframe (even if no wind is defined). This is due to the complex relaxation model used in VUE to model the effects of the wind. In this model, plants subjected to varying wind intensities will “spring back” when the wind stops varying.

Animating Plant Geometry

Plants, like Terrains are special in the sense that you can animate their shape (i.e. their geometry).

To animate the geometry (or shape) of a plant, move the current time slider to the time where you want to create the new geometry keyframe, and open the *Plant Editor* (read [here](#) for details on the *Plant Editor*). Modify the shape of the plant using the controls in the *Plant Editor*, then press **OK**. A message will appear asking whether you want to animate the geometry or not. Click **Yes**. The plant becomes animated, and you can check in the *Timeline* that a new Geometry keyframe has been added at the current time. Changing current time will modify the geometry of the plant as it is interpolated between the previous keyframe and the next one.

If you refuse to animate the geometry of the plant (by answering **No** to the above prompt), the **Forbid animation** option will be activated for that plant (read [here](#)). If you decide later that you would like to animate the geometry after all, you will have to remove the **Forbid animation** option first (using the icon in the **Animation** tab of the *Object Properties* panel, see [here](#)).

You can also animate the materials used by the plant, either directly in the **Aspect** tab of the *Object Properties* panel, or by editing the materials inside the *Plant Editor*.

Amazing effects can be achieved using this plant geometry animation feature. Read the tutorial on [Dying Plants](#) for an example.


Animating Torus Thickness

The torus has an extra animation property than standard primitives like spheres, cubes, etc. This property is the thickness of the torus rim (AKA outer diameter). To animate this property, drag the current time slider to the time where you want to create the new keyframe, and open the *Torus Options* dialog (read [here](#) for details on this). Select a new rim thickness and press **OK**.

The torus automatically becomes animated, and, if you look at the *Timeline*, you will notice that a keyframe has been added to the Thickness property of the torus.



Camera Switching




Sample animation with camera switching

VUE's animation system lets you change the live camera (that is the camera that is used to view the scene) at any point during an animation. This is known as camera switching. Camera switching provides an interesting way to enhance the dynamism of video clips; it is also essential if you want to create storyboards.

Switching cameras is very straight-forward: simply drag the Current Time slider to a new position in the *Timeline*, and select another camera (for instance using the **Previous** and **Next Camera** icons in the *Camera Control Center* – see [here](#)). A new item will automatically be added to the *Timeline*. This item, known as the *Camera Switcher* (see opposite), contains all the cameras that are used throughout the animation. If a camera was already displayed in the list, it will be replaced by the camera switcher. By unfolding the camera switcher, you can gain access to the independent cameras, and adjust all their animation properties independently.

The camera switcher displays a blue line alongside each camera, indicating the period of time during which each camera is live. A keyframe at the end of each line indicates the time at which the switching occurs, and a thin line leads to the new live camera. You can modify the time at which the switching occurs by dragging the corresponding keyframe. The last camera will remain live until the end of the animation, or for 1 second after switching, whichever ends last.

You can prevent camera switching from taking place by making a camera “unswitchable”.

This is done by selecting the **Non switchable camera** option () in the **Aspect** tab of the camera's *Object Properties* panel (see [here](#)). When this option is set, activating this camera will not make it the live camera for rendering, and it will not create a camera switching keyframe. This is particularly useful if you have setup a camera to view your scene from a different point of view (like a director's camera), and you don't want to create a camera switch each time you use that camera.

Cameras don't have to be animated in order to be live, and making them live will also not make them animated. But you can definitely use animated cameras for camera switching. However, because a given camera is not accessible outside its “live” time span, you cannot edit the camera animation by simply dragging the Current Time slider and modifying the



camera settings (because this camera may no longer be the live camera at this new time). To be able to modify camera animation outside the camera’s “live” time span, you first need to select that camera in the camera switcher (unfold the camera switcher’s content and click on the camera to edit). When this is done, the camera will remain active even at times where it isn’t live, thus letting you edit the camera settings at any point in time. Cameras that are selected this way will remain active until they are deselected.


Another interesting aspect of camera switching is that you can easily create cameras that have different post processing settings (see [here](#)). This way, you can easily have one camera film the scene in black and white, while another one films it from another point of view, this time in full color.

Animated Post Processing and Motion-Blur Length

The post processing and motion-blur settings in the *Camera Options* dialog (see [here](#)) can be animated. When you edit these settings at a new time, new post processing and motion blur length keyframes will be automatically added to the *Timeline*.

Animating Post Processing

To animate post processing, simply drag the Current Time slider to a new time and double-click on the active camera to display the *Camera Options* dialog. Adjust the post processing settings and click **OK**. A message will appear asking whether you want to animate the post processing settings. Click **Yes** to animate the post processing. A new post processing keyframe will be added.

If post processing is common to all cameras (the **Post processing applies to all cameras** option is selected in the *Camera Options* dialog), post processing keyframes will appear in the global post processing property (identified by the  pictogram). If post processing is on a per-camera basis, they will appear in the post-processing property that is added to the end of the list of camera animation properties.

You can edit post processing keyframes like any other keyframe, but you have to keep in mind that enabling or disabling post processing options in the *Camera Options* dialog enables or disables these options throughout the entire post processing animation. For instance, if you uncheck the Post processing option to remove post processing at a given time, this will remove all post processing in the entire animation.

Animating Motion Blur Length

To animate motion blur length, simply drag the Current Time slider to a new time and double-click on the active camera to display the *Camera Options* dialog. Adjust motion blur length and click **OK**. A new motion blur length keyframe will be added.



Motion blur length keyframes can be edited like any other keyframe.

Animating Materials

Although there is a material animation property for most objects, material animation is not done in the *3D Views*. It is done directly in the *Material Editor* (read everything about the *Material Editor* [here](#)).

If you want to animate the material of an object, you will first need to open the *Material Editor*; double-click on the material preview in the **Aspect** tab of the *Object Properties* panel (see [here](#)).

There are 4 different ways of animating materials (by order of complexity):

- Material Surface Animation,
- Material Velocity Animation,
- Complete Material Animation, and
- SmartGraph Material Animation.

These different types of material animation will now be presented in detail.

Material Surface Animation

This is the simplest method of animating a material. All you have to do is select the **Animate material surface** option in the *Material Editor*. The **Time Dependent Material** notice becomes visible in the caption of the *Material Editor*. Also, if you check the *Timeline*, you will notice that the material is now listed at the bottom of the list of animated objects. Animated materials always appear at the bottom of the list. No keyframes are available for that material, because it is the same material that is being modified by time (turn to [Complete Material Animation](#) to find out how material keyframes are created).

What this does is replace by the current time the Z component of the position at which the material is being computed. Strange idea? Not quite: since all the procedural noises used in the construction of a Function are defined in three dimensions, replacing one of these dimensions by the time means that these procedural noises will become animated (you can read the section on the [Function Graph](#) if you don't understand this). You'll have undulating waves appear at the surface of water, clouds that change shape over time, moving underwater caustics, and many, many more exciting effects...



The drawback of this type of animation is that it works best on horizontal, flat surfaces. It will look stretched on vertical parts of objects. However, it works perfectly with horizontal planes, so it does a great job of animating the surface of a water plane, or the shape of clouds.

- Read the tutorial entitled [Making Waves](#) for a description of how to animate water.
- Read the tutorial entitled [Animating Clouds](#) for a description of how to animate clouds.
- Read the tutorial entitled [Moving Caustics](#) for a description of how to create and animate underwater caustics.

Material Velocity Animation

The second type of material animation is very similar to the first. It offers you the possibility of moving the origin of the material in time. That way, you can not only create waves that undulate, but also give an overall movement to the waves, like if they were approaching a shore. Creating a material velocity also adds the **Time Dependent Material** notice to the caption of the *Material Editor*. Also, if you check the *Timeline*, you will notice that the material is now listed at the bottom of the list of animated objects. Animated materials always appear at the bottom of the list. No keyframes are available for that material, because it is the same material that is being modified by time (read below to find out how material keyframes are created).

To create a material velocity, switch to the **Effects** tab of the *Material Editor*, and enter the X, Y and Z components of the **Velocity of the material origin**. Alternately, you can drag the current time slider to a new time, and enter X, Y and Z components into the **Origin of material** fields. VUE will automatically compute the corresponding origin velocity.

Note:

Material velocity settings override Material Surface animations. To maintain the effects of Material Surface animation, you should enter 1 as the Z velocity component. Consequently, defining a velocity of (0, 0, 1) is strictly equivalent to doing a Material Surface animation.

If you do a complete material animation (see below), you can also change the material velocity with time...

- Read the tutorial entitled [Making Waves](#) for a description of how to animate water.
- Read the tutorial entitled [Animating Clouds](#) for a description of how to animate clouds.



Complete Material Animation

Using this type of animation, you can define material keyframes that will be interpolated by VUE to produce smooth blending from one material keyframe to the next.

Creating a complete material animation adds the **Animated Material** notice to the caption of the *Material Editor*.

To create a complete material animation, drag the current time slider to the time of the new material keyframe and modify the material. A message will appear giving you the option to animate the material. Click **Yes**. The material is now animated. If you check the *Timeline*, you will notice that the material is now listed at the bottom of the list of animated objects, and that the new keyframe has been added at the current time (animated materials always appear at the bottom of the list). The keyframe is also added to the material property of the animated objects that use the material.

If you move the current time slider, you will notice that the preview of the material is re-rendered to display an updated preview of the material at the current time. The settings in the *Material Editor* are updated to display the settings of the keyframe immediately before the current time.

You can select, move and delete material keyframes just the same as with other object property keyframes.

Read the tutorial entitled [Stoned Frog](#) for an example of an effect that can be achieved using complete material animation.

SmartGraph Material Animation

This is the most complex type of material animation – and also the most powerful. It involves diving into the arcanae of material creation, and driving one or several material parameters using functions that depend on time. That way, you can achieve extremely advanced material animations that could not be achieved using other types of animation. For instance, you could animate the density function of a volumetric material to create swelling smoke effects using the *Function Graph*.

Read the tutorial entitled [Dying Flesh](#) for an example of effects that can be achieved using *SmartGraph* material animation.

This type of animation can also be combined with the other types of material animation for totally amazing visual effects.



Published Parameters for Animated Materials

If you have selected to publish certain parameters for the animated materials, they also appear in the *Timeline* as well as the *Material Editor* (see [here](#) for more information about *Published Parameters for Materials*).

This is particularly useful to animate an EcoSystem population. If you publish origin of a greyscale map and connect to the density, the EcoSystem will be animated correctly.

Animated EcoSystem Population

Dynamic EcoSystems can now vary with time. You can create an animation in a scene using a Dynamic EcoSystem (over a terrain, for instance), and make one or several distribution settings depend on time. If one or several parameters (density, overall color, overall scaling, for example) of the EcoSystem material depend on time, the dynamic population will be recomputed at each frame of the animation, which will give a new kind of animation for EcoSystems.

This can be achieved by editing the corresponding functions in the *Function Graph* – for example by loading an animation map and using it as density, or by editing the material at a non-zero time and answering **Yes** to the question about animating a material. The population will then be recomputed at each frame.

It should be noted that animated EcoSystems will not work well if the density varies smoothly. Instead, the density needs to vary in steps, since flickering may occur due to the fact that plants are being added and removed randomly which could cause some jumping of the population.

Overall color of instances can become animated in the same way.

Please read the tutorial on [Animated Dynamic EcoSystems](#) for an illustration of this topic.



Animating the Atmosphere

You can animate every aspect of a scene in VUE, and that includes the atmosphere. In this section you will find out how this is done.

Basically, there are three different aspects of an atmosphere that can be animated: the atmosphere itself (including sky dome colors, fog and haze, lighting conditions, etc.), the clouds in the sky, and the sun.

Atmosphere Keyframes

Atmospheres are animated by interpolating standard atmosphere settings. That includes all the settings that are not directly connected to cloud materials or to the sun. Please read the section on [Atmospheres](#) for complete details on these settings.

To animate the atmosphere, drag the current time slider to the time where you want to create the new keyframe, and open the *Atmosphere Editor*.

Now modify the atmosphere as required. The atmosphere automatically becomes animated, and the **Atmosphere** keyword appears at the bottom of the list of animated items of the *Timeline*. This keyword always stays at the bottom of the list, so that it can be located easily. An atmosphere keyframe is also added at the current time.

Atmosphere keyframes can be selected, moved and deleted just as other animated object property keyframes.

VUE automatically interpolates the settings in the atmosphere keyframes to produce smooth transitions from one keyframe to the next.

You will find ready-made animated atmospheres in the ***Animated*** collection of atmospheres of the Visual Atmosphere Browser.



Animating Clouds

Besides animating the atmosphere itself, you can also animate the clouds to produce amazing effects such as clouds drifting in the wind, slowly changing shape, or growing increasingly thicker with time!

This is done by animating the cloud material. Using the **Clouds** tab of the *Atmosphere Editor*, select the appropriate cloud layer, then double-click on the cloud material preview to open the *Material Editor*. Use any of the material animation method described in the preceding section about [Animating Materials](#).

When you animate clouds, the Atmosphere keyword in the list of animated items becomes a folder. Animated clouds are automatically appended to this folder.

Other useful parameters used to animate clouds are contained in the **Cloud animation** group.

Using **Direction** and **Velocity** controls, you can make your clouds drift in the wind! The **Rate of change** control is used to set the evolution rate of the cloud layer (whether the shape of the clouds changes slowly or rapidly).

Note:

Those settings remain constant during the animation and can only be set for the first keyframe.

Read the tutorial entitled [Animating Clouds](#) for a nice example of creating a sky with animated layers of clouds in it.

Animating the Sun

The sun is animated as any other animated object in the scene. You can animate the direction of the light, the color of the light and the softness of the shadows cast by the sun.

To animate the sun, use any of the tools described in the section on [Animating Objects](#).

The animation of the sun is done independently from the animation of the atmosphere. The sun will be listed among other animated objects in the scene.

You can also animate the color of the light, and the balance of the light, using atmosphere keyframes (see above).



Working with Animations

Pasting Animation

Pasting animation is a convenient way of applying to one object the animation settings of another object. A typical use could be to copy the animation path of one object onto another one.

To use this command, select an animated object and copy it using the **Edit | Copy** command in the main menu. Then, select the object to which you want to apply the animation settings, and select the command **Edit | Paste Animation** from the same menu. Animation is pasted in such a way that the position, size and orientation of the object at the current time are preserved.

Note:

The **Paste Animation** command attempts to preserve as much of the original animation as possible. For instance, if you paste the animation of a plant onto a sphere, only path, orientation and size will be pasted. But if you paste it onto another plant, wind animation will also be pasted!

Destroying Item Animation

Destroying Object Animation

To destroy the animation of an object, you can either:

- press delete after having selected the object in the list of animated items of the *Timeline*, or
- select **Not animated** from the motion type drop-down list in the **Animation** tab of the *Object Properties* panel (see [here](#)).

The object is removed from the list of animated items of the *Timeline*.

Destroying Material Animation

To destroy the animation of a material, you must:

- If the material is a complete animation, select all keyframes of that material, and delete them.
- If the material is a surface animation, open the *Material Editor* for that material



and uncheck the **Animate material surface** box.

- If the material is a velocity animation, open the *Material Editor* for that material and reset all velocity values.

The material disappears from the list of animated items. If the material is a combination of different types of animation, you will have to destroy all these animations before it is removed from the list. Read [here](#) for details on the different types of material animation.

Destroying Atmosphere Animation

To destroy the animation of the atmosphere, you must select all keyframes of the atmosphere using the *Timeline*, and delete them.

The atmosphere keyword disappears from the list of animated items.

If the atmosphere comprises animated clouds, you will have to destroy the animation of the clouds using one of the aforementioned methods before the Atmosphere folder is removed from the list.

Shifting the Start of an Animation

You can make the animation of a property begin anytime you like by moving the corresponding keyframes. It doesn't have to start at null time.

To do that, select all the keyframes of the property animation you want to shift in time, and drag the keyframes (or Control drag the first keyframe). You can drag the first keyframe to a positive time, or even to a negative time.

The cool thing about this is that you can start a repeating animation when you like. Once it has started, the animation will repeat indefinitely. This is a good way of dephasing identical animations so that they don't look identical. Read the tutorial [Animating a Fish](#) for an example of doing this.

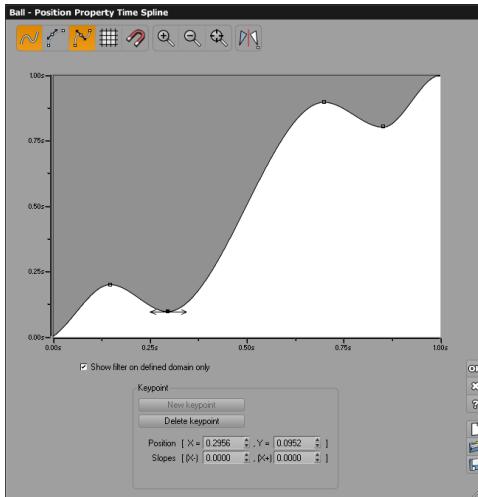
Changing the Duration of an Animation

You can change the duration of the entire animation using the menu command **Scale Animation...** and inputting the new duration of the animation. The animation of individual objects will be automatically scaled to match the requested duration.

You can also change the duration of animation of individual objects using the [Animation Toolbox](#).



Using Time Splines



Time Spline Editor

Time splines are an incredibly powerful tool for whoever wants to gain full control over the animation.

Basically, what time splines do is allow you to control precisely the flow of time. Thus, you can accelerate or slow down events to meet your wildest requirements. You can even invert the flow of time and make your animation play backwards!

For instance, by successively inverting the flow of time, you can make an animation repeat without having to use the [Repeat modes](#). This lets you control how many times an animation repeats, as well as the exact way it repeats...

You will find a set of interesting time splines in the *Time splines* collection of the Visual Filter Browser. Read the [Drop and Bounce](#) and the [Flickering Lights](#) tutorials for examples on how to use time splines.

Time Splines are modified using the *Time Spline Editor*. To edit a Time Spline, simply select **Edit Time Spline** from the popup menu that appears on top of the Time Spline view inside the *Timeline* (or **Control** click on it). The *Time Spline Editor* appears.



Editing Time Splines

The tools you use to modify time splines are very similar to the ones you use to edit filters (see [here](#)). The file format for saving time splines is actually the same as that of filters, and data is interchangeable between both. This is why the preset time splines are placed in the *Filters/Time Splines* collection together with other filter presets.

To open the *Time Spline Editor*, either click on the time spline with the **Control** key pressed, or select **Edit time spline** from the time spline's popup menu.

The *Time Spline Editor* can be resized if you need a more detailed view of a given part of the time spline.

On top of the grid (if it is displayed) you will notice thin lines. These lines indicate the position in time of the keyframes of the animated property, and can be used for reference. The solid vertical line indicates the Current Time.

The range of time covered by the time spline always starts at the first keyframe of the property's animation, and ends at its last keyframe.

Standard time flow is achieved with a slope of 1. If you increase that slope, you are making time flow faster, so you are actually accelerating the animation. If the slope is less than 1, you are making time flow more slowly, so the animation will be slowed down. If the slope becomes negative, time flow is inverted, and the animation will play backwards.

The Curve

The curve is the large area that sits in the middle of the editor, just below the toolbar. This area displays a curve representing the profile of the time spline. You can zoom in and out, and pan the view using standard commands (Right mouse drag to pan, Ctrl + Right mouse drag to zoom).

Time splines are built from **Key Points**, joined together by straight lines or curves. You can modify a Time Spline by adding, moving or deleting key points. The key points are figured by small round handles (🔲) on the curve. These handles appear as soon as the mouse cursor is placed above the curve. All time splines have a key point on the right edge (the corresponding handle can only be moved vertically).

Keyframes can also be added by clicking on the animation curve at the appropriate spot.

Smooth Time Splines

VUE offers two types of time splines: standard (linear) and smooth.


Linear time splines are generated from segments while smooth time splines are generated



from cubic curves.

You can switch from linear to smooth time splines, by clicking the **Smooth curve** icon in the toolbar.

The behavior of a smooth time splines is identical to that of a linear time splines except that you can change the slope of the curve around the key points, yielding a smoother -round- profile.

To modify the slope around a key point, select the key point by clicking on its handle () , or by typing its horizontal position in the **Position X** box. The **Slope** boxes now indicate the slope to the left and to the right of the key point. Type in new slope values. If you selected the handle by clicking on it, the tangents to the curve will appear. You can drag the ends of the tangents to modify the slope.

Selecting the **Smooth joint** icon will ensure that the slope is the same on either side of the key point (the default). If you want to have a different slope on either side of the key point (e.g. to create a crease in the curve), you must deselect this option and then modify the slope.

Toolbar

The *Time Spline Editor's* toolbar is the collection of icons at the top of the editor. The meaning of these icons is as follows:



Smooth curve: this is a toggle icon. If the icon is blue, the time spline is built from straight lines; if it is orange, the time spline will be built from curves. Click on the icon to change the type of time spline.



Auto-tangents: this is also a toggle icon; it is only available when the time spline is smooth. If the icon is toggled, the tangents at newly added key points will be computed automatically in order to modify as little as possible the overall shape of the curve. If you drag a key point when this mode is active, the tangents will be modified dynamically so as to minimize the deformation of the curve.



Smooth joint: this is also a toggle icon; it is only available when the time spline is smooth and a key point is selected. If the icon is toggled, the slope on either side of the key point will be the same, ensuring that the resulting curve doesn't exhibit any sudden changes in slope around that key point. If you click the icon, it will become blue, meaning that the slope on either side of the key point can be modified



independently, resulting in a crease in the curve.



Show grid: this is a toggle icon. When it is orange (enabled) a grid will be displayed on top of the curve. This grid can be used for reference when building a time spline.



Snap to grid: this is a toggle icon. When snapping is on (the icon is toggled), key points will be automatically “attracted” to the grid or nearby keyframes when you approach the mouse cursor from the grid/keyframe. This is useful for setting up time splines with “rounded” values.



Zoom in: click this icon to display a zoomed view of the time spline. This lets you edit detailed portions of the time spline.



Zoom out: click this icon to zoom out of the view of the time spline. This lets you visualize a larger portion of the time spline.



Reset pan/zoom: click this icon to reset the view of the time spline so that the time spline fills up the entire graph exactly.



Flip Vertical Axis/Flip Horizontal Axis: this flips the axis of the graph either horizontally or vertically.

New, Load, Save

Pressing **New** will reset the time spline by deleting all key points.

Press **Load** to load one of the sample time splines using the Visual Filter Browser.

Press **Save** to save the current time spline in a stand-alone file, for use in future scenes. Saved time splines will appear in the Visual Filter Browser like any other of the predefined time splines. By default, time splines are placed in the **Filters** subfolder. This means that they will appear in the **Personal** collection inside the Visual Filter Browser.



Adding Key Points



To create a new key point, you can either:

- double-click in the area where the curve is drawn. The new key point is created at the point you clicked. The curve is redrawn to use the new key point.
- click on the curve where you want the new key point; the coordinates of the clicked point appear in the **Position** boxes; you can edit them if required. To create the new key point, press the **Add key point** button. The curve is redrawn.
- type the coordinates of the new key point in the **Position** boxes, then press the **Add key point** button. The curve is redrawn.

You can't create two key points at the same horizontal position.

Modifying Key Points


To modify a key point, you can either:

- click on the key point's handle () and drag it with the mouse button pressed. If you press Control as you drag the cursor, the movement will be constrained to the closest axis. Each key point must stay between the previous one and the next one. When you select a key point, you can jump to the next one by pressing Tab (Shift Tab jumps to the previous). The selected key point becomes black. You can also modify the position of the key point by using the Up/Down and Left/Right arrow keys.
- click the handle () of the key point you want to modify. The handle becomes black, and the **Position** indicated is now the position of the key point. Type the new position of the key point.
- type the horizontal position of the key point you want to modify in the **Position X** box, then indicate the new vertical position of the key point.

Note:

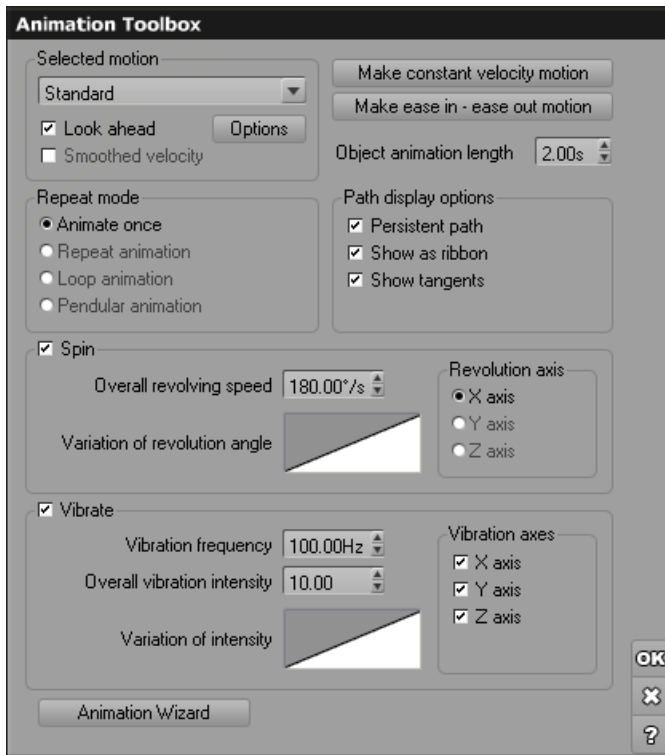
You can't move the horizontal key point position using this method.

Deleting Key Points

To delete a key point, click on the handle () of the key point you want to delete, or type its horizontal **Position** in the box **X**, then press the **Delete key point** button. When you select a key point, you can jump to the next one by pressing Tab (Shift Tab jumps to the previous). You cannot delete the right-most key point.





Animation Toolbox




The Animation Toolbox

The *Animation Toolbox* provides a convenient set of tools to process the animation of your objects. It operates on all the animation properties of an object.

To open the *Animation Toolbox*, you can either:

- Click on the  button in the **Animation** tab of the *Object Properties* panel (read [here](#)). If several objects are selected at the time you press the  button, the Toolbox will operate on all of these objects. If all the selected objects don't have identical settings, these settings will be left blank.
- Select the menu command **Object | Animation Toolbox**. If several objects are selected, the Toolbox will operate on all of these objects. If all the selected objects don't have identical settings, these settings will be left blank.



- Click on the  button of the object you want to open the *Animation Toolbox* for, in the *Properties Timeline*. The *Animation Toolbox* will only operate on the corresponding object.


The **Selected motion** drop-down list lets you select a *motion type*.

Press the **Options** button to display the *Motion Options* dialog and customize the sensitivity of your object to its motion.

The last two controls in the Selected motion group set the **Look ahead** and **Smoothed velocity** properties. These checkboxes are automatically set or reset when you change the type of motion of an object. By overriding these settings you can increase the variety of types of motion available. You can read more about these properties below.

Use the *Repeat mode* group to select the type of **Repeat mode** for the object. The default is **Once**.

The **Make constant velocity motion** and **Make Ease in – ease out motion** buttons automatically reorganize the keyframes of the position property of your object to produce a constant *velocity* motion, or an ease in – ease out motion respectively. The effect of these buttons is identical to setting these options in the *Animation Wizard*.

The **Path display options** group is a mirror of the path options available in the *Properties Timeline* (the  toggle-buttons). The effect of each of these options is explained in the section about the *Timeline*.

Press the **Animation Wizard** button to display the Wizard.

Look Ahead

The Look ahead property is a neat property that makes objects always point in the direction in which they are traveling. You can modify the direction in which the object points relative to that Look ahead direction using the standard rotation tools. You can even animate this relative rotation.

Smoothed Velocity

When this property is set, the velocity of the animated object is automatically processed to ensure smooth acceleration/deceleration between way points. This ensures perfectly smooth motion, and, better still, smooth accelerations!

However, there are some cases when it is not possible to smooth the acceleration of the object, because of very sudden changes in velocity. When this happens, VUE displays a straight line heading away from the path. To correct the problem, you should make the velocities on either side of the faulty way point closer one to the other.



Changing the Duration of an Animation

The **Object animation length** control automatically repositions the keyframes of your object so that the total duration of the object's animation is the time that you indicate (the time lapse between the first keyframe encountered in one of the object's properties, and the last keyframe encountered).

When you create or move a keyframe, its position in time always corresponds precisely to that of a frame of the animation. However, when you resample keyframes using this tool, they may no longer occur at the precise time of a frame of the animation. This is done to maintain a strict equivalence between the initial animation, and the animation with the new duration. If you move a keyframe, it will snap back to the time of a frame.

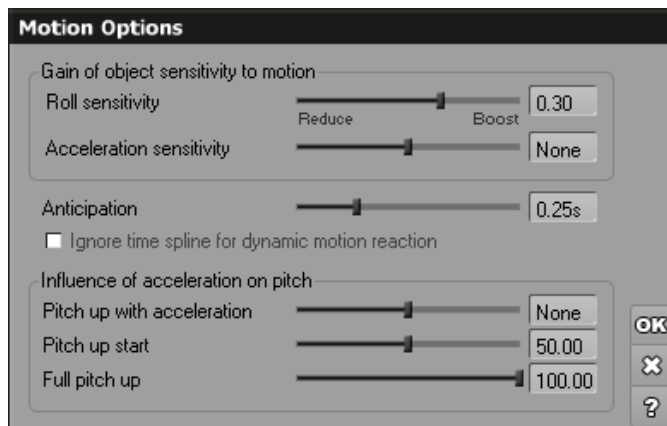
Note:

You can change the [duration of the entire animation](#) by using the **Scale Animation...** command.

Spin and Vibrate

The settings to control the amount of spin and vibration in the animation of the object are identical to those found in the *Animation Wizard*. Please turn [here](#) for details on these effects and the way they are adjusted.

Motion Options Dialog



The Motion Options dialog Airplane Motion

This dialog lets you customize the sensitivity to motion of your animated objects. It is



accessed by pressing the **Options** button in the *Animation Toolbox*.

If your objects tend to “over-react” or not react sufficiently to motion (this happens when you build scenes at unusual scales), this dialog can help solve the problem (you may like to read the troubleshoot entitled [Objects Overreact to Motion](#) for a discussion on this topic).

The default **Roll sensitivity** is 1. To reduce the amount of Roll in an object animation, reduce the corresponding sensitivity by dragging the slider to the left. To boost Roll, drag the cursor to the right. You can revert the effect of roll by changing the sign of the value (e.g. use a negative value where the value was positive).

The default **Acceleration sensitivity** is 1. To reduce the amount of dipping (e.g. **Helicopter** motion) or rising (e.g. **Motorcycle** motion) in an object animation, reduce the corresponding sensitivity by dragging the slider to the left. To boost it, drag the cursor to the right. You can revert the effect of acceleration by changing the sign of the value (e.g. use a negative value where the value was positive).

You can vary sensitivities up to a factor of 10. Although this should cover most requirements, there may be special cases when this is not enough. You can go over this limitation by entering a value directly in the edit fields.

Anticipation: this setting controls the amount of anticipation in the reaction to motion. In reality, an aircraft will start to bank before the turning actually takes place. This setting lets you control this effect. Bigger anticipation values will produce smoother reaction to motion, whereas short anticipation will result in jolty reaction to motion. Keep in mind that an aircraft, for instance, rarely banks completely in half a second!

Ignore time spline for dynamic motion reaction: when this option is checked, the time spline won't be taken into account when processing the motion of your objects. This is useful, for instance, when defining a pendular motion time spline with the look ahead property set, and you don't want your object to turn around before it moves back.

Influence of Acceleration on Pitch

The settings in this group let you control the sudden surge that can take place in reaction to acceleration on certain types of motions (typically Motorbike). With this type of motion, when the acceleration goes beyond a threshold, the motorbike looks up suddenly as if it were raising on its rear wheel.

Pitch up with acceleration: this controls how strong the effect is.

Pitch up start: this controls the acceleration level that starts triggering the pitch-up effect.

Full pitch up: this controls the acceleration level at which full pitch up is achieved.



Linking and Tracking

Linked Hierarchies

VUE supports a type of hierarchical animation entitled forward dynamics. Forward dynamics is a feature that greatly simplifies the animation of complex structures. It lets you build a hierarchy of objects by linking some objects to others (the link parents). When an object is linked, modifying the link parent automatically modifies the linked object.

To set links, you use the **Animation** tab of the *Object Properties* panel.

You can link cascading objects to create complex hierarchical structures. For instance, the [Animating a Fish](#) tutorial or [A Steam Power Train](#) give examples of complex hierarchical animations.

Once you have created a link, you can modify the relative position, orientation and size of the linked object in a standard way. However, if you modify the link parent, the linked object will be affected in some way.


You can link objects to a member of a group, but you cannot link a member of a group to anything else than the group itself. This has an inconvenience: if you want to manipulate a whole hierarchical structure, you will have to manipulate only the topmost parent. All other objects linked to that parent will follow.

If you try to link objects to other objects that depend (directly or indirectly) on that object, you will create a deadlock in the hierarchical structure. VUE detects this situation and displays a warning before it destroys the bad links.

You can animate the relative position, orientation and size of linked objects. This animation will be based on the current conformation of the parent object.

You can even create links to objects that have the **Look ahead** property (see [here](#)), however, it is not possible to create partial links to such objects (see below).

To break an existing link, either:

- select “No link” in the **Animation** tab of the *Object Properties* panel, or
- click on no object (or on an empty part of a toolbar) after having pressed the  button in the **Animation** tab.



Tracking Objects


Using the automatic tracking tool (read [here](#) for instructions on activating this), you can instruct an object to always point in the direction of another object (the track parent). The tracking object points directly at the center of the tracked object.

Once the tracking has been activated, moving the track parent will modify the tracking object so that it keeps pointing at the track parent. Just the same, moving the tracking object will also modify its orientation so that it keeps pointing at the track parent.

You can decide which axis of the object will be pointing at the tracked parent by using the controls in the [Animation](#) tab of the *Object Properties* panel.

You can also rotate the tracking object relative to the direction of tracking. This relative rotation will be maintained if you move the tracked object or the tracking object. However, you cannot animate the relative rotation. Since the tracking object points at the center of the tracked parent, rotating the tracked parent does not affect the tracking object (unless you have defined a pivot for the tracked parent).

To break an existing track relation, either:

- select “No track” in the **Animation** tab of the *Object Properties* panel, or
- click on no object (or on an empty part of a toolbar) after having pressed the  button in the **Animation** tab.

Loose Linking and Tracking

Loose linking and tracking is the ability for VUE to simulate the approximate reactions of a real-world response, and thus avoid the stiff, automatic, and usually jolty effects of linking and tracking in animations.

Note:

Because loose dynamics need to take into account the globality of the animation, you will not see their effect in the interface.

Use the **Response** slider to control the quality of the linking or tracking. Setting the slider to roughly $\frac{1}{4}$ of the range corresponds to a typical human reaction time.

You can also customize this response using the *Forward Dynamic Options* dialog (see below).

Note:



loose dynamics only act on linked object position and tracking orientation. If you want



other object properties to be “loosely related”, you will have to customize the object’s graph.

Forward Dynamics Options

The *Forward Dynamics Options* dialog lets you control the linking between objects, as well as customize the accuracy of the response in case you are using loose dynamics.

To access the *Forward Dynamics Options* dialog, select the menu option **Object | Forward Dynamics Options** or long-click/right click on the **Pick link object** or **Pick tracked object** icons (resp.  and .

Use the **Tracked object** and **Link to** drop-down list boxes to select the tracked and link objects.

Partial Links

You can decide how the linked object will be affected by modifications made to the link parent by checking or unchecking the link options boxes of the **Animation** tab in the *Object Properties* panel.

If you uncheck one of the linking options, the corresponding attribute will not be inherited from the link parent. This type of link is known as a Partial Link

- **Position:** if you uncheck this linking option, the linked object will no longer move with the link parent, but it will keep rotating and changing size as the link parent does.
- **Rotation:** if you uncheck this linking option, the linked object will no longer rotate with the link parent, but it will keep moving and changing size as the link parent does.
- **Size:** if you uncheck this linking option, the linked object will no longer be resized with the link parent, but it will keep moving and rotating as the link parent does.
- **Join:** this is a neat linking option that disables the joining of the linked object’s center to that of the link parent if it is unchecked. This means, for instance, that instead of rotating around the link parent’s center, the linked object will rotate around its own center (while still moving with the parent, if the parent moves...). The same happens with size.

There are two conditions when the use of partial links is not possible:

- You cannot disable one or more linking options for an object that is linked to a parent object having the **Look ahead** property (see [here](#) for details). However, you can disable the linking options prior to declaring the link.



- It is not possible to define a partial link to an object that is itself tracking another one (read further for details on tracking objects).

VUE will issue a warning and cancel the operation.

The tutorial entitled [A Steam Power Train](#) makes an extensive use of partial links. Reading it will get you a better in depth understanding of partial linking.

Loose Dynamics

The **Response** slider replicates the same slider on the **Animation** tab of the *Object Properties* panel (see [here](#)), which lets you control the accuracy of link and track response. A value of 0% indicates that the response of the forward dynamics engine is “perfect”, meaning that there is no error introduced by loose dynamics. Increasing the value gradually “slackens” the accuracy of response to make it more realistic.

You can customize this response by ticking the **Custom response** checkbox. When this option is selected, the following custom response settings become available:

- **Delay:** this settings controls the typical reaction time between a change of the “master” object (the tracked or link object) and the object attached to it. This is similar to the response delay caused by human reflexes (0.2 seconds).
- The three following settings (**Proportional**, **Integral**, **Derivative**) are the three parameters of a standard [PID Controller](#).

Object Graph

As soon as you activate the loose dynamics engine by selecting a non-perfect response in the **Animation** tab of the *Object Properties* panel, an Object Graph is automatically created for that object (see [here](#)). If you edit the graph, you will notice that VUE automatically adds the required nodes to calculate the position, orientation and size of the attached object, based on the properties of the master object.

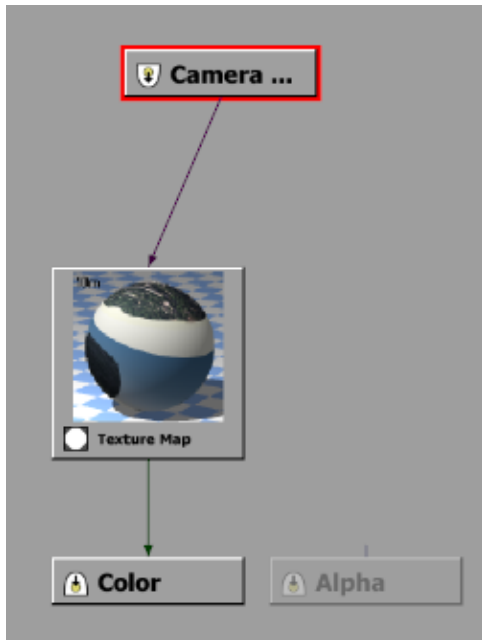
By editing this graph, you can easily create more elaborate types of linking and tracking.

Note:

Object graph-based linking and tracking may not yield exactly the same results as standard forward dynamics, even when **Response** is set to be perfect. The only way to revert to standard forward dynamics after activating loose dynamics is to destroy the link or track relationship and re-create it.



Camera Mapping



Function Graph – Camera Projection node

Camera mapping projects a bitmap through a given camera over some geometry. When rendering an animation where camera movements are of limited amplitude, it can be worth rendering a single high quality static image of a subset of your scene (generally background parts that don't change much through the animation), then removing all corresponding components from your scene, and project this render over a simplified underlying geometry to accelerate subsequent frames rendering. The projection would be done through a secondary static camera that matches the one used to perform the initial render. As long as the viewing camera doesn't shift too much from its original position, the projected render can efficiently replace corresponding complex geometry. Please read the tutorial on [Camera Mapping](#) for an illustration of this topic.

Camera mapping is available through a dedicated projection node in the *Function Graph*. It is therefore material-specific, and should be used to control the color output of the corresponding material:



Edit the material of the object(s) over which your bitmap should be projected

- Open its color production *Function Graph*



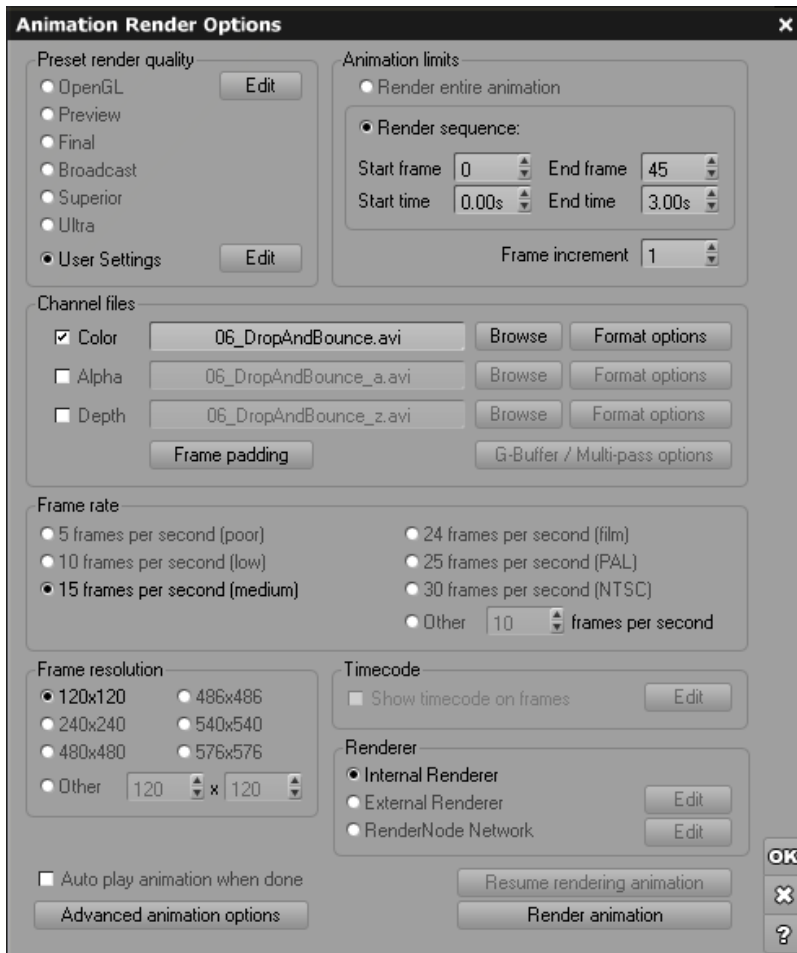
- Create a standard “Texture Map” node where you can load your bitmap
- Connect the texture map node output to the color output
- Replace the default “UV coordinates” input node with a “Camera projection” node
- Select camera through which projection should be performed
- Set the aspect ratio to be the same as your bitmap (the camera needs this information to compute the proper projection)

When performing a reference render that is to be projected via camera mapping, a few rules should be followed to allow for a proper integration without perspective and/or color shifts:



- Always disable all post render effects like natural film response, automatic exposure, lens glare or post processing. These effects should only be applied to the “final” render that uses your reference render via camera mapping (or these would be applied twice!)
- The camera used for the reference render should be stored as is, used as the camera mapping projector, and remain static. So if your viewing camera is animated, it shouldn’t be used as the projector since your reference render is only valid for a given viewpoint.
- Don’t forget to match the Camera projection node aspect ratio to your reference render picture, so the projected render aspect is preserved.
- Use the new **Ignore lighting** () and **Ignore atmosphere** () buttons to disable any external influence over your camera mapping material, as these effects are already included in the reference render.

Rendering the Animation





Animation Render Options dialog

To render an animation, open the *Timeline* and press the **Render animation** icon (). The *Animation Render Options* dialog pops-up. If you activate the alternate action of this icon () , the *Advanced Animation Options* dialog will appear instead. See below for a description of these dialogs.



Animation Render Options

This dialog lets you control the rendering of your animation.

The **Preset Render Quality** group lets you select a [Preset Render Settings](#).

Note:

Motion Blur is only rendered with **Broadcast** or higher render settings. If you select **User settings**, you can fully customize the render engine by pressing the **Edit** button.

Use the [Frame resolution](#) group to indicate the resolution of the frames in the animation (Horizontal x Vertical). The horizontal and vertical resolutions are linked by the aspect ratio of the scene.

The **Animation limits** group lets you indicate the part of the animation that you want to render.

Render complete animation: if this option is selected, the complete animation will be rendered, starting at the beginning of the active part, and ending at the end of the active part (by default, this is the entire animation sequence, starting at 0 and ending at the last keyframe).

Render sequence: if this option is selected, you can enter manually the limits of the animation. You can either enter the limits using Frame numbers, or using Time. The time or frame counterpart is automatically updated. The values in these fields are initialized with the values of the start and end of the active animation part, as defined by the yellow line in the *Timeline* duration bar.

Frame increment: this setting lets you skip frames in an animation to reduce render time without affecting the frame rate. By default, the Frame increment is one, which means that all the frames in the animation are rendered. Entering 2 will render every other second frame; entering 5 will skip 4 frames after each frame rendered (rendered frames will thus be: 0, 5, 10, 15, 20...).

Channel Files

VUE can generate and save the three [channels](#) (Color, Alpha and Depth) of an animation. Using channels, you can easily composite VUE animations with other animations using an external compositing application. You can also generate full G-Buffer information for each frame, for maximum compositing information.

The **Channel files** group of controls let you select the destination files for the channel animations. By default, only the Color channel is saved (Alpha, Depth, G-Buffer and Multi-Pass channel files are disabled).



Animation File Formats

Use the **Browse** buttons to select you want to save the channel animations to, or to change the selected file format. VUE supports the following animation file formats:

- **AVI:** Audio Video Interleaved file format, compressed or uncompressed. Press the **File format options** button to display a standard Codec selection dialog. Using this dialog, you can control how the AVI animation file is compressed.
- **M1V:** Mpeg 1 file format, compressed. Press the **File format options** button to display a standard options dialog. Using this dialog, you can control how the Mpeg 1 animation file is compressed.
- **M2V:** Mpeg 2 file format, compressed. Press the **File format options** button to display a standard options dialog. Using this dialog, you can control how the Mpeg 2 animation file is compressed.
- **BMP, PICT, JPG, GIF, IFF, PCX, PNG, PSD, TGA or TIFF:** set of stand alone pictures of the indicated file formats,
- **RLA, RPF:** set of stand alone Run-Length Encoded (RLA) or Rich Picture Format (RPF) files that contain all the channels of information stored in the G-Buffer (you need to enable G-Buffer rendering to use this option).

Press the **File format options** button to display the standard VUE **Picture Format Options** dialog. The frames of the animation are named after the filename you indicate, with the number of the frame appended to it (e.g. if you save to file *Anim.bmp*, frames will be named *Anim_000000.bmp*, *Anim_000001.bmp*, *Anim_000002.bmp*, etc.).

Frame padding: click this button to bring up the *Frame Name Options* dialog. Using this dialog, you can change the zero-padding of the frame file names.

G-Buffer / Multi-pass options: click this button to configure the creation and gathering of G-Buffer and Multi-pass rendering information while rendering the animation. This option is only available when the **Optimize last render pass** option in the *Render Options* dialog is deselected (see [here](#)). If you click this button, the *Multi-Pass Options (G-Buffer)* dialog will appear, letting you select which rendering components and masks to render.

Note:

Selecting this option doesn't mean the G-Buffer or Multi-Pass information will be saved in the animation file. You need to select the RLA or RPF file formats (for G-Buffer information), or multi-layer PSD file format (for Multi-Pass information) for this to happen.



Frame Rate

This group controls the number of frames that will be rendered for every second of animation. The higher this number, the more smoothly the animation will play back. But the longer it will take to render (and the larger the resulting file).

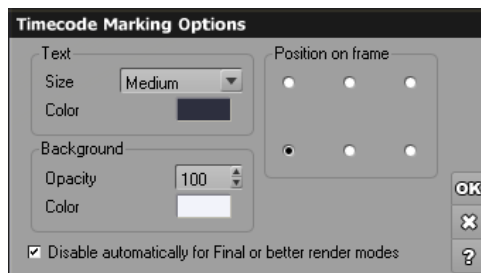
Bear in mind that the human eye is unable to isolate more than 24 frames per second. So there is no real point in rendering more than 24 frames per second (unless you are rendering for TV video, where there are synchronization concerns requiring an increased frame rate).

The default is 15 frames per second, and produces reasonable smoothness at a reasonable expense.

Frame Resolution

This provides a list of typical animation frame resolutions as well as the ability to define custom frame sizes. If the **Other** frame resolution option is selected, entering a resolution in one of the fields will automatically recompute the corresponding resolution for the other field (according to picture aspect ratio, provided that this aspect ratio hasn't been set to **Free (user defined)**). If you want to change the aspect ratio of your frames, press the **Edit** button in the **Preset Render Quality** frame to access the [Render Options](#) dialog.

Timecode



Timecode Marking Options dialog

The **Show timecode on frames** option will automatically add the frame's timecode on the rendered animation frames. If you are saving the frames as multi-layer PSD files, the timecode will be placed on a separate layer, so that it can be hidden in post work. The format of the timecode will be the same as the way the frames are identified in the *Timeline*. This can be changed by using the [Timeline](#) menu (**Show Time As** menu options).

Whenever this option is selected, the **Edit** button becomes active. Clicking on this button will display the *Timecode Marking Options* dialog, letting you configure the color and



location of the timecode on the frames.

Use the **Text** frame of the *Timecode Marking Options* dialog to select the size and color of the text that displays the timecode. Available text sizes are:

- **Small:** very small text, approximately 10 pixels high.
- **Medium:** this is the same as the typical text used in the application's interface. Approximately 13 pixels high.
- **Large:** this is bold text (the same as the dialog captions in the application interface). Approximately 14 pixels high.
- **Extra-Large:** this is very large text. Approximately 20 pixels high.

Double-click on the color control to change the color of the timecode text.

Use the **Background** controls to configure the background of the timecode text. You can adjust the opacity and the color of the background. This is useful to ensure that the timecode text is readable, whatever the frame colors.

Finally, use the **Position on frame** buttons to select the placement of the timecode text on the frame.

Disable automatically for Final or better render modes: when this option is checked, the timecode will be automatically removed when performing the final rendering (Final or better render quality). This can avoid wasting an entire final-render batch just because you forgot to remove the timecodes.


If you have this setting checked, the **Timecode** fields will be grayed out and unavailable when you select a render setting of Final or higher. If you wish to make changes to the **Timecode** stamp settings, perhaps have it render on a Final or higher quality render, switch your render mode back to Preview or OpenGL to enable the **Timecode** stamp, then the **Show timecode on frames** field becomes available so that you can edit it.

Renderer

This option lets you select the renderer to be used for rendering the image or animation:

- **Internal renderer:** select this option to use VUE's internal renderer. This is the best for quick renders that require interactive feedback (e.g. when rendering a quick preview).
- **External renderer:** when this option is selected, VUE will invoke a standalone rendering application that is installed together with VUE. This application will take care of the rendering. Because it is a separate application that is entirely dedicated to rendering, it doesn't have to deal with all the overhead of a graphical interface, and can consequently dedicate more memory to the actual rendering process. The caveat is that the scene has to be sent over to the standalone renderer so the time



it takes to actually start rendering is longer than when using the internal renderer, and also, because rendering is done by a separate application, you do not see the picture appear gradually on screen as it renders. You can configure the external renderer to either render on your computer, or on a network  of *RenderCows*. This option is particularly useful to speed up the rendering of very large pictures because it splits the render load across all available nodes on your network.

Note:

Network rendering is only available when you render to screen or to disk.

- **RenderNode network:** select this option to use a separate render farm administration tool and split the render load across a network of *RenderNodes*. Please turn [here](#) for details on the difference between *RenderCows* and *RenderNodes* and the *Network Rendering Options* dialog.

If you configured your external renderer to use [network rendering](#), the picture will automatically be chopped up into tiny fragments. The *HyperVue™ Network Rendering Manager* will then assign each fragment to a *RenderCow™*. The network manager collects the resulting picture fragments and reassembles them into the final picture. When you press **OK** to begin rendering, the scene is added to the list of queued jobs. If it is the first time you render across a network, the *HyperVue Network Rendering Manager* will appear, letting you configure network rendering.

Note:

Avoid using network rendering for quick renders, because the overhead of managing the render nodes and communicating over the network may actually result in longer render times.

Miscellaneous

Auto-play animation when done: select this option to automatically begin playing the animation in an external player when the rendering completes. This option is only available when rendering to an AVI file on Windows systems, and to a QuickTime MOV file on Macintosh.

Resume rendering animation: if this button is active, press it to resume rendering an animation that was stopped before it completed. VUE restarts rendering from the exact point where it was stopped (thus avoiding any loss of render time) and appends the new frames to the previous animation files (with the exception of compressed AVI files, because AVI doesn't support appending frames to a compressed video stream; the animation will be saved as Part2). To resume rendering an animation, you must use the same filenames as the ones used during the interrupted render. All render options are stored in the resume enabler files, so you don't need to worry about them.

Advanced Animation Options: press this button to display the *Advanced Animation Options* dialog (see below).



Closing the Dialog

Click **OK** to accept the changes and close the dialog. Click **Cancel** to cancel the changes.

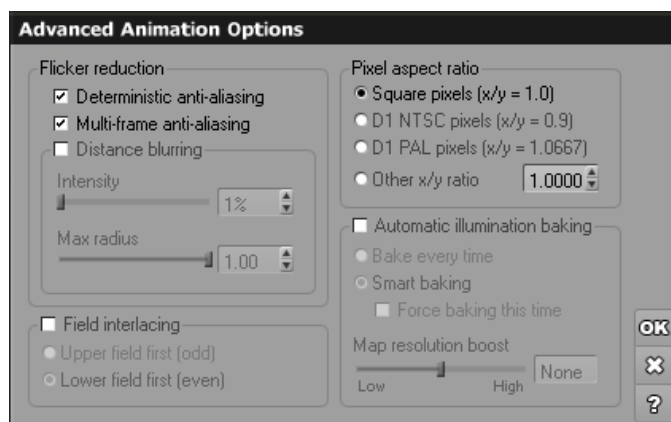
To accept the changes and render the animation with the new settings, click the **Render animation** button.

If you have interrupted a render, the **Resume rendering animation** button will be active. Click on this button to resume rendering the animation.

Note:

Any changes to the render quality will make resuming a render impossible.

Advanced Animation Options



Advanced Animation Options dialog

Using this dialog you can control advanced animation settings, such as flicker reduction, interlacing, pixel aspect ratios as well as illumination baking.

Flicker Reduction

The options in this group let you activate special algorithms to attempt to reduce the dreaded flickering that is so typical of computer graphics animations. While the ultimate solution to eliminate flickering is simply to increase anti-aliasing settings sufficiently, this has a tremendous impact on render times. The methods described below are hacks that will attempt to reduce flickering without having such a dramatic impact on render times:

Deterministic anti-aliasing: when this option is enabled (the default), anti-aliasing



rays are cast in random patterns that are repeatable from one frame to the next. This almost totally eliminates static noise, but may, under certain very specific circumstances, create patterns that would be noticeable to the eye. However, the impact on image quality is, at worst, barely noticeable. This is why it is recommended that you leave this option on by default.

Multi-frame anti-aliasing: when this option is enabled, VUE will compare the current frame to the previous and the next frame, and try to detect areas of strong flickering to concentrate more rendering samples specifically on those areas. This option requires that the last 3 frames be cached before actual completion of each new frame, and hence only works when rendering an animation. It may also produce a slight blurring of the frames.

Distance blurring: this option lets you artificially blur the frames in the animation. Because flickering appears oftentimes on parts of the scene that are far away from the camera – especially with the new EcoSystem technology where you can have very fine geometric details in the distance – this blurring option lets you control the amount of blur according to distance:


The **Intensity** setting controls the influence of the distance on the amount of blur. Low settings will blur only objects that are very far from the camera, while high settings will blur all objects in the scene equally.

The **Max. radius** setting controls the maximum radius of the blurring that is applied to the pixels in the frame (in pixels). You can use this setting in combination with the Intensity setting to fine tune the blurring of the frames.

In the end, the amount of flicker reduction you apply to your frames should be the result of a compromise between the amount of blur or flickering you can tolerate in your renders, and the amount of time you are willing to spend on the rendering of a particular project.



Field Interlacing



Field Interlacing turned on for a sphere rapidly moving left to right

Turn this option on to activate field interlacing. Field interlacing will render every other half of a frame, twice as often. This is due to the way video is played back on TV where the screen is refreshed by halves 60 times per second (50 times for PAL). Use this option to ensure perfect playback on TV – and only when rendering for TV. You can select which field will be the first using the **Upper field first** or **Lower field first** options. Do not use this setting if you are not rendering for playback on TV. Default is off.

Pixel Aspect Ratio

Unlike computer monitors, digital edition systems don't always work with square pixels. You can modify the pixel aspect ratio to render animations that will be compatible with these systems. When played back on a computer monitor, the animation will look squashed or stretched.

- **Square pixels:** this is the default setting, e.g. for computer monitors.
- **D1 NTSC pixels:** select this option if you are rendering for D1 NTSC media.
- **D1 PAL pixels:** select this option if you are rendering for D1 PAL media.
- **Other aspect ratio:** use this option to select an alternate pixel aspect ratio. Enter the desired x/y pixel ratio in the corresponding field.



Automatic Illumination Baking

Select this option to automatically bake the indirect lighting of all the meshes in the scene prior to rendering the animation. Turn [here](#) for details on the concept of illumination baking.

When this option is selected, all meshes that have not forbidden illumination baking will be baked (according to the options of this group and the per-object baking options – see [here](#)) before rendering the animation. This usually results in dramatic reductions of render times, at the expense of potentially very long preparation times.

You can adjust the way automatic illumination baking is handled using the options in this group:

Bake every time: when this option is selected, the illumination will be baked again each time you begin rendering the animation – whatever the current baking status.

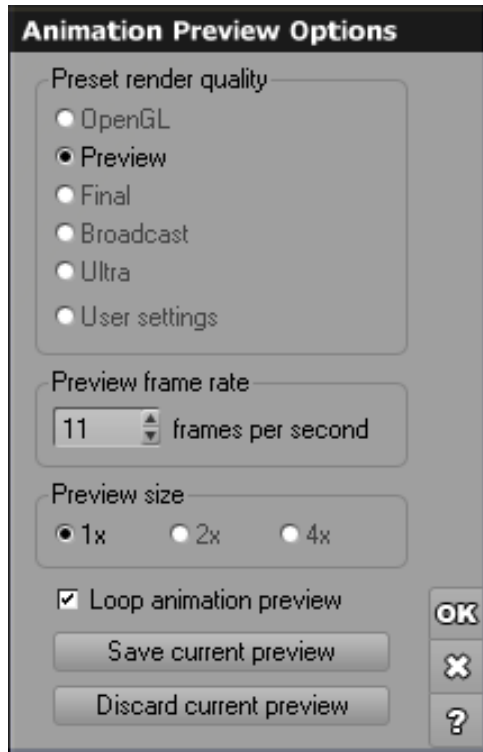
Smart baking: when this option is selected, VUE will check the baking quality of all the meshes in the scene and compare them to the desired rendering quality of the animation. If the current baking quality is greater than required, and if the lighting conditions have not changed, the illumination is not baked for that mesh. If the current baking quality is insufficient, or if VUE determines that lighting conditions have changed, a message will appear asking whether you wish to recompute illumination baking before starting the animation rendering.

Force baking this time: this option is only available when the Smart baking option above is selected. If you check this option, the illumination of all the meshes in the scene will be recomputed this time (the check is automatically removed after completing the baking).


Map resolution boost: this setting controls the overall quality of the baking process. The higher the resolution of the illumination maps, the greater the quality of the baking, and the more detailed the illumination. You can define a base illumination map resolution for each mesh in the scene. This base resolution should be such that, at any time during the animation, the illumination map's resolution will be sufficient to avoid visible artifacts. This setting “boosts” the resolution of the illumination maps of all the objects in the scene by a given boost ratio. This is particularly useful if you decide to increase the output resolution of your animation, because all you have to do is increase the boost factor accordingly. The boost factor works along the principle of octaves (+1 means double resolution, -1 means half resolution).



Animation Preview Options



Animation Preview Options

This dialog lets you customize the quality of the animation preview that is rendered when you press the  icon in the *Timeline*.

To open this dialog you must activate the alternate action of the aforementioned icon .

Use the **Preset render quality** group to select the render setting that will be used to render the preview (read [here](#) for details). Bear in mind that a preview render should be something fast, so you might want to avoid using high quality settings such as **Broadcast** or **Ultra**. The default is **Preview**.

Select the **Preview frame rate** of the animation preview. Default is 5 frames per second, but you may need to increase it to get smoother playback.

Select the **Preview size** which can be **1x**, **2x** or **4x**.



You can loop the preview if you check the **Loop animation preview** box. When this is checked, the animation preview will start again at time 0 at the end of the preview. Stop the preview by using the **Stop** button in the animation control bar.

Press the **Discard current preview** button to get rid of the current preview without having to render another one.

Recovering TMP Files from an Aborted Render

If you are rendering an animation to individual frames and it is interrupted for any reason leaving *.tmp* files in your target directory, there is a way to recover these files. VUE now generates a *recovery.cfg* file containing the necessary information for the recovery.

Just start VUE from the **Run** dialog in Windows or the terminal window on the Mac using this command line:

```
[PathnameofVueProgram] -t\enquote{[path to the tmp files folder]}
```

Examples:

Windows:

```
c:\ProgramFiles\e-on software\Vue 11 Infinite\Vue 11 Infinite.exe -t\enquote{c:\MyDocuments\e-on software\Vue 11 Infinite\Pictures}
```

Mac:

```
/Applications/Vue 11 Infinite/Vue 11 Infinite.app -t\enquote{/myusername/Documents/e-on software/Vue 11 infinite/Pictures}
```

This will convert any *.tmp* files to the finished format you intended. Once finished, Vue will continue to startup as normal.



Technical Notes

Rotating Look Ahead and Track Objects

When you rotate (either manually or using the Numerics panel) an object that has the Look ahead property, or that tracks another object, you must understand that this rotation is always considered to be relative to the orientation of the object as dictated by motion (for Look ahead object) or by the track parent object (for tracking objects).

This is why, when no relative rotation has been defined, 0° rotation angles appear in the Numerics rotation fields, although the object is not oriented that way.

The relative rotation will be maintained along the path of the object (for Look ahead objects) or if you move the track parent (for tracking objects). You can also animate this relative rotation to achieve really cool effects.



Synchronizing Cameras and Lights

VUE offers a powerful synchronization feature in order to make camera, light (or object) animation data exactly match the animation of a scene created in another professional 3D application. Through the use of specific plug-ins, you will have the ability to export animation data from this other application. This data will subsequently be retrieved and used by VUE to synchronize specific components of your VUE scenes. This automatic synchronization feature lets you easily produce animations in VUE and composite them precisely with their counterpart created in your other 3D applications.

VUE offers other Import/Export options as well for your tracking information and animation.

Import Synchronization Data

You can import synchronization data using .fbx and Nuke Channel format (.chan) as well as VueSynch data (.dat). These options are found using the **Animation | Import Synchronization Data** menu command. .fbx and .chan formats are only available in xStream. **VueSynch** is available in VUE Pro versions. You can also import the camera in .abc (Alembic format).

After choosing which objects (among cameras and lights only) should have their animation data exported to VUE, use this function to import them. Camera path is imported as well.

If synchronization data contains animation data for a camera, VUE's main camera will automatically be synchronized with it. If you save your scene, the animation data will be saved with it, in order to be reused later.

Note:

If you refuse to retrieve synchronization data, this data will be destroyed. You will have to generate it again when you want to synchronize your scene.

If you find that only one keyframe is generated for the synch, check your Render settings in the application and be sure the first keyframe starts at 1 and not 0. Then all of the keyframes should be included.



VueSynch

VueSynch ships with VUE. It includes plug-ins to synchronize animation data with the following applications: Autodesk Maya, Autodesk 3DS Max, NewTek LightWave, Maxon Cinema 4D and Autodesk Softimage.

Installing Plug-Ins

Help files are available for each plug-in that explain how to install and use the various plug-ins for each supported 3D application:

- `VueSynch_MayaInstall.txt` for Maya
- `VueSynch_MAXInstall.txt` for 3DS Max
- `VueSynch_LWInstall.txt` for LightWave
- `VueSynch_C4DInstall.txt` for Cinema 4D
- `VueSynch_XSIInstall.txt` for Softimage

These files can be found on the Application CD, in the `Synchro\ Plugins` folder.

Generating Synchronization Data

From your favorite supported 3D application, you will use the corresponding synchronization plug-in to choose which objects (among cameras and lights *only*) should have their animation data exported to VUE. When you are done generating synchronization data, switch to (or launch) VUE. The existence of new synchronization data will be automatically detected and VUE will ask if you want to retrieve it. If you click **Yes**, the synchronization data will be loaded into VUE and used for synchronized objects. This flexible approach lets you touch up the animation in your other application, and then re-synchronize VUE at the touch of a button.

Alternatively, if you want to use the synchronization data on another system, you can save the synchronization data to file for later use. To load a synchronization data file, use the **Animation | Import Synchronization Data** menu command.

Note:

If synchronization data contains animation data for a camera, VUE's main camera will automatically be synchronized with it. If you save your scene, the animation data will be saved with it, in order to be reused later.

Note:

If you refuse to retrieve synchronization data, this data will be destroyed. You will have to generate it again when you want to synchronize your scene.



If you find that only one keyframe is generated for the synch, check your Render settings in the application and be sure the first keyframe starts at **1** and not **0**. Then all of the keyframes should be included.

Synchronizing Objects

Once synchronization data has been retrieved, you can synchronize objects from your VUE scene using this data:

- Select the object you want to synchronize,
- Go to the **Animation** tab of the *Object Properties* panel,
- Choose **Synchronized** motion type in the **Motion** type list box,
- Select the name of the source object in the **Synchronize with** list box. This will instruct VUE to use that source object's animation data for the currently selected object,
- Adjust the **Scale** factor that will be applied to original synchronization position values. This **Scale** factor is the same for every synchronized object in order to ensure homogeneous synchronization.

This is it! Your object is synchronized according to the animation data of the original scene. You can observe the results by moving the current time cursor in the Timeline and see your objects follow the same animation paths as in the original scene. If you want to de-synchronize a specific object, just change its **Motion** type.

Note:

The animation path and orientation of synchronized objects cannot be modified. You can still modify the animation of other animated parameters. For camera objects, not only position and rotation are synchronized, but also focal length and motion blur amount. Therefore, you won't be able to modify these parameters for a synchronized camera.

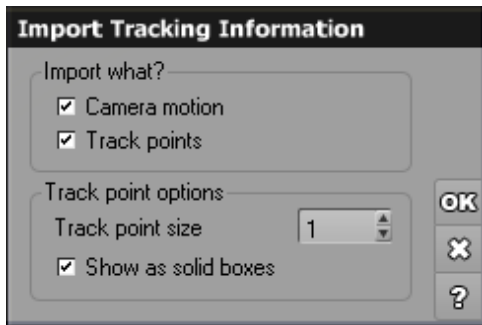
Once you have completed rendering of both animations (the original one and the one in VUE), you will be able to composite items rendered in VUE with your original animation very easily and precisely.

Keep in mind that VueSynch was designed to bring synchronized data from an application into VUE, not from VUE to another application. The conversion of animated VUE cameras into the host application through xStream does not take all parameters into account (the animation of focal, for instance). It is much more effective to use the original native camera instead since the VUE camera synchs to it anyway.

Actually, if you use VueSynch to bring the host camera into VUE, then reimport this scene with animated cameras, it will add a lot of keyframes to match the VUE animation so it is not recommended.



Import Motion Tracking Information



Motion Tracking Import options

You can easily import motion tracking information generated by dedicated software such as RealViz MatchMover, 2d3 Boujou or Anderson Technologies' SynthEyes. Motion tracking information is used to synchronize real footage with CG renders. Once you have imported the motion tracking information, you can create renders that will match the real footage. Usually, this real footage is placed in the background of the render.

Importing motion tracking information is similar to importing synchronization data: select the menu command **Animation | Importing Tracking Information** to import the motion tracking file. Supported file formats are **Max Script (.ms)**, **MatchMover (.rzml)** and **MatchMover Ascii Camera 3D Track (.rz3)**.

Once you have selected the file to import, the *Import Tracking Information* dialog will appear. This dialog lets you configure the import of motion tracking data.

Camera motion: select this option to import the camera motion information contained in the motion tracking file.

Track points: select this option to import all the track points that were used by the motion tracking software to determine the motion information. These tracking points can be used as a reference to place your CG elements relative to the elements in the real footage. Handling all the track points may slow down VUE slightly.

If you have opted to import the track points, the track point options become available:

Track point size: this option controls the size of the track points when displayed inside VUE. Track points appear as cubes that are only visible in the *3D Views*. They have the **Hide from render** option set (see [here](#)), so that they do not appear in the final renders.



Show as solid boxes: when this option is selected, the track points appear as solid boxes. If it is not selected, they will appear as wireframe boxes.

Track points are loaded with the default Track point material. They appear in a separate layer in the *World Browser* (see [here](#)).

When you click **OK**, the motion tracking information will be loaded, and the active camera will pick-up the tracked motion. A message will also appear, proposing to load the real footage animation into the background of the camera. If you accept and the animation information is embedded in the motion tracking file, the animation will be directly loaded into the camera background. If not, you will have to load it manually.

Importing Vertex Keyframe Animation

VUE can also import Vertex Keyframe Animation (*.mdd*) files. Animation in the *.mdd* format is represented by a sequence of baked meshes. Since the *.mdd* file itself contains only lists of baked vertex coordinates for each frame and no information about mesh structure, such as faces, it must always be used with the particular mesh it was created from.

When importing a mesh, VUE detects if there is an *.mdd* file with the same name and tries to load it. Since order and/or number of vertices can be changed by VUE at the object import phase, the correspondence between the mesh vertices and those listed in the *.mdd* file need to remain the same. This is set from their 3D coordinates; vertex coordinates for Frame 1 in the *.mdd* file always correspond to the initial undeformed object. If these coordinates somehow differ, the *.mdd* file will be considered incompatible with the mesh. Therefore, it is important not to scale or deform the mesh before the *.mdd* animation has been loaded.

Later, an *.mdd* file for a given mesh can be loaded or changed in the **Animated Mesh Options** of the *Polygon Mesh Options* dialog. See [here](#) for more information about the *Polygon Mesh Options* dialog.

It is possible to load animations with splitting or exploding objects but the option **Maintain vertex order** must be checked. If this option isn't checked for this type of animation, a warning displays. If this warning appears when loading a simple animation without splitting, it should be simply ignored.

If the animation has no splitting then be sure the **Maintain vertex order** option is unchecked.



Exporting Animation

You can export camera animation/path using the option on the Animation menu. Select the camera (the path is included) and click the option in the menu. Camera animation can be exported in the following formats:

- Maya ASCII script file (.ma)
- Vue Synchronization Data file (.dat)
- Alembic format (.abc)

The following are not available in Infinite.

- .FBX format object (.fbx)
- Nuke .chan file (.chan)





Section 8

Appendices







Mouse and Keyboard Operations


The following is a summary of all mouse and keyboard operations. You can modify the default shortcuts using the *Operations* tab of the *Options* dialog. If you changed the default interface in the *User Interface Presets* dialog, these shortcuts may not apply.

Mouse Operations


Note that the primary and secondary buttons are usually respectively the left and right buttons, and will often be referred as such throughout this documentation, but the actual button mapping may depend on your operating system configuration.

: this symbol identifies the primary mouse button on mice with two or more buttons, or the single button on single-button mice.

: this symbol identifies the secondary mouse button on mice with two or more buttons. You can emulate the secondary button using the *Ctrl* key pressed while using the primary button by enabling the option *Enable 2nd mouse button using Ctrl+First mouse button* in the *Operations* panel.


: this symbol identifies the middle mouse button on mice where the middle wheel is clickable, or on other types of mice with three or more buttons. Lacking a middle button, there is no way to emulate it in VUE. In that case, we recommend to change the few default mappings that use it to another button and key combination.


Inside any dialog and editor


 **click** in an empty space of a scrollable zone (instead of using the scrollbar). A “hand” cursor appears when the mouse is in a scrollable zone



Inside the 2D/3D Views


 **click in inactive view:** activate view.


 **click + Release without moving:** display popup menu.

 **click + Drag:** rotate camera in a 3D view (as a convenience, this trigger has the same effect as the *pan* trigger in a 2D view).


See also the [Mouse Wheel Operations](#) section for the effect of using the mouse wheel or scrolling with a multi-touch device like a trackpad in all 2D and 3D views.


Selecting and Dragging Objects


 **click:** select objects under cursor / deselect all if no objects under cursor.


Shift +  **click:** extend selection.

Ctrl +  **click:** select object inside group.

A +  **click:** select all objects below the mouse cursor, not only the foreground object.

B +  **click:** select the first object *behind* a selected object below the mouse cursor (useful in dense areas of a scene).


 **click + Drag inside objects:** move objects.

 **click + Drag outside objects:** select all objects with center in drag rectangle.

Choosing transformation axes

There are several ways to restrict the transformations of an object or the camera to one or two of the three world axes.






3D Gizmos (left to right): translation mode, rotation mode, resizing mode

The *3D Gizmos* are the red/green/blue manipulation handles that appear centered on the selected object: their are shown as arrows when in translation mode, circles when in rotation mode, and segments with cube handles when in scaling mode. These Gizmos allow you to easily translate, rotate or resize an object along only one or two of the three axes.


The trigger modifier *Constrain to Axis* (default mapping: *X*) allows to constrain movement along one axis when moving objects in a 2D (a.k.a. “orthographic”) view. The trigger *Translate Camera Constrained to Axis* allows to pan the active camera along one axis in a 3D (a.k.a. “perspective”) view. Both follow the convention described in [Moving_the_Perspective_Camera](#) (Documentation/Getting_Started/Framing_and_Picture_Composition).

Finally, when the 3D Gizmos are disabled (by unchecking the menu entry *Display >*

Gizmos > Show Gizmos), this legacy trigger is also available: **Ctrl +  click + Drag inside objects**: movement constrained to nearest axis / move object forwards/backwards if in main view.





Using Bounding Box Handles




Bounding Box Handles: Corners, Middle of Segments, Rotation Handles

 **click + Drag on selection corner dots:** resize selected object(s) keeping all proportions.

Shift +  **click + Drag on selection corner dots:** resize selected object(s) equally along the two view axes, leaving dimension along the third axis unchanged.

 **click + Drag on selection middle dots:** resize selected object(s) along indicated axis.

 **click + Drag on selection rotate handle:** rotate selected object(s) around the axis perpendicular to the view (e.g. for Top view (XY), around vertical axis (Z)).

- When the 3D Gizmos are shown, the *rotate handle* appears when hovering near the sides of the selected object(s)'s bounding box (when not too close to another “hot” area like a corner or segment’s middle point).
- When the 3D Gizmos are not shown, the *rotate handle* is the icon showing a circular arrow with a “Rotate inside view” tooltip, at the right-hand side of the manipulation




icons group.



Manip. Icons Group


Finally, when the 3D Gizmos are *not* shown:


 **click + Drag on selection “4 arrows” handle:** rotate selected object(s) around the 2 axes of the view (e.g. for Top view (XY): around Front (X) and Side (Y) axis).


- The “4 arrows” handle is the icon with the “Rotate around view axes” tooltip at the left-hand side of the manipulation icons group.

Industry-standard navigation triggers

This variant is selected using the *Industry-standard* button in the [User Interface Preset](#) dialog or below the shortcuts list in the [Operations](#) options panel.


Alt +  click + Drag: move (a.k.a. “pan”) the camera or viewpoint along the view axes.


Alt +  click + Drag: change the zoom factor (ie. view extent) in a 2D view, or the camera focal in a 3D view.

Alt +  click + Drag: rotate the camera or viewpoint around the selected object(s) (a.k.a. “orbit”).

Legacy navigation triggers

This variant is selected using the *Legacy* button in the [User Interface Preset](#) dialog or below the shortcuts list in the [Operations](#) options panel.

Shift +  click + Drag: move (a.k.a. “pan”) the camera or viewpoint along the view axes.


Ctrl +  click + Drag: change the zoom factor (ie. view extent) in a 2D view, or the camera focal in a 3D view.

 click + Drag: rotate the camera or viewpoint around the selected object(s) (a.k.a.



“orbit”).


When Editing Paths in 2D/3D Views


 **click on a way point:** select way point / activate path editor.


Ctrl +  **click:** extend way point selection.


Shift +  **click:** select all way points between last selected and this one.


Double-click on way point: select all way points of object.

 **click + Drag way point:** move selected way points.

 **click + Drag on selection corner dots:** resize selected way points keeping proportions.


 **click + Drag on selection middle dots:** resize selected way points along indicated axis.


 **click + Drag on selection 4 arrow handle:** rotate selected way points around the 2 view.

 **click + Drag on selection rotate handle:** rotate selected way points around perpendicular axis to view.

Inside the World Browser

 **click on object:** select object under cursor.


 **click on nothing:** deselect all objects.

Ctrl +  **click on object:** extend object selection.

Shift +  **click on object:** select all objects between last selected object and object under cursor.




 **click on layer:** select all objects in layer.

 **click on layer / group expansion box:** toggle Unfolded / Folded state.

 **click on layer state box:** toggle Active / Locked / Hidden state.


Double-click on layer state box: activate layer.

 **click + Drag:** move selected objects to the release point location.

Ctrl +  **click + Drag:** copy selected objects to the release point location.

 **click:** display popup menu.


Inside the Timeline


 **click on nothing:** deselect all objects.

 **click + Drag:** move ruler.


Ctrl +  **click + Drag:** zoom ruler in / out.

 **click on item:** select item.

 **click + Drag:** move selected item.

 **click on nothing + Drag:** marquee keyframe selection.

Ctrl +  **click on keyframe:** extend selection.

Shift +  **click on keyframe:** select all keyframes between last selected keyframe and this one.

 **click on property:** select keyframe at current time.

Double-click on property: select all property keyframes.

 **click on expansion box:** toggle Unfolded / Folded item state.




Ctrl + Drag selected keyframe: drag all subsequent keyframes together with selected keyframe.

Inside the Animation Wizard Path Editor



Click: add / delete / insert way point.

Ctrl +  click on way point + Drag: move way point.



click + Drag: move view.

Ctrl +  click + Drag: manual zoom.

Space + Drag: move view.

Insert: insert way point at cursor.

Del: delete way point at cursor.

Mouse Wheel Operations

The mouse wheel can be used in several contexts with an intuitive behaviour:

- In a 3D perspective view like a camera view or the 3D preview of a *Terrain*, *Plant* or *Text Editor*, it will move the viewpoint forward (*into* the scene) or backward (*away* from the scene) along the direction the viewpoint is looking at,
- In an orthographic view (top, front or side views), it will shrink (zoom-in) or enlarge (zoom-out) the extent of the scene shown in the view,
 - In this type of view, the center of the zooming action is not the center of the view but the position below the mouse cursor while the wheel is scrolled, making it easy for example to zoom on an object you want to focus on.
- In other types of 1D or 2D views like the *Function Graph*, the *Time Line* or the *Function Node Preview*, it will also zoom in or zoom out the view contents.



Keyboard Operations

You can modify the default shortcuts using the *Options* dialog (see [here](#)).

Note:

The following shortcuts may not apply if you have changed the default interface preset (see [here](#)).

Interface Shortcuts

Ctrl + N: new file.

Ctrl + O: open file.

Ctrl + S: save file.

Ctrl + W: close file.

Ctrl + Q: exit.

Ctrl + Z: undo last operation.

Ctrl + Shift + Z: redo last operation.

Del: delete selected objects / object animation.

Ctrl + X: cut selected objects into clipboard.

Shift + Del: cut selected objects into clipboard.

Ctrl + C: copy objects from clipboard.

Ctrl + Insert: copy objects from clipboard.

Ctrl + V: paste objects from clipboard.

Shift + Insert: paste objects from clipboard.

Ctrl + D: duplicate selected objects.

Ctrl + A: select all objects in scene.



Escape: deselect all.

Ctrl + Num .: store camera.

Ctrl + 0..9: activate camera.

Ctrl + Num +: zoom into view or timeline ruler.

Ctrl + Num -: zoom out of view or timeline ruler.

Ctrl + Shift + M: select objects by material.

Ctrl + Shift + O: select Objects by type.

Ctrl + Shift + W: select objects by preview color.

Tab: walk to next objects in selection.

F1: open help files.

F4: display *Atmosphere Editor*.

F5: load atmosphere.

F6: display *Summary of Materials*.

F7: toggle only display main view.

F8: display last render.

Ctrl + F8: save color picture.

F9: render.

Ctrl + F9: *Render Options*.

Ctrl + Shift + F9: resume render.

F11: display *Timeline*.

Ctrl + F11: display *Animation Wizard*.

Alt + Enter: toggle full screen mode.

0, 1, 2, 3: Select *Main, Top, Front, Side view*.



Object Creation

Shift + W: add Water.

Shift + G: create a Ground plane.

Shift + A: create a Cloud plane.

Shift + S: create a Sphere.

Shift + C: create a Cylinder.

Shift + U: create a Cube.

Shift + O: create a Cone.

Shift + Y: create a Pyramid.

Shift + R: create a Torus.

Shift + P: create a Plane.

Shift + H: create an Alpha Plane.

Shift + T: create a Terrain.

Ctrl + Shift + T: create a terrain inside the editor.

Shift + F: create a Procedural terrain.

Shift + V: create a Plant.

Ctrl + Shift + V: load plant species and create.

Shift + N: create a Planet.

Shift + K: create a Rock.

Shift + Q: add a Directional light.

Shift + L: add a Point light.

Ctrl + Shift + L: add a Quadratic Point light.

Shift + X: add a Spot light.

Ctrl + Shift + X: add a Quadratic Spot light.



Ctrl + L: load Object.

Shift + M: create a Metablob object.

Object Edition

Arrow up: nudge selected objects up / **if no objects selected:** nudge views up / nudge camera down if in *Main camera view*.

Arrow down: nudge selected objects down / if no objects selected: nudge views down / nudge camera up if in *Main camera view*. **Arrow left:** nudge selected objects left / if no objects selected: nudge views left / nudge camera right if in *Main camera view*.

Arrow right: nudge selected objects right / if no objects selected: nudge views right / nudge camera left if in *Main camera view*.

Page up: nudge selected objects closer / if no objects selected: nudge views closer / nudge camera forward if in *Main camera view*.

Page down: nudge selected objects further / if no objects selected: nudge views further / nudge camera backward if in *Main camera view*.

Shift + Nudge key: nudge 1/10th of increment.

Ctrl + Nudge key: resize objects / change focal.

Shift + D: drop selected objects.

Ctrl + G: group selected objects.

Ctrl + Shift + U: make Boolean Union.

Ctrl + Shift + I: make Boolean Intersection.

Ctrl + Shift + D: make Boolean Difference.

Ctrl + U: ungroup selected groups.

Ctrl + M: change selected objects Material.

Ctrl + E: edit selected objects (if applicable).

Ctrl + Shift + L: pick link parent.

Ctrl + Shift + T: pick tracked parent.



Ctrl + Shift + B: display *Animation Toolbox*.



VUE xStream

This section of the reference manual deals with the specifics of running VUE xStream in the integrated mode. In the integrated mode, VUE xStream gives you the ability to create, edit, and render a VUE scene inside a target application.

xStream

If you have trouble installing VUE inside your target application, we suggest you first refer to the online VUE FAQ on the e-on software website. You may find additional tips and advice for specific setups not covered in this manual.

Standalone and Integrated Modes

xStream

VUE Infinite is a standalone application only. It offers all the tools that are required to create, animate and render natural 3D environments, without the need for any other applications. Obviously, it also offers many tools to help integrate your VUE work with other 2D and 3D applications – but using other applications alongside VUE remains optional.

VUE xStream, on the other hand, is primarily designed to run “inside” another 3D application (the target application) as a plugin. These versions let you “host” a VUE environment inside this 3D application, and automatically combines the VUE objects with the native objects of that application. Thus, the VUE environment is “integrated” inside the target application.

Note:

VUE can also be run “outside” of its target application. It is then running in standalone mode and is operated, in all ways, exactly like VUE Infinite.

VUE Licenses

xStream

It means that a single VUE license lets you use the integration plugin in all your compatible 3D applications and renderers. If you want to use VUE with a 3D application, simply make sure the plugin has been installed for the application.



VUE Installation

xStream

During the installation of VUE, you will be asked which versions of the integration plug-in you want to install. All supported applications for the current architecture (64bit) will be listed. Please select the versions for which you want to install the VUE integration plug-in.

If the installer detects the application on your computer, then the corresponding item in the list will be already selected and the path to the application will already be filled. Otherwise, you need to:

- Check the box near the application version (you will be asked to browse to the location where the application is installed)
- Or select the application version, click on the Browse button, and then check the box.

With Windows, the installer will detect if the application you selected is already running, and will ask you to close it before continuing the installation. This is to ensure that the plug-in files are successfully installed and that the configuration files of the host application can be edited. To avoid any mistake during installation, the installer will only allow the installation into a folder if it can detect the application there. But, the installer has no way to tell which version of the application you're selecting in the browser. It is up to you to select the correct path.

Note:

In the case that you haven't selected any compatible application, the installer will ask for confirmation. If you choose to continue the installation, only the standalone application and the core will be installed. You won't be able to use any integration plug-in except if you already installed them (but it's better to use the same version of the plug-ins and the core to be sure you are using the latest improvements and fixes).

32bit versus 64bit Versions

xStream

VUE is only compatible with 64bit applications.



Mental Ray Configuration Files

xStream

When you have finished selecting which versions you want to install, you will be asked to locate the Mental Ray configuration files.

For each application version, you can correct the default path of the Mental Ray configuration file. Unless you have modified the Mental Ray default configuration, you can safely ignore this step. In the other case, be sure to select the correct path to the mental ray configuration file.

If you haven't selected any Mental Ray-compatible application, then this step is skipped by the installer. Only 3DS Max and Maya need their configuration file to be edited. This is not the case for other Mental Ray-compatible applications (like Softimage).

V-Ray Renderer

xStream

With the current version of VUE, V-Ray for 3DS Max and Maya are supported. You must select the correct version of V-Ray during installation (V-Ray 1.5, 2 or 3, depending on the Maya/3DS Max version). The VUE for V-Ray shader is automatically installed in your host application folder, along with the general plug-in. No user action is required to install VUE for V-Ray.

Supported Versions of the Host Applications

xStream...

Only applications listed on our website, on the Requirements page of the VUE xStream product, will be able to load the integration plug-in. If your application is older than the ones listed on the Requirements page, the plug-in won't work.

Note:

On MacOSX, if the plug-in doesn't show up or can't be loaded, make sure that the host application is running in 64-bit mode. VUE is 64-bit only, so an installation of VUE will not show up in a 32-bit version of the host application. You can check that by



opening the information panel on the host application icon (Apple+I, or Get Info: in the contextual menu). Then, enable the 64-bit mode. Remember that VUE is Mac Intel only, not PPC.

...xStream

Adding the VUE Menu and Toolbars

xStream

All of the VUE commands are accessed through the use of a VUE menu in the target application main user interface. VUE introduces toolbars and icons, supported in all applications except LightWave.

Described below are the steps required (if any) to add the VUE menu and toolbars to each of the supported target applications. For most applications, this is done automatically when the target application loads the plug-in.

You don't need to add the VUE menu or toolbars if you won't be using VUE inside your application (e.g. if you install a plug-in for an application that will be used only for network rendering).

3DS Max

xStream...

The plug-in is automatically loaded during 3DS Max startup. The VUE menu is added automatically by the plug-in. You can access the VUE menu in the 3DS Max menu bar. It should be located to the right of the **Help** menu entry.

Check that the plug-in is loaded by opening the *Plugin Manager* (from the **Customize** main menu). There should be at least one entry for VUE:

- *Vue xStream.dlb* is the main VUE file for 3DS Max plug-in,
- *Vue xStream.dlr* is the VUE file for V-Ray plug-in.

To add the toolbar, use the **Customize | Customize User Interface** menu command of 3DS Max to display the customization dialog. Then, select the **Toolbars** tab, and click on the **Load** button. Browse to the ui subfolder of the 3DS Max application and load the *xStream.cui* file.



VUE's Mental Ray shader is directly loaded by Mental Ray, so check the Mental Ray message window to check if VUE is correctly loaded.

...xStream

The VUE for V-Ray plug-in is automatically loaded during startup. If V-Ray is the active renderer at the time you create or load a VUE scene, the 3DS Max integration plug-in will automatically setup everything for the V-Ray renderer. If V-Ray was not the active renderer, and you later want to render with it, you need to open the *Rendering/Environment* dialog in 3DS Max. Then, in the *Atmosphere* section, add the VUE/VRay atmospheric effect to the list.

When this atmospheric effect is present (either added automatically or manually with the step above), and you render with Mental Ray, you will get an error message telling you that this effect is not supported by the Mental Ray renderer. This error has no other consequences other than displaying this message, it won't affect the renderer in any way. If you want to get rid of this message, you can remove the atmospheric effect from the list whenever you switch to Mental Ray, and add it again if you revert to V-Ray.

Depending on the active renderer at the time of the VUE scene creation, the integration plug-in will also try to select corresponding shadow types for VUE proxy lights. This means that if Mental Ray is active, the lights will have ray traced shadows selected; if V-Ray is active, they will have "V-Ray raytraced shadows" selected. So, if you load a VUE scene containing many lights, make sure you first select your preferred renderer in the 3DS Max render settings dialog. This will save you the burden of manually editing each light to select the appropriate shadow type.

Cinema4D

xStream

The plug-in is automatically loaded during Cinema4D startup. The VUE menu is automatically created, and it's located in the **Main Menu** bar.

You can load the VUE toolbars as Cinema4D Palettes using the **Window | Layout | Load Palette** menu of Cinema 4D. You will then have to browse to the `plugins/VuexStream/UI/Palettes/` subfolder of your Cinema 4D application (not in the `user` folder) and select one of the palette files available.

You can also load the complete VUE layout using the **Window | Layout | Load Layout** menu of Cinema 4D. You will then have to browse to the `plugins/VuexStream/UI/Layout/` subfolder of your Cinema 4D application (not in the `user` folder) and select the `xStream.l4d` file there. Please note that this will replace the current layout.



LightWave

xStream

The plug-in is automatically loaded during LightWave startup if you have the AutoScan Plugin option enabled, if not you need to add the plugin manually. This option is under **Preferences | General**, press “o” to open this panel, by default this option on **Layout** is ON.

This menu needs to be added to the interface manually. This operation only needs to be performed once.

Here is how to add the VUE menu:

- Open the **Menu** editor by selecting **Edit | Edit Menu Layout...** or pressing Alt + F10 on LightWave’s Layout.
- In the *Menu Editor*, click on **Main Menu** in the **Menus** list (the list on the right).
- Right click on **Main Menu** and select **Import Branch**. A browser will appear; browse to the VUE application folder, open the **Environment\XStream\Lightwave** subfolder, and select the **VuexStream_2016 Menus.cfg** file.
- This will add a VUE tab menu in the menus list.
- Now you can drag and drop the **VUE** tab to place it where you want. You can, for instance, move it after the **User Tab**.

Due to limitations in the SDK of LightWave, there are no toolbars or icons available.

The current VUE is not compatible with LightWave 9.3.

Problem with Menu Display -- Incomplete Menus

xStream...

If the previous version of xStream menus have been loaded in LightWave, you may not be able to load the new version of the xStream menus correctly.

Before adding the VUE new menu, you must remove the xStream menu from any previous version: In LightWave, open the **Configure Menus** dialog, using the command **Edit | Edit Menu Layout**.

In the list of menus (on the right), select VUE xStream and press the **Delete** button.



Close LightWave and reopen it. Now you can add the VUE new version of the xStream menu.

...xStream...

Warning! Even after removing the old menu with the LightWave interface (see above), the VUE *current* menu can still be incomplete. The bug often happens in Vista/Windows 7. It is certainly caused by a problem in the handling of menus by LightWave.

To fix it, here are the steps to follow:

1. Close LightWave.
2. Delete the configuration files of LightWave, containing the menu layout.

They are located in the user folder:

OS

Version

Path

Windows Vista/Win 7

LightWave *new version*

C:\Users\[user_Name]\LW[xxx].cfg

C:\Users\[user_Name]\LWEXT[xxx].cfg

LightWave *new version*

C:\Users\[user_Name]\.NewTek\LightWave\new\ version\LW[xxx].cfg

C:\Users\[user_Name]\.NewTek\LightWave\new\ version\LWEXT[xxx].cfg

Windows XP

LightWave *current version*

C:\Documents\ and\ Settings\[user_Name]\LW[xxx].cfg

C:\Documents\ and\ Settings\[user_Name]\LWEXT[xxx].cfg

LightWave *new version*

C:\Documents\ and\ Settings\[user_Name]\.NewTek\LightWave\\{NEW_VERSION}\LW[xxx].cfg} \texttt{C:\Documents\ and\ Settings\[user_Name]\.NewTek\LightWave\[NEW_



VERSION\}\LWEXT[xxx].cfg

MacOSX

LightWave *current version*

/Users/[user_Name]/Library/Preferences/LightWave3D/Layout\ 9

LightWave *new version*

/Users/[user_Name]/Library/Application\ Support/NewTek/LightWave/[NEW_VERSION]
\textquotesingle{}/Layout\ 10

/Users/[user_Name]/Library/Application\ Support/NewTek/LightWave/[NEW_VERSION]
/Extension\ 10

Where [xxx] depends on your version of LightWave and your system (for example LW9-64.cfg).

3. Open LightWave again, and add the VUE menu again.

...xStream

Maya

xStream...

The plug-in is not automatically loaded during Maya startup; in order to load the plug-in, open the plugin manager (**Window | Settings/Preferences | Plugin Manager**).

Locate **Vue xStream.mll** entry and check the **Loaded** checkbox to load the plug-in. If you want the VUE plug-in to be loaded each time you start Maya, check the **Auto load** checkbox.

Once the plug-in is loaded, the VUE menu is automatically created. It can be found in the menu bar, at the left of the **Help** menu. If you have checked the **Auto load** option, the next time you start Maya, VUE will be accessible.

Note:

if you want to render with VUE and the Mental Ray renderer, you should ensure that the **Mayatomr.mll** plugin entry (**Maya to Mental Ray** plug-in) is also loaded (and auto loaded) in the **Plugin Manager**. If you load the **Maya to Mental Ray** plug-in *after* having created a VUE scene, the VUE shaders for Mental Ray may not have been successfully created (resulting in a black render when choosing mental ray as the renderer). It is recommended to always load the Mental Ray plug-in during Maya startup.



In order to load the VUE toolbar shelf, please use the **Maya Ui | Load xStream shelf** menu command from the VUE menu.

...xStream

Softimage

xStream

The plug-in is automatically loaded during Softimage startup. The VUE menu is automatically added to the menu bar, between the **Display** and the **Window** menu entries.

If you want to check that the plug-in is correctly loaded inside Softimage, open the plug-in manager (**File | Plug-in Manager...**). The VUE plug-in entry should be located in the **Factory Root** plugin tree.

You can access the VUE toolbars from the **View | Toolbars** menu of Softimage.

xStream Options Dialogs

xStream

xStream Options Dialog

xStream

This dialog is accessed through the **File | Options...** menu command. This is where you will setup the general behavior of VUE.

Some options are not available in some versions of the plugin, because they're not needed or not appropriate: for instance, options related to mental ray are available only in applications that support this renderer.

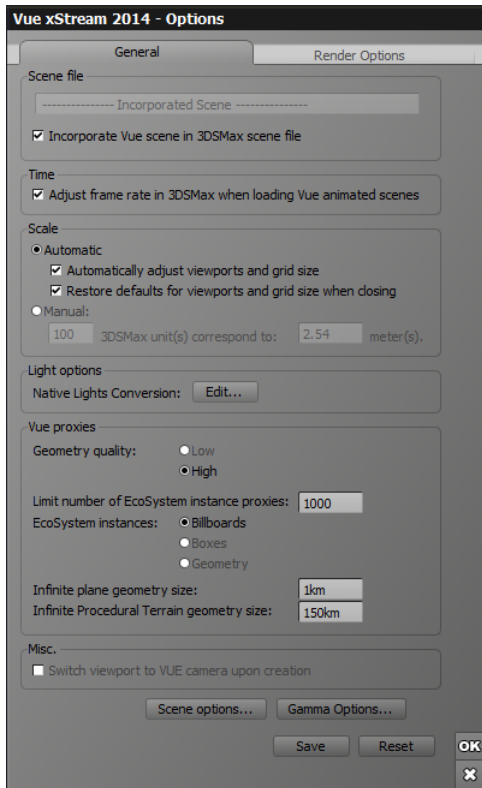
The *Options* dialog has two tabs: **General** and **Render Options**.

xStream...



General Tab

...xStream



xStream Options dialog – General tab

This tab of the Options dialog has options for how VUE and the host application interact.

Scene File

xStream

- **Incorporate VUE scene in native scene file:** When on (which is the default), saving the native scene will create a unique file which contains both the native scene and the VUE scene. (see [Saving your work while using VUE](#)).
- **(LightWave only) Relative to LightWave content folder:** When checked, the path to the VUE scene is no longer absolute, but stored relative to the LightWave content folder path.



Time

xStream

- **Adjust frame rate when loading VUE animated scenes:** This gives you the option of using the frame rate from VUE or from the host application. When loading a VUE scene which contains animated parameters or objects, VUE can load the frame rate of this VUE scene and set it to the same value in the host application. When VUE loads the animation from the VUE scene, it first reads the animation for each frame of the VUE animation. It then recreates all of these keyframes in the native scene. Enabling this option will prevent VUE from creating keyframes at non-integer frames in the native scene. This is especially known to lead to interpolation problems on rotations in Cinema 4D. This option is not enabled by default because it will edit the current frame rate in the host application.

Scale

xStream

- **Automatic or Manual:**
 - if you select **Automatic** (the default), then the conversions between native scene and the VUE scene will be based on the internal units of both applications. For instance, if you have an object that's 1 inch tall in VUE and you open it with the integration plugin it will still be 1 inch tall in the target application, no matter which unit is currently displayed. If you're displaying centimeters, the object will be 2.54 cm tall.
 - If you select **Manual**, you have to specify the ratio between the two internal units with the two editable fields.

Note:

In Softimage, there is no such thing as “real-world units”, an Softimage unit having no particular “meaning”. When **Automatic** is enabled, one Softimage unit is set to one meter.

- **Automatically adjust viewports and grid size:** by default, host applications use smaller scene sizes than VUE. With this option, VUE sets the grid size and viewport configurations to match the VUE default scene size. For better control, if you want to set these options manually, uncheck this option else VUE will override your settings.
- **Restore defaults for viewports and grid size when closing:** if the previous option is enabled, when closing the VUE scene, this option restores the default values for grid size and viewport configuration (clipping planes, orthographic camera position, etc.).



Light Options

- **Native Lights Conversion:** this **Edit** button opens the *Native Lights Options* dialog to let you edit the way the integration plug-in translates the lights from one application to another.

xStream

VUE Proxies

- **Geometry quality:** this option lets you select between **High** quality geometry or **Low** quality geometry. Obviously, the proxy objects look better when the high quality is selected but display in the view ports is faster with the low quality mode. In the high quality mode, if you notice some slowdowns after having loaded a VUE scene or added an object to your VUE scene, then you should switch to the low quality mode.

xStream...

Several notes on this setting:

- It has no effect on the render of the VUE objects. Rendered objects always use the most detailed geometry, or procedurally generated geometry,
- It won't have any effect on EcoSystems instances (please refer to the EcoSystem instance option below),
- It won't have any effect on infinite planes geometry (please refer to the infinite planes geometry size option below).

Changing this option will force the re-generation of the geometry of all objects in the scene, therefore:

- The update of the native scene will take longer than usual,
- Any change you might have applied to the geometry of the proxy objects in the target application will be lost,
- It can be useful to change this setting to reset the geometry.
- **Limit Number of EcoSystem Instance Proxies:** This option lets you edit the number of EcoSystem instances displayed in the viewports of your host application. For instance, if your scene contains large EcoSystems and your application becomes less responsive, you can lower this setting. The limit density setting in the *EcoSystem Painter* dialog has no effect on native (3DS, C4D, Softimage) objects.
- **Infinite Plane Geometry Size:** There is generally no such thing as an infinite primitive in target applications. For this reason, VUE creates a flat polygonal



object instead of a true infinite plane. Depending on the size of your scene, you may get annoyed if the object proxy is too small or too large for your scene: if you want to drop your native objects to the ground you'll have trouble if the ground is not covering your entire scene; it can also affect the display, through the near/far clipping planes.

With this option, you can set the size that VUE will set the proxy objects being used for infinite planes.

...xStream

Note:

This option is not available in Cinema4D, as VUE will use the Cinema4D floor primitive for such proxy objects.

Misc

- **Switch viewport to VUE Camera upon creation:** Checking this option switches the viewport to the VUE camera instead of the viewport/camera currently being used in the native application.
- **Scene Options:** This button will open the *Scene Options* dialog, which allows you to edit the spherical scene options. These options are similar to the ones found in the *Options* dialog of VUE standalone, in the **Units and Coordinates** tab. Please refer to the corresponding section of this manual for more information on spherical scene options. These scene options are specific to the current VUE scene, and are saved directly in the VUE scene file.
- **Gamma Options:** This button opens the *Gamma Options* dialog, where you can enable and adjust the gamma correction. Since you are rendering in the host application, these settings do not affect renders; any gamma corrections for the renders must be made in the host application settings. The settings here affect the VUE Color Editor, Material previews, Color Function previews and the Scalar Function previews.
- **Save:** clicking on this button will save the current options as default. The next time you start VUE, these default options will be used instead of the generic default values. These default values will be used in all versions of the plugin, in case you have installed versions for several target applications. Only scale options are application-specific.
- **Reset:** clicking on this button will reset all options in the dialog to default values (either generic default values or customized default values).

xStream

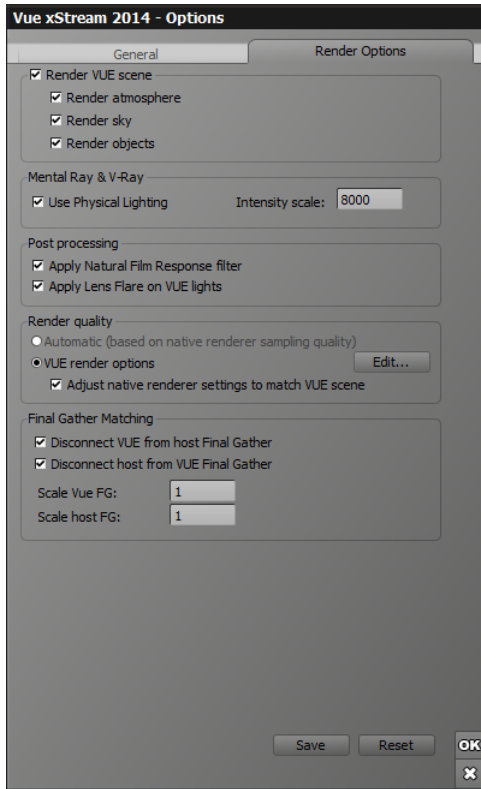


Render Options Tab

xStream

Render VUE Scene

xStream...



xStream Options dialog – Render Options tab

By default, the **Render VUE scene** option and its sub-options are all checked, meaning everything is rendered.

- **Render atmosphere:** uncheck it if you don't want atmospheric effects to be rendered.
- **Render sky:** uncheck this option if you don't want to render the VUE sky. Please note that it won't be rendered for primary rays but also for secondary rays (i.e. it



won't be visible in reflections/refractions). If you want to remove the background, you should use the alpha channel instead of this option.

- **Render Objects:** uncheck it if you don't want to render VUE objects.

If you uncheck the **Render VUE scene** frame checkbox, then VUE won't be used at all during the render of the scene. It can be used to check the render of the native objects alone. In the case of a mental ray render, VUE shaders won't be added to the scene, which means you'll get a 100% VUE-free render.

...xStream

LightWave renderer (LightWave only)

- **LightWave renderer:** check this option to enable the blending of several volumetric shaders in your LightWave scene. For instance, if you use HyperVoxels and VUE, turn this option on. You may have to enable a similar option in your other volumetric shader. Please refer to the documentation of the shader to know how to do this. This option is only compatible with *Spectral* and *Volumetric* atmospheres.

xStream

Mental Ray & Vray (only applications supported for Vray and Mental Ray)

- **Use Physical Lighting:** check this option if you're using a Mental Ray or Vray photographic exposure (this is generally the case if you're using Mental Ray Sun & Sky). This option will scale all values returned by VUE during render by the intensity scale you specify. The default value of 8000 should work well in daylight scenes, but you may have to change it depending on your scene setup. Typical values range from 5000 to 30000.

xStream

Post Processing

- **Apply natural film response filter:** select this option to enable the non-linear reaction to light typical of photochemical films. Please refer to the *Camera Options* for more information about natural film response.

xStream

Please note that if you want to create high-dynamic renders (with floating point values exceeding 1.0), then you should un-check this option, because the natural film response filter tends to reduce the dynamic range of the output image.

- **Apply lens flare on VUE lights:** if checked, then VUE will compute lens flares on VUE lights.



Render Quality

xStream

Select either the **Automatic mode** or the **VUE render options** mode.

- **Automatic (based on native renderer sampling quality):** In this mode, the VUE render quality will be set to match the native renderer settings. If you perform a draft render, the VUE render quality will be set to a low render quality mode in order to speed up the rendering of VUE objects. If you use production quality render settings, the VUE render quality will be increased. The exact render quality depends on several parameters, most noticeably on the Sampling Quality parameter of the native renderer.
- **VUE render options:** In this mode, the render quality preset that will be used is the one that's selected in the *VUE Render Options* dialog. For a quick access to the *Render Options* dialog, click on the **Edit...** button. Alternately, you can access it through the **Render | Render options...** VUE menu.
- **Adjust native renderer settings to match VUE scene:** When enabling this option, the integration plug-in will try to match as closely as possible the VUE render preset you select. For instance, if you select a **Preview** render quality preset in the *VUE Render Options* dialog, the integration plug-in will select a low quality setting in the native renderer. If you select a **Final** preset, the plug-in will increase the quality of the native renderer accordingly. When loading a scene, the render output size will also be matched.

Final Gather Matching

xStream...

- **Disconnect VUE from host Final Gather:** by checking this option, you instruct Mental Ray to ignore the VUE scene when computing Final Gather. This way, you can configure VUE and Mental Ray Global Illumination independently.
- **Disconnect host from VUE Final Gather:** by checking this option, you instruct VUE to ignore the native scene when computing Final Gather.

By using these two options, you can configure VUE and the native renderer Global Illumination independently:

- **Scale VUE FG:** this value can be changed to modify the contribution of VUE's Final Gather in the host Final Gather. A higher value will increase the influence of VUE objects on the global illumination of native objects. The default value is 1 (which should be correct in most cases).
- **Scale Host FG:** conversely to the above, this value will modify the contribution of the host's Final Gather in VUE's Final Gather. A higher value here will make



the native objects have a higher influence on the indirect lighting of VUE objects.

- **Save:** clicking on this button will save the current options as default. The next time you start VUE, these default options will be used instead of the generic default values. These default values will be used in all versions of the plugin, in case you have installed versions for several target applications. Only scale options are application-specific.
- **Reset:** clicking on this button will reset all options in the dialog to default values (either generic default values or customized default values).

...xStream

The only way to reset the default values back to the generic default values is to delete the **VueStream.cfg** file in your **Config** folder, where “xxxxxx” is your product name and version.

It is found in c:\Documents\ and\ Settings\username\Application\ Data\e-on\ software\ Vue\ xxxxxxxx\Config on Windows XP machines.

It is found in your c:\Users\username\AppData\Roaming\e-on\ software\Vue\ xxxxxxxx\Config folder on Vista, Win 7, 8 and Win 10 machines.

It is found in \Users\username\Library\Application\ Support\e-on\ software\ Vue\ xxxxxxxx\Config on a Mac.

VUE Light Options Dialog

xStream...

This dialog is accessed from the VUE Options dialog, using the **Edit...** button in the *Light Options* section.

Using this dialog, you can choose the way you want the integration plug-in to translate lights before the native and the VUE scene.

It is used in the following cases:

- you want to control what VUE will do concerning a specific native light.
- you want to control what the host application will do concerning a specific VUE light.

This dialog is therefore separated in two tabs: **Native Lights** and **VUE Lights**.

Before moving to the description of the options offered by the two tabs, please note that these options are saved in the scenes, and that VUE is also able to uniquely identify

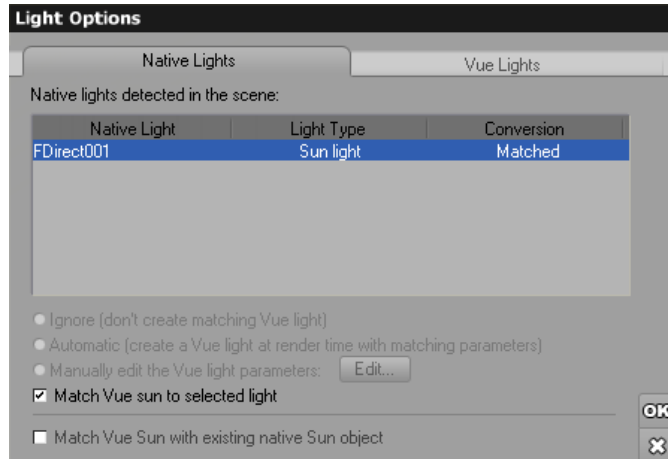


lights, even if you edited the VUE scene separately in the standalone. In this case, newly added lights will use the default settings, while the lights that were already present when you edited the scene using the integration plug-in will use the settings you set at that time.

...xStream

Native Lights

xStream...



xStream Light Options dialog – Native Lights

On this tab, you will see a list of all the native lights detected in the native scene. Each light will be displayed with its name, its type (Point light, Spot light, etc.), and the current mode for this light.

When you select a light in the list, you can change the current mode by using the radio buttons in the bottom of the dialog. The available modes are:

- **Ignore:** when this option is selected native light will be ignored by VUE. No VUE light will be created at render time.
- **Automatic (which is the default mode):** the selected native light will have a corresponding VUE light created at render time by VUE. This means that if you have a red spot light placed in your native scene, VUE will create a corresponding red spot light at the same position and orientation. VUE will try to match it as close as possible to the native light, and will use it to illuminates the VUE scene (which will make your VUE elements appear red if placed under the spot).
- **Manual:** Using this option, you tell VUE that you would like to manually edit



the settings of the VUE light that is used at render time. This is especially useful when the automatic matching doesn't give you the results you're expecting, or if you want to fine tune the matching. To edit the light, use the **Edit** button, and edit the light as you would do for any other VUE light, using the standard *Object Properties* dialog. When using this option, only the position and orientation of the light will still be matched. This means that if you move the native light in the native scene, the corresponding VUE light will be moved accordingly.

...xStream

- **Match VUE Sun to Selected Light:** when this option is checked, the VUE plugin will match any native light to the VUE sun. It's position and orientation will be read and set to the VUE sun.
- **Match VUE Sun with existing native Sun object:** when this option is checked, then the VUE plugin will try to match an existing native sun object. This means that if your native scene contains a sun object, then its position and orientation will be read and set to the VUE sun.
It will also have two important effects, to avoid conflicts between the VUE proxy light and the original native sun light.
 - The VUE sun proxy object will be created in the target application but it won't illuminate native objects or cast any shadows in the native scene, because the native scene is already illuminated by its own sun.
 - The native sun object will be ignored at render time, when creating the temporary lights in the VUE scene that recreates the native scene lights. Please refer to the Rendering section to learn more about temporary lights created at render.

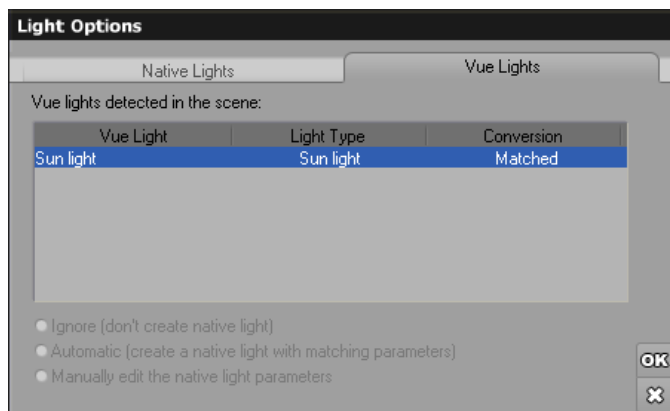
Using this option, you can continue using your native sun without having to tweak your scene for VUE; everything will be handled by the plugin. Of course, if you move the proxy object representing the VUE sun light, the change will be ignored and the proxy object will go back to its original location the next time the scene is updated (please refer to the [VUE Integration](#) section to know more about scene updates).

xStream...



VUE Lights

...xStream



xStream Light Options dialog – Vue Lights

On this tab, you will see a list of all the VUE lights detected in the VUE scene. Each light will be displayed with its name, its type (Point light, Spot light, etc.), and the current mode for this light.

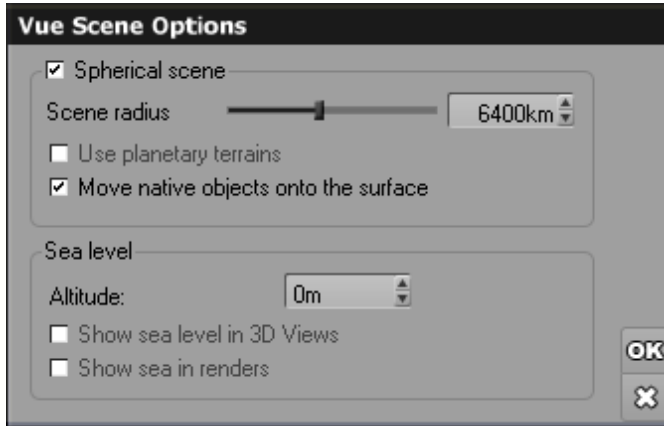
When you select a light in the list, you can change the current mode by using the radio buttons in the bottom of the dialog. The available modes are:

- **Ignore:** when this option is selected, the selected VUE light will be ignored by the integration plug-in. No native light will be created in the native scene. When changing from another mode to this one, the native light that was present in the native scene is deleted. When changing from this mode to another, the native light is re-created in the native scene, with default settings.
- **Automatic (which is the default mode):** the selected VUE light has a corresponding native light created in the native scene. This means that if you have a red spot light placed in your VUE scene, the integration plug-in will create a corresponding red spot light at the same position and orientation when loading the scene. The integration plug-in will try to match it as close as possible to the VUE light, and will use it to illuminate the native scene (which will make your native elements appear red if placed under the spot).
- **Manual:** Using this option, you tell the integration plug-in that you would like to manually edit the settings of the native light. This is especially useful when the automatic matching doesn't give you the results you're expecting, or if you want to fine tune the matching. To edit the native light, just use the regular tools of the host application. When using this option, only the position and orientation of the light will still be matched.



VUE Scene Options Dialog

xStream



VUE Scene Options dialog

This dialog is for setting up a scene with a spherical terrain, either complete planets or partial curved terrains. These properties should probably not be checked as a scene default. This dialog is accessed from the *VUE Options Dialog*.

Spherical scene

xStream

- **Enable spherical scene:** this will enable the spherical scene properties in the scene you are currently working on.
- **Use planetary terrains:** this will reform all of the infinite planes currently in your scene (and any you might add) into a spherical shape.
- **Scene radius:** this sets the size of the terrain you are creating.
- **Move native objects onto the surface:** this places any native objects directly onto the terrain surface.

Sea Level

xStream...

- **Altitude:** this sets the default for the water plane in your scene. This default affects all of the scenes you create using a water plane.
- **Show sea level in 3D Views:** this gives you a visible plane in your views as a



reference. A Sea level plane will show in the *World Browser*, but be invisible in renders.

...xStream

- **Show sea in renders:** this will give you a visible water plane and it will show in the *World Browser* as Sea.

If you don't check either option, sea level is still present and its value is defined by default as $z=0$, or whatever value you give in on this screen.

VUE Render Options Dialog

xStream

This dialog is accessed:

- With the **Render | Render Options** menu command,
- From the *VUE Options* dialog, using the **Edit...** button in the **Render Quality** section.

This dialog is similar to the standalone *Render Options* dialog, with the following changes:

- General render options like image size, aspect-ratio, render area, tile rendering, etc. are not present because these options are set in the target application and have no meaning in the case of the integration plugin.
- Options about which renderer to use (standalone, *RenderCows*, etc.) are obviously not present, and the same applies for the renderer destination (render in main view, etc.).
- The **OpenGL** render quality preset is not available because the VUE render engine is not compatible with this render preset.
- **Object anti-aliasing** options are removed because the target renderer is in charge of object anti-aliasing. **Texture anti-aliasing** is still performed by the VUE renderer on VUE objects so these options are still available.
- If you enable **G-Buffer** and/or **Multi-Pass** rendering, VUE will compute additional information. In order to save this information to a file, use the **Save to disk** option. If this option is not selected, you will not be able to see the computed information (unlike in the standalone VUE application where you can display the computed passes in the VUE interface). While this information is generated, it is not anti-aliased when rendered in the host application. If anti-aliasing is needed, it is recommended that you render using VUE standalone. Please turn [here](#) for full details on **Multi-Pass & G-Buffer** options.



xStream Integration

xStream

VUE will manage all creation, deletion, and synchronization required to edit/render the VUE scene along with your native scene. The plug-in will register several plug-in types when loaded: an object plug-in, a shader plug-in, a scene plug-in, etc. These plug-in types are written specifically to be used internally by the VUE integration plug-in. You may encounter these plug-in types in the target application interface (other than in the VUE menu), but you should never attempt to create these objects or shaders directly. If you do so, an error message telling you to use the VUE menu will be displayed, but you may also experience crashes in the worst cases.

VUE is not an import/export plug-in, so the entire VUE scene won't be converted into native format. Therefore, the plug-in is required during all editing and rendering operations. For the same reason, you can't network render a scene containing VUE content if VUE (or a VUE RenderNode – see [here](#)) is not installed on all the computers participating in the render process (as with any other plug-in).

Note:

If you are using Mental Ray and you uncheck the **Render VUE Scene** in the VUE *Options* dialog, the translated Mental Ray scene will carry no reference to any VUE elements, and you can therefore send it to other computers that don't have VUE installed on them.

In order to integrate into the target application, VUE uses commands, dialogs, and proxy objects to let you specify how you want to edit your scene. VUE also creates toolbars and icons. All VUE commands can be accessed through icons in the host application interface. Whenever we reference a VUE command in this reference manual, you can also use the corresponding icon.

VUE Commands

xStream...

The VUE commands are accessed through the VUE menu in the target application interface. All actions you may want to perform on the VUE scene will be done through the use of this menu. Most of the entries of this menu are directly taken from the main menu bar of the standalone VUE application, so you should already be familiar with them if you've used the standalone application.



You can also edit the VUE scene directly through the target application's view ports, as explained in the [Proxy Objects](#) section.

...xStream...

There are a few additional commands though, which are not found in the standalone version. Below are listed these additional commands, or commands that behave in a slightly different way than in the standalone:

In the **File** menu:

- **New...:** this will add a default VUE scene to your current project, letting you choose the atmosphere you want to add. If no VUE Scene existed yet in your project, this default VUE scene will just be added to the project (this scene only contains a VUE Camera, a Sun light and a Ground). If your project already contained a VUE Scene, the default VUE scene will overwrite the previous scene. Warning: this action cannot be undone.
- **Open...:** this will let you choose an existing VUE scene to load in your current project. If no VUE Scene existed yet in your project, the new VUE scene will be added to the project (VUE will create objects and materials in your native scene to reflect the VUE scene). If the VUE scene you are adding has multiple cameras defined, these cameras will be included in this scene. If your project already contained a VUE Scene, the newly loaded scene will overwrite the previous one. Warning: this action is not undoable.
- **Merge...:** the only difference with the **Open** command is that the newly loaded VUE scene will be merged with the existing VUE Scene, if any, instead of simply overwriting it (it is the same distinction as in VUE standalone). Here again, it will only affect the VUE part of your project, and will not clear you current native scene. And here again, this command is not undoable.
- **Close:** this will remove the VUE Scene from your project, if any. All VUE objects and materials will be removed, and if you save your project afterwards, you will lose them permanently. Here again, this command is not undoable.
- **Export object/sky/scene:** this will export the selected VUE object, the VUE sky or the entire VUE scene. Keep in mind that EcoSystems do not export.
- **Purge memory:** this will automatically reorganize the system's memory ensure memory defragmentation and cleaning up of any data that is not immediately required (for instance, if you delete a very large object, this object stays in memory in case you decide to undo this operation – by purging the memory, the object will be removed from RAM and stored on disk, until it is completely removed when the delete operation goes out of the undo list).
- **Export Animation To VUE Scene:** will open the Export Animation to VUE dialog. Please refer to the corresponding section for more information.



- **Open scene in Standalone VUE:** it will save the current scene, and will open this scene in the standalone application (launching it as a separate process). Please refer to the [Where is the Vue Interface?](#) section.
- **Reload scene from file:** if you have edited the scene from the standalone application, you can use this command to reload the scene file in the integration plugin. Please refer to the [Where is the Vue Interface?](#) section.
- **Options...:** this command is present in the standalone but it opens a totally different dialog in the case of the integration plugins. The VUE Options dialog is indeed the place where you can set all the general VUE options. Please refer to the VUE Options Dialog section below for more information.

...xStream...

In the **Object** menu:

- **Edit Object Material...:** this will open the *Material Editor* to edit the material of the currently selected object. If the currently selected object contains several materials (ie. plants or meshes), a dialog will open to let you select the material you want to edit.
- **Object Properties:** this will open the *Object Properties* panel of the selected object. This panel lets you edit properties such as light color, camera focal, etc.
- **Edit Object Graph:** this opens the VUE *Function Graph* to display the graph for the selected VUE object.
- **Edit Wind on Plant:** this opens the *Wind Editor* so that you can adjust wind amount on the plant.

In the **Windows** menu (not present in the standalone)

- **Show Material Summary:** it will display the *Material Summary* of the scene. For more information about the material summary, please refer to the corresponding section.
- **Show World Browser:** this option gives you easy access to the controls for the VUE scene's materials, bitmaps, and EcoSystems. An icon for this has also been added in the Toolbar (for all applications except Lightwave). For more detailed information about the World Browser, please refer [here](#).

In the **Maya UI** menu (only present in the Maya plug-in):

- **Restore grid and viewports:** this option will restore the default settings of Maya viewports and grid size (please refer to the section below to know more).
- **Load VUE shelf:** This will add the VUE toolbar to the interface of Maya. This



toolbar will allow you to access the VUE commands directly from icons, without the need to use the VUE menu.

...xStream

Other commands act in the same way as in the standalone: for instance, if you select **Object | Create | Sphere** it will obviously add a sphere to the scene.

Scene Display in Host Application Viewports

xStream

In order to allow the edition of large scenes, typically landscapes, the VUE plug-in can change the viewport and grid settings of the host application scene. This behavior can be changed from the VUE *Options* dialog. Before doing any operation, the integration plug-in will ask for your confirmation (you can check the “don’t show this message again” checkbox in the confirmation box if you don’t want to see it again).

Note:

In some cases, especially in the Maya application, in the case Maya was not exited normally (either killed by user or a crash), the default values are not restored, and these modified values will be used by Maya as default values.

You can use the **Maya UI | Restore grid and viewports** menu command to reset these settings. You can also restore original values by deleting your Maya preference files, which are located in your user folder (please refer to Maya documentation to know the exact location depending on the version of Maya and your operating system).

VUE Dialogs

xStream...

Some commands will have a direct effect on the scene (like the creation of a primitive), others will display a dialog: an *Atmosphere Editor*, an *Options* dialog, a *Plant Editor*, etc.

Unlike the standalone application, SOME dialogs opened in the integration plugin are modal, which means you can’t keep these dialogs open while editing the scene. For instance, in the standalone, it’s possible to change the current selected object while editing a material (resulting in the *Material Editor* switching to the material of this other object). This is not possible in the integration plugin; you must first close the *Material Editor*, select another object, and then open the *Material Editor* again.

Most of the editor dialogs (ie. *Plant Editor*, *Terrain Editor*, *Material Editor*, etc.) are amodal in VUE, which means you can continue working on your native scene while these



dialogs are still open. When you validate the change in the editor by clicking on OK, the change will be instantly reflected in the host application viewports.

...xStream

In order to simplify your workflow, and because you will probably access the VUE menu commands often, you should use the VUE Toolbars that are installed along with the plug-ins. You should also consider using keyboard shortcuts to the most often used commands.

Depending on the possible customization of the target application, you should be able to create such keyboard shortcuts for separate entries in the VUE menu. Please refer to your target application documentation to see how to do such a thing.

Proxy Objects

xStream

The integration of the VUE scene in the target application is accomplished with the help of what we refer to as proxy objects. Proxy objects are standard native objects created by the VUE integration plugin to represent the VUE objects, and to allow users to transform (move/rotate/scale) VUE objects directly from the target application interface.

Depending on the type of the VUE object, different proxy objects will be used:

- **VUE cameras:** in this case, a native camera will be used.
- **VUE lights:** a native light with a matching type (point light, spot light, etc.) will be used.
- **VUE objects:** a native polygonal object will be used. The geometry will be read from the VUE object and this geometry will be set in the proxy object mesh. In the case of procedural terrain or plants, an approximate version of the geometry will be used (the same that is used in the standalone VUE application view ports).
- **EcoSystems** are not handled through the proxy objects mechanism.:

As it's a central concept of VUE, proxy objects are further explained in the following section.

VUE Proxy Objects

xStream

In order to understand the concept of proxy objects, we will first take a look at an example of a session with VUE.



Scene Loading Example

xStream

Let's say you created a VUE scene in the standalone VUE application. This scene contains a procedural terrain, 2 plum trees, a camera, and the Sun light.

When you open this scene with the integration plugin, the plugin will first create a VUE scene object, which will act as the “root” of the VUE scene in the native scene. It will then create a proxy object for every object in the VUE scene:

- A proxy object for the terrain: a native polygonal object. The geometry of this polygonal object will represent the terrain. A native material will be created and assigned to this object, with a color matching the color of the VUE material applied to the terrain.
- A proxy object for each of the plum trees: two native polygonal objects. The geometry of the polygonal objects will represent each of the trees. A single native material will be created (plants of the same species have the same material but different geometry), with texture maps for the leaves and the trunk. This material will be assigned to both proxy objects.
- A proxy for the camera: a native camera. The camera properties will be set according to the properties of the VUE camera.
- A proxy for the sun light: a native directional light. Light color and other properties will be set according to the properties of the VUE sun and atmosphere.

When the plugin has finished creating these proxy objects, you will see in the view ports of the target application your whole scene: the native objects and the newly created proxy objects, giving you a fairly good representation of how your scene will look when rendered.

Several important notes about proxy objects:

- Proxy objects for VUE objects other than cameras or lights are flagged as non-renderable in the target application: they will therefore be ignored by the native renderer, avoiding a conflict with the VUE renderer. If you manually reset this option, then proxy objects will be rendered by the native renderer at the same position as the corresponding VUE objects, leading to strange effects.
- Proxy lights will affect the native objects, and will cast shadows depending on the “cast shadows” property of the corresponding VUE light.
- you should never attempt to create such proxy objects by yourself (the same is true for the volumetric shader used by VUE).



VUE Splines in a Host Application

There is no VUE menu command to create a spline in the host application. In most applications, native splines can be used exactly as VUE splines. They are converted internally and appear in the VUE *World Browser* in a “native splines” layer – so, you can then double-click them to add EcoSystems, terrain effects or geometry exactly like in VUE Standalone. The edition of the spline itself, however, must be performed directly in the host, with native tools.

If you load an existing VUE scene which already holds splines, the latter splines will be converted into native splines and will appear in the views. They can be manipulated with native tools, and the modifications will affect the initial VUE spline, and thus, will have an impact on the **Terrain**, **EcoSystem** or **Geometry** effect(s).

Due to limitations in LightWave’s SDK, this feature is not implemented in LightWave.

xStream

Editing the VUE Elements

Now, let’s say you’re not happy with the position of one of the plum trees. You select the plum tree proxy object in the target application’s view port and with the native transformation gizmos you move the object to a different place. If you render the scene, the integration plugin will detect that the plum tree has moved and it will update the VUE scene accordingly, rendering the plum tree at its new location.

Proxy objects can indeed be freely edited or used with the target application tools:

- Transformation gizmos,
- Align tools,
- Scripts (MEL, MaxScript, JScript, VBScript, Python, etc.).

Note:

In Cinema4D, you must be in **Object Tool** mode in order to resize VUE elements. If you are in **Model Tool** mode, the resize gizmo will have no effect on the VUE elements.

xStream

xStream...



Special Case for Lights and Cameras

...xStream...

You should not edit the type of the proxy lights (direction, point, etc.) or proxy cameras (free, target, etc.) created by the integration plugin.

Additionally, you should not edit light or camera settings directly from the host application. You should instead use the **Object | Object Properties...** menu command. This command will open the *Object Properties* dialog of the selected object, letting you making your changes. When closing the dialog, the changes you've made to the object will be reflected on the corresponding proxy object.

Here is the list of parameters set in the proxy lights and cameras:

- Light position and orientation,
- Light color,
- Light Shadow (on or off),
- Spot Light cone angle,
- Camera position and orientation,
- Camera field of view and target point (focus point).

If some parameters in the host application are not set by the plugin, you can freely change them in the host application. For instance, the type of shadows is not set by the plugin, so you can select your preferred method (raytracing, shadow map, etc.).

You will probably notice a difference between the light color set in VUE and the light color applied to the light proxy. This is because the light color of a proxy is actually computed using several different parameters to create the final color that best matches the corresponding VUE light. These parameters include:

- The actual color of the VUE light,
- The VUE atmosphere **Color shift** setting if used for this light,
- The VUE atmosphere **Global Light Balance**,
- A decay that is applied near horizon with spectral atmospheres.



If you want to freely edit the native light settings without changing the VUE light ones, you should refer to the *Native Lights Options* dialog described later in this document. This dialog let you specify the way the integration plug-in translates the settings of lights between the two applications.

...xStream

Animating Your VUE Elements

xStream

Proxy objects can also be freely animated using the animation tools in the target application: ie. you can create a keyframe for the position, rotation or scale of a proxy object. You can for instance select the proxy object of the VUE sun light and animate it with the timeline or with a script.

Animation is entirely managed by the host application, which means keyframes are stored in the native scene, and not in the VUE scene. If you want the animation to be stored in the VUE scene, you can use the **File | Export Animation To VUE Scene**.

Because each application use its own interpolation system, exporting animation from one application to another is done using a keyframe per frame, to ensure the animation of your object is not modified by a different interpolation.

For the same reason, when you load a VUE scene containing animated objects, proxy objects will be animated using a keyframe per frame.

Further Edition of the VUE Scene

xStream...

If you want to create another tree in your scene, simply select the **Edit | Create | Load Plant Species...** command in the VUE menu. A tree will be added to the VUE scene, and the VUE plugin will create a new proxy object for this tree.

Note:

In order to properly duplicate a VUE proxy object, you need to use the commands from the VUE menu (VUE

If you use the host application commands (by using the host application's own menus, or by simply pressing the copy/paste shortcut keys), this will result in an undefined behavior concerning the newly created object.

Explanation: the newly created object will be a regular native object, it won't be



managed by VUE or present in the VUE scene. So, VUE won't render it. So why does the native renderer ignore it as well? Proxy objects have a ignore from render flag enabled (to avoid both renderers from rendering it) which is copied into the duplicate object, which explains why this object is also ignored from native renderer.

...xStream

You can remove a VUE object from the scene by selecting it and deleting it directly in the target application. The deletion of the object will be detected by the plugin, which will remove the corresponding object from the VUE scene.

Proxy objects' geometry can be used for a variety of tasks:

- To drop a native object onto their surface,
- To grow some native hair/fur on VUE objects, etc.

Because EcoSystems are not handled through the proxy objects mechanism, this means that it's not possible to select an EcoSystem instance or use the geometry of an instance. It's also not possible to access the individual instances in the target application.

Display of the Proxy Objects

xStream

Proxy objects are displayed natively by the application. This means that whichever display mode you're choosing for the application (OpenGL, Direct3D, Heidi, etc.) or for a view port (wire frame, smooth, flat, etc.), the display of the VUE object will match the display of other native objects. It also means that the target application is entirely responsible for the display of the VUE objects, without the expense of calling a plugin to display its objects.

In order to get the best looking VUE objects, you should enable transparency (especially for *SolidGrowth* plants) and advanced OpenGL/Direct3D effects.

Below are listed options you should enable to get the best of each target application when using xStream.

3DS Max

VUE is fully compatible with Direct3D and OpenGL, so you can keep your preferred engine (Direct3D generally offers best performance in 3DS Max).

xStream



Cinema4D

It is suggested to enable the **Enhanced OpenGL** option in the Display menu of the view ports. This will especially enable transparency.

xStream

LightWave

No option is required to enable the correct display of VUE objects. But it's possible your VUE scene is too small in LightWave when loading or creating a VUE scene. In such case, you should open LightWave display options and change the grid size. You can also disable the **Fixed Near Clip Distance** option.

xStream

Maya

It is suggested to enable the **Shading | Hardware Texturing** option of the view port to enable the display of texture maps (color and transparency).

xStream

Softimage

There are no specific requirements in Softimage for the proper display of VUE proxy objects.

xStream

Editing Proxy Objects

xStream...

Geometry of the proxy objects can be edited in the target application but it has no effect on the “real” geometry stored in the VUE scene. If you want to edit the geometry of terrains, plants or texts, you can select the **Object | Edit Object** VUE command, which will display the appropriate *Object Editor*. When closing the editor, the geometry of the proxy object will be updated according to the changes you made.

You can still edit the geometry of the proxy object in the target application (by moving the vertices or applying a bend modified for instance) for specific reasons, but some VUE actions (like changing from Low quality to High quality meshes, global scale change, etc.) may trigger a geometry update and therefore overwrite your changes. Furthermore, this kind of edition directly on the proxy object (without using the VUE menu) violates the principle of letting the UE plugin handle the whole life of the proxy objects; you should be really careful doing this.



Also, you can edit the materials applied to the proxy objects, but this won't have any effect on the VUE scene: changing the color of an object is only possible through the VUE menu: **Change Object Material** or **Edit Object Material** entries.

...xStream

Just like for the geometry, you can still edit the native material manually. This is less critical than for the geometry so you should not experience problems changing a texture map or a color in the material. But, because VUE may want to access the material after its creation (even if you modified it), you should never delete the material manually.

It's important that you keep in mind what you can do, and what you cannot do with proxy objects:

You can:

- Edit a proxy object transformation matrix (position/orientation/scale),
- Edit proxy light properties,
- Edit proxy camera properties.

You shouldn't:

- Edit proxy object geometry,
- Edit proxy object materials.

Scene Updates

xStream...

The target application only gives the control to the VUE plugin when specific events are received.

These events are:

- An VUE menu command,
- When a proxy object is deleted (depending on the versions),
- At various stages of a scene loading or saving,
- At various stages of a render.



Apart from these events, the VUE plugin will never get the control of the native scene. And, more importantly, the VUE plugin can never decide on its own to act on the native scene.

...xStream

Before transferring the control back to the application, VUE updates the native scene with any changes that should occur during the processing of the event. These scene updates are therefore not automatic; they are always triggered by a user event.

This behavior has an important side-effect: during normal edition of the native scene, such scene updates will never occur. For instance, if you move a proxy object, the plugin will only detect the scene has changed during the next VUE action (the beginning of a render session, a VUE menu command, etc.). For this reason, a few effects may require you to trigger such an action before showing up, to force a scene update.

Indirect effects fall into this category: changing a property “A” on object 1 has an effect on property “B” of object 2 (object 1 and 2 can be the same). In the standalone application, any change on property “A” automatically triggers an update of the property “B”. For instance, if you move the VUE sun close to the horizon in a spectral atmosphere, the sun light color will change (from very bright white at noon to dark red/brown at dusk). In the integration plugin, if you move the proxy object standing for the VUE sun light close to the horizon, its color won’t change. It’s only after a scene update that the color will be updated.

It can be disturbing at first, but it has no real impact on the scene edition. In the case of the previous example – the sun proxy – if you haven’t forced a scene update after moving the sun, the color would still have been updated if you have launched a render (because initializing a render falls into the category of events triggering an update of the scene). For this reason, you should never *have* to force a scene update (if you really want to, you can still open the VUE *Options* dialog and close it; it won’t have any effect on the scene apart from triggering a scene update).

xStream...



xStream EcoSystem Painter

...xStream

It's possible to paint EcoSystems directly in the view ports of 3DS Max, Cinema 4D and Softimage. At the time of writing, painting is not yet supported in Maya or LightWave due to limitations in the corresponding SDKs. We are working with Autodesk and NewTek so they can add the hooks in their SDKs that are required to enable this feature.

During the painting session, you can freely paint on the surface of every object in your scene, including native objects. You can also paint in all views: orthogonal, perspective, or camera views.

Painting EcoSystems is done by the use of a specific painting tool plugin. During normal edition of a scene, this tool is disabled. Once you start a painting session, this tool becomes enabled and you can start painting in the view ports. There are two ways to start a painting session: you can paint an EcoSystem material or you can paint the (unique) global EcoSystem.

When editing an EcoSystem material using the *Material Editor*, clicking on the **Paint** button will activate the painter tool and open the *EcoSystem Painter* dialog.

You can also paint the global EcoSystem: to do this, simply select the **Edit |Paint EcoSystem** menu entry.

During a painting session, you can freely move around the objects, enlarge a view port, etc. All the actions done through the use of the left mouse button are generally not possible (because painting is done this way) whereas actions using other buttons are still possible.

While painting, just like in the standalone version, the *EcoSystem Painter* dialog will remain visible, to let you adjust the painting settings on the fly. This dialog is used to add species to the EcoSystem, or to set general options. You can for instance select if you want to paint only on the underlying object or on the whole scene (including native objects).


Once you have finished painting (when stopping the painting session), if you want to continue painting, you'll have to re-activate the painter tool by following the above steps.

The VUE *EcoSystem Painter* tool is implemented differently in each supporting application. Here is how to use it once it's activated in each application.



Painting EcoSystems in 3DS Max

3DS Max Painting tools work better if you look at the objects from a reasonable distance. If you are too close to an object (especially if it's a large object), you may not be able to paint on it at all due to a limitation in the 3DS Max Painting system; A good rule to follow when you want to paint on an object is to have it entirely visible in the view port.

On the *EcoSystem Painter* screen, the icon  next to the Airbrush option, to enable/disable the painting on selected native objects because sometimes complex geometries can slow down the painting. When it is enabled, the user can select which objects he wants to be processed in the core of 3DS Max, in order to accelerate the painting.

xStream

Painting EcoSystems in Cinema4D

Select the **Painter** menu entry to activate the painter tools.

xStream

Painting EcoSystems in Softimage

When the painter tool is activated, you can access the painter features using the following keys:

- While keeping the **b** key pressed, you can paint with the mouse by using the left mouse button.
- Press the **v** key to end the paint session

xStream

Notes about Painting

Depending on the implementation of the painter tool in each of the supporting applications, it may not be possible to paint on some objects. In this case, try the following:

- Duplicate the object so you have a backup of your object,
- Convert the clones into polygonal objects,
- Paint your EcoSystem onto these objects,
- Delete the clones or hide them from render if you want to keep them for later re-use.

xStream...



This may or may not be possible to convert the objects to polygonal objects. Please refer to the documentation of your target application to get more information on this operation.

...xStream

Rendering with VUE xStream

xStream

Rendering of the VUE scene elements is done through the use of a single volumetric shader added to the native scene. Depending on the target application, this volumetric shader is either applied to the cameras, to the whole scene/environment, or to a volume object.

The VUE volumetric shader is responsible for rendering VUE objects (and EcoSystems), the VUE atmosphere, and the VUE sky in a single pass. It is possible to enable these elements independently in the VUE *Options* dialog.

The rendering process with VUE is a constant communication between the native renderer and the VUE Rendering Engine. Thus, VUE objects will cast shadows onto native objects; native objects will reflect VUE objects, etc.

Once your VUE scene is loaded, you can render it without any additional setups (other than selecting a VUE-compatible renderer). Some options (like turning ray tracing on) are automatically set when starting a render.

General options such as image size, output file, etc. are entirely controlled by the native renderer, so you don't need to change your regular workflow.

If you selected the **Adjust native renderer settings to match VUE scene** option in the VUE Options dialog, the image size will be based on the image size set in the VUE scene, so you won't have to change it manually after loading a VUE scene.

If you're trying to render with a renderer that is not compatible with VUE, you will either get an error message from the renderer, or the VUE scene won't be rendered without notice. If you're using a compatible renderer, but you're using a non-activated version of VUE, the VUE scene will be rendered with a watermark applied to the image.

xStream...



Render Quality

...xStream...

The main difference between a render in the standalone VUE application and in the integration plugin is that primary rays are cast by the native renderer. This means that the anti-aliasing is entirely handled by the native renderer, and not at all by VUE. This is why changing the render mode used by VUE will have less influence in the integration plugin than for a standalone render.

Sampling quality of the native renderer will affect primary rays that hit VUE objects (they're cast from the camera by the native renderer), but it won't affect sub-rays (reflections, refractions, shadows) cast by VUE objects. These sub-rays are cast by VUE, so their sampling quality depends on the VUE render quality.

If aliasing appears on an object, and you would like to increase the sampling quality to avoid it, keep in mind the following:

VUE objects are always shaded by VUE, and native objects are always shaded by the native renderer.

So:

- If aliasing appears on the shading of a VUE object, increase the VUE renderer sampling quality by selecting a higher quality mode,
- If aliasing appears on the shading of a native object, increase the native renderer sampling quality.

If aliasing appears on a shadow, don't look for the object that's casting it. You only need to look at the actual object that's being shaded (the object which receives the shadow). This is the same for reflections and refractions.

We're talking here about aliasing on the shading of the objects, not on the edges of the object. If you want to reduce edge aliasing, then you need to increase the sampling quality of the native renderer.

To review, if you want to generate a draft render:

- Use a draft preset in the native renderer,
- Use a draft preset in the VUE renderer to avoid casting lots of sub-rays or computing advanced effects.

If you want to generate a final render:



- Use a high quality preset in the native renderer to increase object anti-aliasing and compute advanced effects,
- Use a high quality preset in the VUE renderer, to increase the quality of texture anti-aliasing and advanced effects.

...xStream

To get consistent results, you should therefore leave the automatic render quality option enabled in the xStream options dialog. The preset of the VUE renderer will be based on the current settings of the native renderer.

G-Buffer and Multi-Pass

xStream

On top of the actual picture being rendered by the native renderer, VUE xStream can compute and save additional information about the VUE objects being rendered, using Multi-Pass and G-Buffer rendering. You can use the *Multi-Pass and G-Buffer Options* dialog (accessible from the *Render Options* dialog) to select which additional information you would like to generate.

Several notes:

- Additional information about native objects can't be included this way, because the VUE Render Engine only renders VUE objects.
- Unlike in the standalone VUE application, these passes can't be anti-aliased. Because the integration plugin has no control of the order a render is actually performed, the plugin would have to keep a several gigabytes buffer in memory to be able to anti-alias each channel.

Post-Processing

xStream...

VUE is not compatible with VUE's post-processor or the native renderer's post-process pass.

This means that:

- VUE objects that use post process (like objects with a glowing material) won't be rendered correctly (the object will be rendered without glow in this case),



- VUE Post-Processing Effects won't be applied to an image rendered with VUE,
- Native Post-Processing Effects might not work on VUE objects, depending on the actions performed by these post-process passes.

...xStream

There are two exceptions:

- The natural film response filter applied to each pixel can be enabled,
- The lens flare of VUE lights can be computed.

Both options are available in the VUE *Options* dialog.

Advanced Rendering Effects

xStream

Global Illumination

Global Illumination is computed by both renderers:

xStream...

- The native renderer computes the Global Illumination on the native objects, using the information of the entire scene. VUE objects contribute to the global illumination of the scene as any other native object, except if you checked the **Disconnect VUE from native renderer final gather** option in the VUE *Options* dialog. If you haven't checked this option, you can scale the contribution of the VUE objects onto native renderer Global Illumination from the VUE *Options* dialog.
- VUE's renderer computes the Global Illumination on the VUE objects, casting global illumination rays into the entire scene. Native objects contribute to the global illumination, except if you checked the **Disconnect native scene from Vue final gather** option in the VUE *Options* dialog. If you haven't checked this option, you can scale the contribution of the native objects onto VUE Global Illumination from the VUE *Options* dialog.

Therefore, if you want to use Global Illumination for the entire scene, you need to turn on the Global Illumination in both renderers. In the native renderer, this is done the usual way. In VUE, Global Illumination is specific to the atmosphere, so if the scene you loaded contains a Global Illumination atmosphere, Global Illumination will be computed during render. You can edit the atmosphere settings using the *Atmosphere Editor* (see [here](#)).



The VUE render engine will generally compute its prepass before the native renderer begins to render the image. Because shader plug-ins generally don't have a way to display images directly in the native renderer frame buffer, this (lengthy) operation has no feedback. If possible, the plug-in will display progress information about the prepass computation. In the V-Ray plug-in, the plug-in has access to the frame buffer, and you will see the prepass computation as you would in the VUE standalone.

...xStream

With the Mental Ray renderer, depending on several settings, you may or you may not see the VUE elements during the Mental Ray Final Gather pass. It actually depends on the following settings:

- if VUE is not set to do a prepass (standard illumination mode or preview preset), you will see VUE elements showing up in the render.
- if you're rendering only a part of the image (region render), VUE will do its prepass before Mental Ray Final Gather pass, in which you won't see VUE elements showing up.
- if you're using an existing Final Gather map (final gather lock) in Mental Ray, then VUE will also have to do its entire prepass before actual rendering takes place.
- in all other cases, both Mental Ray and VUE will do their prepass at the same time, which will greatly improve the memory usage. In this case, you will see the VUE elements showing up during the mental ray prepass.

Whether the VUE elements are visible or not during the Mental Ray final gather pass will not impact the final beauty pass.

To sum up, if you don't see VUE elements showing up during the Mental Ray Final Gather pass, there is no need to dig into the settings to change this. You just have to wait for the final gather pass to end to then see your VUE elements in the final pass.

Saving Global Illumination Prepass for Reuse

If you will be rendering over a network using *RenderNodes*, it is possible to set the *RenderNode* to only render the GI prepass and save it in order to reuse it for further render jobs. See [here](#) for rendering commands and more information.

xStream

xStream...



Motion Blur

VUE is compatible with native Ray-Traced Motion Blur. This means Motion Blur on VUE objects works if it is performed by casting additional rays during render. It can therefore greatly increase the rendering time. Because it is not possible for VUE to pass post-processing information up to the native renderer, VUE objects cannot be computed during the Post-Processing pass, so native Scanline Motion Blur or native Image Motion Blur (also known as 2D Motion Blur) will not work on VUE objects.

...xStream

If you activate Ray-traced Motion Blur in the native renderer, VUE objects, as well as VUE shadows and reflections, will automatically display motion blur. Motion Blur of the VUE scene cannot be enabled on a per object basis.

Motion Blur with 3DS Max

In order for the VUE render engine to initialize Motion Blur correctly, motion blur needs to be enabled for the VUE scene object (in the *Object Properties*). If this is not the case and you selected the option **Blur all objects** in the *Mental Ray Settings* dialog, Mental Ray will still try to apply Motion Blur on the VUE scene (resulting in a longer render time), but VUE objects won't be blurred.

xStream

If you don't want to compute Motion Blur on the whole scene, you can do the following:

- Turn on Motion Blur in the *Mental Ray Settings* dialog, and deselect the **Blur all Objects** option,
- Disable Motion Blur for the VUE Scene object (so Motion Blur isn't computed for VUE objects)
- Disable Motion Blur for the rendering camera (so Motion Blur isn't computed for the VUE sky)
- Enable Motion Blur (object Motion Blur) for the Max objects you want to be blurred.

Motion Blur with Cinema 4D

In C4D, motion blur is applied by using the Scene Motion Blur effect in the C4D render settings. It will render several passes, with a time increment, and will then blend all the passes into a final image, therefore creating some blur on moving objects.

xStream



Camera Motion Blur with Mental Ray

In Maya and Softimage, by default, camera motion blur in Mental Ray is not computed using raytracing, but using scanline to speed up rendering. This is not compatible with VUE so you need to make sure the motion blur is computed by Mental Ray's raytracing engine.

In Maya, open the Mental Ray *Render Options* dialog, and in the **Rendering features** section, select **Raytracing** instead of **Scanline** or **Rasterizer**.

In Softimage, in the *Render Manager* dialog, in the **Rendering** tab of Mental Ray settings, select **Raytracing** as the **Type of Primary Rays**.

In 3DS Max, scanline is automatically disabled by the VUE plug-in.

xStream

Motion Blur Settings in Cinema 4D

If you want to add motion blur to renders in Cinema 4D, the setting, which is **Scene Motion Blur**, is found in the Cinema 4D *Render Settings* dialog. Look for the **Effects** setting (in version 10, it is found under **Effects**, the **Post Effects** droplist. In version 11, it is found under **General**, the **Effects** button). It will render several passes, with a time increment, and will then blend all the passes into a final image, therefore creating some blur on moving objects.

xStream

Depth of Field

VUE is compatible with the native renderer's Depth of Field option. You can enable it in the Render options and/or in the camera parameters, depending on the application. Depth of Field will be applied to both VUE and native objects.

xStream

xStream...



Saving Your Work with VUE xStream

...xStream

Information about the VUE scene (render options, position of the proxy object in the native scene, etc.) will be saved along with the file. Please note that unless you selected the option to incorporate the VUE scene inside the target application scene (see the VUE Options dialog), only the path to the file is saved, and not the actual scene.

This means that if you edit the VUE scene externally, it will be modified the next time you load your native scene. But it also means that if you delete or move this VUE scene file, the application will fail to find it when opening your native scene. You will then be prompted to browse to the new location of the VUE scene file. If you re-save the native scene, it will now point to the new location.

The same happens if you share the native scene with another user. For this reason, you should always use the **Incorporation** option when sharing your scene.

Automatic Backup of 3DS Max

xStream

When using the **AutoBackup** feature of 3DS Max, VUE will automatically enable the incorporate option for these backup files, whichever mode is selected for regular saving. This way, backup scenes will contain all the necessary information to prevent anything from being lost.

Incorporating VUE Scenes inside Native Scenes

xStream

If you want to incorporate the VUE scene inside the native scene, you can use the corresponding option from the VUE *Options* dialog. The VUE scene will be stored in a compressed format and will contain all incorporated texture maps so you can easily share the native scene with other users or computers. If you want to network render with plug-ins using xStream render node licenses, you first need to save your scene with this option enabled. You can then freely distribute the scene to the render nodes. When using BackBurner for 3DSMax, VUE will automatically save the Max scene with the VUE scene incorporated, so you don't have to manually select this option.

Additional Content

xStream...

If you use copy-protected content (such as content found on the [Cornucopia3D](#) website) in your VUE scene, other users:



- Will be able to load your scene, but any copy-protected content they don't own will be replaced by dummy objects or materials,
- Will be able to load the entire scene if they acquired the same content for their seat.

...xStream

If you have several VUE licenses installed on one or more computers, each plug-in will be able to load copy-protected content acquired for use in VUE.

Sharing Scenes with Other Users

In order to open a scene containing VUE information (either incorporated or not), a VUE plug-in is required. If this plug-in is not installed in the application that's trying to open the scene, an error message will appear. Non VUE elements, however, should still be read correctly (depending on the standard behavior of the native application).

xStream

Export Animation to VUE

xStream...

This dialog is accessed through the **File | Export Animation To Vue** menu command. It is used to export the animation information of all the proxy objects to the VUE scene. This reads the transformation information of the proxies for each frame of the animation and sets the transformation information in the VUE scene (as keyframes).

It is especially useful in the following cases:

- You have setup an animation on a proxy object using the integration plugin, you want to share your VUE scene with another user, and you want this scene to integrate the complete animation on the corresponding VUE object.
- You want to network render your scene with the mental ray bucket rendering option and you have to copy the VUE scene on all your computers. Please refer to the [Network rendering](#) section for more information on how to setup VUE for use with mental ray bucket rendering.

The dialog consists of three editable fields:

- **Frame Start** (default to 0): set the start of the animation.



- **Frame End** (default to 100): set the end of the animation.
- **Frame Step** (default to 1): set the step of the animation.

...xStream

For example, if you set **Frame Start** to 10, **Frame End** to 20 and **Frame Step** to 5, and click **OK**, the VUE plugin will perform the following:

- It will set the current time to frame 10,
- It will read the transformation information of all the proxy objects present in the native scene. It will set a keyframe in VUE for the corresponding object at the corresponding time.
- It will do the previous step for frame 15, and then for frame 20.

A message will appear at the end of the processing to let you know the result of the operation.

Converting VUE Objects into Native Objects

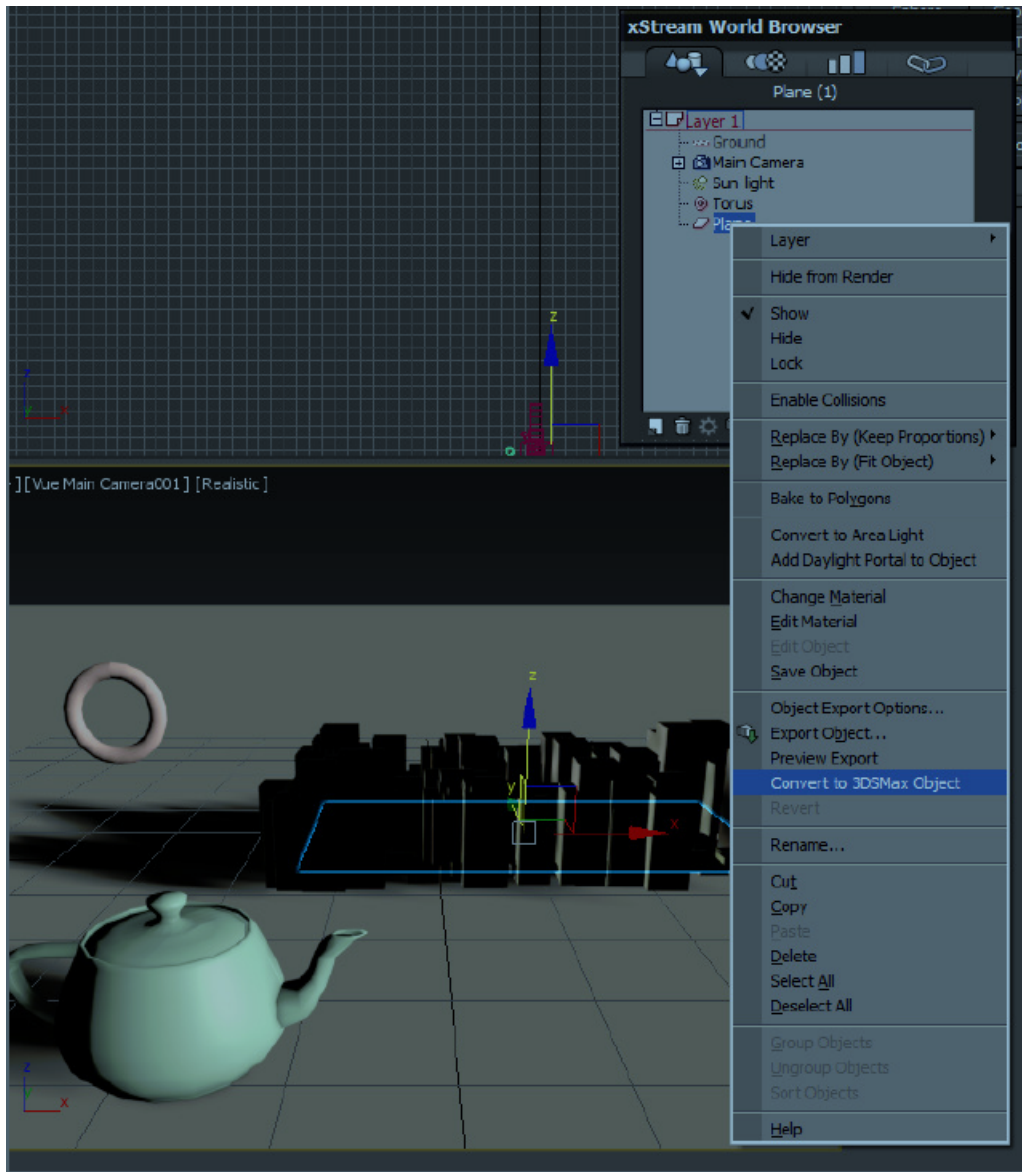
xStream...

You can now convert VUE objects (including EcoSystems) into native host objects.

This reduces the dependency on the VUE xStream plug-in (still useful for network rendering), or to edit the converted objects with your 3D host application's tools.

This feature is available in the VUE xStream plug-in. At this time it is available only in 3DSMax, Maya and Cinema4D.

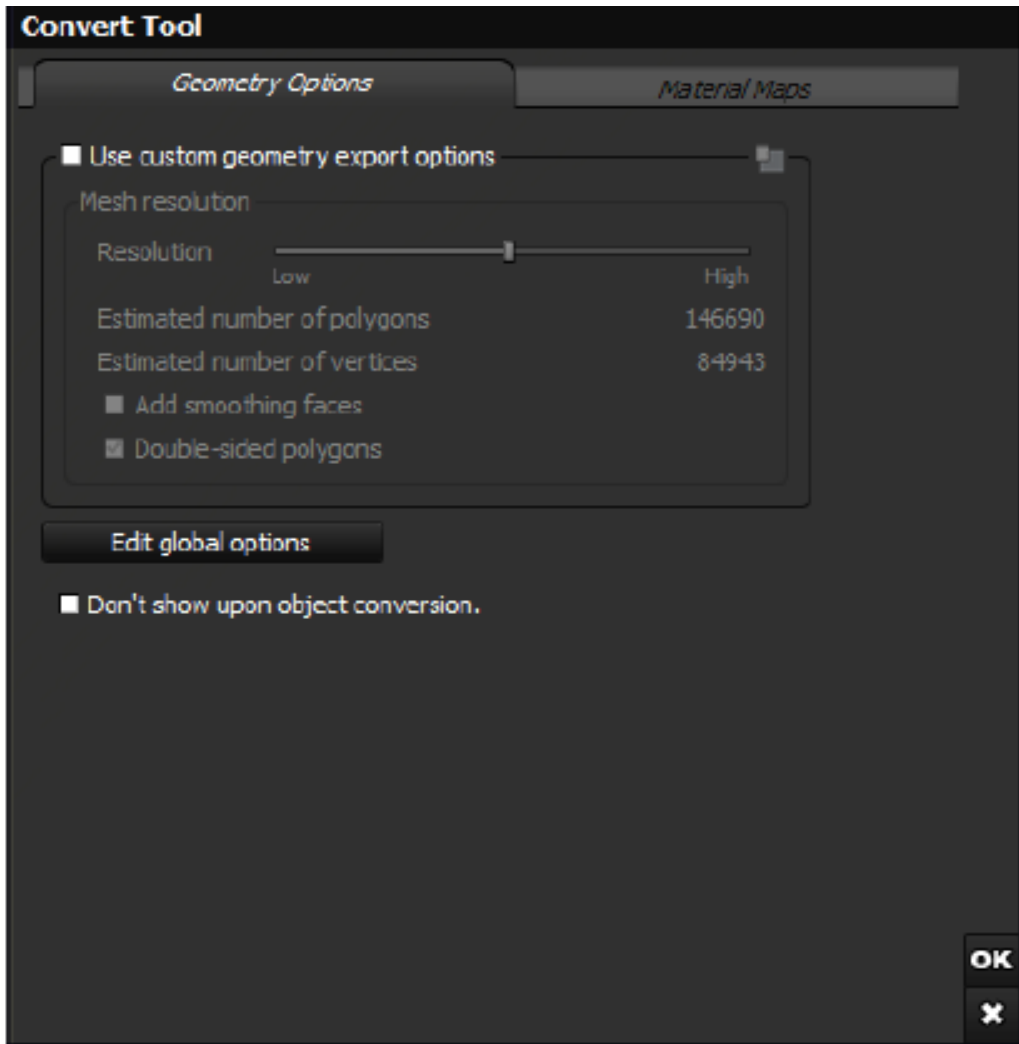




...xStream...

Selecting the Option





...xStream...

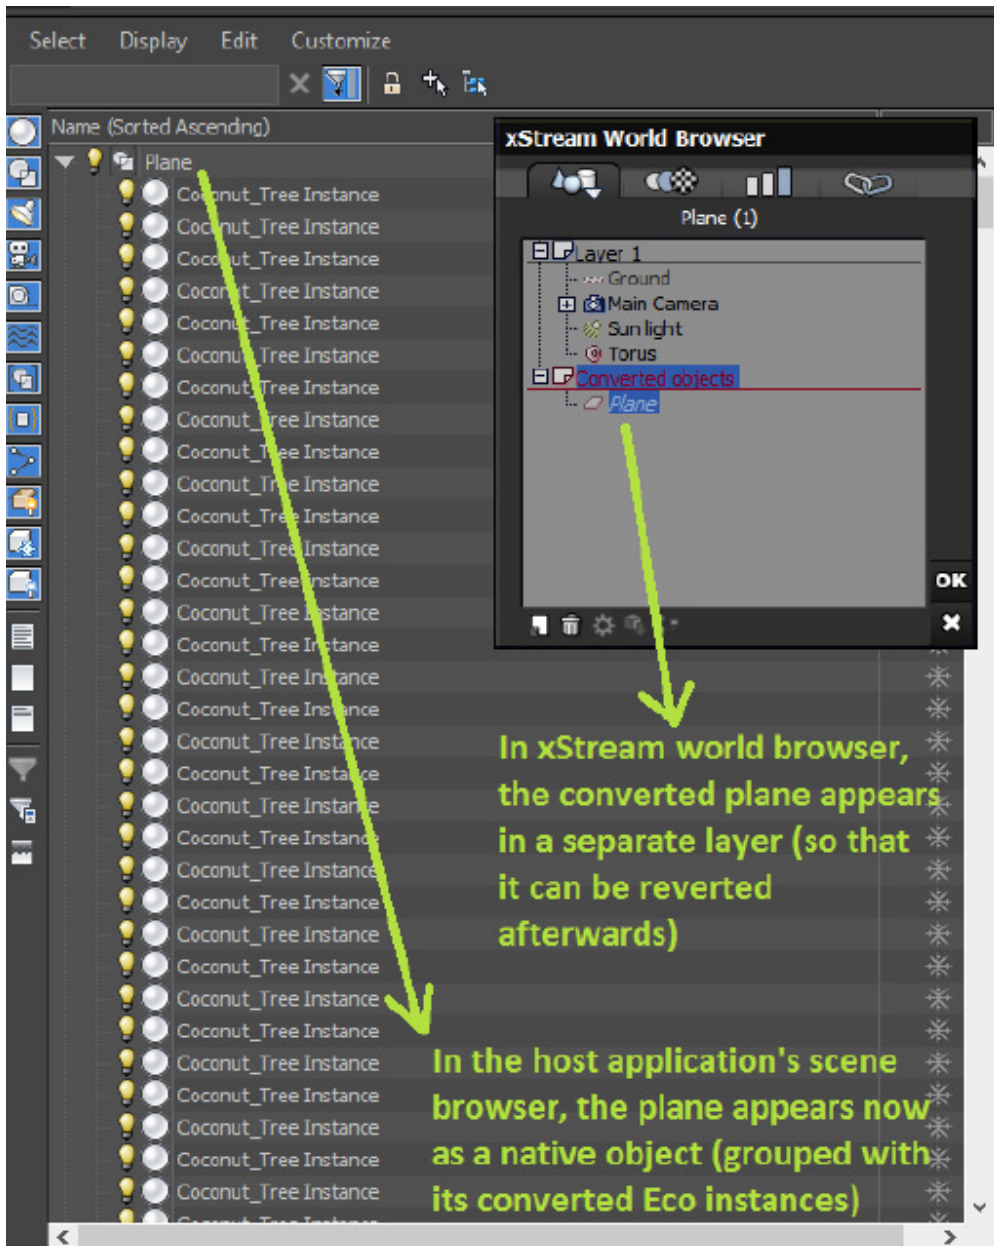
Conversion Options

Open the xStream World Browser and select the object you want to convert, using the contextual menu (-> Convert to XXX Object)

When you convert an object, VUE offers to customize geometry and material maps options. This editor is quite similar to what you get for the [Export option](#).

Once converted, the object appears among the native objects.

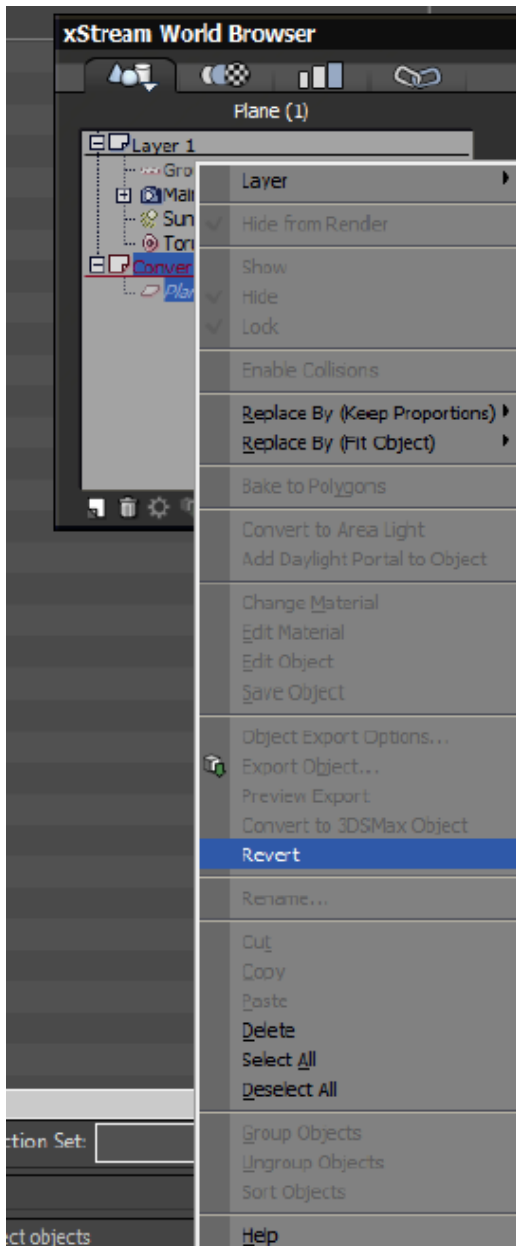




...xStream...

After Conversion





After Conversion

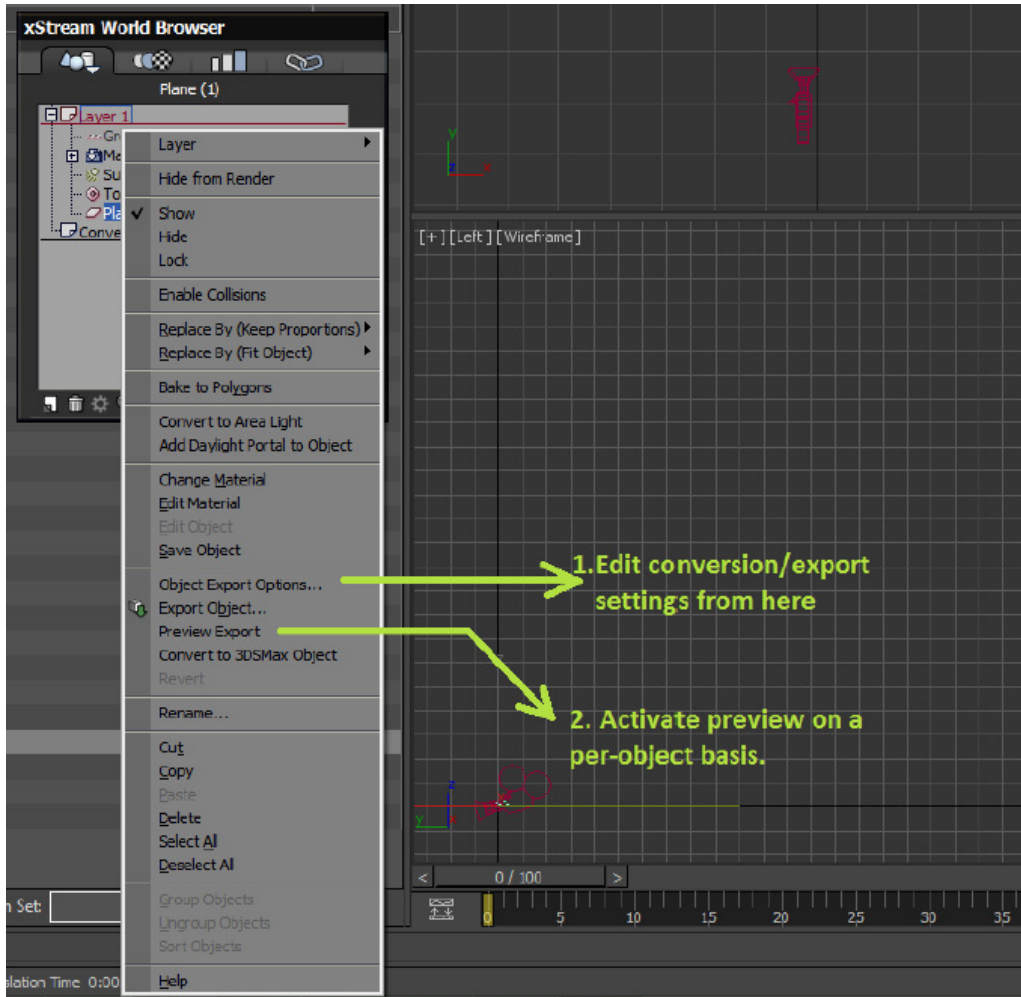
...xStream...



Its EcoSystem instances are now also native objects, and are grouped with it.

The object is no longer rendered by the xStream render engine: it behaves now like any other native object and cannot be edited anymore with VUE tools. With VUE content converted to a native format, you can now use any renderer you wish such as Arnold, V-Ray RT, Octane, Renderman, etc.

...xStream...



Converting back to a VUE object

Afterwards, you can choose to revert the object, so that it becomes a VUE object again. From the World browser, select the object (part of the Converted objects layer), right-



click to open the Object menu, and choose the Revert command.

...xStream

FBX Import

xStream

Objects

xStream

FBX supports have more than one object in a FBX file. So when you import a FBX file in VUE, a group named after the FBX file name will be created containing all objects inside the FBX file (except for Cameras). It is possible to import scenes with textured geometries, but other objects (like lights, characters with skin and bones, morphers, etc...) are not supported. Some materials may not be well imported (it can be caused when the exporter writes data specific to its application only).

Camera

xStream

When importing a FBX containing cameras, the cameras will be added to the camera list.

Camera Synchronization

xStream...

Unfortunately, the camera settings such as Focal Length or Field of View do not mean the exact same thing in each software (Maya, 3ds Max, C4d...). To ensure you have the same image in VUE and in another software, you must ensure the image resolution is the same on the software before you export the FBX.

Example

If you want to make a 1024×1024 render on Maya and VUE from a Maya scene:

- Go to Maya, render settings, select 1024×1024 resolution.
- Go to file, export All.



- Import the generated FBX in VUE.
- Select the camera you imported.
- Go to render options, set resolution to 1024×1024 , render.

...xStream

Supported features from Maya

- Camera Scale is not supported in FBX.
- Film Offset is supported.
- Lens Squeeze ratio is supported.

xStream

Nuke EXR Exporting

xStream

To use this feature, you need to activate the [multipass rendering](#) on the **Render Options** screen. Select to use Multi-Pass Options and click the **Edit** button.

On this screen, select the items you wish to include in the render. Be sure to select the Several Files option.

Compositing

xStream

There are two scripts inside the install files for your VUE. In folder **Synchro Plugins/Nuke**, you can find **VUE Compositing.n**, a compositing script and **VUE Relighting.nk**, a relighting script. This script combines the rendering passes to recreate the image VUE create from the same data.

Just load the multipass EXR files and the “VUE Compositing.nk” script. You can then modify the different contribution of each pass at your convenience.

Note:

there are passes that allow you to make masks on the different items so you can change the contribution for one item without changing the others.



Relighting

xStream

The latest layers such as **World Point Position** allows the Nuke to recreate a 3D image of what is seen on the EXR.

The script **VUE Relighting.nk** uses the relighting nodes from Nuke.

From the .exr file, it creates the 3d image with *PositionToPoints* node, adds a light and the Relighting node.

You can then modify the light position in Nuke and the Image will reflect the changes.

Tutorial

xStream

You can find a tutorial [here](#).



E-on License Server

Node-Locked vs. Floating Licenses

Each license of VUE can perform either as a nodelocked or as a floating license:

- Nodelocked licenses are specific to your computer (they are locked to it). They are typically used if you will be running VUE on a single computer.
- If you want to use the same license of VUE on several computers, you will need an *E-on License Server*. The *License Server* is responsible for issuing temporary licenses as they are requested. VUE licenses that are installed on a *License Server* will automatically become floating licenses. If there are no licenses of the requested kind, or if no licenses of the specified kind remain available, or if the *License Server* is unreachable, the application will run in non-activated mode.

To convert your VUE licenses to floating, you will need to reinstall your software, selecting the “floating” option on the screen where you would normally enter your serial number. You don’t enter a serial number with the “floating” option. Then you need to purchase and install the License Server and add your VUE licenses to it. You can install and use VUE on the same computer that is running the *E-on License Server*.

Note:

you cannot install the same license both as nodelocked and floating. This would be a violation of the EULA.

About the License Server

The License Server can run as a service (being logged into a user account is no longer a requirement). All network services have the ability to generate detailed logs for monitoring and tracking of network defects.

On top of being able to run as a service, the *License Server* offers improved usability features such as automatic detection by VUE seats, easy management of installed licenses via a dedicated graphical interface as well as hardened network operations. If you are running the *License Server* as a service, be sure to shut it down before applying any updates. It can be shut down from the **Administrative Tools | Services** application.



Purchasing a License Server

In order to convert your VUE licenses into floating licenses, you will have to purchase an *E-on License Server*. Please refer to the [e-on software](#) website for further information.

Using a Floating License from VUE

If you haven't already installed VUE, select the **Use a floating license** option during the VUE installation.

If VUE loses the connection to the *License Server*, it will immediately try to reconnect to the license server. If, after several attempts, it still fails to connect, a dialog will popup allowing you to choose another license server. License servers can also be auto-detected.

You can choose to use VUE in non-activated mode until the *License Server* problem is remedied. This will not change the current configuration of the floating license and normal operation will be resumed next time you start VUE.



Network Rendering

This section will show you how to setup and use network rendering.

RenderCows vs. RenderNodes

There are two ways of handling network rendering in VUE Pro: *RenderCows* and *RenderNodes*:

- *RenderCows* are controlled via the *HyperVue* network rendering interface (see below). They are easy to setup and manage, but can only be controlled via the *HyperVue* interface. All VUE versions that incorporate HyperVue include a pack of 5 *RenderCows*. With the Pro versions, you can purchase additional *RenderCow* packs to extend this to as many render nodes as you need. *RenderCows* are a very cost-effective solution for setting up a small-scale VUE-dedicated render farm. *RenderCow* packs will only work with the license of VUE they were purchased for. You can install as many *RenderCows* as you want, on as many computers as you like, but your license of Vue will only be able to render on 5 of these computers simultaneously (unless you have purchased separate *RenderCow Packs*).
- *RenderNodes* are more complicated to set up, but are a lot more flexible and can be integrated into large render farm management systems. They are controlled solely by command line. *RenderNodes* require the installation of the e-on *License Server*. *RenderNodes* are ideal for large production houses that have a render farm that needs to handle jobs for different rendering platforms.




Description

HyperVue™ is VUE's network rendering system. Using *HyperVue™*, you can create your own Render Farm by distributing the rendering of your pictures or animations over a network of computers. Each computer on your network will do its share of the work. That way, if you have, for example, 3 computers (of equivalent power) on your network, the animation will render 3 times faster!

In order to take part in the rendering fun, each computer must receive a rendering node. The render nodes in *HyperVue* are called *RenderCow™* (pun intended).

HyperVue™ Network Rendering Manager



HyperVue Network Rendering Manager

The *HyperVue™ Network Rendering Manager* is the control center that dispatches rendering jobs to all the *RenderCows* on your network, and then collects the resulting images once they are completed.

To enable and configure network rendering, select the **External** renderer, select the **Network** icon and click the **Edit** button alongside it in the *Render Options* or *Animation Render Options*. This opens *HyperVue* so that you can set up your *RenderCow* network.

When you start *HyperVue*, it will try to contact all the *RenderCows* that are referenced on your network. If the *RenderCow* is running, it will appear in the list of available *RenderCows*.



Each *RenderCow*TM appears as a line in the *HyperVue Network Rendering Manager*. The line features the name of the host computer running the *RenderCow*, as well as the current status of the cow. If the *RenderCow* is available for rendering, it will be listed as ‘Idle...’. At the beginning of each line is a checkbox that should be checked if the computer is to take part in the rendering fun.

Initially, no *RenderCows* are referenced by your *HyperVue Network Rendering Manager*, so the list should appear empty (except on Mac OS X systems where *RenderCows* installed on other Mac OS X systems will automatically be detected and added through the use of the Apple’s *RendezVous*TM technology). Please read on for details on how to add other *RenderCows* to your *HyperVue*TM *Network Rendering manager*.

RenderCowTM

RenderCow is the name of VUE’s stand-alone multi-processor rendering engine. A *RenderCow* must be installed on each computer that is to participate in the network rendering.

Basically, the *RenderCow* spends its time waiting for an order from the *HyperVue Network Rendering Manager*. As soon as it receives an order to start rendering a picture, it does so and then returns the completed render. The resulting picture is subsequently retrieved by the *HyperVue Network Rendering Manager*, and a new job is issued to the *RenderCow*.

Additional RenderCow Licenses

VUE ships with 5 *RenderCows*, but you can add the ability for your license of VUE to handle more *RenderCows* by purchasing separate *RenderCow Packs*. This will let you render on more than 5 computers simultaneously.

RenderCow Packs will add the ability for your *HyperVue Network Rendering Manager* to control more than 5 *RenderCows*. They are sold as a separate extension to VUE, and are tied to your license – so they will only enable the handling of more *RenderCows* in the VUE license they were purchased for.

Note:

If you have several VUE licenses and wish to network render with more than 5 *RenderCows* from each one of them, you will need a separate *RenderCow Pack* for each VUE license. In such a case, you might like to consider using *RenderNodes* as an alternative to *RenderCows* (see [here](#) for a comparison of the two solutions).


When you install the *RenderCow Pack* into your copy of VUE, it adds the ability for VUE to control a number of additional *RenderCows* (5, 10, 25 depending on the pack you



purchased). The total number of cows that can be controlled by your copy of VUE is displayed at the top of the *HyperVue Network Manager*.

Please visit the e-on software website for details on purchasing additional *RenderCow Packs*.

Installing RenderCow Packs



Installing RenderCow Pack Licenses

To install a *RenderCow Pack*, simply select the menu command **Help | Register RenderCow Pack**. A dialog will appear with the list of registered *RenderCow Packs* for this license of VUE. Click the **Add** button and enter into the field your *RenderCow Pack* serial number – the number sent to you when you purchase a *RenderCow Pack* license. Make sure you use the *RenderCow Pack* serial number that corresponds to your installation of VUE.

RenderCow Packs can also be installed directly in the [License Server](#).



Setting Up RenderCows

Installing a RenderCow

A *RenderCow*TM must be installed on each computer that is to take part in the network rendering. To install a *RenderCow* on a computer, locate your downloaded installation files on your disk, and double-click the 'Install Vue.app' icon on a Mac or the 'Setup.exe' file on the PC. Follow on-screen instructions to complete installation.

Note:

If you purchased a downloaded version, you can copy the entire installation files folder to a shared network folder and access the *RenderCow* installation from each target computer. Or, burn your installation folder to a DVD and use that to install the *RenderCow* program on your networked machines.

If you get cryptic error messages the first time you launch the *RenderCow*, you will have to change the port number (see below) and restart the *RenderCow*.

Launching RenderCows at Boot

On the last screen of the *RenderCow* installer is an option to automatically launch the *RenderCow* at boot. It is recommended that you check this option so that the *RenderCow* will be automatically loaded each time you reboot your computer. This will save you having to go around your network to start each *RenderCow* manually. Besides, since very little resources are actually used when the *RenderCow* is idle, there is no real down side in doing so. Also, because it runs in the background and only uses extra CPU cycles that would otherwise be lost, the *RenderCow* should have very little impact on overall system performance (even when rendering).



RenderCow Port Numbers



Setting the RenderCow's Port Number

In order to communicate with the *HyperVue Network Rendering Manager*, the *RenderCow* listens to communications on a given port. By default, this is set to 5004. However, there can be cases where another application is already using that port. If this is the case, you should change the port number used by the *RenderCow* to a free value. This is done by picking the **Settings | Port Number** command from the *RenderCow* menu and entering the new port number.

Announce RenderCow: select this option if you want the *RenderCow* to be automatically detected by *HyperVue* managers running on your network. Indicate the port number that is watched by the *HyperVue* manager (5591 by default).

Please note that each *RenderCow* can use a different port number.

CPU Affinity

On machines with multiple processors, you can select how many processors/cores you would like the *RenderCow* to use. This setting is accessed from the *RenderCow* menu (right-click on the *RenderCow* in the taskbar and select **Settings | Processor count**). This setting is defaulted to the number of processors you have on your machine, but can be changed to less, if you prefer. This setting takes affect immediately. There is no need (on Windows machines), to go into the Task Manager and change it there as well. The setting will be visible in the Task Manager, however.



Installing a RenderCow on the Computer Running VUE

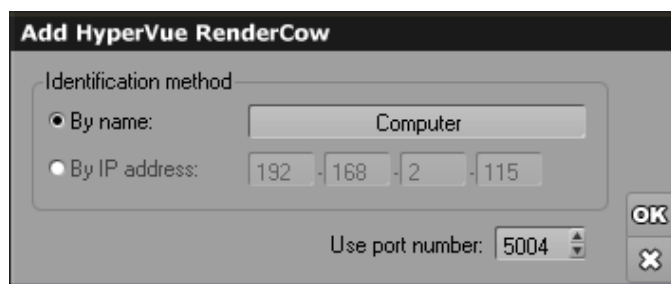
You can also install a *RenderCow* on the computer that is running the VUE program and the *HyperVue* manager. Simply reference the *RenderCow* as any other one in the *HyperVue Network Rendering Manager*.

Because the *RenderCow* runs in the background, it will not slow down the *HyperVue* manager as it manages the distribution of the rendering over your network. If the workload of the *HyperVue* manager is high, that cow may not be able to do much, though.

Rendering With HyperVue™

To access the *HyperVue Network Rendering Manager*, go to the *Render Options* dialog, press the **Edit** button alongside the **Standalone renderer** option and select **Use HyperVue network manager**. Click on the Edit button alongside that option, and the interface of the *HyperVue Network Rendering Manager* will appear (this is a separate application).

Configuring HyperVue™



Adding a RenderCow to HyperVue

Before you can start network rendering, you must first reference all the *RenderCows* running on your network. To add a specific *RenderCow* to the list of available rendering nodes, press the **Add** button. A dialog appears prompting you to identify the rendering node either by entering the host name of the computer, or its IP address. The Port



number should match that used by the *RenderCow*. Unless you had to modify the value because of conflicting uses of the default port number, you should leave this untouched.

Make sure the host is booted and the *RenderCow* is running, then press **OK**. After a short wait, the new host should appear in the list of available *RenderCows*, together with its current status. If the *RenderCow* cannot be contacted for any reason, nothing will be added to the list. One of the typical reasons that the *RenderCow* cannot be contacted is if you have a Firewall blocking access (on either end). If this is the case, you should contact your system administrator so he can allow connections of VUE and *RenderCows* to your local area network on the port numbers you selected for the *RenderCows* (see *RenderCow Port Numbers* above).

If you check the **Auto discover RenderCows** option, *HyperVue* will automatically scan your network on the indicated **Port** number, to identify and add any new *RenderCows* to the *HyperVue Network Rendering Manager* as soon as they become available. This scanning continues while a render is in progress, so new cows can be put to work as soon as they are powered up. You can change the port number at any time, in case you installed your cows on different port numbers (this is a *bad* idea).

You can remove a *RenderCow* from the list by highlighting it and pressing **Remove**. If you just want to temporarily stop a *RenderCow*, click **Stop**.

Press **Select all** to select all *RenderCows* in the list. Press **Deselect all** to deselect all *RenderCows*.

The list of *RenderCows* is stored when you close the *HyperVue* dialog, and automatically restored next time you open it. Next time you open the *HyperVue Network Rendering Manager*, it will contact all the referenced *RenderCows*. Only those that could be contacted will appear in the list, together with their current status.

You can save your list of nodes in a standalone file for future use: press the **Save nodes** to save the list of *RenderCows* to disk. Press **Load nodes** to load an existing *RenderCow* node list.

When rendering, you can press the **Minimize** button to minimize the application.



Putting RenderCows to Work

Before you start rendering, you should select which node is going to participate in the rendering. This is done by checking or unchecking each *RenderCow* on the list. *RenderCows* that are selected will participate, others will be left aside.

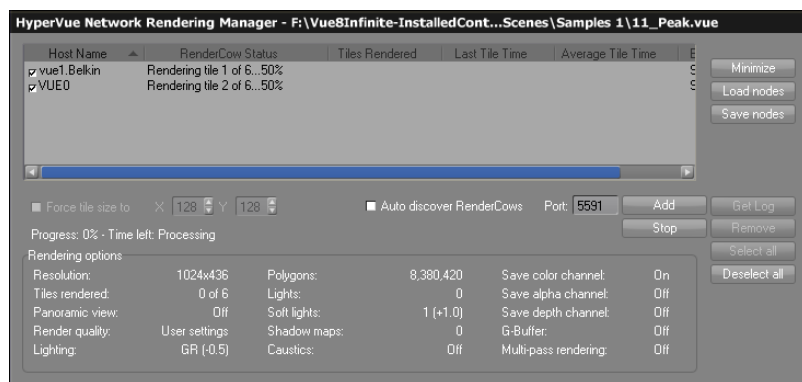
RenderCows can be added even when the rendering process has already begun – see below for details.

Starting a Render

When you are done adding *RenderCows* to *HyperVue*'s list and you have decided which cows would take part in the render, you are ready to start rendering using your network. In the *Render Options* dialog, press **Edit** alongside the **Standalone renderer** option and select **Use HyperVue network manager**, then press **Render** and watch as *HyperVue* contacts each *RenderCow* in turn, first sending the scene to be rendered, then issuing orders to render a given frame and finally collecting the rendered pictures.

Please note that the scene is only sent once to each *RenderCow*. The traffic generated on your network will rapidly drop as soon as all *RenderCows* have received the scene.

Monitoring the Render Farm



Rendering with the HyperVue Network Rendering Manager

When a *HyperVue* render is in progress, the *HyperVue Network Rendering Manager* displays information regarding each of the *RenderCows*:



- **Host Name:** this is the hostname of IP address of the *RenderCow*. Uncheck the checkbox to disable rendering on that *RenderCow*.
- **RenderCow Status:** this indicates the current status of the *RenderCow* (sending scene, sending textures, rendering...).
- **Number of Frames:** this indicates the total number of frames of the animation that have been rendered by this *RenderCow*.
- **Last Frame Time:** this indicates the render time of the last frame.
- **Average Frame Time:** this displays the average render time of all the frames rendered by the *RenderCow* since the beginning of this rendering session.

Aborting a Render

You can abort a render any time by clicking the **Abort** button. You may either opt to save the frames already rendered, or discard them. There will be a short time lapse before the rendering process aborts, because the *HyperVue Network Rendering Manager* must first contact all the *RenderCows* and tell them to stop their work. Discarding the frames that have already been rendered is faster because it doesn't require the *HyperVue Rendering Manager* to save the frames to disk. Also, please note that if you save the frames that have already been rendered, you may have some missing frames in your animation (especially if some computers on your network are significantly slower than others). It isn't possible to resume an aborted network rendering session.

Managing RenderCows™

This section describes operations that are performed on the host computer running the *RenderCow*.



Showing RenderCow Status

You can view the status of a *RenderCow* running on a computer by clicking on the *RenderCow* icon in the task bar, or selecting the menu command **Show** (Windows only). The *RenderCow* window will popup, displaying the current status of the *RenderCow*. Click the **Close** button (or select the menu command **Hide**) to store the *RenderCow* back in the Task bar.

Adding RenderCows During Render

You can easily add *RenderCows* even when a rendering process has already begun. Simply press **Add** and enter the IP or host name of the computer holding the new *RenderCow*. The new *RenderCow* will automatically be put to work and join the rendering process.

Thanks to the use of Apple's *RendezVous*TM technology, *RenderCows* running on Mac OS X systems will be detected automatically. What this means is that you can start a network rendering session, and then turn new Mac systems on. If you setup the *RenderCow* to launch at boot, it will be detected by the *HyperVue Network Rendering Manager* and immediately put to work! This cool e-on technology is known as *SmartCow*TM.

Pausing a RenderCow

RenderCows can be paused/resumed by an option on the *RenderCow* tray icon menu. The status is visible both in the RC status window, and in *HyperVue*. The *RenderCow* can only be resumed from the tray icon menu, however; it cannot be resumed from *HyperVue*.

Shutting Down a RenderCow

If you want to stop a rendering job that is currently running on a *RenderCow*, select the menu command **Shut Down**.

You can safely turn off *RenderCows* or shut down the system hosting the *RenderCow* even when the *RenderCow* is working... To shut down the *RenderCow*, select the **Exit** menu command.



Updating RenderCows

Just like the rest of the application, *RenderCows* require regular servicing (see [here](#) for details on updating your software). However, downloading and installing software updates on all the computers on your network can rapidly become a very time consuming hindrance.

Fortunately, Vue features e-on software's unique network updating technology, called *NewCow™*. This technology automatically maintains your render nodes by remotely installing all required updates on demand.

When an obsolete *RenderCow* is detected by the *HyperVue Network Rendering Manager*, it hands over control to the *NewCow* technology. *NewCow* then dispatches update commands to the *RenderCow* so that it can be updated. Because this updating process takes place in the background, it does not keep up-to-date nodes from beginning their work. The progress of the updating process is displayed in the *HyperVue* window. As soon as the obsolete *RenderCow* is updated, it will join the other nodes in the rendering fun.

RenderNodes

The *RenderNode* is a command line version of VUE devoid of any graphical user interface. *RenderNodes* are ideal for production houses that have a render farm management utility and wish to integrate VUE rendering into this render farm.

RenderNodes are controlled via the command line and by placing appropriate files at the appropriate location and at the appropriate time. Your render farm management utility will have to be configured to interact with the *RenderNodes* (e.g. so that it launches the *RenderNodes* with the required command line parameters so as to perform the desired operation).

When launched, a *RenderNode* parses its command line to find out the operation that needs to be performed, performs the said operation and then quits automatically once the operation is completed.

Some examples of configuration for typical render farm management software are supplied to help you setup your own system. You will have to adapt these examples to the specificities of your setup.




RenderNode Licenses

RenderNode licenses are sold separately. There are two types of *RenderNode*:

- *VUE Infinite RenderNode*: (renders only VUE stand-alone content), and
- *VUE xStream RenderNode*: (renders both VUE stand-alone and VUE hosted content).

You can purchase *RenderNode* licenses from the [e-on software store](#).

RenderNode Network Options



RenderNode Network Options

When you instruct VUE to render using the **RenderNode Network** option, VUE will:

- Save the current scene to a user-specified folder or to its temporary folder, then
- Execute a custom command.

The *RenderNode Network Options* dialog lets you specify the temporary folder and the command line that is executed when you launch a network rendering session using a render farm with *RenderNodes*. This dialog is accessed by pressing the **Edit** button alongside the *RenderNode* option in the *Render Options* dialog.

The command line depends on your particular render farm setup. Please consult your render farm management documentation for the specificities of configuration.

Temporary Folder

Depending on your network manager, it's sometimes possible to attach a file to a job. If this is possible, VUE will save the current scene to a temporary file, and send the file to the network manager which will store it into its own temporary folder before sending it to the nodes. If this case, you can leave this field blank (it's fine for VUE to delete its own temporary file after submitting the job, so VUE can use its own temporary folder).

However, if your network manager does not allow attaching a file to a job, VUE will just



tell the network manager where it can find the scene file. Because the network manager sends the path of the file to the nodes, and not the file itself, the file must exist during the entire rendering session. Also, the path to that file needs to be a valid network path accessible from all the computers taking part in the render.

Because the file must exist during the entire rendering session, VUE should not delete the temporary file, so it needs to be stored outside VUE's standard temporary folder (which is flushed at regular intervals). In this case, click on the **Browse** button and select an alternate temporary folder.

Command to Execute

This setting indicates the command that is executed each time you start a *RenderNode* network render.

As an example to illustrate what this setting does, if you entered `C:/Windows/Notepad.exe` in the command field, VUE would launch the *NotePad* each time you press the render button (of course, we expect something a little grander). If you want VUE to launch your network manager application, all you need to do is enter the path to the application file in the command field. Launching the network manager without any arguments will surely do nothing useful, so, basically the command should contain at least:

- the network manager executable path and file, and
- the scene file path.

It can also contain some additional information required by the network manager, like the description of the job, how to handle it, etc.

Arguments in the Command

Following is a list of arguments that can be added to the network manager command line. When running the command, VUE will dynamically replace them with the appropriate value:

- **[SCENE_NAME]**: the scene name (name of the job).
- **[NUM_FRAMES]**: the number of frames in the animation (may be required by your network manager in order to distribute rendering across several nodes).
- **[FILE_PATH]**: path to the temporary scene file. If you specified an alternate temporary folder in the *RenderNode Configuration* dialog, this folder will be used to store temporary scene files. If not, the default VUE temporary folder will be used. You should specify a folder that's accessible from all the render nodes (e.g. a network path), in case your network manager doesn't send the files to the nodes itself.
- **[UNC_FILE_PATH]**: path of the temporary scene file after conversion to UNC format (e.g. `\\Server\SharedFolder\Folder\scene.vue`). Available on Windows systems only.



If you are using a render farm of Windows computers and the scene files need to be accessed by their paths, you can either:

- select a folder on a network drive (the letter of the network drive must be the same on all the render nodes), or
- use the `[UNC_FILE_PATH]` argument so that any folder path will be automatically converted to UNC (in which case either the folder, one of its parents, or the entire drive must be shared, lest it won't be possible for the system to create the UNC path – for obvious reasons).

Sample Setup for BackBurner

Here is an example of the command to enter for *BackBurner*:

```
C:\Program Files\Autodesk\backburner\cmdjob.exe -jobName: VueJob_[SCENE_
NAME] -jobNameAdjust -description: 'Rendering of Vue scene: [SCENE_NAME]
.vue' -numTasks: [NUM_FRAMES] -workPath: 'C:\Program Files\e-on software\
Vue xxxx Infinite RenderNode\Application' RenderNode.exe -file '[FILE_PATH]
' -job %tn
```

Note:

The path to the *RenderNode* is: C:\Program\ Files\e-on\ software\Vue\ 2014\ xStream\ Application

We used both the `-numTasks:\ [NUM_FRAMES]` and `-job~\%tn` options to distribute the rendering on several *RenderNodes*. Each job sent to a node will correspond to the rendering of one frame of the animation. For this reason, the animation output file can't be a single file (i.e. **AVI** or **MOV**), it needs to be a separate file for each frame.

If we want to render the whole animation on one single computer, we can specify `-numTasks: \ 1` in the command line.

In the example above, the workpath (the folder where the *RenderNode* is installed on the render node) needs to be the same for all the nodes. This is a limitation of custom jobs with *BackBurner*.

BackBurner is used here as an example. The *RenderNode* system was designed to work with all the major network managers. It should be simple to configure any custom network manager to use *RenderNodes*, provided the application was designed to accept custom jobs.



Hints on Setting Up Your Network Rendering Manager

You should of course be familiar with how jobs are submitted to your network manager before trying to setup this feature. Here are a few recommendations before you attempt to configure your network manager:

- Read the network manager reference manual on custom jobs,
- Read and try to reproduce the sample commands that are usually described in the reference manual,
- Run a few simple tests with batch files, or simple applications (the *NotePad* application is perfect for that!)
- Finally, try to setup VUE to use your network manager, first with simple commands, and then with more elaborate commands that include all the required options.

If, after several attempts, it appears that it's impossible to use *RenderNodes* with your network manager, please get in touch with our tech support to try and find a solution.

Setting Up RenderNodes

Unlike *RenderCows*, *RenderNodes* are not tied to a specific VUE license, but have their own licensing system. So each *RenderNode* license that you purchase may be used to render scenes created with any one of your VUE seats (*RenderNodes* are not controlled directly by VUE).

Installing the *RenderNodes* is done easily using the VUE product installer. On the screen where you select to install either Infinite or xStream, select **Network Rendering**. On the next page, select which type of *RenderNode* you will be installing. And continue with the rest of the installation program screens.

In order to operate, the *RenderNode* requires the installation of an e-on *License Server* (see [here](#)). The first thing the *RenderNode* does when it is started is attempt to connect to the *License Server* and get a license. If it cannot retrieve a valid license (either because no more licenses are available, or the network connection is not functional), the *RenderNode* will not start.

Updating RenderNodes

RenderNodes are updated automatically by the e-on *License Server*.



Controlling RenderNodes via Command Line

If you want to build your own Network Rendering Manager, or for any other purpose, you can also control the *RenderNodes* directly by command line. The following options are available:

- `-file\ \textquotesingle{}[FILE_PATH]\textquotesingle{}:` specifies the scene file that will be loaded and rendered, and (optionally)
- `-frame\ X:` indicates which frame to render if you don't want to render the entire animation (if the scene is animated). If you don't use this command, if the scene is animated, the entire animation range will be rendered (the part between the start and end frames),

using the [Animation Render Options](#) stored in the file. If no animation is defined, the picture will be rendered using the [Render Options](#) stored in the file and the `-frame X` option will have no effect.

- `-range\ X\ Y:` indicates the range of frames that will be rendered (from X to Y inclusive).
- `-step\ S:` sets the rendered frame increment (e.g., if set to 10, only frame 0, 10, 20... will be rendered).
- `-job\ N:` renders frame N-1 (for compatibility reasons with *BackBurner* that numbers jobs starting with 1).
- `-output\ \textquotesingle{}[FOLDER_PATH]\textquotesingle{}:` specifies a path to a folder where all output images should be saved (it overwrites the output paths set in the scene, and only keeps the base names).
- `-output_color\ \textquotesingle{}[PATH_OR_FILE_NAME]\textquotesingle{}:` specifies either a path or a file name where color output still images should be saved (it overwrites the value set in the scene).
- `-output_depth\ \textquotesingle{}[PATH_OR_FILE_NAME]\textquotesingle{}:` specifies either a path or a file name where depth output still images should be saved (it overwrites the value set in the scene).
- `-output_alpha\ \textquotesingle{}[PATH_OR_FILE_NAME]\textquotesingle{}:` specifies either a path or a file name where alpha output still images should be saved (it overwrites the value set in the scene).
- `-output_color_anim\ \textquotesingle{}[PATH_OR_FILE_NAME]\textquotesingle{}:` specifies either a path or a file name where color output animated images should be saved (it overwrites the value set in the scene).
- `-output_depth_anim\ \textquotesingle{}[PATH_OR_FILE_NAME]\textquotesingle{}:` specifies either a path or a file name where depth output animated images should be saved (it overwrites the value set in the scene).



- `-output_alpha_anim\ \textquotesingle{}[PATH_OR_FILE_NAME]\textquotesingle{}:` specifies either a path or a file name where alpha output animated images should be saved (it overwrites the value set in the scene).
- `-tile\ X\ Y\ W\ H:` specifies the tile of the image to render. (X,Y) will be the top-left corner of the tile, and (W,H) will be its size (in pixels).

Warning:

The size of the tile will be $(W+1)*(H+1)$ since the first row/column is taken into account. For instance `-tile\ 0\ 0\ 64\ 64` will output 1 65×65 tile.

- `-rendersize\ W\ H:` where W is the width and H is the height of the rendered frame.
- `-cpu\ N:` indicates the maximum number of processors/cores which will be used by the *RenderNode* to render the job. This can be useful if you want to use other applications on the same machine where the *RenderNode* is running (by default, the *RenderNode* will run on all available processors, which can make other applications quite unusable).
- `-output_suffix\ something :` adds something to the output images filenames (before the extension)
- `-auto_tile_suffix :` adds `_tx_ty_tr_tb` to the output images filenames (before the extension), where tx is the left coordinate of the rendered tile, ty the top one, tr the right one and tb the bottom one.

Example: `-file\ L:\scene_test\Vue_Tests\PTx9c.vue\ -output\ E:\Nemo\Desktop\ -currentf`
`tile_suffix` will produce the “output_38_42_107_115.png” ($69+38=107$ and $42+73=115$) on the desktop.

Saving Global Illumination Prepass for Reuse

If you will be rendering over a network using *RenderNodes*, it is possible to set the *RenderNode* to only render the GI prepass and save it in order to reuse it for further render jobs.

Then you add an option to specify where the GI prepass can be found. If the file is found, the GI prepass will be skipped for this render.

To save the GI prepass file, there are now three new options to be added to the command line.

`-gi_output_file\ <path_to_gi_file>`

The rendernode will save the GI prepass to the specified path. You may use the file extension *.dat*, but it will have no importance.



`-gi_only`

The *RenderNode* will stop after the GI prepass is computed and saved using the option `gi_output_file` is used. This option shouldn't be used without the `gi_output_file` command. In the case of an animation, by default the *RenderNode* will save one file per frame (padded like specified in the *Render Animation Options*).

If you want to render only one single frame of GI (the first one), you may add this option:

`-gi_single_frame`

Only the first frame will be computed and the *RenderNode* will stop.

Reusing a GI Prepass File

The command for reusing the GI prepass file is:

`-gi_input_file\ <path_to_gi_file>`

The *RenderNode* will try to load the GI info from the given file (which should have been saved previously with option `-gi_input_file`). If it succeeds, the prepass phase will be skipped in the render process.

If rendering an animation, by default, the *RenderNode* will rebuild a different path for each frame (depending on padding options), exactly like it does when saving GI files.

If you want to use the same file for all frames (this single file should have been saved with the `-gi_single_frame` option as well), the option is:

`-gi_single_frame`

This feature can be useful if you have a big image which you want to distribute over a network, using the `-tile` option.

- First, render the full prepass on one node, with options `-gi_output_file` and `-gi_only`
- Then, distribute the tiles over the network with the `-gi_input_file` to all nodes.

xStream Network Rendering

xStream

xStream...



Saving Native Scenes for Network Rendering

...xStream

If you want to network render a native scene that includes VUE content, you first need to install VUE rendernodes on all the computers participating in the network render. You can then launch the render using your usual network rendering manager. In order for the VUE scene to be loaded correctly on all the computers, it is suggested that you incorporate the VUE scene into the native scene, using the VUE *Options* dialog using the **Incorporate Vue scene in native scene file** option.

If you don't incorporate the VUE scene, it means all plugins will try to load the scene from its absolute path, so you need to copy the VUE scene on the same exact location on all the computers, and you also need to check they use the same version of the VUE scene. This is a little tedious, so this is why it's generally more convenient to use the incorporate option.

In LightWave, you can use the **Relative to LW content folder** option and then save the VUE scene in a subfolder of your LightWave content folder. Then, if all plugins use the same network shared content folder, they will be able to load the VUE scene, without having to incorporate the VUE scene in the LightWave native scene.

VUE Licenses for Network Rendering

xStream

Instead of installing a full license of VUE on all your render node computers, you can install a *RenderNode* license. This license is restricted to render only. It cannot be used to edit a VUE scene, nor to save it.

In order to install a rendernode version of VUE on your computers, you need to launch the VUE installer, and select the **Vue RenderNode** option on the first screen. This will actually install the full VUE application, but will edit the *xStream_Render_Node.cfg* file in the application folder, so that when a plugin initializes, it will search for a *RenderNode* license on your network (*RenderNode* licenses are floating only, so you need to install them on your *License Server*).

You now can choose nodelocked or floating when you install the rendernode. If you choose nodelocked, you will have to enter your serial number during installation and activate it at this time. You will then be able to render with the nodes without having to install a license server.



Setting-up Your Workstation for a Network Render

xStream

You can configure your regular workstation so that it will take part in the network rendering process. To do this, we recommend that you edit the *xStream_Render_Node.cfg* file and change the license to a *RenderNode* one (use any basic *Text Editor* to do that – the file is commented to show where you need to make changes).

If your workstation is using a node-locked license, the plugin – when launched as a *RenderNode* – will continue to use this node-locked license, to avoid taking up a *RenderNode* license on the network. It is suggested to switch to a *RenderNode* so that the plugin will know it should not ask for any user interaction (all warnings, information messages and dialogs will be skipped). Depending on how the network rendering manager operates, the plugin may detect that it is participating in a network render and do the switch dynamically, but it's not always the case. This is why it's better you switch it manually.

Behavior without License

xStream...

Another option for *RenderNode* licenses that you can find in the *xStream_Render_Node.cfg* file is the default behavior when no license can be retrieved. There are three options:

- **Wait:** use this option if you want the plugin to wait indefinitely for a floating license. It's useful in a render farm because plugins not getting a license won't interfere with others.
- **Black:** use this option if you want the plugin to render a black frame if no license is found. Beware that in a render farm, setting this option can lead to undesired effects: a slave with this option will render black frames very quickly and will therefore be sent most of the jobs. You will therefore end up having all your frames black. Use this option only if your rendering manager can detect black frames.
- **Error:** use this option if you want the plugin to trigger an error (a crash of the host application) when no license is found. It is generally not possible for a plugin to abort a network render or to tell the renderer that it should stop rendering, so this option can be useful. Generally, the network rendering manager will be informed of the error, and will stop contacting this node (or do whatever you told it to do in this case).



If no option is selected in the file (which is the case when you install VUE xStream), the default option is **Wait**.

Note:

This option has no effect when the plugin is not configured as a *RenderNode*. In this case, the plugin will display an *Activation* dialog or a *License Server Connection* dialog if no license can be found. If the *Activation* dialog is skipped, the plugin will render watermarked frames.

...xStream

Rendering with Mental Ray Satellites

xStream...

In order to be able to use VUE xStream with Mental Ray satellite rendering (also known as “distributed bucket rendering”), you need to have the VUE scene on all the computers involved in the rendering. As long as you don’t edit the VUE part of the scene, you won’t need to update the scene on your network. But, if you move or rotate a VUE object proxy, or edit the VUE atmosphere, etc. you’ll need to save your scene and copy the file to the other computers. Due to limitations of services (Mental Ray satellite is a service), it’s not possible to put the scene file on a network shared location.

In order to use DBR with Mental Ray, you need to dissociate the VUE “scene file” from the main scene file, and copy it on each satellite (to a similar location on all of the nodes). To dissociate the VUE scene from the main host file, go to xStream options (menu [**Files** | **Options...**]) and untick the option **Incorporate VUE scene in 3DS Max scene file**. Close the *Options* dialog. You will have to choose a path for your VUE scene. Once this is done, save the 3DS Max scene too, and copy the *.vue* file you have just chosen to all nodes before starting the DBR render.

Additional note: you should NOT save the VUE scene to the desktop or “My Documents” directory, as the corresponding absolute path would depend on current user account and OS – for instance, if you save your scene on such a place under Vista/Windows 7, xStream would not be able to find the matching scene on Windows XP, because those places are interpreted differently; instead, choose a basic path such as **C:/xStream_scenes/my_scene_001.vue**.

We suggest you use the Mental Ray satellite option only when rendering large scenes or animations. Because the Mental Ray satellites have to initialize the VUE plugin, and then read the VUE scene, they will suffer a 30s-1min. delay before really starting to render. The local Mental Ray is already loaded into memory, and the VUE scene is already opened, which means it will start rendering a lot faster. It means that for a preview render, the local Mental Ray plugin will have already finished rendering all its



tiles before the satellites even begins to render one tile.

If you animate one of the VUE proxy objects in the native interface, the animation is not stored in the VUE scene, it's only stored in the native scene. Because the VUE object proxies are not rendered by Mental Ray, the native application is not sending the animation information on these objects to the Mental Ray satellites, which means your VUE objects won't be animated as they should. In order to get the proper animation, you'll need to export the animation of the VUE object proxies to the VUE scene. You can do this through the **Export Animation to VUE Scene** menu command. Following is the complete process of a render with xStream and mental ray satellite involving animation:

...xStream

- When your scene is ready to be rendered, select the **File | Export Animation To Vue Scene** menu item.
- In the *Export Animation* dialog, enter the frame start and end of your animation. Then click **OK**. The plugin will retrieve the animation of all the VUE object proxies and will set corresponding keyframes in the VUE scene.
- Then, save the VUE scene, by choosing either **File | Save** or **File | Save As....**. The VUE scene you've just saved now contains the animation you set in the native application.
- In your file explorer, browse to the folder where you saved the scene. Copy this scene file to each of the computers you want to use during the rendering. Please note that the path of the file must be the same on every computer.
- Enable Mental Ray bucket rendering option in the Mental Ray *Render Options*.
- Start the render.
- If you then want to adjust Mental Ray render settings, you don't have to go through all the above process. Simply make your changes and relaunch the render.



Python Scripting

Python is the industry standard, cross-platform, object-oriented application scripting language. It is both sufficiently easy to use and powerful to let you develop complex scripts and expand the capabilities of the VUE Pro versions.

Included in the VUE application folder is a selection of sample Python scripts that demonstrate typical effects that can be achieved with Python. These scripts are fully documented and we recommend that you take a look at the way they are written for a good example of Python scripting. You can find them in the *Python/Scripts* folder of the *Application*.

There are many online documents and tutorials about Python on the internet. If you are not familiar with Python scripting, we recommend that you read general documentation on Python before delving into the intricacies of VUE Python implementation. Please read below for details on the Python Console and interpreting Python commands.

Interesting places to look for Python documentation are:

- **Python homepage:** You will find documentation, tutorial and links to more Python-related websites.
- if you are familiar with scripting or programming, this [page](#) will get you up and running with Python within minutes.

Note:

Because VUE features its own Python interpreter, you do not need to install any additional software in order to create and run Python scripts in VUE.

Python Scripts

Python scripts are text files that contain a list of Python statements and function calls. These files usually carry the extension *.PY*. You can find sample Python scripts in the *Python/Scripts* folder of your *Application files*.

Basically, there are 3 types of Python functions in VUE:

- Functions that create or modify objects,
- Functions that define specific processing and that are called back by VUE at various points during execution; these are known as callback functions,
- Functions that setup the callback functions so that they are actually called back by VUE at the desired point of execution.



The first type of function is typically used in “helper” scripts that could be written to automate certain tasks, such as e.g. generating a patch of grass. The callback functions are used to really “customize” the behavior of VUE and create entirely new rendering geometries and effects.

Running Python Scripts

To run an existing Python script, simply select the menu command **Python | Run Python Script** in VUE. This will open a Standard File Browser letting you browse to the Python script that you would like to run. By default, the File Browser is opened on the *Python/Scripts* sub-folder of your VUE folder.

The first time you run a script, it will be compiled on the fly by the Python interpreter and a compiled Python script will be generated and executed. This is to make sure that the script can be executed with maximum performance. Compiled Python scripts carry the extension *.PYC*. Please note that although the script is compiled and Python performance is very good considering that it is an interpreted script, performance is in no way comparable to hard-coded effects.

The most recently used scripts are listed at the bottom of the Python menu. To re-run a recently used script, simply select it from this list.

Startup Scripts

Startup scripts are script files that are loaded and executed whenever a given scene is loaded. These startup scripts usually setup callback functions, but they could be used for any purpose. For instance, you could easily write a script that creates a new sphere each time a scene is loaded!

Setting up callbacks using a startup script is very useful as it avoids having to set them up manually each time the scene is loaded. For instance, if you look at the scripts in the *Python/Scripts/Filters* folder of your Application CD, you will notice that these scripts setup callback functions to post-process the pictures as they are rendered. But if you save the scene and open it later, you’ll have to run the script again, so that it restores the callbacks. This is where setting a startup script becomes useful:



Now take a look at the scripts in the `Python/Scripts/Quartic` folder of your Application CD. There are two scripts in this folder, one called `Quartic_Builder.py`, the other `Quartic_Startup.py`. If you examine the code in the Builder, you will see that it creates a Python object, and then runs the Startup script. The Startup script sets up several call backs, and then sets itself as the startup script by calling the `SetPythonStartupScriptPath()` function. That way, the next time you load this scene, the Python object won't be created again, but the call backs will be properly setup.

Running Python from Command Line

You can run Python scripts from the command line using the `-p` command. VUE will immediately load and run the indicated Python script after starting up. For instance, you could write a Python script that loads a scene and then renders it before closing. To run a Python script at startup, run VUE from the command line with the option `-p<Name\ of\ Python\ script>` (note: there should be no space between `-p` and the file name; you should add quotes (') around the file name if it contains spaces).

For example, to run `myscript.py` at startup, browse to the VUE application folder and enter (on Windows systems):

```
Vue Infinite xxxx.exe -p'C:\Program Files\e-on software\Vue Infinite xxxx\python\scripts\myscript.py'
```

On Mac OS X system, open a terminal window, go to the VUE Infinite application folder and enter:

```
Vue Infinite xxxx.app/Contents/MacOS/Vue Infinite xxxx -p'/Applications/Vue Infinite :python/scripts/myscript.py'
```



Creating Python Scripts

To create or modify a Python script, all you need is a *Text Editor*. Open the *Text Editor* and start entering a list of Python commands.

Don't forget to save your script with the **.PY** extension, or it won't be listed in the File Browser when you try to run it in VUE. You can check that the script performs as expected by running it on regular occasions.

Python Console

The Python Console can be accessed by selecting the menu command **Python | Display Console**. This console provides feedback on the success of running Python commands or scripts. It is a Python interpreter that can also be used by advanced users to enter commands directly, rather than creating stand-alone Python script files.

You can enter Python statements and function calls at the `>>>` prompt. For instance, entering `VUEInterface::AddSphere()` at the prompt and pressing enter will add a sphere to the current scene.

Python Documentation

You can access the VUE-Python documentation using the **Help | Python Documentation** menu command. This documentation provides a detailed list of all the Python structures and functions used in VUE.



Hot Tips

Resuming Render

If you want to interrupt a render, and be able to resume rendering it in a future session, you must make sure that the **Generate Resume Render** option is selected before starting the render (using the *Options* dialog, or the **Picture** menu). When you stop the render, VUE will generate the required information to resume rendering later. If you save the scene now, you can reload it later and resume rendering exactly where it stopped by selecting the menu command **Render | Resume Render**.

Using Layers for Faster Display

When you build a complex scene, it is good practice to place objects in different layers, each layer corresponding to one part of the picture (e.g. one layer for fuzzy cloud spheres, one layer for plants, etc.). This enables you to lock or hide parts of the scene without affecting render (provided the **Render everything** option is selected in the *Render options* dialog). Locking layers lets you work on other layers without being hindered by the locked ones and keeping them visible for reference. Hiding layers will seriously improve the display speed, and avoid visual clutter. You can also accelerate display by representing complex objects (such as plants) by boxes.




Resuming Animation Render

You can stop rendering an animation anytime, and resume it later without any loss of processing time: Just press the **Resume Rendering Animation** button in the *Animation Render Options* dialog to restart rendering the animation from where you stopped! This button becomes automatically enabled when it is possible to resume rendering.

Using Motion Blur in Still Pictures

You can use motion blur in still pictures to achieve dramatic effects, giving a sense of life and speed to a scene that would otherwise look... still. Objects with motion blur will leave some sort of trail behind them. The blurred cars in the *Bridge* sample scene are a good example of using motion blurring in a still picture.

To achieve this effect, you must first define a motion for the objects that you want to be blurred. To be noticeable, you will probably need to give these objects a very high speed.

Select the object you want to blur and move the **Current time** slider to 1 second. Now slightly move the object in the direction of the motion blur. You shouldn't move the object too much, if not its position will be modified when you do the final render. Now press the  key to reset the Current time to 0. Do a test render to get an idea of the size of the blur trail. You can adjust this trail by dragging the keyframe you just created to the right to reduce the length of the trail, or to the left to increase it.

Broadcast preset render quality may not be sufficient to get perfectly smooth motion blurring in a still picture (read [here](#)). You may want to switch to **Superior** or even **Ultra** (beware the render times...) or increase the number of rays per pixel in the anti-aliasing setting (read [here](#)).




Avoiding Object Detection in Wizard




When you plot the path of an object using the *Animation Wizard*, the altitude of the object is automatically processed to avoid collisions with other objects in the scene. This can sometimes be a problem (e.g. when you want to fly an object through a tunnel).

To avoid the detection of certain objects in your scene, just place them in a Hidden layer. Objects that are not visible are ignored by the Wizard.

Obviously, the other solution is to touch-up the path in the *3D Views*.

Rendering Time Dependent Materials

When you create a scene that contains no animation, apart from time dependent materials, the **Render Animation Preview** icon () only renders one single frame, because it is unable to determine what duration of animation you want to preview (e.g. when solely animating water).

The solution is to define a duration of animation: press the **Render Animation** icon () and select **Render sequence**. Enter the desired start and end of animation, then press **Close**. The **Start of animation** and **End of animation** handles ( and ) are now displayed in the *Timeline*. You can drag them to adapt the duration of the animation. Pressing **Render animation preview** will now preview the selected duration of animation.



Reducing Animation Render Times

Rendering animations is a very processor intensive activity. When you are rendering an animation, you may have to trade off some of the render quality for improved rendering speed. Keep in mind that the frames in the animation will only be visible for a fraction of a second, so the trade off may not even be noticeable. The better results some render features produce may not even be noticeable in the animation. You can probably accelerate the rendering by keeping the following in mind:

- Global illumination and global radiosity are definitely a no-go if you want animations that render quickly,
- Volumetric lights, especially volumetric sunlights are very long to render. Use them only if you have to. Disabling shadows in the volumetric lights will accelerate the rendering quite a bit.
- Avoid depth of field, soft shadows and blurred reflections/transparency as they will increase render times by an order of magnitude. They will also probably add a fair amount of flickering if you don't increase the render quality.

You can seriously reduce rendering time if you are ready to accept some amount of flickering in the animation. Do this by reducing the quality of the anti-aliasing. The flickering is particularly noticeable with highly detailed materials or tiny objects. Keeping the **Flicker reduction** option set helps reduce this effect.

Compressing Video

Since there usually are many frames in an animation, the size of animations on disk can rapidly become daunting. Also, because the files are very large, your computer may experience difficulties reading and displaying the animation smoothly.

Video compression is the solution to these problems. It is available directly in VUE through AVI, QuickTime and Mpeg 1 & 2 compression Codecs. To choose a compression Codec, open the *Advanced Animation Options* dialog, select compressed file format and press the **File format options** button to display a standard system dialog that lets you select the Codec to be used for compression. When you render the animation, it will automatically be compressed using the Codec you selected.



Troubleshooting

The first part of this second appendix gives a list of common problems you may have using the software. The second part may help you find what is wrong in a picture that, for some unknown reason, doesn't look quite right.

Scenes Take Ages to Render

You must keep in mind that some of the effects created by VUE are extremely long to render. If you find that your scene is taking too long to render, go through it and check out for any of the following effects that are not absolutely required (listed by order of importance):

- Global Illumination and Radiosity.
- Volumetric lights (especially directional lights such as the sun, and point lights), godrays. If you absolutely need the volumetric effect, make sure the **Quality boost** setting isn't exaggerated. Check if you can remove the **Cast shadows in volume** option,
- Depth of field (this requires a large amount of anti-aliasing to produce smooth results),
- Displacement mapped materials seen from up close,
- Subsurface materials with low average depth settings,
- Camera is placed in the middle of spectral clouds,
- Mixed materials that mix several transparent materials together,
- Glowing materials (slightly glowing materials that don't show up in the final render),
- Volumetric materials used on large objects (can you use fuzzy materials instead?), volumetric atmospheres,
- Blurred reflections and refractions, soft shadows (even if the blurring/softness amount is very low, it's still going to slow down the render dramatically). Do you need soft shadows on all lights?
- Shadows (do all lights/materials have to cast shadows, do all materials have to receive shadows?), etc.

A good understanding of all the advanced rendering features of VUE is required if you want to keep render times optimal. By getting rid of unnecessary effects you can easily



accelerate the rendering speed by an order of magnitude.

Last, you should keep in mind that the rendering technique used by VUE (i.e. ray-tracing) is an inherently slow technique. This technique was preferred over other techniques because of its vastly superior quality. Since VUE was first released we never stopped optimizing its render engine to make it ever faster, producing one of the fastest ray-tracing engines currently available.

Camera Moves by Itself

If you notice the camera changes position for no apparent reason, check that the height of the camera above ground is not set to be locked (see the **Lock Height** option in the camera's *Object Properties* panel, see [here](#)). This may cause the camera to move following changes in the scene.

Bright Fringes Appear at Wall Base in Radiosity

Sometimes, when rendering with radiosity, it may happen that you see bright areas of light appear in corners or near the floor where there shouldn't be any. This phenomenon is known as "light leak" and is an artifact caused by the photon rendering technology.

This problem occurs when the photon gathering radius is greater than the thickness of the walls. To correct this problem, either increase the **Lighting model Quality boost**, make your walls thicker, or reduce the **Maximum gathering radius** using the **Photon Maps** tab of the *Advanced Effects Options* dialog.



Render Time Estimation Is Pessimistic

You may notice that the render time estimation displayed at the beginning of rendering certain scenes is extremely pessimistic. The reason for this is that, in order to avoid useless computation, a lot of preparation tasks are fragmented in such a way that only those fragments that are indeed required are computed (e.g. displacement mapping, subsurface scattering, etc.). The usefulness of each fragment is determined at render time: each time the renderer needs an uncomputed fragment, it is computed dynamically. This typically happens during the early stages of rendering, when none of the fragments are yet computed. This is especially the case when rendering in preview mode, where the entire surface of the image is evaluated at the first render pass. The render time estimation should rapidly converge to a more realistic evaluation.

Long Preparation Time for Small Image

Preparation time is long, even for small render-areas.

Certain effects, such as radiosity, are “global” to the scene (an object that is not visible in the final image may nonetheless have an influence on the lighting of the scene). Such effects require the same amount of preparation, whatever the portion of the final image that is rendered. That is why this preparation time is not reduced by reducing the rendered area.

Program Crashes Randomly

If you are experiencing frequent random crashes when working with VUE, chances are something is wrong with your video board driver.

Obviously, you might want to enable **Compatibility Mode** (see [here](#)), but before doing this, visit the website of your video board manufacturer and check if there are any recent driver updates available. Depending on the manufacturer, OpenGL drivers in particular will tend to be updated and fixed for a while after the board is released. Also, because VUE uses several threads of OpenGL rendering, it can happen that some drivers that perform well with other applications will not perform correctly with VUE. Obviously, you



should point this out to the manufacturer so they can fix the problem.

In the mean time, here are a number of steps that you can take to improve compatibility (listed by order of importance):

- In the *Options* dialog, uncheck the **Enable background draw thread** option to stop multi-threaded OpenGL.
- Switch to the software implementation of OpenGL by selecting the **OpenGL (software)** option. You will have to restart VUE. The next time you restart the program, the *3D Views* will be drawn using a custom, in-house preview technology.

You should try these options until you find the configuration that works best for your system.

Noise Appears in Volumetric Effects

Grainy Volumetric Lights and Materials.

The reason for this is exactly the same as for getting noise in the sky: the **Preview** render setting is optimized for speed and only produces a rough fake of what the result will finally be. Although you can increase the **Quality boost** setting for the volumetric lights and materials, switching to **Final** render quality (or better) is generally enough.

Atmosphere Is Different in Preview and Final

When you have finished designing your scene (in **Preview** render quality) and switch to **Final** quality, you may notice that the colors of the atmosphere change slightly. Generally, this isn't noticeable, but it may be a problem in certain occasions. The reason is that, just like the noise problem discussed above, the **Preview** render quality isn't optimized the same as the **Final** (or better) render quality. As you switch to **Final**, VUE automatically refines the processing of the atmosphere, resulting in the slight changes you may have noticed.

If colors of the atmosphere are an important issue, you should switch to **Final** render quality before you start fine-tuning it.



Undesired Lens Flares Appear on Lights

If you have unwanted lens flares appearing on your lights (sun, spotlights), double-click on the light in the *World Browser* and uncheck the setting.

The default for this setting – on or off – can be set on this panel.

Missing Details

Some parts of the picture don't look as detailed as they should. It seems as if some fine details are missing.

This is due to the render quality setting you are using. For speedy renders, **Preview** quality optimizes the last render pass, sometimes skipping very fine details.

Preview is fine when you are working on the picture, just to make sure you're going the right way. But for final renderings of a picture, you have to use **Final** render quality.

Vector Graphics Don't Load in Text Editor

When loading vector graphics data, it may happen that nothing is loaded (the text preview remains blank). This typically happens when the Postscript processor was unable to process the vector graphics file. You should use a simpler version of Postscript that doesn't make reference to any external libraries (sometimes the case with Adobe Illustrator documents).



Invisible Objects

Although some objects appear correctly when rendered, they are not visible in the 3D views.

These objects are most certainly placed inside a **Hidden** layer, or have the **Hidden** attribute. To show all objects inside the scene, select **Show All Layers** from the popup menu of the *World Browser*. Alternately, you can use the layer status control to activate the hidden layers. You can reset the Hidden attribute of the object by using the popup menu of the **Preview Options** in the *Object Properties* panel.

Unable to Select Objects

Objects that are placed inside a **Locked** layer, or have the **Locked** attribute, cannot be selected using the *3D Views*. These objects are displayed in gray. Toggle the layer status back to **Active**, or use the *World Browser* to select them. You can reset the Locked attribute of the object by using the popup menu of the **Preview Options** in the *Object Properties* panel.

If you are trying to select an object that is placed behind another one, you can either use another view, or the *World Browser* to select it. Alternately, you can select all the objects under the cursor by **Control – clicking**, and then walking through the selection using **Tab** until the requested object is selected.



Objects Don't Render

Some objects that appear in the 3D views are not visible when rendered.

Make sure these objects aren't made from a totally transparent material. Maybe they are so fuzzy that you can't see them. Or they are black and **Additive** together, which yields invisible results.

Alternately, make sure the **Render everything** option is selected in the *Render Options* dialog.

Close-up Materials Look Like Tiles

Color steps appear when a material made from a mapped picture is seen from close, and when no interpolation has been specified. The pixels of the mapped picture are creating this effect.

Open the *Material Editor* by double clicking on the preview of the material that you want to modify (or select **Edit Material** from the preview's popup menu). Go to the **Colors** tab, and select a **Bilinear**, **Bicubic** or **Normalized** interpolation method.

Look Ahead Objects

Look Ahead Objects Don't Point in the Right Direction

When you animate an object with the **Look ahead** property, the orientation of the object is automatically processed to make it point in the direction of travel. However, depending on the initial orientation of the object, you may find out that your object is in fact pointing at right angles with the direction of travel. This is because the **Main axis** of the object is not correct. Use the **Main axis** drop down list in the *Animation Properties* panel to select the correct axis.

If none of the available options proves satisfactory, this means that the initial orientation of your object is not aligned with one of the axes of VUE. This happens when you import



objects.

You must align the object with one of the axes. First try zeroing all rotations in the **Rotation** sub-tab of the *Numeric Properties* panel. If the object is still not aligned, then rotate it manually until it is aligned. Now ungroup the object and regroup it. You will have to rename it to its old name. The object is now ready for **Look ahead** animation (if the object is not a group/Boolean object, it is necessarily an imported object. You will have to align it in the application that was used to create the object).

Objects Overreact to Motion

Or don't seem to react sufficiently

The algorithms used by VUE to compute the reaction of objects to motion are based on physical rules. While this ensures the most realistic motion, it has the drawback of being sensitive to the scale of your scene (because the speed at which objects travel depends on the scale of the scene).

You will notice this when objects don't seem to react naturally to their motion. You can fix the problem by modifying the sensitivity of objects using the *Motion Options* dialog (press **Options** in the *Animation Toolbox*). Read more on this dialog [here](#).

Materials Don't Move With Objects

Objects seem to move “through” the material

If you notice that the material of an animated object isn't following the object, then it is probably because you are using a **World** space material. To ensure that the material follows the objects, you should always select an **Object** based space mapping in the *Material Editor* (read [here](#)). If the material is mixed and is sensitive to the environment, you should also check the **Object orientation** option in the **Influence of environment** tab.



Objects Keep Getting Animated

If you modify a non animated object at a different time than the time it was created at, it automatically becomes animated. To avoid this problem, you should always modify objects at the time they were created at, or at zero time. All objects can be modified at zero time without becoming animated.

You can remove object animation by selecting **Not animated** from the **Motion type** drop-down list in the **Animation** tab, or by selecting the object in the *Timeline* and pressing **Delete**.

You can definitively prevent an object from becoming animated by selecting the **Forbid animation** option in the **Animation** tab of the *Object Properties* panel.

Animations Flicker

If you notice that rendered animations tend to flicker when played back, check the following:

- First of all, please make sure that the problem is not caused by the playback software or video compression options you have selected. Also, if no compression was used, you might want to check that your computer has sufficient bandwidth to playback the animation smoothly.
- In the *Advanced Animation Options* dialog, check the **Flicker reduction** options and make sure the frame rate is adequate.
- Avoid using high frequency procedural textures, that is textures that have very fine details. The fine details probably won't be visible in the animation anyway...
- Turn on texture anti-aliasing.
- Avoid using soft shadows and blurred effects.
- Increase the render quality by using the **User settings** should be your last resort. In the *Animation Render Options* dialog, make sure that you are rendering at a sufficiently high render quality setting (**Broadcast** render quality is designed for acceptable quality). If render quality isn't sufficient in Broadcast, you'll probably want to switch to **User settings** and select a very high anti-aliasing quality (25 rays per pixel). Make sure the **Motion blur** option is checked.



Keep in mind that there is no limit to the quality of the animations generated, except the time allocated for rendering. If you increase the quality, the rendering time will increase accordingly. So it's all a question of finding the ideal compromise.

Texture Filtering

Texture Filtering is an essential render setting to reduce noise and flickering that can arise because of high frequency textures (components of materials that exhibit very fine detail, usually finer than the size of a pixel). When used properly, it will lower the needs for strong object anti-aliasing (thus speeding up the render), and greatly reduce texture flickering.

Texture Filtering can be accessed via the anti-aliasing options dialog, just above texture anti-aliasing options. Its value is editable through a slider that ranges from 0 to 100%.

How It Works

This value corresponds to the size of the filter applied over textures during render. Ideally, this filter should always have the size of a pixel, so that all texture detail contained in each pixel is properly taken into account during texture evaluation. This corresponds to a value of 50% for Texture Filtering. If you specify a lower value, textures will be sharper but with more noise and/or flickering. If you specify a higher value, noise will be smoothed out but textures will appear blurred.

In Practice

You should tweak the value regarding your specific needs. In practice, the smallest value that yields good enough results should be used. From our own experience, a default value of 33% usually does the trick.



Texture Filtering Will Influence 2 Components at Render

Bitmaps: for each bitmap used in materials, if you edit its **Texture Map** node via the *Function Graph*, you will see a flag named **Allow mip-mapping**, which is checked by default. When this flag is checked, and if **Texture Filtering** has a non-zero value, corresponding bitmaps will be pre-filtered just before rendering. Thus, at render time, distant bitmaps won't exhibit any noise or flickering. This is particularly useful when rendering animated plants, especially for distant ones. You will enjoy much smoother results, and a great reduction in flickering. As specified above, a value for **Texture Filtering** of 33% will generally produce the best results.


Generic texture anti-aliasing: when **Texture anti-aliasing** is enabled, the **Texture Filtering** value will drive the size of the filter used by the texture anti-aliasing process, just like for bitmaps. This is very important because if texture anti-aliasing is enabled but **Texture Filtering** is set to 0%, you won't notice any improvement. Just like for bitmaps, a value of 33% is generally ideal for **Texture Filtering** used along with **Texture anti-aliasing**.

Animations Pulsate

This happens when rendering global illumination or radiosity; it is caused by the way rendering of these effects is optimized and is a common problem to all renderers. Unfortunately, the only way to reduce this effect is to increase the quality boost setting of the global illumination render.




Dark Triangles Appear on Terrains



This is probably an indication that the terrain with the problem doesn't have a high enough resolution. To correct it, double-click on the faulty terrain to open the *Terrain Editor*, and press the **x2** button once. This doubles terrain resolution. Now go to the **Erosion** tab, select a medium to low rock hardness, and press **Diffusive** erosion once to round off the angles in the terrain. This should solve the problem. If it is not the case, try adding more diffusive erosion. If the problem still remains, and especially if the terrain spreads out over a large area, you'll probably need to split it up into smaller terrains.

Another option would be to use a procedural terrain instead of a heightfield terrain.

Wrong Material Scale



If you have different items in your scene that somehow don't seem to fit together correctly, it may be that you have a problem with material scales. It is important that the objects of a scene look in proportion to one another (this doesn't mean the **Scale** control of the material has to be the same...).



This problem often occurs with water (see picture to the right): the scale of the waves on the surface of the water doesn't fit with the scale of the rest of your scene, so particular attention should go into getting this right.

Soft Shadows Look Noisy



When rendering soft shadows, depth of field or blurred transparencies and reflections, the lower preset render qualities (**OpenGL** and **Preview**) approximate the effect by adding noise. This enables you to get an idea of the final result, while not slowing down render.

To render all of these effects properly, super-sampling is needed. This is why they render correctly only under **Final**, **Broadcast** or **Ultra** preset render qualities.

VUE Objects Lose Relative Positions

When you import a set of 3D VUE objects (**.VOB**) it may happen that the relative positions of the different objects get lost. This is because VUE automatically centers objects that you load, so that they fit snugly into the viewports. To get rid of this option, open the *Options* dialog (by selecting the menu command **File | Options**) and uncheck the option called **Center Vue objects when loading**. A similar option is also available when importing objects created with other 3D applications.

Load all the models one after the other by using the menu command **File | Load Object**. Each object will now be positioned as it was when it was saved. If the objects appear too small or too large, select them all (using extended selection) and move/resize them together.



RenderCow Not Responding

If, after having setup and launched a *RenderCow* on a remote computer, you cannot manage to add it in the *HyperVue Network Rendering Manager*, you should check the following:

- Check that no Firewall is blocking access to the RenderCow either on the remote host, or on the computer running *HyperVue*. If you are not sure, contact your system administrator for help.
- Check that no other application is using the same port number on the remote host, or on the computer running *HyperVue*. If so, change port number.
- Check that communication on the selected port number is allowed on your network.

Unable to Export Object

If you are unable to export an object it may be because the object has been forbidden from export. This is often the case with polygon mesh objects, and particularly with all objects imported into earlier versions of VUE, and all Poser objects imported as PZ3. When you try to export such an object, a message will appear informing you that the object cannot be exported.

To forbid exporting a polygon mesh, double-click on it to open the [Polygon Mesh Options](#) dialog and click the **Forbid Export** button. Be advised, however, that you cannot remove the Forbid from export tag once it has been set.



Maya Mental Ray Renders Black

This probably means that VUE was not properly linked to the Mental Ray renderer. Try reloading the VUE scene to re-initialize VUE for Mental Ray.

To avoid this problem, please ensure that the Mental Ray for Maya plug-in (*Maya-atomr.mll*) is loaded before you use VUE, lest VUE won't be correctly linked to Mental Ray when creating the scene. We suggest that you setup Maya to automatically load the Mental Ray extension upon startup (from Maya's Plug-in manager). You can check that Mental Ray has finished loading by waiting for the Output Window to be displayed (it will display the Mental Ray version). Mental Ray state is also displayed in the *Maya Script Editor*.

Native Objects Don't Reflect VUE Objects

This is probably caused by a low ray-tracing depth in either the renderer settings or in the shader settings. You need to set the ray-tracing depth to a value greater than 1 to allow VUE-to-Native and Native-to-VUE reflections and refractions.

For Maya Software, you can increase the **Reflections** slider in the *Render Global Settings* dialog (**Raytracing Quality** tab).

For Mental Ray (Maya, 3DS Max and Softimage), you can increase it in the *Mental Ray Settings* dialog.



Volumetric Plugin Conflict

VUE is considered by V-Ray to be a volumetric plugin. Therefore, if using VUW with another V-Ray volumetric plugin, there may be conflicts. For example, if using a VUE sky in the same scene with a hair creating plugin, the hair may take on the color of the sky. It is probably best to limit the use of other plugins when using VUE, or, just be aware of the possibility of problems and strange effects.



Vue Artist Modules



License Agreement

THIS LICENSE AGREEMENT (“**EULA**”) IS A LEGAL AGREEMENT BETWEEN YOU (EITHER AN INDIVIDUAL OR A SINGLE ENTITY), THE END USER, AND E-ON SOFTWARE, SARL. IT IS A LICENSE AGREEMENT, AND NOT AN AGREEMENT FOR SALE OF THE PRODUCT (“**PRODUCT**”) VUE INFINITE / XSTREAM / RENDERCOW / RENDERNODE / LICENSE SERVER. BY DOING ANY OF THE FOLLOWING, YOU AGREE TO BE BOUND BY THE TERMS OF THIS EULA: INSTALLING, COPYING, DOWNLOADING, ACCESSING OR OTHERWISE USING THE **PRODUCT**, INCLUDING THE ACT OF REGISTERING THE **PRODUCT** (COLLECTIVELY REFERRED TO AS “**USE**”).

THE **PRODUCT** MAY INCLUDE ASSOCIATED MEDIA, PRINTED MATERIALS, AND “**ON-LINE**” OR ELECTRONIC DOCUMENTATION AND FILES. ALL COMPONENTS OF THE **PRODUCT** ARE LICENSED TOGETHER AS A SINGLE PRODUCT. ITS COMPONENT PARTS MAY NOT BE SEPARATED FOR USE ON MORE THAN ONE COMPUTER OR SOLD SEPARATELY.

THE **PRODUCT** AUTOMATICALLY INCLUDES ANY UPDATES, SUPPLEMENTS OR AMENDMENTS TO THE ORIGINAL PRODUCT PROVIDED TO YOU. IF THE UPDATES, SUPPLEMENTS OR AMENDMENTS CONTAIN A SEPARATE AGREEMENT WHICH EITHER AMENDS OR REPLACES THIS EULA, USE OF THE UPDATE, SUPPLEMENT OR AMENDMENT CONSTITUTES YOUR FULL AGREEMENT TO THE AMENDED OR REPLACEMENT **EULA**.

ANY SOFTWARE PROVIDED ALONG WITH THE **PRODUCT** THAT IS ASSOCIATED WITH A SEPARATE END-USER LICENSE AGREEMENT IS LICENSED TO YOU UNDER THE TERMS OF THAT LICENSE AGREEMENT.

PLEASE READ THE FOLLOWING TERMS AND CONDITIONS CAREFULLY BEFORE ANY **USE** OF THE **PRODUCT**. IF YOU DISAGREE WITH ANY OF THE FOLLOWING, IN PART OR IN WHOLE, YOU ARE NOT ALLOWED TO **USE** THE **PRODUCT**, AND SHOULD RETURN THE PACKAGE UNOPENED TO THE DEALER FROM WHOM YOU PURCHASED IT. YOU ARE SUBJECT TO ANY ADDITIONAL RETURN POLICIES AND REQUIREMENTS OF THE DEALER FROM WHICH YOU PURCHASED THE **PRODUCT**.



License

In consideration of the payment of a license fee, you are granted a personal, non-exclusive license to use the Product under the terms stated herein. If any portion of this Agreement is held to be unenforceable, such portion shall be reformed only to the extent necessary to make it enforceable.

You own the physical media on which the Product is provided, under the terms of this Agreement, but all title and ownership of the Product and all other rights not expressly granted to you by this Agreement, remain with e-on software, SARL.

Not For Resale Version (“NFR”). If your copy of the Product is a Not for Resale version, this EULA is non-transferable and may not be resold. If you purchased an NFR version from a reseller, the reseller has violated e-on software’s intellectual property rights under international copyright law. It is not a legal copy of the Product.

Trial Version. If your copy of the Product is a Trial version, this EULA is non-transferable and may not be resold. If you purchased a Trial version from a reseller, the reseller has violated e-on software’s intellectual property rights under international copyright law. It is not a legal copy of the Product.

Academic and Trial versions may not be used for any commercial purpose, including the use of the Academic or Trial version to create works of art, or portions of works of art, for a fee or resale. If you purchased an Academic version and desire to create works of art or portions of works of art for a fee or resale, contact e-on software for upgrade pricing to the full commercial version.

The Product can only be used on a single computer at any one time. You may install and use the Product on a secondary computer (for instance, a laptop), provided that (i) said secondary computer is not operated at the same location as primary computer, and (ii) that the Product is not used simultaneously on the primary and secondary computers. Should you wish to use the Product on more computers, or on several computers simultaneously, you must purchase another copy of the Product. You may not, under any circumstances, have the Product loaded onto the hard drives of more than two computers at the same time.

If you purchased a renewing license for the Product that requires payment of multiple installments, you must pay all installments, even if you decide to stop using the Product or terminate this agreement before the end of the renewal period. You must ensure that payment of each installment is received by e-on software prior to the due date. Failure to pay such installments on time shall result in additional charges of up to one hundred (100) percent of the cost of the late installment. You also agree to pay any and



all legal expenses incurred by e-on software in the process of recovering said payment and additional charges. You also fully understand and agree that your right to use the Product is subject to the complete payment of all installments. Should the payment of one or several of your installments be late by more than 30 days, your license to use the Product will be cancelled retroactively and, aside from owing any and all monies due, you understand and agree that all title, ownership and intellectual property rights on any content, assets or otherwise work created by you using the Product will be automatically transferred to e-on software and you will lose any such rights, title or intellectual property on said content, assets or work.

Copyright

The Product is owned by e-on software, SARL and is protected by U.S. copyright law as well as by international intellectual property conventions. Any reproduction, sale, transfer, transcription, storage in a retrieval system, translation into any language in any form or distribution by any means whatsoever of the Product, in part or in full, without the prior written permission from e-on software, SARL is strictly forbidden. Any such act shall constitute a copyright violation and shall be prosecuted to the fullest extent of the law.

You hereby grant permission to e-on software to communicate publicly, at e-on software's sole discretion, the fact that your company uses the Product, for instance but not by way of limitation by adding your company name and logo to the list of e-on software customers. This permission includes the right to use your company name and logo and is granted on a perpetual basis. Said permission can only be terminated by terminating this Agreement as set forth in the section "Term" of this Agreement.

By law, you, the end user, may do the following: (i) Make one (1) archival copy of the Product, in machine readable or printed form for backup purposes only in support of your use of the Product. You must reproduce and include the e-on software, SARL copyright notice on the backup copy of the Product.

(ii) Transfer the Product onto a hard drive, provided that you keep the original solely for backup purposes.

Any portion of the Product merged into or used in conjunction with another program will continue to be the property of e-on software, SARL and subject to the terms and conditions of this Agreement. You must reproduce and include the e-on software, SARL copyright notice on any portion merged in or used in conjunction with another program.



Restrictions

You may not rent, lend, lease, or assign the Product. If you purchased a time-limited license to use the Product, you may not transfer the Product. If you purchased a perpetual license for the Product, you may however transfer on a permanent basis the entire Product, including all associated media and documentation. If you acquired the Product as an upgrade from another product (the “upgraded product”), you may not (i) transfer the Product without simultaneously transferring the upgraded product, or (ii) transfer the upgraded product without simultaneously transferring the Product. If you acquired the Product as part of a bundle with one or several other products, you may not (i) transfer the Product without simultaneously transferring all other products in the bundle, or (ii) transfer one or several other products in the bundle without simultaneously transferring the Product.

When you make such transfer on a permanent basis, you may not keep any copies of the Product, and you must remove it from your hard drive(s). The person to whom you transfer the Product must agree to the terms of this License agreement. You must notify e-on software in writing of the transfer, including the name, permanent address, phone number and email address of the person to whom you transfer the software.

You may not alter, modify, translate or adapt the Product. In addition, you may not de-compile, disassemble or reverse engineer the Product. You also may not create any derivative works based on the Product. A derivative work is defined as a translation or other form in which the Product may be recast, transformed or adapted.

Content and Content Limitations

You may not distribute or sell VUE content exported into other standard or non-VUE computer graphics formats (such as but not by way of limitation .obj, .3DS) under any circumstances. You may use such content in other third-party computer graphics programs to create images and animations provided that the images and animations produced are not used to create content files (for instance but not by way of limitation 2D billboards, picture tubes, or alpha background planes) for distribution or re-sale.

Notwithstanding the foregoing, content files in VUE format (including, but not by way of limitation, .atm, .cld, .clr, .flt, .fnc, .mat, .mnd, .txs, .veg, .vob, .vue, and/or future proprietary formats) can be distributed and sold provided they meet the requirements



of legal derivative content works. In particular, but not by way of limitation, products that (i) re-use in whole or in part any copyrighted texture maps, (ii) are unsubstantially altered content files contained within standard VUE product distributions or (iii) are rendered standalone 2D images (2D billboards, picture tubes, or alpha background planes) of copyrighted content files are not considered legal derivative content.

Term

This license granted to you is effective until terminated. You may terminate it at any time by returning the Product to e-on software, SARL together with all copies, modifications, and merged portions in any form. The license will also terminate upon conditions set forth elsewhere in the Agreement, or if you fail to comply with any term or condition of this Agreement. You agree, upon such termination, to return the Product and the Documentation to e-on software, SARL together with all copies, modifications and merged portions in any form.

If you purchased a renewing license that requires payment of multiple installments, termination of your license does not void your obligation to pay said installments in full.

Your license to use the Product may be limited to a specific period of time. If this is the case, your right to use the Product shall begin with the first activation of the Product and shall end said period of time after the first activation. Use of the Product once said period of time has expired is not allowed and will constitute a copyright violation. If you wish to continue using the Product after said period of time, you must purchase a new license.

The license to use an Academic or NFR version of the Product is granted for a one year term and will expire on the first anniversary of the first activation of the license. The license to use a Trial version of the Product is granted for a limited term that will be identified with the license. You are not allowed to use the Product after this term. If you wish to continue using the Product after this term, you must purchase a new license.

Upon termination, e-on software, SARL may continue to enforce any rights provided by law. The provision of this Agreement which protects the proprietary rights of e-on software, SARL will continue in force after termination. Termination of this license, either voluntary or involuntary, does not entitle you to a refund of your purchase cost except as provided elsewhere in this License Agreement.



Limited Warranty

E-on software SARL, warrants, as the sole warranty, that the original medium on which the Product is distributed are free from defects in material and workmanship under normal use and conditions for a period of thirty (30) days from the date of the purchase, as evidenced by a copy of your receipt. No distributor, dealer, or any other entity or person has authorization to expand or alter either this warranty or any term of this Agreement. Any such representations will not bind e-on software, SARL.

You agree that in order to qualify for any technical or customer support, including any replacement of any component of the Product under the terms of this Agreement, you need to have registered the Product first with e-on software.

E-on software, SARL does not warrant that the functions contained in the Product will meet your requirements, nor that the operation of the Product will be uninterrupted and error-free. Except as stated above in this section, the Product and Documentation are provided as-is, without any other warranty or representation of any kind, either express or implied concerning the Product, Documentation, or use of the Product, including, but not limited to the implied warranties of quality, merchantability and fitness for a particular purpose. Using the Product is at your own risks, and in no event shall e-on software, SARL be liable for any damage whatsoever (including without limitation, special, incidental, consequential, or indirect damages for personal injury, loss of business profits, business interruption, loss of business information, or any other pecuniary loss) arising out of the use of or inability to use the Product or Documentation, even if e-on software, SARL has been advised of the possibility of such damages.

The warranties and remedies set forth above are exclusive and in lieu of all others, oral or written, express or implied. E-on software, SARL is not responsible for any costs including, but not limited to, those incurred as a result of lost profits or revenue, loss of time or use of the Product, loss of data, the costs of recovering such Product or data, the cost of substitute Product, claims by third parties, of other similar costs.

The warranties provided give you specific legal rights. You may have other rights which vary from state to state. Some states do not allow the exclusion of incidental or consequential damages, or the limitation of how long an implied warranty lasts, so some of the above may not apply to you. In no event will the entire liability of e-on software, SARL, under any provision of this Agreement, ever exceed the amount of the license fee paid by you to use the Product regardless of the form of the claim. This Limited Warranty is void if failure of the Product or Documentation has resulted from accident, abuse, or misapplication. Any replacement Product or Documentation will be warranted for the remainder of the original warranty period or thirty (30) days, whichever is longer.



Remedies

If a defect in the original medium on which the Product is distributed or in the Documentation occurs within thirty (30) days of purchase, you may return the faulty medium or Documentation to e-on software, SARL, along with your name, address, and a dated proof of purchase, and e-on software, SARL will replace it free of charge.

Before returning a product to e-on software, you must contact an e-on software representative to obtain a Return Authorization Number. This Return Authorization Number must figure prominently on the returned package. Products returned without a valid Return Authorization Number will not be considered for replacement or refund.

If you return the medium or Documentation, you must prepay shipping and either insure the Product and Documentation or assume all risk of loss or damage in transit.

US Government Restricted Rights

The Product and Documentation are provided with restricted rights. Use, duplication or disclosure by the US Government is subject to restrictions as set forth in subdivision (c) (1) (ii) of the Rights in Technical Data and Computer Software Clause at (252)227-7013 or subparagraphs (c) (1) and (2) of the Commercial Computer Software – Restricted Rights at 48CFR52.227-19, as applicable. The contractor / manufacturer of the Product and Documentation thereto is e-on software, SARL, 68 Avenue Parmentier, 75011 Paris (France).

E-on software can be contacted at 68 Avenue Parmentier, 75011 Paris, France.



Opening Objects From The Command Line

You can open VUE with a loaded object using the command line. Here are a list of the command options.

- **-software-gl:** Start VUE with the software OpenGL mode
- **-disable-msg-boxes:** Always skip message boxes dialogs and take the saved (or the default) options instead.
- **-autoconf-gl:** Force the openGL configurator to (re)run.
- **-r/-R:** Launch a render
- **-q/-Q:** Set render quality (-q100 set render quality to 100%)
- **-x/-X:** Set X resolution (-x1024 image will be 1024 pixel width)
- **-y/-Y:** Set Y resolution (-y2048 image will be 2048 pixel height)
- **-s/-S:** Set first animation frame rendered (-s12 start animation render on frame 12)
- **-e/-E:** Set last animation frame rendered (-e18 end animation render on frame 18)
- **-f/-F:** Fullscreen (image size will have the same size than the screen)
- **-m/-M:** Resume render
- **-a/-A:** Render animation
- **-o/-O:** Set still image name (-oRendu.png will save render output to Rendu.png)
- **-n/-N:** Set animation image name (-nRendu.png will save animation render output to Render_<frame number>.png)
- **-p/-P:** Launch python script
- **<file>:** Load the specified file. Can be .vue or any supported object format (.obj, .abc, .3ds, ...).



Third Party Licenses

- this software is based in part on the work of the Independent JPEG Group

Alembic File Format

TM & © 2009-2013 Lucasfilm Entertainment Company Ltd. or Lucasfilm Ltd. All rights reserved.

Industrial Light & Magic, ILM and the Bulb and Gear design logo are all registered trademarks or service marks of Lucasfilm Ltd.

© 2009-2013 Sony Pictures Imageworks Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of Industrial Light & Magic nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.



ALEMBIC ATTACHMENT A 6 REQUIRED NOTICES FOR DISTRIBUTION

The Alembic Software is distributed along with certain third party components licensed under various open source software licenses (“Open Source Components”). In addition to the warranty disclaimers contained in the open source licenses found below, Industrial Light & Magic, a division of Lucasfilm Entertainment Company Ltd. (“ILM”) makes the following disclaimers regarding the Open Source Components on behalf of itself, the copyright holders, contributors, and licensors of such Open Source Components:

TO THE FULLEST EXTENT PERMITTED UNDER APPLICABLE LAW, THE OPEN SOURCE COMPONENTS ARE PROVIDED BY THE COPYRIGHT HOLDERS, CONTRIBUTORS, LICENSORS, AND ILM “AS IS” AND ANY REPRESENTATIONS OR WARRANTIES OF ANY KIND, WHETHER ORAL OR WRITTEN, WHETHER EXPRESS, IMPLIED, OR ARISING BY STATUTE, CUSTOM, COURSE OF DEALING, OR TRADE USAGE, INCLUDING WITHOUT LIMITATION THE IMPLIED WARRANTIES OF TITLE, MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT, ARE DISCLAIMED. IN NO EVENT WILL THE COPYRIGHT OWNER, CONTRIBUTORS, LICENSORS, OR ILM AND/OR ITS AFFILIATES BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION), HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THE OPEN SOURCE COMPONENTS, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Boost C++ Libraries

Boost Software License ñ Version 1.0 August 17th, 2003 Permission is hereby granted, free of charge, to any person or organization obtaining a copy of the software and accompanying documentation covered by this license (the “Software”) to use, reproduce, display, distribute, execute, and transmit the Software, and to prepare derivative works of the Software, and to permit third-parties to whom the Software is furnished to do so, all subject to the following:

The copyright notices in the Software and this entire statement, including the above license grant, this restriction and the following disclaimer, must be included in all copies of the Software, in whole or in part, and all derivative works of the Software, unless such copies or derivative works are solely in the form of machine-executable object code generated by a source language processor.



THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR ANYONE DISTRIBUTING THE SOFTWARE BE LIABLE FOR ANY DAMAGES OR OTHER LIABILITY, WHETHER IN CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

MurmurHash3

The MIT License (MIT)

Copyright (c) <year> <copyright holders>

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.



Apache License 2.0

Apache License Version 2.0, January 2004 <http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

“License” shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

“Licensor” shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

“Legal Entity” shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, “control” means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

“You” (or “Your”) shall mean an individual or Legal Entity exercising permissions granted by this License.

“Source” form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

“Object” form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

“Work” shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

“Derivative Works” shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

“Contribution” shall mean any work of authorship, including the original version of the



Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, “submitted” means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as “Not a Contribution.”

“Contributor” shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. **Grant of Copyright License.** Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.

3. **Grant of Patent License.** Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. **Redistribution.** You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:

(a) You must give any other recipients of the Work or Derivative Works a copy of this License; and

(b) You must cause any modified files to carry prominent notices stating that You changed the files; and

(c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and



(d) If the Work includes a “NOTICE” text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.

6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.

7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an “AS IS” BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the



Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.

9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets “[]” replaced with your own identifying information. (Don’t include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same “printed page” as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the “License”); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an “AS IS” BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.



Boost

Boost Software License – Version 1.0 – August 17th, 2003

Permission is hereby granted, free of charge, to any person or organization obtaining a copy of the software and accompanying documentation covered by this license (the “Software”) to use, reproduce, display, distribute, execute, and transmit the Software, and to prepare derivative works of the Software, and to permit third-parties to whom the Software is furnished to do so, all subject to the following:

The copyright notices in the Software and this entire statement, including the above license grant, this restriction and the following disclaimer, must be included in all copies of the Software, in whole or in part, and all derivative works of the Software, unless such copies or derivative works are solely in the form of machine-executable object code generated by a source language processor.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR ANYONE DISTRIBUTING THE SOFTWARE BE LIABLE FOR ANY DAMAGES OR OTHER LIABILITY, WHETHER IN CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

CURL License

The curl license

Copyright (c) 1996 – 2015, Daniel Stenberg, daniel@haxx.se. All rights reserved. Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE AUTHORS



OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.

FBX SDK

Copyright (C) 2001 – 2010 Autodesk, Inc. and/or its licensors. All Rights Reserved.

The coded instructions, statements, computer programs, and/or related material (collectively the “Data”) in these files contain unpublished information proprietary to Autodesk, Inc. and/or its licensors, which is protected by Canada and United States of America federal copyright law and by international treaties.

The Data may not be disclosed or distributed to third parties, in whole or in part, without the prior written consent of Autodesk, Inc. (“Autodesk”).

THE DATA IS PROVIDED “AS IS” AND WITHOUT WARRANTY. ALL WARRANTIES ARE EXPRESSLY EXCLUDED AND DISCLAIMED. AUTODESK MAKES NO WARRANTY OF ANY KIND WITH RESPECT TO THE DATA, EXPRESS, IMPLIED OR ARISING BY CUSTOM OR TRADE USAGE, AND DISCLAIMS ANY IMPLIED WARRANTIES OF TITLE, NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR USE. WITHOUT LIMITING THE FOREGOING, AUTODESK DOES NOT WARRANT THAT THE OPERATION OF THE DATA WILL BE UNINTERRUPTED OR ERROR FREE.

IN NO EVENT SHALL AUTODESK, ITS AFFILIATES, PARENT COMPANIES, LICENSORS OR SUPPLIERS (“AUTODESK GROUP”) BE LIABLE FOR ANY LOSSES, DAMAGES OR EXPENSES OF ANY KIND (INCLUDING WITHOUT LIMITATION PUNITIVE OR MULTIPLE DAMAGES OR OTHER SPECIAL, DIRECT, INDIRECT, EXEMPLARY, INCIDENTAL, LOSS OF PROFITS, REVENUE OR DATA, COST OF COVER OR CONSEQUENTIAL LOSSES OR DAMAGES OF ANY KIND), HOWEVER CAUSED, AND REGARDLESS OF THE THEORY OF LIABILITY, WHETHER DERIVED FROM CONTRACT, TORT (INCLUDING, BUT NOT LIMITED TO, NEGLIGENCE), OR OTHERWISE, ARISING OUT OF OR RELATING TO THE DATA OR ITS USE OR ANY OTHER PERFORMANCE, WHETHER OR NOT AUTODESK HAS



BEEN ADVISED OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

Google-Breakpad

Copyright (c) 2006, Google Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of Google Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Copyright 2001-2004 Unicode, Inc.

Disclaimer

This source code is provided as is by Unicode, Inc. No claims are made as to fitness for any particular purpose. No warranties of any kind are expressed or implied. The recipient agrees to determine applicability of information provided. If this file has been purchased



on magnetic or optical media from Unicode, Inc., the sole remedy for any claim will be exchange of defective media within 90 days of receipt.

Limitations on Rights to Redistribute This Code

Unicode, Inc. hereby grants the right to freely use the information supplied in this file in the creation of products supporting the Unicode Standard, and to make copies of this file in any form for internal or external distribution as long as this notice remains attached.

MIT License

The MIT License (MIT) Copyright (c) <year> <copyright holders> Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions: The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.



Ptex Licensing

Ptex License

PTEX SOFTWARE

Copyright 2009 Disney Enterprises, Inc. All rights reserved

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

The names “Disney”, “Walt Disney Pictures”, “Walt Disney Animation Studios” or the names of its contributors may NOT be used to endorse or promote products derived from this software without specific prior written permission from Walt Disney Pictures.

Disclaimer: THIS SOFTWARE IS PROVIDED BY WALT DISNEY PICTURES AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, NONINFRINGEMENT AND TITLE ARE DISCLAIMED. IN NO EVENT SHALL WALT DISNEY PICTURES, THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND BASED ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.



zlib License

Copyright (c) <year> <copyright holders>

This software is provided “as-is”, without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgement in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

