Autodesk®
Burn® 2012

Installation and User Guide

Autodesk®

Autodesk® Flame® 2012, Autodesk® Flame® Premium 2012, Autodesk® Flare™ 2012, Autodesk® Lustre® 2012, Autodesk® Smoke® 2012, Autodesk® Smoke® for Mac OS® X 2012 software

Trademarks

Disclaimer

Published by: Autodesk, Inc.
111 McInnis Parkway
San Rafael, CA 94903, USA
Title: Autodesk Burn 2012 Installation and User Guide
Document Version: 1
Date: July 4, 2011

Contents

Introduction

About Burn

Autodesk®Burn® is a Linux®-based network processing solution for Autodesk® Inferno®, Autodesk® Flame®, Autodesk® Flint®, Autodesk Flare™, Autodesk® Smoke®, and Autodesk® Backdraft® Conform.

Burn allows you to process images in the background using low-cost Linux systems and frees your workstation for more creative tasks. With Burn, facilities can leverage a larger pool of processing power so that complex visual effects requiring intensive processing can be created more quickly.

As part of the Autodesk® Backburner™ background processing network, a network of Burn render nodes can process multiple tasks in the background, dramatically increasing the efficiency of the visual effects and finishing processes.

Burn render nodes equipped with GPU-accelerated graphics cards are capable of processing complex jobs, such as floating point jobs.

By combining a fast network connection such as InfiniBand® technology with a GPU-accelerated graphics card, Burn can deliver extremely fast background processing of timelines, Batch setups, or any other job that supports background processing.

You can also benefit from the fast background processing provided with Burn when using the Sparks® API to port your custom Autodesk Developer Network Sparks plug-ins to the Linux® environment.

Architectural Overview

The components of the Burn architecture include the following:

- An Autodesk Visual Effects and Finishing application (the render client) that sends jobs to the background processing network
- At least one Linux computer that does the processing (the render node)
- A computer that distributes and manages the jobs running on the background processing network (the Backburner Manager)
- The Burn and Backburner software components running on the render nodes on the background processing network
- At least one computer that monitors the jobs running on the background processing network (the Backburner Monitor)

The following illustration shows the workflow between the components on a background processing network.

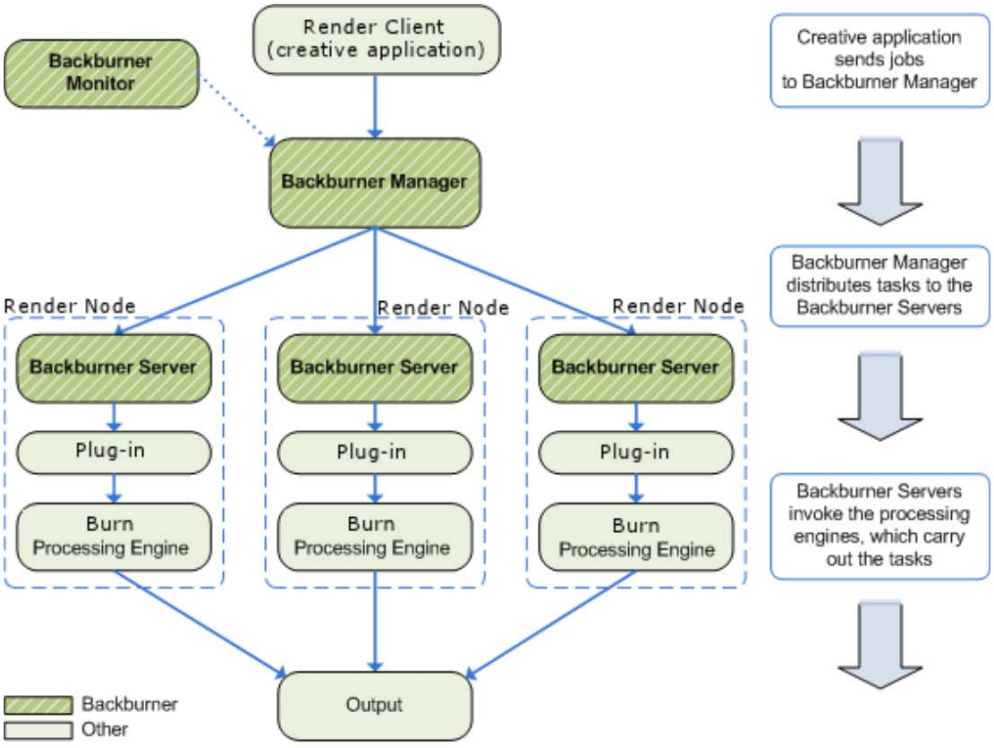

Render Client This is the Autodesk Visual Effects and Finishing application, such as Autodesk® Inferno®, Autodesk® Flame®, Autodesk® Flare®, Autodesk® Flint®, Autodesk® Smoke®, and Autodesk® Backdraft® Conform, running on a Linux workstation, or Autodesk® Smoke® for Mac OS® X running on a Mac OS X workstation. From here, you create and submit processing jobs to be processed by the background processing network. Each setup or clip submitted for processing is called a *job*.

Backburner Manager This is the hub of the background processing network, and it runs on a Windows®, Mac OS X, or Linux workstation. The render client submits jobs to Backburner Manager, which distributes them to the render nodes available on the network, according to the job type.

Backburner Manager runs as a service on Windows, and as a daemon on Linux and Mac OS X. Backburner Manager starts automatically when the system boots, and runs continuously until either the workstation is shut down or the service/daemon is stopped. On Windows, you can also start Backburner Manager manually as an application from the Windows Start menu.

Burn Render Node This is a Linux system on the background processing network that hosts the Burn Processing Engine. Jobs received from a render client via Backburner Manager are assigned to the Processing Engine on the render node via Backburner Server.

Render nodes use common network protocols like TCP/IP and/or Autodesk® Wire® to receive source frames and to return processed frames back to the render client.

A render node's ability to process certain types of jobs depends on its hardware capabilities. Render nodes without GPU-accelerated graphics cards cannot process jobs that require a GPU (such as floating point jobs). They can only process jobs in software mode, using the OSMesa API. Render nodes equipped with GPU-accelerated graphics cards can process both jobs that require a GPU and OSMesa jobs.

Backburner Monitor This is the user interface for the background processing network. Backburner Monitor runs as an application on Windows systems, and as a Web application on Linux and Mac OS X systems. The Web application can be accessed through a Web browser from any computer on your network. Either version allows you to view and control jobs currently being processed. Jobs in the background processing network can be stopped, restarted, reordered, archived, or removed. You can also monitor the overall health of the background processing network and identify any render nodes that are not working.

Backburner Server This is an application that runs on each Linux render node in the background processing network. Backburner Server accepts

commands from Backburner Manager to start and stop the Processing Engine for the assigned processing tasks on the render node.

Burn/Backburner Plug-in This provides the communication link between the Backburner Server and the Burn Processing Engine.

Each Autodesk application uses its own plug-in (such as the Burn/Backburner Plug-in) to communicate with its Processing Engine via Backburner. This architecture allows multiple Autodesk applications to share the same render node for a variety of background processing tasks, such as rendering 3D models or transcoding media between video formats.

There are separate Burn/Backburner plug-ins for GPU-enabled render nodes and for render nodes without GPUs. The Burn installation script detects the presence of a GPU in the render node, and installs the appropriate plug-in.

Processing Engine This is the process that processes frames from jobs submitted from render clients. Specifically, this refers to the Burn application that processes frames from jobs submitted from Inferno, Flame, Flint, Smoke, or Backdraft Conform.

The Processing Engine is installed on each render node. By installing multiple processing engines on a render node, the render node is able to process jobs from different clients.

Wire Autodesk Wire® is a component that enables the high-speed transfer of uncompressed video, film, and audio between Autodesk systems, over industry-standard TCP/IP and InfiniBand networks. Render nodes use Wire to transfer source frames from the render client, and to return the processed frames back again. Wire is only required for applications that use Burn.

Installation Workflows 2

Workflow for Upgrading the Software

Follow this workflow to upgrade your software to a new version, service pack, or extension without reconfiguring your hardware or reinstalling your operating system.

To determine whether you need to install a new version of Linux® for a new software version or hardware platform, see Do I Need to Reinstall Linux? (page 13). If you do need to upgrade Linux, follow the steps in the Workflow for Installing from Scratch or Upgrading Linux (page 6) rather than this workflow.

To upgrade the software:

1. Gather the following materials and documentation for your installation:
   - The *Discreet Kernel Utilities* (DKU) installation package and the software installation package. For major releases, all product software is available on the product USB key. If you are installing a service pack or an extension, download the DKU *tar* file from the link provided in the Release Announcement you received from Autodesk.
   - The *Release Notes* for the software version or service pack you plan to install. The Release Notes contain important procedures that you might need to perform **before** upgrading your software, as well as any system requirements.
2. Make sure the hardware still meets all the system requirements for the new software version or service pack, as specified in Setting up Your Hardware (page 11).
3. Upgrade the Discreet® Kernel Utilities (DKU) to the required version. Refer to the latest Release Notes for the required DKU version and Installing the DKU (page 21) for instructions.

4 Upgrade the Backburner Manager system on your background processing network to the current version. See Installing Backburner Manager (page 28).

5 Upgrade the Visual Effects and Finishing workstations to the same version as the version of Burn you are about to install on the nodes. Each version of Burn is compatible with only one version of Autodesk Visual Effects and Finishing applications. See the *Autodesk Visual Effects and Finishing Installation and Configuration Guide* or the *Autodesk Smoke for Mac OS X Installation and Licensing Guide* for information on upgrading the applications.

6 Install the Burnsoftware on each node. See:
- Installing and Configuring Burn (page 27).

7 License your software. See:
- Network Licensing (page 43)

NOTE Re-licensing is not necessary if you are upgrading to a service pack of the same software version or to a service pack of the same extension.

8 Start using Burn. See Using Burn (page 51).

Workflow for Installing from Scratch or Upgrading Linux

Follow this workflow when you need to completely rebuild your system: from connecting the hardware to installing the operating system, and installing, licensing, and configuring the software components.

Most of this workflow also applies when you need to reinstall or upgrade the operating system without changing your hardware setup. A new distribution of Linux may sometimes be required for a new version or service pack of the software, or for newer hardware platforms. To determine whether you need to upgrade Linux for a new application version or hardware platform, see Do I Need to Reinstall Linux? (page 13).

To install your system from scratch or upgrade Linux:

1 Gather the following materials and documentation for your installation:
- The installation disc for the Linux operating system. If you are setting up a third-party Burn render node, download the required version of CentOS from *www.centos.com*. For information on the version of

Linux required for your hardware platform, see Do I Need to Reinstall Linux? (page 13).

- The *Discreet Kernel Utilities* (DKU) and software installation package. For major releases, the software installation packages are available on the USB key. For service packs and extensions, the software packages are distributed as *tar* files. Links to the *tar* are provided in the Release Announcement you received from Autodesk.
- The *Release Notes* for the software version or service pack you plan to install. .
- The latest version of the *Autodesk Backburner Installation Guide*, and of the *Autodesk Backburner User Guide*. You can download these documents from *www.autodesk.com/backburner-documentation*.

2. Make sure the hardware meets all the system requirements. See Hardware Requirements for Burn Nodes (page 11).
3. Refer to your hardware vendor documentation for information on configuring the hardware and connecting the node to your network.
4. Determine the distribution and version of Linux required for your hardware and for the current version of Burn. See the table in Do I Need to Reinstall Linux? (page 13). If the required Linux distribution for your hardware is CentOS, download the CentOS CD or DVD image from *www.centos.org*.
5. Add the Autodesk kickstart file to the DVD or CD1 of your CentOS distribution. See Preparing your Linux Installation Disc (page 14). This step is not necessary if you are installing the Autodesk distribution of Red Hat® Enterprise Linux on a node purchased from Autodesk.
6. Install the required distribution of Linux. See Linux Installation Workflow (page 15).
7. Perform the tasks in Linux Post-Installation Tasks (page 18) to configure your operating system.
8. Install the required version of the Discreet Kernel Utility (DKU) on each node. Refer to the latest *Release Notes* for the required version and Installing the DKU for installation instructions.
9. Upgrade the Backburner Manager system on your background processing network to the current version. See Installing Backburner Manager (page 28).
10. Upgrade the Visual Effects and Finishing workstations to the same version as the version of Burn you are about to install on the nodes. Each version of Burn is compatible with only one version of Autodesk Visual Effects and Finishing applications. See the *Autodesk Visual Effects and Finishing*

Installation and Configuration Guide or the *Autodesk Smoke for Mac OS X Installation and Configuration Guide* for information on upgrading the applications.

11 Install the Burnsoftware on each node. See:
   - Installing and Configuring Burn (page 27).

12 Configure each node to connect to Backburner Manager. See Connecting the Render Node to Backburner Manager (page 30).

13 Optional: Disable local Stone® and Wire I/O on each node to improve performance. See Disabling Local Stone and Wire I/O on Burn Nodes (page 37).

14 Optional: Group Burn nodes into groups to manage them more efficiently. Consult the *Autodesk Backburner User Guide* for instructions on using Backburner Web Monitor to create node groups.

15 License your software. See:
   - Network Licensing (page 43)

   NOTE Re-licensing is not necessary if you are upgrading to a service pack of the same software version or to a service pack of the same extension.

16 Start using Burn. See Using Burn (page 51).

Workflow for Installing the Smoke for Mac OS X Distribution of Burn

The licensing requirements for the Smoke for Mac OS X distribution of Burn are different from the Linux Visual Effects and Finishing distribution.

NOTE Two distributions of Burn 2012*cannot be installed* on the same node. However, either distribution can process jobs sent from a Mac or Linux product, as long as it is licensed.

To install the Smoke for Mac distribution of Burn:

1 Install and license Smoke for Mac OS X and the Network License Manager. See the *Autodesk Smoke Installation and Licensing Guide*.

2 Write down the Smoke for Mac OS X license server name and host ID.

3 If necessary, set up Burn hardware and install the operating system. See Setting up Your Hardware (page 11) and Installing Red Hat Enterprise Linux or CentOS (page 13).

4 Install and configure Burn.

During installation of the Burn software, enter the license server name and host ID.

This information is recorded in the following file: */var/flexlm/autodesk.lic*

See Installing and Configuring Burn (page 27).

5 Start using Burn. See Using Burn (page 51).

Setting up Your Hardware

3

Overview

You can run the current version of Autodesk Burn on render node hardware purchased from Autodesk and running Red Hat Enterprise Linux, or on third-party nodes running CentOS, as long as they meet the minimum requirements outlined in Hardware Requirements for Burn Nodes (page 11).

NOTE Installing Burn on Visual Effects and Finishing workstations is not supported. Install Burn only on dedicated render nodes.

Check the Red Hat or CentOS Web sites to ensure that the hardware for your render nodes is supported by the operating system.

Hardware Requirements for Burn Nodes

The following table outlines the minimum and recommended hardware requirements for Burn render nodes that were not purchased from Autodesk.

Component	Minimum	Recommended
CPU	Two single-core, or one dual-core 64-bit processor, such as AMD™ Opteron™ or Intel® Xeon®	Two quad-core 64-bit processors, such as the Intel Xeon E5472 CPU at 3.0 GHz
Memory	8 GB or higher	Same amount of memory as the Visual Effects and Finishing workstation

Component	Minimum	Recommended
Hard Disk	120 GB or higher SATA, Ultra-SCSI 320, SAS, or IDE drive. *Note:* The system disk must be a single, physical hard disk drive. It cannot be a logical volume from an array of disks.	
Network card	On-board Gigabit Ethernet adapter	InfiniBand card, if you want to run Burn and Autodesk Incinerator on the same render node
GPU-accelerated Graphics Card	None. *Note:*Burn nodes without a GPU-accelerated graphics card cannot render jobs that require a GPU, such as floating point jobs.	A card from either of the following series: ■ NVIDIA® Quadro® FX 3800 or better ■ NVIDIA® Quadro® 4000 or better

The recommended Burn requirements for Sparks developers are available in the latest *Autodesk Visual Effects and Finishing Configuration Guide for Autodesk Developer Network Sparks Plug-ins*.

NOTE The remaining sections in this chapter apply only to Burn nodes purchased from Autodesk. If you did not purchase your node hardware from Autodesk, contact your vendor for further information on your hardware.

Installing Red Hat Enterprise Linux or CentOS

4

Do I Need to Reinstall Linux?

Burn nodes purchased from Autodesk ship with the correct Autodesk distribution of Red Hat Enterprise Linux already installed and configured. Thus, you should only need to install Linux in the following situations:

- You are setting up a Burn render node that was not purchased from Autodesk. CentOS 5.3 is the required operating system for render nodes not purchased from Autodesk. Make sure your render node hardware is supported by CentOS.

  NOTE If you have older Burn 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.
- Your system experienced an unrecoverable hard disk failure and you had to replace the system disk.

  NOTE If you must replace the hard disk of your license server system, you need to obtain a new license. Contact Customer Support for assistance.

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 HP ProLiant DL140 G3	Custom Autodesk distribution of Red Hat Enterprise Linux WS 4, Update 3
Burn nodes not purchased from Autodesk	■ CentOS 5.3 on new nodes ■ The existing installation of Fedora Core or CentOS 4.6 on already configured nodes

To determine the version of Red Hat Enterprise Linux or CentOS your system is currently running, open a terminal and type:

cat /etc/redhat-release

The Linux version appears. For example:

```
Red Hat Enterprise Linux Client release 5.3 (Tikanga)
```

Preparing your Linux Installation Disc

Before installing Linux, you must add the Autodesk kickstart file to the DVD or first CD of your CentOS distribution. This file is necessary for the Linux installer to install some packages required by Burn.

NOTE If you are reinstalling Red Hat Enterprise Linux on a render node purchased from Autodesk, skip this section. The custom Autodesk DVD of Red Hat Enterprise Linux already contains the Autodesk kickstart file.

This section describes how to create a new DVD or first CD for your CentOS distribution so that it contains the required Autodesk kickstart file.

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

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

2 Insert the DVD or first CD of your Linux distribution into the drive. You do not need to mount it at this time.

NOTE If you downloaded your Linux distribution as an *iso* image, skip to step 5.

3 In a terminal, extract an ISO image of the disc by typing:

dd if=/dev/`<CD/DVD device>`**of=/**`<destination path for the extracted ISO image>`

For example:

dd if=/dev/cdrom of=/tmp/Centos5.iso

4 Eject the disc.

5 Access the *dist/kickstart* subdirectory of your Burn installation package.

The directory contains a kickstart file, *RHEL5_CentOS5_kickstart.cfg*, as well as a script that adds the kickstart file to an ISO image.

6 Run the *build_kickstart_cd* script to add the kickstart file to the ISO image of your Linux distribution DVD or first CD:

./build_kickstart_cd RHEL5_CentOS5_kickstart.cfg`<original ISO image name> <new ISO image name>`

For example:

./build_kickstart_cd RHEL5_CentOS5_kickstart.cfg /tmp/Centos5.iso /tmp/Centos5_KS.iso

7 Burn the new ISO image to a blank disc using a tool such as **cdrecord**. For example:

cdrecord -v speed=2 dev=0,4,0 /tmp/Centos5_KS.iso

NOTE The `dev` value depends on your hardware configuration. Type **cdrecord -scanbus** to determine the address of your CD or DVD writer.

8 Use the new disc as the DVD or first CD of the CentOS distribution you plan to install on the node.

Linux Installation Workflow

Perform the following procedure to prepare your system and perform a fresh install of the customized distribution of Red Hat Enterprise Linux or of CentOS.

To install Linux:

1. Locate the installation media.

   For Burn nodes that were not purchased from Autodesk, use the CentOS disc that you added the Autodesk kickstart file to.

   For Burn nodes purchased from Autodesk, the DVD containing the customized Autodesk distribution of Red Hat Enterprise Linux is included with your shipment. The customized Autodesk distribution installs certain Linux packages that are required by Autodesk applications but are not installed by the commercial distribution of Red Hat Enterprise Linux.

   WARNING Even though your shipment may also include the disc set for the commercial distribution of Red Hat Enterprise Linux, **DO NOT** install the commercial distribution. Your software only works on the custom Autodesk distribution of Red Hat Enterprise Linux.

2. Make sure your system is in the following state:
   - Mouse, keyboard and graphics monitor are connected, and the graphics monitor is powered on.
   - The DVD or CDROM drive is set as the primary boot device in the workstation BIOS.

     For information on configuring your workstation BIOS, refer to the documentation for your hardware.
   - If you are using a KVM switch, it is switched to the system on which you want to install Linux.
3. Insert the disc of the Autodesk distribution of Red Hat Enterprise Linux, or of the CentOS distribution. This is the disc you added the Autodesk kickstart file to in the previous procedure.
4. Restart the system.

   The system should boot to the Linux installation disc. If it does not, review your BIOS settings to make sure the DVD / CDROM drive is set as the primary boot device.
5. At the boot prompt in the Red Hat Linux or CentOS installation menu, type one of the following commands and then press `Enter` to launch the Linux installation.

Type:	To install on:
burn	A Burn node using the Autodesk custom DVD of Red Hat Enterprise Linux

Type:	To install on:
linux ks=cdrom	A Burn node using CentOS with the Autodesk kick-start file

WARNING You must use this exact command to start the Linux installation. If you use the default options presented by the Linux installation prompt, the installation proceeds, but the Autodesk kickstart file is not read by the installer, and some important packages required by Autodesk software are not installed.

The command launches the Linux installation. The system spends several minutes loading drivers. The installer guides you through the rest of the process.

NOTE If you experience problems with the graphical Linux installer, reboot your computer and run the installer in low resolution VESA mode by typing **linux ks=cdrom xdriver=vesa**. If you still experience problems, reboot again and start the installer in text-only mode by typing **linux ks=cdrom text** at the Linux installation prompt.

6 The installation process may ask you to initialize the system disk. Follow the prompts to initialize the disk, if necessary.

7 Insert the remaining Linux distribution discs if prompted.

8 The remainder of the installation process is automated.

When the installation completes, you are prompted with "Congratulations, the installation is complete".

NOTE If the installation drops into text mode or to a blank screen just before completing, press `CTRL+ALT+F6` to return to graphical mode. At this point, the installation should be finished, and you should see the "Congratulations..." message and the Reboot button.

9 Eject the disc and click Reboot to reboot the system. After the system reboots, perform the post-installation tasks in the following section.

NOTE If after rebooting, Linux fails to start in graphical mode on the node, log in as root at the text mode login prompt, and install the DKU. See Installing the DKU (page 21). The default root password for a fresh Linux installation on a Burn node is *password*.

Linux Post-Installation Tasks

After booting into your new Linux installation, perform the following post-installation tasks:

1. Change the default root password to secure the system.
   - Log into your system as *root*, using the default password *password*.
   - Open a terminal and type:

     passwd
   - Enter your current password when prompted, and then enter the new password.
2. If you did not receive your render node from Autodesk, check for firmware or driver updates for your hardware after CentOS is installed. Refer to the CentOS and/or hardware manufacturer Web sites to ensure you have the correct firmware and drivers.
3. Configure the time zone for your geographic location. The automated Autodesk installation sets the time zone to North American Eastern Standard Time (EST) by default.
4. Configure the network settings for your system to match the ones used in your facility. The automated Linux installation sets a default IP address and hostname for your system, which may not be suitable for your facility network. See Configuring Basic Network Settings (page 18).
5. Install the Discreet Kernel Utilities (DKU) on each node. See Installing the DKU (page 21).
6. Optional: Configure the render node to use an InfiniBand network if you plan to integrate the render node into an InfiniBand-connected background processing network. See Configuring an InfiniBand Card (page 22).

Configuring Basic Network Settings

This section provides instructions on configuring your system network settings to match those of your facility.

Before configuring your network settings, gather the following information from your network administrator:

- An available IP address on your network, to assign to your system. The IP must not change over time, and must not be used by any other system on your network.
- A unique host name for your system.
- The IP address of the gateway system on your network.
- The subnet mask of your network.
- The IP addresses of one or more DNS servers.

The following files contain the network settings of your Linux system.

Networking Configuration File	Description
/etc/sysconfig/network	Contains global networking settings, such as the system host name, Gateway IP address, and NIS domain name, if applicable.
/etc/resolv.conf	Contains DNS server information.
/etc/hosts	Contains IP/host name pairs for host name resolution.
/etc/sysconfig/network-scripts/ifcfg-eth<x>, where <x> specifies the ethernet port number, usually 0 for the first on-board port.	Contains interface-specific configuration.

The following procedures provide instructions on setting up each configuration file. Log is as root before editing these files. When you are done changing the network settings, restart your system for the changes to take effect.

To configure the system hostname and default Gateway:

1. Open the */etc/sysconfig/network* file in a text editor.
2. Make sure the `NETWORKING` keyword is set to *yes*.
3. Configure the `HOSTNAME` keyword with the hostname you want to use for your system. For example:

   ```
   HOSTNAME=burn1
   ```

4 Configure the GATEWAY keyword with the IP address of the gateway system on your network. For example:

```
GATEWAY=192.168.0.1
```

This is the global default gateway that will be used by the system if no gateway is defined in a network port's individual configuration file.

5 Save and close the file.

To configure the system IP address and netmask:

1 Open the */etc/sysconfig/network-scripts/ifcfg-eth*<x> file in a text editor, where <x> is the number of the network port your system uses to connect to the facility network, typically 0.

2 Make sure the BOOTPROTO keyword is set to *static*.

```
BOOTPROTO=static
```

3 Set the IPADDR keyword to the IP address you want to use. Make sure this IP is not used by any other system on your network. For example:

```
IPADDR=192.168.0.100
```

NOTE It is not recommended to change this IP address once you have installed your Autodesk software.

4 Set the NETMASK keyword to the netmask used by your network. For example:

```
NETMASK=255.255.255.0
```

5 Make sure the ONBOOT keyword is set to yes.

```
ONBOOT=yes
```

6 Optional: If you have not specified a global default gateway for your system, or if you want to use a different gateway for this particular network port, add a GATEWAY keyword in this file, and set it to the IP address of the gateway. For example:

```
GATEWAY=192.168.0.1
```

7 Save and close the configuration file.

To configure host name resolution:

1 Open the */etc/hosts* file in a text editor.

2 Locate the loopback setting in the *hosts* file. By default, the loopback setting is identified by the IP address 127.0.0.1, for example:

```
127.0.0.1 vxfhost.localhost.localdomain localhost
```

3 Remove the host name preceding `localhost.localdomain` from the loopback setting.

   For example, if you remove the host name `vxfhost` in the example above, the loopback setting should appear as follows:

   ```
   127.0.0.1 localhost.localdomain localhost
   ```

4 On a new line, type the IP address of the system, press `Tab`, and then enter the explicit host name of the system.

   The new line defines the network IP address and host name of the system. For example, if the IP address is 192.168.0.100 and the host name is *burn_02*, your *hosts* file should be as follows:

   ```
   127.0.0.1 localhost.localdomain localhost
   192.168.0.100 burn_02
   ```

5 Optional: Add hostname / IP address pairs for other workstations on your network. For example:

   ```
   192.168.0.150 flame1
   ```

6 Save and close the file.

To configure DNS server information:

1 Open the */etc/resolv.conf* file in a text editor.

2 Add a line for each DNS server IP address you want to use. The syntax is:

   ```
   nameserver <DNS_server_IP>
   ```

   For example:

   ```
   nameserver 22.10.50.101
   nameserver 22.10.50.102
   ```

   NOTE Be careful when editing the *resolv.conf* file. The syntax is very strict, and any extra spaces or carriage returns can render the file unusable by the system.

Installing the DKU

Before installing your software, you must install the required version of the DKU on all nodes. See the Release Notes for the required DKU version for the current version of Burn.

NOTE DKU installation is mandatory whether you are using Red Hat Enterprise Linux or CentOS.

For major releases, the DKU is available on the application DVD or as a download from Autodesk. For extensions and service packs, the DKU is only available for download. The download link is provided in the release announcement you received from Autodesk.

To install the DKU:

1. Open a terminal and log in as root.
2. If you are upgrading an existing application, check the currently installed DKU version by typing:

   head -n1 /etc/DKUversion

   If the DKU version output by the command does not match the version required for the new application version, perform the remaining steps in this procedure.
3. Mount the USB key (page 23) or extract the tar file (page 25) to access the DKU installation directory.
4. Go to the DKU installation directory (*/mnt/usbdisk/DKU-<version>*, if you are installing from the application USB key) and launch the DKU installation script:

   ./INSTALL_DKU
5. If you installed from a disc, return to the root directory, and eject the disc by typing:

   eject
6. Reboot the system. Type:

   reboot

Configuring an InfiniBand Card

If you are planning to use the render node in an InfiniBand-connected background processing network, the render node you are configuring must be equipped with an InfiniBand network adapter.

The precompiled QuickSilver (QLogic) InfiniServ 9000 HCA adapter drivers for the Red Hat Enterprise Linux kernel are included in the *dist/ib* subdirectory of the Burn installation package.

If you are using CentOS, you need to manually compile the InfiniBand driver for your version of the Linux kernel. The source files for the driver are located in the *src/infiniband* subdirectory of the latest DKU installation package. Refer to your CentOS documentation for information on compiling software.

NOTE Refer to the README file located inside the driver *tar* file for instructions on how to install the driver and configure your InfiniBand interface.

Mounting the Product USB Key

For major releases, all your product software is distributed on a USB key.

You must mount the USB key before you can install the product software.

NOTE If Red Hat is newly installed (or re-installed), automount is enabled. When the USB key is connected to the system, the USB device is mounted automatically. However, you cannot install the DKU because the Red Hat installation does not allow for the execution of commands from a USB device. Before you install the software, you must unmount the USB key and then follow the instructions below for mounting this USB device. To unmount the USB key, type **eject /<mountpoint> or eject /dev/<device id>** .

To mount the product USB key:

1. Plug in the USB key to your workstation.
2. Log in to your workstation as root and open a terminal.
3. List the disk devices attached to your system by typing:

   fdisk -l | grep "/dev/sd"

   In the output of the command, identify the SCSI device name that was assigned to the USB device. It should look similar to */dev/sds* , */dev/sde* , etc.

   For example:

   ```
   Disk /dev/sda: 250.0 GB, 250059350016 bytes
   /dev/sda1 * 63 208844 104391 83 Linux
   /dev/sda2 208845 4401809 2096482+ 82 Linux swap / Solaris
   /dev/sda3 4401810 488392064 241995127+ 83 Linux
   …
   …
   Disk /dev/sdf: 8011 MB, 8011120640 bytes
   ```

```
/dev/sdf1 62 12128687 6064313 c W95 FAT32 (LBA)
```

4. Alternatively, you can use the *dmesg* command right after you have connected the USB drive. This will list the most recent USB device connected. For example:

```
usb-storage: device found at 30
usb-storage: waiting for device to settle before scanning
Vendor: TOSHIBA Model: TOSHIBA USB DRV Rev: PMAP
Type: Direct-Access ANSI SCSI revision: 00
SCSI device sdf: 31277056 512-byte hdwr sectors (16014
MB)
sdf: Write Protect is on
sdf: Mode Sense: 03 41 80 00
sdf: assuming drive cache: write through
SCSI device sdf: 31277056 512-byte hdwr sectors (16014
MB)
sdf: Write Protect is on
sdf: Mode Sense: 03 41 80 00
sdf: assuming drive cache: write through
sdf: sdf1
```

5. In the example above, the sdf device was the last device to be connected.
6. Create a mount point directory for the disk. Type:
   mkdir /mnt/usbdisk
7. Mount the disk on the new directory:
   mount -o shortname=winnt /dev/<device>1 /mnt/usbdisk
   where <device>1 is the usb device name. For example:
   mount -o shortname=winnt /dev/sdf1 /mnt/usbdisk

NOTE You can also mount the disk using the GUI. Double-click the Home icon on the desktop, go to the root folder, expand the 'dev' folder, right-click the USB device and select 'Mount' from the context menu.

Extracting an Installer from a TAR Package

At extensions and service packs, your product software packages are distributed as TAR files. Links to the TAR files for your product are supplied in the Release Announcement.

You must extract an installer from the TAR package before you can install your software.

1. Log in to your workstation as root and open a terminal.
2. Download the *tar* file from Autodesk.
3. After downloading the *tar* file, verify its integrity using the Linux md5 checksum. In the directory where you saved the *tar* file and the checksum file, type **md5sum** `<filename.tar.gz>`, and make sure the checksum displayed matches the Linux md5sum listed in the checksum file.
4. Unpack the downloaded application *tar* file to a temporary directory:

   tar -zxvf`<filename.tar.gz>`

   The file is uncompressed and an installation directory is created on your system.

Installing and Configuring Burn

5

Overview

This section outlines the general workflow for installing or upgrading Burn and Backburner.

To install or upgrade Burn and Backburner:

1. Install or upgrade Backburner on the system that will act as the Backburner Manager for submitted Burn jobs. See Installing Backburner Manager (page 28).
2. Install or upgrade Burn on each render node to be used in the background processing network. See Installing Burn on Render Nodes (page 28).
3. Activate each render node so it can receive and process jobs from Backburner Manager. See Connecting the Render Node to Backburner Manager (page 30).
4. Configure Burn-related settings for the Autodesk Visual Effects and Finishing application. See Configuring Your Visual Effects and Finishing Application for Burn (page 31).
5. If this is the first time you installed Burn on this render node, enable multicasting. See Configuring Multicasting (page 36).
6. Optional: Install additional fonts on the render node, if necessary. See Installing Additional Fonts (page 37).
7. Optional: Disable local Stone and Wire I/O operations on the render node to improve performance. See Disabling Local Stone and Wire I/O on Burn Nodes (page 37).

8 Optional: If you are also running versions of Burn earlier than 2.0 on the same render node, synchronize the version of the Burn client and Burn server with the version of the Autodesk Visual Effects and Finishing application you are running. See Running Multiple Versions of Burn on the Same Node (page 38).

9 License the systems on your background processing network. See Network Licensing (page 43).

10 Start using Burn. See Using Burn (page 51).

Installing Backburner Manager

Backburner Manager acts as the communications hub for the background processing network.

If you create a background processing network that serves a single Visual Effects and Finishing workstation, you can run Backburner Manager on the workstation itself. Backburner components are installed automatically with Visual Effects and Finishing applications.

If the background processing network serves multiple workstations, it is recommended to run Backburner Manager on a dedicated machine on your network. Backburner Manager works on Windows, Linux, and Mac OS X systems.

Refer to the *Autodesk Backburner Installation Guide* for detailed system requirements and instructions on installing and configuring Backburner Manager. You can find this document at *www.autodesk.com/backburner-documentation*.

Installing Burn on Render Nodes

Install Burn on each render node. You can install Burn2012 alongside earlier versions of Burn.

NOTE Installing Burn on Visual Effects and Finishing workstations is not supported. Install Burn only on dedicated render nodes in your facility.

To install Burn:

1 Log in to the render node as root, and open a terminal.

2 Mount the USB key (page 23) or extract the tar file (page 25) to access the Burn installer.

3 Go to the Burn installation directory.

NOTE The same Burn installation package is used for Red Hat Enterprise Linux and CentOS.

4 Run the installation script, by typing:

./INSTALL_BURN

The Burn, and Backburner Server packages are installed.

5 If you are installing the Smoke for Mac OS X edition of Burn, the following prompts appear:

- **Enter the license server name or address** Type the server name or IP address.
- **Enter the license server MAC address** Type the MAC address of the license server

For more information on licensing Smoke for Mac applications, see the *Autodesk Smoke Installation and Licensing Guide*.

6 If graphics are enabled, the following prompts appear:

- **Backburner Manager configuration: Do you want to automatically run the Backburner Manager on this machine?** Click No.
- **Backburner Server configuration: Do you want to automatically run the Backburner Server on this machine?** Click Yes
- **Backburner Server configuration: Do you want to enter the manager for this server?** Click Yes.

  The *manager.host* file appears. If you know the host name or IP address of your Backburner Manager, enter it here. Otherwise, close it without saving.

  NOTE You may edit this file later. See Connecting the Render Node to Backburner Manager (page 30).

  Burn is installed on the render node.

Configuring Burn

After the Burn software is installed on the render node, perform the following tasks to configure it.

Connecting the Render Node to Backburner Manager

With Burn installed, validate the connection to Backburner Manager using the following procedure to check settings in the *manager.host* configuration file. Repeat this procedure for each render node on the network.

NOTE You do not need to perform this procedure if you already specified the correct manager host name when installing Burn on each node.

To specify the Backburner Manager connection:

1. Log in as root on the render node.
2. In a text editor, open the */usr/discreet/backburner/cfg/manager.host* configuration file.
3. Replace the text in the file with the host name or IP address of the system running Backburner Manager. For example:

   172.16.25.2

   Contact your network administrator if you are unsure of these values.

   NOTE The name or IP address you specify here for Backburner Manager must be the same as the one specified in the configuration of the Visual Effects and Finishing application. See Configuring Burn-related Keywords for Linux Applications (page 33) or Configuring Burn-related Settings in the Smoke Setup Utility on Mac OS X (page 31). Note also that the *manager.host* file cannot contain comments.
4. Save and close the *manager.host* file.
5. Restart Backburner Server on the render node by typing:

   /etc/init.d/backburner_server restart

Configuring Your Visual Effects and Finishing Application for Burn

Your Autodesk Visual Effects and Finishing application can use Burn by default. No additional software is required. You must, however, update the Visual Effects and Finishing application to the same version as Burn, and to configure the application to use Burn.

To configure your application and workstation to use Burn, do the following:

- For Smoke for Mac OS X, configure Backburner settings in the Smoke Setup utility. See Configuring Burn-related Settings in the Smoke Setup Utility on Mac OS X (page 31).
- For Linux Visual Effects and Finishing applications, configure the Backburner keywords in the software initialization configuration file (*init.cfg*). See Configuring Burn-related Keywords for Linux Applications (page 33).

Configuring Burn-related Settings in the Smoke Setup Utility on Mac OS X

NOTE This section only applies to Smoke for Mac OS X. If you are configuring Burn with a Visual Effects and Finishing application running on Linux, you must manually edit the application's *init.cfg* configuration file. See Configuring Burn-related Keywords for Linux Applications (page 33).

To connect Smoke to Backburner Manager and the background processing network, you must configure four settings in the Smoke Setup utility on Mac OS X workstations.

To configure Backburner settings for Smoke for Mac OS X:

1. Run the Smoke Setup utility from the *Applications / Autodesk / Smoke 2011 / Smoke Utilities* folder.
2. Click the Backburner tab.

   Backburner settings are grouped into Application Settings and Local Server Settings. To use Smoke with Burn, you only need to configure the four fields in the Application Settings section.

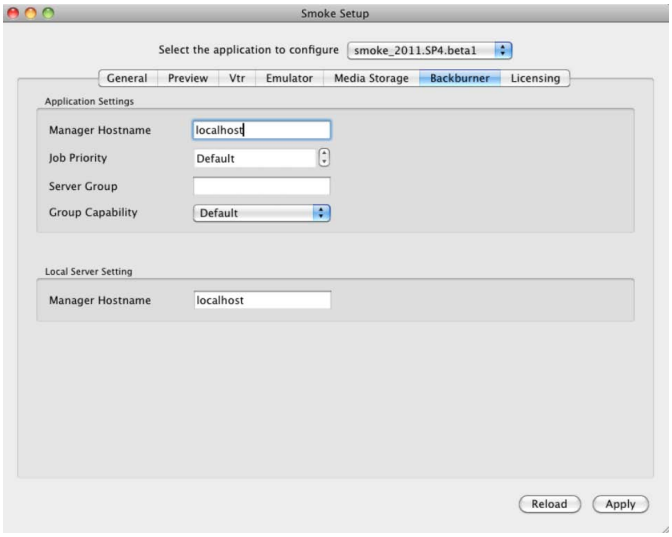

3 In the Manager Hostname field, enter the hostname or IP address of the Backburner Manager system that will handle background jobs submitted by Smoke.

NOTE Setting the Manager Hostname also enables the Background Wire® and Background Proxies buttons in the application; they are otherwise greyed out.

4 In the Job Priority field, enter the priority for jobs sent by Smoke to the Backburner Manager on your rendering network.

Job priority can range from 0 to 100, where 0 is the highest priority and 100 is the lowest. The default is 50.

5 In the Server Group field, enter the name of the group of render nodes to which jobs created in Smoke will be submitted. Refer to the latest *Autodesk Backburner User Guide* for information on creating groups.

6 From the Group Capability drop-down list, select whether the nodes in your group are equipped with GPU-accelerated graphics cards or not. Set this value as follows:

- If none of the nodes in your rendering network are equipped with GPUs, select *software*. Smoke will not send jobs that require a GPU to the rendering network, but only jobs that can be rendered in software mode by the render nodes.

- If **ALL** the nodes in your rendering network are equipped with GPUs, select *gpu*. Smoke will send all jobs to the GPU-equipped nodes, even if some jobs do not specifically require a GPU render node.

  NOTE If your rendering network also contains render nodes without a GPU, but this value is incorrectly set to *gpu*, all jobs are sent only to GPU-equipped render nodes, and the nodes without a GPU are never used.

- If your rendering network contains a mix of nodes with GPUs and without GPUs, select *hybrid*. The application sends all jobs to the rendering network, and Backburner Manager distributes each job to the appropriate type of render node.

  NOTE Select *gpu* or *hybrid* only if you are sure that at least one node in your rendering network is equipped with a GPU. Attempting to submit a job that requires a GPU to a rendering network with no GPU-equipped render node results in the job being stuck in the rendering queue indefinitely.

7 Click Apply to save the new settings, and then close the Smoke Setup utility.

Configuring Burn-related Keywords for Linux Applications

NOTE This section only applies to Visual Effects and Finishing applications running on Linux workstations. If you are configuring Burn with Smoke for Mac OS X, see Configuring Burn-related Settings in the Smoke Setup Utility on Mac OS X (page 31).

To connect a Visual Effects and Finishing application to Backburner Manager and the background processing network, you must set up four keywords in the application software initialization configuration file (*/usr/discreet/*`<product_home>`*/cfg/init.cfg*).

Perform the following procedure for each application installed on your workstation.

NOTE Some of these keywords are already enabled if you are using Background Wire and Background Proxies in your Visual Effects and Finishing application.

To set up Burn-related keywords for your application:

1. Exit the Visual Effects and Finishing application and open a terminal.
2. Open the software initialization configuration file */usr/discreet/*`<product_home>`*/cfg/init.cfg* in a text editor.
3. Locate the `BackburnerManagerHostname` keyword.

   This keyword identifies the workstation running the Backburner Manager on the background processing network.

   Uncomment this keyword and set its value to the host name or IP address of the system running Backburner Manager on the background processing network.
4. Locate the `BackburnerManagerPriority` keyword.

   This optional keyword sets the priority for all Burn jobs submitted by the application to the Backburner Manager, from 1 (most important) to 100 (least important). The Manager uses job priority to determine the order in which jobs are processed. By default, all submitted jobs are given the same priority, which allows Backburner Manager to manage jobs and network resources automatically.

   Uncomment this keyword and change its value to override the Manager's default behavior when processing your Burn jobs.

   WARNING Changing the `BackburnerManagerPriority` keyword may compromise access to the background processing network for other users in your facility. Autodesk does not recommend changing this value, except if you are in a client session and/or facing a deadline.
5. Locate the `BackburnerManagerGroup` keyword.

   This optional keyword specifies a Server group (a preset group of render nodes) used to process all Burn jobs submitted by the application. By default, Backburner Manager assigns a job to all available render nodes capable of processing it.

   If you have a dedicated group of render nodes for processing Burn jobs, uncomment this keyword and set its value to the name of the render node group.
6. Scroll through the file and locate the `BackburnerManagerGroupCapability` keyword.

   This keyword specifies whether the nodes in your background processing network are equipped with GPUs or not. Based on the value of this keyword, the Visual Effects and Finishing application enables or disables the submission of jobs that require a GPU (such as floating point jobs) to the background processing network.

Configure the `BackburnerManagerGroupCapability` keyword according to the GPU capabilities of the nodes in your background processing network:

- If none of the nodes in your background processing network are equipped with GPUs, set the value of the `BackburnerManagerGroupCapability` keyword to *software*. The application will not send jobs that require a GPU to the background processing network, but only jobs that can be processed in software mode (using OSMesa) by the Burn render nodes.
- If ALL the nodes in your background processing network are GPU-enabled, set the value of the `BackburnerManagerGroupCapability` keyword to *gpu*. The application will send all jobs to the GPU-equipped Burn nodes in the background processing network, even if some jobs do not specifically require a GPU Burn node. The GPU-equipped render nodes will render jobs that require a GPU, as well as OSMesa jobs.

  NOTE If your background processing network also contains Burn nodes without a GPU, but the `BackburnerManagerGroupCapability` keyword is wrongfully set to *gpu*, all jobs are sent only to GPU-equipped render nodes, and the nodes without a GPU are never used.

- If your background processing network contains a mix of nodes with GPUs and without GPUs, set the `BackburnerManagerGroupCapability` keyword to *hybrid*. The application will send all jobs to the background processing network, and Backburner Manager will distribute each job to the appropriate type of Burn render node. Jobs that require a GPU are sent only to GPU-equipped Burn nodes, while jobs that do not require a GPU are sent to any available Burn render node (GPU or non-GPU), to be processed in software mode.

NOTE Set the `BackburnerManagerGroupCapability` keyword to *gpu* or *hybrid* only if you are sure that at least one node in your background processing network is equipped with a GPU. Attempting to submit a job that requires a GPU to a background processing network with no GPU-enabled Burn nodes results in the job being stuck in the queue indefinitely.

7 Save and close the *init.cfg* file, and restart your Visual Effects and Finishing application.

Configuring Multicasting

Perform the following procedure on each render node to enable multicasting in Stone and Wire. You do not need to perform these tasks if you upgraded an existing installation of Burn.

1. Open the */usr/discreet/sw/cfg/sw_probed.cfg* configuration file on the render node in a text editor.
2. Enable the `SelfDiscovery` option by uncommenting it and setting its value to *Yes*. By setting this option to Yes, *sw_probed* will run in self-discovery mode and it will automatically probe the network for other systems. This option is set to Yes by default when Stone and Wire is installed on the render node.
3. Use the `Scope` parameter to define the scope for the multicast. For networks with one subnet, set the `Scope` parameter to *LinkLocal*. For networks with subnets, use a value that is appropriate for your requirements and router configuration. Available values are *LinkLocal*, *OrganizationLocal*, and *GlobalWorld*.

   The scope is set to *LinkLocal* by default.

   NOTE This value must be the same for all machines on your network.
4. Save and close the *sw_probed.cfg* configuration file.
5. If the workstations and Burn nodes in your facility are on separate networks connected through routers, use the `ttl` parameter in the file to specify the number of router hops for a multicast. Burn-related transfers across multiple routers may cause bottlenecks at network bridges, especially with jobs involving film frames. Using the `ttl`parameter may reduce multicast-related traffic and improve general network performance in your facility.

   NOTE Consult your network administrator for guidance on setting the appropriate values for your network.
6. Restart the *sw_probed* daemon. Type:

   /etc/init.d/stone+wire restart

Installing Additional Fonts

When you install Burn, the same fonts that are installed by default with your Autodesk Visual Effects and Finishing application are also installed on the render nodes.

However, if you installed additional fonts on the workstation, that are not provided with your application, you must also install those fonts on each render node. Contact your third-party font supplier(s) for information about Linux support for those fonts.

Make sure that any 3D Text fonts used with Action nodes in the Batch setups you submit to Burn are installed.

Disabling Local Stone and Wire I/O on Burn Nodes

If the Burn render node has a slower connection to a shared storage device (such as a SAN) than the Visual Effects and Finishing workstation, it is preferable for Stone and Wire I/O operations to be performed on the Visual Effects and Finishing workstation, rather than on the render node.

Perform the following procedure to disable all local Stone and Wire I/O operations on a Burn render node.

To disable local Stone and Wire I/O operations:

1. Open a terminal on the Burn render node, and log in as root.
2. Stop Stone and Wire by typing:
   /etc/init.d/stone+wire stop
3. Open the */usr/discreet/sw/cfg/stone+wire.cfg* configuration file in a text editor.
4. Locate the *DisableLocalIO* keyword in the *[Initialization]* section.
5. Uncomment the keyword, and set its value to true:
   `DisableLocalIO=True`
6. Save and close the configuration file.
7. Restart Stone and Wire by typing:
   /etc/init.d/stone+wire start

Running Multiple Versions of Burn on the Same Node

Each version of an Autodesk Visual Effects and Finishing application includes its own version of Burn, and each version of Burn is compatible with only one version of Autodesk Visual Effects and Finishing applications.

Because your facility may have different versions of Autodesk Visual Effects and Finishing applications running with their corresponding version of Burn, you can run Burn2012 alongside earlier versions of Burn.

After you have installed the current version of Burn on the node, perform the following tasks to run multiple versions of Burn:

- Determine the Burn client and server version needed for the version of the Autodesk Visual Effects and Finishing application that you are running. See Assessing Compatibility Between Client and Server Versions (page 38).
- Ensure there is a Burn server running on the background processing network that is capable of handling the jobs for the version of Burn you are running. See Managing Multiple Burn Servers on a Render Node (page 39).
- For Autodesk applications that require the Burn 1.x client, synchronize the Burn client with the application. See Synchronizing a Burn 1.x Client with the Autodesk Application (page 40).

Assessing Compatibility Between Client and Server Versions

For Burn versions 2.0 and later, your Visual Effects and Finishing application sends jobs directly to the background processing network, removing the need to synchronize the client with the server.

For Burn versions earlier than 2.0, you must synchronize the version of the Burn client and Burn server with the version of the Autodesk Visual Effects and Finishing application you are running.

The following table lists the compatibility between versions of Autodesk Visual Effects and Finishing applications and versions of the Burn client and server earlier than 2.0.

Burn Client Version	Burn Server Version	Visual Effects and Finishing Version
Burn client 1.0.x	Burn 1.0.x	Inferno 5.3.2, Flame/Flint 8.3.2
Burn client 1.5	Burn 1.5	Inferno 5.5, Flame/Flint 8.5
Burn client 1.6	Burn 1.6	Inferno 6.0, Flame/Flint 9.0, Fire® 6.5, Smoke 6.5
Burn client 1.7	Burn 1.7	Inferno 6.2, Flame/Flint 9.2, Fire 6.7, Smoke 6.7

Managing Multiple Burn Servers on a Render Node

You can have multiple versions of the Burn server installed on a render node to handle jobs from different Burn clients. For example, you can run the Burn 1.6 and Burn2012 servers to allow the same render node to handle jobs from the Burn 1.6 client used by Flame 9.0 and Smoke 6.5, as well as jobs from other Autodesk applications that use Burn2012.

Use the *select_burn* script to help you manage the Burn servers running on a render node. You can use this script to do the following:

- List the Burn server versions installed on a render node.
- Select the version of a Burn server responsible for processing jobs for a particular version of Burn.
- Test the configuration of a Burn server to make sure it is working.

Use the following procedures to list the Burn server versions installed on a render node and set the type of Burn jobs that each server handles.

To list the Burn server versions installed on a render node:

1. On the render node, log in as root, and then open a terminal.

2 Navigate to the *bin* directory of the latest Burn version. For example:

cd /usr/discreet/burn_2011/bin

3 Run the *select_burn* script in the directory to list the Burn server versions that are installed on the render node. Type:

./select_burn -l

All versions of the Burn server installed on the render node appear, showing their current status. Burn servers listed as *Enabled* process jobs of their type. If necessary, continue to the next procedure to select a specific version of the Burn server responsible for processing specific types of jobs on the render node.

TIP Type **./select_burn -h** to learn more about the options and usage of this script.

To select the version of the Burn server used for processing jobs on a node:

1 On the render node, log in as root, and then open a terminal.

2 Navigate to the *bin* directory of the latest Burn server. For example:

cd /usr/discreet/burn_2011/bin

3 Run the *select_burn* script to enable a particular version of the Burn server on the render node by typing the following:

./select_burn -c `<job type, server version>`

You must specify the type of Burn job and the job's corresponding server version to enable the correct server. For example, to enable the Burn 2010 SP1 server on a render node to handle Burn 2010 jobs, type:

./select_burn -c 2010,2010.SP1

The Burn server version is enabled.

4 To ensure that the Burn server version is enabled properly, review the status of all Burn servers installed on the render node. Type:

./select_burn -l

The status of the Burn server listed in step 3 now appears as *Enabled*.

Synchronizing a Burn 1.x Client with the Autodesk Application

Autodesk Visual Effects and Finishing applications that used Burn 1.x included a daemon called the Burn client. The Burn client linked the Autodesk

application with the Backburner Manager. Because a Burn 1.x client can only submit processing jobs to its version of Burn, you must manually change the Burn client if you run different versions of Inferno, Flame, Flint, or Smoke on the same workstation.

If the workstations in your facility are running multiple versions of Autodesk applications, check the following list of versions that included the Burn client with Burn:

- Inferno 5.3.2 to Inferno 6.2
- Flame/Flint 8.3.2 to Flame/Flint 9.2
- Fire/Smoke 6.5 to Fire/Smoke 6.7

If a workstation in your facility is running one of the above applications, perform the following procedure to synchronize the Burn client version with the application. Otherwise, you will be unable to submit Burn jobs from these applications.

To synchronize the Burn client version with its application:

1. If necessary, exit the Autodesk application and log in as root on the workstation.
2. Stop the Burn client currently running. Type:

   /etc/init.d/burnclient stop
3. Switch to the *burnclient* directory by typing:

   cd /usr/discreet
4. View the contents of this directory. Type:

   ls

   This directory contains the following:

   - A sub-directory for the current version of the Burn client.
   - Sub-directories for earlier versions of the Burn client, named *burnclient.previous.<number>*. For example, */usr/discreet/burnclient.previous.1* contains the last installed version of the Burn client.
   - A symbolic link called *burnclient* that points to the directory of the current version of the Burn client.
5. Remove the existing symbolic link to the current version of the Burn client. Type:

   rm burnclient

6 Create a new link to the Burn client version that you want to run by typing:

ls -s `<target directory>` **burnclient**

where `<target directory>` is the name of the sub-directory containing the version of the Burn client you want to run.

7 Start the previous version of Burn client. Type:

/etc/init.d/burnclient start

The Burn client is started. You can now use a previous version of Burn with the Autodesk application.

NOTE Starting a Burn client with the **-C** option cleans the *burnclient* directory of completed Batch setups or timeline clips.

Network Licensing 6

Two Burn Licensing Scenarios

The licensing requirements for the Smoke for Mac OS X distribution of Burn are different from the Linux Visual Effects and Finishing distribution.

- To license Burn for Visual Effects and Finishing Linux workstations, see Licensing Workflow for the Linux Visual Effects and Finishing Distribution of Burn (page 43).
- For information on licensing the Smoke OS X distribution of Burn, see Workflow for Installing the Smoke for Mac OS X Distribution of Burn (page 8).

Two distributions of Burn 2012*cannot be installed* on the same node. However, either distribution can process jobs sent from a Mac or Linux product, as long as it is licensed.

Licensing Workflow for the Linux Visual Effects and Finishing Distribution of Burn

Although you can install Burn without a license, you must license it before you can use it.

Burn uses a "floating" license system, made up of the following components.
License Server A Linux daemon that provides concurrent licenses to Burn nodes on your network as needed.

Licensing clients Each Burn node on the network that requests a license from the License Server.

The following workflow outlines how to set up Burn floating licenses on your network.

To license a Burn network:

1. Install the license server. See Installing the License Server Software (page 45).

   NOTE If you already have an Autodesk license server in your network, skip this step and use the existing license server to distribute licenses to your application.

2. Obtain a license code for the license server. See Obtaining License Codes (page 44).
3. Add the Burn license code to the license server, and configure the server to distribute licenses to Burn nodes. See Creating a License File for a Remote License Server (page 46).
4. Configure each Burn node to retrieve a license from the license server. See Configuring Burn Nodes to Retrieve Licenses (page 48).
5. Optional: To avoid conflicts with other applications or license servers in your facility, you may need to change the default port setting used by the Burn license server. See Changing the Default Port Used by the License Server (page 49).

Obtaining License Codes

NOTE You get your license codes for the Smoke for Mac OS X distribution of Burn when you register Smoke. For details, see Workflow for Installing the Smoke for Mac OS X Distribution of Burn (page 8).

You can obtain license codes for Burn by registering the application with the Autodesk Media and Entertainment Licensing Department by e-mail or telephone.

All license codes obtained by e-mail or telephone are temporary 30-day licenses that you use until your permanent license is confirmed and delivered.

To obtain license codes for Burn, you must obtain the unique host ID of the license server. This ID is used to confirm your Burn license and issue license codes.

To obtain license codes for Burn:

1. Log in as root on the license server system.
2. Obtain the unique Discreet host ID for the system. Open a terminal and type:

   /usr/local/bin/dlhostid

   A message appears indicating the *dlhostid* of the machine. For example (your value will differ):

   ```
   The Discreet host ID of this machine is
   "DLHOST01=25231AEF83AD9D5E9B2FA270DF4F20B1"
   ```

3. Send the Discreet host ID (including the *DLHOST01*= part) to the Autodesk Media and Entertainment Licensing Department using one of the following methods to register Burn and obtain license codes:
   - **By E-mail** To acquire a license code by e-mail, submit a request with the host ID of the system to *me.licensing@autodesk.com*.
   - **By Telephone** You can speak to a licensing representative by calling the Licensing Department toll-free in North America at 1-800-925-6442 between 8 AM and 8 PM EST. Outside of North America, call 1-514-954-7199.

   NOTE 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.

4. Once you receive your license codes from the Licensing Department, add the licenses to the license server.

   WARNING The license string is locked to the hardware of the system you use as a license server. If you replace a major hardware component of the license server system, or if you decide to use another system as a license server, repeat this entire procedure to obtain the new *dlhostid* for the license server, and a new license from the Autodesk Media and Entertainment Licensing Department.

Installing the License Server Software

NOTE For the Smoke for Mac OS X distribution of Burn, you use the same license server as the application. See the *Autodesk Smoke for Mac OS X Installation and Licensing Guide* for information.

The license server is a Linux daemon that provides concurrent licenses to Burnnodes on the network, as needed.

You can install the license server on a Burnnode 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.

NOTE 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:

1. Log in as root to a system designated as a license server.
2. Access your Burn software installation directory.
3. Launch the license server installation script by typing:

   ./INSTALL_LICSERV

   The license server is installed.

Creating a License File for a Remote License Server

NOTE If you are using a Linux Network License Manager with the Smoke for Mac distribution of Burn, you must save the license file you receive from Autodesk in the following location: */opt/flexnetserver/adsk_server.lic*. For more information, see *Autodesk Smoke for Mac OS X Installation and Licensing Guide*.

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*.

Keyword	Description
DAEMON	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 by typing:
   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:

   ```
   SERVER server DLHOST01=886C2B75E8E57E4B03D784C3A2100AC0
   ```

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

   ```
   DAEMON discreet_l discreet_l
   USE_SERVER
   FEATURE burn_x86_64_r_2011 discreet_l 2011.999
   30-nov-2009 4 \
                 6D7AE3402ECB46174B70 ck=47
   ```

   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 Burn 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 nodes on your network.

Configuring Burn Nodes to Retrieve Licenses

NOTE This procedure is not required for the Smoke for Mac OS X distribution of Burn. See Workflow for Installing the Smoke for Mac OS X Distribution of Burn (page 8).

You must configure each Burn node to contact the license server to obtain its license. You do this by creating a license file on each Burn node to identify the license server. The node contacts the license server to obtain its license.

NOTE If the license server for your network is running on a Burn node, make sure you perform the following procedure on this node as well. Otherwise, this node is able to distribute Burn licenses to other nodes, but is unable to retrieve a license for itself.

To configure the Burnnode license:

1. Log in as root to the node.
2. Navigate to the licenses directory by typing:

   cd /usr/local/flexlm/licenses

   NOTE If the directory does not exist, create it by typing:

   mkdir -p /usr/local/flexlm/licenses

3. Create a file called *DL_license.dat* in the directory by typing:

   touch DL_license.dat
4. Open the *DL_license.dat* file in a text editor and copy the SERVER, DAEMON, and USE_SERVER lines into the license file. The strings for these keywords are provided by the Autodesk Licensing Department.

   The resulting license file should look similar to the following example:

   ```
   SERVER server DLHOST01=25231AEF83AD9D5E9B2FA270DF4F20B1
   DAEMON discreet_l discreet_l
   USE_SERVER
   ```
5. Repeat the preceding steps for each node.

Changing the Default Port Used by the License Server

NOTE This procedure is not required for the Smoke for Mac OS X distribution of Burn. See Workflow for Installing the Smoke for Mac OS X Distribution of Burn (page 8).

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 Burnnode on the network.

To change the default port used by a license server:

1. Log in as root to the system where the license server is installed.
2. Open the */usr/discreet/licserv/licenses/DL_license.dat* file in a text editor.
3. 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.

   Single license server example:

   ```
   SERVER server DLHOST01=886C2B75E8E57E4B03D784C3A2100AC0
   ```

   Redundant license servers example:

   ```
   SERVER host01 DLHOST01=886C2B75E8E57E4B03D784C3A2100AC0
   27005
   SERVER host02 DLHOST01=886C2B75E8E57E4B03D784C3A2100AC1
   27005
   SERVER host03 DLHOST01=886C2B75E8E57E4B03D784C3A2100AC2
   27005
   ```
4. 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, to use port 62222, add 62222 to the end of the line:

   ```
   SERVER server DLHOST01=886C2B75E8E57E4B03D784C3A2100AC0
   62222
   ```
5. Save and close the file.

6 Verify that the new port settings are correct. Type the following commands:

/etc/init.d/license_server stop

/etc/init.d/license_server start

cat /usr/discreet/licserv/log/license_server.log

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

7 Look for messages similar to the following examples in the output, and verify that the port numbers are what you requested:

```
15:08:49 (lmgrd) lmgrd tcp-port 6222215:08:49 (lmgrd)
Starting vendor daemons ...15:08:49 (lmgrd) Using vendor
daemon port 12344 specified in license file15:08:49
(lmgrd) Started discreet_l (internet tcp_port 12344 pid
5013)
```

8 Log in as root to each node on the network, open the file */usr/local/flexlm/licenses/DL_license.dat* in a text editor, and repeat steps 3 and 4 in this procedure, using the same port as the one you set for the license server.

Using Burn 7

Overview

Once Burn is installed and licensed, send jobs from your Autodesk Visual Effects and Finishing applications to the background processing network. The *background processing network* refers to all the nodes on the physical network that are used for background processing.

The following procedure provides a general overview for using Burn to do background processing and assumes that the network is configured properly, including the TCP/IP settings.

To process jobs using Burn:

1. Prepare and submit a job from the application. See Submitting Jobs to the Background Processing Network (page 51).
2. During an active session, monitor and manage jobs that are being processed on the network. See Monitoring and Managing Burn Jobs (page 52).
3. Preview the result of the submitted job, if necessary. See Previewing Results (page 52).

Submitting Jobs to the Background Processing Network

When the Visual Effects and Finishing application is properly configured for Burn, a Burn button appears in the clip library or in various application modules. Click this button to submit a setup, timeline, or clip as a background processing job. Refer to the application help for details.

NOTE The Burn button appears when Burn-related settings are configured in the *init.cfg* file of Linux applications, or in the Smoke Setup utility of Smoke for Mac OS X. However, this button's appearance does not mean that the background processing network is operational. Always check the status of the network using the Backburner Monitor before submitting jobs to Burn.

Monitoring and Managing Burn Jobs

During an active render session, you can monitor and manage jobs using the following:

- The Background I/O window in your Autodesk Visual Effects and Finishing application. See your application help.
- The browser-based Backburner Monitor. See the latest *Autodesk Backburner User Guide* for details.
- The Windows-based Backburner Monitor application. See the latest *Autodesk Backburner User Guide* for details.
- The browser-based Backburner Monitor included in Autodesk WiretapCentral. See the WiretapCentral chapter in your application help for details.

Previewing Results

You can preview results for Burn jobs by viewing the *_Burn_* library in the clip library of an Autodesk Visual Effects and Finishing application. As you are previewing your result, the clip may appear semi-rendered. Frames are not processed sequentially, so it is important to wait until the processing tasks are finished and the clip is completely processed.

To preview a Burn result:

1. Open the clip library. See your application help.
2. Open the *_Burn_* library.

   Submitted Burn Jobs appear with a name that identifies the machine, the date and time of the job, as well as the job type. Jobs that require a render node equipped with a GPU contain the string “gpu” in their name.

   The *_input* reel contains jobs and source clips that are sent as input to Burn for processing over the network. The *_output* reel indicates the

rendered result that is copied back to the Linux workstation. A clip labelled PENDING RENDER indicates that the clip is being rendered by Burn.

3 To refresh the view of the rendered result, press the F hot key.

Troubleshooting Burn 8

Overview

Troubleshooting problems with jobs on an Autodesk background processing network can be complex because Burn and Backburner components interact with Windows, Mac OS X, and Linux platforms over the network.

The following procedure outlines the general workflow for troubleshooting Burn and Backburner problems.

To troubleshoot Burn problems:

1. Shut down and restart all Burn and Backburner components, including the following:
   - All Autodesk Visual Effects and Finishing applications that submit jobs to Burn. See the user guide for your application.
   - Backburner Manager. See Restarting Backburner Manager and Backburner Server (page 56).
   - The Backburner Server running on each node processing Burn jobs. See Restarting Backburner Manager and Backburner Server (page 56).

   In some cases, simply restarting these components is enough to fix problems with Burn. Otherwise, use the remaining steps to diagnose and correct the problem.
2. Verify that the problem with Burn is not being caused by a larger issue with the background processing network, such as a lack of network connectivity. See Troubleshooting Problems in the Background Processing Network (page 60).

3 Use the *GATHER_BURN_LOGS* script to compile and review the Burn and Backburner logs from render nodes on the network. See Reviewing Burn and Backburner Log Files (page 69).

   Reviewing the Burn and Backburner logs helps you determine the following:

   - The component of the background processing network that is failing
   - The conditions under which the component fails

4 If these logs indicate that the problems may be due to background processing network components such as Wire networking, use scripts to troubleshoot the background processing network. See Testing Network Components (page 71).

5 If these logs indicate that the problems may be due to memory issues on the render nodes, or if you suspect that the render nodes do not meet the graphics card requirements for a certain job type, see Assessing Compatibility between Jobs and Render Nodes (page 73).

6 If neither the logs nor the included scripts help you troubleshoot the problem with the background processing network, contact *Autodesk Media and Entertainment Customer Support*.

Restarting Backburner Manager and Backburner Server

Backburner Manager and Backburner Server must be running before you can submit jobs to the background processing network.

These components are set to automatically start with the system they are installed on, thus, under normal circumstances, you do not need to manually start them. However, if you are experiencing issues with Backburner Manager and Backburner Server, perform the following tasks to manually restart these services.

NOTE Make sure only one Backburner Manager is running on the background processing network at any given time. Otherwise, Burn jobs submitted to the network may not be processed.

Restarting Backburner Manager on Linux Systems

By default, Backburner Manager is installed as a Linux service (daemon) that automatically starts with the system.

Use the following procedure to check if Backburner Manager is running on a Linux system, and restart it if needed.

To restart Backburner Manager on a Linux system:

1. Log in as root and open a terminal.
2. Check that the Backburner Manager daemon is running. Type:

   ps -ef | grep -i Backburner

   An output similar to the following appears:

   ```
   root 2997 1 0 Dec13 ? 00:00:03
   /usr/discreet/backburner/backburnerManager
   root 24204 24091 0 11:00 pts/4 00:00:00 grep -i
   backburner
   ```

   If `backburnerManager` does not appear in the output, continue to the next step to start this daemon.
3. Start the Backburner Manager daemon manually by typing:

   /etc/init.d/backburner_manager restart

   Backburner Manager starts.

   To view details on the Manager's status, check the backburner log files in the */usr/discreet/backburner/Network* directory.

Restarting Backburner Manager on Mac OS X Systems

By default, Backburner Manager is installed as a Mac OS X service (daemon) that automatically starts with the system.

Use one of the following procedures to check if Backburner Manager is running, and restart it if necessary. The first procedure applies to workstations where Smoke for Mac OS X is installed, while the second procedure is for stand-alone Backburner Manager Mac OS X systems.

To restart Backburner Manager on a Smoke Mac OS X workstation:

1. Run the Service Monitor application from the *Applications / Autodesk / Smoke Common Utilities* folder.

   The version and status of Backburner Manager are displayed in the Components tab.

2 If Backburner Manager is not running, click Start.

Backburner Manager starts.

To view details on the Manager's status, run the Console application from *Applications / Utilities*, and look for entries containing *com.autodesk.backburner_manager*.

To restart Backburner Manager on other Mac OS X systems:

1 Run the Terminal application from the *Applications / Utilities* folder.

2 Check that the Backburner Manager service is running. Type:

sudo ps -ef | grep -i backburner

3 Enter your Mac OS X user password when asked.

An output similar to the following appears:

```
0 358 1 0 0:00.03 ?? 0:00.33
/usr/discreet/backburner/backburnerManager 0 380 281 0
0:00.00 ttys000 0:00.00 grep -i backburner
```

4 If `backburnerManager` does not appear in the output of the command, restart the Backburner Manager service manually by typing:

sudo /usr/discreet/backburner/backburner_manager restart

Backburner Manager starts.

To view details on the Manager's status, run the Console application from *Applications / Utilities*, and look for entries containing *com.autodesk.backburner_manager*, or check the backburner log files in the */usr/discreet/backburner/Network* directory.

Restarting Backburner Manager on Windows Systems

By default, Backburner Manager is installed on Windows as a separate application that you must start manually. Use the following procedure to start Backburner Manager on a Windows system.

TIP You can also configure Backburner Manager to run as a Windows service that is started automatically with the system. See the *Autodesk Backburner Installation Guide* for instructions.

To restart Backburner Manager on a Windows system:

1. Go to the Windows workstation designated as the Backburner Manager for the background processing network.
2. From the Windows Start Menu, select Programs, Autodesk, Backburner, and then Manager.

   Backburner Manager starts and its application window appears. The message "Starting Network Manager" as well as server connection messages are displayed in the window.

   To view details on the status of Backburner Manager, leave its application window open or view the log files in the *C:\Program Files\Backburner\log* folder. See the *Autodesk Backburner Installation Guide*.

Restarting Backburner Server on Render Nodes

Perform the following procedure if you need to restart Backburner Server on each Burn render node on the background processing network.

To restart Backburner Server:

1. Log in as root to each render node.
2. Start Backburner Server by typing:

   /etc/init.d/backburner_server restart

   Backburner Server starts and searches the network for the Backburner Manager application specified in its */usr/discreet/backburner/cfg/manager.hosts* file. Once Backburner Manager is found, a message is logged to indicate that Backburner Server has successfully established communication with the Manager.

Backburner Manager should also detect the Backburner Server running on each node on the background processing network. If error messages indicate that Backburner Manager cannot find Backburner Server or vice versa, see Connecting the Render Node to Backburner Manager (page 30).

3 If you are still experiencing problems after restarting the Backburner components, perform the remaining procedures in this chapter.

Troubleshooting Problems in the Background Processing Network

Problems with processing jobs may not be directly related to Burn but to the general setup or operation of the background processing network. For example, problems in any of the following areas may cause Burn to appear to stop working:

- General network connectivity
- Stone and Wire connectivity and access
- Licensing issues

You can use basic tools like *ping* to check that these areas are not causing a problem for Burn. The following general troubleshooting procedure can be used to verify that the background processing network is working properly. More detailed procedures are provided in the following sections.

To verify that the background processing network is working properly:

1 Ensure that Burn network components are communicating with each other. See Ensuring Basic Network Connectivity Using ping (page 61).

2 Check that the mount points to workstations can be seen and accessed by render nodes. See Checking Mount Points on the Background Processing Network (page 62).

3 Verify that storage devices connected to workstations are available on the Wire network. See Verifying Stone and Wire Connectivity from the Background Processing Network (page 66).

4 Check that the Burn license for the background processing network is valid and is being used. See Checking Burn Licensing for Render Nodes (page 67).

Ensuring Basic Network Connectivity Using ping

The following information and procedure applies to all components of a background processing network.

Burn works only when components on an Autodesk background processing network are communicating with each other. For example, if Backburner Manager cannot contact the render nodes over the network, jobs cannot be rendered remotely.

You can use the *ping* command to test connectivity between components on the background processing network. The following guidelines suggest ways to test this connectivity using the *ping* command:

- From the workstation, use *ping* to test communication between the Backburner Manager and each render node on the network.
- From the Backburner Manager workstation, use *ping* to test communication between the workstation and each render node on the network.
- From the render nodes on the network, use *ping* to test communication between the Backburner Manager and the workstations that submit jobs to the background processing network.
- If your network infrastructure supports jumbo packets, check that jumbo packets can be transferred between the workstation and the render nodes. See Ensuring Network Connectivity for Jumbo Frames (page 61).

NOTE Refer to the documentation included with your operating system to learn how to use the *ping* command.

Ensuring Network Connectivity for Jumbo Frames

Jumbo frames increase the efficiency of a background processing network, but must be supported throughout the network infrastructure. Burn problems can be caused by jumbo frames being sent to a switch or network adapter that is misconfigured or cannot handle these frames.

If you know that your network infrastructure supports jumbo frame switching, use the following procedure to test if jumbo frames can be sent between the workstations and render nodes.

To test network connectivity using jumbo frames:

1. On a workstation or render node, open a terminal.

2 Run *ping* using the **-s** option to set the packet size used for network communications. Type:

ping -s 50000 `<hostname>`

where `<hostname>` is the hostname or IP address of the workstation or render node you are trying to reach.

3 If *ping* fails with the command in step 2, ensure basic network connectivity between the same two network components by typing:

ping `<hostname>`

where `<hostname>` is the hostname or IP address of the workstation or render node you are trying to reach.

If step 3 works while step 2 failed, recheck the configuration for your network adapter and/or switch to ensure the following:

- Jumbo frames are supported by both network components.
- Both the network adapter and switch are properly configured for jumbo frame support.

The network infrastructure in your facility may not support jumbo frame switching. Consult your network administrator if you are unsure if your network supports jumbo frames.

Checking Mount Points on the Background Processing Network

Certain directories on Visual Effects and Finishing workstations, such as */usr/discreet/clip* and */usr/discreet/project*, must be accessible as NFS mount points to render nodes, so that material for processing jobs can be retrieved by render nodes.

The directories shared by each Visual Effects and Finishing workstation must be mounted on the */hosts/*`<workstation_hostname>` directory on render nodes, where `<workstation_hostname>` is the host name of the workstation submitting jobs to Burn.

Under normal circumstances, these directories are mounted automatically when you install the Visual Effects and Finishing application on the workstations, and the Burn software on the render nodes.

Perform the following procedure to check these mount points and ensure that render nodes can access material.

To check that mount points are accessible to the background processing network:

1. Log in as root to a render node, and open a terminal, if necessary.
2. Connect to the */usr/discreet/clip* directory on a workstation that submits processing jobs. Type:

   cd /hosts/`<hostname>`**/usr/discreet/clip**

   where `<hostname>` is the name of the workstation submitting jobs to Burn.
3. If the previous step fails, configure the *amd* automounter service on each render node, and make sure the *amd* and *NFS* services are set to run automatically on each render node and workstation. See Configuring the amd Automounter on Render Nodes (page 64) and Configuring the NFS and amd Services to Start Automatically (page 64).
4. Try again to connect to the *usr/discreet/clip* directory of the workstation that is submitting jobs to the background processing network, and create an empty file by typing:

   touch /hosts/`<workstation_hostname>`**/usr/discreet/clip/export_test**

   where `<workstation_hostname>` is the hostname or IP address of the workstation.

   If the *export_test* file cannot be created, permission to write to the required directories may not be assigned to the render nodes.
5. Check the permissions assigned to the file you created in the previous step. Type:

   ls -al /hosts/`<workstation_hostname>`**/usr/discreet/clip**

   where `<workstation_hostname>` is the hostname or IP address of the workstation.

   File system details for the *export_test* file you created appear.
6. Ensure that the owner of the file is `root` and the group ID is `sys`.

   If permissions are not correct, perform the tasks in Enabling Export Permissions (page 65).

Configuring the amd Automounter on Render Nodes

The *amd* automounter is automatically installed on each render node by the custom Linux *kickstart* file.

Perform the following tasks to modify the *amd* configuration file to redirect the mount point to */hosts* instead of */net*.

To modify the amd automounter configuration file:

1. Log in as root to the render node.
2. Stop the *amd* automounter daemon by typing:
   /etc/init.d/amd stop
3. Open the */etc/amd.conf* configuration file in a text editor.
4. Locate the following lines:
   ```
   #DEFINE AN AMD MOUNT POINT
   [ /net ]
   ```
5. Change */net* to */hosts*.
   The lines should now be the following:
   ```
   #DEFINE AN AMD MOUNT POINT
   [ /hosts ]
   ```
6. Save and close the file.
7. Restart the *amd* daemon. Type:
   /etc/init.d/amd start

Configuring the NFS and amd Services to Start Automatically

By default, the *NFS* and *amd* services are set to start automatically on workstations and render nodes.

Perform the following procedure to check these services, and reconfigure their startup mode if necessary.

To configure NFS and amd as start-up services:

1. Log in as root to the workstation or render node.
2. Open a terminal, and check that the *amd* and *NFS* services are running by typing:

   chkconfig --list | egrep 'nfs|amd'

   The output of the command should contain the following lines:

   ```
   nfs             0:off 1:off 2:on 3:on 4:on 5:on 6:off
   amd           0:off 1:off 2:on 3:on 4:on 5:on 6:off
   ```

3. If the services are not set to on for run levels 2 through 5, configure them to be started automatically by typing the following commands:

   chkconfig nfs on

   chkconfig amd on

4. Reboot the workstation or render node for the changes to take effect. Type:

   reboot

Enabling Export Permissions

Edit the */etc/exports* file on the Visual Effects and Finishing workstation so that render nodes on the background processing network can access the media storage for the jobs submitted to Burn.

Enable access to the following directories.

Directory	Description
/usr/discreet/clip	Enables remote access to clip libraries for Wire.
/usr/discreet/project	Enables remote access to projects by Burn at Burn start-up.

To edit the */etc/exports* file:

1. Log in as root, and open a terminal.
2. Open the */etc/exports* file in a text editor.
3. Enable access to the */usr/discreet* directory on your workstation by adding the following line to the end of the file:

   /usr/discreet *(rw,sync,no_root_squash)

NOTE The previous line allows any remote system access to the */usr/discreet* directory. You can restrict access by adding the host names for each render node to the line, separated by a colon (:). For example, to restrict access to the */usr/discreet* directory to nodes *burn01* to *burn03*, type: **/usr/discreet burn01:burn02:burn03(rw,sync,no_root_squash)**.

4 Save and close the file.

5 Apply the changes by typing:

/usr/sbin/exportfs -va

Verifying Stone and Wire Connectivity from the Background Processing Network

The following information and procedure applies to render nodes on a background processing network.

Render nodes access frames on storage devices attached to the workstation using the Wire network. Do the following procedure to ensure that these storage devices are available to the render node.

To verify Stone and Wire connectivity from a render node:

1 Log in as root to a render node on the background processing network and open a terminal, if necessary.

2 View all storage devices available to the render node using the *sw_framestore_dump* utility. Type:

/usr/discreet/sw/tools/sw_framestore_dump

All storage devices attached to the Wire network appear.

3 If a storage device does not appear in the list of devices available to the render node, check the following:

- The workstation is on the same network as the render node.
- The workstation to which the device is attached is available on the network and can be contacted using the *ping* command.
- The *probed* daemon running on the workstation is using the same port as the rest of the network. See the entries for *sw_probed* and *sw_probed.cfg* in the latest *Autodesk Visual Effects and Finishing Installation and Configuration Guide*.
- Verify that the *sw_framestore_map* file on the workstation contains only the local media storage. This procedure is necessary for the

self-discovery mechanism in Stone and Wire to work correctly. See the following procedure for instructions.

To verify the *sw_framestore_map* file:

1. Log in as root to the workstation and open a terminal.
2. Open the file */usr/discreet/sw/cfg/sw_framestore_map* in a text editor.
3. Verify that only the local media storage appears in this file. If other media storage devices appear, delete them, unless server self-discovery is not enabled for Stone and Wire.

   NOTE Media storage devices must be explicitly listed in the *sw_framestore_map* file if server self-discovery is not enabled for Stone and Wire, such as when you are using legacy disk arrays. Use the **sw_framestore_dump** command to see if storage devices for other workstations on the network can be viewed.

4. Save and close the file.

Checking Burn Licensing for Render Nodes

The following information and procedures apply to all render nodes as well as the workstation or render node being used as the License Server for the background processing network.

Render nodes must get Burn licenses from the License Server to process jobs. If the License Server is not working, or if Burn licenses cannot be checked out, jobs are not rendered. Do the following procedure to ensure that licensing for Burn is working on the render nodes.

To check Burn licensing on the License Server workstation or render node:

1. Log in as root to the render node or workstation running the License Server for the background processing network.
2. Check if the License Server is set to start automatically. Type:

   chkconfig --list | grep license_server
3. If the License Server is not set to start running automatically, configure it to do so. Type:

   chkconfig license_server on

   If the License Server was not started, start it by typing:

/etc/init.d/license_server start

4 View the *boot.log* file for the License Server to ensure that the Burn license is being read by the server. Type:

more /usr/discreet/licserv/log/boot.log

The *boot.log* file appears, showing the license information for Burn. Check the following information in this file to see if Burn is being licensed properly on the network:

- The versions of Burn that are licensed for the network
- The start and expiry dates for Burn licenses
- The workstation or render node running as the License server for the background processing network
- The availability of each Burn license for render nodes

5 View the *DL_license.dat* file to see the number of Burn licenses purchased. Type:

more /usr/discreet/licserv/licenses/DL_license.dat

The *DL_license.dat* file appears, showing the number of Burn licenses purchased and available to the network. Check to see if the number of purchased licenses are enough for your network. Otherwise, render nodes on your network may be unable to render Burn jobs due to a lack of licenses.

To check Burn licensing on a render node:

1 Log in as root to a render node that is not running the License Server.

2 Contact the render node or workstation running the License Server using the *ping* command by typing the following:

ping `<hostname>`

where `<hostname>` is the hostname or IP address of the render node or workstation running the License Server.

If this step fails, check the network connection between the render node and the render node or workstation running the License Server, then retry the *ping* command. Otherwise, continue to the next step.

3 View the *DL_license.dat* file to check that the render node is licensed for Burn. Type:

more /usr/local/flexlm/licenses/DL_license.dat

The *DL_license.dat* file appears. This file should look similar to the following:

```
SERVER exuma-001
DLHOST01=886C2B75E8E57E4B03D784C3A2100AC0
DAEMON discreet_l discreet_l
USE_SERVER
```

Contact Customer Support if the general format or number of lines in the *DL_license.dat* file for your render node is different than the example above.

4 Repeat steps 1 to 4 for the remaining render nodes on the background processing network.

NOTE If the License Server for your network is running on a render node, make sure you perform this procedure on this node as well. Otherwise, this node is able to distribute Burn licenses to other render nodes, but is unable to retrieve a license for itself.

Reviewing Burn and Backburner Log Files

The following information and procedure applies to render nodes on a background processing network.

Review logged events in Burn and Backburner to identify the source of problems with the background processing network.

Since a background processing network may include many render nodes, a script called *GATHER_BURN_LOGS* is provided to collate and organize the log files from many render nodes into a single text file.

The script collects the Burn log files from the */usr/discreet/log* directory, and the Backburner logs from the */usr/discreet/backburner/Network* directory of each Burn node.

Use the following procedure to troubleshoot your background processing network, using the *GATHER_BURN_LOGS* script.

To review Burn and Backburner logs from your background processing network:

1 If necessary, create a list of render nodes from which Burn and Backburner logs should be collected.

2 On your Autodesk Visual Effects and Finishing workstation, log in to the account for your Autodesk application and open a terminal.

3 Navigate to the *bin* directory for your application by typing:

cd /usr/discreet/<product_home>**/bin**

where <product_home> is the directory where your application is installed (such as flame_2011).

This directory contains scripts and binary files for the application, including those used for submitting jobs to the background processing network.

4 Run the *GATHER_BURN_LOGS* by typing the following:

./GATHER_BURN_LOGS[options]<render_node_1 render_node_2 ...>

where:

<render_node_1 render_node_2 ...> is the list of render nodes on your network that you created in step 1.

[options] include the following options for this script.

Option	Function
-**h**	Shows the usage for the script.
-**v**	Shows the version of the *GATHER_BURN_LOGS* script.
-**V**	Enables additional verbosity.
-**r**<version>	Retrieves logs for the specified version of Burn, such as 2012.
-**b**<build>	Retrieves logs for the specified build of Burn.
-**u**<filename>	Writes the collected Burn logs to a single output file, ordered by each log file's timestamp. By default, the output file is named *burn.log* and is created in the */bin* directory of the application.
-**a**<filename>	Writes the collected Backburner logs to a single output file, ordered by each log file's timestamp. By default, the output file is named *backburner.log* and is created in the */bin* directory of the application.
-**e**<filename>	Writes the events in the collected Burn logs to a single output file, ordered by each event's timestamp. By default, the output file is named *burnevents.log* and is created in the */bin* directory of the application.

Option	Function
-l<username>	Uses the specified user name to log in to the render node. If this option is omitted, you may be prompted to provide a username and password.

The *GATHER_BURN_LOGS* script collects Burn and Backburner logs from each render node listed in the script and collates these logs to a single text file.

NOTE The *GATHER_BURN_LOGS* script uses the *rsh* protocol to access Burn render nodes. To avoid being prompted for login information while running this script, create a user account on the workstation and populate the *.rhosts* file on each render node with the hostname of the workstation followed by the username of this account. You can then use the -l option to run the script as this user.

Testing Network Components

Reviewing Burn and Backburner log files can help to indicate the source of a problem on the background processing network. For example, log messages showing network timeouts for jobs may indicate a problem with the TCP/IP protocol or the Stone and Wire network.

Use the following scripts that are installed with Burn to test background processing network components and identify the source of the problem:

- The *verifySWConn* script is used to test Wire network connectivity in the background processing network. See Testing Stone and Wire Connectivity for Burn (page 71).
- The *verifyBurnServer* script is used to test whether a render node meets minimum hardware requirements for processing Burn jobs on a network. See Testing Render Node Hardware for Burn (page 73).

Testing Stone and Wire Connectivity for Burn

The following information and procedure applies to workstations and render nodes on a background processing network.

Use the *verifySWConn* script to see if problems are due to Stone and Wire. This script can be run from a workstation or from a render node on the background processing network using the following procedure.

To test Stone and Wire connectivity:

1. Log in to the workstation or render node as follows:
   - For workstations, log in using the account for your application and open a terminal.
   - For Burn render nodes, log in as root and open a terminal, if necessary.
2. Navigate to the *bin* directory for your application by typing:

   cd /usr/discreet/`<product_directory>`**/bin**

   where `<product_directory>` is one of the following:
   - For workstations, the directory where your application is installed, such as `flame_2011`
   - For Burn render nodes, the directory containing the Burn version used on the background processing network, such as `burn_2011`

   This directory contains Burn-related scripts, including the *verifySWConn* script.
3. Run the *verifySWConn* script by typing the following:

   ./verifySWConn `<options> <Workstation1 Workstation2 ...>`

   where:

   `<Workstation1 Workstation2 ...>` are the workstations or Burn render nodes on which Stone and Wire are to be tested. You must specify at least two workstations or render nodes to be tested.

   `<options>` include the following options for this script.

Option	Function
-v	Enables additional verbosity.
-l`<username>`	Uses the specified user name to log in to the render node. If this option is omitted, you may be prompted to provide a username and password. The username must be valid on the test host.

The *verifySWConn* script runs, testing Stone and Wire functionality between each network component specified in the script.

Testing Render Node Hardware for Burn

Use the *verifyBurnServer* script to see if a Linux system has the hardware requirements to be used as a Burn render node.

The script also checks if a render node has a supported GPU-accelerated graphics card and graphics driver version, necessary for processing jobs that require a GPU.

To test render node hardware for Burn:

1. Log in to the render node as root and open a terminal.
2. Navigate to the *bin* directory for Burn by typing:

   cd /usr/discreet/`<burn_version>`**/bin**

   where `<burn_version>` is the directory containing the Burn version used on the background processing network, such as `burn_2011`.
3. Run *verifyBurnServer* by typing the following:

   ./verifyBurnServer

   NOTE This script takes no options.

   The *verifyBurnServer* script checks the hardware of the system to ensure it meets the requirements for Burn render nodes, and displays the results.

Assessing Compatibility between Jobs and Render Nodes

Read the following sections if you suspect problems with Burn may be due to memory issues on the render nodes, or that render nodes do not meet the graphics card requirements for a certain job type.

Processing Jobs that Require a GPU

Some of the jobs created in your Visual Effects and Finishing application (for example, floating point jobs, such as unclamped colors in Action, directional RGB blur, radial RGB blur) require a GPU-accelerated graphics card in order to be rendered.

While your workstation is equipped with a GPU-accelerated graphics card, and can render such jobs locally, your background processing network is unable to render these types of jobs if no Burn node is equipped with a GPU.

To see if a render node has the hardware capabilities to process jobs that require a GPU, use the *verifyBurnServer* script, Backburner Monitor, or Backburner Web Monitor.

If you attempt to submit a job that requires a GPU to a background processing network where no render node is equipped with a GPU, one of the following situations occurs:

- If the `BackburnerManagerGroupCapability` keyword in the application's *init.cfg* file is set up correctly, the application does not attempt to submit the job to the background processing network, and an error message is displayed. You must render the respective job locally on the Visual Effects and Finishing workstation.
- If the `BackburnerManagerGroupCapability` keyword is not set up properly, no error message is displayed and the application attempts to send the job to the background processing network. Since no render node can process the job, the job will be stuck in the queue indefinitely.

  Use Backburner Monitor or the application's Background I/O window to remove the job from the queue, and then set the `BackburnerManagerGroupCapability` keyword properly to reflect the hardware capabilities of your background processing network. See Monitoring and Managing Burn Jobs (page 52), and Configuring Burn-related Keywords for Linux Applications (page 33).

To avoid further problems, before attempting to submit a job that requires a GPU to your background processing network, make sure at least one of the render nodes is equipped with a GPU, and that the `BackburnerManagerGroupCapability` keyword in the application's *init.cfg* file is set up correctly. See Hardware Requirements for Burn Nodes (page 11), and Configuring Burn-related Keywords for Linux Applications (page 33).

Diagnosing Memory Issues

Inferno 2012, Flame 2012, Flint 2012, Smoke 2012, and Backdraft Conform 2012 are all 64-bit applications, and can thus make full use of up to 16 GB of memory.

As a general rule, render nodes should have the same amount of RAM as the Visual Effects and Finishing workstation you are sending jobs from.

A Burn server running on a render node equipped with less memory than what is installed on your Visual Effects and Finishing workstation, may fail when processing these jobs due to their higher memory demands. However, do not assume that every problem on render nodes with less memory than your workstation is exclusively caused by memory issues.

This section explains how to diagnose and address problems that are caused by jobs submitted from workstations with more memory than the render node.

If you suspect that a render node has failed due to a Burn job exceeding the Burn nodes memory capacity, use the following procedure to view the Burn and Backburner server logs for memory-related errors.

To check the Burn server log on a render node:

1. If you are running graphics on the render node, log in as root and open a terminal. Otherwise, just log in as root.
2. Navigate to the */usr/discreet/log* directory.

   This directory contains logs of events for the Burn servers installed on the render node. You need to view the log created at the time the server failed.
3. Identify the Burn log file from the time of the Burn server failure using one of the following methods:
   - If the render node has just failed, look for the following file: *burn<version>_<render_node_name>_app.log*.
   - If the render node failed previously and was brought back online, type **ls** and then look for the file *burn<version>_<render_node_name>_app.log.##* created around the time of the render node's failure.
4. View the contents of the Burn log file directly in the terminal or by using a text editor.
5. Review the messages in the log file for entries similar to the following:

   ```
   [error] 8192 PPLogger.C:145 01/24/06:17:06:16.998 Cannot load video media in node "clip17" for frame 2
   ```

   ```
   [error] 8192 PPLogger.C:145 01/24/06:17:06:17.210 Out of memory for image buffers in node "clip6" (76480512 bytes).
   ```

   ```
   Increase your memory token.
   ```

These entries may indicate that the render node was experiencing memory problems at the time of failure. Next, check the Backburner Server log file to corroborate these findings.

6. Navigate to the */usr/discreet/backburner/log* directory.
7. Identify the Backburner Server log file *backburnerServer.log* from the time of the Burn server's failure, using the methods listed in step 3.
8. Review the messages in the Backburner Server log file in a text editor, looking for entries similar to the following:

```
[notice] 16387 common_services.cpp:45
01/24/06:17:06:10.069 Launching 'burn'

[error] 16387 common_services.cpp:37
01/24/06:17:06:48.182 Task error: burn application
terminated (Hangup)

[error] 16387 common_services.cpp:37
01/24/06:17:06:48.182 burn application terminated
(Hangup)

[notice] 16387 common_services.cpp:45
01/24/06:17:06:48.524 Application is down
```

   These log entries confirm that a Burn server failure occurred on the render node. Since you know that the Burn server failed around this time, you can deduce that the memory problem caused the Burn server to fail.

9. Optional: Identify the workstation running the application that submitted the job to Burn, and then look at the Batch setup, timeline segment, or clip to try and determine why the Burn server failed.

   Problems that cause the Burn server to fail due to lack of memory on a render node, usually arise due to:

   - The size of images used in a project. For example, projects using higher resolution HD, 2K, and 4K images require more memory to store and render than SD projects.
   - The complexity of the effect sent for processing. For example, a complex Batch setup with many layers and effects requires more memory to render than a simple Batch setup.

Knowing what factors caused the render node to fail may help you to gauge what jobs your render nodes can handle. It can also give you ideas about how to deal with this problem.

Addressing Memory Issues

If Burn2012 servers on your render nodes are failing while processing jobs, increase the amount of RAM set aside for processing Burn jobs. You must repeat this procedure on each render node on your network running the Burn server.

To configure Burn to reserve a set amount of RAM for jobs:

1. Log in to the render node as root and open a terminal.
2. Stop the Backburner Server on the render node by typing the following:

   /etc/init.d/backburner_server stop
3. Open the */usr/discreet/burn_<version>/cfg/init.cfg* file in a text editor.
4. Uncomment the *MemoryApplication* keyword.

   This keyword sets the amount of RAM in megabytes (MB) to be reserved for Burn jobs. This keyword is disabled by default so Burn can dynamically adjust the amount of RAM used for each job based on the resolution of the project. When you enable this keyword, Burn reserves the same amount of memory for each job regardless of the project's resolution.
5. If necessary, change the value for the *MemoryApplication* keyword to set the amount of RAM (in MB) to be reserved for each Burn job up to 1400 (about 1.4 GB).

   For example, to reserve 1 GB of RAM for Burn jobs, type the following:

   MemoryApplication 1024

   NOTE Setting the *MemoryApplication* keyword so that the (total render node memory) - (value of *MemoryApplication*) is less than 2600 MB may adversely affect the stability of the render node.
6. Save and close the *init.cfg* file.
7. Restart the Backburner Server on the render node by typing:

   /etc/init.d/backburner_server start

   Burn is restarted and reserves the amount of RAM set in the *MemoryApplication* keyword.
8. Optional: Implement the following guidelines for processing Burn jobs:
   - If you know that the size of images in your projects may cause render node failure, enforce guidelines about what can and cannot be sent to the Burn render nodes. For example, if you know that 2K and 4K

images with Batch setups exceeding six layers may cause the render nodes to fail, ensure these setups are not sent to Burn.

- If you know that the complexity of the effects sent for processing may cause render node failure, simplify effects by creating multiple Batch setups or by processing memory-intensive effects locally. For example, if you know that complex Batch setups with multiple logic ops and colour correction may cause render nodes to fail, render these locally instead.

  NOTE Although these guidelines are not mandatory, following them may help increase the success rate while processing Burn jobs on render nodes with limited memory resources.

If, after following these guidelines, your render nodes still fail because of low memory, you might consider adding memory to the render nodes. Matching the amount of memory on the render nodes with the amount of memory found on your Visual Effects and Finishing workstation is the most effective solution to memory issues.

Index

J

L

M

N

O

P

R

S

T

W