

Autodesk®
Flame® Premium 2013
A Discreet® creative finishing product



Installation and Configuration Guide for Linux® Workstations

Legal Notices

Autodesk® Flame® Premium 2013

© 2012 Autodesk, Inc. All rights reserved. Except as otherwise permitted by Autodesk, Inc., this publication, or parts thereof, may not be reproduced in any form, by any method, for any purpose.

Certain materials included in this publication are reprinted with the permission of the copyright holder.

Portions related to MD5 Copyright © 1991-2, RSA Data Security, Inc. Created 1991. All rights reserved. License to copy and use this software is granted provided that it is identified as the "RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing this software or this function. License is also granted to make and use derivative works provided that such works are identified as "derived from the RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing the derived work. RSA Data Security, Inc. makes no representations concerning either the merchantability of this software or the suitability of this software for any particular purpose. It is provided "as is" without express or implied warranty of any kind. These notices must be retained in any copies of any part of this documentation and/or software.

The following are registered trademarks or trademarks of Autodesk, Inc., and/or its subsidiaries and/or affiliates in the USA and other countries: 123D, 3ds Max, Algor, Alias, Alias (swirl design/logo), AliasStudio, ATC, AUGI, AutoCAD, AutoCAD Learning Assistance, AutoCAD LT, AutoCAD Simulator, AutoCAD SQL Extension, AutoCAD SQL Interface, Autodesk, Autodesk Homestyler, Autodesk Intent, Autodesk Inventor, Autodesk MapGuide, Autodesk Streamline, AutoLISP, AutoSketch, AutoSnap, AutoTrack, Backburner, Backdraft, Beast, Beast (design/logo) Built with ObjectARX (design/logo), Burn, Buzzsaw, CAiCE, CFdesign, Civil 3D, Cleaner, Cleaner Central, ClearScale, Colour Warper, Combustion, Communication Specification, Constructware, Content Explorer, Creative Bridge, Dancing Baby (image), DesignCenter, Design Doctor, Designer's Toolkit, DesignKids, DesignProf, DesignServer, DesignStudio, Design Web Format, Discreet, DWF, DWG, DWG (design/logo), DWG Extreme, DWG TrueConvert, DWG TrueView, DWFx, DXF, Ecotect, Evolver, Exposure, Extending the Design Team, Face Robot, FBX, Fempro, Fire, Flame, Flare, Flint, FMDesktop, Freewheel, GDX Driver, Green Building Studio, Heads-up Design, Heidi, Homestyler, HumanIK, IDEA Server, i-drop, Illuminate Labs AB (design/logo), ImageModeler, iMOUT, Incinerator, Inferno, Instructables, Instructables (stylized robot design/logo), Inventor, Inventor LT, Kynapse, Kynogon, LandXplorer, LiquidLight, LiquidLight (design/logo), Lustre, MatchMover, Maya, Mechanical Desktop, MIMI, Moldflow, Moldflow Plastics Advisers, Moldflow Plastics Insight, Moldflow Plastics Xpert, Moondust, MotionBuilder, Movimento, MPA, MPA (design/logo), MPI, MPI (design/logo), MPX, MPX (design/logo), Mudbox, Multi-Master Editing, Navisworks, ObjectARX, ObjectDBX, Opticore, Pipeplus, Pixlr, Pixlr-o-matic, PolarSnap, PortfolioWall, Powered with Autodesk Technology, Productstream, ProMaterials, RasterDWG, RealDWG, Real-time Roto, Recognize, Render Queue, Retimer, Reveal, Revit, RiverCAD, Robot, Scaleform, Scaleform GfX, Showcase, Show Me, ShowMotion, SketchBook, Smoke, Softimage, Softimage|XSI (design/logo), Sparks, SteeringWheels, Stitcher, Stone, StormNET, Tinkerbox, ToolClip, Topobase, Toxik, TrustedDWG, T-Splines, U-Vis, ViewCube, Visual, Visual LISP, Voice Reality, Volo, Vtour, WaterNetworks, Wire, Wiretap, WiretapCentral, XSI.

Adobe, Flash and Reader are either trademarks or registered trademarks of Adobe Systems Incorporated in the United States and/or countries.

Automatic Duck and the duck logo are trademarks of Automatic Duck, Inc.

Lucida is a registered trademark or trademark of Bigelow & Holmes in the U.S. and other countries.

Python is a registered trademark of Python Software Foundation.

All other brand names, product names or trademarks belong to their respective holders.

Disclaimer

THIS PUBLICATION AND THE INFORMATION CONTAINED HEREIN IS MADE AVAILABLE BY AUTODESK, INC. "AS IS." AUTODESK, INC. DISCLAIMS ALL WARRANTIES, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE REGARDING THESE MATERIALS.

Contents

Chapter 1	Installation workflows	1
Chapter 2	System requirements	3
Chapter 3	Hardware configuration	5
	Typical Configuration Overview for Visual Effects and Finishing Applications	6
	Typical Configuration Overview for Lustre	6
	Peripheral connections overview	9
	Connect video components	10
	Connect audio components	10
	Connect media storage	12
	Connect Archive Storage (not Lustre)	20
	Configure BIOS	21
Chapter 4	Install Linux	23
	Configure basic network settings	24
	Configure an InfiniBand card	25
	Install the DKU and the AJA OEM-2K firmware	25
	Configure storage	26
Chapter 5	Installing Your Software	35
	Prepare the installation media	35
	Install Flame Premium	35
	Install Lustre stand-alone	37
	Uninstall	37
Chapter 6	Software configuration	39
	Configure Visual Effects and Finishing Software	39
	Configure Backburner services	42

	Configure Lustre slave rendering	43
	Lustre background rendering	44
	Configure Lustre BrowseD	51
	Advanced Configuration and Troubleshooting	52
Chapter 7	Licensing and Starting	79
	Node-locked licensing	79
	Network licensing	80
	Start the software	85
	Index	89

Installation workflows

1

This guide covers the following:

- Flame Premium
- Flame
- Smoke for Linux
- Lustre
- Inferno
- Flint
- Backdraft

User prerequisites for installation

- Knowledge Linux, preferably Red Hat.
- Familiarity with computer hardware, storage and networking in a professional video/film production environment.
- Root access to your system. The default root account password on an Autodesk workstation is password.

Software upgrade only

To upgrade the software without updating the hardware or OS:

- 1 [Install the DKU and the AJA OEM-2K firmware](#) (page 25)
- 2 [Upgrade projects to the current version](#) (page 41) Unnecessary if you are upgrading to a service pack of the same application version or to a service pack of the same extension.
- 3 License your software. If you are not on subscription use [Node-locked licensing](#) (page 79). On subscription you can use node-locked or [Network licensing](#) (page 80). Unnecessary if upgrading to a service pack of the same software version or to a service pack of the same extension.

Install from scratch or upgrade Linux

To reinstall or upgrade the operating system and/or change your hardware or storage setup:

- 1 If you are installing your hardware make the following connections:
 - Peripherals such as mouse, keyboard, pen tablet, graphics monitor, house network.
 - Autodesk Wire network.
 - VTR and a broadcast monitor.
 - Audio hardware for your workstation.

- Storage arrays to the workstation, but do not power them on before having installed Linux, to prevent the Linux installer from attempting to format the arrays and use them as system drives.
- 2 [Configure BIOS](#) (page 21).
- 3 [Install Linux](#) (page 23).
- 4 [Install the DKU and the AJA OEM-2K firmware](#) (page 25)
- 5 If you are using a new Stone Direct storage array, [Configure storage](#) (page 26)
- 6 [Install Flame Premium](#) (page 35)
- 7 [Configure Visual Effects and Finishing Software](#) (page 39).
- 8 License your software. If you are not on subscription use [Node-locked licensing](#) (page 79). On subscription you can use node-locked or [Network licensing](#) (page 80). Unnecessary if upgrading to a service pack of the same software version or to a service pack of the same extension.

Concepts and Terminology

If this is the first time you are configuring an Autodesk Visual Effects and Finishing filesystem, familiarize yourself with the following key concepts.

Stone and Wire The software package that encompasses local management of media and the transferring of media between Visual Effects and Finishing workstations. This package is installed automatically with the application.

Partition In Stone and Wire, a partition (also referred to as “volume”) is defined as a volume of media storage. When creating a project, the Visual Effects and Finishing applications permanently associate it to one of the available partitions. This association means that the project inherits the rules and media preferences of the partition. You can define up to eight partitions.

Managed Media Media is said to be managed when the media assets, typically DPX files, are managed or “owned” by the Visual Effects and Finishing application. The application deletes managed media it no longer needs. All managed media is stored in a Managed Media Cache.

Unmanaged Media Unmanaged media refers to media assets that are used by, but not exclusively owned by, the Visual Effects and Finishing application. Soft-imported clips are an example of unmanaged media. This is relevant in shared access workflows, where numerous applications are using the same media. Unmanaged media is used in projects and clip libraries in the same way as managed media. However, the application does not delete it when it is no longer needed.

Managed Media Cache A directory residing on a standard FS volume, to which all managed media is written. This includes intermediates, imported media, captured video, audio, and proxies. Although the Managed Media Cache can be accessed by standard operating system tools such as defragmentation and backup tools, it is not meant to be accessed by any other application or user than the Visual Effects and Finishing and Wiretap applications.

Media Database The standard filesystem media database plays a central role in the management of media residing on standard filesystems. It is responsible for mapping frame IDs to the location of the media on disk. Database entries are created for all managed media and soft-imported media. There is one standard filesystem media database file per volume.

Clip library A clip library is a catalog maintained locally by its Visual Effects and Finishing application. It is used, internally by the application, to store clip metadata, including frame IDs. The artist acts upon portions of it indirectly, via the application User Interface.

System requirements

2

Hardware requirements

Supported Workstations

Workstation	Supported Software	Supported BIOS Versions	Operating System
HP Z800	All Creative Finishing software	3.07	Red Hat Enterprise Linux Desktop 5.3 with Workstation Option (64-bit)
HP xw8600	All Creative Finishing software	1.35	Red Hat Enterprise Linux WS 4, Update 3 (64-bit)
HP xw9400	Not supported for Lustre	3.05	Red Hat Enterprise Linux WS 4, Update 3 (64-bit)
HP xw8400	All Creative Finishing software	2.26	Red Hat Enterprise Linux WS 4, Update 3 (64-bit)

Memory Requirements

Minimum RAM	Recommended RAM
12.0 GB for HP Z800 8.0 GB for other platforms	12.0 GB for HP Z800 16.0 GB for other platforms (highly recommended)

For Visual Effects and Finishing products, you can adjust the MemoryApplication keyword.

OS & software requirements

Nodes purchased from Autodesk ship with the correct Autodesk distribution of Red Hat Enterprise installed and configured. You should only need to install Linux in the following situations:

- You are setting up a render node that was not purchased from Autodesk, in which case you need CentOS 5.3. If you have older nodes running Fedora Core or CentOS 4.6, you do not need to upgrade them to CentOS 5.3.
- You plan to upgrade your software on an older supported render node from Autodesk running a version of Red Hat Enterprise Linux no longer supported by the current version of the software. In this case, you must perform a fresh install of the required Red Hat Enterprise Linux version.

- You replaced the system disk. In this case you must get a new license. Contact Customer Support for assistance.
- To ensure optimal performance, it is recommended that only Autodesk Creative Finishing applications and required Red Hat Enterprise Linux packages be installed on the workstations.

Use the following table to determine the correct version of Red Hat Enterprise Linux required for your hardware platform.

Hardware Platform	Linux Version
HP ProLiant DL160se G6	Custom Autodesk distribution of Red Hat Enterprise Linux Desktop 5.3 with Workstation Option
HP ProLiant DL160 G5 or HP ProLiant DL140 G3	Custom Autodesk distribution of Red Hat Enterprise Linux WS 4, Update 3

To determine the version of Linux system is running, open a terminal and type: `cat /etc/redhat-release`

Browser support for documentation

- The offline HTML-based documentation requires a browser that supports frames and JavaScript.
- Firefox, Internet Explorer and Safari are supported.
- Google Chrome is not supported.

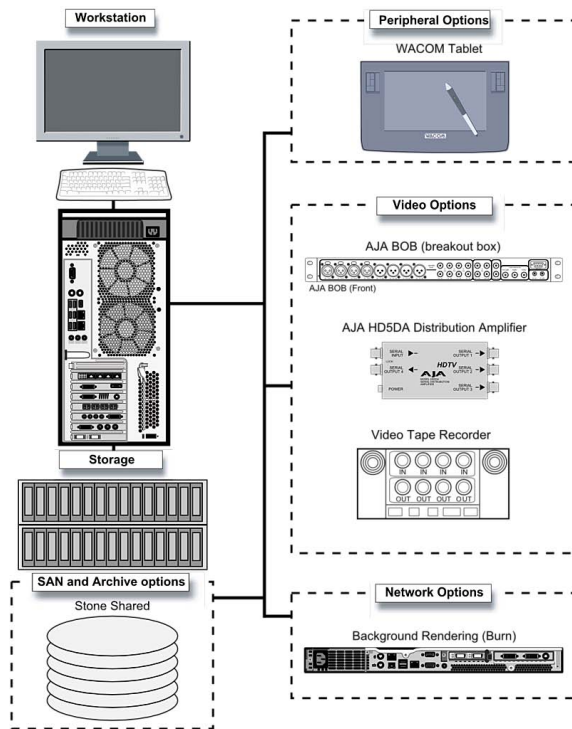
Hardware configuration

3

If you are only upgrading an existing application, you do not need to reconfigure your hardware.

This section contains wiring diagrams for HP Z800 workstation with the NVIDIA 6000 graphics card and the AJA KONA 3G. It does not contain information and diagrams for the older workstations supported in this version. If you have an HP Z800 with an NVIDIA 5800 graphics card, see the HP Z800 Workstation Hardware Setup Guide PDF.

Typical Configuration Overview for Visual Effects and Finishing Applications



Typical configuration

Typical Configuration Overview for Lustre

You may or may not have all of the components listed here in your Lustre setup. The type of workgroup, as well as the feature set you purchase for each of the workstations in that workgroup, determine the components in your installation.

There are two kinds of Lustre workstations:

- The Master Station is designed for GPU-accelerated sessions where the colorist works together with the cinematographer. It includes a toolset for visual design and grading, using up to 4K resolution and 16-bit files and for completing tasks like dust busting, conforming, rotoscoping, and capture/payout. It also includes SD and HD I/O, dual link and HSDL video formats, and the film workflow features, which consist of infrared channel dust removal and support for all standard input and output resolutions and bit-depths. The Master Station can support up to three panels of the Autodesk Control Surface and, with an additional license, a Slave Renderer.
- The optional Lustre Station offers a way to improve the efficiency and cost-effectiveness of a pipeline by offloading tasks that do not require the full feature set of the Master Station. Multiple Lustre Stations can work in parallel to increase throughput and can be used for tasks such as dust-busting, preparatory work, fine-tuning creative sessions, conforming data from EDLs, updating editorial changes, and mastering to different formats using the real-time deliverables function. The Lustre Station includes the film workflow features, and the ability to create geometries and masks. Primary colour grading on the Lustre Station requires a separate license. The following features also require an additional license: SD and HD I/O, dual link and HSDL video formats. The Lustre Station does not support the Slave Renderer.

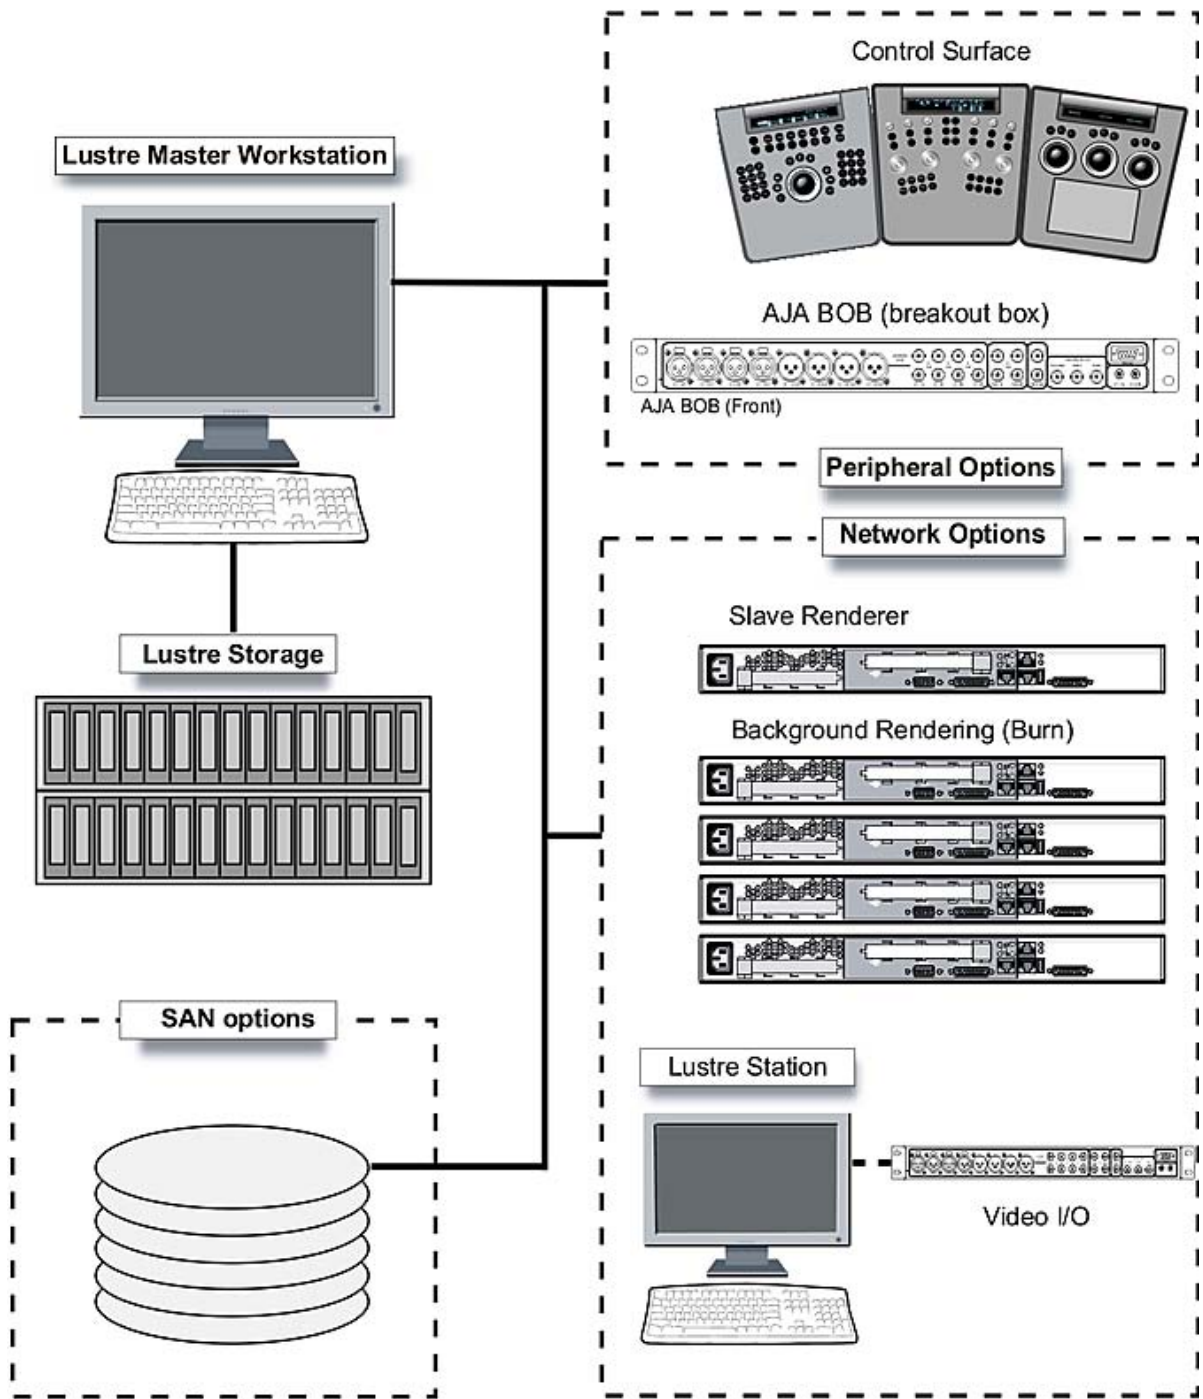
Optional Components

Control Surface Provides improved interactivity when colour grading film and video footage. You can perform many of the same tasks you do in the Lustre user interface using the Control Surface. You connect the Autodesk Control Surface to a Master Station and configure the Control Surface on the workstation to which it is connected.

Video I/O Card and Breakout Box On the HP Z800 workstation, video I/O is provided by the AJA card. The video I/O card consists of an HD/SD card and a breakout box. This configuration provides real-time SDI input and output of uncompressed 8-bit or 10-bit HD or SD video in both YUV (4:2:2) and RGB formats (4:4:4 or 4:2:2). For a list of supported video formats, see the [Autodesk Lustre User Guide](#). The drivers for the video card and breakout box are installed by the Discreet Kernel Utility (DKU).

Slave Renderer The Slave Renderer frees system resources by automatically off-loading render tasks. Connect the Slave Renderer to a switch on your house network. Note that only one Lustre Master station can use the Slave Renderer. The Slave Renderer software package is installed on the Slave Renderer.


Background Renderer Background rendering frees up Lustre workstations for colour grading. You can use up to eight background rendering machines to process your final frames. Background rendering software is installed on each background render node



A typical Lustrer workgroup configuration built around a Master Station.

Peripheral connections overview

HP Z800 Workstation

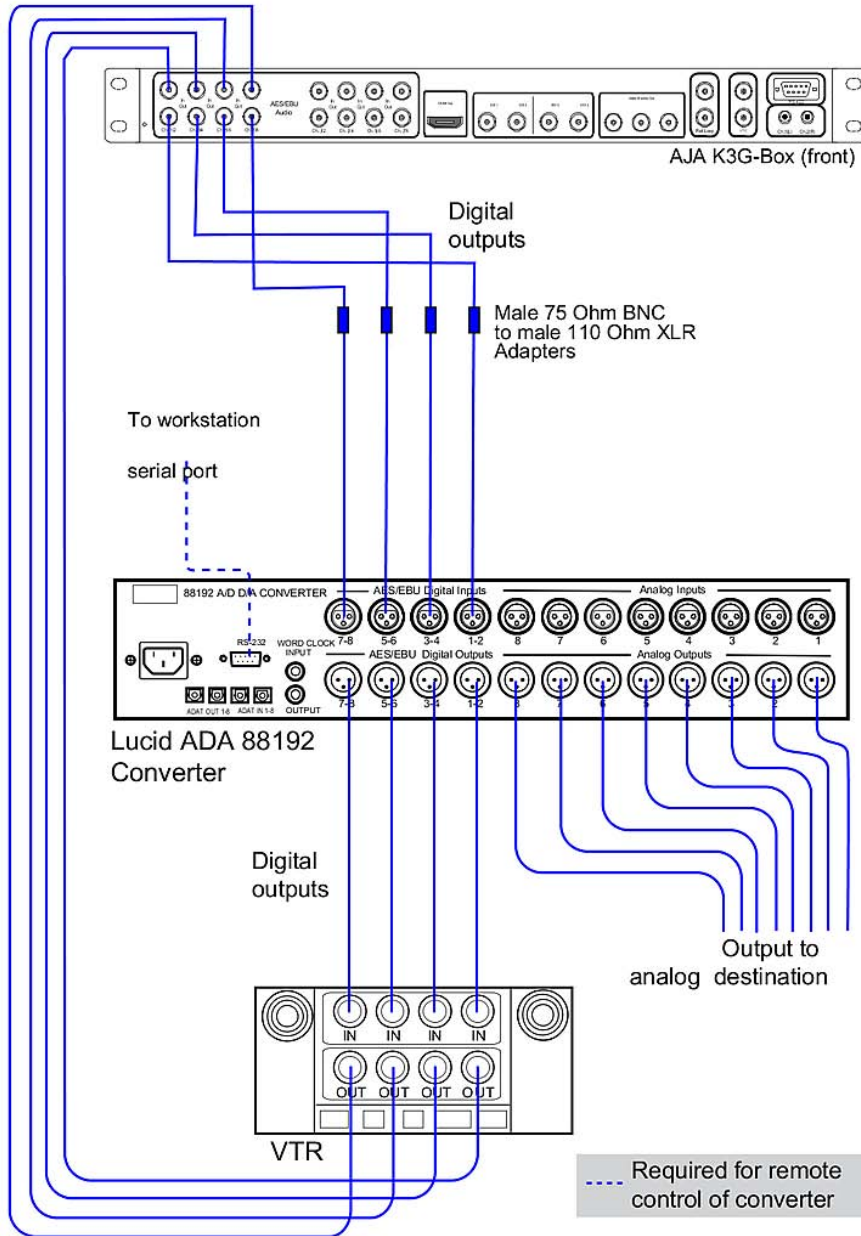


Adapter card positions and peripherals connections

This diagram depicts the HP Z800 workstation with a 2-port GigE adapter in slot 1 (top to bottom), and a Mellanox QDR InfiniBand / 10-GigE adapter in slot 7. Optionally, your workstation can be set up with a second ATTO Fibre Channel adapter in slot 1.

Lucid ADA 88192 Audio Converter Converts signals between the workstation and all digital or analog audio I/O devices.

AJA K3G-Box Balanced Audio breakout box and AJA Kona 3G adapter card The Balanced Audio breakout box is the audio component of the AJA breakout box. It provides connections for audio I/O. This breakout box connects to the AJA Kona 3G adapter on your workstation. The Kona 3G adapter provides real-time input and output of uncompressed SD and HD video signals as well as audio data at 24-bit resolution. The Kona 3G adapter handles balanced AES/EBU audio signals from the Balanced Audio breakout box.



Connect the Discreet Native Audio hardware components to the AJA breakout box

Connect media storage

WARNING Do not power on your storage before installing Linux. If you do, the Linux installer might attempt to format the storage array and use it as the system disk.


The workstation can be connected to two types of media storage:

- Stone Direct disk arrays: Autodesk's high-performance direct-attached storage (DAS) for individual workstations.
- A SAN: infrastructure that allows multiple workstations to share simultaneous access to a central storage enclosure.

Connecting the Fibre Channel Adapter

On HP Z800 workstations, connections to media storage are made using the ATTO Celerity FC84-EN fibre channel adapter, located in slot five (top to bottom).

The fibre channel adapter is equipped with four ports, called *loops*. You can connect your storage enclosure to the fibre channel cards using either 2 loops or 4 loops. A 2-loop device can be connected to either the inner pair of ports or the outer pair of ports. A 4-loop device requires both the inner pair and the outer pair.



NOTE If you are not using all the ports of a fibre channel adapter, terminate the two unused ports with the FC loopback couplers that shipped with your system. Terminating these ports significantly decreases the time required to boot the system.

Connecting Stone Direct Storage


The diagrams in this section illustrate how to connect the fibre channel loops to the Stone Direct storage enclosures supported with your workstation. For technical details on each storage model, see the Storage chapter in the hardware setup guide.

Connecting Dot Hill AssuredSAN 3003 Storage

The Dot Hill storage solution consists of one main disk enclosure, plus an optional expansion enclosure. The main disk enclosure contains two controllers: Controller A at the top, Controller B at the bottom.

Connecting an expansion: if you are using an expansion enclosure, connect the SAS port of controller A to the SAS port of the top I/O module on the expansion, and the SAS port of controller B to the SAS port of the bottom I/O module of the expansion.

Connecting the FC loops to the main enclosure controllers: each controller has 4 FC ports, numbered 0 to 3. The Autodesk recommended configuration uses only ports 0 and 2 of each controller.



The storage can be connected to the workstation through either 2 FC loops or 4 loops.

2-loop configuration: On ATTOcards 2 micro chips handle fiber traffic to the four ports: chip A handles ports 1 and 2 and chip B handles ports 3 and 4. To get maximum storage performance in a 2 loop configuration, the fiber loops should be distributed across the two chips.

Autodesk recommends the following connections:

- 1 ATTO Port 2 to Dot Hill port A0
- 2 ATTO Port 3 to Dot Hill port B0

You can also use the two outer ports on the FC adapter card, if they're not used to connect to a SAN or archiving device.

- 1 ATTOPort 1 to Dot Hill port A0
- 2 ATTOPort 4 to Dot Hill port B0

4-loop configuration: A 4-loop configuration uses all four ports available on the ATTO fibre channel adapter card. If you want to also connect your workstation to a SAN or an archiving device, you will need to install a second FC adapter card into your workstation.

Connect the FC loops as follows:

- ATTOPort 1 to Dot Hill port A2
- ATTOPort 2 to Dot Hill port A0
- ATTOPort 3 to Dot Hill port B0
- ATTOPort 4 to Dot Hill port B2

Connecting XR 6500 Storage

The following diagrams illustrate 2-loop and 4-loops connections for XR 6500 series storage assemblies. Cable your storage exactly as illustrated to ensure proper functionality.

A XR 6500 RAID enclosure supports a maximum of seven XE expansion enclosures. Configurations with two XR RAID enclosures are not supported.

Single XR 6500 RAID enclosure, 2 loops

To FC adapter on workstation / LMS (port 1)

To FC adapter (port 2)


XR RAID Enclosure




XE (EBOD) Expansion Enclosure




XE (EBOD) Expansion Enclosure



XE (EBOD) Expansion Enclosure

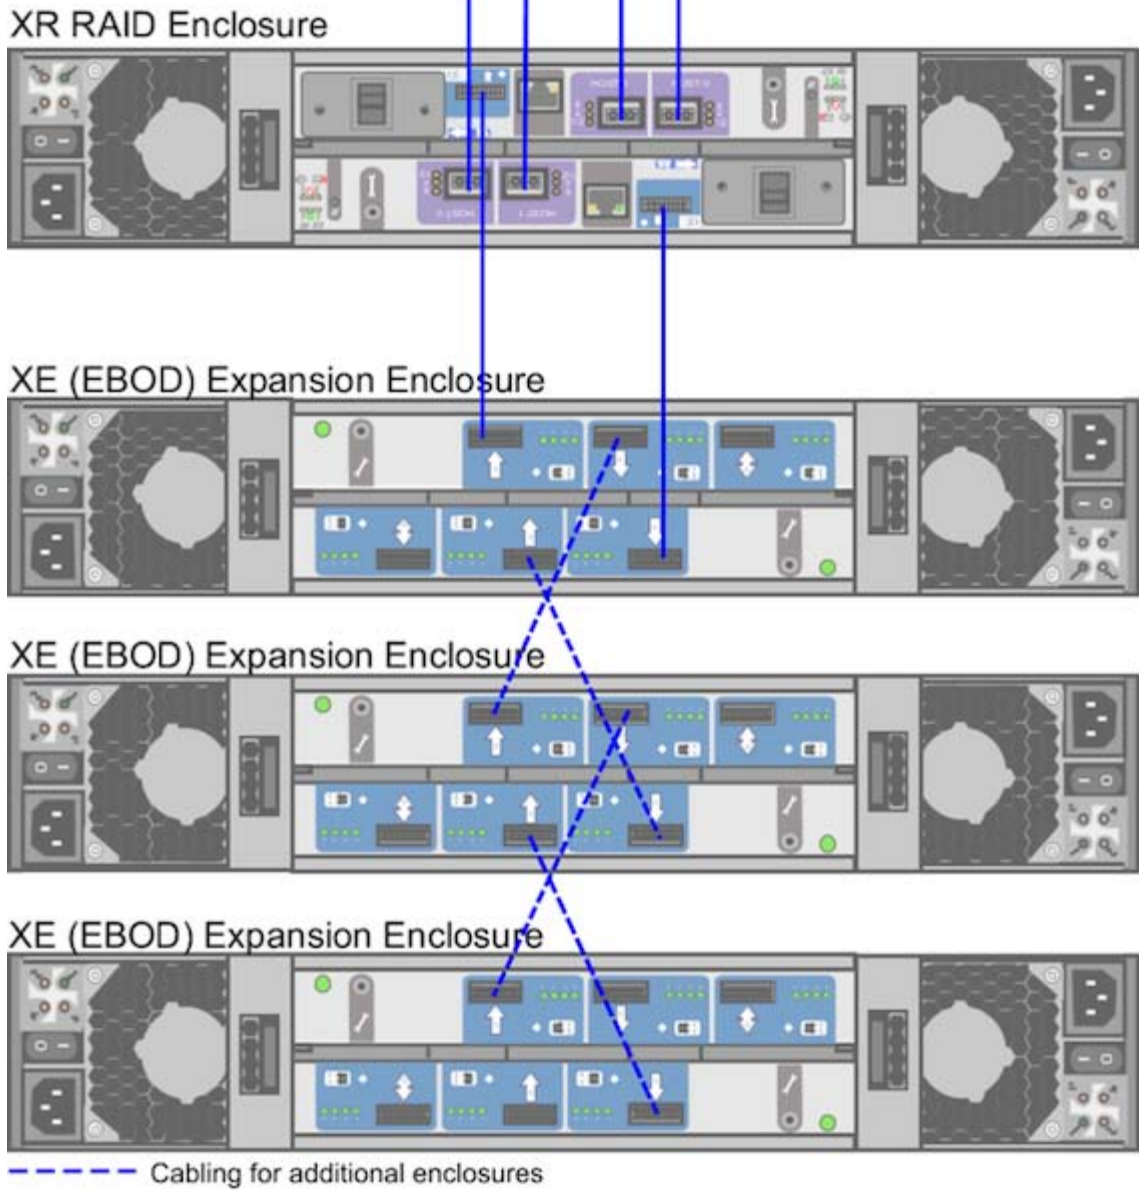


XE (EBOD) Expansion Enclosure



----- Cabling for additional enclosures

- To FC adapter on Workstation / LMS (port 1)
- To FC adapter (port 2)
- To FC adapter (port 3)
- To FC adapter (port 4)



Notes:

- In a 4-loop configuration, you need a minimum of one XE expansion enclosure attached to the XR 6500 RAID enclosure.

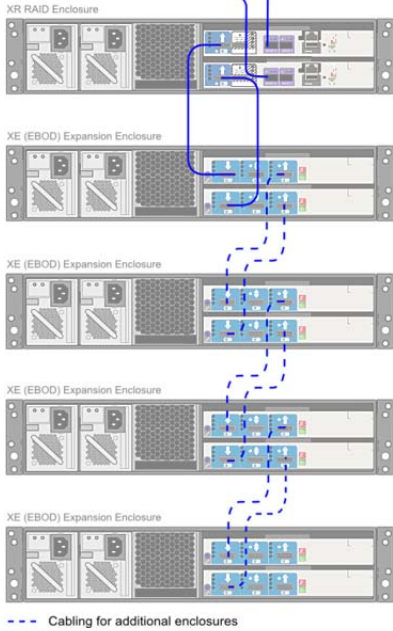
- The total number of enclosures must be an even number.

Connecting XR 6412 Storage


The following diagrams illustrate 2-loop and 4-loop connections for XR 6412 series storage assemblies. A XR 6412 RAID enclosure supports a maximum of seven XE expansion enclosures. Configurations with two XR RAID enclosures are not supported.

Single XR 6412 RAID enclosure, 2 loops

To FC adapter on Workstation / LMS (port 2)
To FC adapter (port 3)



Single XR 6412 RAID enclosure, 4 loops



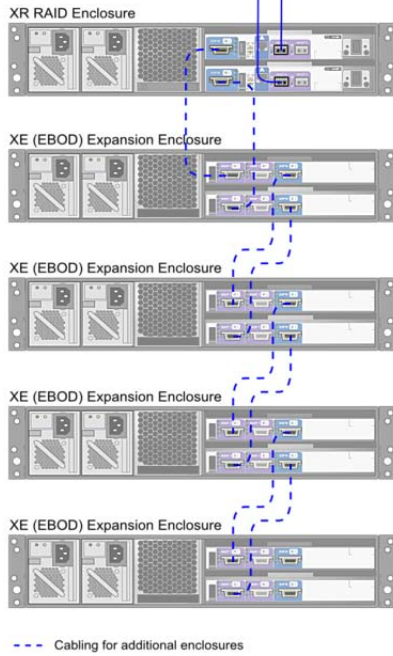
NOTE In a 4-loop configuration with a XR 6412 RAID enclosure, you need a minimum of one XE expansion enclosure attached to the XR RAID enclosure.

Connecting XR 5402 and XR 5412 Storage

The following diagrams illustrate 2-loop and 4-loop connections for XR 5402 and XR 5412 series storage assemblies. XR 5402 and XR 5412 series storage assemblies support 2-loop configurations with one XR RAID enclosure, and 4-loop configurations with two XR RAID enclosures.

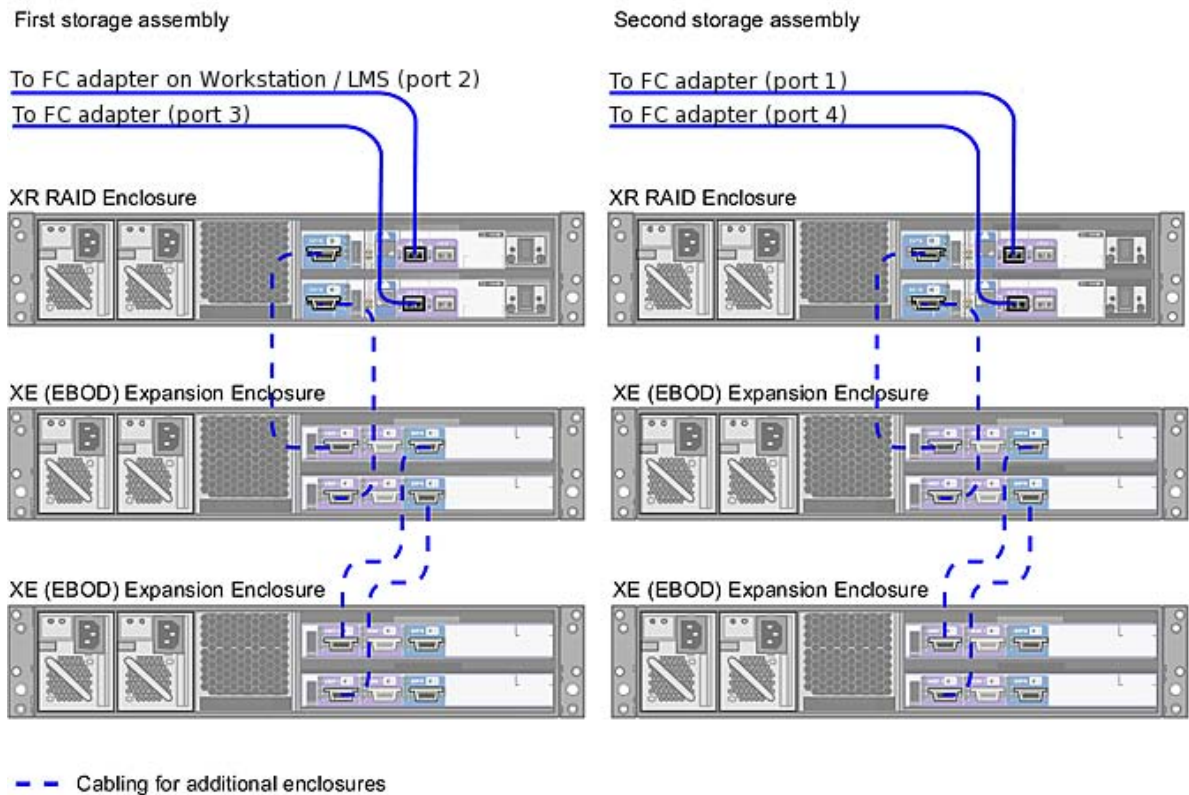
Single XR 5402 / 5412 RAID enclosure, 2 loops

To FC adapter on Workstation / LMS (port 2)
To FC adapter (port 3)



NOTE The XR 5402 and XR 5412 RAID enclosures support a maximum of four XE expansion enclosures.

Two XR 5402 / 5412 RAID enclosures, 4 loops



NOTE In a configuration with two XR RAID enclosures, the number of XE extension enclosures per XR RAID enclosure must be the same. The total number of enclosures in the storage assembly must be an even number.

Connect Archive Storage (not Lustre)

Archiving to USB 2.0, FireWire (IEEE 1394) and fibre channel devices is supported. This includes filesystems, tape drives, and VTRs. For information on connecting a VTR, see [Connect video components](#) (page 10).

Filesystem Devices

USB 2.0 and FireWire (IEEE 1394) attached devices must be use one of the following filesystems: XFS, ext2, ext3, or HFS (Mac). NTFS is not supported.

NOTE Care must be taken to ensure that the filesystems used to store archives are reliable, stable and properly backed up so as to protect the archive integrity. This functionality is not provided by the application.

SCSI Tape Drive Devices

SCSI tape drives can be connected to a fibre channel port of the workstation by way of a fibre-channel-to-SCSI bridge, such as the Atto FibreBridge 2390D. The tape drive device must meet two criteria. First, the device's driver must accept standard UNIX tape device calls. Second, the device must be certified for use with the currently supported version of the workstation operating system and kernel.

While Autodesk Media and Entertainment makes no certification statement about any tape drive device, you can use the “Tape Drive Test” plug-in in Autodesk SystemCentral to test SCSI tape drives for compatibility. For more information, see the Autodesk SystemCentral User Guide.

Configure BIOS

The BIOS of your workstation is correctly configured prior to shipping, so you normally do not need to perform the procedures in this section.

This section contains information on the certified BIOS settings for the HP Z800 workstation, as well as instructions on updating your BIOS and restoring it to the default factory settings. For information on updating your BIOS version, or reverting the BIOS to factory defaults, consult the Hardware Setup Guide.

WARNING If BIOS settings are not set up correctly before installing Linux, you will need to reinstall the OS.

Settings

After wiring the workstation and before installing Linux, turn the workstation on and check the BIOS version and BIOS settings.

The certified BIOS version for your workstation is 3.07.

Ensure the BIOS version number that appears on screen while booting the workstation matches the certified version. If the versions do not match, update the BIOS to the certified version, after configuring BIOS settings. BIOS update instructions are at the end of this chapter.

To adjust system BIOS settings, press F10 while booting the workstation.

Autodesk certified BIOS settings. Unlisted items have factory default settings.

BIOS Menu	Submenu	Item	Value
Storage	Storage Options	SATA Emulation	RAID+AHCI
	Boot Order	Optical Drive	
		USB Device	
		Hard Drive	<hard drive model>
			Integrated SATA
Power	OS Power Management	Runtime Power Management	Disable
		Idle Power Saving	Normal
		MWAIT Aware OS	Disable
		ACPI S3 Hard Disk Reset	Disable
	Hardware Power Management	SATA Power Management	Disable
		Intel Turbo Boost Technology	Disable

BIOS Menu	Submenu	Item	Value
Advanced	Processors	Hyper-Threading	Enable
	Chipset/Memory	Memory Node Interleave	Enable
		NUMA Split Mode	Disable
	Device Options	S5 Wake on LAN	Disable
		Internal Speaker	Disable
		NIC PXE Option ROM Download	Disable
	Slot 1 PCIe2 x8(4)	Option ROM Download	Disable
	Slot 3 PCIe2 x8(4)	Option ROM Download	Disable
	Slot 4 PCIe2 x16(8)	Option ROM Download	Disable
	Slot 6 PCI32	Option ROM Download	Disable

Install Linux

4

To prepare your system and perform a fresh install of Red Hat Enterprise Linux.

- 1 If reinstalling Linux on an existing system, back up all user settings, project settings, and media from the system disk to other media. The Linux installation formats the system disk, resulting in total data loss. In addition, back up the following directories:
 - `/usr/discreet` (for software setup and configuration files)
 - `/etc` (for Linux configuration files, networking, etc)
 - `/usr/local/flexlm/licenses` (for the license file)
- 2 Ensure your system is in the following state:
 - Power off all storage enclosures connected to the system. If left on, the Red Hat installer may reformat them as part of the installation process, resulting in loss of data on the device. Turn off the RAID controllers first, then the expansion enclosures.
 - Mouse, keyboard and graphics monitor are connected, and the graphics monitor is powered on.
 - Other peripherals, such as the tablet and video breakout box are connected and powered on.
 - If you are using a KVM switch, it is switched to the system on which you want to install Linux.
 - The storage array is powered off.
 - The DVI-Ramp (if any) is connected and powered on.
- 3 Insert the DVD of the custom Autodesk distribution of Red Hat Enterprise Linux in the DVD-ROM drive, and power off the system.

WARNING If your system is equipped with internal media storage, after turning the system off, open the side panel and disconnect the main connector for the internal storage drives. Do not disconnect the system drive.
- 4 Check the BIOS settings
 - 1 Reboot the system.
 - 2 On startup press `F10` to check the BIOS settings. If the keyboard is connected to the system through a KVM switch, you may have to press `F10` repeatedly or continuously.
 - 3 Ensure the BIOS settings are correct, save them if necessary, and reboot the system.
 - 4 If the system does not boot from the Linux installation disc, review the primary boot device setting in the BIOS.
- 5 When the Red Hat installation screen appears, press `F8` to display the Autodesk Red Hat Linux installation menu.
- 6 At the boot prompt in the installation menu, enter the command `iffesbc` (do not use the default options) to launch the installation. The system may spend several minutes loading drivers. The installer guides you through the rest of the process.

- 7 If prompted, initialize the system disk. The test of the installation is automatic.
- 8 When the installation has completed, eject the disc and reboot the system. After the system reboots, you may be prompted by the Kudzu hardware setup utility to set up new hardware detected for your system. You can ignore these prompts and allow the utility's count-down to expire since the hardware for the system is unchanged.
- 9 Boot into your new Linux installation and change the root password: Login as root, with the default password password. In a terminal use the `passwd` command to change the password.
- 10 By default the time zone is set to EST. To change it, login to the desktop as root, right click on the clock, and select Timezones.
- 11 [Configure basic network settings](#) (page 24).
- 12 [Configure an InfiniBand Card](#) (page 25)
- 13 [Install the DKU and the AJA OEM-2K firmware](#) (page 25)
- 14 [Configuring Storage](#) (page 26)

Configure basic network settings

Login as root to edit the files described below in a text editor, and reboot the system for the new configuration to be used.

You'll need the following from your network administrator:

- A unique static IP address and host name for your system
- The network gateway IP address.
- The subnet mask of your network.
- DNS server IP address(es).

`/etc/sysconfig/network`

Sample snippet from `/etc/sysconfig/network`.

```
NETWORKING=yes
HOSTNAME=workstation1
GATEWAY="10.1.0.25"
```

The GATEWAY value is used if no GATEWAY is defined in a network port's configuration file.

`/etc/resolv.conf`

Sample snippet from `/etc/resolv.conf`

```
nameserver 192.9.201.1
```

`/etc/hosts`

You may need to edit the loopback setting which may look like `127.0.0.1 vxfhost.localhost.localdomain localhost` by default. Optionally add hostname / IP address pairs for other workstations on your network.

Sample snippet from file:

```
127.0.0.1 localhost.localdomain localhost
192.168.0.100 workstation1
```

`/etc/sysconfig/network-scripts/ifcfg-ethn`

Edit `/etc/sysconfig/network-scripts/ifcfg-ethn`, where *n* specifies the ethernet port number, usually 0 for the first on-board port. If your workstation has an add-on ethernet card installed, the Linux operating system may assign ports eth0 through eth3 to the add-on card. In such a case, the on-board Ethernet ports become eth4 and eth5.

Optionally set the `GATEWAY` if not set in `/etc/sysconfig/network`.

Sample snippet from file:

```
DEVICE="eth0"  
BOOTPROTO="static"  
IPADDR="192.168.1.100"  
NETMASK="255.255.0.0"  
ONBOOT="yes"  
GATEWAY=192.168.0.1
```

You'll need the following from your network administrator:

- A unique static IP address and host name for your system
- The network gateway IP address.
- The subnet mask of your network.
- DNS server IP address(es).

Configure an InfiniBand card

If the card was not present when you last ran the DKU installation, run it again to setup the drivers for the card. After changing the configuration, you must reboot for them to take effect.

If your workstation is equipped with a Mellanox QDR InfiniBand / 10 GigE adapter card, or a QLogic 9000 DDR InfiniBand adapter card, perform the procedures in this section to configure the IP address of the InfiniBand port.

Mellanox QDR InfiniBand / 10 GigE adapter card

- Settings are stored in `/etc/sysconfig/network-scripts/ifcfg-ib0`
- You can update the file by running, as root, the script `/usr/discreet/DKU/InfiniBand/ofed/install.pl`
- You may need to edit the IP address, netmask, broadcast address and base network address.

QLogic DDR InfiniBand card

- Settings are stored in `/etc/sysconfig/network-scripts/ifcfg-ib1`
- You can update the file by running, as root, the script `/sbin/iba_config`
- You may need to edit the IP address.

Install the DKU and the AJA OEM-2K firmware

Check the DKU version

In a terminal, as root, check the currently installed DKU version with the command `head -n1 /etc/DKUversion`. If the version does not meet the requirements for the new application version, install the DKU.

WARNING Do not attempt to install the current version of the DKU on unsupported platforms, including older Creative Finishing systems. This will result in the installation script only installing a generic configuration which can render your operating system unusable.

Install the DKU

WARNING Do not attempt to install the current version of the DKU on unsupported platforms, including older Creative Finishing systems. This will result in the installation script only installing a generic configuration which can render your operating system unusable.

To install the DKU:

- 1 [Prepare the installation media](#) (page 35).
- 2 Run the install script (for example from the USB device): `/mnt/usbdisk/DKU-<version>/INSTALL_DKU`. When the DKU installation script completes, a warning to update the AJA card or DVI-Ramp firmware may appear and you are returned to the command prompt.
- 3 If you installed from a disc, return to the root directory, and eject the disc with the command `eject`.
- 4 Before rebooting the system, ensure the Wacom tablet is connected. If storage devices (external disk arrays or internal storage) were previously disconnected or turned off as part of an operating system installation, reconnect and power them up. Power on the XE expansion enclosures first and the XR RAID controller units last. This ensures the RAID controllers detect the other units in the Stone Direct storage. Wait for all the disks in the storage devices to fully power up before rebooting the system. The green light for each disk stops flashing once it is fully powered up.
- 5 Reboot the system with the command `reboot`. After the workstation reboots and until you install the application, the Wacom tablet might not behave correctly. Installing the application, then logging out of the KDE desktop and logging back in should eliminate this problem.

Check the AJA Firmware

Before upgrading the AJA firmware, make sure you have upgraded the DKU, and that you have restarted the system.

- 1 In a terminal, as root, run `/usr/discreet/DKU/current/Utils/AJA_firmwareUpdate/AJAfw_update`. If you have an AJA OEM-2K card, you must start the firmware update by typing `Y` and then pressing Enter. If you have a AJA Kona 3G, the upgrade begins automatically. While the AJA firmware and drivers are being updated, your workstation appears to be frozen and your mouse and keyboard do not work. This is normal and indicates that the firmware is being updated. Once the firmware update is complete, you are returned to the terminal.
- 2 Shut down the system with `shutdown -g0`. If the system does not prompt you to power down, press the power button for 10 seconds to force a power down.
- 3 Disconnect the power cord, wait 10 seconds, reconnect the power cord, then restart the system.

Configure storage

Before you can use a new storage array, create logical units (LUNs) on it, and then create the XFS filesystem on the logical volume.

The easiest way to configure your storage is to use the DH configuration script (for Dot Hill storage) or the XR Configuration script (for XR storage) available in the DKU.

This can be done with version 1.5 of the Autodesk XR Configuration Utility `/usr/discreet/DKU/current/Utils/Storage/XR/XR_config.pl`, available in versions 4.5.1 and newer of the DKU. For help, run `/usr/discreet/DKU/current/Utils/Storage/XR/XR_config.pl --help`.

[Manually configure media storage](#) (page 29) is also possible.

The workstation must be running the custom Autodesk distribution of Red Hat Enterprise Linux, and the correct DKU.

If you are configuring an XR storage array with two sets of enclosures, connect only one set of enclosures, and then perform the filesystem creation procedures in the next section. Then disconnect the storage, and repeat the entire process for the second set of enclosures.

Workflow

- 1 Set up any Storage Area Network (SAN), before setting up a DAS.
- 2 When configuring a storage array, you can disconnect other storage devices (DAS or SAN) from the workstation so that you do not unintentionally format the wrong device.
- 3 If configuring arrays with two XR enclosures, ensure only one set of enclosures is connected to the workstation.
- 4 Create LUNs. If you have more than one XR enclosure, create all of the LUNs individually, then create the XFS file system on them all at once.
- 5 Create the XFS filesystem

Configure Dot Hill storage

- 1 Ensure the storage enclosures are connected to the workstation as documented. Connect an Ethernet cable to the Ethernet port of the top storage controller (controller A) and to an available network port on the workstation.
- 2 Configure your workstation's eth port to the same subnet as the storage controller. The default IP address of storage controller A is 10.0.0.2.
- 3 In a terminal, as root, run `/usr/discreet/DKU/current/Utils/Storage/current_DH/DH_config.sh`.
- 4 Enter the IP or hostname of the top controller of the storage enclosure.

WARNING If your storage is connected to the network switch rather than directly to the workstation, make sure you are entering the correct IP address or hostname at this step. Entering the information of another storage enclosure will lead to loss of data on that storage.

- 5 The main menu is displayed. Select option 2 to partition the storage, create the XFS filesystem and mount it. Confirm whether you want to delete event logs from the storage.
- 6 Select the number of loops to configure: 2 or 4.
- 7 If the script detects a previous configuration on the storage, it will ask you whether you want to keep it. Answer No if you want to start with a fresh configuration.
- 8 The script partitions the storage, creates the XFS filesystem and the mount point and then mounts the storage. When the operation completes you are returned to the main menu. Type 6 to exit the configuration utility.

After configuration, the storage spends a few hours initializing. During this time, performance is not optimal. You can monitor the status of the drives by logging into the RAIDar Storage Management utility.

- 1 In a web browser enter the IP address of the top storage controller, for example: `http://10.0.0.2`.
- 2 Log in. The default username is "manage" and the default password is "!manage"
- 3 Click on each of the Vdisks. The Current Job row in the table indicates the progress of the initialization operation.

Configure XR storage

Create LUNs

- 1 In a terminal, as root, run `/usr/discreet/DKU/current/Utils/Storage/XR/XR_config.pl`.

The XR Configuration Utility exits without configuring your storage if any of the following is true:

- The total number of disks is not a multiple of 12.
 - One of more of the enclosures does not have the correct firmware.
 - When using two sets of enclosures, the number of expansion chassis attached to each XR RAID enclosure is not the same.
 - An odd number of enclosures in a 4-loop configuration. Only even numbers of enclosures are supported.
- 2 The script displays information about the storage, and if a LUN configuration already exists on the storage, you are prompted for confirmation to overwrite that configuration.
WARNING LUN setup destroys data on the device.
 - 3 The utility detects the number of enclosures and drives and presents you with a list of options. Choose 2 to create LUNs with a sector size of 512 bytes. This is the optimal sector size for XFS DAS (direct-attached) storage of Creative Finishing applications.
 - 4 Choose 2-loop or 4-loop configuration. 4-loop configurations are only supported for XR 6412 and XR 6500 storage. The utility creates LUNs on your storage. This process might take a few minutes.
 - 5 When the LUN creation process completes, you are prompted to create the XFS filesystem.
 - If configuring arrays with two XR enclosures, type N . If you answer Y (the default) by mistake, you get another chance to answer no after the system rescans the disks. Do not create the filesystem until you have create LUNs on all enclosures.
 - If you have only one XR enclosure choose Y to create the XFS filesystem,

Create the XFS filesystem

- 1 After choosing Y, and type Yes to confirm the detected devices. The configuration utility partitions the storage array, groups the partitions into a logical volume, creates the XFS filesystem on the volume, and mounts the volume to `/mnt/StorageMedia` on an HP Z800, or `/mnt/stoneMedia` on older workstations. If the storage driver does not reinitialize quickly enough after LUN creation, the filesystem creation step may fail, and the XR Configuration Utility posts an error like: "Can't open /dev/sdb1 exclusively. Mounted filesystem?" The filesystem will appear mounted. Type `pvdisplay -v` in the terminal to verify that the physical volume, volume group and logical volume were created. If the physical volume, volume group and logical volume were not created, reboot your workstation, and perform the steps in [To perform only the filesystem creation operations:](#) (page 29).
- 2 Exit the configuration script with `x`.
- 3 Even though your storage array is usable at this point, the drives are still initializing following the LUN creation. Your storage operates at reduced performance until the initialization process completes. This might take a few hours. To check the progress of the initialization run the command `/usr/discreet/DKU/current/Utils/Storage/XR_config.pl --arrayStatus`
- 4 Once the initialization process completes, create a back-up of the storage configuration with the command: `/usr/discreet/DKU/current/Utils/Storage/XR_config.pl --saveConfig`.

Arrays with two enclosures

- 1 After creating the LUNs on the first set of enclosures, power off the workstation, disconnect the storage enclosures, and connect the second set of enclosures, then power the workstation back on.

- 2 Create the LUNs.
- 3 Power off the workstation again, connect both sets of enclosures to the workstation.
- 4 Power the workstation back on and Create the XFS filesystem.

To perform only the filesystem creation operations:

- 1 In a terminal, as root, run `/usr/discreet/DKU/current/Utils/Storage/current/XR_config.pl --createFS`. The configuration utility partitions the storage array, groups the partitions into a logical volume, creates the XFS filesystem on the volume, and mounts the volume. This is the same as running the `XR_config.pl` with no parameter, and choosing option 4 from the main menu.
- 2 Type Yes to confirm.
- 3 Following LUN creation, drive initialization may take a few hours. The storage array is usable during this time but operates at reduced performance. To check the progress of the initialization, run `/usr/discreet/DKU/current/Utils/Storage/XR_config.pl --arrayStatus`
- 4 Once the initialization process completes, back up the storage configuration:
`/usr/discreet/DKU/current/Utils/Storage/XR_config.pl --saveConfig`.

How to power storage on or off

Power your system and storage up or down in the proper sequence. An incorrect power up sequence can mean your system does not recognize all drives.

Power on a system:

- 1 Ensure your workstation is shut down.
- 2 Power up the XE expansion enclosures.
- 3 Power up the XR RAID controller enclosures.
- 4 Wait about 90 seconds for all the drives to spin up. Their lights are solid green when they are spun up.
- 5 Power on your workstation.

Power off a system:

- 1 Shut down your workstation.
- 2 Shut down the XR RAID controller enclosures.
- 3 Shutdown the XE expansion enclosures.

Manually configure media storage

Overview

This section provides advanced instructions for manually configuring a storage array. Normally, you do not need to perform any of these tasks, as they are all automatically performed by the `XR Configuration Utility`. This utility is installed with the DKU and automates the LUN creation process.

NOTE Make sure you do not re-partition the system drive or any other disks that you want to preserve as they are. Partitioning destroys all data on the disks.

Creating Hardware LUNs

LUNs, also referred to as Logical Units or Logical Drives, are groups of disk drives that are striped together to provide optimal performance and RAID protection. Once configured, LUNs are seen by the Linux operating system as if they were single disk drives.

For systems with two sets of enclosures, you have to configure one set at a time with the `XR Configuration Utility`. Connect the first set of enclosures, and use the utility to configure it. When done, disconnect the first set and connect the second set. When the second set of enclosures is configured, re-connect both sets.

To configure LUNs on XR-series storage:

- 1 Open a terminal and log in as root and run `/usr/discreet/DKU/current/Utils/Storage/current/XR_config.pl`. The utility detects whether a LUN configuration exists on the storage attached to that workstation.
- 2 If a LUN configuration already exists on the storage, you are prompted for confirmation to overwrite that configuration.

WARNING LUN configuration is destructive. Make sure you want to overwrite an existing configuration before you confirm.
- 3 After the script detects the number of enclosures and drives, it prompts you to indicate the filesystem your storage uses. Type 2.
- 4 When asked if you have a 2-loop or a 4-loop configuration, select the option that applies to your storage. The `XR Configuration Utility` configures your storage.
- 5 Type `x` to exit the `XR Configuration Utility`.
- 6 Reboot your workstation, so that the newly-created LUNs are rescanned by the operating system.

The `XR Configuration Utility` exits without configuring your storage if any of the following is detected:

- An incorrect number of disks. The total number of disks must be a multiple of 12.
- One or more of the enclosures do not have the correct firmware.
- In a dual RAID enclosure environment, the number of expansion chassis on each RAID enclosure is not the same.
- An odd number of enclosures in a 4-loop configuration. Only even numbers of enclosures are supported.

Partitioning Disks or LUN devices as Primary Partitions

To achieve optimal performance, each disk or LUN in the array should be partitioned as a single primary partition.

On storage arrays with 450 GB drives, use the `parted` utility to create GPT (GUID Partition Table) type partitions. On arrays with smaller drives, use the `fdisk` utility to create Linux LVM type partitions.

To partition disk or LUN devices with 450 GB drives or larger:

- 1 Reboot your system to reload the fibre channel adapter drivers.
- 2 Open a terminal, and log in as root and view a list of disks or LUN devices detected by the operating system, using the following command: `fdisk -l | grep dev`. Identify the disk or LUN devices that are part of the storage array to be configured with a standard filesystem. These devices will be re-partitioned.
- 3 Use the `parted` command to re-partition each disk device identified in the previous step: `/sbin/parted -s -- <disk name> mklabel gpt mkpart primary 0 -1` where `<disk name>` is the name of a disk device identified in step 1, without a partition number, such as `/dev/sdb`.

- 4 Repeat for each disk.

To partition disk or LUN devices with drives smaller than 450 GB:

- 1 Reboot your system to reload the fibre channel adapter drivers.
- 2 Open a terminal, and log in as root and view a list of disks or LUN devices detected by the operating system: `fdisk -l | grep dev`. Identify the disk or LUN devices that are part of the storage array to be configured with a standard filesystem. These devices will be re-partitioned.
- 3 If you plan to configure a standard filesystem on a former Stone FS storage array, delete the volume label and volume table on each LUN device that is part of the storage array. Type the following command for each LUN device: `dd if=/dev/zero of=<LUN device> count=4096` Where <LUN device> is the device name of a LUN in your storage array, such as `/dev/sdc`.

WARNING When using the `dd` command, be very careful to not delete your system drive (usually `/dev/sda`) or any other drive aside from the LUNs in your storage array.

- 4 Use `fdisk` to re-partition each disk device identified in the previous step: `fdisk <disk name>` where <disk name> is a disk device name without a partition number, such as `/dev/sdf`. The `fdisk` utility starts, checks the disk device, and then displays its prompt.

NOTE When `fdisk` starts, a warning about the number of disk cylinders may appear. You can disregard this warning.

- 5 Type `n` to display the New partition creation menu. `fdisk` displays the type of partitions you can create (primary or extended).
- 6 Create a primary partition on the disk device by typing `p` at the prompt.
- 7 When prompted to enter a partition number, type `1` to make the primary partition the first one on the LUN.

NOTE You may have to delete pre-existing partitions by entering `d` when prompted, and repeating step 3.

- 8 When prompted to set the starting cylinder number, press `Enter` twice to accept the defaults, which are the first and last cylinder on the device. The `fdisk` prompt reappears.
- 9 Type `t` to set the partition type. You are prompted to enter the hexadecimal code of the partition type to be created on the LUN.
- 10 Type `8e` to set the partition type to Linux LVM. `fdisk` sets the partition as Linux LVM and the following output appears: `Changed system type of partition 1 to 8e (Linux LVM)`
- 11 Type `w` to save the new partition table.
- 12 Repeat steps 2 through 9 for each disk or LUN device identified in step 1.

Assembling the Disk or LUN Devices into a Logical Volume

After you have formatted each disk or LUN device as a partition, you must assemble the LUNs into a single logical volume on which you create the XFS filesystem. This procedure does not cover creating fault-tolerance and assumes that the LUNs are RAID-protected, as is the case with Stone Direct XR-series arrays.

To assemble a logical volume:

- 1 Verify that the disk or LUN devices are detected by the operating system: `fdisk -l | grep dev` All devices appear in a list similar to the following example (your values may vary):

```
/dev/sdb1 1 180482 1449713663+ ee EFI GPT
/dev/sdc1 1 180482 1449713663+ ee EFI GPT
/dev/sdd1 1 180482 1449713663+ ee EFI GPT
/dev/sde1 1 180482 1449713663+ ee EFI GPT
```

Partitions created with the `parted` command for arrays with 450 GB disks are marked “EFI GPT”. Partitions created in `fdisk` for arrays with smaller capacity disks are marked “Linux LVM”. Other devices of different types may be listed before and after the GPT or LVM devices.

- 2 Create a physical volume on each of the devices: `pvcreate <list of devices>` where `<list of devices>` is a list of all the devices in the storage array. For example, if you have four devices, ranging from `/dev/sdb1` to `/dev/sde1`, you would type: `pvcreate /dev/sdb1 /dev/sdc1 /dev/sdd1 /dev/sde1`. The physical volumes are created.

TIP You can use the command `pvremove` to delete any erroneously entered devices.

- 3 Verify that the physical volumes were initialized correctly: `pvscan -v`. A list of all of the physical volumes you created appears. The following sample output is for the previous example of 4 physical volumes created on devices `/dev/sdb1` through `/dev/sde1`:

```
Walking through all physical volumes
PV /dev/sdb1 lvm2 [2.03 TB / 2.03 TB free]
PV /dev/sdc1 lvm2 [2.03 TB / 2.03 TB free]
PV /dev/sdd1 lvm2 [2.03 TB / 2.03 TB free]
PV /dev/sde1 lvm2 [2.03 TB / 2.03 TB free]
Total: 4 [8.10 TB] / in use: 0 [0 ] / in no VG: 4 [8.10 TB]
```

- 4 Create the volume group “vg00” from the physical volumes you created in the preceding step: `vgcreate vg00 <list of volumes>` where `<list of volumes>` is the list of physical volumes you created in the preceding step.

TIP You can use the command `vgremove` to delete any erroneously entered volume.

- 5 Verify the volume was created and obtain the value of the “Free PE / Size” field: `vgdisplay -v`. In the output, find the line that contains the “Free PE / Size” field and write down the value of the “Free PE”. For example, in the following example output the “Free PE” value is 2124556. Free PE / Size 2124556 / 8.10 TB

- 6 Create a new logical volume on “vg00”: `lvcreate -l <Free_PE_value> -i <#_of_physical_volumes> -I 32 -n lv011 vg00` where `<Free_PE_value>` is the “Free PE” value you noted in the preceding step and `<#_of_physical_volumes>` is the number of physical volumes. If we continue with the example used in the previous steps, you would type: `lvcreate -l 2124556 -i 4 -I 32 -n lv011 vg00`. The output confirms the creation of the logical volume: Logical volume “lv011” created

NOTE If the command outputs several lines about a file descriptor leaked on `lvdisplay` invocation, ignore them.

- 7 Check if the `adsk_lvm` startup script has been installed by the DKU to enable automatic logical volume reassembly upon reboot: `chkconfig --list | grep adsk_lvm`. If the script is properly configured, the command output is: `adsk_lvm 0:off 1:off 2:on 3:on 4:on 5:on 6:off`. If the command output is different, enable the script with:

```
chkconfig --add adsk_lvm
chkconfig adsk_lvm on
```

Creating the XFS Filesystem on the Logical Volume

After having created the logical volume, you are now ready to create and mount the XFS filesystem.

To create and mount an XFS filesystem:

- 1 Identify the optimal `agsize` value for your array by running the `mkfs.xfs` command: `mkfs.xfs -d agcount=128 -f /dev/vg00/lv011`. This command displays diagnostics information similar to the following (your values may differ):

```

meta-data=/dev/vg00/lvol1 agcount=128, agsize=1066667 blks
          =                sectsz=512 attr=0
data      =                bsize=4096 blocks=134400000, imaxpct=25
          =                sunit=16   swidth=64 blks, unwritten=1...

```

- 2 From the diagnostic information printed in the previous step, note: `agsize` on the first line, `sunit` and `swidth` on the fourth line.
- 3 Depending on the values of `sunit` and `swidth`, calculate a new `agsize` value using one of the following three methods:
 - If the values of `sunit` and `swidth` are both equal to 0, multiply the `agsize` value by 4096. For example (your values will differ): $1066667 * 4096 = 4369068032$. Proceed using the value calculated above as the new `agsize` value.
 - If the command displays a warning message about the `agsize` being a multiple of the stripe width, multiply the `agsize` value by 4096, and subtract the `sunit` value multiplied by 4096. For example (your values will differ):

```

1066667 * 4096 = 4369068032
16 * 4096 = 65536
4369068032 - 65536 = 4369002496

```

Continue using the value calculated above as the new `agsize` value.

- If the values of `sunit` and `swidth` are not equal to 0, and no warning message appears, proceed to step 4 using the `agsize` value displayed by the `mkfs.xfs` command in step 1.
- 4 Run `mkfs.xfs` again to create the XFS filesystem on the device `/dev/vg00/lvol1` using the value calculated in one of the previous steps: `mkfs.xfs -d agsize= <new agsize> -f /dev/vg00/lvol1`. The filesystem is created on the storage array.

NOTE If the command fails, redo your calculations starting from step 1.

- 5 Verify that the storage can be mounted by running one of the following commands:
 - For HP Z800 systems: `mount /mnt/StorageMedia`
 - For older systems: `mount /mnt/stoneMedia`

The storage should mount, as the DKU installation script should have created the mount point directory for your storage (`/mnt/StorageMedia` on HP Z800 workstations, or `/mnt/stoneMedia` on older workstations), as well as the corresponding entry in the `/etc/fstab` file. If you receive an error message and the storage does not mount, follow the instructions in the next section to manually mount the storage.

Manually Creating a Mount Point and Mounting the Storage

If the mount point directory for your storage was not created automatically by the DKU, or if the storage does not mount, create the mount point and mount the storage manually:

- 1 Create the directory that will serve as the mount point for the filesystem, if it does not exist. For example: `mkdir /mnt/StorageMedia`

WARNING Do not use the word “stonefs” as the name for your mount point directory. “Stonefs” is a reserved word, and can cause issues if used as the mount point directory name.

- 2 Mount the XFS filesystem from the logical volume `/dev/vg00/lvol1` on the directory you created in the previous step. For example: `mount -av -t xfs -o rw,noatime,inode64 /dev/vg00/lvol1 /mnt/StorageMedia`. The filesystem is mounted as `/mnt/StorageMedia`.
- 3 Confirm that the storage is now mounted: `df -h`. The output should list `/dev/mapper/vg00-lvol1` mounted on your mount point directory.

- 4 Append a line to `/etc/fstab` so the filesystem is mounted at startup, for example: `/dev/vg00/lvol1 /mnt/StorageMedia xfs rw,noatime,inode64`
- 5 Optional: Confirm that the filesystem can mount automatically by rebooting the workstation and using the command `df -h` again.

Installing Your Software

5

Prepare the installation media

Check the release announcement to find out on what media the installers are available.

Major releases are distributed on a USB device. To mount a USB device:

- Attach the device. Log in to the terminal as root.
- Use the `dmesg` command to output something like `sdf: sdf1` to list a recent device connected. Or list the devices attached to your system with the command: `fdisk -l | grep "/dev/sd"`. The device assigned to the USB device should look similar to `/dev/sds` or `/dev/sde`
- On a fresh Red Hat installation, automount is enabled, but execution from the device is disabled. So you must unmount the USB drive with `eject /<mountpoint>` or `eject /dev/<device id>`.
- Create a mount point directory for the disk with `mkdir /mnt/usbdisk`.
- Mount the device `mount -o shortname=winnt /dev/<device>1 /mnt/usbdisk`. Or from the GUI browse to `/dev`, right-click the USB device, and choose Mount.

Software is sometimes distributed as tar files. To extract from a tar file:

- 1 In a terminal, as root, use the `md5sum` command to verify the checksum matches the `md5sum` listed in the checksum file.
- 2 Extract the from the tar archive with `tar -zxvf filename.tar.gz`.

Install Flame Premium

Prerequisites

- If you need to change your system date or time, do it before installing the application.
- Archive existing media on the framestore (highly recommended). Refer to the “Archiving” topic in the application help.
- [Prepare the installation media](#) (page 35) to access the install directory.

Install the software:

- 1 Log in to your workstation as root and open a terminal. Go to the application installation directory. Start the application installation script. Use with the command `./INSTALL_APPLICATION`, where `SOFTWARE` is the name of the software to be installed, e.g. `./INSTALL_FLAMEPREMIUM`. If installing an educational version, add the switch `--edu`, e.g. `./INSTALL_FLAMEPREMIUM --edu`. The installation script verifies the installation packages and checks your system for previously installed versions of the application. This process may take a few minutes.
- 2 If earlier installed software is detected you are prompted to choose None for a clean installation of the software. Choose one of the previous versions to copy its custom resource and setup directories to the new version.
- 3 You are prompted to create an application user: choose Yes to create a Linux user account for the application, or No to run the application with one or more existing user account(s). Use CTRL+click to select multiple entries.
- 4 When prompted to configure the software initialisation file (*init.cfg*), click OK. The software initialisation file contains general settings the application reads on startup. Check the keywords. They are fully documented in the User Guide. It is recommended to verify that the following subset of keywords are correct for your hardware configuration.
 - If installing the for the first time, `init.cfg` file is opened in a text editor. If the file opened in a text editor, edit and save `init.cfg`.
 - If upgrading, the old and new `init.cfg` files are opened on the left and right sides respectively of an `xxdiff` window. Each difference between the old and the new files is highlighted. Click the correct value for each keyword, regardless of whether it is in the left or right panel. Make sure you select a value for each highlighted difference. Then open the File menu and choose Save as Right. After the file is saved, close `xxdiff`.

Keyword	Setting
Video	The video card used by your workstation. Can be either <code>AJA_KONA3G</code> or <code>AJA_OEM2K</code>
DVIRamp	Used if your workstation is equipped with a Miranda DVI-Ramp.
VideoPreviewDevice	Used for workstations with a Miranda DVI-Ramp or an NVIDIA SDI card. You can use the <code>VideoPreviewDevice</code> keyword to enable alternate SD or HD broadcast monitors.
TabletDriver	<code>wacom</code>
VTR	Set the VTR keywords that have settings corresponding to the VTRs you intend to use.
Audiodevice	<code>AJA</code>

- 5 If any changes are detected in the `/etc/X11/xorg.conf` file, you are prompted to configure this file. The `/etc/X11/xorg.conf` file contains specific parameters for the graphics display, tablet, and other external devices. It is customized prior to shipping to reflect your configuration. Editing this file is risky and can result in an inability to launch the user interface and other problems. Back up the file before changing it. Use `xxdiff` to move entries from one version of the file to another.
 - To keep the old configuration settings, when the file appears in the `xxdiff` editor, in the Global menu, choose Select Left, and then, in the File menu, choose Save as Right. If prompted to overwrite, click OK.

- To use the new configuration file without adding the settings from the old configuration file, just close the program.
- 6 If you modified the `xorg.conf` file, you are prompted to restart the X server. When the installation script completes, log out of the Linux desktop and then log in again to restart the X server. The application, as well as additional components (WiretapCentral, Wiretap Gateway, Backburner Server, Backburner Manager, Backburner Media I/O Adapter) are installed on your workstation. By default, Backburner Server and Backburner Manager are set to automatically run on the workstation. The manager for the local Backburner Server is set to `localhost`. You can reconfigure these services after the application is installed. See [Configure Backburner services](#) (page 42).

Install Lustre stand-alone

Lustre is included in the [Install Flame Premium](#) (page 35) process. The following is for a stand-alone Lustre installation.

Prerequisites

- [Prepare the installation media](#) (page 35) to access the install directory.
- If you need to change your system date or time, do it before installing the application.

To install Lustre:

- 1 In a terminal, as root, go to the install directory and run `./INSTALL_LUSTRE`
- 2 If prompted with a message asking if you would like to automatically run Backburner Server on this machine, click Yes. The Lustre application, as well as supporting software (Backburner, Wiretap, Wiretap Gateway, the server) are installed. A new user account is created for the version of Lustre you just installed.
- 3 Configure Lustre.

Uninstall

- 1 If you are logged in as the application user in KDE, log out and log back into KDE as root.
- 2 From the KDE menu, choose Autodesk > Remove Software.
- 3 Select the packages you want to uninstall in the RPM list on the left (click Select All to select all the packages), then click the arrow to move them to the RPM uninstall list on the right then click Next.
- 4 In the Choose folders window select the application directories you want to remove from the `/usr/discreet` directory, and click Next.
- 5 In the confirmation dialog click Uninstall & Remove to confirm the removal of the selected packages and directories. The uninstallation starts and displays details on the process.
- 6 Optional: You can also delete the log files associated with a given application version in `/usr/discreet/log`, or `/var/log/` for Flame.

Configure Visual Effects and Finishing Software

This is necessary for new installations. For an upgrade you do not need to repeat these procedures, unless adding new media storage devices. The only task you might need to perform when upgrading from a previous version is upgrading projects to the current version and transferring clip libraries.

Configure a Direct Attached Storage or SAN as Media Storage

You can use your application with the following storage configurations (but not the system drive):

- A standard UNIX-compatible filesystem on a DAS (such as an Autodesk-recommended Dot Hill or XR-series disk array)
- A SAN. When attached to a CXFS SAN declared as a standard filesystem partition to Stone and Wire, Visual Effects and Finishing workstations running the current release have shown optimal (real-time) performance with version 4.02 of the CXFS client, and the following mount options for the CXFS volume:
`rw,noatime,filestream,inode64.`
- A standard UNIX-compatible filesystem on a Network Attached Storage (NAS) based on the Network File System (NFS) protocol. The fact that all I/O operations occur over the network makes it impossible to reliably predict steady performance for NAS filesystems. Also that the NFS protocol features built-in cache and time-out mechanisms, which may cause several workstations working on the same media source to appear de-synchronized.

Configure your Application to Use the Media Storage Filesystem

- 1 Ensure the disk array or SAN storage is connected to your workstation and powered up, and that a UNIX-compatible filesystem exists on the storage.
- 2 In a terminal, as root, stop Stone and Wire with the command: `/etc/init.d/stone+wire stop`
- 3 Create the Managed Media Cache directory or directories:
 - Create a point for your storage e.g. `mkdir -p /mnt/StorageMedia`. Do not use the reserved word “stonefs” in the directory name. For a DAS, the DKU installer should have already created a mount point for it at `/mnt/StorageMedia` (for HP Z800 workstations) or `/mnt/stoneMedia` (for older workstations).
 - Mount the filesystem e.g. `mount -av -t xfs -o rw,noatime,inode64 /dev/vg00/lv011 /mnt/StorageMedia`. To automatically mount at boot, create an entry for it in the `/etc/fstab`.

- On the mounted filesystem create the Managed Media Cache directory. The name should include the partition name and the hostname. Set permissions for the directory to be readable, writable, and searchable by the owner and its group, and readable and searchable by all:

```
1 mkdir -p /mnt/StorageMedia/smoke1/p7
2 chown -R root:users /mnt/StorageMedia/smoke1/p7
3 chmod -R 775 /mnt/StorageMedia/smoke1/p7
```

- Optionally create separate mount points and Managed Media Cache directories for each additional volume. The DKU only creates the mount point and `/etc/fstab` entry for the initial volume `/dev/vg00/lvol1`.

- 4 Configure Stone and Wire by editing the the following keywords in `/usr/discreet/sw/cfg/stone+wire.cfg`. Usage is documented in comments in the file:

- Path
- Name
- Shared
- [DefaultFileFormats]
- SymlinkAcrossFilesystems

- 5 If this is the first filesystem you are configuring for this workstation: Get the `FRAMESTORE` ID, e.g. `grep "FRAMESTORE" /usr/discreet/sw/cfg/sw_framestore_map` and use the ID value to update `/usr/discreet/sw/cfg/sw_storage.cfg` e.g.

```
[Framestore]
ID=myworkstation
```

- 6 Optionally [Configure bandwidth reservation](#) (page 60).
- 7 Restart Stone and Wire with: `/etc/init.d/stone+wire restart`
- 8 Check the filesystem is mounted: `/usr/discreet/sw/sw_df`.

Any project created in your application after this point has the new standard filesystem partition as its primary partition. Any pre-existing projects are unaffected by this setting and remain associated to their respective partitions.

Set up the Wire Network

- Wire is a component that enables the high-speed transfer of uncompressed video, film, and audio between Autodesk workstations, over industry-standard TCP/IP and InfiniBand networks.
- Wiretap is a networking protocol that allows products that do not run Stone and Wire to access the filesystem of a Visual Effects and Finishing product.

Stop all Stone and Wire daemons before editing any of the following configuration files.

- 1 In a terminal, as root, stop Stone and Wire services with: `/etc/init.d/stone+wire stop`
- 2 Specify the network interfaces to be used by Wire in `/usr/discreet/sw/cfg/sw_framestore_map` according to the instructions in the `sw_framestore_mapman` page.
- 3 Configure Wire self-discovery in `/usr/discreet/sw/cfg/sw_probed_cfg` file in a text editor. Ensure the value of the `Port` keyword matches the value used by the other workstations in your network and `SelfDiscovery` is set to "yes". If these values are set properly, Stone and Wire uses its self-discovery mechanism to locate all other Wire hosts on your network and make their framestores accessible to your application.

- 4 To enable remote access to clip libraries, in `/etc/exports` specify all directories you want remote hosts to be able to access. You can export the entire `/usr/discreet` directory to enable remote access to all clip libraries, projects, user preferences and setups. Or you can export only certain subdirectories, such as `/usr/discreet/clip` or `/usr/discreet/project`. For example, to export the `/usr/discreet/clip` directory, add the following line to the end of the `exports` file: `/usr/discreet/clip *(rw, sync, no_root_squash)`. Save the updated `exports` file, and run the command `exportfs -a` to apply the changes.
- 5 To configure the Wiretap server edit `/usr/discreet/wiretap/cfg/wiretapd.cfg`. The house network IP address should already be present in the file, as it is added automatically during installation. Specify any additional IP addresses Wiretap should broadcast for self-discovery, for example: `IP1=10.10.0.1IP2=10.10.11.200`. Optional: If your network comprises Windows or Mac OS X systems that will use Wiretap to access media from your workstation, configure the Wiretap Path Translation service so that media paths are accurately translated between operating systems. See [Visual Effects and Finishing applications Wiretap path translation](#) (page 54).
- 6 To restart and test Wire: `/etc/init.d/stone+wire restart`. To test Wire is working: `/usr/discreet/sw/tools/sw_framestore_dump` which should show all Wire hosts discovered on your network.

Upgrade projects to the current version

You do not need to perform this task if you are upgrading to a service pack of the same version or to a service pack of the same extension.

Use the `copyProjects` command-line tool to upgrade your projects to the current version. Using this tool instead of manually re-creating your projects is the recommended approach, as it reduces the time needed to migrate multiple projects, and prevents human error. `copyProjects` does not copy any media. It only re-creates projects and copies their settings and setups to the current version. You still have to transfer your media using the network library. :

- 1 In a terminal, as root, run `/usr/discreet/<product_home>/bin/copyProjects`.
- 2 It detects all projects from previous versions, and prompts you to continue.
 - If the version returned by the script is the newly installed version, answer Yes and confirm. It creates a copy of each project for the current version. The string `_copy` is appended to the end of the new project name. Then use the network library to transfer media into the copied projects.
 - Otherwise, change the application version to which projects are copied:
 - 1 In a terminal, as root, Open a terminal and log in as root, run `/usr/discreet/<product_home>/bin/selectWiretapVersion.py`.
 - 2 Select the version you want to copy projects to.
 - 3 When finished, run the `copyProjects` script.

Transfer clip libraries to current projects

To access clip libraries from projects created with any Visual Effects and Finishing software between versions 2007 to 2011, you must transfer the clip libraries to a project created in the current version. You can install the current version and use Wire to transfer the clip libraries to a current project. See your application help. For a Linux system where you do not wish to install the current version, or for IRIX systems, you can also use the following procedure to transfer media to the current version.

- 1 Use Backdraft Conform 2007 or any 2008, 2009, 2010, or 2011 application to create a project and a clip library on the 2013 system. See your application help. The project and library is formatted for the software version it was created with.

- 2 Use the 2007, 2008, 2009, 2010, or 2011 application to move clips over the network into the previously created project on the 2013 system. You can use Background Wire to help free up the resources of the workstation running the older version of the application. See your application help.
- 3 Create a new project on the current version system. If you have many projects, use the `copyProjects` script.
- 4 Use the current application to open the library created in step 1 (the library is read-only).
- 5 Use Dual View to drag the clips into the 2013 project and library.
- 6 Use the older application version to remove the project created in step 1. 2013 applications cannot remove projects created with previous versions because they are read-only.

TIP If you have previous versions from 2007 installed on your current version system, you can access the older system remotely to transfer the media using Wire without interrupting a session on that station. See the application help.

Configure Backburner services

Background Wire transfers and proxy generation rely on Backburner. Backburner can also be used to manage background processing jobs sent to Burn render nodes, and distributed transcoding jobs to import certain media formats, such as REDCODE RAW or OpenEXR.

Backburner Server and Backburner Manager are installed with the Autodesk Visual Effects and Finishing application, and are set to run automatically when the workstation starts. By default, the Backburner Server installed on the workstation communicates with the locally installed Backburner Manager.

Use the following procedure to configure Backburner services on your workstation.

To configure Backburner services:

- 1 Make sure Backburner Server is set to run automatically on your workstation. Type:

```
chkconfig --list | grep -i backburner_server
```

The command should output the following information:

```
backburner_server 0:off 1:off 2:on3:on4:on5:on6:off
```
- 2 If `backburner_server` is not set to “on” for run levels 3, 4, and 5, type the following commands to enable it:

```
chkconfig backburner_server on  
/etc/init.d/backburner_server restart
```
- 3 Decide if you want to use the locally installed Backburner Manager, or a centralized Manager installed on another system in your network. It is recommended to use a centralized Manager if several workstations on your network use the same farm of Burn render nodes.

Consult the [Autodesk Backburner Installation Guide](http://www.autodesk.com/backburner-installation-guide) at www.autodesk.com/backburner-documentation for detailed guidelines.
- 4 If you do not want to use the local Backburner Manager installed on your workstation, type the following commands to disable it:

```
chkconfig backburner_manager off  
/etc/init.d/backburner_manager stop
```
- 5 If you stopped the local Manager, open the `/usr/discreet/backburner/cfg/manager.host` file on your workstation in a text editor, delete “localhost”, and enter the hostname of the system where Backburner Manager is installed.

- 6 Make sure the following Backburner keywords in the `/usr/discreet/<product_home>/cfg/init.cfg` file are configured properly.

Keyword	Setting
BackburnerManagerHostname <hostname>	Set <hostname> to the hostname or IP address of the workstation running Backburner Manager. This hostname or IP address must match the one in the <code>/usr/discreet/backburner/cfg/manager.host</code> file. If it does not, background Wire transfers will fail.
BackburnerManagerPriority <priority>	Set <priority> to a value from 0 to 100. The default is 50.
BackburnerManagerGroup <groupname>	Set <groupname> to the name of a group of computers on a Burn® rendering network. For example, if the name of the group is "renderfarm1", you would set this keyword to <code>BackburnerManagerGroup renderfarm1</code> .

Configure Lustre slave rendering

The Slave Renderer renders modified frames when the artist moves to the next shot on the timeline from the Master Station. With slave rendering, playback is enabled without compromising the interactivity of the Master Station during creative sessions.

Prerequisites

- 1 [Install Linux](#) (page 23) and the DKU.
- 2 Connect the Slave Renderer to a switch on your house network.

After installing Linux on the Slave Render, configure the IP of the ethernet port that connects the Slave Render to your house network. The address you choose must not conflict with any of the other IP addresses on the network. In a Lustre installation in which you have not altered any of the IP addresses used by default, you can use 193.1.1.2 as the Slave Renderer IP address.

- 1 Start the Slave Renderer and log in as root. Assign an IP address to the ethernet port that connects the Slave Renderer to the network switch. The IP address must be in the same IP class as your network. See `/etc/sysconfig/network-scripts/ifcfg-eth#`.
- 2 To install the Slave Renderer software, start the Slave Renderer and log in as root. Navigate to the Lustre install directory whether or a USB device or from an unpacked tar file and run `./INSTALL_LUSTRE_SLAVE`.
- 3 To install the Slave Renderer license: get a license for the Slave Renderer. You will need the `dlhostid` of the Slave Renderer. Log in to the Slave Renderer as root and create the directory for the license file, e.g: `mkdir -p /usr/local/flexlm/licenses/`. Create a file called `DL_license.dat` in the new directory and copy the license code into it.
- 4 If not Using BrowseD with the Slave Renderer, configure read and write access on the storage filesystem, and mount it on the slave renderer. The storage filesystem is mounted as `/mnt/StorageMedia HP Z800`, or `/mnt/md0` for older systems:
 - 1 Log in to the Master Station as root. Open `/etc/exports` and add a new line for the mount point, e.g.: `/mnt/StorageMedia *(rw,no_subtree_check,async) OR /mnt/md0 *(rw,no_subtree_check,async)`. This creates a network share for the mount point for the storage filesystem. Other users on the network will access the storage devices by mounting its network share.

- 2 Activate the new network share for the storage device with `exportfs -a`. The list of network shares is refreshed from the `exports` file.
- 3 Confirm the storage devices are available to be mounted with `exportfs`.
- 4 Create a mount point on the Slave Renderer and set permission: on the Slave Renderer, as root in a terminal, for example


```
mkdir /mnt/StorageMedia
chmod 666 /mnt/StorageMedia
mount <Master Station>:/mnt/StorageMedia/ /mnt/StorageMedia
```
- 5 To configure the Slave Renderer to automatically mount the storage filesystem edit `vi /etc/fstab` and add a line for the mount point you just created, e.g.: `<IP address>:/<exported filesystem>/<mount point> nfs rw,bg,hard,intr 0 0`. Restart the Slave Renderer machine. The storage filesystem should mount automatically.
- 5 After you have configured the Master Station and the slave rendering machine to communicate, configure the “HostName string=” keyword with the slave rendering machine IP address in the Slave Renderer section of the `init.config` file.
- 6 Enable slave rendering for the project.
- 7 Render shots as you work. See [Autodesk Lustre User Guide](#)

Lustre background rendering

During background rendering, a shot on the timeline is rendered by a background rendering network. This is different from the Slave Renderer, which renders shots on a shot-by-shot basis as they are colour graded to enable improved playback performance.

Background rendering in Lustre is performed using Burn for Lustre, also known as the Lustre Background Renderer. This application is specific to Lustre and provides asynchronous background processing of Lustre render jobs. By off-loading rendering activities to remote Linux servers, Lustre stations are freed up for interactive colour grading, while background rendering is sped up by splitting the task amongst multiple hosts.

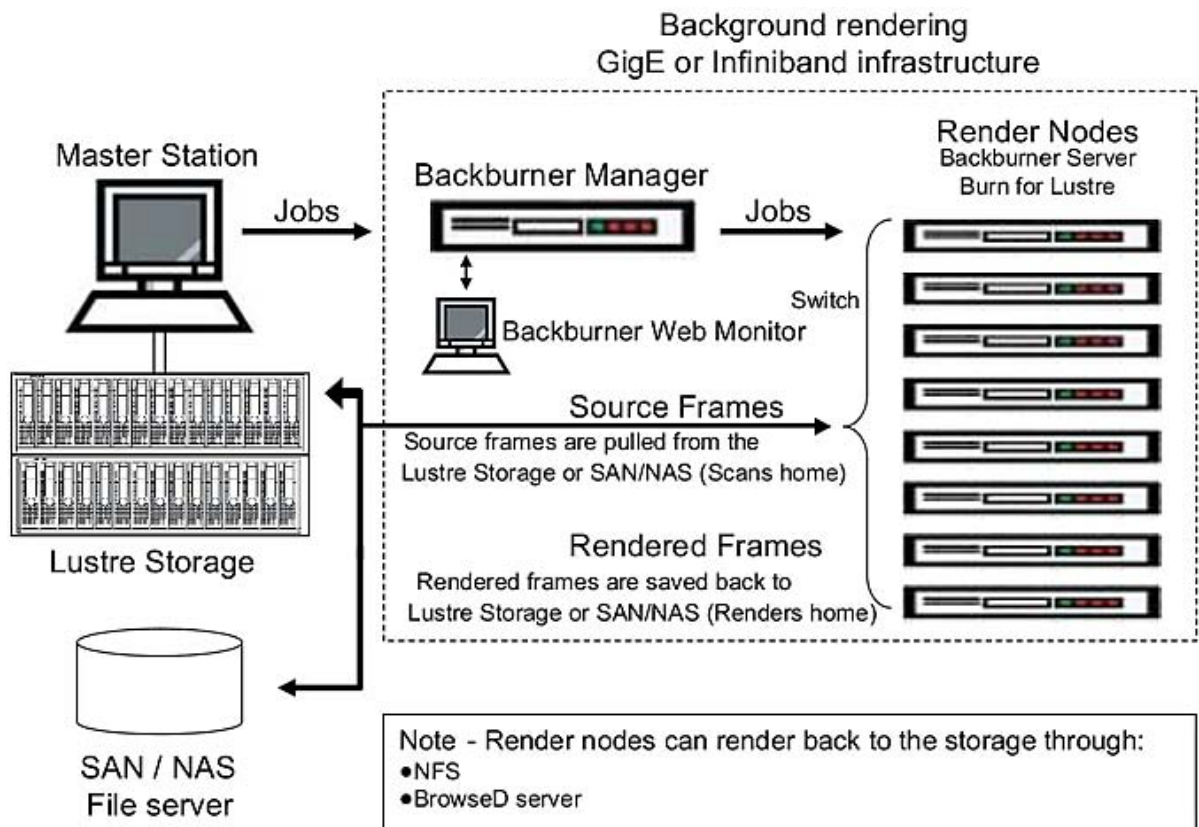
General workflow for installing and configuring background rendering:

- 1 If you are not using `BrowseD`, [Share the storage for rw access from background render nodes](#) (page 46).
- 2 [Install Backburner Manager and Backburner Web Monitor](#) (page 46).
- 3 [Configure Lustre to detect Backburner Manager](#) (page 50).
- 4 [Set up render nodes](#) (page 47).
- 5 [Specify the Background Rendering path in Lustre](#) (page 50).

Background rendering components

The components of the basic background rendering package include Lustre, a background management and monitoring application (such as Backburner™ Web Monitor, or the Backburner Monitor in Autodesk WiretapCentral™), and several render nodes running on Linux servers. The Lustre system and all background rendering nodes are connected over a dedicated background TCP/IP network. Render nodes can access media through NFS mount points, or by using the faster `BrowseD` service. Using `BrowseD` is the recommended approach. See [Configure Lustre BrowseD](#) (page 51).

The background rendering components are illustrated as follows.



NOTE You can have up to eight render nodes on the background rendering network.

The other background rendering components are described as follows.

Lustre application This is the client application. Running on a Linux workstation, Lustre rendering jobs are submitted for background rendering through the Render | Backburner menu.

Backburner Manager This is the hub of the background rendering workgroup. Backburner Manager can reside on a Master Station or Lustre station, or can be installed on a dedicated machine in your network. When jobs are submitted from Lustre to Backburner Manager, Backburner Manager breaks each submitted job into tasks and distributes the tasks to the rendering servers on the network. To view the progress of the tasks, use Backburner Monitor.

Backburner Web Monitor This is the user interface for the Backburner rendering network. It is automatically installed with Backburner Manager, and can be accessed through a Web browser from any workstation on the network. It allows you to view and control jobs currently being processed. You can stop, restart, reorder or remove jobs completely using the Monitor. You also use Backburner Web Monitor to identify any render nodes that are not working and check the overall status of the rendering network.

Backburner Server This is a server that runs on render nodes, and is automatically installed with Burn for Lustre. Backburner Server accepts commands from Backburner Manager to start and stop rendering tasks. Backburner Server communicates through a plug-in to execute rendering tasks.

NOTE Consult the latest [Autodesk Backburner Installation Guide](#) and [Autodesk Backburner User Guide](#) for detailed information on Backburner. Use the guides to obtain Backburner information that does not relate directly to Lustre background rendering, such as details on setting up Web monitoring and troubleshooting tips. You can find these guides on the [Backburner documentation webpage](#).

Burn for Lustre This is the Linux rendering engine that renders one or more frames from Lustre render jobs.

NOTE Burn for Lustre is a different software than Burn for Visual Effects and Finishing software. While they can co-exist on the same node, the installation and configuration procedures are different. For information about how to set up Burn for Visual Effects and Finishing software, consult the [Autodesk Burn Installation and User Guide](#).

Shared storage mount point This is the mount point on each Linux server that allows Burn for Lustre to transfer rendered frames/files to the Lustre storage system.

NOTE The storage does not have to be mounted on the render nodes if you are using `BrowseD` for background rendering. See [Using BrowseD for Rendering with Burn for Lustre](#).

Share the storage for rw access from background render nodes

To allow read and write access, the storage must be exported from the system to which it is connected. This makes it possible for the background rendering components on the network to remotely mount the storage system.

NOTE Skip this section if you are using `BrowseD`.

To make the storage system visible over the network:

- 1 Log in to the Lustre system (or the central storage system) as root. The storage filesystem is mounted as `/mnt/StorageMedia` (on HP Z800 workstations) or `/mnt/md0` (on older workstations).
- 2 Open for editing `/etc/exports`. The file appears, lists the available network shares (filesystems mountable over the network).
- 3 Append in the file a line for the mount point of the storage filesystem: `/mnt/StorageMedia *(rw,no_subtree_check,async)` OR `/mnt/md0 *(rw,no_subtree_check,async)`. Other users on the network will access the storage devices by mounting its network share.
- 4 Save and close the file.
- 5 Activate the new network share for the storage device. Type: `exportfs -a` The list of network shares is refreshed from the `exports` file.
- 6 Verify the storage filesystem mount points are listed in the output of `exportfs`.

Install Backburner Manager and Backburner Web Monitor

On Lustre stations, Backburner components are automatically installed with the Lustre software. The installer asks you if you want to run Backburner components on the station.

See the [Autodesk Backburner Installation Guide](#) for detailed information on configuring Backburner.

To avoid performance degradation, it is recommended to avoid running Backburner components on your Lustre stations, and to relocate Backburner Manager on a dedicated Linux machine in your network. Render jobs handled by Backburner can be viewed and managed in Backburner Web Monitor from any workstation on the network.

To install Backburner Manager and Backburner Web Monitor on a Linux system:

- 1 In a terminal, as root, access the Autodesk Lustre software installation package (from a downloaded tar file or on USB key).
- 2 Install components (Server, Manager and Web Monitor): `rpm -Uvh backburner*`.

- 3 If the system on which you are installing Backburner runs Linux in graphical mode, you are asked if you want to automatically run Backburner Manager and Backburner Server on the machine. Answer Yes to the question about the Manager, and No to the question about the Server.
- 4 Run the Backburner Web Monitor configuration script:
`/usr/discreet/backburner/backburnerConfigWeb.`

Verify that Backburner components are running properly:

- 1 On the system on which you installed Backburner Manager, as root, run: `ps -ef | grep -i backburner.` The output of the command should include a line containing the path to the Backburner Manager application, similar to: `root 14989 10 Nov29 ? 00:02:36 /usr/discreet/backburner/backburnerManager` If Backburner Manager is not running, the Backburner log file `/usr/discreet/backburner/Network/backburner.log` may be a useful source of information for troubleshooting.
- 2 To test that Backburner Web Monitor is installed and configured properly, open in a Web browser on the Manager machine the URL <http://localhost/backburner>. Backburner Web Monitor opens. If you are asked for a user name and password, use `backburner` for both.

Set up render nodes

You can set up as many as eight render nodes for background rendering with Lustre. On each system intended for background rendering, you must verify and perform the procedures below, in order.

Render node hardware and software requirements

Render nodes purchased from Autodesk come with all the necessary hardware and software components preinstalled.

Minimum hardware and OS to use Burn for Lustre on nodes that were not purchased from Autodesk

Processor:	1 or 2 Xeon Intel Processor DP
Memory:	2 GB
Hard drive:	20 GB
Ethernet:	100/1000 Mbps
OS:	Red Hat Enterprise Linux 5.3 with Workstation option, customized with the Autodesk kickstart file

Install Linux on nodes

Render nodes purchased from Autodesk ship with the correct Linux distribution. If you did not purchase your node from Autodesk, get your own 64-bit distribution of Red Hat Enterprise Linux Desktop 5.3 with Workstation option, customize it using the Autodesk kickstart file, and install it. The kickstart is used to install the packages required for Burn, some of which are not installed as part of a general Linux installation. The kickstart file can be used to automate the Linux installation process for multiple render nodes.

To copy the Autodesk kickstart file to the DVD of your Linux distribution:

- 1 On a computer running Linux and with a DVD burner, log in as root.

- 2 Insert the first installation disc of your Linux distribution into the drive. You do not need to mount it at this time.
- 3 Extract an ISO image of the disc: `dd if=/dev/ <optical_disc_device> of=/ <destination path for the extracted ISO image>`
- 4 Eject the disc.
- 5 In the Lustre installation package run `dist/build_kickstart_cd` the kickstart file is at `dist/ks_centos5_rh5.cfg`. This script adds the kickstart file `dist/ks_centos5_rh5.cfg` to the ISO image of your Linux distribution DVD: `./build_kickstart_cd ks_centos5_rh5.cfg <original ISO image name> <new ISO image name>`.
- 6 Once the new ISO image of the Linux distribution DVD is created, burn it to a blank disc. The new disc will contain the Autodesk kickstart file and replaces the disc in the Linux distribution.

Mount storage on the render nodes

You can use BrowseD to allow background render nodes to access the workstation's storage. See [Configure Lustre BrowseD](#) (page 51). Otherwise, mount the media storage on each render node using an NFS mount point.

If the system that stores the render files is different from the system that stores the Project Home directories, you need two (or more) mount points. For example, if, in the Project Setup menu, the Project Home path is `/mnt/md0/SCENES` and the Renders Full Home path is `/sanserver:/RENDERS`, located on a central storage system, you must define a mount point for each path.

Mount the storage on a render node:

- 1 Determine the IP address of the computer that manages the storage. If you are using direct attached storage only, this is the IP address of the Lustre Master Station. You must also determine the IP address of your centralized file server if you are also using a server area network (SAN) or network attached storage (NAS).
- 2 Create a directory for the mount point. For example, type: `mkdir /mnt/md0`. The directory name must be the same for the mount point on each node. It should not be in `/root`, but in `/mnt`.
- 3 Change the permissions on this directory to allow read/write access.
- 4 Set up the mount point to access the Lustre system storage or the central system storage. You must use the mount point as defined in the Lustre Project Setup menu. For example, if the Project Home is `/mnt/md0/SCENES`, and the Linux mount point directory is named `/mnt/md0`, the mount command will be: `mount <hostname> :/mnt/md0/SCENES /mnt/md0`

To add an entry in fstab:

- 1 Open `/etc/fstab` in a text editor and then add a line for your new mount point: `<IP address>:<exported filesystem> /<mount point> nfs rw,bg,hard,intr,nfsvers=3 0 0`
- 2 Save and close the file and restart the render node.

Install and Configure Burn for Lustre on render nodes

- Install Burn for Lustre.
- Add the IP address of the machine where Backburner Manager is installed to the `manager.host` file on each render node.
- Start the Backburner Server on each render node.
- License Burn for Lustre.

Install Burn for Lustre on render nodes

When you install Burn for Lustre, the necessary Backburner components are also installed on the render node.

Perform the following procedure to install the Burn for Lustre software on each node.

- 1 Open a terminal to the render node, and log in as root.
- 2 If you are installing from a USB key, insert the key navigate to the Lustre install directory.
- 3 If you are installing from a downloaded tar file, unpack the tar file by typing: `tar zxvf <file_name>.tar.gz`. The installation package is unpacked into a new directory. Navigate to the new directory.
- 4 Start the Burn installation script: `./INSTALL_LUSTRE_BURN`. The Burn for Lustre package is installed. For instructions on obtaining and installing a Burn license, see [Licensing Burn for Lustre](#) (page 49).

Licensing Burn for Lustre

Burn uses a floating license scheme where licenses are centralized on a server. The license server provides a license to all registered machines. The license server machine can be any of the render nodes on the rendering network.

- 1 Open a terminal on the license server and log in as root. Run: `dlhostid`. A message appears that includes a line indicating the `dlhostid` of the machine. For example: `The Discreet host ID of this machine is "DLHOST01=25231AEF83AD9D5E9B2FA270DF4F20B1"`
- 2 Write down the `dlhostid` (including the "DLHOST01=" part).
- 3 [Request license codes](#) (page 81).
- 4 If you are installing from a USB key, insert the key navigate to the Lustre install directory. Or if you are installing from a tar file, unpack the tar file.
- 5 Run the license server installation script: `./INSTALL_LICSERV`. The license server is installed.
- 6 Once you have received the license codes, go to the `/usr/discreet/licserv/licenses/` directory, and open the `DL_license.dat` file in a text editor. If the file does not exist, create it by typing: `touch /usr/discreet/licserv/licenses/DL_license.dat`
- 7 In this file, enter all the information submitted to you by the Licensing Department upon your registration. Save and close the file.
- 8 Start the license server: `/etc/init.d/license_server start`. The license server starts/stops automatically when the machine is booted/shut down. You can stop and restart the server manually by running one of the following commands: `/etc/init.d/license_server stop` and then `/etc/init.d/license_server start`

Enable render nodes to contact the license server.

- 1 Log in to each render node as root, open for editing `/usr/local/flexlm/licenses/DL_license.dat`.
- 2 Copy the first three lines of the information submitted to you by the Autodesk Licensing Department upon your registration. For example:

```
SERVER burn-01 DLHOST01=25231AEF83AD9D5E9B2FA270DF4F20B1
VENDOR lustre
USE_SERVER
```

- 3 Save and close the file.

Configure backburner server to detect Backburner Manager

Backburner Server needs to be able to detect the location of Backburner Manager to provide status information concerning the render jobs:

- 1 On the Backburner Manager system, open a terminal and log in as root.
- 2 Determine which IP address the Backburner Manager workstation uses to connect to the network. You can also use the hostname of the machine, if it has been properly defined in the `/etc/hosts` file.
- 3 On each render node, open a terminal and log in as root and open for editing: `cd /usr/discreet/backburner/cfg/manager.host`. Enter the IP address or hostname of the Backburner Manager machine.
- 4 Save and close the file.

Starting Backburner Server

You must start the Backburner Server daemon on each Linux system for the first time. Once it is started the first time, the daemon will start automatically.

- 1 Log in to your Linux system as root. In a terminal, type: `service backburner_server start`. You can verify that Backburner Server is properly configured: `tail -f /usr/discreet/backburner/Network/backburnerServer.log`. The file should contain lines similar to the following: `<INF>: Loading plugin: Lustre renderer plugin`
- 2 Repeat these steps to start Backburner Server on each node in the background rendering network. If you cannot start the Backburner Server, contact Autodesk Customer Support.
- 3 If you need to stop or restart Backburner Server, use `/etc/init.d/backburner_server stop` or `/etc/init.d/backburner_server restart`.

Configure Lustre to detect Backburner Manager

For Backburner Manager to receive render jobs, you must set up the Lustre application to connect to the system on which Backburner Manager is running:

- 1 To configure new projects to use background rendering, you must set the IP address of the Backburner Manager workstation in the `init.config` file. Locate the `Backburner` keyword. In the `HostName` line, set the `string` parameter to the hostname or IP address of the system where Backburner Manager is installed. For example: `<HostName string="172.19.23.161" />`

NOTE Do not forget the quotation marks.

- 2 Save and close the configuration file.

Specify the Background Rendering path in Lustre

Specify the correct mount points in the settings of each project that uses background rendering:

- 1 In the Lustre Setup menu, go to Settings.
- 2 Select your project in the Project drop-down list, and click Edit.
- 3 In the Project Settings menu click Network Rendering, then click Backburner.
- 4 Enter the location of the Project Home, Scans Full Home, Scans Half Home, Renders Full Home and Renders Half Home, as seen from the Linux render nodes. You only need to enter those locations that are defined for the project in the local project setup, located in the Setup > Project menu. The path

required in these fields is comprised of the directory defined for the mount point and the actual storage folder. For example, if the Project Home on the Lustre workstation is `E:\SCENES/mnt/md0/SCENES` in the Project Setup menu, and the mount point for the Lustre workstation storage on the Linux render nodes is `/mnt/Lustre_storage`, the Project Home path to enter in this menu is:

```
/mnt/Lustre_storage/SCENES
```

- 5 Press Enter.
- 6 Verify that you have entered the proper path for each project, look at the project configuration file in `/usr/autodesk/lustre<application_version>/project/<project name>/project.config`. The value in the `burn_project_home` line is sent to the render node. This value must match a mount point on the background render nodes.

For help setting up and managing projects, refer to the Lustre Help.

Configure Lustre BrowseD

About BrowseD

BrowseD is a Lustre file server with its own high level protocols that optimize network transfers. BrowseD provides Lustre workstations, Slave Renderers, and background render nodes with high-speed access to centralized storage, especially if your network uses InfiniBand, without mounting the storage on the render nodes.

Consider using BrowseD if your facility uses a storage area network (SAN), network attached storage (NAS), or a network file server. In these configurations, you install and configure BrowseD on the computer that is connected to the centralized storage and configure all workstations to use the BrowseD server to access Lustre project files.

The recommended Lustre configuration in a digital film or high-resolution workflow is to store the full-resolution images on a SAN, NAS, or file server, and the proxies locally.

You can run BrowseD on the Lustre Master Station to provide render nodes or Slave Renderers high-speed access to local storage for background rendering. However, this is not recommended. The BrowseD process requires extra bandwidth from the storage and adds extra load on the host system's CPU. Real-time functionality cannot be guaranteed with this setup—2K playback speed and video input and output will not function correctly when BrowseD is running on the Lustre workstation, and is serving client requests.

If you are working with StandardFS DPX or Cineon footage coming from the local Wiretap Server, you must mount your storage on the render nodes.

BrowseD startup and configuration:

- 1 The BrowseD service starts automatically after you install it. You can manually start the service using the command line interface. The service starts automatically whenever the computer running BrowseD restarts.

```
/etc/init.d/browseed_ <version> stop Stops BrowseD.
```

```
/etc/init.d/browseed_ <version> start Starts BrowseD.
```

```
/etc/init.d/browseed_ <version> restart Restarts BrowseD if it is not currently running.
```

```
/etc/init.d/browseed_ <version> condrestart Restarts BrowseD if it is already running.
```

- 2 Configure the `init.config` file for all machines that will use the BrowseD server to access centralized storage.

Username The administrative user on the BrowseD server.

Password The password for Username.

Port All computers on the BrowseD network must use the same port to communicate. Set to 1055, the default. If configuring a render node or a workstation running on a GigE network, set this as 1044. If configuring a render node or a workstation running on a high-speed InfiniBand network, set this as 1055.

BrowsedServer Defines the IP address or DNS host name for a specific BrowseD Server.

UseInfi Enable this option if the networking protocol to use with BrowseD is InfiniBand.

- 3 Making BrowseD Directories Accessible from the Lustre Browser: When using centralized storage, make the directories on the BrowseD server accessible to all Lustre workstations through the browser. You do this by configuring the `BrowsedServerGroup` keyword in the `init.config` file. The mapped directory will appear in the Lustre browser.
- 4 You must configure the Lustre Render Backburner paths with the IP address of the BrowseD server to use BrowseD to render with Burn for Lustre. To use BrowseD for rendering with Burn for Lustre:
- 5 In Lustre, open the Setup menu, and select Settings.
 - 1 Select your project from the Project list, and click Edit.
 - 2 Click Project, and set the local project paths. For information on configuring your projects, see the [Autodesk Lustre User Guide](#).
 - 3 Click Network Rendering, and then click the Backburner tab.
 - 4 Enter the IP address of the BrowseD server to each path, or click the ellipsis button to browse to the BrowseD path you want to use. For example, if your Project_Home is set to the `bernice` directory on the mount point `/mnt/StorageMedia` of a station that has an IP address of `172.17.20.146`, the Project_Home path in the Render > Backburner menu should read `172.17.20.146: /mnt/StorageMedia/project/bernice`.
- 6 To use BrowseD with the Slave Renderer:
 - 1 In Lustre, open the Setup menu, and select Settings.
 - 2 Select your project from the Project list, and click Edit.
 - 3 Click Project, and set the local project paths. For information on configuring your projects, see the [Autodesk Lustre User Guide](#).
 - 4 Click Network Rendering, and then click the Slave Render tab. Enter the IP address of the BrowseD server to each path, or click the ellipsis button to browse to the BrowseD path you want to use. For example, if your Project_Home is set to the `bernice` directory on the mount point `/mnt/StorageMedia` of a station that has an IP address of `172.17.20.146`, the Project_Home path in the Render > Backburner menu should read `172.17.20.146:/mnt/StorageMedia/project/bernice`.

Advanced Configuration and Troubleshooting

Event triggers

Overview

You can set up your Visual Effects and Finishing software to execute custom external commands when certain events take place inside the application, for example, when the project or the video preview timing is changed by the user.

The `/usr/discreet/<application_home>/python/hook.py` Python script contains a number of predefined functions that are called by the Visual Effects and Finishing application when certain events occur. These

functions are referred to as “hooks”. You can enter Python commands directly inside these functions, or use the functions to call an external script or executable.

NOTE To be able to run shell commands or external script using the `os.system` Python statement, add the following line in the `hook.py` script, before the definition of the first hook:

```
import os
```

Note the following about hook functions:

- Functions in `hook.py` receive certain predefined parameters from the application. These parameters contain specific details on the event that occurred in the application, for example, the name of the project loaded. Use these details to better customize the response action.
- Functions in `hook.py` are executed in a separate thread from the application. Thus, the application does not wait for a hook function to complete.
- When a function is called by the application, it executes with the privileges of the user account from which the application was launched. Make sure your custom code in the `hook.py` file does not call external scripts or executables that require more elevated user privileges.
- Functions do not return any value to the application that called them. Thus, the application cannot provide any feedback on whether your custom code executed properly or not.
- Actions defined in the `hook.py` file for one version of an application are not automatically used by other versions or applications. To carry your code over to another application or version, you must manually copy your customized `hook.py` file into the `python` subdirectory of that application’s home directory.

This chapter describes the functions available in the hook API. More functions may be added by Autodesk in subsequent versions. Check the latest `hook.py` file after installing new versions of the software, to see any new functions that may have been added.

NOTE You must be logged in as `root` to edit the `hook.py` file.

void projectChanged(string projectName)

This hook is called by the Visual Effects and Finishing application when the project is changed in the Project Management section of the Preferences menu. This hook is not triggered when a project is loaded at application start-up.

This function receives the following parameter from the application.

projectName A string that specifies the name of the loaded project. For example, `My Project`.

Usage example:

```
def projectChanged(projectName):  
    print projectName
```

When the project is changed in the application, this example function outputs the name of the project in the application terminal.

void previewWindowConfigChanged(string description, int width, int height, int bitDepth, string rateString, string synchString)

This hook is called by the Visual Effects and Finishing application when the video preview timing is changed in the software.

This function receives the following parameters from the application. The values of the parameters are extracted by the application from the `VideoPreviewDevice` keyword in the `init.cfg` software initialization configuration file.

- **description** A string that describes the video preview timing, as listed in the `init.cfg` file. For example, `1920x1080@5994i_free`.
- **width** An integer that specifies the width of the preview timing. For example, `1920`.
- **height** An integer that specifies the width of the preview timing. For example, `1080`.
- **bitDepth** An integer that specifies the bit depth of the preview timing. For example, `12`.
- **rateString** A string that specifies the refresh rate of the preview timing. For example, `6000i`.
- **syncString** A string that specifies the sync source of the preview timing. For example, `reesync`.

Usage example:

```
def
  previewWindowConfigChanged(description,width,height,bitDepth,rateString,syncString):

    print os.system("/usr/bin/mySyncChangeScript --sync="+syncString)
```

When the video preview timing is changed in the application, this example function calls a custom external script, `mySyncChangeScript`, and passes the sync value to it. That external script could use the sync value, for example, to trigger a sync change on the sync generator. The output of the external script is displayed in the application terminal.

Visual Effects and Finishing applications Wiretap path translation

About Wiretap Path Translation

The Wiretap path translation service converts the syntax of a file path referred by a Wiretap server running on a Linux workstation to a format that the client can understand. This service allows Wiretap clients, who might be running Windows or Mac OS X, to understand file paths that are referenced by an Autodesk Creative Finishing product to clips that have been soft-imported from a NAS, SAN, or any remote mounted storage.

For example, Lustre running on a Windows workstation can work with an Autodesk Visual Effects and Finishing product's soft-imported clip on a SAN or NAS.

For the Windows workstation, the syntax of the path to the media files may resemble:

```
N:\myclips\clip1\frame1.dpx
```

On a Linux workstation, the path to the same media files may resemble: `/CXFS1/myclips/clip1/frame1.dpx`

When the Wiretap client tries to locate this clip, the path translation service converts the syntax of the path so that the client can access the soft-imported clip from a remote storage device it has mounted locally.

Configuring the Path Translation Database

The path translation service uses a translation database. The path translation database contains a set of rules that specify how to convert (or “translate”) a file path referred to by a source workstation (the host running the Wiretap server) into a path that the client can understand. These rules are contained in the `/usr/discreet/wiretap/cfg/sw_wiretap_path_translation_db.xml` file on each Autodesk Creative Finishing workstation. This configuration file is read by the Wiretap server.

The translation database is an XML file that must be updated by hand as the network configurations change. The file should be created and maintained by a system administrator who is aware of the configuration of the network.

To configure the path translation database:

- 1 Consider your network environment and how you can organize hosts to simplify the path translation process.
Do you have many hosts accessing the same storage device? Do they mount the central storage using the same syntax and drive names? Are these hosts running the same application? Are they running the same operating system?
- 2 Open `/usr/discreet/wiretap/cfg/sw_wiretap_path_translation_db.xml` in a text editor.
- 3 Create the rules for translating the file path. The rules are applied in the order that they appear in the file. The first rule is applied. If that rule matches, then the translation is complete. If the rule doesn't match, the next rule is examined.
If your network includes groups of hosts that mount their storage using the same syntax, you may be able to use a group or platform rule to translate the syntax for all similar hosts. If, however, all of your hosts use different syntax and mount points for their storage, you will have to create a host/path rule for each source and destination host.
See the sections that follow for a description of the possible path translation rules you can create.
- 4 Save the file.
- 5 You do not have to restart Stone and Wire to apply changes in the path translation database. The Wiretap server periodically checks the timestamp of the file and, if there are changes, it updates itself accordingly. However, if you want to test the results of your changes, you can restart Stone and Wire manually. To restart Stone and Wire manually, type: `/etc/init.d/stone+wire restart` Stone and Wire restarts.

About XML Syntax

You must respect XML syntax standards when editing the path translation file. For example, the values assigned to attributes must not include ampersands (&) or left angle brackets (<) in their literal form. All instances of these characters, except left and right angle brackets used to delimit XML tags, should be replaced with `&` and `<`, respectively. Do not concern yourself with the white-space between attributes; it is ignored when the file is parsed.

Creating a Host/Path Rule for Host-to-Host Translation

Create a host/path rule to translate the path syntax used by the source workstation (the workstation running the Wiretap server) to the path syntax used by the destination workstation.

The syntax of the host/path rule is as follows: `<map src_host="<src_host>" src_path="<src_path>" dst_host="<dst_host>" dst_path="<dst_path>" ">`

Enter a value for each attribute. See the following table for a description of the valid values for each attribute.

Attribute	Description
src_host	The host name or IP address of the workstation to which a client connects. If you create a group rule for source workstations, the value of this attribute can be the group name.
src_path	The path to the remote storage as referred to by the source host (or hosts, when mapping a group).

Attribute	Description
dst_host	The host name or IP address of the client that mounts the same storage referred to by the source, using a different syntax. If you create a group rule for destination workstations, the value of this attribute can be the group name.
dst_path	The path to the remote storage as referred to by the destination host (or hosts, when mapping a group).

Example:

```
<map src_host="flame1" src_path="/CXFS1"
  dst_host="windows1" dst_path="N:\">
```

Setting up a host-to-host translation for every machine in your network is inefficient. If you have many hosts that will use the same path translation rule, you can set up a group or platform rule. See the sections that follow.

Creating a Rule for Groups of Hosts

Use the group rule when you have several workstations that will use the same rule to convert the path syntax of a single source or group of source workstations. This rule is effective in larger installations where numerous workstations mount storage using identical syntax.

All hosts in a group rule must mount directories using exactly the same syntax.

NOTE Group names must be unique and must not conflict with host names or platform names.

The syntax of the group rule is as follows:

```
<group name="<group_name>" os="<os_name>">
  [<host name="<host_name>">]
</group>
```

You must enter a value for each attribute. See the following table for a description of the valid values for each attribute.

Attribute	Description
group name	Identifies the name of the group. Create a group name of your choosing. Each group name must be unique. Use the value of this attribute in a host-to-host rule to map all members of the group to the same storage mount point.
host name	Identifies the name of a host that is in the group.
os	This attribute is optional. It allows you to link all of the hosts in the group to a single operating system. You can then use a platform rule to map all the hosts in a group to the same mount point. The os attribute is restricted to the following values: Linux, Windows NT, and MacOSX.

Example:

```
<group name="Autodesk_Stations" os="Linux">
```

```

    <host_name="flame1">
    <host_name="inferno1">
    <host_name="smoke1">
</group>
<group name="Lustre_Windows_Stations" os="WindowsNT">
    <host_name="lustre1">
    <host_name="lustre2">
    <host_name="lustre3">
</group>

```

After you have defined a group, you can use the host/path rule to map all the hosts in the group to the same path by using the group name for either the `src_host` or `dst_host` attribute value.

In the following example, all hosts in the `Lustre_Windows_Stations` group map the `Autodesk_Stations` storage path to `N:\myclips`, a mount point they all share:

```

<map src_host="Autodesk_Stations" src_path="/CXFS1/myclips"
    dst_host="Lustre_Windows_Stations" dst_path="N:\myclips">

```

You can also use a platform rule to map all the hosts in a group to the same destination path based on the operating system they all run.

In the following example, the platform rule is used to map all Windows users to `N:\` when translating `/usr/` from a Linux system:

```

<map src_os="Linux" src_path="/usr/"
    dst_os="WindowsNT" dst_path="N:\">

```

Creating a Rule for a Group of Hosts Running on the Same Operating System

The platform rule is similar to the group rule in that the rule is applied to all hosts running the same operating system.

All hosts running the same operating system must mount directories using exactly the same syntax. For example, all Windows workstations must mount the NAS on the `N:\` mount point to use the same path translation rule for the NAS.

NOTE Platform names must be unique and must not conflict with host names or group names.

The syntax of the platform rule is as follows:

```

<map src_os="<src_os>" src_path="<src_path> "
    dst_os="<dst_os>" dst_path="<dst_path>">

```

You must enter a value for each attribute. See the following table for a description of the valid values for each attribute.

NOTE The `src_os` and `dst_os` attributes are restricted to the following values: Linux, WindowsNT, and MacOSX.

Attribute	Description
<code>src_os</code>	The operating system of the workstation to which a client connects

Attribute	Description
src_path	The path to the remote storage as referred to by all hosts using the same operating system
dst_os	The operating system of the client that mounts the same storage referred to by the source
dst_path	The path to the remote storage as referred to by the hosts running the destination operating system

Example:

```
<map src_os="Linux"      src_path="/usr/"
      dst_os="WindowsNT" dst_path="N:\">
```

Sharing the Database Among Multiple Wiretap Servers

When you have many Wiretap servers, you must update the path translation database on each server every time there is a change to the network configuration.

To simplify the management of the database, you can create a symbolic link from the path translation file on all hosts running the Wiretap server to one machine that contains the actual file.

Testing the Path Translation Database

After the path translation database is updated by the system administrator, it must be tested. The `sw_wiretap_translate_path` command verifies that a path requested from the Wiretap server on the local machine can be translated into a path that can be understood by the client machine.

The syntax of the command is the following:

```
sw_wiretap_translate_path[-h <host name>] [-p <path>] [-f <file containing paths>] [-H
<destination host name>] [-O <destination OS>]
```

The command options are described in the following table.

Option	Description
-h	Specifies the host name or IP address of the Wiretap server. The default is <code>localhost</code> .
-p	Specifies the path on the Wiretap server host to translate.
-f	Specifies the file containing the paths on the remote host to translate to the path on the local host, delimited by new lines.
-H	Specifies the destination host name. The default is <code>localhost</code> .
-O	Specifies the destination operating system (Linux, Windows NT, Mac OSX).

NOTE Either `-p` or `-f` must be specified, but not both.

To test the path translation database:

- 1 Change to the Stone and Wire tools directory: `cd /usr/discreet/sw/tools`
- 2 Type the `sw_wiretap_translate_path` command with the appropriate options.

For example, suppose you have a SAN mounted as “/CXFS1” on your Flame workstation (flame1). A Windows PC (windows1) will mount the SAN as “N:\”. The path translation file contains the following entries:

```
<map src_host="flame1" src_path="/CXFS1"
    dst_host="windows1" dst_path="N:\">
```

From the workstation, open a new terminal and type:

```
./sw_wiretap_translate_path -h flame1 -p /CXFS1/myclips -H windows1
```

The result should return the value of the destination path, which represents how “windows1” will mount “/CXFS1” on “flame1”.

```
N:\myclips
```

NOTE If the result returns the value of the source path, the translation database entry was not successfully applied.

Visual Effects and Finishing Standard Filesystem Maintenance

Performance of a filesystem depends on a number of factors, including the following:

- Hardware and operating system
- Fragmentation. Fragmentation is a phenomenon that consists of a file being broken into several blocks, which are stored in non-contiguous sectors on the disk. This decreases filesystem performance, as it takes a lot of time to retrieve and reassemble the scattered pieces of the file before making it usable.
- I/O sizes
- Filesystem and tuning parameters. Standard filesystems tuned for a specific file format can under-perform for other file formats. For instance, a filesystem optimized and capable of delivering real-time performance for 2K 10-bit DPX files delivers lower real-time performance for SD media.
- Free space available and fill rate
- Locality of reference. This principle states that files used in the context of the same application or procedure (for example, consecutive frames in a timeline) should be as close to each other as possible on the disk, in order to minimize seek times between two successive I/O operations. Locality of reference is good in the case of frames written on an empty partition, and decreases as the filesystem fills up and becomes fragmented.
- Concurrent I/O and metadata operations
- Number of I/O threads and processes in use at a given time. Multi-threaded I/O is the use of multiple parallel I/O threads within a disk read process. This method allows the RAID array to optimize the fetching of multiple blocks of data.

Control fragmentation

Filesystem fragmentation is directly related to the amount of mixing and interleaving of blocks of data of different sizes, and is aggravated by multiple I/O clients concurrently writing data to the partition.

Although some standard filesystems (such as XFS) feature advanced space allocation techniques, which make them more resistant to fragmentation, this problem still occurs as the filesystem's free space decreases and frequent file deletion, creation, and resizing operations take place.

Here are some steps you can take to counter filesystem fragmentation:

- Avoid storing mixed resolution projects on a standard filesystem partition. Workflows where many different media formats are used simultaneously tend to become fragmented more rapidly than if a single resolution media format is used consistently.

- Separate your media and metadata partitions. The worst contributors to fragmentation on a media storage device are the files used to store metadata. These include clip libraries, setup files, and databases. It is recommended to keep all metadata on a separate partition.

NOTE You may keep the audio data on the same partition as the frames. This has no significant impact on performance or filesystem fragmentation.

- Use the defragmentation tools provided by your filesystem vendor on a regular basis.

NOTE Defragmentation tools are often unable to realign files based on their locality of reference. Therefore, it is advisable to reinitialise your partition periodically, by archiving your projects, formatting the partition, and restoring the archives.

Limit concurrent usage

A filesystem performs best when only one process is accessing it. In some workflows this is not always possible, which leads to a temporary decrease in performance. Furthermore, mixing read and write operations can have a detrimental effect on throughput.

Autodesk strongly advises against allowing any third-party application to access the storage at the same time as your Visual Effects and Finishing applications. The I/O operations from such applications may drastically decrease the bandwidth available to your Visual Effects and Finishing application and also contribute to filesystem fragmentation.

NOTE Unlike direct attached storage (DAS), a Storage Area Network (SAN) is typically designed to handle concurrent usage and will perform better under such situations.

Configure bandwidth reservation

Stone and Wire provides a mechanism to reserve storage bandwidth for Visual Effects and Finishing applications and tools that use a local mount point to a standard filesystem (DAS or SAN) as their media storage. This ensures that the local application gets the bandwidth it requires and that real-time playback on the local system is not jeopardized by requests from concurrent processes, including access from remote hosts.

NOTE Bandwidth reservation policies apply only to I/O requests from Visual Effects and Finishing applications and tools. They cannot protect your storage bandwidth from I/O requests coming from third-party processes or user interactions. It is your responsibility to avoid using third-party tools with the frame storage. See [Limit concurrent usage](#) (page 60).

Bandwidth requests are managed by the Stone and Wire Bandwidth Manager, based on the parameters in the `/usr/discreet/sw/cfg/sw_bwmgr.cfg` configuration file.

When Stone and Wire starts up, the Bandwidth Manager automatically creates a `[Device]` section in the configuration file for each partition declared in `/usr/discreet/sw/cfg/stone+wire.cfg`. The Manager sets total theoretical read and write bandwidth values for each device, as well as the default reservation values.

Adding a New Device to the Bandwidth Manager

Use the instructions in this section if you want additional standard filesystem mount points (such as mount points to SAN volumes or to any other standard filesystem volume that is not present in the `/usr/discreet/sw/cfg/stone+wire.cfg` file) to be managed by the Bandwidth Manager. You may also use these instructions as guidelines for understanding and editing the settings in the predefined `[Device]` sections of the configuration file.

Each partition managed by the Bandwidth Manager must have a [Device] section associated to it in the `/usr/discreet/sw/cfg/sw_bwmgr.cfg` file.

To add a new device to the Bandwidth Manager:

- 1 Open the `/usr/discreet/sw/cfg/sw_bwmgr.cfg` file in a text editor.
- 2 Add a new [Device] section at the end of the file, using the following syntax:

TIP You may also copy and paste an existing [Device] section, and change its parameters.

```
[Device<device_number>]
Path<path_number>=<device_mount_point>
TotalAvailableReadBandwidth=<total_read_bandwidth>
TotalAvailableWriteBandwidth=<total_write_bandwidth>
DefaultReservation=<default_reservation>
```

where:

- `<device_number>` represents the number of the [Device] section, starting at 0. For example:
`[Device0]`
 - `Path` specifies the partition's mount point. Since a partition can have several paths, `<path_number>` represents the number of the current path, starting at 0 for each device. For example:
`Path0=/mnt/XYZ`
`Path1=/usr/local/ABC`
 - `TotalAvailableReadBandwidth` represents the estimated total reading bandwidth of the device, in megabytes per second. Based on this value, and on the amount of active reservations for the current partition, the Bandwidth Manager decides whether new reservation requests are granted or rejected. For example:
`TotalAvailableReadBandwidth=200`
 - `TotalAvailableWriteBandwidth` represents the maximum estimated writing bandwidth of the device, in megabytes per second. If this parameter is not specified, its value defaults to the value of `TotalAvailableReadBandwidth`. For example:
`TotalAvailableWriteBandwidth=150`
- NOTE** The total bandwidth parameters are estimates of the theoretical maximum bandwidth of the partition. The actual bandwidth is affected by several factors, including multiple applications trying to concurrently read or write to it. The Bandwidth Manager continuously measures partition performance and dynamically maintains the actual total available bandwidth for each partition.
- `DefaultReservation` sets the bandwidth assigned by default to applications that are denied a reservation request or that don't have individual or group reservations associated to them. The value is expressed in megabytes per second. For example:
`DefaultReservation=10`

- 3 After setting up all the partition parameters described above, you can start defining reservations for the current partition.
- 4 Save and close the configuration file, and restart Stone and Wire by typing:
`/etc/init.d/stone+wire restart`


Reserving Bandwidth for an Application

You may define a bandwidth reservation for each Visual Effects and Finishing application or tool that uses a certain partition. By default, a group reservation for the major Visual Effects and Finishing applications (Inferno, Flame, Flint, Smoke, and Backdraft Conform) running locally is already defined.

Remote applications are identified by application-specific tokens, workstation host name, as well as the user name under which the application is executed. You can use any or all of these parameters to create detailed bandwidth reservation rules.


You can also customize the way bandwidth is redistributed in low bandwidth situations. In low bandwidth situations, instead of reducing all reservations proportionally towards zero, the bandwidth manager first reduces the bandwidth of each process towards the low-bandwidth value specified for that process, and attempts to maintain that minimum for as long as possible. If the device bandwidth keeps degrading, then the bandwidth manager starts reducing bandwidth proportionally from the low-bandwidth values towards zero.

In the following example, low-bandwidth values are not specified in the Bandwidth Manager configuration file. The diagram illustrates the way the Bandwidth Manager redistributes device bandwidth in such a case, as the total available bandwidth decreases from 800 MB/s to 420 MB/s and then to 320 MB/s. Note how bandwidth is proportionally reduced for all processes.



In the following example, low-bandwidth values are configured for each process (300 MB/s for Flame, 100 MB/s for Flare, 10 MB/s for Wiretap and 10 MB/s for Wire). The diagram illustrates the way the Bandwidth Manger redistributes device bandwidth as the total available bandwidth decreases from 800 MB/s to 420 MB/s and then to 320 MB/s.

Note how the Bandwidth Manager keeps the bandwidth for each application at the low bandwidth watermark. When total available device bandwidth falls under the sum of the low bandwidth watermarks, the Bandwidth Manager decreases bandwidth for all processes proportionally towards zero.



Perform the steps in the procedure below to set up an optimal bandwidth reservation for the local application, as well as for remote workstations, based on your system configuration.

To set up bandwidth reservation:

- 1 Open a terminal and log in as root.
- 2 Open the `/usr/discreet/sw/cfg/sw_bwmgr.cfg` file in a text editor.
- 3 Locate the `[Device]` section that corresponds to the standard filesystem partition (by default `[Device0]`), and uncomment it if it is commented out.
- 4 Add a `Reservation` line for each local or remote application, using the following syntax:

```
Reservation<number>=<application_name> [<user_name>] [@<workstation_hostname>]
<reading_bandwidth> [( <low_reading_bandwidth>)]
[<writing_bandwidth>] [ <low_writing_bandwidth>)]
```

where:

- <n> is the ID of the reservation, starting at 1 for each device.
- <application_name> represents the name of the application that needs the reserved bandwidth. This parameter can take one of the following predefined values: “inferno”, “flame”, “flint”, “flare”, “smoke”, “imcopy”, “stonifiseTool”, “publishTool”, “S+W Server”, “IFFFS Wiretap Server”, or “IFFFSTool” for other Visual Effects and Finishing command-line tools. The “smoke” token also includes Backdraft Conform.

NOTE Application names must use quotes if they contain spaces.

- <user_name> is the Linux user account under which the remote application is running. This parameter is optional.
- <workstation_hostname> is the hostname of the remote workstation that is accessing the storage. This parameter is optional.
- <reading_bandwidth> represents the minimum reading bandwidth required by the application, expressed in megabytes per second. If more bandwidth is available on the device, the Bandwidth Manager gives the application as much bandwidth as possible.
- <low_reading_bandwidth> represents the minimum value towards which reading bandwidth for this application is reduced in low-bandwidth situations. The Bandwidth Manager will attempt to maintain this minimum bandwidth allocation for as long as possible. If the device bandwidth keeps degrading, then the bandwidth manager starts reducing bandwidth proportionally from the low-bandwidth values towards zero for all applications.
- <writing_bandwidth> represents the minimum writing bandwidth required by the application, expressed in megabytes per second. If more bandwidth is available on the device, the Bandwidth Manager gives the application as much bandwidth as possible. If this parameter is not specified, the Bandwidth Manager automatically calculates a default value, based on <reading_bandwidth> and on the total configured read and write bandwidth values of the device.
- <low_writing_bandwidth> represents the minimum value towards which writing bandwidth for this application is reduced in low-bandwidth situations. The Bandwidth Manager will attempt to give the application this minimum amount of bandwidth for as long as possible. If the device bandwidth keeps degrading, then the bandwidth manager starts reducing bandwidth proportionally from the low-bandwidth values towards zero for all applications.

For example:

```
Reservation1=flame 500 (300)
Reservation2=flare artist@flare1 200 (100)
```

- 5 Save and close the configuration file, and restart Stone and Wire by typing:

```
/etc/init.d/stone+wire restart
```

NOTE Bandwidth for an application is actually reserved only when that application is running. Processes that do not have a bandwidth reservation setting in the `sw_bwmgr.cfg` file fall under the default combined reservation of 10 MB/s.

Defining Groups of Applications

The Stone and Wire Bandwidth Manager allows you to set up bandwidth reservations for custom groups of applications. All the applications in a group will share the same reservation.

By default, the configuration file already contains a definition for the “IFFFS” group, which includes the “inferno”, “flame”, “flint”, and “smoke” application tokens.

To set up bandwidth reservation for a group of applications:

- 1 Open the `/usr/discreet/sw/cfg/sw_bwmgr.cfg` file in a text editor.
- 2 In the `[Groups]` section, add a line for each group of applications you want to define. The syntax of the line is as follows:

```
<group_name>=<application_name> <application_name>
```

where:

- `<group_name>` is the custom name of the group. The group name must not contain spaces and must not be the same as one of the predefined application names.
- Each `<application_name>` represents the name of an application in the group. This parameter can take one of the following predefined values: “inferno”, “flame”, “flint”, “smoke”, “imcopy”, “stonifiseTool”, “publishTool”, “S+W Server”, “IFFFS Wiretap Server”, or “IFFFSTool” for other Visual Effects and Finishing command-line tools.

For example:

```
[Groups]
IFFFSTools=imcopy stonifiseTool publishTool
```

- 3 After you have defined the group, go to the `[Device]` section for the appropriate partition and add a `Reservation` line for the new group. The syntax is the same as when adding a reservation for a single application. For example (using the “IFFFSTools” group from the previous step):

```
Reservation1=IFFFSTools 50 50
```

- 4 Save and close the configuration file, and restart Stone and Wire by typing:

```
/etc/init.d/stone+wire restart
```

For additional information on all the parameters in the Bandwidth Manager configuration file, and for a list of application names you can use when defining reservations, see the comments inside the configuration file.

Monitor partition fill rate

As each disk in a storage array gets filled near capacity, once a certain threshold is reached, its overall read and write performance decreases.

Free up space on your array by archiving and deleting projects. Or reinitialise your partition, making sure that you restore from archive only the projects you need to work on.

For an XFS filesystem on direct-attached storage, keep the partition fill rate below 90% for best performance. The Stone and Wire database may experience corruption if there is too little available space on the disk where clip metadata is stored.

Configure the Stone and Wire database to automatically shut off when there is too little free space:

- 1 As root, open the `/usr/discreet/sw/cfg/sw_dbd.cfg` file in a text editor.
- 2 In the `[Disk Check]` section uncomment the `Enabled` and `Minimum` keywords.
- 3 Set the value of the `Minimum` keyword to the minimum allowed amount of free space on the disk (in Megabytes). If the amount of free space on the disk falls below this threshold, the Stone and Wire database daemon automatically shuts off.
- 4 Optional: uncomment the `Timeout` keyword, and set how often the Stone and Wire daemon should check the available disk space. The default value is 30 seconds.

- 5 Save and close the configuration file.
- 6 Restart Stone and Wire: `/etc/init.d/stone+wire restart`.

Use multi-threaded direct input output

Most filesystems perform best when the I/O is parallelised across multiple threads/processes and sent asynchronously. This allows the device to buffer I/O operations and reorganize requests for optimal performance.

Some applications perform better than others on the same storage device, based on how they perform their I/O. Applications that use single-threaded buffered I/O can be slow.

Your Visual Effects and Finishing applications use multi-threaded asynchronous direct I/O when capturing, playing back, and performing other I/O intensive operations.

Make sure that any third-party application you use in conjunction with your Visual Effects and Finishing application or to handle your media files uses direct I/O. This prevents needless CPU-intensive in-memory data copying, thus maintaining optimum storage performance and freeing resources for your Visual Effects and Finishing application.

Enable media pre-allocation

By default, disk space for video files is allocated on-the-fly as the files are rendered or written. Concurrent I/O operations from multiple clients or background processes can generate frames with poor locality of reference, which leads to clip fragmentation and decreases playback performance.

To prevent this problem, Stone and Wire provides media file pre-allocation mechanisms that reserve the necessary space on the disk before the actual writing of the media files takes place. Pre-allocation reserves as many available contiguous disk sectors as needed for the file sequence to be written.

Depending on the filesystem you are using, you may choose between two media pre-allocation methods.

FALLOCATE The disk space needed for the file is allocated and filled with zeroes. On most filesystems, the performance impact of this operation is as serious as when writing the actual file.

RESERVE The disk space needed for the file is reserved, but not written to. The apparent size of the frame remains zero. This operation is much faster than **FALLOCATE**, but is supported only on certain filesystems, such as XFS.

By default, pre-allocation is turned off in Stone and Wire (the method is set to `NONE`). In order to enable it, you must modify a parameter in the `stone+wire.cfg` configuration file.

NOTE This is not a global setting. Pre-allocation must be individually enabled for each partition.

To enable media pre-allocation for a partition:

- 1 Open the `/usr/discreet/sw/cfg/stone+wire.cfg` file in a text editor.
- 2 Uncomment the `Preallocation` parameter in the section that configures the respective partition (for example `[Partition7]`).
- 3 Change the parameter value from the default `NONE` to the pre-allocation method you want to use. For example: `Preallocation=RESERVE`
- 4 Save the configuration file and restart Stone and Wire.

NOTE If you are using XFS and unwritten extent flagging was disabled when creating the filesystem, the RESERVE pre-allocation method is not available for it. For more information about unwritten extent flagging, consult the Man page for `mkfs.xfs`. Type `man mkfs.xfs` in a terminal.

Test filesystem performance

Each standard filesystem comes with its own set of tools to measure performance.

XFS comes with the `xfs_db` command line tool for troubleshooting various aspects of the filesystem, including fragmentation. For information on using the `xfs_db` tool, consult the man page for `xfs_db`.

Stone and Wire comes with a command line tool to measure filesystem performance called `sw_io_perf_tool`. This tool simulates I/O requests (audio and/or video) that would be issued by the Visual Effects and Finishing application, and reports the throughput and maximum latency of all the requested I/O operations.

Use `sw_io_perf_tool` to measure read and write performance of a volume declared and mounted by Stone and Wire. Several options are available to simulate many file formats and writing/reading patterns. Launching the command will display a list of available options.

To list all options available for `sw_io_perf_tool`, open a terminal and type
`/usr/discreet/sw/tools/sw_io_perf_tool -h`.

To test a standard filesystem partition's performance for a specific resolution, Run
`/usr/discreet/sw/tools/sw_io_perf_tool <-p #> <options>` where `<-p #>` specifies a partition ID (0 to 7), and `<options>` can be a combination of several media format attributes.

```
./sw_io_perf_tool -p7 -v 1600,900
```

The command displays the throughput (in MB/sec) and maximum IO latency detected (in ms), as shown below:

```
Running Stone+Wire IO Performance Test (Partition 7)
```

```
Video: 1 stream(s) of HDTV @ 10bit
```

```
Proxy: None
```

```
Audio: None
```

```
This test will take approximately 11 seconds
```

```
*** Read Test ***
```

```
Bandwidth: 275.456 MB/s
```

```
Frame Rate: 39.351 Frames/s
```

You may also use `sw_io_perf_tool` with custom resolutions. For example, to test read performances for frames with a width of 1600 pixels and a height of 900 pixels: `Max Latency: 335.284 ms`

Notes:

- The command can be run while other operations are ongoing to simulate how the application will behave under these conditions.
- The media used by `sw_io_perf_tool` is allocated and freed automatically.
- `sw_io_perf_tool` will not reserve any bandwidth on the storage device.
- Partitions/Volumes must be properly configured in the Stone and Wire configuration file.

Visual Effects and Finishing Troubleshooting

Killing an Application

If a Visual Effects and Finishing application hangs, but you can still interact with the operating system, you can terminate it using either a custom command alias for your application, or the generic Linux `kill` command in conjunction with the application's process ID.

WARNING These procedures should only be used if the application is hung or otherwise not responding. Killing processes can cause system problems.

When you install a Visual Effects and Finishing application, a special command alias is defined that can be used to kill all Visual Effects and Finishing application processes at once:

- 1 Open a terminal, and type one of the following commands:
 - `kf` to kill Flame, Flint or Flare
 - `ks` to kill Smoke
 - `ki` to kill Inferno
 - `kc` to kill Backdraft Conform
- 2 To verify that all application-related processes have been killed, type: `\ps -fu <application account name>`. This displays information about all active processes, including the application with which each process is associated, and a process ID number. If all application processes have been successfully killed, this command should return no results.

WARNING Never restart an application if there are still processes left over from a previous session. If some processes persist after using the kill alias, perform the following procedure to kill them individually.

Use the `kill` Linux command to kill processes individually, using their process ID number:

- 1 Press Alt+F3 to send the application to the background.
- 2 Open a terminal and type: `ps -fu <application account name>`. This command displays information about all active processes, including the application with which each process is associated, and a process ID number.
- 3 Determine the process ID number of the Visual Effects and Finishing application.
- 4 At the command line, type: `kill <process number>` where `<process number>` is the process number you determined in the previous step. This command terminates the Visual Effects and Finishing process that is currently executing. There may be more than one Visual Effects and Finishing process running at any time. For example, there may be one process per CPU, plus some additional processes to manage the framestore. Kill each of these processes.
- 5 Restart the application.

Troubleshoot the filesystem

This section describes some common filesystem problems and steps you can take to solve them. When troubleshooting storage or wire issues, start by verifying that Stone and Wire processes are running properly, and by checking the log files.

Verifying that Stone and Wire Processes Are Running

There are five processes that must be running for Stone and Wire to work:

- `sw_serverd`
- `sw_probed`
- `sw_dbd`
- `sw_bwmgr`
- `ifffsWiretapServer`. This process is started by the `sw_wiretapd` script.

To verify that Stone and Wire processes are running:

- 1 Run: `ps -ef | grep sw_`. This command should return several lines similar to the following:

```
root 9447 1 0 Feb25 0:00:35 /usr/discreet/sw/sw_probed --daemonize
root 9455 1 0 Feb25 0:00:35 /usr/discreet/sw/sw_dbd -d
root 9463 1 0 Feb25 0:08:11 /usr/discreet/sw/sw_bwmgr -d
root 9472 1 0 Feb25 0:08:11 /usr/discreet/sw/sw_serverd -d
root 9513 1 0 Feb25 0:00:00 /bin/sh /usr/discreet/sw/sw_wiretapd
```

- 2 To verify that the `ifffsWiretapServer` process is running, type: `ps -ef | grep Wire`. The command should return a line similar to the following:

```
root 9522 9513 0 Feb25 00:03:29 /usr/discreet/wiretap/2010/ifffsWiretapServer -c
/usr/discreet/wiretap/cfg/wiretapd.cfg
```

- 3 If Stone and Wire processes are not running, start them.

Starting Stone and Wire Processes

All Stone and Wire processes should start when you restart your system. If they are not running, check that they are configured to run and then restart Stone and Wire. To start Stone and Wire processes, you must modify a text file listing all processes.

- 1 Use a text editor to open the Stone and Wire configuration file `/etc/sysconfig/stone+wire`.

The file should include the following:

```
dl_stone=on
dl_wire=on
dl_wiretap=on
dl_vic=on
```

- 2 If any of the processes you need are set to “off”, change them to “on”. Restart Stone and Wire. To restart Stone and Wire, type `/usr/discreet/sw/sw_restart`. If any Stone and Wire process is still not running, check the error logs.

Checking Stone and Wire Error Logs

Every Stone and Wire process has a log file. You can find these log files at `/usr/discreet/sw/log`.

The current log file is named `<process>.log`, where `<process>` is the name of the Stone and Wire process or daemon. The next time Stone and Wire creates a log file for the process, it renames the previous log file by adding a number to the file name.

For example, the `sw_served` process log file is named `sw_served.log`. The next time the process is launched, the first log file is renamed to `sw_served.log.1`. Each time Stone and Wire creates a new log file, the old log file is renamed by incrementing the number that appears at the end of the filename.

You can also use the terminal to list the most recent version of the log files. Type the following in a terminal:

```
ls -altr
```

Enabling Verbose Stone and Wire Driver Error Reporting

Use the `swr_set_debug` utility to enable verbose reporting of Stone and Wire driver errors.

When you turn on error reporting, Stone and Wire outputs error messages in the following locations:

- `/var/log/messages`
- Console window

NOTE You must disable verbose error reporting after you have finished troubleshooting. The messages generated by this process can fill up the system drive.

To enable verbose Stone and Wire driver error reporting:

- 1 Run: `/usr/discreet/sw/tools/swr_set_debug <level>` where `<level>` is the number representing the verbosity of the report. Choose from one of the following levels:
 - **1** - Turns on error reporting. This is the default level.
 - **2** - Provides more detailed reporting. All operations and errors are logged.
 - **0** - Disables error reporting.
- 2 Check the terminal and the log files for error information.

Troubleshooting Missing Frames

If the application cannot find some of the frames referenced by your clips, the following error message is displayed when the application starts:

- **In the terminal:** `WARNING: <nnnn> <volume_type> frames are missing on this volume; see the application installation guide for corrective measures. Where <nnnn> represents the number of missing frames, and <volume_type> can be stonefs or standardfs.`
- **In the application start-up screen:** `VOLUMEMGT: WARNING: Some frames not found; see the application installation guide for corrective measures.`

By default, the message only appears in the terminal. To have it appear also in the start-up screen, open the `~/.cshrc` file in a text editor and add the following on a new line: `setenv DL_IC_NOTIFY_MISSING 1`. Then, log out of the KDE desktop, and log back in. The error message appears in one or both of the following situations:

- Some clips in your libraries or desktops reference missing frames on a remote volume.
- Some clips in your libraries or desktops reference local frames that no longer exist.

To identify the type of issue:

- 1 Open a terminal and type: `vic -v <volume_name> -s remote -s lost` where `<volume_name>` is the name of the affected volume, for example, `stonefs` or `standardfs`.

The output of this command should be similar to the following:

```
Checking libraries for remote and lost frames...
/usr/discreet/clip/stonefs/My_Project1/editing.000.desk has none
/usr/discreet/clip/stonefs/My_Project1/Default.000.clib references 30 missing frames.
/usr/discreet/clip/stonefs/My_Project2/editing.000.desk has none
/usr/discreet/clip/stonefs/My_Project2/from_caplan.000.clib references 70 remote frames
```

- 2 Depending on the result of the previous command, do one of the following:
 - For clips with frames listed as missing, contact your local Autodesk support representative for help on resolving this issue.
 - For clips with frames listed as remote, perform the following procedure.

To recover remote frames:

- 1 Archive all clip libraries that are reported as containing remote frames. In the previous example, the fourth library (from_caplan) contains remote frames.
- 2 Rename the libraries you just archived, such as by adding the remote suffix to the library name.
- 3 Restore the clip libraries from the archive. All the frames referenced by the clips should now be stored locally.
- 4 Delete the original libraries.

Verifying that the Disks are Available to the Operating System

If your disks or logical drives are not visible in the output of the `sw_df` command, verify that they are visible to the operating system.

- 1 Run the command that will provide you with a list of all the controllers or disks on your system. `cat /proc/scsi/scsi`. A listing of all controllers and hard disks attached to your system appears.
- 2 If any controllers or hard disks in your configuration are not visible, check the connections and cabling to your storage device, and ensure that your storage hardware is running.

If you must reconnect or tighten connections to your storage hardware, it is a good idea to rescan your hardware to ensure that the operating system is able to communicate with it.

- 1 Unload the Stone and Wire driver: `/etc/init.d/stone+wire unload`
- 2 Unload and reload the fibre channel drivers. Type one of the following sequence of commands, according to your fiber channel card model.

- If you have a QLogic-based card, type:

```
rmmod qla2300
rmmod qla2400
rmmod qla2xxx
rmmod qla2xxx_conf
modprobe qla2xxx
modprobe qla2300
modprobe qla2400
```

- If you have an ATTO-based card, type one of the following sequences of commands:

For HP Z800 workstations, type:

```
rmmod celerity8fc
modprobe celerity8fc
```

For older HP workstations, type:

```
rmmod celerityfc
modprobe celerityfc
```

NOTE Depending on the storage you are running, your system might not use all of the drivers listed. If your system does not use a driver listed, the commands to unload or reload the drivers will fail. You can ignore these failures. They just indicate that the driver is not required by your system.

- 3 Reload the Stone and Wire driver: `/etc/init.d/stone+wire reload`. Your filesystem should now be mounted.

Verifying Disk Usage

You cannot write to storage partitions that are completely full.

To see if the filesystem is full, run `/usr/discreet/sw/sw_df`.

If your filesystem is full or almost full, consider deleting or archiving older material from clip libraries to free disk space. VTR archiving in Visual Effects and Finishing applications requires free space on the framestore to generate slates, test patterns, headers, and audio encoding. You can also clear the Undo Buffer to free up space.

Clearing the Undo Buffer

The Undo Buffer takes up space on your filesystem. Consider clearing it if your filesystem is full or almost full to free disk space.

WARNING Clearing the Undo Buffer cancels any undoable operation. Consider what operations might be stored in the Undo Buffer before you clear it.

To clear the Undo Buffer:

- 1 Open your Visual Effects and Finishing application, and go into the Preferences menu.
- 2 Click Clear Buffer in the Undo section.

Solving a Framestore ID Mismatch

The framestore ID must match the filesystem ID for the workstation to mount your filesystem.

If the IDs do not match, you must change your framestore ID to ensure that all hosts on your Wire network have a unique ID.

- 1 Open a terminal as root, and run `cat /usr/discreet/sw/cfg/sw_storage.cfg`. The output of the command includes the `[Framestore]` section, which lists the framestore ID.
- 2 Open `/usr/discreet/sw/cfg/sw_framestore_map` in a text editor. The `sw_framestore_map` file includes a section similar to the following:

```
[FRAMESTORES]
FRAMESTORE=london HADDR=172.16.129.32 ID=32
```

The Framestore ID value in the `sw_framestore_map` file must match the Filesystem ID value listed in the `sw_storage.cfg` file.

NOTE If the framestore ID in the `sw_storage.cfg` file is not the same as the one in the `sw_framestore_map` file, the ID listed in `sw_storage.cfg` has priority.

- 3 Change the Framestore ID value in the `sw_framestore_map` file to match the Filesystem ID listed in the `sw_storage.cfg` file.

NOTE The last sequence of numbers in the IP address defined by the `HADDR` keyword in the `sw_framestore_map` file does not have to match the Framestore ID. These values are often the same by default, but it is not a requirement for Stone and Wire operation.

- 4 Save and close the file.
- 5 Restart Stone and Wire: `/usr/discreet/sw/sw_restart`

- 6 If you continue to get error messages, contact Customer Support.

Solving a Partition ID Conflict

Each partition must have a different partition ID. If there are multiple partitions using the same ID, you will receive an error message.

To solve a Partition ID conflict:

- 1 Note the ID listed in the error message.
- 2 Stop Stone and Wire: `/etc/init.d/stone+wire stop`
- 3 Open the `/usr/discreet/sw/cfg/stone+wire.cfg` file in a text editor.
- 4 Find the [PARTITION#] section, where # is the ID listed in the error message. Change the # for an unused ID. Partition identifiers go from 0 to 7, giving a total of 8 possible partitions. Stone filesystem and standard filesystem partitions share the same pool of identifiers, but only standard filesystem ones are listed in `stone+wire.cfg`. It is recommended to start numbering standard filesystem partitions at 7, decrementing with each additional standard filesystem partition.
- 5 Save and close the `stone+wire.cfg` file.
- 6 Start Stone and Wire. Type: `/etc/init.d/stone+wire start`. A message appears in the terminal indicating that Stone and Wire is restarted.
- 7 Start your Visual Effects and Finishing application to verify that the partition ID conflict is resolved.

Removing Invalid Frame References from the Stone and Wire Database

Stone and Wire has a database inconsistency detection mechanism. If you change the framestore ID of a Stone filesystem or standard filesystem partition in your system without reformatting the Stone and Wire database, the inconsistency check mechanism detects invalid frame references in the database, and prevents the Stone and Wire database server (`sw_dbd`) from starting.

If you receive an error message that `sw_dbd` could not be started when you start a Visual Effects and Finishing application or Stone and Wire, perform the following steps to determine if the problem is related to invalid frame references, and to remove those references from the Stone and Wire database.

To remove invalid frame references from the Stone and Wire database:

- 1 Open a terminal and log in as root.
- 2 Search the `/usr/discreet/sw/log/sw_dbd.log` file for the text “Bad framestore ID” e.g.: `cat /usr/discreet/sw/log/sw_dbd.log | grep -i bad/ framestore/ id`
 - If the search returns one or more lines, the problem is related to invalid frame references. Perform the remaining steps in this procedure to clean up the Stone and Wire database.
 - If the search returns no results, the problem is not related to invalid frame references. Read the `/usr/discreet/sw/log/sw_dbd.log` file to determine the cause of the problem.
- 3 Stop Stone and Wire: `/etc/init.d/stone+wire stop`
- 4 Clean up the Stone and Wire database by typing: `/usr/discreet/sw/sw_dbd -C`. A message is displayed in the terminal asking you to confirm the deletion of invalid entries from the database.
- 5 Type `Y` to confirm the operation. Invalid entries are removed from the Stone and Wire database.
- 6 Restart Stone and Wire: `/etc/init.d/stone+wire start`

Troubleshoot Wire network problems

This section describes some common Wire problems and steps you can take to solve them.

Testing Wire Using `sw_framestore_dump`

Use `sw_framestore_dump` to identify the Wire hosts that are reachable from one workstation. You should do this if you are having any problems with Wire or after you configure Wire for the first time. You must have at least two Wire hosts set up to test Wire connectivity:

- 1 Run: `/usr/discreet/sw/tools/sw_framestore_dump`. A list of all the available framestores appears in the terminal. Review the output and verify that all the framestores on your Wire network appear in the list. You can also verify that each framestore on your network is using a unique framestore ID.
 - If You are unable to see other Wire hosts on your network, check:
 - The filesystem and networking configurations in `sw_framestore_map`.
 - The `sw_probed.cfg` port number.
 - You may also have problems with your Wire network. Use `sw_ping` to test that you can connect to other Wire hosts.
 - If you see only some of the Wire hosts (as opposed to all or none), check that each framestore has a unique Framestore ID.
- 2 Repeat this procedure on each Wire host.

Using ping to Test Network Communication

Try to ping your local host from a client machine. If this works, ping all other machines that should be accessible through Wire. :

- 1 Type the following command: `ping <host_name>`.
- 2 If ping fails, try using the machine's IP address (for example, 172.16.100.23) instead of its hostname. If this is successful, verify how the machine resolves host names on the network. You should set the order of host name resolution to first look at the local setup file, then validate on the network. The `/etc/nsswitch.conf` file should include a "hosts" line that has the name validation process in the following order: `hosts: files nis dns`

Using `sw_ping` to Test Wire Communication

If you are unable to connect to a Wire host, use `sw_ping` to determine if you can connect to the other Wire host over the Wire network:

- 1 Run `sw_ping`: `/usr/discreet/sw/sw_ping -framestore <framestore_name> -r -w -size <packetsize> -loop <n>` where `<host_name>` is the name of the host to ping, `<count>` is the number of times to execute this test, `<packetsize>` is the size of the read/write buffer (in bytes). Results similar to the following are reported to the terminal:

```
Using Stone+Wire API 2009 [ProductSW_1166] (64 bits).
Stone+Wire TCP port set to 7000.
Ping successful to host tunisia (186)
Buffer size = 9043968
Sending data to tunisia, filesystem id 186
MB/sec: 111.460127
Ping successful.
Buffer size = 9043968
Sending data to tunisia, filesystem id 186
```

```

MB/sec: 111.509057
Ping successful.
Buffer size = 9043968
Sending data to tunisia, filesystem id 186
MB/sec: 111.500577
Ping successful.
Buffer size = 9043968
Sending data to tunisia, filesystem id 186
MB/sec: 111.516597
Ping successful.
Minimum MB/sec: 111.460127
Maximum MB/sec: 111.516597
Average MB/sec: 111.496590

```

- If an `sw_ping` from machine 1 to machine 2 is successful and you still cannot access the remote framestore through your Visual Effects and Finishing application, check that the remote system has exported its directories so that the local system has read/write access. See [Verifying Remote Clip Library Access and Permissions](#)
- If an `sw_ping` from machine 1 to machine 2 is unsuccessful, validate the `sw_framestore_map` file on both machines and check that `sw_probed` and `sw_serverd` are running on both machines.

Using `sw_ping` to Test Network Performance

Use the `sw_ping` command to test network performance. For more significant results, run the test 100 times.:

- 1 Start `sw_ping`: `/usr/discreet/sw/sw_ping -framestore <framestore_name> -r -w -size <packetsize> -loop <n>`

Option:	Description:
<code>-framestore <framestore_name></code>	Is the name of the framestore to ping.
<code>-r</code>	Reads a buffer from the remote framestore.
<code>-w</code>	Writes a buffer to the remote framestore (non-destructive).
<code>-size <packetsize></code>	Reads/writes a buffer of packetsize bytes.
<code>-loop <n></code>	Executes this test n times.

The following is an example for an NTSC (720 x 486 x 3) frame format, sent 100 times in read and write mode to framestore `my_framestore`:

```
/usr/discreet/sw/sw_ping -framestore my_framestore -read -write -size 1049760 -l 100
```

Verifying Remote Clip Library Access and Permissions

Verify that the local host has write permissions to the clip library directory of the remote host:

- 1 Try to access the clip library directory of the remote host. Type: `cd /hosts/<remote_machine>/usr/discreet/clip`
- 2 If an error message appears on the client machine, check to see that the required network daemons are on. If you have access to `/hosts/<remote_machine >/usr/discreet/clip`, check you have root

privileges to read and write to this directory: `touch /hosts/ <remote_machine>
/usr/discreet/clip/deleteme .`

- 3 If error messages appear, verify the permissions in `/etc/exports` on the server machine. Read and write permission must be given to the hosts. The appropriate line should look similar to the following:
`/usr/discreet/clip *(rw, sync, no_root_squash`

Verifying that NFS and Automounting Daemons are Running

These daemons are installed and running by default. Use the following procedures to verify that these daemons are running. Network File System (NFS) and automounting daemons (AMD on Linux) must be running for Wire transfers to work. To start NFS and AMD:

- 1 Run: `chkconfig --list | grep nfs` and then `chkconfig --list | grep amd`
- 2 Verify that NFS and AMD are set to “on” for run levels 3,4, and 5.
- 3 If NFS or AMD is off on any of those run levels, run: `chkconfig nfs on` and `chkconfig amd on`
- 4 Restart your network: `/etc/init.d/network restart`. You might also consider rebooting your workstation.

Checking the Status of Network Interfaces

If you continue to have problems with your network, you should verify that your network interfaces are up and running:

- 1 Run: `ifconfig`
 - If your network interface is up and running, an “up” appears in the broadcast report for the interface. The report includes a line similar to the following: `UP BROADCAST RUNNING
MULTICASTMTU:1500Metric:1`
 - If your network interface is not up and running, check the connections on your network card. A green light appears when there is a good connection between your network card and its destination. If you must reconnect cables on Linux, you must restart the network interface. To restart a network interface, run: `ifconfig <interface_name> up`

Viewing InfiniBand Statistics

This section describes how to view port statistics for the workstation and InfiniBand switch, and how to obtain InfiniBand log files. For more information about configuring, monitoring, and managing the InfiniBand switch and network, see the documentation provided by the manufacturer of your InfiniBand switch.

Viewing InfiniBand Information for a Workstation

The InfiniBand driver installed on your workstation provides two commands that output statistics and information about InfiniBand ports. Log in as root, and type the following commands in a terminal:

```
cat /proc/iba/ <driver_id> /1/port <x> /stats  
cat /proc/iba/ <driver_id>/1/port <x> / info
```

where `<x>` is the port number on the device, and `<driver_id>` is the HCA driver ID for your device, for example: `mt25218`. A report appears in the terminal for each command. To find out the HCA (Host Channel Adapter) driver ID, type `ls /proc/iba/ | grep mt`. The driver number, beginning with `mt`, will be returned.

Generating InfiniBand Log Files for Support

The script `/sbin/iba_capture` gathers information about your system and InfiniBand configurations that can be useful for troubleshooting InfiniBand issues. This script outputs a GZIP file.

As root, open a terminal on your workstation and run `/sbin/iba_capture <path and name of output gzip file>`. A GZIP file is generated that includes a number of log files from your system. The help for this command incorrectly indicates that the output is a TAR file, when it is in fact a GZIP file.

Viewing Port Statistics for a Mellanox IS5030 Switch

The Mellanox IS5030 switch is the recommended switch model for QDR InfiniBand networks. It comes with a Web-based management console that you can use to configure the switch and monitor the InfiniBand network.

When you suspect that performance is not optimal on your QDR switch, it is a good idea to view statistics for each connected port, to determine if they are running normally. The nominal speed for DDR InfiniBand connections is 5 Gbps, and the nominal speed for QDR connections is 10 Gbps.

To access the management console, your browser must be on the same subnet as the InfiniBand switch, which has a default IP address of 10.10.10.252.

To view switch port statistics:

- 1 Open a Web browser and enter the IP address of the InfiniBand switch. The default IP for switches shipped from Autodesk is 10.10.10.252. The login page appears.
- 2 Enter your user name and password. The defaults are `admin / admin`. The Summary page appears.
- 3 Click Ports in the menu at the top. The Ports page appears, displaying an overview of the switch. Connected ports are displayed in green.
- 4 Click a port to view information and statistics on it. If you have ports with DDR connections that appear to be running at SDR speed (2.5 Gbps instead of 5 Gbps), unplug the cable and then plug it back in. The connection should run at normal DDR speed afterwards. This issue occurs because of a bug in the switch firmware.

Viewing Port Statistics for an InfiniCon InfinIO 9024 Switch

The InfinIO 9024 InfiniBand switch is the recommended switch model for DDR InfiniBand networks. The InfinIO switch comes with a Web-based application, the InfiniView Device Manager, that you can use to configure, monitor, and manage the InfiniBand network. When you suspect that performance is not optimal on your QDR switch, it is a good idea to view statistics for each connected port, to determine if they are running normally. The nominal speed for DDR InfiniBand connections is 5 Gbps. To access the Device Manager, your browser must be on the same subnet as the InfiniBand switch, which has a default IP address of 10.10.10.252.

To view switch port statistics:

- 1 Open a Web browser and enter the IP address of the InfiniBand switch, usually 10.10.10.252. The main page of the Device Manager appears.
- 2 Click Port Stats in the menu on the left. The IB Port Statistics page appears. The page displays the status of each port on the switch.
- 3 Optional: Click Refresh to update the data on the page.

Licensing and Starting

7

Node-locked licensing

Node-locked licenses are locked to one computer.

- All customers receive a temporary node-locked license when they purchase their software.
- Subscription customers receive a temporary node-locked license when they receive their extension software.

All customers must request permanent licenses:

- Subscription customers have the choice of requesting permanent node-locked or network licenses.
- Customers not on subscription are entitled to only node-locked licenses.

1 If you are installing the application for the first time:

- 1** Request temporary license codes. For emergencies, you can acquire an immediate temporary emergency license using the emergency license generator at <http://melicensing.autodesk.com/templicensing/>. A 4-day license code is e-mailed to the address you provide.
 - 1** Start the software you want to license. The License Wizard starts automatically when you launch an unlicensed application. You can also start the Wizard manually, from the Autodesk section in the KDE system menu.
 - 2** In the License Wizard select “No, request a License Code”, and click Next.
 - 3** Select Node-Locked License and click Next.
 - 4** Enter your contact information in the license code request form. The host ID and system serial number of your workstation are automatically filled in. Click Next.
 - 5** Log in to support portal: <https://edge.autodesk.com/LicenseAssignment>. If you do not have an account on the support portal, email me.support@autodesk.com to make a request.
 - 6** Follow the links to get the a node-locked license. Your license codes are emailed to you.
- 2** Install the temporary or permanent license codes.
 - 1** Double-click the application icon on the Linux Desktop to start the application. After a few moments, the License Wizard appears.
 - 2** Select “Yes, install the License Code”, and click Next.

- 3 Select Node-Locked License and click Next.
 - 4 Paste the license codes you received from Autodesk in the next window. Be very careful to paste the license codes exactly as you have received them. Any incorrect or missing characters will prevent your application from recognizing the license.
 - 5 Click Install to install the license code. Your application is now licensed. License codes are unique for each workstation, and are tied to the system hardware. If you replace major system components in your system, you must request and install a new license code for your software.
 - 6 To add additional license codes you receive from the Licensing Department (such as the DI license), run the License Wizard from the Autodesk section in the KDE system menu.
 - 7 When you receive permanent license codes, install those codes using the same procedure.
- 3 [Start the software](#) (page 85) and accept the license agreement.

Network licensing

With a network license, you can deploy your Creative Finishing software on any supported workstation in your facility. A license server manages the number and types of licenses you have available. Network licenses are available to customers with subscription to Inferno, Flame, Flame Premium, Smoke Advanced, and Lustre Master Station. Network licenses are also available for Autodesk Flare, Autodesk Burn, and Autodesk Burn for Lustre. All other customers must use node-locked licenses.

Choose one of the following network license server models:

- Single license server. The license server daemon is installed and runs on a sole server. A single license file represents the total number of licenses available on the server. This single license server can reside on a remote workstation or on the same workstation as your Creative Finishing application. Characteristics of the Single License Server Model :
 - Of the two license server models, this configuration requires the least maintenance.
 - Because all license management takes place on a single server, you have just one point of administration and one point of failure. On the other hand, if the single license server fails, the Autodesk product cannot run until the server is back online.

- Redundant license server model. In the redundant license server model, you use three Linux servers to authenticate a single license file. One server acts as the master, while the other two provide backup if the master server fails. With this configuration, licenses continue to be monitored and issued as long as at least two servers are still functional. The license file on each server is the same. You install the license server daemon on each server. Characteristics of the Redundant License Server Model:
 - If one of the three servers fails, all licenses that are managed in the server pool are still available.
 - If more than one server fails, no licenses are available.
 - All three servers must reside on the same subnet and have reliable network communications. The redundant server pool does not provide network fault tolerance.
 - If you reboot one of the license servers or restart one license server daemon you must also reboot/restart all three license servers at the same time so the license servers are synchronized.
 - If one of the three servers is replaced, you must request a new license key from Autodesk Creative Finishing Support

Workflow for Configuring a License Server

- 1 Select a license server. You can install a license server on your Creative Finishing workstation or any other machine on your network. To see if you already have a license server running on a machine, open a Terminal, login as root, and type: `chkconfig --list | grep lic`. For redundant license servers, select three machines. All must be set to the same date and time.
- 2 [Install the license server software](#) (page 81) on all machines that will be servers.
- 3 [Get the unique host ID of a license server](#) (page 81) for all servers.
- 4 [Request license codes](#) (page 81).
- 5 Create the workstation license. [License the workstation with a local or single remote license server](#) (page 82). [Configure the workstation to use a set of redundant license servers](#) (page 83) if using the redundant server model.
- 6 If using one or more remote license servers: [Create a license file for a remote license server](#) (page 82).
- 7 [Start the license server](#) (page 84)(s).
- 8 To avoid conflicts with other applications or license servers in your facility, you may need to [Change the default port used by the license server](#) (page 84)(s).

Install the license server software

The license server is a Linux daemon that provides concurrent licenses to Creative Finishing Softwareworkstations.

You can install the license server on a Creative Finishing Softwareworkstation in your network, or on any 64-bit system running Red Hat Enterprise Linux Desktop 5.3 with Workstation Option (64-bit), Red Hat Enterprise Linux WS 4, Update 3 (64-bit), CentOS 5.3, or CentOS 4.3.

For a redundant network license server configuration, you must install the license server software on all three workstations selected as license servers.

To install the license server, as root, run on the `./INSTALL_LICSERV` from the software installation directory.

Get the unique host ID of a license server

To get license numbers, you need the host ID of the license server, or servers, in the case of a redundant network configuration.

To get the unique host ID of the license serve, in a terminal run `/usr/local/bin/dlhostid`. The number should look something like `25231AEF83AD9D5E9B2FA270DF4F20B1`

Request license codes

Request licensing codes from the Autodesk M&E Edge support portal: <https://edge.autodesk.com/LicenseAssignment>.

If you do not have an account on the support portal, email me.support@autodesk.com to make a request.

To request license codes:

- 1 Make sure you have the unique ID of your single license server (local or remote) or a set of three unique IDs for each server of a redundant network license server configuration ready before you request your license codes. See [Get the unique host ID of a license server](#) (page 81).

- 2 Log in to support portal: <https://edge.autodesk.com/LicenseAssignment>.
- 3 Follow the links to set up your license server assets and select the software that you want to access using a network license.
Your license codes are emailed to you.

License the workstation with a local or single remote license server

Create license files

- For a redundant network license server configuration, you must create the license files for the workstation and license server manually. See [Create a license file for a remote license server](#) (page 82) and [Configure the workstation to use a set of redundant license servers](#) (page 83).
- To license the workstation with a local license or a single remote license server, use the license wizard:
 - 1 Start the application. If no license file is found, the License Wizard opens.
 - 2 Select “Yes, install the license code” and click Next.
 - 3 In the License Installation dialog, do one of the following and click next:
 - If the license server is on the same machine as your Creative Finishing application, select Local Network License.
 - If the license server is on another machine, select Remote Network License.
 - 4 Open the license file you received from Autodesk in a text editor or display its contents in the shell.
 - 5 In the Local Network License Installation dialog, select the license text and copy and paste the text into the field. The text should resemble the following:

```
SERVER servername DLHOST03=2BB094B4722BDBECFA218BD3A6BB8F83
VENDOR discreet_1
USE_SERVER
PACKAGE smokeadvanced_x86_64_r discreet_1 \
```

Ensure the servername (in the first row) is replaced with the actual name of the computer your license server is on.

- 6 Click Install.
- 7 [Start the license server](#) (page 84).

The license wizard creates the following license files:

- Workstation license: `/usr/local/flexlm/licenses/DL_licenseNetwork.dat`
- For a local license server, it creates the license server license:
`/usr/discreet/licserv/licenses/DL_license.dat`.
- For a remote license server, you must create the license file for the license server manually. See [Create a license file for a remote license server](#) (page 82).

Create a license file for a remote license server

After you receive your license codes, edit the `/usr/discreet/licserv/licenses/DL_license.dat` license file on the license server (or all three servers, if you are configuring redundant license servers), and enter the license codes you received from Autodesk.

The license information contains keywords that identify the license server, as well as the license strings for products or features.

Keyword	Description
SERVER	Specifies the hostname of the license server from which to obtain the license, followed by its dlhostid.
VENDOR	Specifies the daemon that is serving the license.
USE_SERVER	Indicates whether the system should obtain its license from a license server.
FEATURE	License strings for the software and feature entitlements.

To create the license server file on a license server:

- 1 Log in as root to the license server.
- 2 Navigate to the licenses directory by typing: `cd /usr/discreet/licserv/licenses`
- 3 If the file `DL_license.dat` does not exist in the directory, create it: `touch DL_license.dat`
- 4 Open the file `DL_license.dat` in a text editor.
- 5 Enter the information provided by Autodesk in this file. It should be similar to the following shortened example:

NOTE If you are setting up redundant a network license, there are three SERVER strings.

```
SERVER server DLHOST01=886C2B75E8E57E4B03D784C3A2100AC0
VENDOR discreet_l
USE_SERVER
FEATURE flameEDU_x86_64_b_2011_1 2011.99 02-oct-2010 \
```

NOTE If you are working in an existing `DL_license.dat` file, make sure you do not accidentally alter any of the existing text in the file when entering the Creative Finishing Software license information.

- 6 Save and close the file. This file sets up the network licenses available for distribution by the license server to the Creative Finishing workstations on your network.

Configure the workstation to use a set of redundant license servers

To configure the workstation to use a set of redundant license servers, edit as root `/usr/local/flexlm/licenses/DL_licenseNetwork.dat`. Create it if it doesn't exist. Add the SERVER, VENDOR and USE_SERVER strings provided by Autodesk. It should be similar to the following shortened example:

```
SERVER host01 DLHOST01=886C2B75E8E57E4B03D784C3A2100AC0 27005
SERVER host02 DLHOST01=886C2B75E8E57E4B03D784C3A2100AC1 27005
SERVER host03 DLHOST01=886C2B75E8E57E4B03D784C3A2100AC2 27005
VENDOR discreet_l
USE_SERVER
INCREMENT flameEDU_x86_64_b_2011_1 2011.99 02-oct-2010 \
```

This file sets up the floating licenses available for distribution by the license server to the Creative Finishing workstations on your network.

Change the default port used by the license server

To avoid conflicts with other applications or license servers in your facility, you may need to change the default port setting used by the license server, or license servers in the case of a redundant configuration. This requires a minor change to the `DL_license.dat` file on your license server(s), as well as on every Creative Finishing Softwareworkstation.

To change the default port used by a license server:

- 1 Log in as root to the license server and open `/usr/discreet/licserv/licenses/DL_license.dat` for editing.
- 2 Find the `SERVER` line. By default, no port number is specified at the end of the `SERVER` line for a single license server and the license server uses default port number in the range of 27000-27009. By default, redundant license servers are set to port 27005.
- 3 Enter a different port at the end of the `SERVER` line. The license server can use a default port number in the range of 27000-27009. For example: `SERVER server`
`DLHOST01=886C2B75E8E57E4B03D784C3A2100AC0 62222`
- 4 Save and close the file, then verify that the new port settings are correct:

```
/etc/init.d/license_server stop
/etc/init.d/license_server start
cat /usr/discreet/licserv/log/license_server.log
```

For redundant license servers, reboot each server in close sequence to properly restart the license system.

- 5 Look for messages similar to the following examples in the output, and verify that the port numbers are what you requested, e.g.:

```
15:08:49 (lmgrd) lmgrd tcp-port 62222
15:08:49 (lmgrd) Starting vendor daemons ...
15:08:49 (lmgrd) Using vendor daemon port 12344 specified in license file
15:08:49 (lmgrd) Started discreet_1 (internet tcp_port 12344 pid 5013)
```
- 6 Repeat with `/usr/local/flexlm/licenses/DL_license.dat` for each workstation, using the same port as the one you set for the license server.

Start the license server

NOTE For redundant license servers, reboot each server in close sequence to properly restart the license system.

To start the license server:

- 1 Type the following in a shell: `/etc/init.d/license_server start`

WARNING The license server cannot start unless the license is entered correctly in `DL_license.dat`. Check the `boot.log` file to make sure the license server is started and working properly.

NOTE The license server starts/stops automatically when the machine is booted/shut down. You can stop and restart the server manually by running one of the following commands: `/etc/init.d/license_server stop` or `/etc/init.d/license_server restart`

Start the software

You must license Inferno, Flame Premium, Flame, Flint, Smoke Advanced, Smoke and Backdraft Conform, and where applicable, Lustre, before you can start.

If you use the pen and tablet while the application is starting, the tablet will fail to initialise. Press `Shift+T+Insert` to initialise the tablet if it does not function during a work session. To start the software for the first time:

- 1 Log into your workstation and do one of the following:
 - **WARNING** The `-v` option deletes all material on the framestore. Use this option only if you have no material that you want to preserve on the framestore.

If this is the first time you are starting the software on a new system, in a terminal run `/usr/discreet/<product_home>/bin/startApplication -v` to initialize the framestore and start the application. For information on all the command-line options available, see [Visual Effects and Finishing Command Line Start-up Options](#) (page 85).
 - To start the application without initializing the framestore, double-click the application icon on the Desktop.
- 2 The Autodesk Master Systems Software License Agreement appears. To view the License Agreement in another language, click a language tab at the bottom of the License Agreement window.
- 3 Read the text carefully, and click I Accept if you agree to all the terms in the License Agreement.
- 4 Use the Project Management menu to create a project and a user. For information on creating projects and users, see the application Help.
- 5 Click Start or press `Enter`.
- 6 When prompted to confirm the framestore initialisation, answer “yes”. After a few moments, the following message appears: `Startup complete`. You are ready to start working in the application. If you see a splash screen rather than the application interface, click anywhere on the screen.

Starting Lustre for the First Time

- To start Lustre, log into the application’s Linux user account and double-click the Lustre icon on the KDE desktop. Lustre starts. The license you installed determines the type of station you are using (Lustre Master Station or Lustre Station).
- When the Lustre UI appears, you can begin working. For a new project, this includes:
 - Creating the project and setting its parameters (such as logarithmic or linear colour space). For information on creating projects and users, see the “Project Management” topic in the Lustre Help. To open the Help, press `Shift+F1`.
 - Specifying the location of footage.

Visual Effects and Finishing Command Line Start-up Options

Overview

There are a number of options you can use with the application start-up command under special circumstances.

NOTE More than one start-up option can be specified on the command line. These start-up options are case-sensitive.

To start the application using a command line option, open a terminal and type:

```
<application_name>-<option>
```

For example, **smoke -h**.

To see a list of all start-up options, use the **-h** option.

NOTE Running the application start-up command with the **-h** option does not launch the application.

Media Storage-Related Start-up Options

The following table lists the start-up options that enable you to initialize the storage, delete clips from a project, or troubleshoot missing frames.

Type:	To:
-v	Initialise the media storage volume (all projects). WARNING When you use the -v option, all data on the storage is deleted, with no possibility of recovery. NOTE Initialising the storage also removes all audio clips from the project.
-H	Specify the remote host name.
-V	Specify the media storage name.
-i	Remove all clips from the project.
-r	Initialise only the Desktop or Edit Desk for the selected project.
--red-frames	Activate repair mode. This option enables <i>vic</i> to search the media library and replace missing frame ids with a LOST frame label. In the library, the names of impacted clips appear red. When these clips are loaded into the record timeline, the names of the affected segments also appear red. Lost frames are displayed with the LOST label over them. The flag set on clips is not permanent, so resaving the clip will remove the red from the names.

Project Management-Related Start-up Options

The following table lists the start-up options that enable you to specify a custom application configuration file, as well as choose the default user and project the application starts with.

Type:	To:
-c<filename.cfg>	Use a configuration file other than the default, where <i><filename.cfg></i> is the name of the file. The specified file supersedes the default configuration file. If you do not use this option, the application looks first for <i>init.cfg</i> , and if that is not found, it looks for <i>flame.cfg</i> , <i>inferno.cfg</i> , <i>smoke.cfg</i> , or <i>conform.cfg</i> .
-J<project>	Define the project that appears in the Project Management menu when you start the application, where <i><project></i> is the name of the project you want to use. If you start the application with both the -J and -U options, the application bypasses the Project Management menu.

Type:	To:
-U <user>	Define the user that appears in the Project Management menu when you start the application, where <user> is the name of the user you want to use. If you start the application with both the -J and -U options, the application bypasses the Project Management menu.

Miscellaneous Start-up Options

The following table lists start-up options that enable you to install new fonts, handling log files, managing the amount of memory and number of CPU cores used by the application.

Type:	To:
-F	Force the application to install new fonts that you added to the <i>/usr/lib/DPS/outline/base</i> directory (and the <i>/usr/lib/DPS/AFM</i> directory, if you have also installed the corresponding font metric file). See the application help.
-B	Prevent the broadcast monitor from blanking when an image viewer or Player is not displayed. The broadcast monitor then displays the screen content that is displayed in the corresponding region. NOTE This option is hardware dependent and may not work on all hardware configurations.
-L	Disable writing to the <i>.log</i> file. The application keeps a log of your actions in a text file (<i>.log</i>) in your application home directory. This file is usually cleared and recreated every time the application starts. (Note that the log file is useful when reporting problems to Autodesk Customer Support.)
-I	Append the current session to the existing log file. (See description of the log file above.)
-M	Define the amount of memory assigned to frame buffering. This option overrides the default settings, or if it has been uncommented, the value set by the <i>MemoryApplication</i> keyword in the <i>init.cfg</i> file.
-p <n>	On multi-processor machines, start the application using less than the full array of processors, where <n> is the number of processors to be used. Normally, the application uses all available processors.
-f <filename>	Use a custom menu file, where <filename> is the name of the menu file. For information on custom menus, see the application help.
-J <project> -U <user> -b <batch_setup> -N	Launch the application and specify a Batch setup to render. The application launches, renders the setup, and exits. NOTE Use the -N option only if you want to use Burn for rendering. WARNING To avoid corrupting the render, which uses the graphics hardware, you should not use the graphics console of the system (such as command shells) while the render is taking place.
-Z	Starts the application, and skips the integrity check.

Index

A

- AMD
 - starting daemon 77
- application
 - entering license codes 79
 - installing 35
 - licensing 79
 - starting 85
 - uninstalling 37
- audio
 - Discreet Native Audio 10
 - hardware components 10
 - wiring workflow 10
- audio converter 11
- Audiodevice keyword 35
- AutoFS
 - starting daemon 77

B

- Backburner Manager
 - configuring 46
 - detecting in Lustre 50
 - detecting with Backburner Server 50
 - installing 46
 - overview 45
- Backburner Monitor
 - installing 46
 - overview 45
- Backburner Server
 - detecting Backburner Manager 50
 - overview 45
 - starting 50
- BackburnerManagerGroup keyword 43
- BackburnerManagerHostname keyword 43
- BackburnerManagerPriority keyword 43
- background rendering
 - components 44
 - configuring in Lustre 50
 - overview 44
 - setting up render nodes 47
 - setup, workflow 44
 - storage, mounting 48
 - with BrowseD 51
- background Wire transfers 42
- broadcast monitor, wiring 10
- BrowseD
 - accessing directories from Lustre 51
 - configuring, workflow 51

- overview 51
 - remote rendering 51
 - starting and stopping 51
 - using 51
- browsed.config
 - configuring 51
 - Burn for Lustre
 - hardware and software requirements 47
 - installing and configuring 48
 - license 49
 - overview 46

C

- clip library
 - definition 2
- commands, start-up options 85
- compatibility
 - transferring media from previous versions 41
- configuring
 - Backburner Manager 46
 - browsed.config 51
 - Burn for Lustre 48
- connecting
 - audio 10
 - video I/O components 10
 - video I/O devices 10
 - workstation peripherals (illustration) 9
- CXFS filesystem 39

D

- definition
 - clip library 2
 - managed media 2
 - Managed Media Cache 2
 - unmanaged media 2
- Discreet Kernel Utility version 26
- Discreet Native Audio 10
- disk failure
 - operating system access 72
- disk usage 73
 - clearing undo buffer 73
- DKU version 26
- DVIRamp keyword 35

E

- Effects and Finishing application
 - typical configuration (illustration) 5
- error logs 70
- error reporting
 - errors logs 70
 - verbose 71

F

- filesystem
 - disk usage 73
 - problems 69
- filesystems
 - CXFS 39
 - NFS 39
- floating licenses
 - background rendering 49
- fonts
 - installing 86
- framestore ID
 - mismatch, resolving 73

G

- group rule
 - Wiretap path translation 56

H

- hardware
 - requirements, Burn for Lustre 47
- hardware configuration guidelines 5
- hardware platforms 1
- host/path rule
 - Wiretap path translation 55

I

- InfiniBand
 - log files 77
 - management 77
 - port statistics 77, 78
- init.cfg, configuring 35
- installation
 - overview 6
- installing
 - Backburner Manager 46
 - Backburner Monitor 46
 - Burn for Lustre 48
 - Lustre (Linux) 37
- installing the application 35

K

- keywords 36
 - Audiodevice 35

- BackburnerManagerGroup 43
- BackburnerManagerHostname 43
- BackburnerManagerPriority 43
- DVIRamp 35
- TabletDriver 35
- Video 35
- VTR 35

- Kill process command 69

L

- license
 - background rendering, floating 49
 - Burn for Lustre 49
- license codes
 - editing the license file 79
 - entering application 79
 - requesting 79
- license server
 - background rendering 49
 - creating the license file 82
- licensing
 - creating the license server license file 82
- log files
 - InfiniBand 77
- logs
 - See error logs
- Lucid ADA 88192. See audio converter
- Lustre
 - accessing BrowseD directories 51
 - detecting Backburner Manager 50
 - installing (Linux) 37

M

- managed media
 - defining 2
- Managed Media Cache
 - defining 2
- mount points
 - overview 46
- multi-OS networks
 - Wiretap path translation 54
- multicasting
 - Linux 24, 25

N

- NAS
 - using BrowseD 51
- network interfaces
 - status 77
- network performance
 - sw_ping 76
- network tests
 - ping 75
- networking
 - multicasting, Linux 24, 25

NFS 39
starting daemon 77

O

operating system rule
Wiretap path translation 57

P

path translation database
XML syntax 55
ping 75
platforms in this release 1
port statistics, InfiniBand 77, 78
powering on or off, sequence for 29
problems
filesystem 69
Wire 75
processes
displaying, Stone and Wire 70
starting, Stone and Wire 70
processes, stopping 69
products in this release 1
projects
setting up slave rendering 44
proxy generation 42

R

render nodes
setting up 47
rendering
with BrowseD 51
requirements
DKU 26

S

SAN
using BrowseD 51
scanning storage hardware 72
shutdown sequence 29
slave rendering
configuring, workflow 43
setting up Lustre projects 44
software
requirements for Burn for Lustre 47
standard filesystem
testing performance 68
start-up options 85
starting the application 85
status
network interfaces 77
Stone and Wire processes
displaying 70
starting 70

Stone and Wire utilities
swr_set_debug 71
storage
scanning hardware 72
workflow for configuring 26
sw_framestore_dump
testing Wire 75
sw_ping
network performance 76
testing Wire 75
swr_set_debug 71
system
freezes, what to do 69

T

tablet
initialising 85
TabletDriver keyword 35
testing
Wiretap path translation 58
testing Wire
sw_framestore_dump 75
sw_ping 75
troubleshooting
common filesystem problems 69
common Wire problems 75
troubleshooting Stone filesystem
operating system access 72

U

undo buffer
clearing 73
uninstalling the application 37
unmanaged media
defining 2

V

verbose error reporting 71
video I/O, wiring 10
Video keyword 35
video wiring 10
VideoPreviewDevice keyword 35
VTR
connecting 10
VTR keyword 35

W

Wire
clip library access, verifying 76
problems 75
Wire transfers, background 42
Wire troubleshooting
network interface status 77

- Wiretap
 - configuring path translation 54
 - path translation 54
- Wiretap path translation
 - configuring 54
 - group rule 56
 - host/path rule 55
 - operating system rule 57
 - sharing the database 58
 - testing 58
- wiring video I/O and broadcast monitor 10
- workflow
 - configuring BrowseD 51
 - configuring slave rendering 43
 - licensing 79
 - setting up background rendering 44
 - storage configuration 26
- workflows
 - audio wiring workflow 10

X

- X server
 - configuring 35
 - restarting 35
- XML syntax
 - path translation database 55
- xxdiff, using 35